Enhancing Energy Systems with Privacy-Aware Data Sharing and Collaborative Intelligence

A Thesis Submitted to

School of Science and Technology, Department of Engineering,

City St George's, University of London, UK

In Partial Fulfilment of the Requirement for the Degree

Doctor of Philosophy (PhD)

In Information Engineering

By

Veniamin Boiarkin

December 2025

Supervised By

Professor Muttukrishnan Rajarajan

To my mom...

Declaration

I hereby declare that no portion of the work contained in this document has been submitted in support of an application for a degree or qualification of this or any other university or other institution of learning. All verbatim extracts have been distinguished by quotation marks, and all sources of information have been specifically acknowledged.

I, Veniamin Boiarkin confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Veniamin Boiarkin
2025

Acknowledgements

The past four years have been an extraordinary journey, one that has shaped me far beyond the bounds of academic research. Along the way, I have learned the value of consistency, resilience in the face of challenges, effective time management, and many other essential skills. This thesis is the culmination of sustained effort and dedication, made possible through the support and encouragement of many individuals. This section serves as a small gesture of gratitude to those who have accompanied and supported me throughout this journey.

First, I would like to thank my PhD supervisor, Professor Muttukrishnan Rajarajan, for his support throughout all aspects of my PhD journey—from publishing my first conference paper to presenting our collaborative project at Level39, Canary Wharf. He consistently acted as a mentor, offering guidance and support across various areas, including research, business, and life in general.

I have been fortunate to work alongside outstanding researchers and mentors throughout my PhD journey. I would like to express my sincere gratitude to the following individuals for their invaluable support and guidance. I am especially thankful to Dr. Bruno Bogaz Zarpelao for his encouragement and contributions to the development of innovative ideas and solutions. I also extend my appreciation to Dr. Jafar Al-Zaili for his consistent support, insightful guidance, and the many in-depth theoretical discussions we shared over the years. In addition, I am grateful to my colleagues Safwana, Tooba, Akanksha, Subhajit, Muthupavithran, and Hamid for their collaboration, support, and enriching discussions related to our work.

Throughout every stage of my PhD journey—through both the good days and the challenging ones—my family has been a constant source of unwavering support. I am deeply grateful to my mother and brother, whose belief in me and in the value of education never wavered. Despite the physical distance, they were always just a phone call away, offering

encouragement, strength, and understanding whenever I needed it. This work would not have been possible without their love and support. Thank you for always being there for me.

Abstract

The growing demand for electricity and the increasing complexity of energy systems have necessitated innovative approaches to efficient, secure, and sustainable energy management. Energy systems are undergoing a transformative shift driven by Smart Grid technologies that integrate renewable energy sources and distributed energy sources, as well as advanced data-driven technologies. These innovations aim to enhance energy management, reduce environmental impact, and empower consumers as active participants in energy markets. However, traditional energy systems face challenges such as inefficiencies in pricing and energy trading, centralization risks, and data privacy concerns. Recent research highlights the limitations of centralized systems, emphasizing the need for secure, scalable, and user-centered approaches that preserve data privacy while enabling efficient energy management. Emerging technologies, such as blockchain, differential privacy, and federated learning, offer promising solutions to address these challenges. This work focuses on innovative pricing models for energy trading, advanced privacy-preserving techniques for data-sharing, and secure collaborative frameworks for energy demand forecasting to enhance the functionality, security, and equity of modern energy systems. To this end, several contributions are presented.

The first contribution is a novel dynamic pricing model tailored for a microgrid of prosumers with photovoltaic panels. The proposed model introduces mathematical frameworks for determining equilibrium prices based on supply-demand ratios and incorporates mechanisms for calculating energy usage costs, profits, and penalties for participants who deviate from predicted energy profiles. The effectiveness of the model is validated using real-world energy profiles, showcasing its potential to reduce energy costs.

The second contribution is a blockchain-based data-aggregation scheme for a microgrid of prosumers. This scheme ensures prosumers' privacy by concealing their real energy usage data, thereby mitigating risks like eavesdropping and man-in-the-middle cyber

attacks. The proposed scheme utilizes on-chain and off-chain techniques to protect user privacy while efficiently reducing the blockchain size. The results show that the proposed technique achieves high accuracy in data aggregation while ensuring user privacy.

The third contribution is a user-centric data-sharing scheme that leverages local differential privacy techniques to preserve the privacy of end-users. The proposed scheme provides the end-user with control over the utility of their data, with the level of privacy being calculated from individual utility preferences. The results show that the proposed scheme allows keeping the utility within the boundaries defined by the end-user, while providing the maximum possible level of privacy.

Lastly, a privacy-preserving federated learning framework incorporating local differential privacy is designed. This framework enables several parties to collaboratively train a central machine learning model without sharing their private datasets. Compared to the state-of-the-art schemes that propose a fixed privacy setting based on a number of simulations, in the proposed scheme, the privacy configuration changes over time to match the pattern of the desired level of accuracy of the central model.

Publications

The following publications have been produced from the research conducted in this thesis.

International Conference Papers

- **Publication I**: **V. Boiarkin**, W. Asif and M. Rajarajan, ”Decentralized Demand Response Power Management System for Smart Grids,” 2020 IEEE 8th International Conference on Smart Energy Grid Engineering (SEGE), Oshawa, ON, Canada, 2020, pp. 70-74, doi: 10.1109/SEGE49949.2020.9182024.
- **Publication II**: **V. Boiarkin**, M. Rajarajan and W. Asif, ”Fair Pricing Model for Smart Grids,” 2021 9th International Conference on Smart Grid and Clean Energy Technologies (ICSGCE), Sarawak, Malaysia, 2021, pp. 9-14, doi: 10.1109/ICSGCE52779.2021.9621701.
- **Publication III**: **V. Boiarkin** and M. Rajarajan, ”A novel Blockchain-Based Data-Aggregation scheme for Edge-Enabled Microgrid of Prosumers,” 2022 Fourth International Conference on Blockchain Computing and Applications (BCCA), San Antonio, TX, USA, 2022, pp. 63-68, doi: 10.1109/BCCA55292.2022.9922099.
- **Publication IV**: **Boiarkin, V.**, Zarpelao, B. B., Rajarajan, M., Roy, R., Tapper, K., ”Local Differential Privacy-Based Data-Sharing Scheme for Smart Utilities,” 202320th International Conference on Manufacturing Research, Aberystwyth, UK, 2023, doi: 10.3233/atde230930

Journal Articles

- **Publication V**: **V. Boiarkin**, M. Rajarajan, J. Al-Zaili, and W. Asif, ”A novel dynamic pricing model for a microgrid of prosumers with photovoltaic systems,” in Applied Energy, Volume 342, 2023, doi: 10.1016/j.apenergy.2023.121148.

- **Publication VI**: **V. Boiarkin**, B. Bogaz Zarpelao, J. Al-Zaili and M. Rajarajan, ”A Privacy-Preserving User-Centric Data-Sharing Scheme,” in IEEE Access, vol. 12, pp. 149977-149987, 2024, doi: 10.1109/ACCESS.2024.3476209.

Table of Contents

List of Figures

List of Tables

Abbreviations

AMI Advanced Metering Infrastructure.

BAN Building Area Network.

BSC Binance Smart Chain.

CDF Cumulative Distribution Function.

CIS Customer Information System.

DAOs Decentralized Autonomous Organization.

dApps Decentralized applications.

DeFi Decentralized finance.

DER Distributed Energy Resources.

DEX Decentralized Exchange.

DLT Distributed Ledger Technology.

DP Differential Privacy.

DPoS Delegated Proof of Stake.

DR Demand Response.

ECC Elliptic Curve Cryptography.

ECDSA Elliptic Curve Digital Signature Algorithm.

EdDSA Edwards-Curve Digital Signature Algorithm.

ESCO Energy Service Company.

EV Electric Vehicle.

EVM Ethereum Virtual Machine.

FAN Field Area Network.

FL Federated Learning.

GDPR General Data Protection Regulation.

HAN Home Area Network.

HIPPA Health Insurance Portability and Accountability Act.

HMG Home Microgrid.

IAN Industrial Area Network.

ICT Information and Communication Technologies.

IEA International Energy Agency.

IoT Internet of Things.

IPFS Interplanetary File System.

ISO Independent System Operator.

LDP Local Differential Privacy.

MAE Mean Absolute Error.

MDMS Meter Data Management System.

MGO Microgrid Operator.

NAN Neighborhood Area Network.

NFTs Non-Fungible Token.

NIST National Institute of Standards and Technology.

OMS Outage Management System.

P2G Peer-to-Grid.

P2P Peer-to-Peer.

PBFT Practical Byzantine Fault Tolerance.

PDF Probability Density Function.

PKI Public Key Infrastructure.

PLC Power Line Communication.

PMU Phasor Measurement Unit.

PoA Proof-of-Authority.

POS Proof of Stake.

PoSA Proof-of-Staked Authority.

PoW Proof of Work.

PV Photovoltaic.

RES Renewable Energy Sources.

RMSE Root Mean Square Error.

RTO Regional Transmission Operator.

SDR Supply and Demand Ratio.

SGO Smart Grid Operator.

TA Trusted Authority.

WAN Wide Area Network.

WSN Wireless Sensor Network.

Chapter 1

Introduction

1.1 Overview

The demand for electricity is constantly growing worldwide [1], which is caused by different factors, including but not limited to the increasing number of energy consumers and the rising number of different electric appliances emerging. According to the World Electricity Consumption forecast, the global demand for electricity is predicted to reach 40 TWh by 2040, followed by a rise to 45 TWh in 2050 (Fig. 1.1).

According to the data in the Global Energy Review report, more than half (68%) of electricity in China in Q1 2020 was produced using coal-fired energy generation compared to only 28% generated by Renewable Energy Sources (RES) [2]. In the United States, around 22% of electricity was generated by RES and around 41% using gas. On the other hand, 75% of electricity in India was generated using coal, whereas only 17% using RES. In the European Union, more than one-third (42%) of electricity was generated by RES, followed by 27% generated using nuclear sources, whereas 10% of the energy output was derived from coal. Global coal consumption reached its peak in 2023, with around 8.7 billion tonnes used to generate energy, a 10% increase from 2014 [3]. In 2024, coal continues to play a substantial role in energy generation, particularly in developing countries like China and India, which are the largest producers and consumers of coal. Oil and natural gas continue to play a crucial role in the global energy mix, especially in sectors like transportation and industry. According to the International Energy Agency (IEA) report [4], demand for oil and natural gas is expected to reach its peak by 2030, but current usage remains high. The reliance on fossil fuels persists despite significant investments in renewable energy and international commitments to reduce greenhouse gas emissions.

One of the key global trends is the increase in energy generation from renewable sources, such as solar, wind, and hydropower, driven by the need to reduce carbon emissions and achieve energy sustainability. According to the Renewable Energy Directive, The European Union has set a target to produce at least 42.5% of its energy from renewable sources by 2030, with an aspirational goal of reaching 45% [5]. To promote net-zero growth, the United Kingdom has set an ambitious target to increase the share of renewable energy in its electricity generation mix with the goal of achieving at least 95% clean power by 2030 [6].

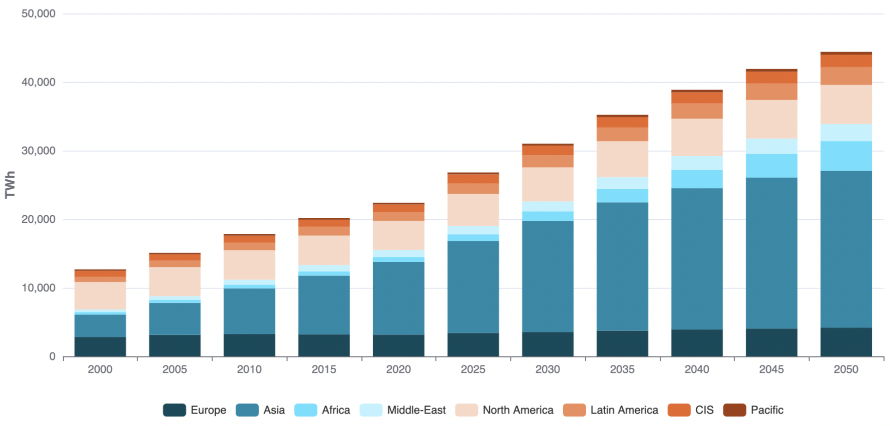

Figure 1.1: Worldwide electricity forecast consumption [Source: Electricity Forecast Consumption][1]

This global push for renewable energy is reflected in recent trends and projections, highlighting year-over-year growth in renewable energy production and the increasing dominance of photovoltaic and wind energy in future energy systems. The growth in energy generated by renewable sources can be observed in Fig. 1.2. There was a noticeable increase in Photovoltaic (PV) energy production in 2020, approximately 16% higher than in 2019. The overall growth in renewable energy production was around 5% compared to 2019. According to a net-zero report [7], around 90% of global energy production in 2050 will come from RES, whereas around 70% of energy will be produced by PV and

[1] https://eneroutlook.enerdata.net/forecast-world-electricity-consumption.html

wind together.

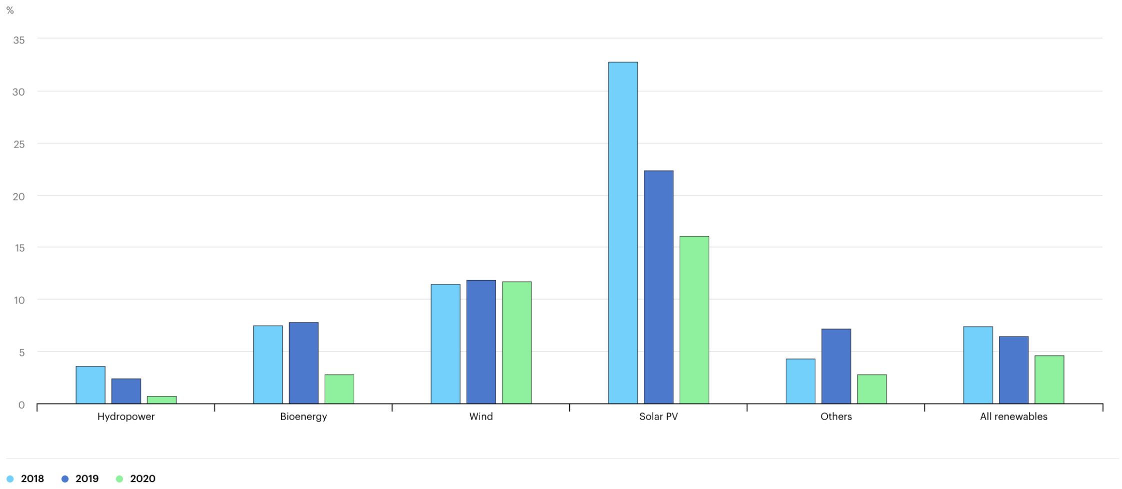

Figure 1.2: Annual growth for renewable electricity generation [Source: Global Energy Review 2020][2]

Wholesale electricity prices in many European economies, including Germany, France, the Netherlands, Spain, and Poland, have reached their highest levels as of November 2024 [8]. For instance, electricity prices in Germany have skyrocketed by 280% compared to the 2016-2019 period. The U.S. has seen a rise in electricity prices due to increased demand from sectors like artificial intelligence and data centers, which has led to higher utility bills [9]. The International Energy Agency (IEA) forecasts a 3.4% annual increase in global electricity demand through 2026, driven by economic growth and the electrification of various sectors, which can exert upward pressure on electricity prices [10].

The trend of increasing electricity prices, coupled with growing energy demand, has incentivized consumers to seek alternative energy solutions, such as photovoltaic systems, enabling them to transition from mere consumers to prosumers who can produce and potentially sell surplus energy. Prosumers can lower their energy bills by using photovoltaic systems to generate electricity, allowing them to cover a portion of the energy demand, which results in the reduced amount of energy bought from the main grid. In case the energy demand is fully covered by the energy generated using photovoltaic systems and there is still excess energy after self-consumption, the surplus can be sold to the main grid.

[2] https://www.iea.org/reports/global-energy-review-2020/electricity

Currently, despite the integration of renewable energy sources into power networks worldwide and a portion of the demand being met by local photovoltaic generation, the trading of renewable energy remains inefficient. For instance, when a prosumer generates more energy than he consumes, the surplus is sold to the main grid at a low price. Meanwhile, regular consumers must purchase energy from the main grid at a significantly higher price.

With recent advancements in Information and Communication Technologies (ICT) and massive integration of renewable energy sources into power networks, concepts of a smart grid and microgrid have become alternative solutions for efficient energy management. A Smart Grid concept refers to an advanced, interconnected electricity network that uses digital communication and advanced technologies to manage electricity flow efficiently. It operates on a large scale and aims to improve energy efficiency, reliability, and sustainability throughout the power grid. A Microgrid is a smaller, localized energy grid compared to a Smart Grid that can operate independently or in conjunction with the main grid. It operates within a specific area, such as a community of prosumers, campus, or industrial complex, with the goal of providing localized, reliable, and resilient power.

Energy transmission inevitably involves energy transfer losses, making it crucial to establish a method for allocating these losses among prosumers, thereby determining the additional costs each prosumer needs to pay. Most existing pricing models for energy trading in a microgrid focus on price formation and do not take into account the loss allocation problem. Microgrids bring possibilities for increasing efficient energy management, which leads to a reduced amount of energy imported from the main grid and a decrease in power losses during transmission. It is necessary to explore how to design and integrate a loss allocation mechanism into the pricing model, ensuring that energy prices remain more advantageous than those of the main grid, even when accounting for the additional costs associated with covering transmission losses. An effective energy pricing model would motivate prosumers to trade and share energy with the microgrid. Therefore, the pricing underlying any energy trading scheme serves as the most effective mechanism to incentivize prosumers to actively participate in the microgrid.

By providing more beneficial prices to prosumers compared to the main grid, the energy can be managed more efficiently within the microgrid. Pricing mechanisms usually depend on the supply-demand ratio, namely the overall levels of energy production (supply) and consumption (demand) within a community of prosumers [11, 12]. Prosumers report their energy usage data to a network operator that aggregates energy production and consumption data and calculates the supply-demand ratio, based on which the prices for electricity are calculated. Hence, a data aggregation mechanism directly affects pricing. In other words, the more accurate a data aggregation mechanism, the more beneficial prices can be provided to prosumers.

In a conventional scenario, prosumers report their energy usage patterns to a network operator, which becomes a single point of failure in this case. Another noticeable issue is that all the data are aggregated at one point (network operator), as well as prices for electricity and energy bills are determined on the network operator's side. This approach may cause delays in a power network's operation and problems with resiliency.

Recently, the concept of edge-computing has emerged, which allows to bring the computations closer to the customers (prosumers in a microgrid). Thus, instead of performing all tasks on the network operator's side, some tasks, including but not limited to determining the prices for electricity and energy bills, can be done closer to an edge (a microgrid). It may reduce the workload on the network operator's side, increase performance in the operation of a power network, and make real-time applications suitable for microgrids.

Both a conventional approach — where prosumers and ordinary consumers communicate directly with a network operator — and an approach where customers interact with an edge server within a local community can be vulnerable to various cyber threats. Under the General Data Protection Regulation (GDPR), energy usage data is considered personal data when it can be linked to an identifiable individual [13]. This includes energy usage data collected from smart meters, which can reveal insights into a prosumer's daily activities and behaviors. For instance, analyzing energy consumption patterns can indicate when a prosumer is occupied or when specific appliances are in use, providing information about the occupants' routines and thus disclosing their privacy. In addition, an adversary may intercept communication between two parties (a prosumer and a network operator) and may observe, alter, delete, or inject data without the knowledge of the sender or receiver, which can compromise the integrity and confidentiality of the data being transmitted. This may affect or even cause a disruption in the operation of a power network (a smart grid or a microgrid). Ensuring secure data exchange between prosumers and the network operator becomes challenging, as it requires safeguarding the privacy of prosumers and protecting sensitive information from unauthorized access.

These vulnerabilities in communication and data handling not only threaten the integrity and functionality of power networks but also underscore the broader privacy and security concerns in smart grids, particularly as they increasingly rely on sensitive prosumer data for efficient network operation and demand forecasting, which allows to predict future electricity consumption patterns. Accurate forecasts help balance supply and demand, ensuring grid stability, minimizing energy waste, and supporting dynamic pricing models that promote cost-effective consumption. Traditional centralized models for forecasting energy demand rely on collecting sensitive energy usage data from households, businesses, and microgrids, raising concerns about protecting private information. Similarly, demand-response programs, where customers are incentivized to adjust their energy usage during peak times, require analyzing individual consumption patterns, which could compromise privacy. Additionally, centralizing energy usage data for analysis increases the risk of exposing sensitive prosumer information, highlighting the need for secure and privacy-preserving solutions in smart grid systems.

Blockchain technology [14, 15] has become a promising solution to address security and privacy issues due to its key features, including non-repudiation, immutability, and robust cryptographic encryption. Additionally, certain blockchain frameworks, such as Ethereum and HyperLedger, offer support for smart contracts. Smart contracts are automated code-based agreements that execute when predefined conditions between parties are met. The use of blockchain eliminates the need for a trusted third party, streamlines interactions, and eradicates the risk of a single point of failure due to its decentralized architecture. However, blockchain technology comes with its own limitations, including restricted throughput and the continuously expanding size of the blockchain, driven by the increasing number of blocks in the chain.

Differential Privacy (DP) [16,17] has gained significant attention as one of the most popular techniques to protect individuals' data, offering mathematical guarantees of privacy. By injecting controllable noise into the data, DP ensures that analytics can be performed on aggregated datasets without exposing specific details about any individual user's data. This feature of DP is particularly valuable when utilities or third parties need to perform analytics on customer data for tasks such as energy consumption analysis. Differential privacy mechanisms can be applied to the energy usage datasets before they are shared with other parties, ensuring that even if the data is accessed, no single user's behavior can be identified. The use of DP balances the utility of the data with strong privacy guarantees, making it a practical solution for secure data sharing. However, implementing differential privacy effectively requires careful tuning of the DP mechanism's parameters to balance the trade-off between the utility and the level of privacy.

Federated Learning (FL) [18,19] has emerged as a powerful technology to address privacy and data security concerns in smart grid systems, thanks to its ability to enable collaborative machine learning without sharing raw data. By keeping data on local devices, such as prosumers' smart meters or edge servers within a local community, FL ensures that sensitive information, such as energy usage patterns, remains private while still contributing to the development of robust predictive models. For instance, edge servers equipped with FL capabilities can participate in training machine learning models locally, safeguarding customer privacy while enhancing the accuracy of the energy demand prediction. In addition, FL's decentralized approach reduces the risks associated with centralized data storage and processing, making it a compelling solution for privacy-preserving smart grid operations. Despite its advantages, FL has its own challenges, such as increased computational demands on local devices and the need for secure aggregation of distributed model updates.

The ongoing transition toward renewable energy sources, coupled with the rising demand for electricity and advancements in digital technologies, underscores the urgency for efficient, secure, and privacy-preserving energy management systems. As the energy landscape evolves, consideration of possible applications of innovative technologies, such as blockchain, differential privacy, and federated learning, will be pivotal in achieving the dual objectives of sustainability and consumer empowerment while safeguarding the integrity of modern electricity networks.

The following sections of this chapter detail the research challenges in the smart grid ecosystem, the research objectives, the contributions of the research in achieving the stated objectives, and an overview of the thesis structure.

1.2 Research Challenges

The evolution of modern energy systems, catalyzed by the proliferation of renewable energy sources, such as photovoltaic systems, and advancements in digital technology, has introduced unprecedented opportunities and challenges. Inefficient energy trading schemes lead to high energy bills and energy transfer losses during transmission, while centralized energy usage data aggregation approaches may disclose individuals' privacy. At the same time, according to GDPR, users must have control over their energy usage data and privacy while the data are transmitted to a network operator in its raw form. Similarly, the same situation can be observed at the smart grid level, where future energy demand is forecasted by aggregating raw energy usage data. The primary challenges associated with inefficient energy management can be listed as follows:

- ***Inefficient Energy Trading***: The inefficiencies in energy trading within microgrids primarily stem from conventional pricing models that do not adequately motivate prosumers to trade energy within a local community. When a prosumer has excess energy after self-consumption, it is usually sold back to the main grid at a low price, while ordinary consumers within a community purchase energy from the main grid at a higher rate. This disparity discourages prosumers from actively participating in energy trading within microgrids and underscores the need for dynamic pricing models that can maximize the benefits for all stakeholders.

- ***Centralized Data Aggregation***: In a traditional setup, prosumers share their energy usage data with a network operator, which becomes a single point of failure and a bottleneck for real-time applications in smart grids. Since the supply-demand ratio is a key parameter in most pricing models for smart grids, accurate and efficient data aggregation is crucial for enabling efficient energy management. Centralized data aggregation approaches lack the scalability to handle growing data volumes from smart grids and may be vulnerable to various cyber threats, such as man-in-the-middle and insider attacks. This poses a requirement for novel data aggregation mechanisms to ensure the high accuracy of aggregated data while maintaining privacy.

- ***Privacy Concerns***: The increasing reliance on user energy usage data for operational efficiency raises significant privacy concerns. In a conventional scenario, prosumers share raw energy usage data with a network operator for analysis. This centralized system fails to provide users with control over how their data is used. According to GDPR, energy usage patterns of prosumers are considered personal data, while users must have control over their data and privacy. Therefore, there is a need for novel privacy-preserving data-sharing approaches to empower users to manage their data preferences while enabling secure and efficient data exchange.

- ***Energy Demand Forecasting***: Currently, energy demand forecasting is performed by aggregating users' raw energy usage data. Similarly, at the smart grid level, command centers within local communities (microgrids) report energy usage data to a network operator for energy demand forecasting purposes. This approach fails to provide users with control over their data and privacy, as required by the GDPR. Aggregating the raw data of users, even utilizing encryption, may lead to various cyber threats and disclose users' privacy. Therefore, there is a need for novel approaches to energy demand forecasting that would enable achieving a high level of accuracy while protecting user privacy.

Addressing these challenges associated with modern energy systems is pivotal to fostering efficient energy management. Inefficient energy trading schemes necessitate innovative pricing approaches to incentivize participation within local communities. At the same time, the limitations of centralized data aggregation underscore the urgency for scalable and privacy-preserving approaches to ensure efficiency and resilience against cyber threats. Privacy issues and compliance with regulations, such as GDPR, further amplify the need for robust data-sharing approaches that empower users with control over their data and privacy. Achieving accurate energy demand forecasting while protecting user privacy remains a critical challenge, demanding novel techniques to balance operational efficiency with stringent data protection standards. Together, these challenges highlight the need for a comprehensive and multidisciplinary approach to innovate and transform energy systems for a sustainable future.

1.3 Research Objectives

As detailed in the previous sections, modern energy systems face different challenges that hinder their efficient and sustainable operation. Traditional energy management systems, dominated by centralized frameworks, reveal weaknesses in data privacy, scalability, and resiliency. Addressing these challenges requires innovative approaches for more efficient energy management.

The primary objective of this work is to design, propose, and evaluate novel solutions to address the identified challenges within microgrid and smart grid environments. This work examines the potential of applying advanced digital technologies and concepts, including local differential privacy, federated learning, and blockchain, within the smart grid ecosystem.

Taking into account the relatively large number of research challenges, this thesis identifies the most relevant concerns in the smart grid ecosystem, including energy trading inefficiencies, centralized data aggregation issues, privacy concerns, and energy demand forecasting. To successfully address identified challenges, the following objectives must be fulfilled:

- **Objective 1: Dynamic pricing model for a microgrid of prosumers**
  To develop a novel dynamic energy trading framework that incentivizes the participation of prosumers within a microgrid environment. The framework will incorporate a dynamic pricing model, designed to offer electricity prices that are more economically advantageous than those of the main grid, while accounting for additional factors such as energy transfer losses. This development will involve the design, implementation, and evaluation of pricing mechanisms that adapt to real-time supply-demand conditions and prosumer behavior.

- **Objective 2: Privacy-preserving data aggregation scheme for a microgrid of prosumers**
  To design a novel, privacy-preserving, and scalable data aggregation framework for microgrids composed of prosumers. The framework will focus on achieving high aggregation accuracy while preserving individual user privacy and mitigating cyber threats such as man-in-the-middle attacks. The development and evaluation of a data-aggregation mechanism will support more reliable and efficient energy management within the microgrid environment.

- **Objective 3: Privacy-preserving user-centric data-sharing scheme**
  To enable end-users to control the utility-privacy trade-off, this objective focuses on developing a novel user-centric, privacy-preserving data-sharing framework for microgrids comprising prosumers. The scheme will help end-users to choose a privacy configuration considering the error in utility they would tolerate. A core component of this framework is a data-sharing mechanism that ensures privacy-preserving data exchange by protecting against external cyber threats and maintaining compliance

with data protection regulations, such as the GDPR. This mechanism will be built on top of a new probability distribution and a noise injection mechanism that will allow adding more noise without decreasing the utility.

- **Objective 4: Privacy-preserving energy demand forecasting scheme**
  To ensure the strong protection of users' sensitive energy data while maintaining accurate energy demand forecasting results, this objective focuses on developing a novel privacy-preserving energy demand forecasting framework. The core component of the framework is a theoretical model that links privacy configurations and machine learning algorithms, promoting the use of dynamic and adjustable privacy settings. In the proposed framework, the central server can suggest a privacy configuration to the clients, which, if followed, may ensure that the accuracy of the learned model would stay as desired. In addition, the novel scheme will include an algorithm that adjusts clients' privacy configurations over time to ensure the accuracy of the central model behaves as desired during training.

1.4 Research Contributions

Designing robust and efficient solutions for energy trading, privacy-preserving data sharing, and energy demand forecasting in smart grids is challenging because it must achieve competitive results compared to the state of the art. This thesis presents a collection of innovative approaches that can be integrated to form an efficient energy management framework that fulfills the aforementioned objectives and addresses the limitations of existing methods presented in the literature. This thesis, therefore, makes the following original contributions:

- **Proposes a dynamic pricing model for a microgrid of prosumers**:
  Pricing formation is demonstrated to be a well-established field in the smart grid domain, with different approaches proposed to optimize energy trading. However, it was found that most of the existing schemes do not take into account the energy

transfer losses during the transmission of energy. In addition, very few pricing mechanisms incorporate penalty mechanisms that are based on the prosumers' behavior in terms of energy consumption. This contribution, therefore, proposes a novel dynamic pricing model to optimize energy trading within microgrids, emphasizing the importance of reducing dependence on the main grid by incentivizing prosumers to trade energy within the microgrid by providing competitive energy prices. Furthermore, the proposed model introduces penalty mechanisms based on deviations from predicted energy usage patterns and considers energy transfer losses to ensure equitable cost distribution among participants. This contribution has been published in [20].

- **Develops a data-aggregation scheme for a microgrid of prosumers**:
  Data aggregation is a well-researched area across different domains, including smart grids, with different technologies and approaches used to optimize the aggregation process. However, it was found that most approaches face scalability issues due to computationally intensive cryptographic methods employed. Therefore, this contribution proposes a novel blockchain-based data aggregation scheme for edge-enabled microgrids. This scheme leverages on-chain and off-chain components to reduce blockchain size while maintaining high privacy standards. In addition, the proposed approach achieves high aggregation accuracy and ensures resistance against common cyber threats, such as eavesdropping and man-in-the-middle attacks. This contribution has been published in [21].

- **Proposes a privacy-preserving user-centric data-sharing scheme**:
  Privacy-preserving data sharing is an important research area across different application domains. The most crucial question in this area is how to ensure the balance between end-user privacy and data utility. Traditional DP-based methods allow data aggregators to inject noise into the aggregated data to protect user privacy, but this poses security risks, such as data breaches due to centralized noise injection and possibilities of a key leakage and insider attacks. While Local Differential Privacy (LDP) enables end-users to perturb their data locally, existing approaches, such

as the Laplace-based mechanism, face challenges in maintaining high data utility without significant privacy degradation. This contribution, therefore, proposes a privacy-preserving user-centric data-sharing scheme, which conforms with LDP requirements. The proposed scheme utilizes a novel bimodal probability distribution (the Boiarkin distribution) and allows users to set their utility error boundaries, enabling the scheme to adjust the level of privacy dynamically. This contribution has been published in [22].

- **Develops a privacy-preserving framework for collaborative energy demand forecasting**:
  The problem of accurate forecasting is prevalent in various business domains nowadays, and maintaining a high level of accuracy in the central model while preserving the privacy of participants is one of the most notable challenges. Most existing schemes tend to introduce accuracy-compromising noise, posing challenges in balancing strict privacy requirements with forecasting precision. It is common that proposed solutions imply the use of a fixed privacy configuration determined based on a number of simulations. Additionally, some models increase communication and computational burdens, which may complicate deployment at resource-constrained grid nodes. This contribution, therefore, proposes a federated learning framework in which the global privacy configuration is determined by the central server, and the privacy configurations used by participating clients change over time, which is mathematically dependent on the desired level of accuracy at every particular training round.

1.5 Outline of the Thesis

The remainder of this dissertation is structured as follows:

Chapter 2 presents a survey of the state-of-the-art approaches relevant to the objectives of this thesis, followed by a discussion of the current challenges and limitations.

Chapter 3 presents the technological foundation essential for comprehending the core components of this thesis, including Distributed Ledger Technology, Edge Computing, Differential Privacy, Local Differential Privacy, and Federated Learning.

Chapter 4 presents a novel dynamic pricing model for a microgrid of prosumers with photovoltaic systems. It presents the design, a mathematical formulation, and evaluation results of the proposed pricing scheme. This work is presented in Publication V.

Chapter 5 presents a blockchain-based data-aggregation scheme for an edge-enabled microgrid of prosumers. This chapter outlines the design, findings, and evaluation results of the proposed scheme. This work is presented in Publication III.

Chapter 6 presents a novel user-centric privacy-preserving data sharing scheme. It presents the design, mathematical formulation, and evaluation results of the proposed model. This work is presented in Publication VI.

Chapter 7 presents a novel privacy-preserving federated learning framework incorporating local differential privacy. It outlines the design, mathematical formulation, and evaluation results of the proposed scheme.

Chapter 8 finally concludes the work presented in this dissertation. It describes how the research objectives have been addressed and outlines the future work directions.

Chapter 2

Related Work

2.1 Overview

The advancement of smart grid technologies, coupled with the rise of decentralized energy production, presents both opportunities and substantial challenges in managing energy efficiently and securely. Traditional approaches to energy trading, data aggregation, and energy demand forecasting are increasingly inefficient due to their limited scalability and privacy vulnerabilities. These shortcomings underscore the need for innovative frameworks and mechanisms that not only ensure operational efficiency but also empower end-users with control over their energy data and privacy preferences.

This chapter presents a literature survey of existing research publications relevant to the objectives of this thesis, focusing on energy trading mechanisms, data aggregation techniques, privacy-preserving data-sharing schemes, and privacy-aware energy demand forecasting in smart grids. Specifically, this chapter surveys the state-of-the-art in energy trading models, exploring how different techniques have been employed to improve smart grid participant engagement and cost efficiency. It then examines a range of data aggregation frameworks, including blockchain-based, cryptographic, and edge-computing approaches, aimed at achieving privacy-preserving and accurate data handling. The review continues with an analysis of data-sharing models that incorporate differential and local differential privacy, evaluating their efficacy in balancing the trade-off between data utility and privacy. Finally, it presents a review of existing privacy-preserving techniques for energy-demand forecasting, with a focus on a federated learning framework that enables collaborative machine learning model training without exposing the raw data of participants.

By reviewing the current state-of-the-art across these domains, this chapter identifies research gaps and limitations of existing approaches, setting the foundation for the novel contributions of this thesis. The remainder of this chapter is organized as follows: Section 2.2 reviews energy trading schemes in smart grids; Section 2.3 discusses data aggregation methods; Section 2.4 focuses on privacy-preserving data-sharing techniques; Section 2.5 examines privacy-aware energy demand forecasting approaches; and Section 2.6 outlines current limitations and open challenges.

2.2 Energy Trading Schemes in Smart Grids

The increasing demand for electricity, coupled with the transition toward renewable energy sources, has reshaped the energy landscape and underscored the importance of efficient energy management. Smart grids and microgrids have emerged as promising solutions to enhance energy efficiency, reliability, and sustainability. Within these systems, prosumers play an important role in enhancing energy utilization and reducing reliance on the main grid. Fostering the active participation of prosumers has become a focal point of research, driving the development of innovative energy trading schemes such as game-theoretic, optimization-based, and auction-based approaches.

A large number of game theory-based Peer-to-Peer (P2P) energy trading schemes have been proposed [23]. A P2P energy trading system for a clustered microgrid is proposed in [24]. Multi-objective game-theoretic optimization is applied to find the most suitable sizes of the players and optimized payoff values. Alhasnawi *et al.* [25] have proposed a consensus algorithm-based coalition game-theoretic approach for multi-agent smart microgrids to minimize energy mismatching, energy bill, and load energy waste. Ullah *et al.* [26] have proposed a two-level Peer-to-Peer-to-Grid energy management framework. To minimize the prosumers' total operational costs, a Distribution Locational Marginal Price (DLMP) solution with Alternating Direction Method of Multipliers (ADMM) is designed. Prosumers trade energy with the utility grid based on the DLMP pricing signals, whereas a game-theoretic model is used when prosumers trade energy with each other. In [27], a novel scheme for P2P energy trading in a smart grid has been proposed. When demand matches supply, both selling and buying prices of energy are set according to the mid-market rate, which may lead to a higher buying price of energy in the microgrid compared with the utility grid. A novel model for real-time P2P energy trading in a community microgrid has been proposed in [28], where game theory is used to model the interactions among prosumers. A two-stage bidding strategy for P2P energy market has been proposed in [29] to enhance the utilization of local renewable energy. Prosumers can adjust the energy quantities before trading, which requires constant involvement of the end-user.

Constrained optimization is another popular technique to design P2P energy trading schemes. A near-optimal algorithm, named Energy Cost Optimization via Trade (ECO-Trade), for P2P energy trading among the smart homes in a microgrid has been proposed in [30]. An unfair cost distribution problem is addressed by assuring Pareto optimality among the prosumers. To maximize the renewable energy consumption, He *et al.* [31] have proposed a P2P energy trading scheme for a community sharing market, where the interactions between participants are modeled as by a leader-follower framework. To incentivize users to participate in the energy sharing, a novel dynamic pricing mechanism is developed, whereas prosumers and consumers are assumed to adjust their energy consumption. Mehdinejad *et al.* [32] have proposed a novel energy market, where prosumers and retailers can trade energy among each other. Each consumer can buy energy from local prosumers or retailers, whereas each prosumer has the right to sell energy to local consumers or retailers. It is challenging to achieve the effective utilization of the energy generated in the local area since prosumers and consumers have the right to choose a party to trade energy with. A market-clearing mechanism for smart grid has been proposed in [33]. The P2P energy trading is formulated as a social welfare maximization problem taking into account energy losses and network utilization fees. Alskaif *et al.* [34] have proposed a fully P2P energy trading model for Residential Energy Systems to minimize the overall cost for all households. To determine the bilateral trading preferences of prosumers, two strategies are proposed, where the first one is based on matching between energy demand and supply of participants, whereas, the second one is based on the distance between the households. Liu *et al.* [35] have proposed a P2P energy trading platform for residential houses to coordinate Demand Response (DR) schemes in the hour-ahead market.

In addition to the works mentioned above, auction theory is also used to model P2P energy trading markets. In [36], an integrated model has been proposed for P2P multi-energy sharing for Home Microgrids (HMGs). To achieve the optimal coordination of energy and heat systems in each HMGs, a double auction-based multi-energy sharing mechanism is developed. Xu *et al.* [37] have proposed a novel iterative uniform-price auction mechanism for P2P energy trading in a microgrid. Prosumers iteratively adjust their bids to determine the trading price, whereas an auction is used to match surplus and deficit among seller and buyer prosumers. Haggi *et al.* [38] have proposed a framework for P2P energy trading in active distribution networks using a combination of a multi-round double auction and an average pricing mechanism, where excess energy may be sold to the utility grid because the agreement has not been reached during negotiation, which leads to inefficient energy utilization in the microgrid. A double auction-based energy trading mechanism for a community sharing market consisting of prosumers and consumers has been proposed in [39]. A non-cooperative game is applied to determine the final equilibrium spot price. Yu *et al.* [40] have proposed a continuous group-wise double auction scheme to coordinate energy transactions among prosumers in a distribution-level market.

2.3 Privacy-preserving Data-Aggregation in Smart Grids

The growing complexity of modern power networks, driven by the proliferation of renewable energy sources and the increasing participation of prosumers, has made secure and efficient data aggregation an important aspect of smart grid operations. Existing literature has explored a variety of privacy-preserving aggregation techniques, ranging from cryptographic methods to blockchain- and edge-enabled solutions, each addressing different aspects of accuracy, performance, and privacy.

DP has become a cutting-edge data privacy preserving technique, which is widely used in the domain of smart grid. Zheng *et al.* [41] have proposed a decentralized mechanism for privacy-preserving computation in smart grid that utilizes the differential privacy. Laplace noise is injected to the smart meter readings in a distributed manner, whereas a random permutation algorithm is used to shuffle the energy usage data sequence. To ensure that the prosumers' smart meter data are protected, Gough *et al.* [42] have proposed a novel DP algorithm that takes into account the privacy preferences of the end-users. By adding Laplace noise to the smart meter readings, the privacy of prosumers is ensured, whereas the higher level of privacy introduced a higher level of noise should be added to the original data. The use of the proposed algorithm increases the system costs by a maximum 5.6%. Adding noise to the smart meter readings may ensure prosumers' privacy but there is always a trade-off between the privacy and accuracy. Thus, the higher level of privacy, the higher level of noise needs to be added to the energy usage data.

Some works have proposed different data segmentation and fragmentation techniques to ensure secure data aggregation. Kapusta *et al.* [43] have proposed a keyless efficient algorithm for data protection by means of fragmentation. Initial data are fragmented into n fragments, any k of which are sufficient for data recovery, whereas, a randomly taken fragment does not provide any information about the initial data. Cryptographic strenght of the proposed algorithm depends on the use of a seed that generates k pseudo-random values, which means that the proposed algorithm is secure if the seed generates pseudo-random values in a secure way. Guan *et al.* [44] have proposed a blockchain-based dual-side privacy-preserving multi-party computation scheme for edge-enabled smart grid. To ensure the security of multi-party computation, a data segmentation method is adopted. First, the data of each participant are divided into two parts, the first one of which is kept on the participant's side, whereas, the second part is encrypted and sent to another participant. Each participant adds up a part of its own data and a part of the received data;

thus, getting the mixture of data parts, which is called data registration, whereas the sum of all data registrations is equal to the original data of the participants. The ring signature is used to hide the identity of the data sender, whereas the computational overhead of the ring signature increases with the number of public keys used. Thus, there is a lack of scalability in the proposed approach.

A number of approaches that utilize cryptographic techniques have been proposed to ensure secure communication among parties. To ensure secure data exchange between smart meters and a neighborhood area network gateway, Garg *et al.* [45] have proposed a lightweight authentication scheme for smart grid by combining fully hashed Menezes-Qu-Vanstone key exchange mechanism, one-way hash functions, and Elliptic Curve Cryptography (ECC). Initially, a smart meter and a gateway authenticate each other and then establish a session key for communication, whereas a new session key is generated with every key agreement phase. To enhance the security during the communication between a smart meter and a service provider, Srinivas *et al.* [46] have proposed an anonymous signature-based authenticated key exchange scheme for Internet of Things-enabled smart grid by combining ECC and Schnorr's signature scheme. The proposed scheme requires the presence of a Trusted Authority (TA), which is responsible for generating public and private keys for smart meters and service providers, whereas a smart meter and a service provider establish a session key for secure communication. Since smart meters are very limited in terms of computational capacity, it is challenging to run complicated cryptographic algorithms on the smart meters' side. To provide efficient data aggregation, while maintaining data integrity and privacy, Qian *et al.* [47] have proposed two secure lightweight data aggregetion schemes for smart grid by using Lattice-based homomorphic encryption and Lattice-based sequential aggregate signatures schemes. TA is used to issue IDs for smart appliances and public keys for a whole residential area, whereas, to resist a replay attack, the timestamp is used. In this approach, a smart appliance reports its energy usage data once per day, which is not suitable for real-time applications.

In addition to the works mentioned above, many blockchain-based data-aggregation schemes have been proposed. To protect user's privacy and ensure data integrity, Luo *et al.* [48] have proposed a blockchain-based data aggregation scheme for a microgrid. Smart meters encrypt energy usage data using Paillier homomorphic encryption and send encrypted data to the aggregator, which aggregates encrypted data and sends it to a bookkeeper, where the data are decrypted and written to the blockchain through smart contracts. Practical Byzantine Fault Tolerance (PBFT) is used as a consensus mechanism. In the proposed scheme, all microgrids share the same power chain (blockchain), whereas they all have different control layers. The overall size of the blockchain can be reduced if each microgrid will have its own blockchain. Yang *et al.* [49] have proposed a privacy-preserving blockchain-based energy management system for IoT-aided smart homes. To run the blockchain nodes on smart meters, a modified PBFT consensus protocol is proposed, which reduces the communication complexity from $O(N^2)$ to $O(N)$. To ensure privacy, users' private information is processed locally on the smart meters and not revealed to the blockchain. Since the computational capacity of smart meters is limited, it becomes challenging to run a blockchain node on the smart meter's side.

To reduce the communication overhead, edge computing has become a promising solution by bringing services close to the end-users. Su *et al.* [50] have proposed a lightweight and communication-efficient data aggregation scheme for smart grid. Each smart meter encrypts its energy usage data and sends it to a gateway that collects and aggregates data from smart meters and sends it to the command center. A smart meter is assumed to be tamper-resistant, whereas, the secret value of a smart meter is stored in the smart meter and cannot be disclosed or changed. To improve security, performance, and robustness of the edge layer, Lu *et al.* [51] have proposed a lightweight privacy-preserving data aggregation scheme for a smart grid by integrating blockchain, edge-computing, homomorphic Paillier encryption, and one-way hash chain technique. The energy usage data are sent from a smart meter to the edge node, whereas the distributed voting authorization certificate consensus mechanism is used to confirm a new block in the blockchain. Edge nodes vote to select candidate nodes that negotiate a master node with a block right ownership. It is challenging for a smart meter to execute complicated cryptographic algorithms that may slow down the data-aggregation process. To ensure data privacy, Wang *et al.* [52] have proposed a lightweight privacy-preserving Q-learning-based scheme for a smart grid. The energy usage data from smart meters are sent directly to the command center that then sends the data to the two edge nodes for training a decision-making model. Edge nodes perform computations without revealing any original energy usage data and send the results back to the command center, which has to add the results together to recover a plaintext. In this scheme, the command center becomes a single point of failure.

2.4 Privacy-Preserving Data-Sharing

With the increasing reliance on digital technologies and sensitive prosumer data, privacy-preserving data sharing has become a critical component of modern smart grid systems. Traditional centralized data-sharing approaches not only have scalability issues but also expose users to serious privacy risks. To enable efficient energy management while safeguarding individual data, existing literature has explored various advanced techniques, including differential privacy, local differential privacy, and federated learning.

To preserve end-user privacy, DP-based models are widely used across different domains that include but are not limited to machine learning [53], Industrial Internet of Things [54], and healthcare [55]. Most of the privacy preserving data-sharing schemes based on DP focus on injecting noise into end-users' aggregated data on the service provider's side. For example, Wei *et al.* [56] have proposed a DP-based genetic matching scheme to achieve effective genetic matching and protect genetic data before outsourcing it to an untrusted cloud server for diagnosing patients' diseases. End-users share their sensitive data with a gene provider that injects noise into the aggregated data, which means that the end-users' private data may be disclosed in case of an attack on the gene provider's side. In [56], the level of privacy is chosen by following other works.

To protect the location privacy of both workers and tasks in a location-based crowdsourcing service, Wei *et al.* [57] have proposed a novel DP-based location protection scheme. End-users share their sensitive information, including precise locations, with a cellular service provider. The provider then adds noise to the data and sends the noisy data to a server for processing. However, this process carries the risk of personal data leakage. The level of privacy is determined by the service provider, which means that end-users do not have control over their privacy. In [58], a novel traffic estimation scheme using DP is proposed to protect the vehicles' data in vehicular cyber-physical systems. A vehicle shares its location data with the roadside unit after the authentication using Public Key Infrastructure (PKI). The roadside unit perturbs aggregated data and submits it to the central server for future analysis. A key leakage attack on the roadside unit's side may cause the leakage of vehicle's data. In this case, the level of privacy is determined by the roadside unit (service provider). To achieve privacy and high estimation accuracy, the model's parameters must be set properly, but there is still no clear understanding of the relationship between privacy level and accuracy.

Instead of injecting noise on the aggregator's side, the noise can be injected on the end-user's side using a LDP-based mechanism. For example, Zheng *et al.* [59] have proposed a novel recommendation system scheme by combining a matrix factorization algorithm with LDP to prevent the leakage of end-users sensitive information. End-users' sensitive data are perturbed using the Laplace mechanism on the end-user's side based on personalized privacy requirements and then sent to an aggregator. Similar to other works, the accuracy of the model increases with decreasing level of privacy (amount of noise). End-users may adjust their level of privacy, whereas it is still unclear how the change in the end-user's privacy level affects the accuracy of the result. In [60], a novel game-theoretic federated learning framework using DP is proposed to prevent malicious clients from compromising private information of other parties through inference attacks. After performing local training, end-users perturb the trained model parameters and submit these data to a central server. To enhance the global model accuracy, the amount of noise injected on the end-user's side should be reduced, resulting in a low level of privacy. To reduce the risk of industrial data leakage in the process of deep model training, Jiang *et al.* [61] have proposed a new federated edge learning scheme using hybrid DP for industrial data processing. After training a local model, the edge terminal (end-user) generates and injects noise into the parameters of the local model, after which the noisy data are sent to the central server that generates the parameters of the new global model. To achieve better accuracy of the training model, the level of privacy chosen is relatively low compared to other works. As expected, the training loss decreases with the decreasing level of privacy (amount of noise). An optimal level of privacy is chosen (adjusted) based on the results of several simulations. To protect node feature and graph structure information against a malicious data curator, Lin *et al.* [62] have designed a novel privacy-preserving framework for decentralized network graphs based on graph neural networks using edge LDP. The central server sends a query to the end-users, whereas each end-user sends an obfuscated answer (noisy data) back to the server. The proposed model achieves high accuracy for the predefined level of privacy, which means that the level of privacy (model's parameters) is chosen based on the number of simulations. A novel privacy-preserving data aggregation scheme satisfying LDP is presented in [63] to prevent the disclosure of end-user's electricity usage habits and daily activities in the smart grid. A smart meter, which is deployed on the end-user's side, measures the electricity consumption, perturbs electricity usage data using randomized response, and submits noisy data to the aggregator.

The use of DP-based schemes implies that end-users share their sensitive data with a service provider that determines the level of privacy according to which the noise is injected into the aggregated data. Thus, end-users do not have control either over their privacy or utility. Because of sharing raw data with a third party, the end-users privacy may be disclosed due to a key leakage or insider attack on the aggregator's side. On the contrary, the use of LDP schemes allows each end-user to determine an individual level of privacy, based on which the noise is injected into the data on the end-user's side before sharing it with a third party. Both DP and LDP-based data-sharing schemes allow control over the level of privacy. However, since there is no theoretically defined relationship between privacy level and model accuracy, users have no control over the resulting data utility, even when privacy levels are predefined. Most existing schemes rely on the results of simulations and select the most appropriate level of privacy (model's parameters) that should be used to achieve high accuracy.

2.5 Privacy-Preserving Energy Demand Forecasting

Modern energy systems heavily rely on accurate energy demand forecasting to balance supply and demand, enable dynamic pricing, and ensure grid stability. The emergence of advanced techniques, such as federated learning and differential privacy, offers promising pathways to enhance forecasting accuracy while safeguarding user data.

A federated learning framework, FedGrid, has been proposed in [64] for renewable energy prediction and electricity load forecasting while ensuring data privacy. By leveraging federated learning, FedGrid enables collaborative model training without directly sharing sensitive data from participants. Experimental results demonstrate accurate predictions for residential and commercial energy consumption, supporting efficient resource utilization. However, the proposed framework does not fully address the statistical heterogeneity challenge among grid participants, potentially impacting forecasting accuracy. Iqbal *et al.* [65] have introduced a privacy-preserving split learning-based framework for load forecasting, where a deep neural network model is split between local grid stations and central service providers. The proposed solution ensures raw data remains local, only sharing intermediate model updates. Empirical experiments demonstrate this method achieves comparable or superior forecasting accuracy relative to central models. However, this approach can incur high communication overhead due to frequent transmission of intermediate results, particularly in large-scale deployments. Lin *et al.* [66] have proposed a privacy-preserving distributed energy management framework that uses vertical federated learning to collaboratively clean and process energy datasets without compromising privacy. The approach enables effective cooperation among multiple entities, enhancing data quality and operational efficiency. The framework significantly supports better energy management decisions in smart grids. However, the vertical federated learning design primarily addresses privacy for vertically partitioned data and is less effective when applied to horizontally partitioned datasets. In [67], an attention-based deep learning model integrated with federated learning to enhance the accuracy of smart grid load forecasting has been introduced. The federated learning approach maintains privacy by collaboratively training the model without sharing raw energy consumption data. The attention mechanism significantly improves the detection of temporal patterns, yielding better forecasting accuracy. On the other hand, the added complexity from the attention mechanism significantly increases computational demands, limiting efficiency at resource-constrained edge devices. A personalized federated learning for smart grid load forecasting under heterogeneous (non-IID) data distributions has been proposed in [68]. Leveraging meta-learning, the method adapts individual client learning rates, improving gradient efficiency and forecasting accuracy. Simulation results reveal the proposed approach outperforms traditional federated learning methods. Nonetheless, this personalized approach may have limited generalizability in highly dynamic grid environments or when data distributions drastically change over time. In [69], a privacy-preserving forecasting method using recurrent neural networks (LSTM and GRU) combined with differential privacy to predict electricity consumption has been proposed. The approach maintains individual privacy through added noise while achieving competitive forecasting accuracy. Experimental validation confirms a balanced trade-off between privacy and utility. However, the differential privacy mechanism's noise introduction may negatively impact forecasting accuracy, particularly under stringent privacy requirements. Riedel *et al.* [70] have proposed a federated learning model for predicting feed-in power in smart grids. Federated learning enables privacy-preserving collaboration across distributed data sources, enhancing grid operational effectiveness. The proposed scheme demonstrates improved forecasting stability and accuracy via advanced aggregation methods. Nonetheless, security risks, such as membership attacks inherent to federated learning, are not sufficiently discussed. In [71], Chang has proposed an anonymous, demand-driven privacy-preserving mechanism for data sharing within smart grids. This scheme enables grid operators to manage and optimize operations using consumption data without exposing user identities. It effectively balances operational insights and user privacy requirements. A federated deep reinforcement learning solution for managing shared energy storage systems in smart buildings while preserving privacy has beed proposed in [72]. The federated approach allows multiple buildings to cooperatively optimize energy management without sharing sensitive energy data directly. Simulation results validate effective energy optimization in heterogeneous environments. However, the proposed federated reinforcement learning solution demands significant computational resources and may encounter convergence difficulties in large-scale or complex energy systems.

Most DP-based approaches tend to introduce accuracy-compromising noise, posing challenges in balancing strict privacy requirements with forecasting precision. Additionally, some models increase communication and computational burdens, thereby complicating deployment at resource-constrained grid nodes.

2.6 Current limitations and challenges

The exploration of existing work in energy trading, data aggregation, privacy-preserving data-sharing, and privacy-aware energy demand forecasting within smart grids reveals several limitations. While recent advancements in these areas demonstrate significant potential, key gaps remain that pose barriers to practical, scalable, and privacy-aware solutions.

Most energy trading mechanisms rely on game theory, auction theory, or constrained optimization. However, these models typically assume continuous user involvement in bidding or adjustment processes, which is not realistic in real-world scenarios. Furthermore, most existing energy trading schemes do not account for energy transfer losses. This results in suboptimal or unfair trading outcomes. Furthermore, the lack of scalability and high communication overhead in real-time trading environments hinders their effectiveness in large-scale deployments.

The majority of existing data aggregation mechanisms are centralized, creating a single point of failure and introducing potential vulnerabilities to cyber threats such as man-in-the-middle or insider attacks. While some decentralized methods have been introduced, these often bring additional computational and communication overheads. Furthermore, high reliance on encryption or blockchain-based mechanisms poses scalability concerns, especially when assumed to run on resource-constrained devices such as smart meters.

Numerous privacy-preserving data-sharing schemes, particularly those utilizing differential and local differential privacy, suffer from a fundamental trade-off between privacy and data utility. The injection of noise to ensure privacy can degrade the quality of data analysis and decision-making. Most existing solutions do not empower end-users to control their privacy preferences and instead rely on static and predefined privacy parameters, often determined based on simulation results, which raises questions about the applicability of the solution in real-world settings and across other application domains. This not only violates the GDPR's emphasis on user-centric data control but also limits adaptability across diverse user scenarios.

Current energy demand forecasting frameworks, including federated learning and differential privacy-based approaches, offer privacy-preserving alternatives to centralized model training. However, conventional federated learning faces challenges of information leakage from the shared model updates, while differential or local differential privacy on its own adds additional communication overhead. More importantly, most existing methods that combine federated learning and local differential privacy techniques lack a theoretical foundation that links privacy configurations and model training algorithms, instead proposing the use of static and predefined privacy parameters determined based on a number of use case-specific simulations.

While the literature presents a wide range of solutions addressing various challenges in smart grids, their real-world applicability is often constrained by scalability, user participation requirements, computational limitations, and unresolved trade-offs between privacy and data utility. These limitations highlight the need for novel solutions that align with evolving regulatory and operational demands in modern energy systems.

2.7 Summary

This chapter has provided a review of existing approaches in energy trading, data aggregation, privacy-preserving data-sharing, and energy demand forecasting in smart grids. The analysis reveals that many existing solutions fall short of addressing the core challenges of scalability, user-centricity, and regulatory compliance, particularly with respect to GDPR requirements.

Energy trading schemes predominantly employ game theory, auction models, and constrained optimization. These models often require constant user interaction and overlook energy transmission losses, which hampers real-world applicability. Similarly, current data aggregation schemes are mostly centralized or dependent on resource-intensive cryptographic techniques, creating single points of failure and scalability concerns. While data-sharing mechanisms provide a layer of protection, they often impose a trade-off between privacy and data utility. The majority of existing approaches enforce static privacy configurations without empowering end-users to control their data-sharing preferences, thereby contradicting the GDPR's user-centric principles. Although frameworks that combine federated learning and differential privacy show promise in decentralized model training, most existing schemes lack a theoretical connection between privacy configurations and model training algorithms, promoting the use of predefined privacy settings.

The following chapter presents the background information central to this research, followed by the chapters that detail the novel research aimed at overcoming some of the identified challenges.

Chapter 3

Technological Foundation

This chapter presents the background information for understanding the technological underpinnings central to this research. First, an overview of the smart grid ecosystem is presented, detailing its conceptual framework, key components, and communication infrastructure. Next, foundational aspects of distributed ledger technologies, with a particular focus on blockchain, its architecture, and storage types, are discussed. Finally, this chapter introduces the theoretical aspects of differential and local differential privacy, and concludes with the foundational concepts of federated learning.

3.1 Smart Grid Ecosystem Overview

The Smart Grid represents a profound shift from traditional centralized energy systems toward efficient, distributed, and digitalized power networks. The concept of Smart Grid integrates advanced communication, sensing, control, and information technologies to enhance the operation of traditional power networks, including generation, transmission, distribution, and consumption of electricity [73]. The main objective of the Smart Grid is to enhance efficiency, support the integration of RES, improve grid reliability, and empower ordinary consumers and prosumers with greater visibility and control of their energy usage. The Smart Grids ecosystem goes beyond just physical infrastructure such as power lines and substations and incorporates innovative digital software solutions, data analytics approaches, smart devices, as well as a wide range of stakeholders, including utilities, regulators, technology providers, and prosumers [74].

By seamlessly integrating the latest advancements in ICT with the physical infrastructure, smart grids open doors for more dynamic energy markets, novel pricing schemes, and interactive demand response programs. Advanced devices, such as smart sensors in cooperation with real-time data analytics, allow utilities to monitor, predict, and respond to fluctuations in energy supply and demand. The utilization of advanced digital infrastructure, real-time data exchange, and distributed energy resources lead to the emergence of new types of participants in the Smart Grid ecosystem, including prosumers, aggregators, Electric Vehicles (EVs), and energy service companies, thus enabling decentralization in the power sector. With the emergence of the Smart Grid, the entire energy system is evolving into a more sustainable and efficient ecosystem, enhancing resiliency against outages, enabling smoother integration of green energy sources, improving overall power quality, and providing consumers with greater control over their energy usage.

3.1.1 Smart Grid Conceptual Model

The National Institute of Standards and Technology (NIST) developed the Smart Grid Conceptual Model [75] (Fig. 3.1) to provide a general idea of how the Smart Grid operates, as well as how different components interact. The proposed model serves as a foundational framework to guide the development of Smart Grid technologies, promote security and interoperability, and enhance understanding among different stakeholders.

The NIST Smart Grid Conceptual Model highlights the importance of standardization, seamless integration, and cybersecurity across the entire energy network. According to this model, the Smart Grid is not a monolithic system anymore but a dynamic energy ecosystem consisting of legacy systems and infrastructures, advanced and smart devices, communication platforms, Distributed Energy Resources (DER), and innovative solutions and services. Thus, the NIST model serves as a framework for planning, implementing, and managing smart grids.

NIST categorized the smart grid into seven domains, with their roles and functions described as follows:

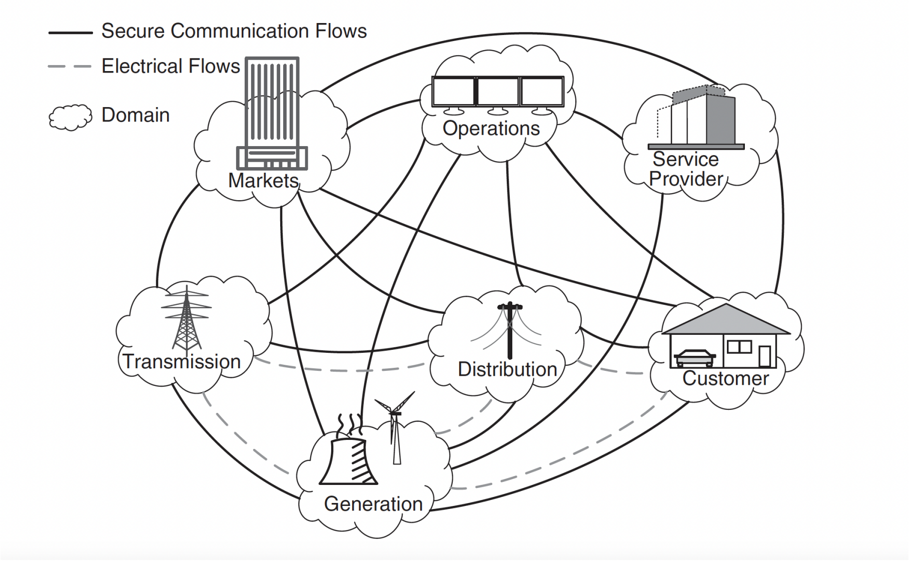

Figure 3.1: NIST conceptual model of Smart Grid [Source: [75]]

- The **Generation including DER** domain outlines the aspects of producing electricity from different energy sources, which is the first step in the process of delivering energy to customers. Traditionally, energy generation is dominated by power plants utilizing fossil fuels, hydroelectric systems, and nuclear energy. However, the generation domain has become more diverse in terms of energy mix, including wind, solar, and other energy sources. This transformation promotes cleaner and more sustainable energy generation, posing requirements for network operators to efficiently manage the diverse set of renewable energy sources. The Generation domain includes a wide range of energy resources, including but not limited to solar, wind, hydro, and biomass power plants. In addition, the generation goes beyond just producing electricity by incorporating technologies for real-time monitoring and demand forecasting to efficiently maintain the balance of supply and demand.
- The **Transmission** domain represents the high-voltage infrastructure used for bulk transfer from generation sources to the distribution network, which includes transmission lines, transformers, substations, and other electrical equipment that enables the efficient movement of energy over long distances. The main objective of the trans-

mission network, which is operated by Regional Transmission Operators (RTOs) or Independent System Operators (ISOs), is to ensure efficient energy transfer and maintain grid stability. This domain plays an important role in supporting the integration of DER and energy markets, while cybersecurity measures protect the reliability and integrity of the system. By utilizing advanced technologies for sensing and control, the Transmission domain transforms into a more flexible backbone of modern smart grids.

- The **Distribution** domain represents medium and low-voltage network infrastructure used to deliver electricity from the transmission system to customers. This domain includes transformers, distribution substations, metering points, and interconnects Transmission, Markets, and Customer domains. The emergence and integration of DER, such as solar panels, has transformed the distribution network into a dynamic and complex system. Integration of advanced technologies, such as sensing, monitoring, and smart meters, enables utility companies to isolate faults quickly, reroute power, and better management of power flows. The integration of DER and market-driven consumption behaviors underscore the need for advanced cybersecurity and privacy-enhancing technologies to maintain data privacy and trust. Through the integration of high-speed communication systems and advanced operational functionalities, the Distribution domain facilitates adaptive load management, enhances system reliability, and enables the seamless incorporation of emerging technologies.
- The **Operations** domain spans the processes and systems that ensure the reliable, secure, and efficient electricity flow across the smart grid. With the introduction of the smart grid, the operations domain has evolved from traditionally managing grid conditions, dispatching resources, and balancing supply and demand to leveraging advanced solutions for real-time data analytics, predictive modeling, and automation, enabling the efficient management of a dynamic and complex energy network. Integration of data from smart sensors, smart meters, and Phasor Measurement Units (PMUs) enables network operators to identify emerging problems, such as line overloads and frequency deviations, and respond to them in order to prevent their

escalation. Cybersecurity is a foundational aspect of the Operations domain, providing protection against potential cyberattacks and data breaches that may disrupt the integrity and functionality of the grid.

- The **Markets** domain describes the economic and regulatory frameworks that enable the trading of electricity and related services. It spans retail, capacity, and wholesale markets, involving the participation of customers, aggregators, retailers, and generators. This domain plays an important role in balancing supply and demand and ensures efficient resource allocation. The use of advancements in ICT enables transparent and timely price signals that guide decisions related to energy generation and consumption. Customers with DER, such as photovoltaic systems, can sell surplus after self-consumption back to the grid, while dynamic pricing enables the incentivization of customers to shift their loads to off-peak periods. Pricing schemes, including flat rates, time-of-use, and real-time pricing, provide cost signals to customers to align energy usage with grid efficiency.

- The **Service providers** domain encompasses entities that deliver energy-related products to energy producers, distributors, and customers, including utility companies, Energy Service Companys (ESCOs), aggregators, cybersecurity providers, and vendors that offer a variety of solutions, such as metering infrastructure and data analytics. Service Providers are responsible for core business processes, such as billing and energy management, with the aim of enhancing grid performance and customer experience. One of the functions of Service Providers is to drive innovation and economic growth by creating and adapting new services, such as personalized energy management or real-time insights, offered by utility companies or third parties entering the market. In addition, Service Providers enable utility companies to adopt dynamic pricing strategies and effectively respond to emergencies.

- The **Customers** domain is a cornerstone of the conceptual model consisting of residential, commercial, and industrial customers. Traditional customers, namely energy consumers, have become active participants in the smart grid, who can consume, manage, and even generate electricity by investing in distributed generation and storage, such as photovoltaic and battery systems. By responding to price signals, customers can optimize their energy usage, shifting loads to off-peak periods, whereas owners of EVs contribute to the grid stability by returning energy back to the grid during peak periods. The Customer domain integrates with other domains through Advanced Metering Infrastructure (AMI), which enables features such as energy usage monitoring, remote load control, and generation management. The active participation of customers in energy management transforms a conventional grid into a collaborative and dynamic energy ecosystem.

The NIST conceptual model of smart grids represents a comprehensive framework of the modern energy ecosystem, underscoring the complexity and dynamics of the smart grid. Each domain in the conceptual model plays an important role in achieving goals of efficiency, reliability, and sustainability. These domains enable the integration of advanced technologies and solutions, supporting the emergence of DER and empowering customers to actively participate in energy management. This conceptual model serves as a technical guide and a strategic vision to facilitate the transition from the traditional one-way power system to an advanced, interactive, and data-driven energy ecosystem. By highlighting the importance of cybersecurity, interoperability, and scalability, the model enables seamless collaboration of different stakeholders with the aim of safeguarding energy infrastructure.

3.1.2 Essential components of the Smart Grid

The Smart Grid is the next step in the evolution of the electrical grid that integrates advanced technologies for communication, monitoring, and control to improve efficiency, reliability, and sustainability in energy generation, transmission, and distribution. The integration of traditional network infrastructure with advanced sensors, communication technologies, and smart systems enables dynamic and real-time management of modern power networks. The essential components of the smart grid include smart sensors, PMUs, smart meters, and wireless sensor networks, each contributing to the adaptability and resilience of the energy grid. These components with their roles and functions are described as follows:

- **Smart Sensors and Sensor Networks** form the foundation of smart grids, providing accurate data for system optimization. These smart devices measure parameters, such as voltage, current, and frequency, across generation, transmission, and distribution networks, enabling real-time monitoring and control. By integrating smart sensors, smart sensor networks ensure the confidential and dependable transfer of measurement data to centralized or distributed management systems, alongside high efficiency. These components enable early fault detection, outage management, and power rerouting, ensuring grid stability under various circumstances.
- **Phasor Measurement Units** are crucial for efficient monitoring and control, providing measurements of electrical waves using GPS signals and offering insights into important parameters, such as voltage and current across the smart grid. By enabling wide-area monitoring, PMUs facilitate real-time response to disruptions in the smart grid, which is pivotal in grid stability analysis and performance assessment.
- **Smart Meters** represent an integral component of the smart grid, enabling advanced monitoring and management of energy usage by providing a two-way flow of energy usage data between utility companies and customers, which empowers both parties to optimize energy utilization and improve the efficiency of the smart grid. By monitoring electrical parameters, such as current and voltage, smart meters can track energy usage and detect anomalies like energy surges. In addition, smart meters support dynamic pricing mechanisms, which allows utility companies to adjust electricity prices based on supply and demand, whereas customers can adjust their energy consumption patterns to reduce costs. Smart meters play an important role in the process of transformation of traditional power networks into smart grids, enabling demand-side management and supporting the integration of renewables.
- **Wireless Sensor Networks** enable seamless communication among different components of the smart grid, connecting smart sensors, smart meters, and other devices

across wide geographical areas using communication standards such as Zigbee, Wi-Fi, and LTE. By reducing the need for extensive wiring, supporting real-time data acquisition, system monitoring, and the integration of RESs and EVs, Wireless Sensor Networks (WSNs) ensure a dynamic and adaptive smart grid infrastructure.

The integration of these essential components transforms traditional power networks into dynamic and intelligent smart grids, enhancing the smart grid's ability to respond efficiently to changes in supply and demand under various conditions and supporting the seamless integration of emerging technologies and services.

3.1.3 Communication in Smart Grid

Smart Grids represent an advancement in energy networks by integrating traditional power grids with cutting-edge information and communication technologies. The communication network enables seamless data exchange among various components of the smart grid. By utilizing wireless and wired communication methods, the smart grid facilitates the development of robust and scalable solutions.

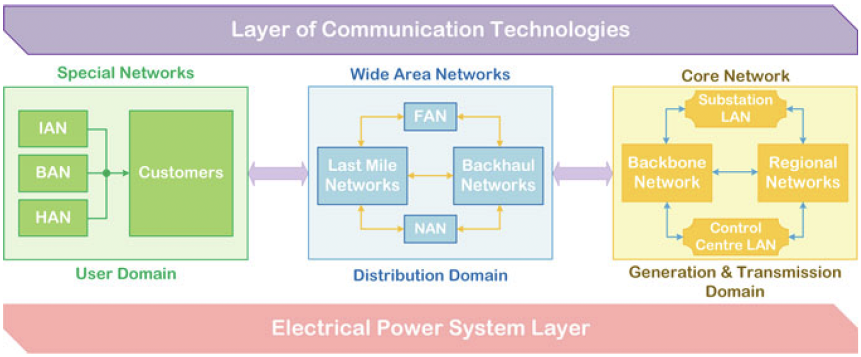

Figure 3.2: Smart grid communication architecture (IEEE 2030 standard [76]) [Source: [77]]

The communication architecture of the smart grid [76] is fundamental to ensure efficient data communication in the modern energy ecosystem. This architecture (Fig. 3.2) consists of three domains, including the Generation and Transmission Domain, Distribution Domain, and User Domain. The User Domain includes Home Area Networks (HANs), Industrial Area Networks (IANs), Building Area Networks (BANs) to facilitate local data exchange and management. The Distribution Domain consists of Neighborhood Area Networks (NANs) and Field Area Networks (FANs), facilitating the aggregation and transmission of data from local networks to centralized systems. The Wide Area Network (WAN) is at the core of the communication architecture of the smart grid, which connects generation and distribution centers with control systems. This layered architecture enables the efficient operation of advanced applications, such as AMI.

A **Home Area Network** enables the integration of various smart devices within a home, providing real-time monitoring and control of energy usage. By employing technologies such as ZigBee, Wi-Fi, and Bluetooth, HAN facilitates communication among different components, such as home appliances, energy management systems, and smart meters.

The **Industrial Area Network** supports communication among various industrial machinery and smart meters within industrial facilities, enabling efficient monitoring, control, and automation. IANs use robust and high-speed communication technologies, such as Ethernet and industrial wireless networks, ensuring optimized energy consumption and reduced operational costs.

Building Area Networks are responsible for managing communication among various sensors, smart meters, and control devices by connecting them to central energy management systems within commercial or residential buildings. By utilizing low-latency technologies such as Power Line Communication (PLC) and Ethernet, BANs are used for the purposes of monitoring and controlling heating, ventilation, and lighting systems, enabling the automation and optimization of energy usage within a building.

Neighborhood Area Networks connect HANs, BANs, and other local systems to centralized data concentrators and substations, enabling the aggregation of data from various smart meters and transferring it to utility companies for analysis and billing by utilizing technologies like ZigBee, Wi-Fi, and PLC. NANs can cover ranges of several square kilometers, supporting efficient data exchange required for applications such as outage management and power quality monitoring.

Field Area Networks serves as a bridge in the communication between distribution substations and control centers, enabling real-time monitoring and control of substations. FANs are used to support applications such as voltage control and fault isolation by utilizing a mixture of communication technologies, including WiMAX, LTE, and fiber optics, ensuring high reliability and coverage.

Wide Area Networks is the backbone communication infrastructure that connects power generation plants, substations, and utility control centers to ensure a seamless flow of information across the grid. To ensure efficient data exchange over large geographical areas, WANs utilize optical fiber and cellular technologies that provide low-latency and high bandwidth capabilities, which are essential for the long-distance transmission of critical grid data.

By employing modern communication technologies, smart grids address various requirements of monitoring, control, and data transmission, ensuring efficiency, scalability, and resilience.

3.1.4 Advanced Metering Infrastructure

AMI is an important part of modern smart grid networks, which enables real-time data exchange between customers and utility providers by providing a comprehensive framework for electricity consumption data collection, management, and analysis. AMI utilizes advanced communication technologies to ensure a two-way flow of information, enabling utility companies to optimize resource allocation while empowering customers to make informed decisions to manage their energy usage. In addition to the real-time data exchange, AMI contributes to improved billing, power quality monitoring, and outage detection, thus enhancing the operational efficiency of modern smart grids. With the evolvement of energy systems, AMI becomes a fundamental technology to support the development of a resilient and intelligent energy ecosystem.

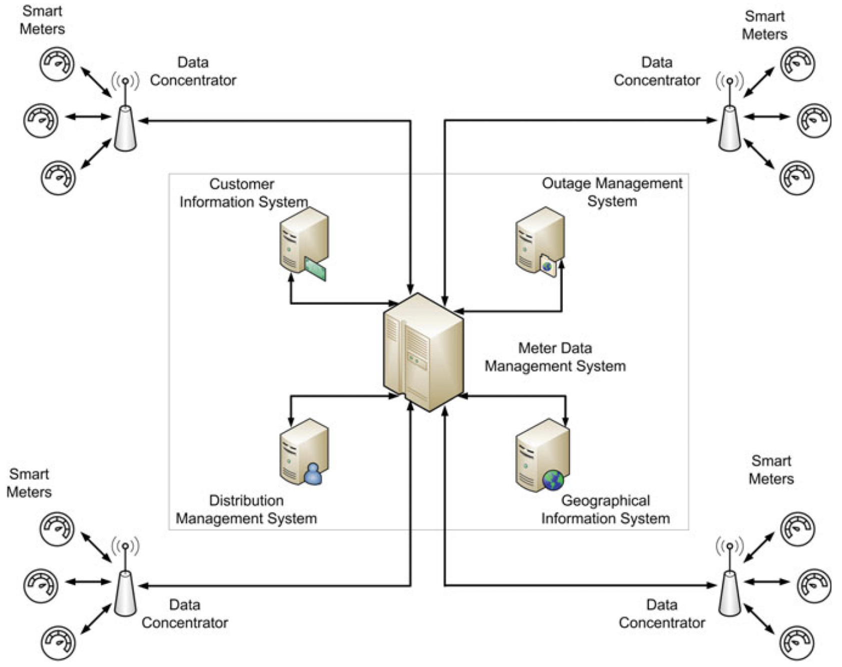

Figure 3.3: AMI infrastructure and its components [Source: [77]]

AMI represents a sophisticated infrastructure (Fig. 3.3), the purpose of which is to collect and transmit electricity consumption data, enabling utility companies to analyze and manage electricity usage more efficiently. The core elements of AMI include smart meters, communication networks, data concentrators, and Meter Data Management Systems (MDMSs). Smart meters, which are deployed at customer premises, measure real-time energy usage, enabling functionalities such as fault detection and time-of-use pricing. Communication networks connect smart meters to data concentrators by utilizing technologies like PLC and ZigBee, whereas data from multiple smart meters are aggregated by a data concentrator and transmitted to MDMS. MDMS is the central hub for processing and managing smart metering information. By integrating with utility management systems, including Customer Information Systems (CISs) and Outage Management Systems (OMSs), MDMS provides valuable insights and improves decision-making.

AMI supports services in several important areas, such as dynamic pricing models that enable customers to adjust their energy consumption patterns based on real-time electricity prices. Additionally, AMI contributes to the enhancement of smart grid reliability through automated fault detection and outage management. By identifying anomalies in energy usage profiles, AMI helps to reduce energy theft. The integration of advanced smart metering technologies with communication systems enables real-time monitoring of energy usage, fostering a more interactive and efficient energy ecosystem.

3.2 Foundations of Distributed Ledger Technology

Distributed Ledger Technology (DLT) has become a transformative technology, enabling a decentralized and tamper-proof way of managing data [78]. In a traditional centralized system, data is stored and controlled by a single authority, whereas DLT enables distributed storage of data across multiple nodes in a peer-to-peer network. Each participant (node) maintains a copy of the ledger to ensure that the system remains operational even in the event of a single node failure [79].

DLT provides a transparent and immutable record of transactions, where each transaction is validated by consensus among participants (nodes) before it is appended to the ledger, thus creating an irreversible chain of transactions. To safeguard the ledger and make unauthorized alterations computationally prohibitive, DLT utilizes advanced cryptographic techniques, such as a cryptographic hashing function SHA-256 that generates a fixed-size hash from input data. To create digital signatures that are used in the process of verification of transactions' authenticity without revealing private keys, techniques like Elliptic Curve Digital Signature Algorithm (ECDSA) are widely used among different DLT implementations. This technology is widely used across different domains, including finance, energy systems, supply chain management, and healthcare, to ensure trust, transparency, and accountability. One of the most recognizable implementations of DLT is Blockchain, where data is stored in sequentially linked blocks protected using advanced cryptographic techniques, making it highly resistant to fraud and single points of failure.

3.2.1 Blockchain Technology

Blockchain serves as a foundation for trustworthy, decentralized, and transparent transaction processing. Blockchain is a chain of blocks, where each block contains a set of transactions and a cryptographic reference to the previous block in the chain. This mechanism ensures data integrity and immutability because altering one block requires changes to all subsequent blocks, which is computationally infeasible.

In the Blockchain, multiple nodes participate in transaction validation and confirmation and maintain synchronized copies of the blockchain. Cryptographic techniques like SHA-256 and ECDSA are fundamental to verifying the integrity of data and ensuring trust among nodes (participants), whereas consensus mechanisms like Proof of Work (PoW) and Proof of Stake (POS) further enhance the robustness of blockchain by enabling decentralized agreement on the validity of transactions.

3.2.2 Essential Components of Blockchain

Blockchain implementation represents an integration of different components that are interconnected and work together. Nodes are individual devices, such as computers and servers, that store, validate, and propagate data across the blockchain network. Blocks in blockchain represent data structures that consist of a set of transactions, a timestamp, and a cryptographic hash reference to the previous block. Transactions in the blockchain are fundamental operations that represent the exchange of information or assets, and they are recorded on the blockchain after they are validated by nodes according to the logic of a consensus mechanism. Together, these components enable the core features of blockchain, namely data integrity, transparency, and decentralization.

Nodes represent critical components of the distributed network that enable decentralization in the blockchain. A node can be any device, such as a computer, server, or laptop, capable of running the blockchain's protocol software, thus participating in the blockchain network.

The most popular types of nodes that are used by different blockchain implementations can be categorized as follows:

- **Full nodes** are essential enablers of the data integrity and decentralization features in the blockchain. They maintain a complete copy of the ledger, validate new blocks according to the rules of a consensus protocol, and propagate new transactions across the blockchain network. By validating new blocks and ensuring compliance with consensus protocol, these nodes prevent fraud, such as double-spending. The most popular examples of full nodes are Bitcoin Core for the Bitcoin blockchain and Geth for the Ethereum blockchain.
- **Light nodes** rely on full nodes and require less storage and computational requirements to run, which is suitable for resource-constrained devices, such as Internet of Things (IoT) devices and smart meters. Instead of storing the entire copy of the blockchain, these nodes store only block header information, which makes them more resource-efficient.
- **Miner Nodes** perform computational work by solving cryptographic puzzles, which is specific to PoW consensus-based blockchains, such as Bitcoin. In addition, these nodes validate transactions, propose new blocks to be added to the blockchain, and compete for rewards.
- **Validator Nodes** responsible for validating the legitimacy of transactions and blocks submitted to the network. Usually, these nodes verify digital signatures and ensure that the sender of a transaction has sufficient balance (Bitcoin, Ethereum). Validator nodes are used by different consensus protocols, including POS, Delegated Proof of Stake (DPoS), and PBFT.

Blocks represent the core components of a blockchain that reliably store and organize data with integrity. Each block consists of three main elements, including the block header, transaction data, and Merkle root. The **block header** serves as a container for metadata information, such as a unique cryptographic hash, a hash reference to the previous block, a timestamp, and consensus protocol-specific metadata, such as the nonce in PoW-based blockchains. The **transaction data** consists of a set of validated transactions that may represent an exchange of assets, smart contract executions, or other activities. The **Merkle root** is a hash summary of all transactions in the block, which ensures data integrity and enables efficient transaction verification. Each block in the blockchain is linked to the previous block using cryptographic hashes, creating an immutable chain (Fig. 3.4), ensuring that altering one block would require recalculation of cryptographic hashes of all subsequent blocks, which is computationally impractical.

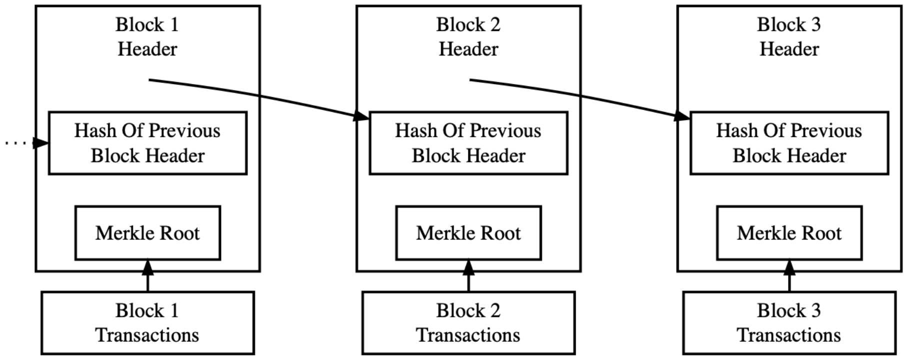

Figure 3.4: Blocks in a blockchain [Source: [80]]

The first block in the blockchain is the **genesis block** that serves as an immutable starting point of the ledger and does not reference a previous block. This block sets the foundational parameters of the blockchain, such as the initial hash, and may include the initial set of transactions. To ensure that all nodes can recognize and validate the genesis block, it is hardcoded into the blockchain implementation. Any changes to the genesis block would invalidate the entire chain, which highlights its critical role.

Transactions represent the fundamental operations of assets exchange across the decentralized network. A transaction is a digital record that contains essential parameters, such as the public keys of a sender and recipient, the value of an asset, and additional optional metadata. Each transaction is signed by the sender using his private key and cryptographic algorithms like ECDSA or Edwards-Curve Digital Signature Algorithm

(EdDSA), which ensures authenticity and prevents alterations. Blockchain network nodes use a cryptographic signature of a transaction to verify its integrity and the ownership of the sender.

Transactions are broadcast to the blockchain network and aggregated into blocks by miner or validator nodes, depending on the consensus protocol used by the network. Before adding a transaction to a block, it goes through the validation process, which includes checking for sufficient balance in the sender's account and conformance to consensus protocol rules. Transactions in blockchains like Ethereum may also include information regarding the execution of smart contracts or interaction with Decentralized applicationss (dAppss). Once a transaction is validated, confirmed, and added to the blockchain, it cannot be modified, highlighting the immutability feature of the blockchain.

Consensus mechanisms represent protocols that are used to reach an agreement among blockchain participants (nodes) on the state of the blockchain ledger, ensuring the integrity and protection of the chain. Consensus mechanisms are pivotal in decentralized systems, allowing participants to validate transactions, prevent double spending, and maintain trust. These protocols ensure that all nodes in the blockchain network reach an agreement, even in the presence of malicious nodes or network failures. While each consensus mechanism is designed with trade-offs in mind, balancing decentralization, robustness, and scalability, depending on the context, the most widely used protocols can be listed as follows:

- **Proof of Work** is a consensus mechanism used in the Bitcoin blockchain, where miner nodes compete with each other to solve computationally intensive cryptographic puzzles. If a node solves a computational puzzle, it gets the right to append a block to the blockchain and receives a reward. Although PoW is one of the consensus mechanisms with the highest resistance to attack, it is energy-intensive and has scalability issues.
- **Proof of Stake** is a consensus protocol where validator nodes are selected to validate and propose new blocks based on the amount of cryptocurrency they hold and are

willing to stake as collateral. Due to the fact that blocks are validated by a selected number of nodes, POS significantly reduces energy consumption and enables higher throughput compared to PoW.

- **Practical Byzantine Fault Tolerance** is a consensus protocol that is designed to ensure that participants in the blockchain network can reach an agreement even in the presence of malicious nodes. PBFT reaches consensus through three phases, where nodes exchange messages to validate and confirm a proposed transaction. This protocol requires low-latency communication between the nodes and is commonly used in permissioned blockchains, such as Hyperledger Fabric. PBFT can tolerate up to a third of malicious nodes, ensuring reliability as long as two-thirds of the nodes operate honestly.
- **Proof of Authority** is a consensus mechanism that relies on the reputation of selected validator nodes to validate transactions. This protocol does not require intensive computations or staking assets, whereas validator nodes are identified and pre-approved through their real-world identities, enabling high throughput. Proof-of-Authority (PoA) is widely used in permissioned blockchains where participants are known and trusted, such as consortiums of organizations.

3.2.3 Blockchain Types

By introducing a decentralized and immutable digital ledger, Blockchain has redefined the way data is stored, shared, and protected. To meet different requirements across different domains, industries, and applications, various blockchain types have emerged. These blockchain types can be listed as follows:

- **Public blockchains** represent decentralized and open networks where anyone with an internet connection can become a participant (node). These blockchains operate without a central authority, whereas PoW and POS are the most widely used consensus mechanisms by this type of blockchains. In a public blockchain, everyone can see all transactions and records, which makes it fully transparent. The most notable examples of public blockchains include Bitcoin, which uses PoW, and Ethereum that supports smart contracts and dAppss. Another example is Polkadot, designed to enhance interoperability between blockchains.

- **Private blockchains** are permissioned networks of nodes where access is given to known participants, making them ideal for businesses and consortiums of organizations that prioritize data confidentiality and efficiency. These blockchains are controlled by a single entity or a designated group of entities, enabling fast transaction processing. Private blockchains are widely used in industries such as finance, supply chain, and healthcare, where sensitive data needs protection and only designated participants can join the network. The most notable example of a private blockchain is Hyperledger Fabric, which is used by enterprises for various business applications. Corda has been designed for financial institutions, enabling reliable and direct transaction flow between parties while maintaining privacy.

- **Consortium blockchains** are semi-decentralized blockchains governed by a group of organizations that can access or validate transactions. These blockchains enable enhanced robustness, scalability, and efficiency, which makes them an ideal solution for industries requiring collaboration and shared governance, such as banking, supply chain management, and healthcare. Some of the popular examples of consortium blockchains include R3 Corda, which is used in banking for efficient financial transactions, and Quorum, developed by JPMorgan for high-speed financial processing.

- **Hybrid blockchains** offer a solution for business cases where both transparency and confidentiality are required by combining features of public and private blockchains. In hybrid blockchains, some data and transactions can be publicly accessible, whereas other information is kept private, which enables organizations to comply with regulations while leveraging the transparency of blockchain technology. These blockchains can connect with public blockchains, which enables seamless communication and data sharing across different ecosystems. Some examples of hybrid blockchains include Dragonchain, which is used for enterprise applications with customizable privacy settings, and XDC Network, which supports selective data exposure and is used in finance applications.

3.2.4 Smart Contracts

Smart Contracts have become a core innovation in the blockchain domain, combining decentralized networks, programmable logic, and cryptography. These contracts represent software programs stored on a blockchain, such as Ethereum or Binance Smart Chain (BSC), and enable the automatic execution of predefined actions (blocks of code) when specific conditions are met. Different blockchain implementations utilize different programming languages to support smart contracts, such as Solidity and Vyper, which are used in the Ethereum blockchain [81], and Rust in the Solana network.

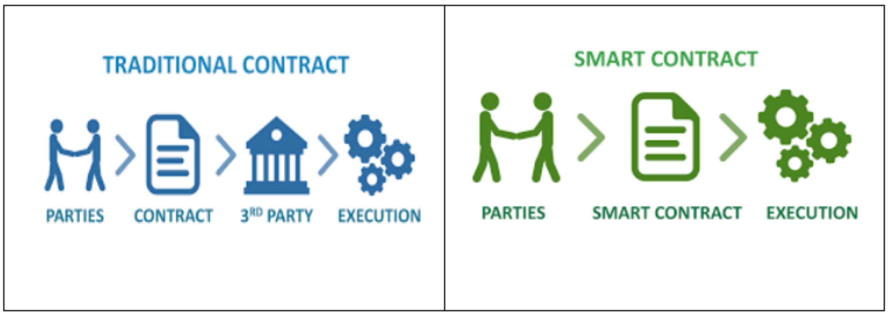

Figure 3.5: Traditional contracts versus Smart Contracts

The main idea of smart contracts is to eliminate the requirement for intermediaries, which enables reduced costs, speeds up the process of transaction execution, and enhances trust. The difference between a traditional contract and a smart contract is shown in Fig. 3.5. Smart contracts support a wide range of features, such as access control, asset transfer, and event-triggered payments, becoming a crucial enabler in dAppss world.

Blockhain's features, such as immutability, cryptographic protection, and consensus protocols, make smart contracts resistant to tampering, fraud activities, and unauthorized alterations. Most of the existing blockchain platforms employ rigorous testing and deployment processes before a smart contract is executed, which enables the identification and mitigation of vulnerabilities in smart contracts. To ensure that the outcome of a smart contract's execution is predictable and consistent across all nodes in the network, blockchain platforms enforce deterministic execution of contracts, which means that given the same input and network state, the contract will produce the same output, regardless of where or when it is executed. Since all blockchain participants must agree on the smart contract's result before appending transactions to the blockchain, deterministic execution is a required feature to maintaining integrity and consensus. At the same time, smart contracts can be vulnerable to coding errors or exploits and must be properly designed and tested before being written to the blockchain.

Real-world examples highlight the potential of smart contracts. A Decentralized Exchange (DEX), Uniswap, utilizes smart contracts for automated cryptocurrency trading, eliminating the need for intermediaries. One example in the supply chain domain is IBM Food Trust, which uses smart contracts to trace food from a farm to a store, thus ensuring transparency. One of the most popular Decentralized finance (DeFi) platforms, Aave, uses smart contracts to enable automated lending and borrowing, eliminating the need for banks. The growing adoption of smart contracts across different industries underscores their importance in technology-driven innovation.

3.2.5 On-chain and Off-chain Storage

On-chain type of storage refers to storing data (transactions and blocks) directly on the blockchain. Transactions are added to blocks that are sequentially stored on the blockchain, which become an immutable part of the ledger, ensuring transparency, decentralization, and data integrity. Data that is stored on-chain represents critical and sensitive information, such as smart contract codes, account balances, or transaction records. Although storing data on-chain guarantees data integrity, it becomes expensive due to the blockchain scalability and cost. One of the blockchain implementations, Ethereum, introduced fees (gas fees) for transaction execution based on the data size (size of a smart contract) to be processed, making large-scale data storage impractical for cost-sensitive applications. Moreover, poorly designed and written smart contracts may lead to increased gas fees, affecting the cost of executing transactions. Due to the fact that every participant (node) in the blockchain network must replicate and store on-chain data, thus maintaining a relevant copy of the ledger, performance trade-offs, such as transaction speed and storage demand, should be carefully considered.

On the other hand, in the off-chain approach, data are stored in external systems outside the blockchain, such as cloud, centralized servers, databases, or decentralized file storage mechanisms like Interplanetary File System (IPFS) and Filecoint. Storing data off-chain is suitable for large volumes of data, such as images, videos, or documents, which do not require direct verification. One of the possible ways to maintain a link between on-chain and off-chain data is to use a cryptographic reference (hash) to an off-chain file, whereas the reference is included in a transaction that is stored on-chain. The use of off-chain storage enhances the scalability, cost-efficiency, and flexibility of applications, whereas data availability, trust, and data integrity must be carefully considered when using off-chain storage.

3.2.6 Blockchain Platforms

Various blockchain implementations have been developed for different business applications, ranging from finance to supply chain management. These platforms differ in their architectures, use different consensus mechanisms, and enable different functionalities, making them suitable for specific application domains. This section explores the most notable blockchain platforms, their applications, and challenges.

Bitcoin is the first blockchain platform that serves as a peer-to-peer digital currency and is mainly used for digital payments. The decentralized nature of Bitcoin makes it an alternative solution for remittances, especially in geographical areas with unstable financial systems. Bitcoin lacks native support for smart contracts and dAppss, which limits its functionality to only the transfer of cryptocurrency between participants. Bitcoin uses PoW energy-expensive consensus protocol and suffers from scalability issues, where the transaction throughput is around seven transactions per second.

Ethereum is a blockchain platform that gained its popularity by introducing smart contracts functionality, which enables software developers to build and deploy dAppss. The most popular applications utilizing the Ethereum platform include DeFi systems like Uniswap and Aave, Non-Fungible Tokens (NFTss) for digital ownership, and Decentralized Autonomous Organizations (DAOss). Ethereum has its own cryptocurrency, Ether (ETH), which is used as a currency for digital payments and as ”gas” to cover fees for transaction execution. Despite its benefits, Ethereum has its own disadvantages, such as high transaction fees and high entry barriers for software developers due to the complexity of the Solidity language.

Binance Smart Chain is a high-performance blockchain platform optimized for dAppss, DeFi, and decentralized cryptocurrency exchanges. This platform is Ethereum Virtual Machine (EVM) compatible, enabling seamless migration of dAppss from Ethereum to BSC. The most popular use cases of the BSC include decentralized exchanges of cryptocurrencies (Binance), staking platforms, and NFTss marketplaces. Despite its advantages, BSC is controlled by the Binance decentralized exchange, raising concerns about robustness and decentralization. In addition, BSC uses Proof-of-Staked Authority (PoSA) consensus protocol, prioritizing speed and cost efficiency while compromising decentralization.

Hyperledger Fabric is a permissioned blockchain platform that is widely used by consortiums of organizations, ensuring that only authorized participants can access the system. Fabric supports private channels only accessible by defined parties, which enables protected and confidential data sharing between those parties. Fabric uses a a pluggable consensus protocol, which can be customized by participating organizations based on their specific needs. By default, Fabric uses the Raft consensus mechanism, where a leader node is responsible for managing the order of transactions while follower nodes validate and replicate transactions. As a permissioned blockchain, Fabric lacks common features of public blockchains, such as decentralization, and requires substantial technical expertise to set up and maintain the network.

3.3 Differential Privacy and Local Differential Privacy Concepts

In today's world, a lot of personal data are collected and analyzed to improve services, make decisions, and drive innovation. However, this reliance on data raises concerns about privacy and the potential misuse of sensitive information [82]. DP offers a solution to analyze datasets without exposing individual identities. Instead of simply removing names or other identifiers, DP ensures that even if someone's data is included or excluded, the overall results remain nearly the same [83]. This means that individuals can contribute to data-driven studies, such as medical research or public policy analysis, without fear that their personal information could be leaked or traced back to them. By injecting controllable noise into the data, DP provides a strong mathematical guarantee of privacy, balancing the need for useful insights with the protection of personal data [17]. As governments and companies increasingly adopt DP, it represents a major step toward ethical and privacy-preserving data analysis.

3.3.1 Differential Privacy

DP is a mathematical framework designed to ensure that the inclusion or exclusion of any single individual's data in a dataset does not significantly affect the output of a computation. This concept is crucial in modern data privacy, particularly when analyzing sensitive information while maintaining end-users' confidentiality. Introduced by Dwork *et al.* in 2006 [84], differential privacy provides a formalized approach to quantifying privacy risks and mitigating them through controlled noise addition. By incorporating carefully calibrated randomness into computations, DP enables the extraction of valuable insights from data while offering rigorous privacy guarantees.

At its core, differential privacy is defined using the privacy parameter ϵ called privacy budget, which controls the level of privacy protection. A lower ϵ value signifies stronger privacy, meaning that an individual's participation in a dataset has minimal influence on the results. In contrast, a higher privacy budget value allows for more accurate statistical analyses but offers weaker privacy guarantees. The trade-off between privacy and utility is central to DP and requires careful tuning to balance data utility and privacy. DP mechanism is often implemented through techniques such as the Laplace mechanism or randomized response, each of which introduces statistical noise in different ways to achieve the desired privacy guarantees.

One of the defining features of DP is its robustness against various attacks. Unlike traditional privacy-preserving techniques such as anonymization, which are vulnerable to re-identification, DP ensures that even an adversary with additional background knowledge cannot infer sensitive details about individuals with high confidence. This resilience makes DP particularly suitable for applications in machine learning, statistical analysis, and database queries, where data privacy is paramount. As differential privacy gains traction, its adoption is expanding across various domains, including smart grids, healthcare, finance, and government data releases.

The data publishing and analysis process is presented in Fig. 3.6 and consists of two stages. In the first stage, the data collection stage, individuals submit their personal information to a data curator. The second stage is the data publishing or analysis stage, where the dataset is shared with public users. Both data publishing and analysis bring social benefits, such as providing better services, publishing official statistics, providing data mining or machine learning tasks.

Most of the collected datasets are personally related and may contain private or sensitive information. Privacy violations may occur in both stages. If a data curator is untrustworthy, personal information will be disclosed directly in the data collection stage. Even though

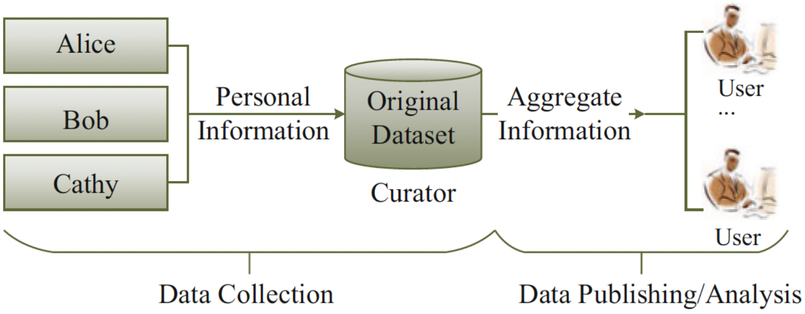

Figure 3.6: Privacy-preserving data publishing and analysis [Source: [85]]

a data curator can be trusted and applies several simple anonymization techniques, when it publishes aggregate information to the public for analysis, personal information may be disclosed as the public is usually not trustworthy.

A basic attack model is shown in Fig. 3.7. A data curator aims to preserve the privacy of n records in an original dataset D, whereas an adversary has background information about $n-1$ records except the x_n record. An adversary makes a query on the original dataset to get aggregated information about n records, whereas by comparing the difference between the query result with the background information, the information of record x_n can be easily identified by the adversary.

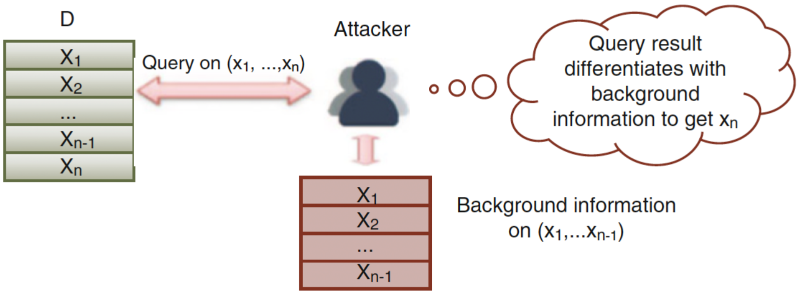
D
X1
X2
...
Xn-1
Xn
Query on (x1, ...,xn)
Attacker
Query result differentiates with background information to get xn
X1
X2
...
Xn-1
Background information on (x1,...xn-1)

Figure 3.7: Basic attack model [Source: [85]]

DP can be used to preserve privacy and resist the attack by not revealing too much information about any individual record in the dataset (Fig. 3.8). Assume that two datasets differ in only one record (neighboring datasets), and by performing the query on both datasets, an adversary will get the same result with a very high probability and cannot identify which dataset contains an extra record. If the difference cannot be determined, the privacy of a particular record is preserved. If the same applies to all records in the original dataset, then the privacy of all records is preserved.

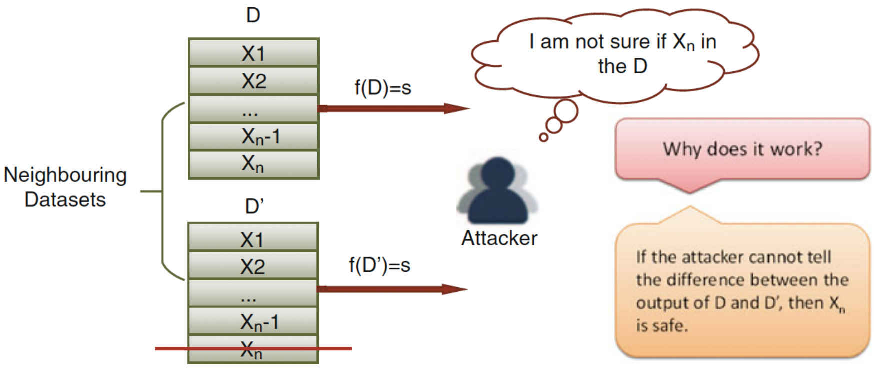
D
X1
X2
...
Xn-1
Xn
Neighbouring Datasets
D'
X1
X2
...
Xn-1
~~Xn~~
f(D)=s
f(D')=s
I am not sure if Xn in the D
Attacker
Why does it work?
If the attacker cannot tell the difference between the output of D and D', then Xn is safe.

Figure 3.8: Using Differential Privacy to preserve privacy of individuals [Source: [85]]

3.3.2 Local Differential Privacy

In a world where personal data fuels technological progress, maintaining individual privacy while extracting meaningful insights is a challenge. LDP has emerged as a promising solution, ensuring that sensitive user information remains protected before sharing. Unlike traditional privacy-preserving methods that rely on a centralized authority (a data curator) to safeguard data, LDP shifts the privacy protection to the user's side. This means that data is protected before leaving the individual's control, making it impossible to trace it back to any specific person. LDP has gained widespread adoption, particularly in large-scale applications where user trust is paramount, such as data analytics, recommendation systems, and mobile data collection.

Traditionally, DP is implemented in a centralized fashion, where a trusted data curator collects raw data from end-users and applies a noise injection mechanism before publishing aggregated results. This approach ensures that the presence or absence of any single individual's data does not significantly influence the overall outcome, making it difficult for adversaries to infer personal information. However, trusting a central authority to handle raw data confidentially is not always feasible. This limitation has given rise to the Local Differential Privacy approach, which eliminates the need for a trusted curator by ensuring that privacy mechanisms are applied on the end-users' side before sharing the data.

LDP operates under a fundamentally different paradigm from standard DP. Instead of collecting raw data and injecting noise at the data curator's side, LDP introduces noise at the individual level (Fig. 3.9). Each end-user perturbs their sensitive data locally before sharing it with a third party, ensuring that even if the data is intercepted or stored improperly, no meaningful personal information can be extracted. This approach provides a stronger privacy guarantee, as no entity has access to the raw data. On the other hand, enhanced privacy protection comes with an additional cost. Due to the fact that the noise is introduced at the individual level, the accuracy of aggregated results can be affected.

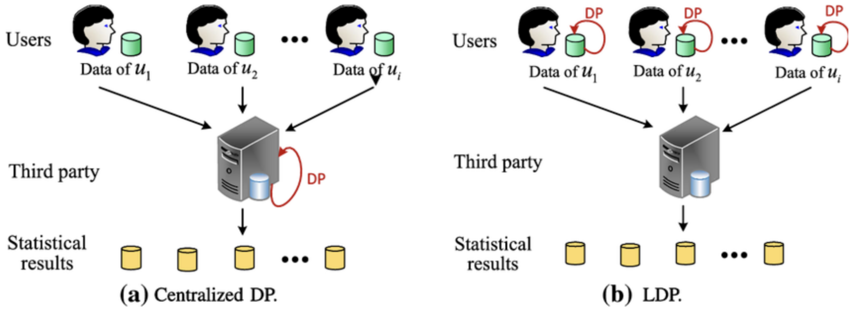

Figure 3.9: Differential Privacy (**a**) and Local Differential Privacy (**b**) approaches [Source: [86]]

One of the key differences between DP and LDP is the trust model. In the traditional DP approach, a centralized data curator is assumed to be trustworthy and responsible for applying the privacy-preserving mechanisms, namely injecting noise into the aggregated data. This assumption may not hold in many applications where organizations collect data from a considerable number of users without the ability to store raw personal information. LDP addresses this challenge by shifting the responsibility of privacy protection to the data owner, ensuring that privacy is maintained even in the presence of an untrusted data curator. This makes LDP particularly useful in large applications like smart grids, web browsing analytics, and personalized recommendations, where data is collected from a vast number of users.

3.4 Foundations of Federated Learning

In today's digital world, vast amounts of data are generated daily by individuals, businesses, and connected devices. From smart meters to healthcare records and financial transactions, this data holds valuable insights that can drive innovation and improve services. However, the collection of such data poses concerns about user privacy, protection, and regulatory compliance. Traditional approaches, including machine learning, often require gathering all data in a centralized location, raising risks of data breaches, unauthorized access, and regulatory violations. These concerns have led to an urgent need for new approaches that allow service providers to leverage data-driven insights without compromising individual privacy. One of the most notable approaches that has gained significant attention is FL, a decentralized method of training machine learning models while keeping data on end-users' devices.

Federated Learning represents an innovative approach to machine learning that redefines how data is processed and utilized. Unlike traditional centralized models that require collecting data in a single location for training, FL enables multiple devices or entities to collaboratively train a machine learning model while keeping end-users data on their devices. Thus, instead of transferring sensitive information to a central server, the machine learning process occurs directly on local devices, with only model updates being shared between a central server and clients. This method offers significant advantages in terms of data confidentiality, protection, and efficiency. The growing interest in Federated

Learning technology is due to the increasing concerns about data privacy and regulatory requirements that restrict data sharing. By empowering end-users to contribute to model improvements without exposing their raw data, FL aligns with global data protection regulations, such as the GDPR.

Federated Learning aims to build a joint machine learning model based on the data of several end-users (devices). In the process of training a global machine learning model, participating parties share only the model's updates with a central server, thus not revealing any protected private portions of the data of each participant. Federated Learning is a machine learning framework for building machine learning models that can be characterized by the following features:

- There are at least two participants jointly training a machine learning model. Each party holds some data and contributes it to the training of the global model.
- In the training process, confidential data never leaves the participant's system (device), thus preserving their privacy.
- The accuracy of the global machine learning model is close to the model trained with all data transferred to a central server.

An example of a Federated Learning system is shown in Fig. 3.10. A central server distributes an initial machine learning model to all participants (data owners). Each participant trains a model using their private dataset and sends the model updates to the central server (aggregator). When a central server receives model updates from all participants, it combines all updates and produces a new version of the global model, which is then distributed to all participants. This process continues until a certain level of accuracy or the maximum number of training iterations is reached. In this approach, a private dataset of each participant is never exposed to other parties. This ensures the privacy of participants and reduces the communication overhead needed to share heavy datasets.

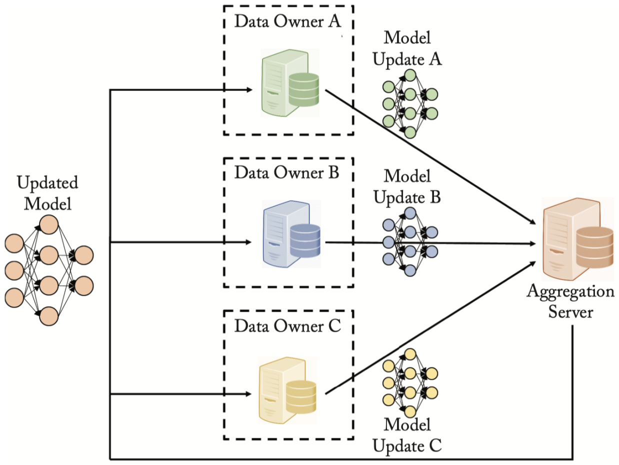

Figure 3.10: An example of a Federated Learning system [Source: [87]]

Beyond privacy, FL also improves efficiency by leveraging the power of edge computing. Instead of relying on centralized infrastructure, FL utilizes the processing power of distributed devices, reducing bandwidth consumption and enabling real-time model updates. This is particularly useful for applications like demand forecasting and financial fraud detection, where timely insights are crucial. Moreover, FL fosters collaborative intelligence by allowing multiple parties to contribute to a shared machine learning model without exposing their private data. FL opens doors for scalable and privacy-preserving machine learning, making it a valuable technology for future applications across different domains.

Federated learning and differential privacy can be effectively combined to mitigate membership inference attacks, where an attacker aims to determine whether a specific data point was used in the training of a machine learning model. While federated learning enhances privacy by keeping private datasets on local devices and only sharing the updates to the machine learning model with the central server, adversaries can still analyze the model updates to infer sensitive information. Differential privacy could be used to prevent attackers from inferring sensitive information by analyzing model updates. By injecting carefully calibrated noise into the model updates before they are shared with the central server, it limits the amount of information any single data point contributes to the model. Thus, the combination of both techniques provides stronger privacy guarantees than federated learning alone.

3.5 Summary

This chapter provided an overview of the key technologies that form the basis of this research. Initially, the foundations of the smart grid ecosystem were presented, including its conceptual model, essential components, and communication mechanisms. Next, the fundamentals of distributed ledger technologies, with a particular focus on blockchain, its architecture, and storage approaches, were outlined. Moreover, this chapter covered theoretical aspects of differential and local differential privacy as critical techniques for safeguarding user data. Lastly, the foundational elements of federated learning, highlighting its relevance for collaborative, privacy-preserving model training across distributed systems, were presented. These foundational elements collectively support the design and implementation strategies outlined in the subsequent sections of this thesis.

Chapter 4

Energy Trading Scheme for a Microgrid of prosumers with photovoltaic panels

4.1 Overview

Due to the increasing number of end-users, as well as the growing number of appliances used, the demand for electricity is constantly growing. The global demand for electricity is predicted to rise to 40,000 TWh by 2040 [88], which leads to an increase in energy generation on the utility side, where fossil fuels are the main sources of electricity production. For example, 68% of electricity was generated using coal-fired energy generation in Q1 2020 in China, whereas only 28% using RES. The same situation is observed in India, where 77.5% of electricity was produced using coal in Q1 2020. On the other hand, in the US, about 40% of electricity was generated by gas, while in the European Union, around 40% was supplied by renewables [89].

Due to high electricity prices and high penetration of DER, some consumers utilize PV systems to provide their homes with electricity. Thus, ordinary consumers who can produce and consume electricity become prosumers. When a prosumer produces more electricity than it needs to cover its demand, excess energy is sold to the utility grid at a lower price compared with the buying price at which ordinary consumers within a neighboring area buy energy from the utility grid.

With recent advancements in Information and Communication Technologies (ICT) and massive deployment of DER, the concept of a microgrid has emerged as an alternative solution for coordinating local groups of prosumers. Microgrids bring the possibilities to increase efficient electricity utilization, which leads a reduced amount of energy imported from the utility grid. To achieve this, prosumers must be motivated to share energy with each other in the microgrid. Thus, pricing is one of the most important mechanisms for motivating prosumers to interact with each other in the microgrid.

Some works [28], [90] have approached pricing formation in the microgrid in terms of optimizing prosumers' behavior, which is possible but unlikely in the real-world scenario. Due to the random prosumers' behavior, it becomes challenging to design a pricing model that motivates prosumers to share energy with each other and contributes to a decrease in energy imported from the utility grid. A new pricing model should provide more beneficial buying and selling prices of energy compared with the utility grid that will motivate prosumers not to sell excess energy to the utility grid and buy energy from the microgrid. Thus, the energy produced in the microgrid can be utilized more efficiently, which in turn will contribute to the reduction in the demand from the utility grid.

In this chapter, a novel dynamic pricing model for a microgrid of prosumers is proposed, which is able to provide lower buying and higher selling prices for electricity compared to the utility grid that decreases the demand from the utility grid, as well as allows to reduce the energy usage cost for a prosumer.

The primary design goal of the proposed pricing model is economic efficiency, aiming to reduce energy costs for prosumers while increasing the efficiency of local energy utilization. Additionally, the scheme aims to reduce dependency on the main utility grid by incentivizing local energy trading within a microgrid. To demonstrate the achievement of these goals, the evaluation framework utilizes real-world energy datasets to simulate various trading scenarios. The evaluation compares the proposed model against baseline utility grid trading and existing approaches from the literature. Specifically, the achievement of economic efficiency is measured by the percentage reduction in total energy usage costs and the increase in prosumer profits.

4.2 Proposed Model

In this section, a novel dynamic pricing model for a microgrid of prosumers with photovoltaic systems is proposed.

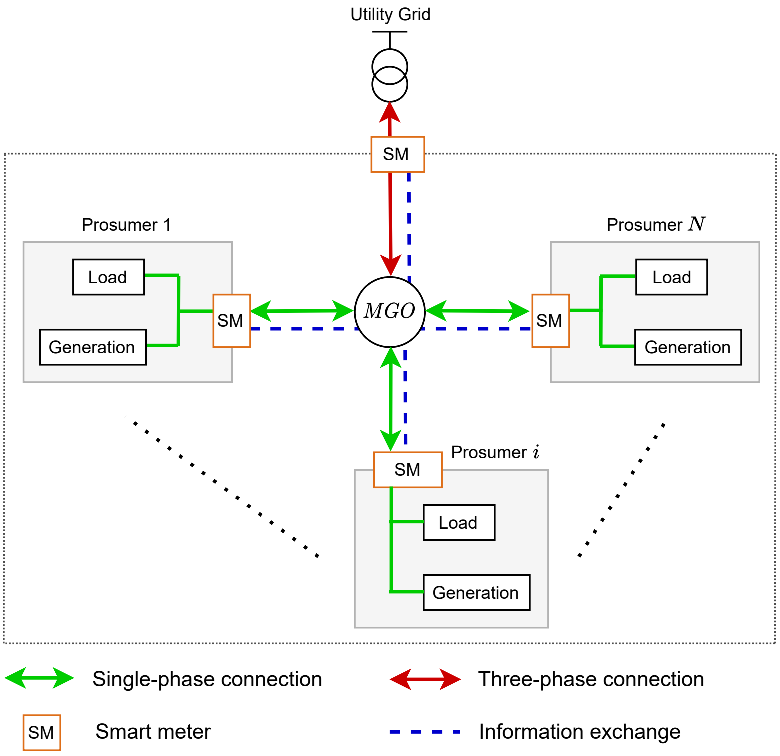

Figure 4.1: System model of a microgrid of prosumers

4.2.1 System Model

Fig. 4.1 shows the structure of a microgrid consisting of N prosumers. Each end-user is expected to consume (buy) and produce (sell) energy. Thereby all end-users in a microgrid can be represented as prosumers, whereas the only difference between them is the energy usage profiles. Prosumers are connected to the Microgrid Operator (MGO) using a single-phase electricity connection, whereas the electricity connection between the MGO and the utility grid is three-phase. It should be noted that the type of the connection does not affect the overall performance of the proposed pricing model. However, the relevant assumption is that the connection between the households and the MGO is limited to 4 kW, which is a level of electricity consumption for a typical household. Thus, the MGO operates as an aggregator, which means all prosumers can sell excess energy to the MGO and buy the shortage from the MGO. All prosumers communicate with the MGO through a bidirectional communication channel, which is mainly used to submit energy usage data to the MGO. A smart meter, which is deployed at each prosumer, measures the amount of energy produced and consumed and sends the corresponding information to the MGO. In turn, the MGO aggregates energy usage data and then calculates the cost of electricity used for each prosumer.

In the proposed model, it is assumed that there is a MGO located within a microgrid that aggregates and dispatches energy within the renewable energy community (a microgrid). First, the energy consumption and production profiles are predicted for all participants in a microgrid, which gives the information on total energy consumption and production in a particular time slot. Thus, the MGO provides an independent aggregator [91] with the information whether it will require some energy to be supplied to the microgrid or whether it will be able to supply some energy to the aggregator in a particular time slot. The MGO, which is an aggregator itself, may exchange energy with other aggregators outside the microgrid according to the EU regulations [92].

It should be emphasized that the interaction with the utility grid and the terms import/export and buy/sell from the utility grid are considered from an economic point of view only. For instance, import from the utility grid means electricity was bought at the utility grid's selling price. The same applies to export to the utility grid. Whereas the interaction with the microgrid means that the prosumer buys/sells at the microgrid's prices. It should be noted that irrespective of whether the electricity is bought/sold from/into the utility grid or the microgrid, for the same demand and generation profiles the physical energy flow remains the same.

The main aim of this work is to design a pricing model for a microgrid of prosumers. All other aspects, including smart meters operation, communication between prosumers and the MGO are beyond the scope of this work.

In this work, a microgrid consists of N prosumers. Let $\mathcal{N} = \{1, 2, 3, ..., N\}$ denote the set of prosumers in the microgrid, where n is the prosumer index and $n \in \mathcal{N}$, whereas, the total number of prosumers is given by $N \triangleq |\mathcal{N}|$. Energy usage cost for all prosumers is calculated at the end of a time slot, where each time slot is of one hour. Let $\mathcal{H} = \{1, 2, 3, ..., H\}$ denote the set of all time slots, where h is the time slot index and $h \in \mathcal{H}$, whereas, the total number of time slots is given by $H \triangleq |\mathcal{H}| = 24$. Let QD_{P_n} denote the predicted energy consumption profile for a prosumer n for one day, and it is defined as follows:

$$QD_{P_n} = \{QD_{P_{n_1}}, QD_{P_{n_2}}, ..., QD_{P_{n_H}}\}, \quad n \in \mathcal{N} \tag{4.1}$$

where $QD_{P_{n_h}}$ is the predicted level of energy consumption for a prosumer n in a time slot h and $QD_{P_{n_h}} \in QD_{P_n}$. Let QS_{P_n} denote the predicted energy production profile for a prosumer n for one day, and it is defined as follows:

$$QS_{P_n} = \{QS_{P_{n_1}}, QS_{P_{n_2}}, ..., QS_{P_{n_H}}\}, \quad n \in \mathcal{N} \tag{4.2}$$

where $QS_{P_{n_h}}$ is the predicted level of energy production for a prosumer n in a time slot h and $QS_{P_{n_h}} \in QS_{P_n}$. Smart meters may submit the prosumers' energy usage data to the MGO at different time intervals (5 min / 10 min / 15 min). Let QD_{R_n} denote the actual energy consumption profile for a prosumer n for one day, and it is defined as follows:

$$QD_{R_n} = \{QD_{R_{n_1}}, QD_{R_{n_2}}, ..., QD_{R_{n_H}}\}, \quad n \in \mathcal{N} \tag{4.3}$$

where $QD_{R_{n_h}}$ is the actual level of energy consumption for a prosumer n in a time slot h and $QD_{R_{n_h}} \in QD_{R_n}$. Let QS_{R_n} denote the actual energy production profile for a prosumer n for one day, and it is defined as follows:

$$QS_{R_n} = \{QS_{R_{n_1}}, QS_{R_{n_2}}, ..., QS_{R_{n_H}}\}, \quad n \in \mathcal{N} \tag{4.4}$$

where $QS_{R_{n_h}}$ is the actual level of energy production for a prosumer n in a time slot h and

$QS_{R_{n_h}} \in QS_{R_n}$. It should be noted that the prosumers' energy usage profiles are not the same, rather they are selected randomly from publicly available datasets [93, 94].

The first priority for prosumers is self-consumption; thus, each prosumer consumes its own PV energy to cover its demand. If there is not enough PV energy produced on the prosumer's side, the rest is bought from the MGO. On the other hand, if a prosumer can cover its demand by consuming its own PV energy and has excess energy to share, the surplus is sold to the MGO. Let NP_{n_h} denote the net power for a prosumer n in a time slot h, and it is defined as follows:

$$NP_{n_h} = QS_{R_{n_h}} - QD_{R_{n_h}} \tag{4.5}$$

Let QS_h^T denote the total excess energy that is available in a microgrid after self-consumption in a time slot h, and it is calculated as follows:

$$QS_h^T = \sum_{n=1}^{N} NP_{n_h}, \quad NP_{n_h} > 0 \tag{4.6}$$

Let QD_h^T denote the total shortage of energy in a microgrid after self-consumption in a time slot h, and it is calculated as follows:

$$QD_h^T = \sum_{n=1}^{N} |NP_{n_h}|, \quad NP_{n_h} < 0 \tag{4.7}$$

where NP_{n_h} is the net power for a prosumer n in a time slot h.

After self-consumption, if the total shortage of energy is greater than the total excess energy, the rest is bought (imported) from the utility grid. On the other hand, excess energy is sold (exported) to the utility grid. Hence, the state of the microgrid can be identified as the ratio of the total excess energy to the total shortage of energy. Let SDR_h denote the Supply and Demand Ratio (SDR) in a time slot h. Combining (4.6) and (4.7), SDR in a time slot h is calculated as follows:

$$SDR_h = \frac{QS_h^T}{QD_h^T}, \quad \text{where} \quad QD_h^T > 0 \tag{4.8}$$

where QS_h^T is the total excess energy available in the microgrid after self-consumption in a time slot h. Thus, when $SDR_h = 1$, the microgrid is in the state "Standalone", and there is no interaction with the utility grid. If $SDR_h > 1$, the microgrid operates in the state "Seller", which means that excess energy is sold to the utility grid. On the other hand, when $SDR_h < 1$, the microgrid operates in the state "Buyer", and the shortage is bought from the utility grid. The case when there is no demand ($QD_h^T = 0$) and all the produced energy in the microgrid is sold to the utility grid is not considered in this work; thus, in (4.8), the only case when $QD_h^T > 0$ is considered.

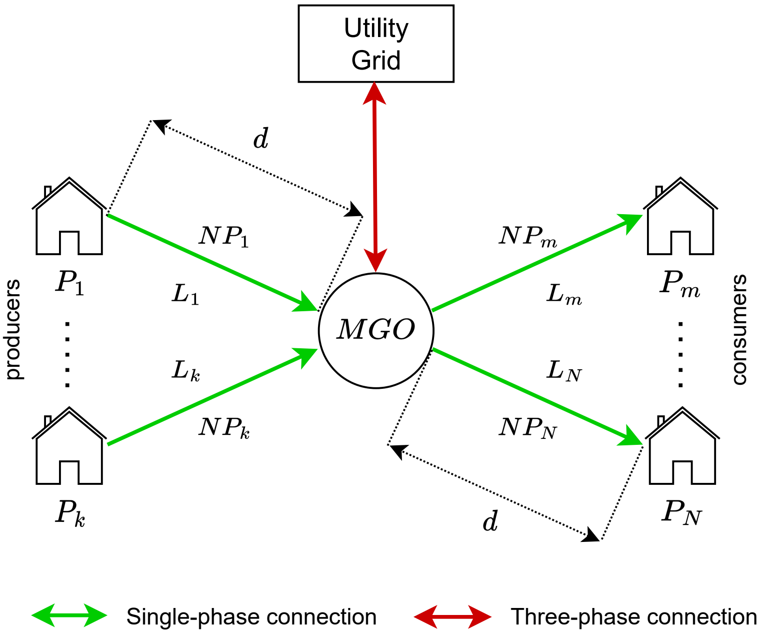

Figure 4.2: Energy transfer loss model

In this work, the energy transfer loss is taken into account. Fig. 4.2 shows that in each time slot prosumers are divided into two categories, namely producers and consumers. During the transfer of energy from a producer to a consumer (between prosumers), the energy transfer loss is unavoidable. For simplicity, all the prosumers are located on the same distance from the MGO, which is defined as d. Moreover, all prosumers are connected to the MGO using a single-phase electricity connection and the same type of wire, the resistance of which is defined as ρ. Let k_L denote the energy transfer loss coefficient, and it is calculated as follows:

$$k_L = \frac{\rho * d}{V} \tag{4.9}$$

Let L_{n_h} denote the energy transfer loss during the transfer of energy (NP_{n_h}) between a prosumer n and the MGO in a time slot h, and it is calculated as follows:

$$L_{n_h} = k_L * NP^2_{n_h} \tag{4.10}$$

Let λ_b denote the buying price of energy from the utility grid, whereas, λ_s denote the selling price of energy to the utility grid. However, the prices λ_b, λ_s may vary over time. Let Λ_b denote the set of buying prices of energy from the utility grid for one day, and it is defined as follows:

$$\Lambda_b = \{\lambda_{b_1}, \lambda_{b_2}, \lambda_{b_3}, ..., \lambda_{b_h}\} \tag{4.11}$$

where λ_{b_h} is the buying price of energy from the utility grid in a time slot h and $\lambda_{b_h} \in \Lambda_b$. Let Λ_s denote the set of selling prices of energy to the utility grid for one day, and it is defined as follows:

$$\Lambda_s = \{\lambda_{s_1}, \lambda_{s_2}, \lambda_{s_3}, ..., \lambda_{s_h}\} \tag{4.12}$$

where λ_{s_h} is the selling price of energy to the utility grid in a time slot h and $\lambda_{s_h} \in \Lambda_s$.

The overall methodology of the proposed approach is as follows. Smart meters measure the energy consumption and production of a prosumer and send energy usage data to the MGO at a time interval that is equal to one hour in this work, whereas the predicted energy usage profiles for all prosumers are predicted by the MGO. At the end of each hour, the MGO aggregates energy usage data sent by prosumers to calculate the supply and demand ratio (SDR_h). Furthermore, the process of calculating the energy usage cost and profit for all prosumers is based on the MGO's side.

4.2.2 Assumptions

There are six major assumptions in this work: (i) To simplify the energy transfer loss calculation and allocation, the distance between any prosumer and the MGO is taken as a constant, which is equal to 100 m ($d = 100\ m$). (ii) The energy transfer loss when interacting with the utility grid is not considered, as these losses have already been accounted for in the utility grid's prices. (iii) All prosumers are connected to the MGO using a single-phase electricity connection with the same type of wire; thus, the electrical wire resistance (ρ) is the same for all single-phase connections ($\rho = 0.01\ Ohm/m$) between prosumers and the MGO. (iv) The MGO predicts energy usage profiles for each prosumer n for each time slot h. Although an energy usage forecasting mechanism is beyond the scope of this part of the work, the dependency of the prosumer's energy usage cost on it's absolute deviation from the predicted energy consumption is analyzed in Section 4.3. (v) Since the data aggregation mechanism is beyond the scope of this part of the research, it is assumed that energy usage data of each prosumer n are received from smart meters and aggregated at the MGO during each time slot h, so that actual energy usage profiles of all prosumers are available at the end of each time slot. (vi) There is a continuous supply of energy from the utility grid so that the shortage can be bought from the utility grid at any time, as well as excess energy can be sold to the utility grid at any time.

4.2.3 Pricing Model

In this section, a novel pricing model for the microgrid is proposed to reduce the energy usage cost for all prosumers and to increase the profit from selling energy, which increases efficient energy utilization in the microgrid and contributes to the reduction in the demand from the utility grid. Instead of selling energy to the utility grid, excess energy can be utilized in the microgrid to cover some part of the microgrid's demand. Hence, it is important to motivate as many prosumers as possible to interact with each other in the microgrid. To motivate prosumers, the total energy usage cost for each prosumer in the microgrid should be less than or equal to the cost of buying energy from the utility grid. In an ideal scenario, namely when demand matches supply i.e. $SDR_h = 1$, there is no need to interact with the utility grid, and all the energy is sold and bought in the microgrid at the same price. Thus, instead of defining both internal buying and selling prices in the microgrid [95], only one internal price is used.

Initially, the price, at which the energy is sold and bought in the microgrid is called an equilibrium price that is calculated based on the SDR [95, 96]. In case when $QD_h^T = 0$, there is no demand in the microgrid and all excess energy produced by prosumers (QS_h^T) is sold to the utility grid at a price λ_{s_h}. It should be noted that there could also be a case when $QS_h^T = 0$ and $QD_h^T \approx 0$ simultaneously, which means that all prosumers cover their demand by their own PV energy or there is no demand at all, as well as there is no excess energy to share. In case when $SDR_h = 0$, it means that there is no excess energy in the microgrid and energy required to cover the demand of the microgrid is bought from the utility grid at a price λ_{b_h}. When $SDR_h > 1$, it means that there is enough excess energy to cover the demand in the microgrid, as well as to sell some part to the utility grid at a price λ_{s_h}. The most challenging case is when $SDR_h = 1$, namely the total excess energy is equal to the total shortage of energy ($QS_h^T = QD_h^T$). Prosumers may decide the price they agree with in advance, or it may be set by the MGO. To ensure the lowest buying price in the microgrid, the selling price of energy to the utility grid is used as the internal equilibrium price in this scenario. Moreover, when $0 < SDR_h < 1$, it means that there is not enough excess energy to cover the demand in the microgrid and the shortage is bought from the utility grid at a price λ_{b_h}.

In [95], internal buying and selling prices are not the same, whereas in the proposed approach, the equilibrium price is represented by a collapse of two curves (grey curves) to a single line (red line) (Fig. 4.3). It should be noted that the final price for a prosumer may be changed (increased or decreased) due to the penalties added, as well as the cost to cover the energy transfer loss, which is discussed further in this section.

Thus, the internal equilibrium price depends on SDR_h and is calculated according to the formula of a straight line passing through two points. The first point is (0; λ_{b_h}), namely when $SDR_h = 0$, and the shortage is bought from the utility grid at a price λ_{b_h}. The second point is (1; λ_{s_h}), and it represents the case when $SDR_h = 1$, and excess energy is sold to the utility grid at a price λ_{s_h} (Fig. 4.3). Let λ_{e_h} denote the internal equilibrium price in a time slot h, which is calculated as follows:

$$\lambda_{e_h} = SDR_h * (\lambda_{s_h} - \lambda_{b_h}) + \lambda_{b_h} \tag{4.13}$$

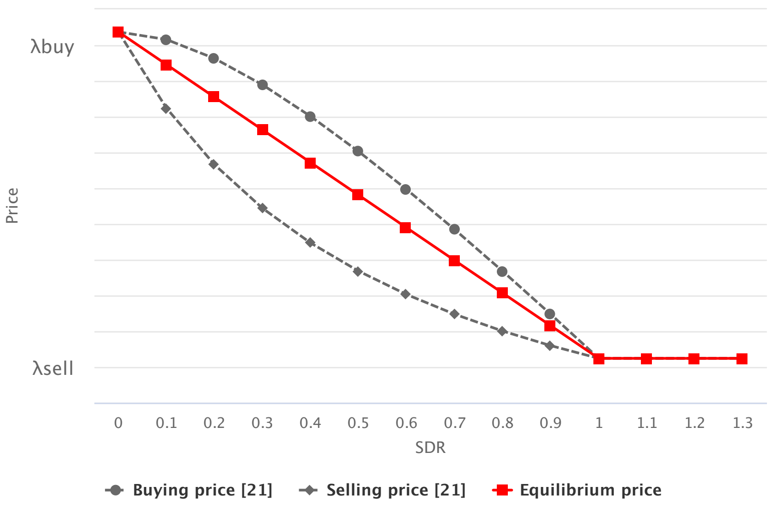

Figure 4.3: An internal equilibrium price

It is highly unlikely that a prosumer's energy usage profile can be accurately predicted in a real-world scenario [35]; thus, there always will be a difference between predicted and actual energy usage profiles. In this work, the MGO is responsible for predicting the energy usage profiles for all prosumers in a microgrid. The proposed approach focuses on determining the energy usage cost and profit from selling energy and not on determining a particular buying or selling price of energy. In other words, prosumers' behavior is considered as random, whereas, the energy usage cost for a prosumer is calculated taking into account a possible absolute deviation from the predicted level of energy consumption. The profit from selling energy for a prosumer is calculated taking into account a possible absolute deviation from the predicted level of energy production. The absolute value of the prosumer's deviation is used to understand what the prosumer's contribution to the total absolute deviation is. It does not matter whether a prosumer's actual value of electricity consumption (production) is increased or decreased compared to the predicted one. In other words, a prosumer will be penalized in both scenarios. This is due to the fact that the microgrid operator sends the information on the predicted energy consumption (production) to an independent aggregator. This means that the amount of energy to be supplied from the independent aggregator to the microgrid is identified in advance based on the prediction. If we define the total deviation as the difference between the Real Consumption and Predicted Consumption (without the absolute value), then that the larger the number of prosumers the smaller is the total deviation. Next, the three most common microgrid's scenarios are explained.

Scenario 1: The total excess energy is equal to the total shortage of energy

In this scenario, the total shortage of energy (QD_h^T) is equal to the total excess energy (QS_h^T) in the microgrid in a time slot h ($SDR_h = 1$). All energy is sold and bought in the microgrid, and there is no need to interact with the utility grid until the demand including the energy transfer loss does not exceed the excess energy. Let $C_{n_h}^B$ denote the initial cost of buying energy from the microgrid for a prosumer n in a time slot h, which is calculated as follows:

$$C_{n_h}^B = QD_{R_{n_h}} * \lambda_{e_h} \tag{4.14}$$

In case when two prosumers consumed the same amount of energy, while one of them did not follow its predicted level of energy consumption, then the energy usage costs will not be the same for these prosumers. Let $\gamma_{n_h}^{QD}$ denote the absolute deviation from the predicted energy consumption profile of a prosumer n in a time slot h, which is calculated as follows:

$$\gamma_{n_h}^{QD} = \left| QD_{R_{n_h}} - QD_{P_{n_h}} \right|, \quad \text{where} \quad QD_{R_{n_h}} > 0 \tag{4.15}$$

Prosumers cannot be charged if they do not consume energy in a time slot h; thus, in (4.15), an absolute deviation is only calculated for prosumers with non-zero consumption.

Prosumers are penalized by adding a penalty to the initial cost ($C^{B}_{n_h}$) based on the prosumer's contribution to the total absolute deviation from the predicted level of total energy consumption. Hence, more a prosumer deviates from the predicted energy consumption profile, the greater the penalty. Let Γ^{QD}_{h} denote the total absolute deviation from the predicted level of total energy consumption in a time slot h, which is calculated as follows:

$$\Gamma^{QD}_{h} = \sum_{n=1}^{N} \gamma^{QD}_{n_h} \tag{4.16}$$

Let $\Delta^{QD}_{n_h}$ denote the prosumer's contribution to the total absolute deviation from the predicted level of total energy consumption in a time slot h. Combining (4.15) and (4.16), $\Delta^{QD}_{n_h}$ is calculated as follows:

$$\Delta^{QD}_{n_h} = \frac{\gamma^{QD}_{n_h}}{\Gamma^{QD}_{h}}, \quad \text{where} \quad \Gamma^{QD}_{h} > 0 \tag{4.17}$$

If prosumers do not deviate from the predicted energy consumption profiles, the total absolute deviation will be equal to 0 ($\Gamma^{QD}_{h} = 0$), and there is no point in calculating prosumer's contribution because it will be equal to 0; thus, in (4.17), the prosumer's contribution is only calculated for the case when at least one prosumer deviates from its predicted energy consumption profile. If $\Gamma^{QD}_{h} \approx 0$, which means that all prosumers followed predicted energy consumption profiles, the total cost for a prosumer n in a time slot h is calculated according to (4.14). To motivate a prosumer not to deviate from the predicted energy consumption profile, a penalty is added to the initial cost. On the other hand, to be able to provide a lower buying price of energy in the microgrid compared with the utility grid, the total cost for a prosumer cannot exceed the baseline cost (cost of buying energy from the utility grid). Let $C^{UG}_{n_h}$ denote the baseline cost for a prosumer n in a time slot h, which is calculated as follows:

$$C^{UG}_{n_h} = QD_{R_{n_h}} * \lambda_{b_h} \tag{4.18}$$

Let $C^{F}_{n_h}$ denote the total cost of buying energy from the microgrid for a prosumer n in a time slot h, which is subject to the constraint:

$$C^{F}_{n_h} \leq C^{UG}_{n_h} \tag{4.19}$$

More precisely, no fixed value can be added to the initial cost because the total cost may exceed the baseline cost. Let $\xi^{QD}_{n_h}$ denote a penalty that is added to the initial cost for a prosumer n in a time slot h, which is subject to the constraint:

$$\xi^{QD}_{n_h} \leq C^{UG}_{n_h} - C^{B}_{n_h} \tag{4.20}$$

In this work, the difference between the baseline cost ($C^{UG}_{n_h}$) and the initial cost ($C^{B}_{n_h}$) is considered as the starting point for calculating a penalty. Let $\psi^{QD}_{n_h}$ denote the difference between the baseline and initial costs for a prosumer n in a time slot h, which is calculated as follows:

$$\psi^{QD}_{n_h} = C^{UG}_{n_h} - C^{B}_{n_h} \tag{4.21}$$

The difference between the baseline and initial costs ($\psi^{QD}_{n_h}$) cannot be used as a penalty because it will lead to the case when the total cost ($C^{F}_{n_h}$) equals to the baseline cost ($C^{UG}_{n_h}$). To ensure (4.19), only some part of $\psi^{QD}_{n_h}$ can be used as a penalty. Combining (4.17) and (4.21), the penalty for a prosumer n in a time slot h (when buying energy) is calculated as follows:

$$\xi^{QD}_{n_h} = \Delta^{QD}_{n_h} * \psi^{QD}_{n_h} \tag{4.22}$$

During the transfer of energy between a prosumer and the MGO, the energy transfer loss is unavoidable. Let $L^{s}_{n_h}$ denote the energy transfer loss for a prosumer n (producer) transferring NP_{n_h} amount of energy to the MGO in a time slot h, and it is calculated as follows:

$$L^{s}_{n_h} = NP^{2}_{n_h} * k_L, \quad NP_{n_h} > 0 \tag{4.23}$$

Let $L^b_{n_h}$ denote the energy transfer loss for a prosumer n (consumer) buying NP_{n_h} amount of energy from the MGO in a time slot h, and it is calculated as follows:

$$L^b_{n_h} = NP^2_{n_h} * k_L, \quad NP_{n_h} < 0 \tag{4.24}$$

Let $L^s_{T_h}$ denote the total energy transfer loss for all prosumers (producers) transferring energy to the MGO in a time slot h, and it is calculated as follows:

$$L^s_{T_h} = \sum_{n=1}^{N} L^s_{n_h} \tag{4.25}$$

Let $L^b_{T_h}$ denote the total energy transfer loss for all prosumers (consumers) buying energy from the MGO in a time slot h, and it is calculated as follows:

$$L^b_{T_h} = \sum_{n=1}^{N} L^b_{n_h} \tag{4.26}$$

Let L^T_h denote the total energy transfer loss in the microgrid in a time slot h, and it is calculated as follows:

$$L^T_h = L^s_{T_h} + L^b_{T_h} \tag{4.27}$$

Thus, the total excess energy available to share in the microgrid is decreased by $L^s_{T_h}$, whereas the total shortage of energy in the microgrid is increased by $L^b_{T_h}$. Let QD^{UG}_h denote the amount of energy has to be bought from the utility grid to cover the energy transfer losses in a time slot h, and it is calculated as follows:

$$QD^{UG}_h = L^s_{T_h} + L^b_{T_h} \tag{4.28}$$

The energy to cover the total energy transfer loss in the microgrid is bought from the utility grid at a price λ_{b_h}. Let C_{L_h} denote the cost of buying energy from the utility grid in a time slot h to cover the energy transfer losses in the microgrid (QD^{UG}_h), and it is calculated as follows:

$$C_{L_h} = QD_h^{UG} * \lambda_{b_h} \tag{4.29}$$

Thus, all prosumers exchanging energy with the MGO have to pay an additional cost to cover the cost of buying energy from the utility grid to cover the losses. Since all prosumers are located on the same distance from the MGO and the type of the wires is the same for all prosumers, the cost each prosumer has to pay is based on the amount of energy exchanged with the MGO. Let λ_{L_h} denote the price of exchanging 1 kW of energy with the MGO for a prosumer n in a time slot h, and it is calculated as follows:

$$\lambda_{L_h} = \frac{C_{L_h}}{L_h^T} \tag{4.30}$$

Let $\tau_{n_h}^L$ denote the cost that has to be paid by a prosumer n in a time slot h to cover the energy transfer losses, and it is calculated as follows:

$$\tau_{n_h}^L = \lambda_{L_h} * L_{n_h}^{b(s)} \tag{4.31}$$

It should be noted that in (4.31), the energy transfer loss for a prosumer n has a superscript $b(s)$ ($L_{n_h}^{b(s)}$), which means that the cost to cover the losses may be calculated using (4.31) for both producers (4.23) and consumers (4.24).

Finally, by combining (4.14), (4.22) and (4.31), the total cost of buying energy from the microgrid for a prosumer n in a time slot h is calculated as follows:

$$C_{n_h}^F = C_{n_h}^B + \xi_{n_h}^{QD} + \tau_{n_h}^L \tag{4.32}$$

Thus, the total energy usage cost for a prosumer n in a time slot h ($C_{n_h}^F$) consists of the initial cost ($C_{n_h}^B$), a penalty ($\xi_{n_h}^{QD}$) that depends on the prosumer's contribution to the total absolute deviation from the predicted level of total energy consumption, and the cost that has to be paid to cover the losses ($\tau_{n_h}^L$).

Let $P^{B}_{n_h}$ denote the initial profit from selling energy to the microgrid for a prosumer n in a time slot h, which is calculated as follows:

$$P^{B}_{n_h} = QS_{R_{n_h}} * \lambda_{e_h} \tag{4.33}$$

In case when two prosumers produced the same amount of energy, while one of them did not follow its predicted level of energy production, then the profits of selling energy will be different for these prosumers. Let $\gamma^{QS}_{n_h}$ denote the absolute deviation from the predicted energy production profile of a prosumer n in a time slot h, which is calculated as follows:

$$\gamma^{QS}_{n_h} = \left| QS_{R_{n_h}} - QS_{P_{n_h}} \right|, \text{ where } \quad QS_{R_{n_h}} > 0 \tag{4.34}$$

Prosumers cannot get any profit if they do not produce energy in a time slot h; thus, in (4.34), an absolute deviation is only calculated for prosumers with non-zero production. Prosumers are penalized by subtracting a penalty from the initial profit ($P^{B}_{n_h}$) based on the prosumer's contribution to the total absolute deviation from the predicted level of total energy production. Hence, the more a prosumer deviates from the predicted energy production profile, the greater the penalty. Let Γ^{QS}_{h} denote the total absolute deviation from the predicted level of total energy production in a time slot h, which is calculated as follows:

$$\Gamma^{QS}_{h} = \sum_{n=1}^{N} \gamma^{QS}_{n_h} \tag{4.35}$$

Let $\Delta^{QS}_{n_h}$ denote the prosumer's contribution to the total absolute deviation from the predicted level of total energy production in a time slot h. Combining (4.34) and (4.35), $\Delta^{QS}_{n_h}$ is calculated as follows:

$$\Delta^{QS}_{n_h} = \frac{\gamma^{QS}_{n_h}}{\Gamma^{QS}_{h}}, \text{ where } \quad \Gamma^{QS}_{h} > 0 \tag{4.36}$$

If prosumers do not deviate from the predicted energy production profiles, the total absolute deviation will be equal to 0 ($\Gamma_h^{QS} = 0$), which means that the contribution to the total absolute deviation will be equal to 0 for all prosumers; thus, in (4.36), the prosumer's contribution is only calculated for the case when at least one prosumer deviates from its predicted energy production profile. If $\Gamma_h^{QS} \approx 0$, which means that all prosumers followed predicted energy production profiles, the total profit for a prosumer n in a time slot h is calculated according to (4.33). To motivate a prosumer not to deviate from the predicted energy production profile, the initial profit is decreased by a penalty. On the other hand, to be able to provide a higher selling price of energy in the microgrid compared with the utility grid, the total profit for a prosumer cannot be less than the baseline profit (profit from selling energy to the utility grid). Let $P_{n_h}^{UG}$ denote the baseline profit for a prosumer n in a time slot h, which is calculated as follows:

$$P_{n_h}^{UG} = QS_{R_{n_h}} * \lambda_{s_h} \tag{4.37}$$

Let $P_{n_h}^{F}$ denote the total profit from selling energy to the microgrid for a prosumer n in a time slot h, which is subject to the constraint:

$$P_{n_h}^{F} \geq P_{n_h}^{UG} \tag{4.38}$$

The initial profit cannot be decreased by any fixed value because the total profit may become less than the baseline profit. Let $\xi_{n_h}^{QS}$ denote a penalty that is subtracted from the initial profit for a prosumer n in a time slot h, which is subject to the constraint:

$$\xi_{n_h}^{QS} \leq P_{n_h}^{B} - P_{n_h}^{UG} \tag{4.39}$$

Similar to (4.21), the difference between the initial and baseline profits is used to calculate a penalty. Let $\psi_{n_h}^{QS}$ denote the difference between the initial and baseline profits for a prosumer n in a time slot h, which is calculated as follows:

$$\psi_{n_h}^{QS} = P_{n_h}^{B} - P_{n_h}^{UG} \tag{4.40}$$

The difference between the initial and baseline profits ($\psi_{n_h}^{QS}$) cannot be used as a penalty because it will lead to the case when the total profit ($P_{n_h}^{F}$) equals to the baseline profit ($P_{n_h}^{UG}$). To ensure (4.38), only some part of $\psi_{n_h}^{QS}$ can be used as a penalty. Combining (4.36) and (4.40), the penalty for a prosumer n in a time slot h is calculated as follows:

$$\xi_{n_h}^{QS} = \Delta_{n_h}^{QS} * \psi_{n_h}^{QS} \tag{4.41}$$

Finally, by combining (4.33), (4.41), and (4.31), the total profit from selling energy to the microgrid for a prosumer n in a time slot h is calculated as follows:

$$P_{n_h}^{F} = P_{n_h}^{B} - \xi_{n_h}^{QS} - \tau_{n_h}^{L} \tag{4.42}$$

In this particular scenario, when $SDR_h = 1$, the prices of selling energy to the utility grid and to the microgrid in a time slot h are the same ($\lambda_{s_h} = \lambda_{e_h}$) according to (4.13), which means that there will be no penalties for prosumers when calculating the profit from selling energy. Even if some prosumers deviate from the predicted energy production profiles, the microgrid continues working in the standalone mode; thus, prosumers are not penalized for any changes in energy production profiles, and the total profit for a prosumer n in a time slot h consists of the initial profit ($P_{n_h}^{B}$) and the cost that has to be paid to cover the losses ($\tau_{n_h}^{L}$).

Scenario 2: The total excess energy is greater than the total shortage of energy

In this scenario, the total excess energy (QS_h^T) is greater than the total shortage of energy (QD_h^T) in the microgrid in a time slot h ($SDR_h > 1$). The total cost ($C_{n_h}^{F}$) of buying energy from the microgrid (except the cost that has to be paid to cover the energy transfer losses) for each prosumer n in a time slot h is calculated according to (4.32). The energy transfer losses may be covered by PV energy produced in the microgrid or by buying energy from the utility grid. If there is enough excess energy in the microgrid, the energy transfer losses are covered by the energy produced in the microgrid. On the other hand, if the energy transfer losses cannot be covered by the excess energy in the microgrid, the shortage is bought from the utility grid. Let QD_h^{UG} denote the amount of energy that has to be bought from the utility grid to cover the losses, and it is calculated as follows:

$$QD_h^{UG} = (QD_h^T + L_{T_h}^b) - (QS_h^T - L_{T_h}^s) \tag{4.43}$$

The price of exchanging 1 kW of energy with the MGO for a prosumer n in a time slot h is calculated as follows (taking into account (4.43)):

$$\lambda_{L_h} = \begin{cases} \dfrac{QD_h^{UG} * \lambda_{b_h}}{L_h^T}, & QD_h^{UG} > 0 \\ \dfrac{L_h^T * \lambda_{e_h}}{L_h^T}, & QD_h^{UG} < 0 \end{cases} \tag{4.44}$$

Taking into account (4.44), the cost that has to be paid by a prosumer n in a time slot h to cover the energy transfer losses is calculated according to (4.31).

The total profit ($P_{n_h}^F$) from selling energy to the microgrid is calculated in a different manner compared with (4.42). Consider a simple example. Let P_1, P_2 be the prosumers who only consume energy buying it from the microgrid, whereas the total energy consumed $QD_{R_h}^T$ = 4 kW in a time slot h. Let P_3, P_4 be the prosumers who only produce energy and sell it to the microgrid, where both P_3 and P_4 produced the same amount of energy ($QS_{R_{3_h}} = QS_{R_{4_h}}$ = 4 kW) in a time slot h. In this example, $SDR_h = 8/4 = 2$. If the total demand in the microgrid ($QD_{R_h}^T$ = 4 kW) is covered by only energy produced by P_4 (4 kW), it will lead to the case when P_3 sells all the energy produced ($QS_{R_{3_h}}$) to the utility grid at a price λ_{s_h}, whereas, P_4 sells $QS_{R_{4_h}}$ (4 kW) to the microgrid at a price λ_{e_h} ($\lambda_{e_h} \geq \lambda_{s_h}$).

To eliminate such a scenario, each prosumer sells at least some part of the energy produced to the microgrid, whereas the rest is sold to the utility grid. Let $\mu_{n_h}^{MG}$ denote the part of the energy produced by a prosumer n in a time slot h, which is sold to the microgrid, and it is calculated as follows:

$$\mu_{n_h}^{MG} = QS_{R_{n_h}} * SDR_h^{-1} \tag{4.45}$$

Let $\mu_{n_h}^{UG}$ denote the part of the energy produced by a prosumer n in a time slot h, which is sold to the utility grid, and it is calculated as follows:

$$\mu_{n_h}^{UG} = QS_{R_{n_h}} * (1 - SDR_h^{-1}) \tag{4.46}$$

In this scenario, the initial profit consists of two parts, namely the profit from selling energy to the microgrid and the profit from selling energy to the utility grid. Thus, the initial profit for a prosumer n in a time slot h is calculated as follows:

$$P_{n_h}^{B} = \mu_{n_h}^{MG} * \lambda_{e_h} + \mu_{n_h}^{UG} * \lambda_{s_h} \tag{4.47}$$

To motivate a prosumer not to deviate from the predicted energy production profile, the initial profit is decreased by a penalty. It should be noted that the profit from selling energy to the utility grid cannot be decreased. Hence, it is only possible to decrease the profit from selling energy to the microgrid. Thus, taking into account (4.38), the penalty that is subtracted from the initial profit (4.47) is subject to the constraint:

$$\xi_{n_h}^{QS} \leq \mu_{n_h}^{MG} * \lambda_{e_h} - \mu_{n_h}^{MG} * \lambda_{s_h} \tag{4.48}$$

In this scenario, the difference between the profits is calculated only for the part of the energy that is sold to the microgrid ($\mu_{n_h}^{MG}$). Thus, the difference between the profits for a prosumer n in a time slot h is calculated as follows:

$$\psi_{n_h}^{QS} = \mu_{n_h}^{MG} * \lambda_{e_h} - \mu_{n_h}^{MG} * \lambda_{s_h} \tag{4.49}$$

The initial profit (4.47) is decreased depending on (4.36). Thus, combining (4.36) and

(4.49), the penalty that is subtracted from the initial profit (4.47) is calculated as follows:

$$\xi^{QS}_{n_h} = \Delta^{QS}_{n_h} * \psi^{QS}_{n_h} \tag{4.50}$$

Combining (4.45), (4.46), (4.50), and (4.31) (taking into account (4.44)), the total profit from selling energy to the microgrid for a prosumer n in a time slot h is calculated as follows:

$$P^{F}_{n_h} = \mu^{MG}_{n_h} * \lambda_{e_h} + \mu^{UG}_{n_h} * \lambda_{s_h} - \xi^{QS}_{n_h} - \tau^{L}_{n_h} \tag{4.51}$$

Scenario 3: The total excess energy is less than the total shortage of energy

In this scenario, the total shortage of energy (QD^T_h) is greater than the total excess energy (QS^T_h) in the microgrid in a time slot h ($SDR_h < 1$). The total profit ($P^F_{n_h}$) from selling energy to the microgrid for each prosumer n in a time slot h is calculated according to (4.42). The total cost ($C^F_{n_h}$) of buying energy from the microgrid is calculated in a different manner compared with (4.32). Consider a simple example. Let P_1, P_2 be the prosumers who only produce energy and sell it to the microgrid, whereas the total energy produced $QS^T_{R_h} = 5$ kW in a time slot h. Let P_3, P_4 be the prosumers who only consume energy buying it from the microgrid, where both P_3 and P_4 consumed the same amount of energy ($QD_{R_{3_h}} = QD_{R_{4_h}} = 5$ kW) in a time slot h. In this example, $SDR_h = 5/10 = 1/2$. If all the energy produced in the microgrid ($QS^T_{R_h} = 5$ kW) is sold to P_3, it will lead to the case when P_3 covers the demand ($QD_{R_{3_h}} = 5$ kW) buying energy from the microgrid at a price λ_{e_h}, whereas, P_4 covers the demand ($QD_{R_{4_h}} = 5$ kW) buying energy from the utility grid at a price λ_{b_h} ($\lambda_{b_h} \geq \lambda_{e_h}$). To eliminate such a scenario, each prosumer covers at least some part of the demand by the energy produced in the microgrid, whereas the rest is bought from the utility grid.

Let $\omega^{MG}_{n_h}$ denote the part of the energy consumed by a prosumer n in a time slot h, which is bought from the microgrid, and it is calculated as follows:

$$\omega^{MG}_{n_h} = QD_{R_{n_h}} * SDR_h \tag{4.52}$$

Let $\omega_{n_h}^{UG}$ denote the part of the energy consumed by a prosumer n in a time slot h, which is bought from the utility grid, and it is calculated as follows:

$$\omega_{n_h}^{UG} = QD_{R_{n_h}} * (1 - SDR_h) \tag{4.53}$$

In this scenario, the initial cost consists of two parts, namely the cost of buying energy from the microgrid and the cost of buying energy from the utility grid. Thus, the initial cost for a prosumer n in a time slot h is calculated as follows:

$$C_{n_h}^{B} = \omega_{n_h}^{MG} * \lambda_{e_h} + \omega_{n_h}^{UG} * \lambda_{b_h} \tag{4.54}$$

To motivate a prosumer not to deviate from the predicted energy consumption profile, a penalty is added to the initial cost. It should be noted that the cost of buying energy from the utility grid cannot be increased. Hence, it is only possible to add a penalty to the cost of buying energy from the microgrid. Thus, taking into account (4.19), the penalty that is added to the initial cost (4.54) is subject to the constraint:

$$\xi_{n_h}^{QD} \leq \omega_{n_h}^{MG} * \lambda_{b_h} - \omega_{n_h}^{MG} * \lambda_{e_h} \tag{4.55}$$

In this scenario, the difference between the costs is calculated only for the part of the energy that is bought from the microgrid ($\omega_{n_h}^{MG}$). Thus, the difference between the costs for a prosumer n in a time slot h is calculated as follows:

$$\psi_{n_h}^{QD} = \omega_{n_h}^{MG} * \lambda_{b_h} - \omega_{n_h}^{MG} * \lambda_{e_h} \tag{4.56}$$

The initial cost (4.54) is increased depending on (4.17). Thus, combining (4.17) and (4.56), the penalty that is added to the initial cost (4.54) is calculated as follows:

$$\xi_{n_h}^{QD} = \Delta_{n_h}^{QD} * \psi_{n_h}^{QD} \tag{4.57}$$

Combining (4.52), (4.53), (4.57), and (4.31), the total cost of buying energy from the microgrid for a prosumer n in a time slot h is calculated as follows:

$$C^F_{n_h} = \omega^{MG}_{n_h} * \lambda_{e_h} + \omega^{UG}_{n_h} * \lambda_{b_h} + \xi^{QD}_{n_h} + \tau^L_{n_h} \tag{4.58}$$

By combining the methodologies of calculating the total cost and total profit for a prosumer in different microgrid scenarios, the total cost for a prosumer n in a time slot h is calculated according to the Algorithm 1, whereas, the total profit for a prosumer n in a time slot h is calculated according to the Algorithm 2.

Moreover, the actual price at which a prosumer n bought energy from the microgrid in a time slot h can be determined by dividing the total cost ($C^F_{n_h}$) by the amount of energy consumed ($QD_{R_{n_h}}$). Let $Pr^B_{n_h}$ denote the actual buying price of energy from the microgrid for a prosumer n in a time slot h, which is calculated as follows:

$$Pr^B_{n_h} = \frac{C^F_{n_h}}{QD_{R_{n_h}}}, \text{ where } \; QD_{R_{n_h}} > 0 \tag{4.59}$$

In (4.59), the actual buying price can be calculated only for a prosumer with a non-zero level of energy consumption. Similarly, the actual price at which a prosumer n sold energy to the microgrid in a time slot h can be determined by dividing the total profit ($P^F_{n_h}$) by the amount of energy sold ($QS_{R_{n_h}}$). Let $Pr^S_{n_h}$ denote the actual selling price of energy to the microgrid for a prosumer n in a time slot h, which is calculated as follows:

$$Pr^S_{n_h} = \frac{P^F_{n_h}}{QS_{R_{n_h}}}, \text{ where } \; QS_{R_{n_h}} > 0 \tag{4.60}$$

In (4.60), the actual selling price can be calculated only for a prosumer with a non-zero level of energy production.

Algorithm 1 Algorithm for calculating the total cost

```
**Input:** $h$, $\mathcal{N}$, $\lambda_{b_h}$, $\lambda_{s_h}$, energy usage profiles
Calculate $SDR_h$ according to (4.8)
Calculate $\lambda_{e_h}$ according to (4.13)
Calculate $\lambda_{L_h}$ according to (4.30)
for all $n \in \mathcal{N}$ do
    Calculate $\gamma_{n_h}^{QD}$, $\Gamma_h^{QD}$ according to (4.15), (4.16)
    Calculate $\Delta_{n_h}^{QD}$ according to (4.17)
    Calculate $\tau_{n_h}^{L}$ according to (4.31)
    if $SDR_h \geq 1$ then
        Calculate the initial cost $C_{n_h}^{B}$ according to (4.14)
        Calculate $\psi_{n_h}^{QD}$, $\xi_{n_h}^{QD}$ according to (4.21), (4.22)
        Calculate the total cost $C_{n_h}^{F}$ according to (4.32)
    else if $SDR_h < 1$ then
        Calculate $\omega_{n_h}^{MG}$, $\omega_{n_h}^{UG}$ according to (4.52), (4.53)
        Calculate the initial cost $C_{n_h}^{B}$ according to (4.54)
        Calculate $\psi_{n_h}^{QD}$, $\xi_{n_h}^{QD}$ according to (4.56), (4.57)
        Calculate the total cost $C_{n_h}^{F}$ according to (4.58)
    end if
end for
```

Algorithm 2 Algorithm for calculating the total profit

```
**Input:** $h$, $\mathcal{N}$, $\lambda_{b_h}$, $\lambda_{s_h}$, energy usage profiles
Calculate $SDR_h$ according to (4.8)
Calculate $\lambda_{e_h}$ according to (4.13)
Calculate $\lambda_{L_h}$ according to (4.30)
for all $n \in \mathcal{N}$ do
    Calculate $\gamma_{n_h}^{QS}$, $\Gamma_h^{QS}$ according to (4.34), (4.35)
    Calculate $\Delta_{n_h}^{QS}$ according to (4.36)
    Calculate $\tau_{n_h}^{L}$ according to (4.31)
    if $SDR_h \leq 1$ then
        Calculate the initial profit $P_{n_h}^{B}$ according to (4.33)
        Calculate $\psi_{n_h}^{QS}$, $\xi_{n_h}^{QS}$ according to (4.40), (4.41)
        Calculate the total profit $P_{n_h}^{F}$ according to (4.42)
    else if $SDR_h > 1$ then
        Calculate $\mu_{n_h}^{MG}$, $\mu_{n_h}^{UG}$ according to (4.45), (4.46)
        Calculate the initial profit $P_{n_h}^{B}$ according to (4.47)
        Calculate $\psi_{n_h}^{QS}$, $\xi_{n_h}^{QS}$ according to (4.49), (4.50)
        Calculate the total profit $P_{n_h}^{F}$ according to (4.51)
    end if
end for
```

4.3 Evaluation and Discussion

This section presents the results of simulations to evaluate the proposed pricing model for a microgrid of prosumers. The microgrid consists of 100 prosumers ($N = 100$) that are capable of producing energy using photovoltaic systems.

To conduct the simulations, two real datasets are used. First, the energy consumption profiles for 100 prosumers are extracted from the first dataset [93] that contains energy consumption readings of London households between November 2011 and February 2014. Then, the energy production profiles for 100 prosumers are extracted from the second dataset [94] that contains energy data from domestic premises with high uptake of solar photovoltaic (PV) embedded generation between July 2013 and November 2014. Next, the energy consumption and production profiles of 100 prosumers are combined together to form the input data for the proposed pricing model. For a single-phase electricity connection, the voltage (V) is taken as 230V. The amount of energy any prosumer exchanges with the MGO does not exceed 3.7 kW. The electrical wire resistance ρ is equal to 0.01 Ohm/m, whereas the distance d between any prosumer and the MGO is equal to 100 m. Since the energy usage forecasting mechanism is beyond the scope of this part of the research, predicted energy usage profiles are generated by adding noise to the real data.

To highlight the efficiency of the proposed pricing model, the simulations are conducted for two different tariffs:

- **TARIFF 1:** Based on [97], the buying price of energy from the utility grid (λ_{b_h}) in the UK is taken as 14.37 pence/kWh, whereas, the selling price of energy to the utility grid (λ_{s_h}) in the UK is taken as 5.24 pence/kWh [98]. In this case, the buying and selling prices of energy from (to) the utility grid do not vary during the day.
- **TARIFF 2:** The Green Energy UK's TIDE tariff [99] is used, where the buying price of energy from the utility grid (λ_{b_h}) varies during the day from 7.5 pence/kWh to 32.55 pence/kWh (Table 4.1). The selling price of energy to the utility grid (λ_{s_h})

is taken as an average (4.04 pence/kWh) of different selling prices of energy based on the data from different suppliers [100] (Table 4.2). In this case, the selling price of energy to the utility grid does not vary during the day.

Table 4.1: Green Energy UK's TIDE tariff (weekday)

Time	Buying Price (pence/kWh)
Midnight - 7am	7.5p
7am - 4pm	16.44p
4pm - 8pm	32.55p
8pm - Midnight	16.44p

Table 4.2: Selling prices of energy for different suppliers

Supplier	Selling price (pence/kWh)
Social Energy	5.6p
Octopus Energy	5.5p
E.ON	5.5p
Bulb	5.38p
SO Energy	5p
OVO Energy	4p
Scottish Power	4p
EDF Energy	3.5p
Shell Energy	3.5p
SSE	3.5p
British Gas	3.2p
Avro Energy	3p
Utilita	3p
Utility Warehouse	2p

To evaluate the proposed pricing model and emphasize its effectiveness in the total energy cost reduction is independent of the buying and selling prices of energy to (from) the utility grid, simulations are conducted for two different tariffs (**TARIFF 1**, **TARIFF 2**). For each tariff, the simulation results are compared with the baseline scenario when prosumers interact with the utility grid only (baseline), an approach in [95], and when prosumers interact with the microgrid using the proposed method. The simulations are conducted using PHP language on the machine with Intel Core i5 CPU @ 1.30 GHz and 4.00 GB RAM. It takes around 0.05 seconds on average to calculate the total energy usage cost and profit for 100 prosumers.

Fig. 4.4 shows the total actual and predicted energy consumption and production of all

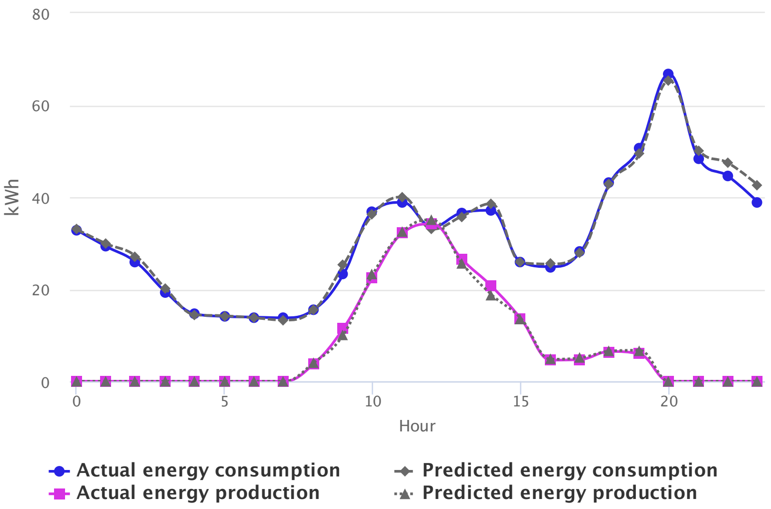

Figure 4.4: Total predicted and actual energy consumption and production in the microgrid

prosumers in the microgrid for one day. The increase in energy consumption can be observed during peak periods, namely from 8 a.m. to 3 p.m. and from 5 p.m. to 10 p.m. It can be observed that prosumers deviate from the predicted energy consumption profiles. For example, the total actual energy consumption is greater than the total predicted energy consumption in a time slot $h = 12$, whereas, in a time slot $h = 14$, the total predicted energy consumption is greater than the total actual energy consumption. The increase in energy production using photovoltaic systems can be observed from 8 a.m. to 4 p.m. It can be seen that prosumers deviate from the predicted energy production profiles. For example, the total predicted energy production is greater than the total actual energy production in a time slot $h = 12$, whereas the total actual energy production is greater than the total predicted energy production in a time slot $h = 14$.

Fig. 4.5 shows the supply and demand ratio (SDR_h) that depends on the total excess energy (QS_h^T) and the total shortage of energy (QD_h^T) in the microgrid. It can be seen that with the increase in energy production in the microgrid from 8 a.m. to 4 p.m., SDR also increases and reaches its peak in a time slot $h = 12$ ($SDR_{12} = 1.039$). It should be noted that in a time slot $h = 12$, excess energy (except the amount of energy that is needed to cover the losses) is sold to the utility grid, which also can be observed in Fig. 4.6 in a time slot $h =$

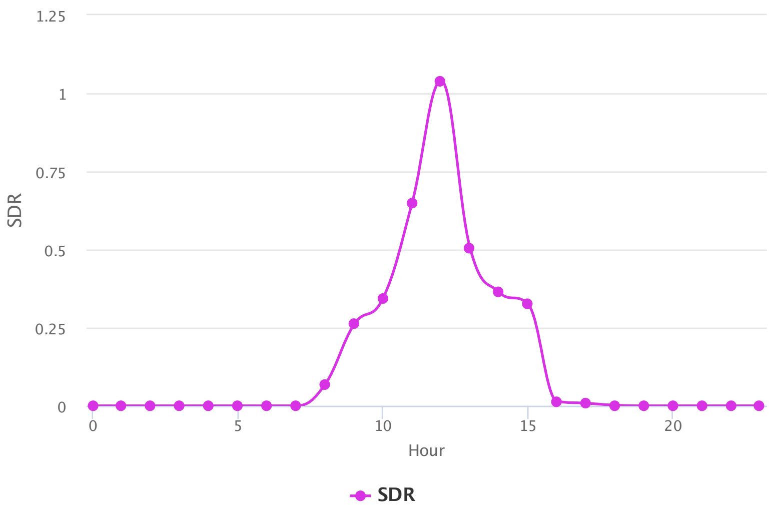

Figure 4.5: Supply and demand ratio

12.

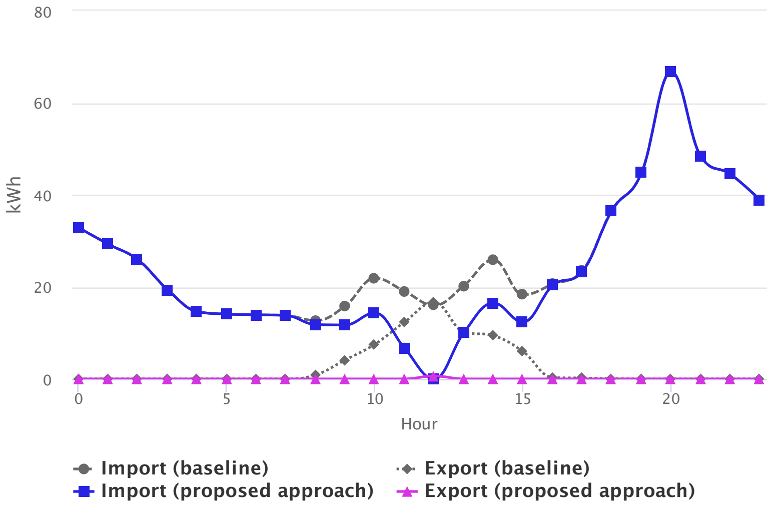

Figure 4.6: Total amount of energy imported and exported from (to) the utility grid for the baseline scenario and using the proposed approach

By utilizing energy produced in the microgrid more efficiently, the demand for electricity from the utility grid can be reduced. Fig 4.6 shows the total amount of energy imported and exported from (to) the utility grid for the baseline scenario and using the proposed approach. Due to the increase in energy production in the microgrid from 8 a.m. to 4 p.m. (Fig. 4.4), it can be observed that the amount of energy exported to the utility grid decreases during this time interval because of utilizing produced energy in the microgrid. In addition, when $SDR_h > 1$ ($h = 12$), the only small part of produced energy in the microgrid is sold to the utility grid (taking into account the losses). Thus, Fig. 4.6 emphasizes that the application of the proposed pricing model may reduce the amount of energy imported and exported from (to) the utility grid by utilizing energy produced in the microgrid more efficiently.

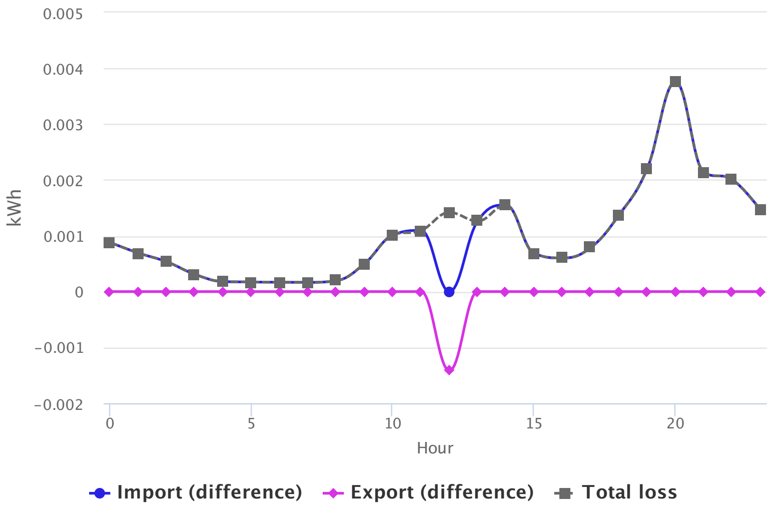

Figure 4.7: Total energy transfer loss and the difference between the amount of energy imported and exported from (to) the utility grid using the proposed approach and an approach in [95].

Fig. 4.7 shows the total energy transfer loss in the microgrid. It can be observed that the total loss reaches its peak in a time slot $h = 20$ because of the high demand for energy (high amount of energy exchanged with the MGO). It can be seen that there is no difference in energy imported from the utility grid in a time slot $h = 12$ because the excess is sold (exported) to the utility grid. On the other hand, there is a slight difference in energy imported from the utility grid from midnight to 11 a.m. and from 1 p.m. to midnight because the proposed approach takes into account the energy transfer loss, whereas the approach in [95] does not. In addition, there is no difference in energy exported to the utility grid because $SDR_h < 1$ from midnight to 11 a.m. and from 1 p.m. to midnight. The amount of energy exported to the utility grid using the proposed approach is less by 0.0014 kW compared with the approach in [95] because of the energy transfer losses.

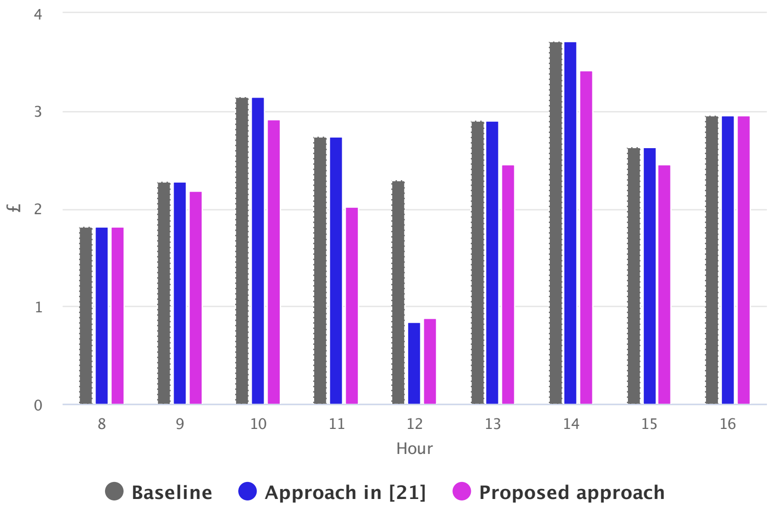

Figure 4.8: Total energy usage cost of all prosumers (**TARIFF 1**) for the baseline scenario, an approach in [95], and the proposed approach from 8 a.m. to 4 p.m.

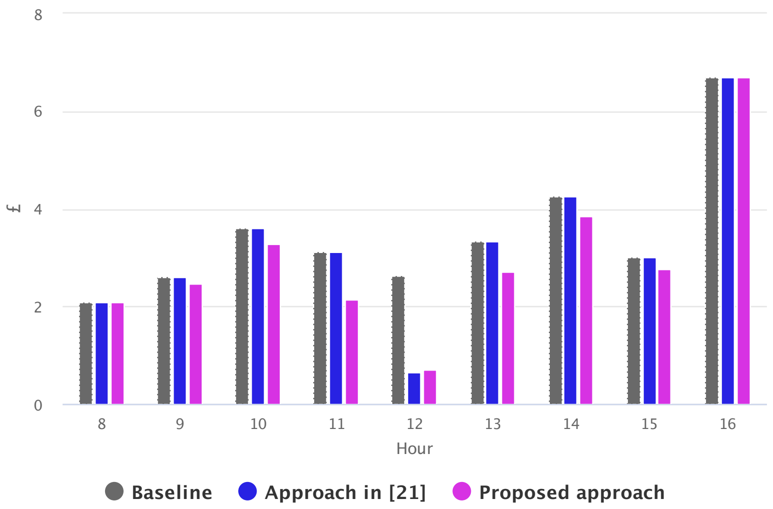

Figure 4.9: Total energy usage cost of all prosumers (**TARIFF 2**) for the baseline scenario, an approach in [95], and the proposed approach from 8 a.m. to 4 p.m.

Fig. 4.8 shows the total energy usage cost of all prosumers (**TARIFF 1**) in the microgrid for the baseline scenario, an approach in [95], and the proposed approach. It can be observed that with the increase in the energy production in the microgrid from 8 a.m. to 4 p.m., the total energy usage cost decreases because prosumers buy energy from the microgrid at a lower price (λ_{e_h}) compared with the buying price of energy from the utility grid (λ_{b_h}). In addition, it can be seen that using the proposed approach, the total energy usage cost of all prosumers is less in most time slots compared with the baseline scenario and the approach in [95]. In a time slot $h = 12$, the total energy usage cost is greater compared with the approach in [95] because of the energy transfer losses. The same pattern can be observed for the second tariff (**TARIFF 2**) in Fig. 4.9.

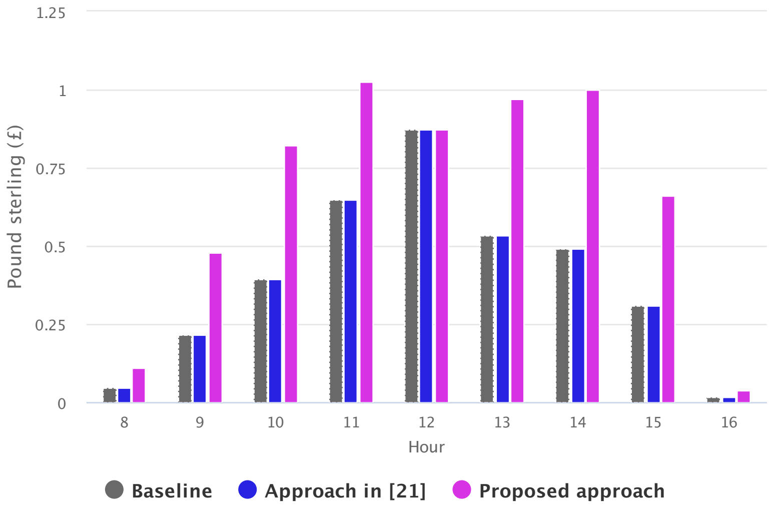

Figure 4.10: Total profit of all prosumers (**TARIFF 1**) for the baseline scenario, an approach in [95], and the proposed approach from 8 a.m. to 4 p.m.

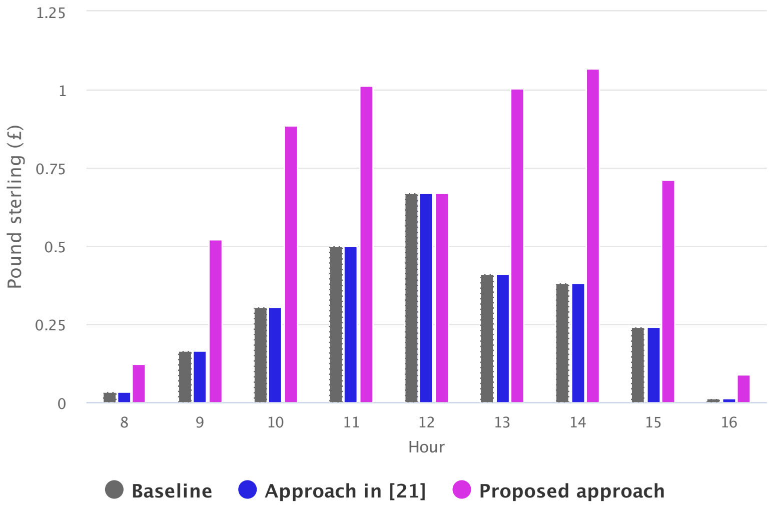

Figure 4.11: Total profit of all prosumers (**TARIFF 2**) for the baseline scenario, an approach in [95], and the proposed approach from 8 a.m. to 4 p.m.

Fig. 4.10 shows the total profit of all prosumers in the microgrid for the baseline scenario, an approach in [95], and the proposed approach. Prosumers sell energy to the microgrid at a higher price (λ_{e_h}) compared with the selling price of energy to the utility grid (λ_{s_h}). Thus, the total profit of all prosumers when interacting with the microgrid using the proposed approach is greater compared with the baseline scenario and the approach in [95]. However, the total profit of all prosumers interacting with the microgrid is the same as the baseline scenario and the approach in [95] in a time slot $h = 12$ ($SDR_{12} = 1.039$) to make prosumers continue to participate in the microgrid. The same dynamics can be observed for the second tariff (**TARIFF 2**) in Fig. 4.11.

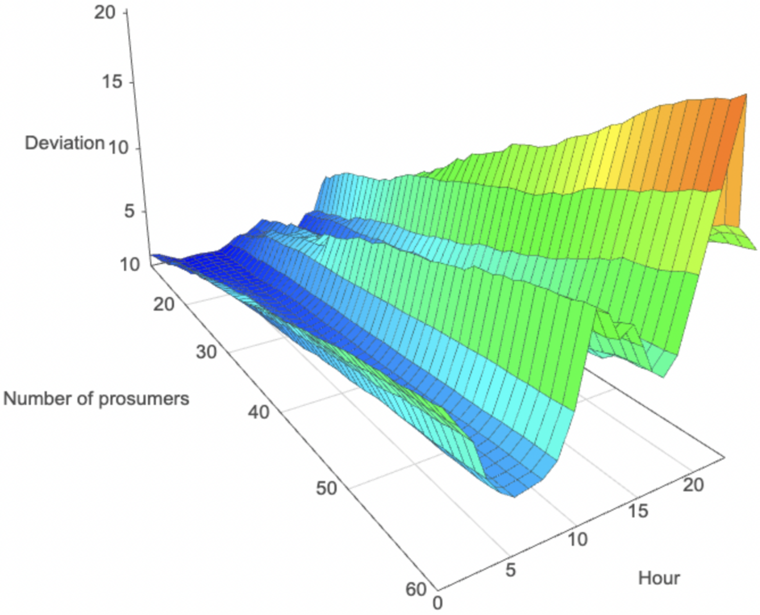

Figure 4.12: Dependency of the total absolute deviation from the total predicted energy consumption on a number of prosumers in the microgrid

Fig. 4.12 shows the dependency of the total absolute deviation from the total predicted energy consumption on the number of prosumers in the microgrid. It can be seen that with the increasing number of prosumers in the microgrid, the total absolute deviation also increases and reaches its peak (20.3182 kW) in a time slot $h = 20$ when there are 60 prosumers in the microgrid, whereas the minimum absolute deviation (0.5077 kW) can be observed in a time slot $h = 3$ for the microgrid that consists of 10 prosumers. Thus, the more prosumers in the microgrid, the more the total absolute deviation from the predicted energy consumption.

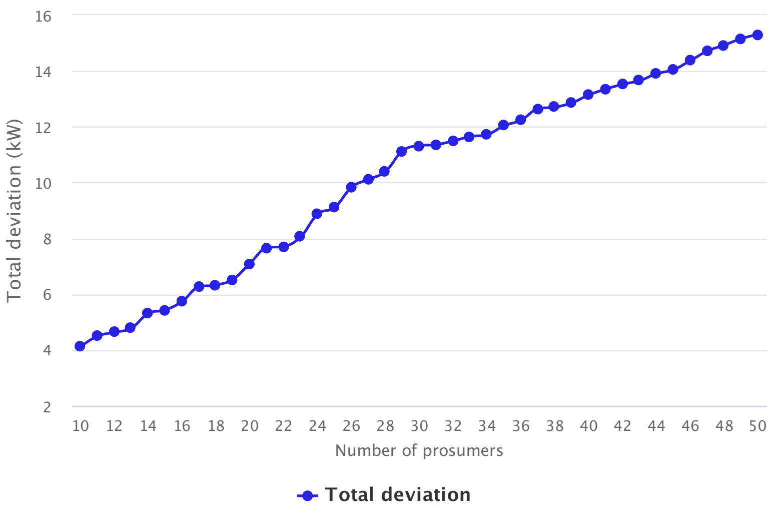

Figure 4.13: Dependency of the total absolute deviation from the total predicted energy consumption on a number of prosumers in the microgrid in a time slot $h = 11$

Fig. 4.13 shows the dependency of the total absolute deviation from the total predicted energy consumption on the number of prosumers in the microgrid in a time slot $h = 11$. It can be seen that with the increasing number of prosumers, the total absolute deviation also increases.

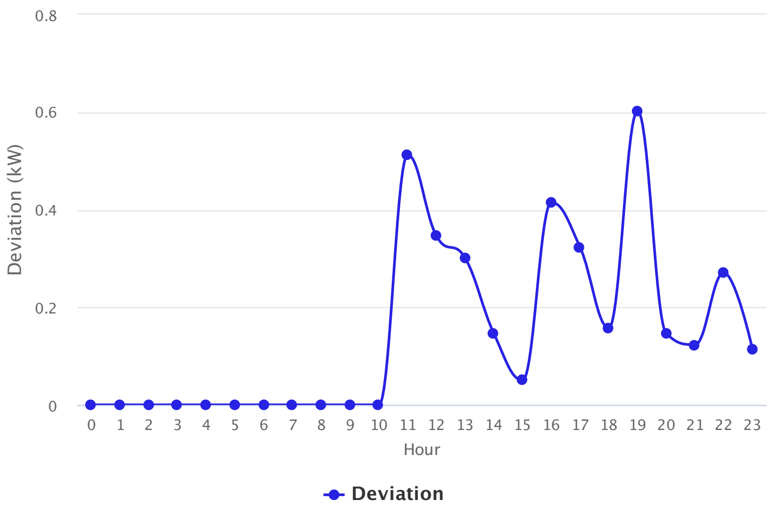

Figure 4.14: Absolute deviation from the predicted energy consumption for a randomly chosen prosumer

Fig. 4.15 shows the dependency of the randomly chosen prosumer's contribution to the total absolute deviation from the total predicted energy consumption on the total

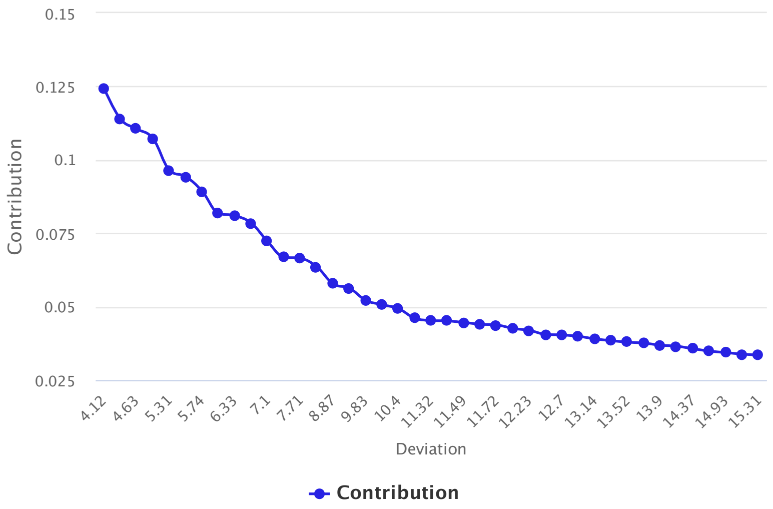

Figure 4.15: Dependency of the randomly chosen prosumer's contribution to the total absolute deviation from the total predicted energy consumption

absolute deviation from the total predicted energy consumption. It can be observed that the contribution not only depends on the prosumer's absolute deviation from its predicted energy consumption (Fig. 4.14) but also depends on the contribution (deviation) of other prosumers. With the increasing total absolute deviation , the prosumer's contribution to the total absolute deviation decreases because the prosumer's absolute deviation does not change with the increasing number of prosumers (total absolute deviation).

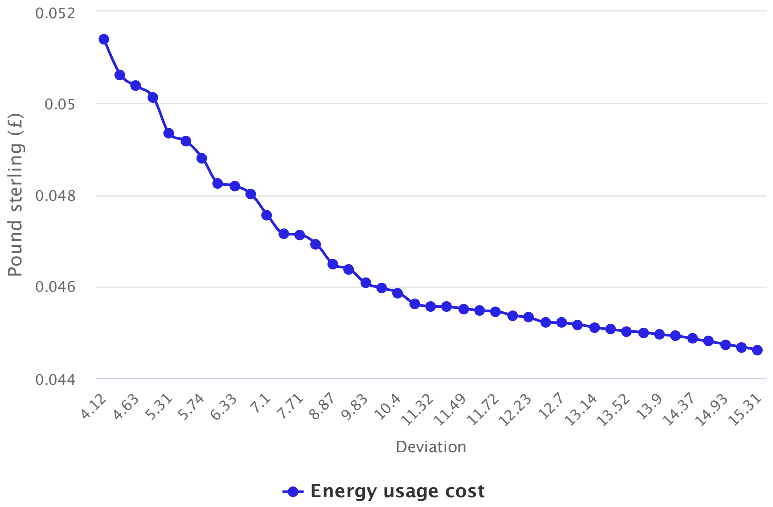

Figure 4.16: Dependency of the total energy usage cost for a randomly chosen prosumer on the total absolute deviation from the predicted energy consumption in a time slot $h = 11$

Fig. 4.16 shows the dependency of the energy usage cost for a randomly chosen prosumer on the total absolute deviation from the total predicted energy consumption in the microgrid in a time slot $h = 11$. It can be seen that with the increasing total absolute deviation , the energy usage cost for a prosumer decreases because the prosumer's contribution to the total absolute deviation decreases (Fig. 4.15), which affects the amount of penalty (4.22) that is added to the cost.

Table 4.3: Decrease in the total energy usage cost of all prosumers (%)

Hour	TARIFF 1	TARIFF 2
8	0.38	0.35
9	5.3	4.96
10	10.61	8.76
11	55.33	31.19
12	61.41	73.37
13	36.41	18.64
14	14.11	9.67
15	9.82	7.7
16	0.02	0.02

The percentage decrease in the total energy usage cost of all prosumers (from 8 a.m. to 4 p.m.) is shown in Table 4.3. It can be seen that with the increase in the energy production in the microgrid (Fig. 4.4) from 8 a.m. to 4 p.m., the value of the decrease in the total energy usage cost grows. For the **TARIFF 1**, the maximum decrease in the cost is around 61.41% in a time slot $h = 12$, whereas for the **TARIFF 2**, the maximum decrease in the cost is around 73.37% in a time slot $h = 12$.

Fig. 4.17 shows the dependency of the total energy usage cost on the penetration rate of prosumers in the microgrid in a time slot $h = 12$. It can be observed that with the increasing number of prosumers (participants who can produce PV energy), the total energy usage cost in the microgrid decreases because more energy produced in the microgrid, which is sold and bought in the microgrid at a price λ_{e_h} instead of buying it from the utility grid at a price λ_{b_h}.

To assess the proposed pricing model, it is also compared with the existing works. The total energy usage cost of all prosumers in the microgrid with DR using the method proposed

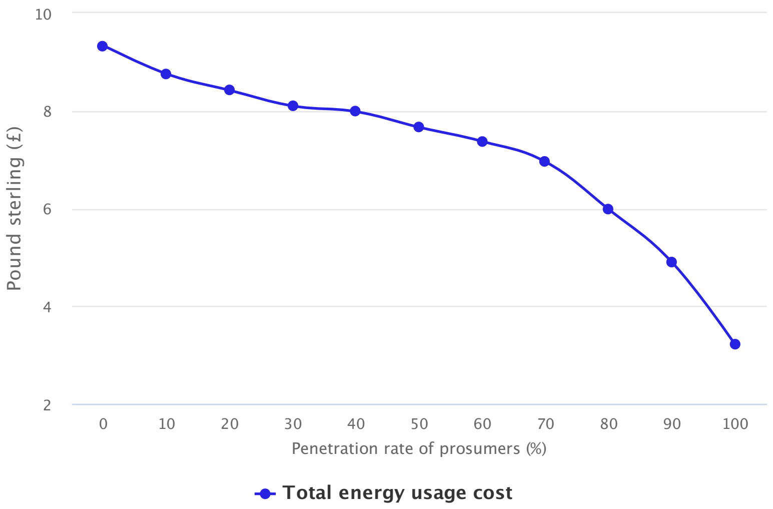

Figure 4.17: Dependency of the total energy usage cost on the penetration rate of prosumers in the microgrid in a time slot $h = 12$

in [28] is equal to 88.13% of the energy usage cost in Peer-to-Grid (P2G) energy trading, which means that the decrease in the total energy usage cost is around 11.87%. The cost for each prosumer can be decreased maximum by 23.77% using the approach in [101]. The model in [102] can achieve the reduction in the cost of around 18.4% for low-demand customers only, whereas the approach in [103] achieves the decrease of around 11%. In contrast, the maximum reduction in the total energy usage cost of all prosumers using the proposed approach that was observed is around 73.37%. It should be noted that the proposed approach takes into account the energy transfer losses, which means that the energy usage cost for a prosumer includes the cost to cover the energy transfer loss.

Sensitivity analysis

This section presents the sensitivity analysis of the proposed pricing model. First, the dependency of the energy usage cost for a prosumer on the absolute deviation from it's predicted energy consumption profile is analyzed. Secondly, the dependency of the total energy usage cost on the utility grid's buying and selling prices is assessed.

Fig. 4.18 shows the dependency of the total energy usage cost for a randomly chosen prosumer on it's absolute deviation from the predicted energy consumption. It can be

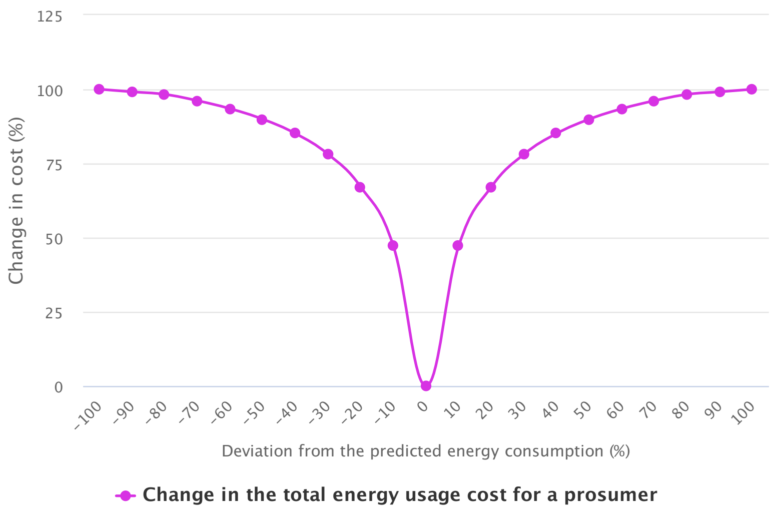

Figure 4.18: Dependency of the total energy usage cost for a randomly chosen prosumer on it's absolute deviation from the predicted energy consumption

seen that when the prosumer does not deviate from it's predicted energy consumption profile ($\gamma_{n_h}^{QD} = 0$), there is no change in the total energy usage cost. In contrast, with the increasing absolute deviation from the predicted energy consumption, the absolute deviation grows that leads to the increased total energy usage cost for a prosumer, which can be observed when the actual energy consumption decreases ($\gamma_{n_h}^{QD}$ changes from -10% to -100%) or increases ($\gamma_{n_h}^{QD}$ changes from 10% to 100%). For example, a prosumer n was expected to consume 1 kW of energy in a time slot h, whereas it's actual energy consumption is 1.1 kW ($\gamma_{n_h}^{QD} = +10\%$), which means that the total energy usage cost will be 50% higher compared to the scenario when a prosumer does not deviate from the predicted energy consumption. The same energy usage cost will be if a prosumer's actual consumption is 0.9 kW ($\gamma_{n_h}^{QD} = -10\%$). Similarly, if the difference between the predicted and actual energy consumption reaches the level of 30%, the total energy usage cost will be 75% higher. Thus, the more a prosumer deviates from it's predicted energy consumption profile, the more the total energy usage cost for this prosumer.

Fig. 4.19 shows how the change in the utility grid's buying (λ_{b_h}) and selling (λ_{s_h}) prices affects the total energy usage cost for all prosumers in the microgrid. First, the blue curve shows how the total energy usage cost for all prosumers depends on the change in the

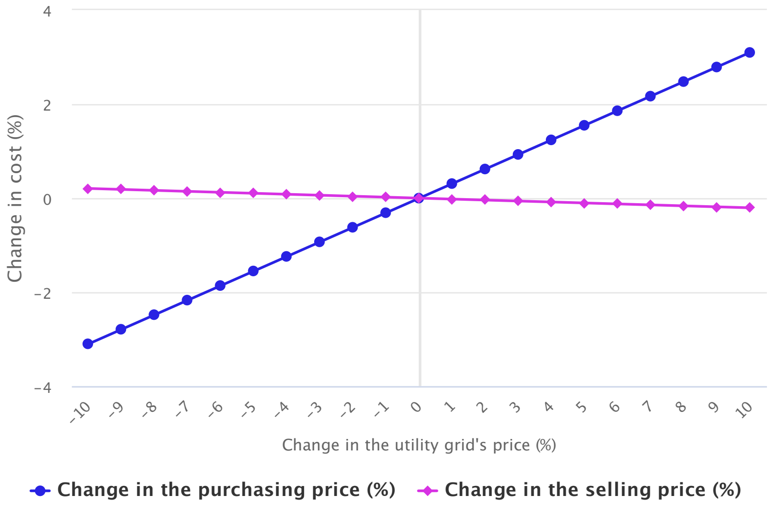

Figure 4.19: Dependency of the total energy usage cost for all prosumers on the utility grid's buying and selling prices

utility grid's buying price of energy, while the selling price does not change. It can bee seen that when the buying price of energy decreases, the total energy usage cost decreases. For example, if the buying price decreases by 5%, the total energy usage cost decreases by around 1.7%. When the buying price increases, the total energy usage cost also increases. For example, if the buying price increases by 10%, the total energy usage cost increases by around 3%. Secondly, the pink curve shows how the total energy usage cost depends on the change in the utility grid's selling price, while the buying price does not change. It can be observed that when the selling price decreases, the total energy usage cost increases, whereas with the increase in the selling price, the total energy usage cost decreases. It can be observed that the change in the buying price affects the total energy usage cost more than the change in the selling price. It happens because in a scenario, when $SDR_h < 1$, the shortage is bought from the utility grid, which increases the total energy usage cost, whereas the change in the selling price mainly affects the profit from selling energy. Due to the fact that the excess energy may be sold to the utility grid only when $SDR_h > 1$, most of the time the energy is bought from the utility grid to cover a part of the microgrid's demand or to cover the losses. Thus, the change in the buying price has more influence on the total energy usage cost of all prosumers.

4.4 Summary

The increasing penetration of DER, particularly PV systems, has given rise to a new class of energy consumers—prosumers—who can both consume and produce electricity. In this evolving energy landscape, the efficient utilization of locally produced renewable energy is an important aspect for reducing dependency on utility grids and fostering energy sustainability. This chapter addressed the challenge of incentivizing energy sharing among prosumers within a microgrid by proposing a novel dynamic pricing model that improves local energy utilization.

The first major contribution of this chapter is the formulation of a dynamic pricing mechanism that provides more favorable buying and selling prices within the microgrid compared to the main grid. By introducing an internal equilibrium price based on the SDR, the model enables effective pricing for prosumers and reduces the need for interaction with the main grid. The equilibrium price adjusts dynamically, creating financial incentives for prosumers to participate in energy trading within the microgrid.

Secondly, the chapter presents a detailed cost and profit allocation strategy that accounts for deviations from predicted levels of energy consumption and production. The model includes penalty mechanisms based on the prosumer's contribution to the total deviation, thereby encouraging predictability and stability within the microgrid. Furthermore, the model incorporates energy transfer losses and allocates additional cost equitably among participants, ensuring a fair and realistic assessment of energy costs and profits.

Through a structured approach that differentiates between three operational scenarios—standalone, seller, and buyer states—the model demonstrates flexibility and adaptability to real-world microgrid dynamics. Mathematical formulations for calculating total cost and profit for a prosumer provide a practical blueprint for implementation.

In conclusion, the research in this chapter has presented a new dynamic energy trading framework for a microgrid of prosumers, incorporating cost and profit allocation algorithms, penalty mechanisms, and loss-aware cost allocation to enhance prosumer engagement and improve the overall efficiency and sustainability of localized energy systems.

Chapter 5

A blockchain-Based Data-Aggregation scheme for a Microgrid of Prosumers

5.1 Overview

Due to the ever-increasing number of end-users in electricity networks, as well as the growing number of different devices are used, the demand for electricity is constantly growing, whereas the management and maintaining of electricity networks is becoming more challenging.

With recent developments and advancements in ICT, the concept of a microgrid has come to the limelight, which is considered as a promising solution for the management of local groups of electricity users and maintaining local electricity networks. The idea of a microgrid brings the opportunities to achieve efficient energy utilization within a local community by using the end-users' energy usage data for energy supply scheduling and planning, as well as to provide more beneficial prices for local participants compared to the main grid.

The main source of energy usage data of a prosumer is a smart meter that measures the amount of energy consumed and produced. The use of smart meters' data opens up new possibilities to increase the efficiency of energy utilization, supply planning and maintaining a microgrid network. While this benefits all parties in a microgrid, there are some concerns about the privacy of the end-users' energy usage data. In addition, to provide better services to the end-users, the aggregated energy usage data should be very accurate, which is challenging to achieve by utilizing existing approaches. Blockchain technology has become a promising solution to ensure privacy of end-users' data. Thus, it is valuable to utilize blockchain when developing privacy-preserving data-aggregation schemes for a microgrid.

This chapter addresses three specific design goals: privacy, accuracy, and scalability. The scheme is designed to ensure privacy by preventing the exposure of raw energy readings to any single entity, including edge nodes and the microgrid operator. Accuracy is critical to ensure that aggregated totals remain mathematically precise despite the obfuscation techniques. Finally, scalability is prioritized to prevent the blockchain size from growing unmanageably. The evaluation framework assesses these goals through a simulation of an edge-enabled microgrid. Scalability and efficiency are demonstrated by measuring the average computation time for data reporting and aggregation, as well as monitoring the growth of the blockchain size relative to the number of prosumers. Privacy and accuracy are validated by verifying that the aggregated energy consumption derived from the blockchain exactly matches the sum of the original inputs, while confirming that individual inputs remain concealed.

5.2 Proposed Model

In this section, a novel blockchain-based data-aggregation scheme for edge-enabled microgrid of prosumers is proposed.

5.2.1 System Model

Fig. 5.1 shows the structure of a microgrid consisting of N Edge Nodes (EN), M households, a MGO, and a TA that is responsible for generating and securely distributing cryptographic keys, specifically Public and Private key pairs, among all participants in a microgrid. Each edge node has its own local database (DB) to store off-chain transactions, whereas all edge nodes form a blockchain network to store data on-chain. Each end-user (household) consumes and produces energy; thus, all end-users in a microgrid are prosumers (consumers, who can consume and produce energy). A smart meter is deployed at each prosumer to measure the amount of energy consumed and produced. Each smart meter can communicate with any edge node by invoking a smart contract that is deployed at every edge node to report prosumer's energy usage data. Each edge node stores received energy usage data from prosumers in an off-chain local storage (DB), whereas at the end of each time interval, the energy usage data are selected from the DB, aggregated and written to the blockchain.

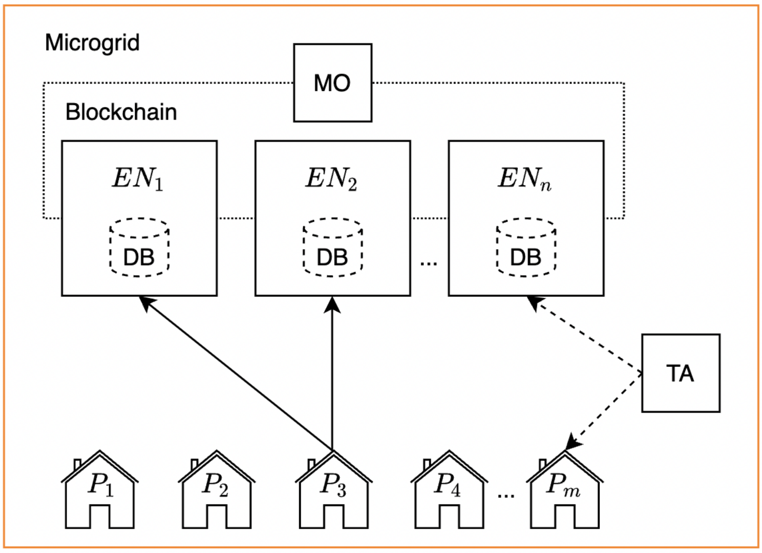

Figure 5.1: System model of a microgrid

In this work, a microgrid consists of N edge nodes and M prosumers. Let $\mathcal{N} = \{1, 2, 3, ..., N\}$ denote the set of edge nodes in the microgrid, where n is the index of the edge node and $n \in \mathcal{N}$, whereas, the total number of edge nodes is given by $N \triangleq |\mathcal{N}|$. Let $\mathcal{M} = \{1, 2, 3, ..., M\}$ denote the set of prosumers in the microgrid, where m is the prosumer index and $m \in \mathcal{M}$, whereas, the total number of prosumers is given by $M \triangleq |\mathcal{M}|$. Prosumers report smart meter readings to the edge nodes during a time interval (time slot), where each time slot is of one hour. Let $\mathcal{H} = \{1, 2, 3, ..., H\}$ denote the set of all time slots, where h is the time slot index and $h \in \mathcal{H}$, whereas, the total number of time slots is given by $H \triangleq |\mathcal{H}| = 24$.

Let QD_m denote the energy consumption profile for a prosumer m for one day, and it is defined as follows:

$$QD_m = \{QD_{m_1}, QD_{m_2}, ..., QD_{m_H}\}, \quad m \in \mathcal{M} \tag{5.1}$$

where QD_{m_h} is the level of energy consumption for a prosumer m in a time slot h and $QD_{m_h} \in QD_m$.

In this work, the only aggregation of energy consumption data is considered because the energy production data is aggregated in the same way.

5.2.2 Proposed Scheme

In this section, a novel privacy-preserving data-aggregation scheme for edge-enabled microgrid of prosumers is presented. In this work, there are N edge nodes in a microgrid that participate in the blockchain. Each edge node has its own local database (off-chain storage) to store prosumers' energy usage data during a time slot h, after which the data are aggregated and written to the blockchain, and the energy usage data are removed from the local database. Thus, the size of the data stored in the local database does not grow, whereas the size of the blockchain does not grow too fast. A smart contract that implements the logic of the proposed scheme is deployed at each edge node. Prosumers communicate with the edge nodes by invoking the smart contract's methods.

The algorithm of the proposed data-aggregation scheme consists of seven stages (Fig.5.2):

Stage 1: Initialization.
First of all, TA generates unique Elliptic Curve Cryptography (ECC) key pairs for all participants. The private keys are distributed to the Smart Meters and Edge Nodes via a secure off-chain channel during the device deployment phase. These keys allow end-users to digitally sign transactions and authenticate themselves when communicating with the edge nodes by invoking the smart contract. The information about all edge nodes is stored

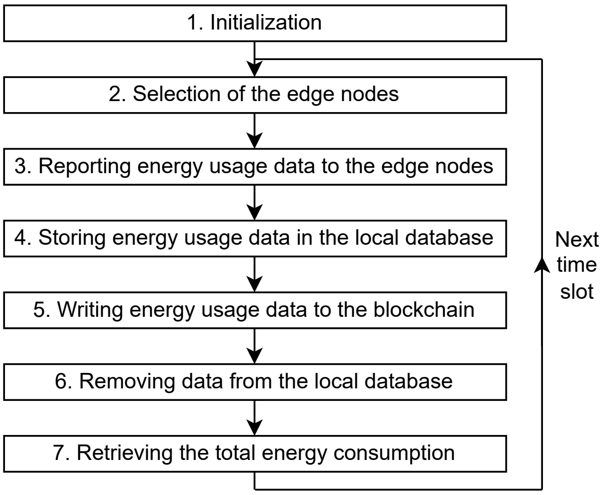
1. Initialization
2. Selection of the edge nodes
3. Reporting energy usage data to the edge nodes
4. Storing energy usage data in the local database
5. Writing energy usage data to the blockchain
6. Removing data from the local database
7. Retrieving the total energy consumption
Next time slot

Figure 5.2: Algorithm of the proposed scheme

in the blockchain, whereas each record with the information about the edge node contains ID (edge node ID) and PORT (port number to which requests are sent from the smart meters).

Stage 2: Selection of the edge nodes.

Before reporting the energy usage data to the edge nodes, each prosumer randomly selects two edge nodes to which the energy usage data will be sent. First, a smart meter invokes a smart contract's method *GetAllEdgeNodes* to get all available edge nodes, and then two edge nodes are randomly selected from the list of all available edge nodes. It should be noted that a smart meter randomly selects two edge nodes in each time slot; thus, the edge nodes, to which the energy usage data are sent may be different in all time slots.

Stage 3: Reporting energy usage data to the edge nodes.

Initially, a smart meter that is deployed at each prosumer, measures the amount of energy consumed during a time slot h (one hour). At the end of each time slot, a smart meter prepares a value (QD_{m_h}) that represents the amount of energy consumed during a time slot h. Before reporting prosumer's energy usage data to the edge nodes, a smart meter applies some modifications to the smart meter readings.

Let r_{m_h} denote the random value that is generated by a smart meter at the end of a time slot h before sending prosumer's energy consumption data QD_{m_h} to the edge nodes. The random value r_{m_h} is generated as follows:

$$r_{m_h} = rand(min, QD_{m_h}) \tag{5.2}$$

where min is the minimum value of r_{m_h}, which is equal to 0 in this work, whereas QD_{m_h} is the maximum value of r_{m_h}. Thus, r_{m_h} is randomly generated within the interval $(0, QD_{m_h})$.

Next, a smart meter generates two different values that will be sent to the edge nodes. Let $QD^{+}_{m_h}$ denote the amount of energy consumed by a prosumer m in a time slot h, which is calculated by adding a random value r_{m_h} to the original consumption value as follows:

$$QD^{+}_{m_h} = QD_{m_h} + r_{m_h} \tag{5.3}$$

Let $QD^{-}_{m_h}$ denote the amount of energy consumed by a prosumer m in a time slot h, which is calculated by subtracting a random value r_{m_h} from the original consumption value as follows:

$$QD^{-}_{m_h} = QD_{m_h} - r_{m_h} \tag{5.4}$$

Instead of sending the original consumption data, a smart meter reports two different values $(QD^{+}_{m_h}, QD^{-}_{m_h})$ to randomly selected edge nodes. To send the energy usage data to the edge nodes, a smart meter invokes the smart contract's method *ReportEnergyUsageData*, whereas the first request is sent to the first randomly selected edge node with the value $QD^{+}_{m_h}$, while the second request is sent to the second edge node with the value $QD^{-}_{m_h}$. Thus, other parties do not know what value ($QD^{+}_{m_h}$ or $QD^{-}_{m_h}$) a prosumer sends, and to what edge node the value is sent. Even if an attacker intercepts a message sent to the edge node, the original value of energy consumption (QD_{m_h}) cannot be restored from the modified value ($QD^{+}_{m_h}$ or $QD^{-}_{m_h}$).

Stage 4: Storing energy usage data in the local database.

When an edge node receives the energy usage data from a smart meter, it stores received data in the local database. Let $QD^{s}_{m_h}$ denote the amount of energy consumed that is generated from the original level of consumption QD_{m_h} in a time slot h by adding or subtracting the random value r_{m_h}. It should be noted that even the edge node does not know what value it receives ($QD^{+}_{m_h}$ or $QD^{-}_{m_h}$). Thus, the original energy consumption of a prosumer in a time slot h cannot be revealed during the eavesdropping or man-in-the-middle attacks. Finally, the edge node stores a tuple (h, $QD^{s}_{m_h}$) in the local database.

Stage 5: Writing energy usage data to the blockchain.

At the end of each time slot h, each edge node aggregates the energy usage data stored in the local database. Let $QD^{T}_{n_h}$ denote the aggregated value of all smart meter readings received by the edge node n during a time slot h, which is calculated by executing a query to the local database as follows:

$$QD^{T}_{n_h} = \sum QD^{s}_{m_h} \tag{5.5}$$

Next, the edge node n invokes a smart contract's method *StoreAggregatedConsumption* with the value of $QD^{T}_{n_h}$. In this way, at the end of each time slot h, the smart contract's method *StoreAggregatedConsumption* is invoked N times (one time by each edge node). Under the hood, the smart contract adds a new record to the smart contract's state that contains a time slot h and the aggregated energy usage data $QD^{T}_{n_h}$.

Stage 6: Removing data from the local database.

After the aggregated data are written to the blockchain, each edge node n deletes energy usage data stored in the local database. Thus, the size of the local database at each edge node does not grow along with the time. After deleting data from the local database, process starts all over again (from ***Stage 2***) for the next time slot ($h + 1$).

Stage 7: Retrieving the total energy consumption.

Since all edge nodes add their own records (h, $QD^{T}_{n_h}$) to the smart contract's state at ***Stage***

5, the total energy consumption cannot be just queried from the smart contract's state. Let QD_h^T denote the total energy consumption in a microgrid in a time slot h, which is calculated by accessing the data in the smart contract's state for a particular time slot h as follows:

$$QD_h^T = \frac{\sum_{n=1}^{N} QD_{n_h}^T}{2} \tag{5.6}$$

It should be noted that it is necessary to devide the sum of all aggregated values by 2 because each prosumer reports its original energy consumption two times ($QD_{m_h}^+$ and $QD_{m_h^-}$). Thus, the total energy consumption in a microgrid in a time slot h can be expressed as follows:

$$QD_h^T = \frac{\sum_{n=1}^{N} \sum QD_{m_h}^s}{2} \tag{5.7}$$

By combining (5.7), (5.3), and (5.4), QD_h^T can be represented as follows:

$$QD_h^T = \frac{\sum_{m=1}^{M}((QD_{m_h} + r_{m_h}) + (QD_{m_h} - r_{m_h}))}{2} \tag{5.8}$$

To make it clear that the data stored in the blockchain represent the total energy consumption of all prosumers in a microgrid, (5.8) can be rewritten as follows:

$$QD_h^T = \frac{2 * \sum_{m=1}^{M} QD_{m_h} + \sum_{m=1}^{M}(r_{m_h} - r_{m_h})}{2} \tag{5.9}$$

It can be seen from (5.9) that the microgrid operator can get the total energy consumption of a microgrid for any time slot by invoking the smart contract's method *GetMicrogrid-Consumption* that returns the actual level of energy consumption in a requested time slot.

By utilizing the proposed data-aggregation scheme in a microgrid, the prosumers' energy usage data are not revealed to other parties, while the microgrid operator can get the total energy consumption in a microgrid, which is equal to the sum of the original levels of consumption of all prosumers without added noise. Thus, the proposed scheme ensures privacy of prosumers, while achieving high level of accuracy.

5.2.3 Threat Model and Assumptions

The proposed blockchain-based data aggregation scheme aims to preserve the privacy of individual prosumers' energy consumption profiles while ensuring the integrity of aggregated values stored on the blockchain. In this subsection, the participating entities, their behaviour, the considered adversarial capabilities, and the assumptions under which the scheme is analysed are defined.

Participating entities and behavioural assumptions

- **Trusted Authority (TA) and Microgrid Operator (MGO).** The TA and MGO are assumed to be fully trusted entities that correctly execute the protocol. The TA generates the cryptographic keys and public parameters, and the MGO reads the aggregated energy usage from the blockchain for operational purposes. Neither the TA nor the MGO is compromised or malicious.
- **Smart meters.** Smart meters are trusted devices that honestly measure each prosumer's energy consumption and correctly implement the randomisation procedure in Stage 3 of the scheme. They generate random values and compute the modified readings. The random values remain secret and are not accessible to other entities.
- **Edge nodes.** Edge nodes are modelled as honest-but-curious: they follow the protocol, correctly store the received values in their local databases, perform aggregation per time slot, and invoke the smart contract to write aggregated values to the blockchain. However, they may attempt to infer individual prosumers' consumption values from the data available to them. It is assumed that edge nodes do not collude when they

receive shares corresponding to the same prosumer and time slot.

- **Communication channels.** Communication between smart meters and edge nodes is authenticated using the ECC key pairs generated by the TA, such that an adversary cannot forge messages on behalf of another participant. It is assumed that standard transport-layer encryption is used; therefore, the focus is on privacy properties provided by the masking mechanism rather than on channel security.

Adversarial goals and capabilities.

The goals of adversaries are as follows:

- to reconstruct an individual prosumer's original energy consumption from intercepted messages, edge-node databases, or blockchain data; and/or
- to tamper with the aggregated values stored on the blockchain.

Adversaries may perform the following actions:

- passively eavesdrop on all network traffic between smart meters, edge nodes, and the MGO;
- observe all data stored on the blockchain; and
- compromise at most one edge node in a given time slot, gaining full access to its local database and blockchain view.

Under these capabilities, each edge node observes only one of the masked values (either $QD^{+}_{m_h}$ or $QD^{-}_{m_h}$) for a given prosumer and time slot. Communication between smart meters and edge nodes is encrypted and authenticated, and we assume that at most one non-colluding edge node can be compromised. Consequently, no single adversary can access both $QD^{+}_{m_h}$ and $QD^{-}_{m_h}$ for the same prosumer and time slot and, without knowledge of the corresponding random value r_{m_h}, the original consumption QD_{m_h} cannot be reconstructed.

Collusion and Sybil assumptions.

If both edge nodes that receive $QD^{+}_{m_h}$ and $QD^{-}_{m_h}$ for the same prosumer and time slot collude, they can reconstruct QD_{m_h} by combining the two shares. Detecting or mitigating such collusion is not addressed in this work and is treated as a deployment-time assumption; designing mechanisms to tolerate collusion is left for future work.

In the proposed scheme, a permissioned blockchain is used, in which edge nodes are registered and authenticated by the TA. This registration process prevents an adversary from creating an unbounded number of fake edge-node identities; therefore, Sybil attacks at the blockchain layer are not considered. Attacks on the underlying consensus protocol (e.g., majority attacks) and denial-of-service attacks on communication links are beyond the scope of this work.

5.3 Evaluation and Discussion

The simulations are conducted using a permissioned blockchain with smart contracts and off-chain MySQL databases deployed in Docker containers. The microgrid consists of 20–100 prosumers and a fixed number of edge nodes; prosumers send masked readings once per hour over a 24-hour period. For each configuration, the average latency is measured for four operations: (i) selecting two edge nodes via the smart contract, (ii) reporting masked readings to the selected edge nodes (including database writes), (iii) aggregating local readings and writing aggregates to the blockchain, and (iv) deleting raw readings from the local database. The resulting number of blocks in the blockchain is recorded to measure ledger growth.

Fig. 5.3 shows the average time it takes for a prosumer to randomly select two edge nodes, to which the energy usage data will be sent.

It can be seen that the average time to randomly select two edge nodes, which implies invoking the smart contract's method *GetAllEdgeNodes* and random selection of two edge

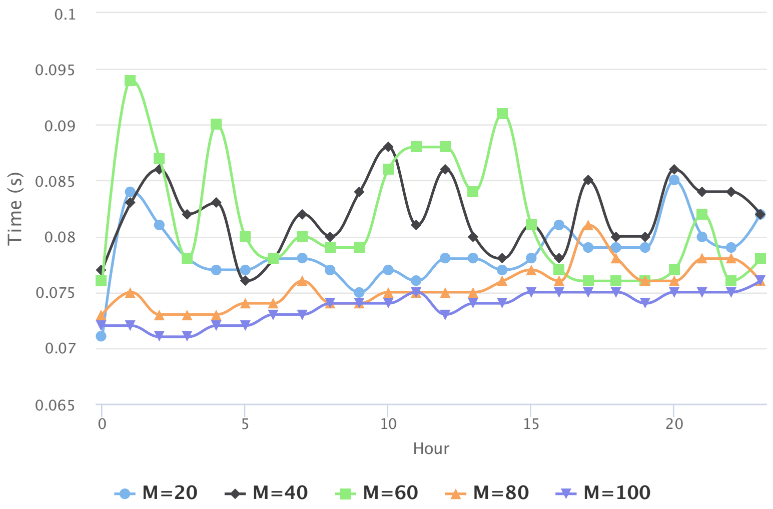

Figure 5.3: Average time for the random selection of two edge nodes for a different number of prosumers (M)

nodes from the list of the available edge nodes, does not exceed 0.2 seconds on average, whereas there are some peaks can be observed that caused by the delays in communication between Docker containers.

Fig. 5.4 shows the average time it takes for a prosumer to report energy usage data to the randomly selected edge nodes.

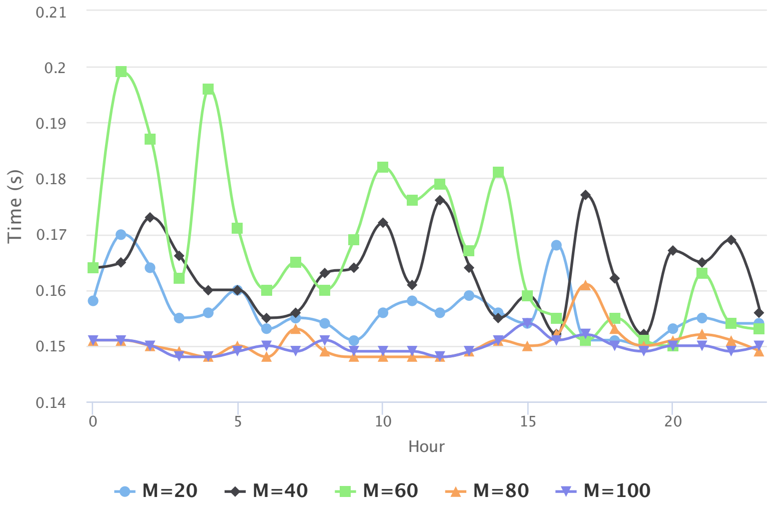

Figure 5.4: Average time for reporting energy usage data to the edge nodes for a different number of prosumers (M)

Reporting the energy usage data to the edge nodes implies invoking the smart contract's method *ReportEnergyUsageData* two times by each prosumer, whereas the smart contract stores received data in the local database (MySQL); thus, the time is measured for complete execution of the requests to the smart contract including the time it takes to write data to MySQL. It can be seen that the average time it takes for one prosumer to submit the energy usage data to the edge nodes changes with slight fluctuations but does not exceed 0.5 seconds.

Fig. 5.5 shows the average time it takes for one edge node to aggregate energy usage data, namely to retrieve the data from the local database and write the aggregated data to the blockchain.

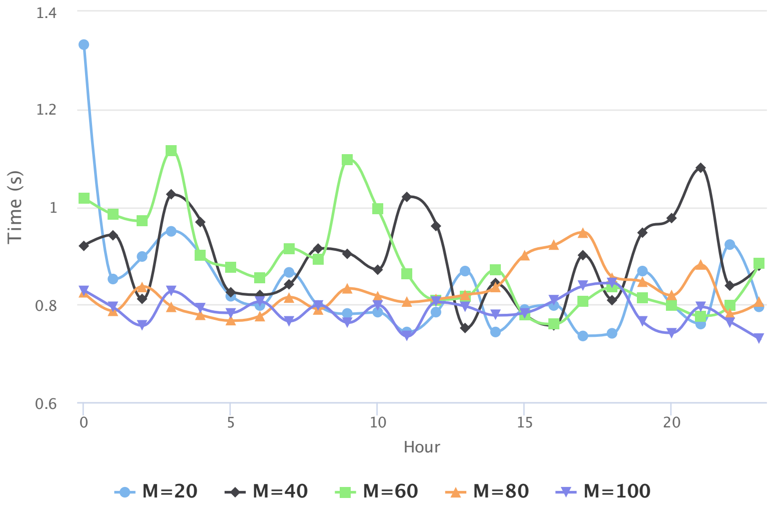

Figure 5.5: Average time for aggregating the energy usage data and writing the aggregated data to the blockchain for a different number of prosumers (M)

Each edge node queries the local database (MySQL) to get the sum of all received smart meter readings from prosumers in a time slot h, and then invokes the smart contract's method *StoreAggregatedConsumption* to store the aggregated data in the blockchain. It can be seen that the average time for querying the local database and invoking the smart contract for each edge node fluctuates slightly but does not exceed 3 seconds on average, whereas there are three peaks (around 4.2 and 9.5 seconds) can be observed that may be explained by the delay in the communication between Docker containers.

Fig. 5.6 shows the average time it takes for each edge node to delete the energy usage data from the local database (MySQL).

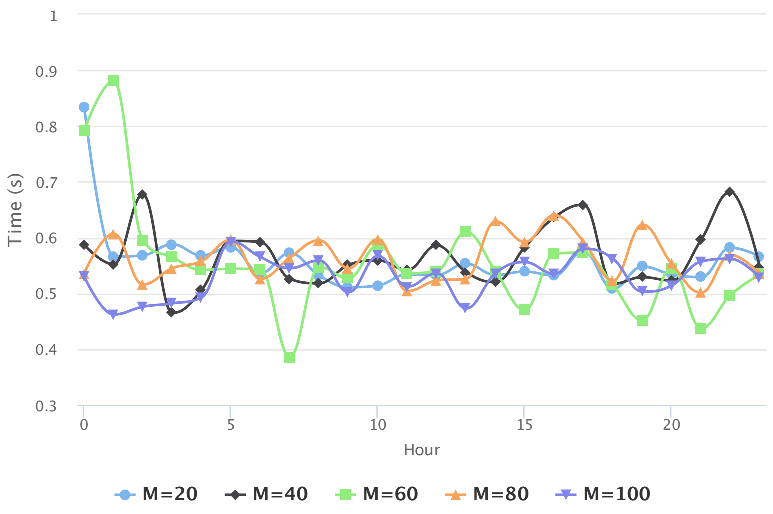

Figure 5.6: Average time for deleting the energy usage data from the local database for a different number of prosumers (M)

It can be seen that the average time it takes for each edge node to delete the energy usage data from the local database does not exceed 2 seconds on average. There are some peaks can be observed that may happen from time to time that depend on the configuration and the type of the off-chain storage selected for the local database.

One of the most important indicators to evaluate a blockchain-based application is the size of the blockchain, namely the number of blocks. Fig. 5.7 shows the average number of blocks in the blockchain for a different number of prosumers for one day.

It can be seen that the number of blocks in the blockchain does not depend on the number of prosumers in a microgrid, whereas it depends on the number of edge nodes because each edge node invokes the smart contract's method *StoreAggregatedConsumption* to write the aggregated data in the blockchain at the end of each time slot; thus, to write the aggregated energy usage data to the blockchain, the smart contract is invoked N times in each time slot. Hence, increasing the number of prosumers in a microgrid does not affect the size of the blockchain.

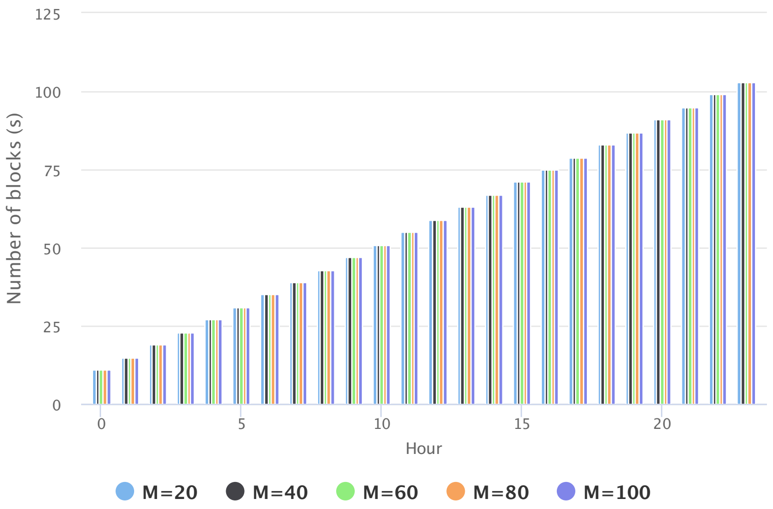

Figure 5.7: Number of blocks in the blockchain for a different number of prosumers (M)

5.4 Summary

The evolution of smart energy systems and the growing number of prosumers in microgrids demand privacy-preserving mechanisms for data aggregation. While smart meters enable fine-grained energy usage monitoring, conventional data aggregation approaches raise concerns around data privacy and the efficiency of data handling. This chapter presented a blockchain-based data aggregation scheme specifically designed for edge-enabled microgrids of prosumers.

The proposed model introduces a decentralized architecture consisting of smart meters, edge nodes, a microgrid operator, and a trusted authority. Each smart meter reports modified energy consumption values to two randomly selected edge nodes, ensuring that the original consumption data cannot be revealed during transmission or storage. By distributing these values and aggregating them at the edge nodes, the system ensures privacy while maintaining accurate aggregation of energy consumption data. The proposed approach ensures that energy usage data is stored off-chain temporarily and only aggregated values are committed to the blockchain, optimizing both privacy and storage efficiency. Each edge node aggregates received values per time slot and contributes to the blockchain with minimal performance overhead.

In conclusion, the research in this chapter has presented a new privacy-preserving data-aggregation scheme for microgrids, leveraging blockchain, off-chain storage, and edge computing to ensure data confidentiality, reduce storage overhead, and maintain high aggregation accuracy—thereby promoting reliable and efficient energy management in decentralized smart grid environments.

The proposed blockchain-based data-aggregation scheme complements and extends prior data aggregation methods for smart grids summarized in Section 2.3. In contrast to DP-based schemes that inject Laplace noise into smart meter readings and thus trade off privacy against the accuracy of aggregated values [41, 42], the proposed design preserves exact aggregation, as only masked values and their sums are exposed to the microgrid operator. Compared to fragmentation- and multi-party computation-based solutions [43, 44], which rely on complex data segmentation and ring signatures whose communication and computational overhead grow with the number of participants, the dual-reporting mechanism uses only simple masking operations at each smart meter and constant-size reports per time slot. Finally, unlike blockchain-based architectures that either store fine-grained meter readings on-chain or require each smart meter to act as a blockchain node—leading to unbounded ledger growth and heavy consensus participation—the proposed scheme stores only per-slot aggregates on-chain. As a result, blockchain size scales with the number of edge nodes and time slots rather than the number of prosumers, delivering strong privacy guarantees with low smart-meter overhead.

Chapter 6

Privacy-preserving User-Centric Data-Sharing Scheme for a Microgrid of prosumers

6.1 Overview

The variety of sensitive data generated by end-users grows over time while the number of data-driven services also increases gradually. To use a service provided by a third party, customers need to share their sensitive information. For example, households need to share their electricity consumption data with the energy supplier that calculates the electricity bill. Another example is when patients share their health information with a hospital to predict genetic diseases [56]. The number of use cases where sensitive end-user data may be used is enormous and constantly expanding.

Although having access to sensitive end-user data benefits both parties - the end-user gets more accurate results, and a service provider can better manage its business - it also poses some privacy-related issues. For example, access to electricity consumption profiles of households can help local authorities make better decisions about where to install charge points for electric vehicles [104], whereas by analyzing these data, an adversary may gain some knowledge about the customers behavior [105]. According to regulations such as GDPR [106] and Health Insurance Portability and Accountability Act (HIPPA) [107], end-users must have control over the use of their sensitive data. Thus, it is essential to preserve end-user privacy when sharing data with third parties.

DP has become one of the most widely used techniques to preserve privacy by injecting noise into sensitive data. The main idea of DP is that end-users report sensitive data to an aggregator through an encrypted channel. Then, the aggregator adds controllable noise to the aggregated data and shares it with a service provider for processing. The most notable problem of DP is that end-users still share raw data with a third party, which may cause a leakage of sensitive information in case of a key leakage or insider attack on the aggregator's side. Secondly, it cannot be used when a service provider needs access to each end-user's individual data to provide personalized services.

LDP is a variation of DP that addresses the issue of sharing raw data with an aggregator. LDP enables end-users to inject noise into the data on their side before sharing it with a third party. As a result, in the event of a key leakage or insider attack on the third party's side, an adversary would only gain access to the noisy data.

Many DP and LDP-based data-sharing schemes utilize the Laplace randomized mechanism to perturb sensitive data with noise drawn from the Laplace probability distribution. These schemes allow end-users to specify the desired level of privacy, based on which a controllable amount of noise is injected into the data. However, they do not allow end-users to specify utility (the result accuracy), which could then determine the level of privacy. Instead of adjusting privacy levels directly, end-users might prefer to set limits (boundaries) on the extra cost incurred from added noise, i.e., the cost they are willing to pay for enhanced privacy. The privacy level would then be determined based on this extra cost or utility preference. This approach would make it easier for end-users to control utility while the level of privacy is automatically adjusted based on their utility preferences.

Another challenge in DP and LDP-based schemes is obtaining more privacy without reducing utility proportionally. In the Laplace mechanism, an intuitive approach to increase noise levels, and consequently enhance privacy for a given privacy budget, is to shift the distribution so that its mean is not centered around zero. However, this can lead to a significant reduction in utility. An alternative approach involves using a multimodal distribution, where random variables with opposite signs are drawn with equal probability. This means that similar amounts of noise with opposite signs, drawn from the probability distribution, can cancel each other out over time—a concept referred to as ’noise compensation’ in this work. Noise compensation can ensure that, over time, the accuracy of queries performed on noisy data does not significantly deviate from the results obtained by querying the original data. To address the above-mentioned challenges, this chapter proposes a novel user-centric data-sharing scheme that preserves privacy using LDP.

The main design goal of this scheme is to empower end users to define their own utility boundaries (acceptable error in utility) rather than rely on arbitrary privacy budgets. A secondary goal is to preserve utility through noise compensation, ensuring that while individual data points are perturbed for privacy, the long-term statistical utility remains high. Moreover, the scheme aims to provide the maximum possible level of privacy for the user-defined utility boundaries. The evaluation framework assesses these goals through a simulation of data sharing scenarios using real household energy data. The user-centricity and utility goals are evaluated by measuring the relative error in utility and confirming it stays within the user-defined boundaries. The privacy goal is assessed using the Mutual Information score to quantify the independence between the noisy and original data profiles, comparing the performance of the proposed scheme against the standard Laplace mechanism.

6.2 Preliminaries

In this section, the theoretical background and mathematical formulation of LDP, sensitivity of a function, and the Laplace mechanism are presented.

DP is one of the most popular privacy-preserving techniques, where a trusted centralized aggregator (data curator) accesses sensitive data of end-users, aggregates those data, and adds controllable noise to the aggregated data. In a real-world setting, it is challenging to determine whether a centralized aggregator operates honestly or not, as well as to guarantee that end-users’ sensitive data are not shared (accessed) with malicious actors during the aggregation process. LDP has become a solution to overcome these limitations of DP, so that end-users locally perturb their sensitive data using a LDP mechanism, after which end-users share noisy data with a centralized aggregator. Thus, an adversary cannot obtain the sensitive data of end-users.

Definition 1 (ϵ-local differential privacy). Let x and y denote two neighboring datasets, where y can be produced by adding, removing, or modifying exactly one entry from x. A randomized mechanism $\mathcal{M} : D \longrightarrow S$ satisfies ϵ-local differential privacy iff for any output $s \in S$, and two neighboring datasets $x, y \in D$:

$$\frac{Pr[\mathcal{M}(x) = s]}{Pr[\mathcal{M}(y) = s]} \leq e^{\epsilon} \tag{6.1}$$

where S is the set of all possible outputs that a mechanism $\mathcal{M}$ can produce, $Pr[\mathcal{M}(x) = s]$ is the probability of a randomized mechanism $\mathcal{M}$ outputting the result s given the input x, and ϵ is the privacy budget (level of privacy) that bounds the probability of $\mathcal{M}$ outputting the same result for any pair of neighboring datasets x, y [108]. A smaller value of ϵ provides stronger privacy guarantee, whereas large ϵ provides weak privacy guarantee [109].

The Laplace mechanism has become one of the most popular techniques to preserve end-users privacy. Let $f : D \longrightarrow \mathbb{R}^k$ denote the function that maps datasets (D) to real numbers. For example, $f(\cdot)$ may be a function that takes a dataset $x \in D$ as an input and calculates the mean $f(x) \in \mathbb{R}$. To introduce controllable noise to the result of $f(\cdot)$, the Laplace mechanism relies on the sensitivity ($\mathcal{L}_1$-sensitivity) of $f(\cdot)$.

Definition 2 ($\mathcal{L}_1$-sensitivity). Given a query function $f(\cdot)$, its $\mathcal{L}_1$-sensitivity Δf is the maximum $\mathcal{L}_1$ distance between the results of $f(\cdot)$ over any pair of neighboring datasets x and y, which is defined as follows [108]:

$$\Delta f = \max_{x,y} \|f(x) - f(y)\|_1 \tag{6.2}$$

The Laplace mechanism uses the Laplace probability distribution to generate noise. The

Probability Density Function (PDF) of the Laplace distribution centered around 0 with the scale factor $b = \Delta f / \epsilon$ is defined as follows [108]:

$$Lap(x|b) = \frac{1}{2b} exp\left(-\frac{|x|}{b}\right) \tag{6.3}$$

where the scale factor b is calibrated according to the Δf. The Laplace mechanism is ϵ-differentially private [42], and has the following definition:

$$\mathcal{M}_L\left(x, f(\cdot), \epsilon\right) = f(x) + Lap\left(\frac{\Delta f}{\epsilon}\right) \tag{6.4}$$

6.3 Proposed Model

Preserving end-users' privacy is important, but in real-world applications, it is equally important to maintain control over the utility of the protected data. The main question in the DP domain is recognized to be the trade-off between privacy and utility [110]. Adding more noise to the data increases the level of privacy and may decrease the utility. End-users may find it challenging to understand how added noise will affect data utility. As a result, selecting an optimal privacy budget (ϵ) becomes difficult. Instead of adjusting ϵ, end-users could simply set the maximum change (error) in utility they can bear, which is usually more palpable for them.

This chapter proposes a novel user-centric data-sharing scheme utilizing LDP. In the proposed scheme, end-users determine the boundaries for utility change, that is the relative error in utility they are willing to tolerate. The noise within the randomized mechanism is generated using a novel probability distribution that allows to determine positive and negative ranges of random variables from which the noise will be drawn with high probability. Keeping the amount of noise within the specified ranges helps to keep the relative error in utility within the boundaries. Thus, based on the end-user's utility preference, the proposed scheme adjusts its parameters to inject the right amount of noise and to provide the maximum possible level of privacy.

6.3.1 Boiarkin probability distribution

The proposed privacy-preserving scheme relies on a new probability distribution, named the Boiarkin distribution. When the Laplace distribution is used in differential privacy mechanisms, its mean (μ) is set to 0 [41, 42]. As a result, the noise values drawn from the distribution are centered around zero. In the Boiarkin distribution, we have two non-zero means with opposite signs. For this reason, the noise values drawn from the mechanism based on the Boiarkin distribution are higher than those from the Laplace-based mechanism, but the utility is maintained since the opposite signs compensate each other in the long term. Additionally, the Boiarkin distribution includes some hyperparameters that allow fine-tuning, providing more precise control over the generated values. The rationale behind the proposed probability distribution is to provide control over the probability of a random variable $r = 0$ to be generated ($Pr[r = 0]$), as well as to control the amount of noise drawn from the distribution by increasing the probability $Pr[r \neq 0]$.

The PDF of the Boiarkin probability distribution is defined as follows:

$$Boi(x|b, \psi, \mu) = q \cdot exp\left(-\frac{\left|\psi - |x - \mu|\right|}{b}\right) \tag{6.5}$$

where b is the scale of the distribution, ψ is the spread of the modes, μ is the shift, and q is the normalizing factor, which is defined as follows:

$$q = \frac{1}{2b\left(2 - exp\left(-\frac{\psi}{b}\right)\right)} \tag{6.6}$$

The Boiarkin distribution is a symmetric bimodal probability distribution centered at μ (mean of the distribution) and has two modes, namely $\lambda_1 = -\psi$ and $\lambda_2 = \psi$. The purpose of having two modes is to determine the ranges from which random variables (with opposite signs) can be drawn with high probability. Since continuously generating non-zero random variables may seriously affect the utility, it is important to have the noise compensation feature. When using the Boiarkin probability distribution with increased $Pr[r \neq 0]$, more non-zero random variables will be generated around the modes λ_1 and λ_2. Since the negative and positive modes are on the same distance (spread ψ) from the mean $\mu = 0$, many non-zero random variables will be generated around λ_1 and λ_2 but with opposite signs, which means that they will compensate each other. Thus, the Boiarkin distribution enables the noise compensation feature, while providing control over the intervals from which random variables can be drawn with high probability.

Let c (Fig. 6.1) denote the probability of a random variable $r = 0$ to be generated ($c = Pr[r = 0]$), and q denote the probability of a non-zero random variable $r = \lambda_1 = \lambda_2$ to be generated ($q = Pr[r = \lambda_1] = Pr[r = \lambda_2]$). To be able to control the ratio of probabilities of generating zero (c) and non-zero (q) random variables, the ratio p ($p \in (0; 1]$) of these probabilities is introduced. The ratio p allows to fine-tune the amount of noise drawn from the probability distribution and is defined as follows:

$$p = \frac{c}{q} = exp\left(\frac{\lambda_1}{b}\right) \tag{6.7}$$

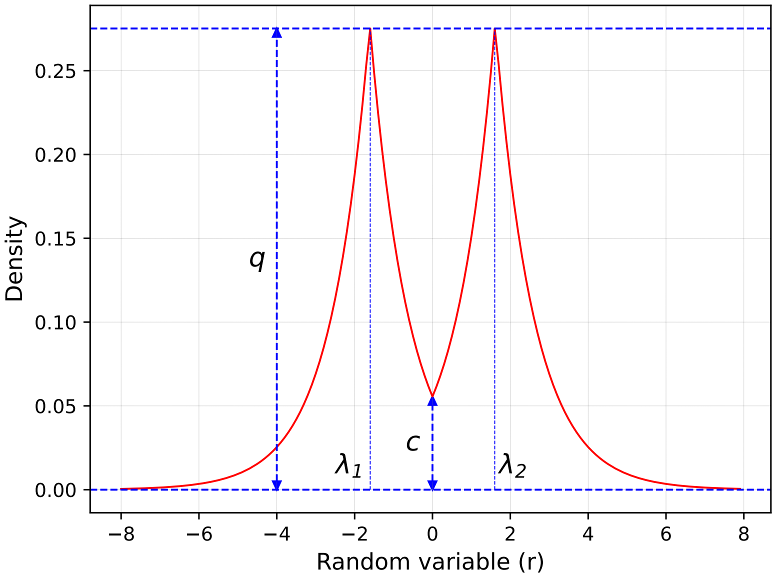

Figure 6.1: Probability density function of the Boiarkin probability distribution with scale factor $b = 1$ and $p = 0.2$

Fig. 6.1 shows the PDF of the Boiarkin probability distribution with the scale factor $b = 1$, and $p = 0.2$. It can be seen that the probabilities around the modes λ_1 and λ_2 are higher compared to the mean ($\mu = 0$). More precisely, the probabilities $Pr[r = \lambda_1]$ and $Pr[r = \lambda_2]$ are 5 times higher than $Pr[r = 0]$, which is controlled by p.

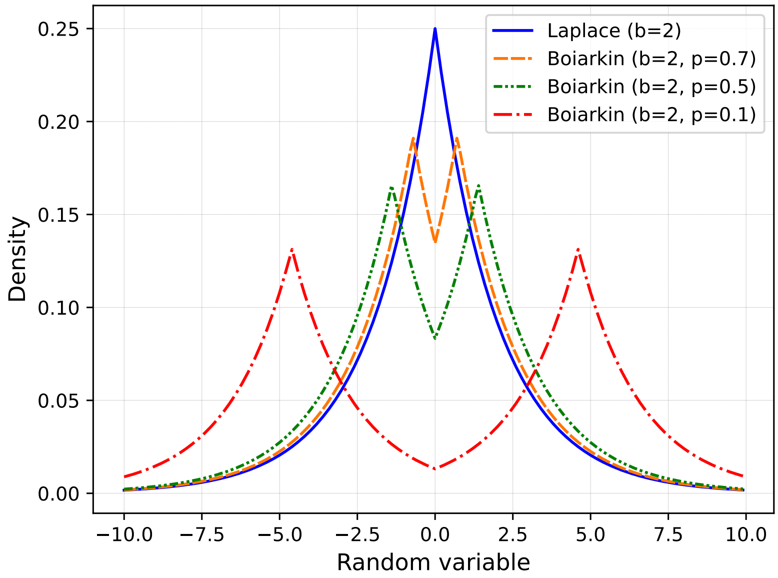

Figure 6.2: Probability density functions of the Laplace distribution with the scale factor $b = 2$ and Boiarkin distribution with the scale factor $b = 2$ and $p = 0.1, 0.5, 0.7$

Fig. 6.2 shows the PDF of the Laplace distribution with the scale factor $b = 2$ and Boiarkin distribution with the scale factor $b = 2$ and $p = 0.1, 0.5, 0.7$. By decreasing p, $Pr[r = 0]$ decreases, whereas the probabilities $Pr[r = \lambda_1]$ and $Pr[r = \lambda_2]$ increase. Thus, by decreasing p, more random variables are generated around λ_1 and λ_2. When $p = 1$, i.e., the $Pr[r = \mu] = Pr[r = \lambda_1] = Pr[r = \lambda_2]$, the Boiarkin distribution becomes the Laplace distribution with the scale factor b centered at μ.

6.3.2 Boiarkin mechanism

In this section, to preserve end-user privacy, a novel ϵ-LDP mechanism is introduced, which utilizes the proposed Boiarkin probability distribution to generate noise.

Definition 3 (The Boiarkin Mechanism). Given any function $f : D \to \mathbb{R}^k$, the Boiarkin mechanism is defined as follows:

$$\mathcal{M}_{\mathcal{B}}(x, f(\cdot), \epsilon, \psi) = f(x) + (Y_1, ..., Y_k) \tag{6.8}$$

where Y_i are i.i.d. random variables drawn from $Boi(\Delta f/\epsilon, \psi, 0)$ (6.5).

Theorem 1. *The Boiarkin mechanism $\mathcal{M}_{\mathcal{B}}(x, f(\cdot), \epsilon, \psi)$ preserves ϵ-LDP for any end-user with personal privacy budget ϵ.*

Proof. Let $x \in \mathcal{D}$ and $y \in \mathcal{D}$ be any neighboring datasets, differing in one entry, and $f(\cdot)$ be some function $f : \mathcal{D} \rightarrow \mathbb{R}^k$. Let $P_x(z)$ denote the probability density function of $\mathcal{M}_{\mathcal{B}}(x, f(\cdot), \epsilon, \psi)$, and let $P_y(z)$ denote the probability density function of $\mathcal{M}_{\mathcal{B}}(y, f(\cdot), \epsilon, \psi)$. To prove ϵ-local differential privacy [108], it is shown that the ratio $P_x(z)/P_y(z)$ is bounded by $exp(\epsilon)$ at any arbitrary point $z \in \mathbb{R}^k$.

$$\frac{P_x(z)}{P_y(z)} = \prod_{i=1}^{k} \left(\frac{exp\left(-\frac{\epsilon\big|\psi - |z_i - f(x)|\big|}{\Delta f}\right)}{exp\left(-\frac{\epsilon\big|\psi - |z_i - f(y)|\big|}{\Delta f}\right)} \right)$$

$$= \prod_{i=1}^{k} exp\left(-\frac{\epsilon}{\Delta f}\big|\psi - |z_i - f(x)|\big| + \frac{\epsilon}{\Delta f}\big|\psi - |z_i - f(y)|\big|\right)$$

$$= \prod_{i=1}^{k} exp\left(\frac{\epsilon}{\Delta f}\left(\big|\psi - |z_i - f(y)|\big| - \big|\psi - |z_i - f(x)|\big|\right)\right)$$

$$\leq \prod_{i=1}^{k} exp\left(\frac{\epsilon}{\Delta f}\big|\psi - |z_i - f(y)| - \psi + |z_i - f(x)|\big|\right)$$

$$= \prod_{i=1}^{k} exp\left(\frac{\epsilon}{\Delta f}\big||z_i - f(x)| - |z_i - f(y)|\big|\right)$$

$$\leq \prod_{i=1}^{k} exp\left(|z_i - f(x) - z_i - f(y)|\right)$$

$$= \prod_{i=1}^{k} exp\left(\frac{\epsilon}{\Delta f}|f(y) - f(x)|\right)$$

$$= exp\left(\frac{\epsilon}{\Delta f}\sum_{i=1}^{k}|f(y) - f(x)|\right)$$

$$= exp\left(\frac{\epsilon}{\Delta f}\|f(y) - f(x)\|_1\right)$$

$$\leq exp(\epsilon) \tag{6.9}$$

where the first and second inequalities are triangle inequalities, and the last inequality follows from the definition of sensitivity (6.2). Therefore, the Boiarkin mechanism ensures that the probability of a given outcome is nearly the same (bounded by $exp(\epsilon)$) for any neighboring datasets $x, y \in D$ differing in one entry.

6.3.3 User-centric data-sharing scheme

Most of DP privacy-preserving mechanisms imply the adjustment or use of a predefined privacy budget ϵ to control the trade-off between the level of privacy and utility. The question regarding this trade-off has been studied by many researchers [111–113], whereas there is still no clear answer on how to choose the right ϵ.

In this section, a novel user-centric privacy-preserving mechanism is proposed, which enables the end-users to choose an acceptable error in utility, whereas the level of privacy (an optimal privacy budget ϵ) is automatically calculated by the scheme. The proposed scheme utilizes the Boiarkin ϵ-LDP mechanism, which relies on the Boiarkin probability distribution to fine-tune the amount of generated noise. Let $U : D \rightarrow \mathbb{R}$ denote the utility function that maps input datasets (D) to real numbers, and it is defined as follows:

$$U(d) = g(d), \quad d \in D \tag{6.10}$$

where $g(\cdot)$ is the function that takes d as an input and outputs the result for a particular use case by performing mathematical operations on d. For example, $g(\cdot)$ may be a function that just returns d, or it may be a function that outputs the energy usage cost for a household by multiplying the average energy consumption per day (d) by the number of time slots and purchasing price for energy per kWh.

Let δ denote the maximum acceptable relative error in utility, which is expressed as follows:

$$\delta = \frac{U(d') - U(d)}{U(d)} \cdot 100\% \tag{6.11}$$

where d is the result of a query function $f(\cdot)$ (original data), and d' is the output of a randomized mechanism $\mathcal{M}(\cdot)$. Thus, (6.11) may be rewritten as follows:

$$\delta = \frac{U(\mathcal{M}(x)) - U(f(x))}{U(f(x))} \cdot 100\% \tag{6.12}$$

where x is the input dataset. Taking into account that a randomized mechanism adds noise to the result of a query function (6.8), (6.12) may be rewritten as follows:

$$\delta = \frac{U(f(x) + r) - U(f(x))}{U(f(x))} \cdot 100\% \tag{6.13}$$

where r is the noise drawn from a probability distribution used by a randomized mechanism. Expressing the variable r (6.13), the maximum acceptable noise can be found for the chosen error δ. Note that an equation for r would be different depending on $U(\cdot)$.

To make sure that the noise drawn from a probability distribution does not exceed the maximum acceptable noise, the boundaries for this probability distribution are calculated using its Cumulative Distribution Function (CDF). The CDF of the Boiarkin probability distribution centered at 0 is defined as follows:

$$F_X(x) = \begin{cases} q \cdot b \cdot exp\left(\dfrac{x - \lambda_1}{b}\right), \ x \leq \lambda_1 \\ \\ \dfrac{1}{2} - q \cdot b \cdot \left(exp\left(\dfrac{\lambda_1 - x}{b}\right) - exp\left(\dfrac{\lambda_1}{b}\right)\right), \ \lambda_1 \leq x \leq 0 \\ \\ \dfrac{1}{2} + q \cdot b \cdot \left(exp\left(\dfrac{x - \lambda_2}{b}\right) - exp\left(\dfrac{-\lambda_2}{b}\right)\right), \ 0 \leq x \leq \lambda_2 \\ \\ 1 - q \cdot b \cdot exp\left(\dfrac{\lambda_2 - x}{b}\right), \ \lambda_2 \leq x \end{cases} \tag{6.14}$$

To find a point x at which the CDF of the Boiarkin distribution gives the required level of the CDF's accuracy α ($\alpha = 0.9999$), x has to be expressed from the last equation in (6.14). Taking into account that $p = exp(\lambda_1/b)$ and $p = exp(-\lambda_2/b)$, x is expressed as follows:

$$x = -b \cdot log(p) - b \cdot log(2(1 - \alpha)(2 - p)) \tag{6.15}$$

By replacing x with the maximum acceptable noise, scale b or spread ψ of the Boiarkin probability distribution may be adjusted, so that a random variable r drawn from the probability distribution lies within the acceptable interval $[-x, x]$. For the Laplace probability distribution, x would be expressed as follows:

$$x = -b \cdot log(2(1 - \alpha)) \tag{6.16}$$

Based on the end-user's preference regarding the maximum acceptable error in utility (δ), parameters of the probability distribution used by the randomized mechanism can be adjusted, so that δ (6.12) is within the acceptable range (e.g., $\delta \leq 50\%$). Note that end-users set the maximum acceptable error δ, which affects the parameters of the randomized mechanism, more specifically the probability distribution used by it. This means that the boundaries for the probability distribution are calculated based on the end-user's preference regarding δ, and this configuration of the probability distribution gives the maximum possible level of privacy (range of random variables).

6.3.4 Threat Model and Assumptions

The proposed user-centric data-sharing scheme is designed to protect individual prosumers' energy usage data when it is shared with an external service provider, while allowing users to define acceptable utility boundaries. In this subsection, the threat model and the assumptions under which the scheme is analysed are outlined.

Participating entities and behavioural assumptions

- **End-users.** Each end-user holds their own energy usage data locally and applies the proposed mechanism before sharing any value with an external party. Users are assumed to honestly follow the data-sharing protocol and correctly specify their utility boundary. Malicious users who deliberately inject arbitrary or adversarially crafted values (data poisoning) are not considered.
- **Aggregator / service provider.** The data aggregator and downstream service provider are modelled as honest-but-curious. They correctly follow the proposed protocol and use the noisy data for the intended service (e.g., billing or analytics). However, they may attempt to infer individual users' original data from the noisy reports they receive.
- **Communication channels.** It is assumed that communication between end-users and the aggregator is protected by standard authenticated encryption. Accordingly, an external network adversary can at most obtain the same view as the honest-but-curious aggregator, namely the noisy reports.

Adversarial goals and capabilities.

The main adversarial goal is to reconstruct an individual user's energy usage profile, given access to:

- all noisy values reported by that user over time;
- noisy values from other users; and
- full knowledge of the randomized mechanism (probability distribution, scale, and utility boundaries).

An adversary is assumed to know all other auxiliary information about the dataset except the contribution of the target user. Under this model, the proposed mechanism provides $\epsilon - LDP$ guarantees, bounding the adversary's inference capabilities irrespective of such side information.

6.4 Evaluation and Discussion

This section presents the simulation results to evaluate the proposed data-sharing scheme. In this work, a smart grid environment is studied as a specific use case, whereas the model can be applied to other scenarios.

In the simulated use case, a smart meter deployed on the end-user's side measures the electricity consumption of a household and submits the average electricity consumption per hour to the energy supplier once per day. In the smart grid, end-users energy usage data may be shared with different third parties, including the electricity system operator that controls the balance between supply and demand. In this case study, the data are only shared with the energy supplier and used exclusively for billing. All other scenarios, including the influence on the network operation, are out of the scope of this part of the research. End-users' individual electricity consumption data is classified as personal data under GDPR. For this reason, end-users must have control over who can access their electricity consumption data, how often, and for what purposes, except when the data access is mandated [114]. In addition, in case of an insider attack or a key leakage attack on the energy supplier's side, an adversary would be able to access the raw electricity consumption data of the end-users.

To make sure that the privacy of end-users is not disclosed, as well as to prevent the consequences of cyber-attacks on the energy supplier's side, households may choose to use a LDP scheme to protect their privacy and data, whereas the final energy usage cost may be increased due to the noise injected into the electricity usage data. Thus, end-users decide the extra cost they are willing to pay to increase their privacy level, based on which the noise will be added to the data. The real electricity consumption data for 100 households in London are taken from [93], which contains smart meter readings from 2011 to 2014. The buying price of energy from the utility grid in the UK is taken from [97] and is equal to 14.37 pence/kWh. The objective of this use case is to evaluate (i) whether the relative error in utility remains within the user-defined boundaries, (ii) the additional expected cost caused by noise, (iii) the mutual information between original and perturbed values as a privacy metric, (iv) how the proposed mechanism performs compared to the Laplace mechanism. The simulations are conducted on a machine with Apple M1 CPU @3.2GHz and 8.0GB RAM using Python programming language.

Let $f(x)$ denote the function that calculates the average electricity consumption per hour for a household, which is defined as follows:

$$f(x) = \frac{1}{n}\sum_{i=1}^{n} x_i \tag{6.17}$$

where x is the vector that contains electricity consumption for each time slot (30 minutes) in a day. Let $g(f(x))$ denote the function that takes the average electricity consumption of a household and calculates the energy usage cost for one day, which is defined as follows:

$$g(f(x)) = f(x) \cdot T \cdot \gamma \tag{6.18}$$

where T is the number of time slots in a day ($T = 48$), and γ is the buying price of energy per kWh (γ=14.37 pence/kWh).

By combining (6.17) and (6.18), the maximum acceptable relative error in utility δ (6.12) is calculated as follows:

$$\begin{aligned}\delta &= \frac{(f(x)+r)\cdot T\cdot\gamma - f(x)\cdot T\cdot\gamma}{f(x)\cdot T\cdot\gamma}\cdot 100\% \\ &= \frac{r}{f(x)}\cdot 100\%\end{aligned} \tag{6.19}$$

where r is the random variable (noise) drawn from a probability distribution used by a mechanism $\mathcal{M}(\cdot)$. When the end-user decides on the acceptable relative error in the energy usage cost, the boundaries for the probability distribution can be calculated as follows:

$$r = \frac{\delta\cdot f(x)}{100} \tag{6.20}$$

By combining (6.15) and (6.20), an adjusted scale for the Boiarkin probability distribution is calculated as follows:

$$b = -\frac{\delta\cdot f(x)}{100(log(p)+log(2(1-\alpha)(2-p)))} \tag{6.21}$$

By combining (6.16) and (6.20), the scale for the Laplace distribution is adjusted as follows:

$$b = -\frac{\delta\cdot f(x)}{100(log(2(1-\alpha)))} \tag{6.22}$$

Therefore, based on the end-user's preference regarding δ, the scale of a probability distribution used by a mechanism $\mathcal{M}(\cdot)$ is adjusted to make sure that the noise is within the boundaries, which helps to control the change (error) in utility.

Let σ denote the actual relative error in utility. The dependency between the relative error in utility σ and the maximum acceptable error δ set by the end-user using the Laplace and Boiarkin mechanisms for randomly chosen end-users is shown in Fig. 6.3. To demonstrate how p affects the relative error σ, first we set $p = 0.7$. This p value close to 1 makes the modes of the Boiarkin distribution close to 0, which means that most of the noise will be concentrated around 0. The second value picked for p ($p = 0.2$) is closer to zero, which increases the spread of the modes of the Boiarkin distribution. As a result, most of the

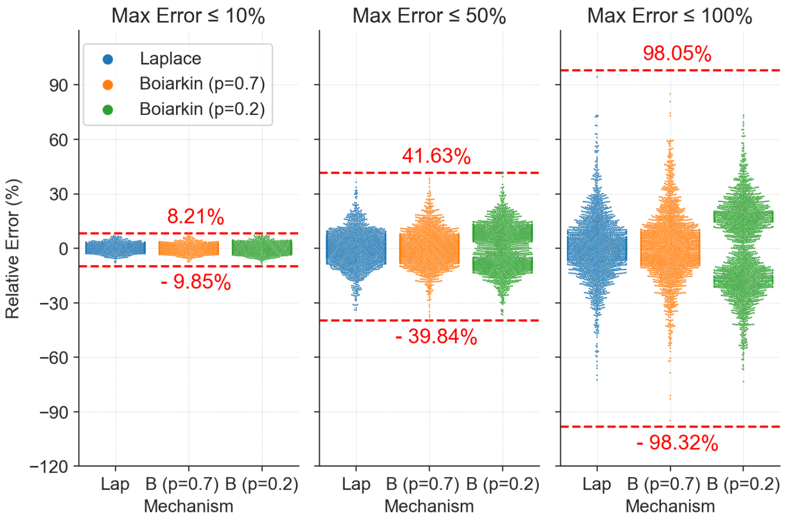

Figure 6.3: Dependency between the relative error in utility σ and the maximum acceptable error δ for randomly chosen end-users for 1 day using the Laplace mechanism and the Boiarkin mechanism with $p = 0.7$ and $p = 0.2$

noise will be concentrated around λ_1 and λ_2. Since end-users send data to the energy supplier once per day, Fig. 6.3 shows the relative error in the energy usage cost (σ) for one day depending on the maximum acceptable error (δ), 10%, 50%, or 100%. It can be seen that σ does not exceed δ. The Laplace mechanism results in a concentration of relative errors around 0% because the noise (a random variable r) drawn from the Laplace distribution has the highest probability around 0. The Boiarkin mechanism results in similar relative error as the Laplace mechanism when p is set close to 1 because the modes λ_1 and λ_2 of the Boiarkin distribution are close to 0. When decreasing p for the Boiarkin mechanism, the concentration of the relative errors around 0% decreases because the noise is concentrated around the modes λ_1 and λ_2, which can be clearly observed when the maximum acceptable error $\delta \leq 100\%$ and $p = 0.2$. Compared to the Laplace mechanism, the Boiarkin mechanism allows to increase the spread of relative errors while keeping it within the boundaries by fine-tuning the amount of generated noise using the parameter p. Note that Fig. 6.3 shows the results for only one day (iteration), so there is no room for noise compensation. When $p = 0.2$, although the noise is frequently not zero, positive and negative values have the same magnitude, which can result in compensation in the long term.

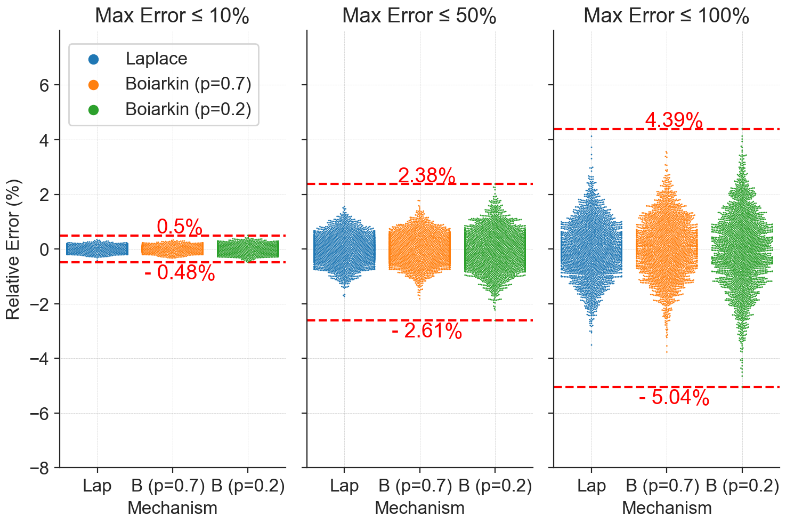

Figure 6.4: Dependency between the relative error in utility σ and the maximum acceptable error δ for randomly chosen end-users over a period of 300 days using the Laplace mechanism and Boiarkin mechanism with $p = 0.7$ and $p = 0.2$

To show the effect of noise compensation, the simulation was conducted for randomly chosen end-users for a period of 300 days using the Laplace mechanism and the Boiarkin mechanism with $p = 0.7$ and $p = 0.2$. The dependency between the relative error in utility σ and the maximum acceptable error δ is shown in Fig. 6.4. It can be clearly observed that the relative error in utility is smaller than the maximum acceptable error because of the noise compensation for both the Laplace and Boiarkin mechanisms. The concentration of the relative errors for the Laplace and Boiarkin mechanisms is around 0%, whereas the spread of relative errors using the Boiarkin mechanism is slightly higher but does not exceed the maximum acceptable error δ. When a randomized mechanism is executed only once (Fig. 6.3), $U(\mathcal{M}(x))$ may result in both higher and lower values compared to the result of $U(x)$, which may result in positive or negative relative error. Over time, with the mechanism being executed multiple times, the relative error becomes smaller because of the noise compensation. Unlike the Laplace mechanism, the Boiarkin mechanism allows to fine-tune the amount of noise injected into the data, which results in a slightly higher relative error when the mechanism is executed only once (see Fig. 6.3), while over time, due to the noise compensation, both the Laplace and Boiarkin mechanisms yield comparable relative errors, all of which fall within the user-defined boundaries (see

Fig. 6.4). Thus, fine-tuning the parameter p of the Boiarkin mechanism allows to inject more noise, while resulting in the same relative error over time compared to the Laplace mechanism.

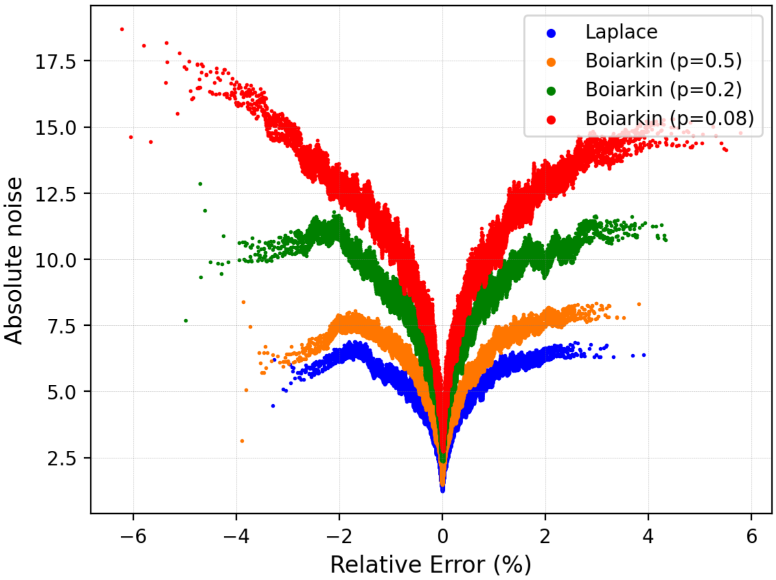

Figure 6.5: Dependency between the relative error σ and the absolute amount of noise added to the data for randomly chosen end-users over a period of 300 days using the Laplace mechanism and the Boiarkin mechanism with different p, where the maximum acceptable error in utility δ is set to 100%.

To show the dependency between the amount of noise injected by a randomized mechanism and relative error σ, the simulation was conducted for randomly chosen end-users over a period of 300 days using the Laplace mechanism and the Boiarkin mechanism with $p = 0.5$, $p = 0.2$, and $p = 0.08$ (Fig. 6.5). The maximum acceptable error in utility δ is set to 100%, which means that the maximum energy usage cost the end-user is willing to pay should not exceed the double of the original cost. It can be observed that the Boiarkin mechanism allows to inject more noise compared to the Laplace mechanism for the same relative error, which may provide better privacy for the end-user with the same utility. When p is set close to 1, the amount of noise is slightly higher compared to the noise injected by the Laplace mechanism. With the decreasing p, the amount of noise increases because the probability of a random variable $Pr[r = 0]$ decreases, and the probabilities around the modes (λ_1 and λ_2) of the Boiarkin distribution increase. Since the Boiarkin probability distribution has two modes, around which the probabilities are concentrated, and taking into account the effect of noise compensation, it is possible to inject more noise and produce the same error in utility as the Laplace mechanism.

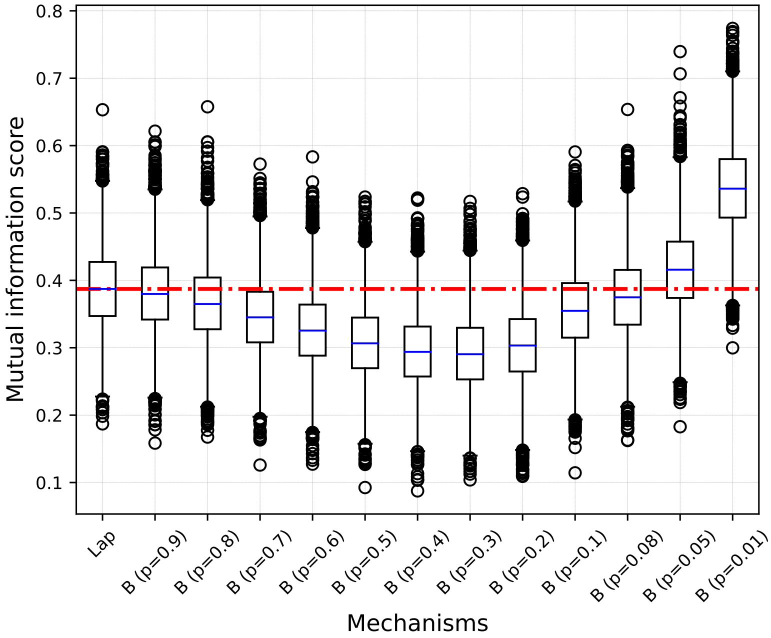

Figure 6.6: Mutual information score between the noisy and original electricity consumption profiles for different end-users over a period of 300 days using the Laplace mechanism and the Boiarkin mechanism with different p, where the maximum acceptable error in utility δ is set to 100%.

The results above showed that the Boiarkin mechanism allows to inject more noise compared to the Laplace, whereas the final error in utility is the same. To check whether more noise means more privacy, the mutual information score between the noisy electricity consumption profiles and the original profiles for randomly chosen end-users over a period of 300 days is evaluated for the Laplace and Boiarkin mechanisms (Fig. 6.6). The maximum acceptable error in utility δ set by end-users is equal to 100%. The mutual information score reflects the extent to which the noisy electricity consumption profile is similar to the original profile of the end-user. The lower the mutual information score, the more independent electricity consumption profiles are. It can be seen that for the Boiarkin mechanism, the mutual information score decreases when p decreases ($p \geq 0.2$), whereas when p becomes too small ($p < 0.2$), the mutual information score increases compared to the Laplace mechanism. Thus, when increasing the probability around the modes (λ_1 and

λ_2) of the Boiarkin distribution, more noise can be injected into the data, which makes noisy and original electricity consumption profiles of the end-user more independent. To keep noisy and original electricity consumption profiles of the end-user more independent and preserve the privacy of end-users, it is suggested to use $p \geq 0.2$. These results also show that the parameter p of the Boiarkin mechanism allows to fine-tune the amount of noise injected into the data, while still keeping the relative error within the boundaries defined by the end-user. For $p = 0.2$, the Boiarkin mechanism provides more privacy compared to the Laplace mechanism since the mutual information score between the noisy and original profiles is lower than for the Laplace mechanism. When $p = 0.9$, the Boiarkin mechanism provides slightly better privacy compared to the Laplace mechanism because the modes λ_1 and λ_2 (Fig. 6.2) of the Boiarkin distribution are close to 0.

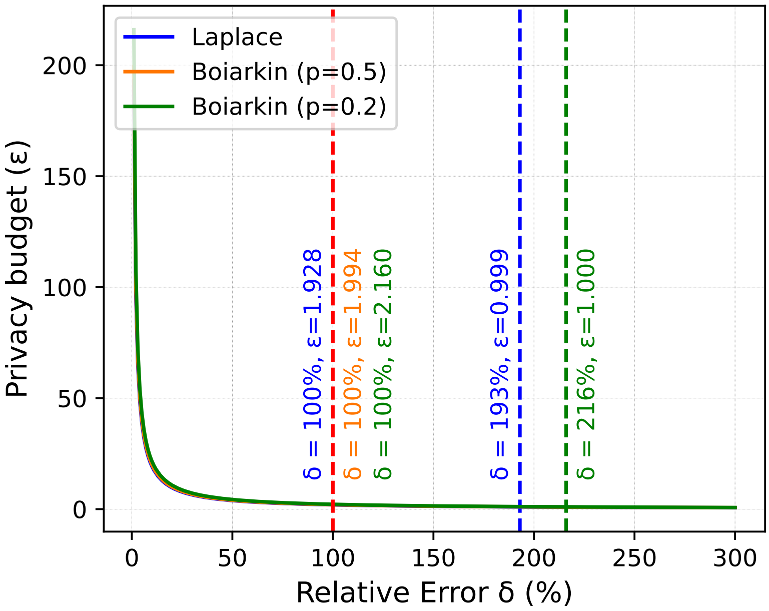

Figure 6.7: Dependency between the privacy budget ϵ and the maximum acceptable relative error in utility δ for a randomly chosen end-user for one day using the Laplace mechanism and the Boiarkin mechanism with $p = 0.5$ and $p = 0.2$.

To understand how the end-user's preference regarding δ affects the privacy budget ϵ used by the Laplace and Boiarkin mechanisms, the simulation was conducted for a randomly chosen end-user for one day using the global sensitivity of a query function $f(\cdot)$ (6.17) (Fig. 6.7). The global sensitivity reflects the maximum distance between two neighboring datasets. The level of electricity consumption for a typical household is around 4kWh, and by changing one entry in an empty dataset (no electricity consumption) to the maximum level of consumption (4kWh), the global sensitivity is calculated as $\Delta f = 4/48 = 0.08333$, where 48 is the number of elements (time slots) in the dataset. After calculating the sensitivity of the function Δf, our next step is to determine the appropriate scale, denoted as b. This scale will be determined based on the maximum acceptable error in utility δ. Subsequently, we calculate ϵ, taking into account the global sensitivity and the adjusted scale ($\epsilon = \Delta f / b$). It can be seen that to keep the relative error around 0%, the privacy budget has to be very high ($\epsilon \approx 200$) because the scale of the probability distribution used by the randomized algorithm should produce random variables within a very limited range, which depends on the sensitivity and may be different for other applications. If the end-user sets the maximum acceptable error $\delta = 100\%$, the privacy budget for the Laplace mechanism $\epsilon = 1.928$, whereas $\epsilon = 1.994$ and $\epsilon = 2.160$ for the Boiarkin mechanism with $p = 0.5$ and $p = 0.2$ respectively. When $\epsilon \leq 1$, the relative error in utility increases, namely the Laplace mechanism with $\epsilon = 0.999$ results in 193% error, whereas the Boiarkin mechanism with $p = 0.2$ and $\epsilon = 1$ results in 216% error.

Although the privacy budget ϵ is slightly higher for the Boiarkin mechanism, the amount of noise injected by the Boiarkin mechanism is higher, as well as the mutual information score between noisy and original data is lower compared to the Laplace mechanism, which means that the Boiarkin mechanism may provide better privacy for the same ϵ. In the scheme proposed in this work, when the end-user defines the boundaries for utility (δ), the privacy budget ϵ is calculated automatically based on the sensitivity and the scale of the probability distribution. Since the scale of the probability distribution is adjusted based on δ, which allows to control the amount of noise, there is only one ϵ that can be calculated for the given δ and Δf. This provides the maximum possible level of privacy for the defined δ. On the other hand, if the end user had to choose the privacy budget ϵ instead of the maximum tolerated error, the error in the utility they would get in the future would be unclear. By utilizing the proposed data-sharing scheme, end-users have control over the utility, as well as better understanding how the level of privacy affects utility.

6.5 Summary

The increasing demand for data-driven services has led to growing concerns about preserving the privacy of end-users' sensitive data, particularly in contexts such as smart grids, where energy usage data is considered personal data under the GDPR. While DP and LDP have gained popularity for their ability to obscure sensitive information, most existing approaches commonly rely on fixed privacy budgets and offer limited flexibility to end-users. Furthermore, these schemes often lack mechanisms that allow users to control the trade-off between utility and privacy in a practical, intuitive manner.

This chapter introduced a novel, user-centric privacy-preserving data-sharing scheme that addresses identified limitations through three core contributions. First, a new bimodal probability distribution—termed the Boiarkin distribution—was proposed. This distribution enables fine-tuned control over noise injection by concentrating probability mass around two symmetric modes, thereby facilitating a noise compensation effect over time. Secondly, the chapter presents a novel privacy mechanism utilizing the proposed probability distribution that satisfies ϵ-local differential privacy. The mechanism allows greater control over noise characteristics compared to traditional Laplace-based methods. Finally, a flexible data-sharing framework was designed that lets end-users specify their acceptable level of utility error, from which the system automatically determines the highest feasible privacy budget.

Existing DP- and LDP-based data-sharing mechanisms predominantly rely on the Laplace mechanism and expose the privacy budget ϵ as the primary tuning parameter [41, 42]. In contrast, the framework proposed in this chapter offers three key advantages. First, it is user-centric: end-users specify an acceptable relative error in utility (δ), and the mechanism derives the corresponding privacy budget ϵ automatically, making the configuration more intuitive for end-users. Second, by leveraging the proposed bimodal probability distribution, the mechanism injects more noise than the Laplace mechanism while maintaining the same long-term relative error in utility, due to the noise-compensation effect (Sections

6.3.1–6.4). Third, the evaluation shows that, for a fixed utility bound δ, the proposed mechanism yields lower mutual information between original and perturbed data than the Laplace mechanism, indicating stronger privacy for the same utility level. These properties position the proposed mechanism as a competitive alternative to Laplace-based schemes.

Simulation results demonstrated that the proposed mechanism could inject more noise while preserving utility over time through compensation effects. Additionally, the system showed better privacy protection—as measured by reduced mutual information score between original and noisy data—compared to Laplace mechanisms. These results confirm that the proposed approach can significantly enhance both privacy and usability in practical deployments.

In conclusion, the research in this chapter has presented a new user-centric data-sharing scheme based on a novel probability distribution and randomized mechanism, offering end-users greater control over utility while achieving stronger privacy guarantees compared to traditional approaches.

Chapter 7

Privacy-preserving Energy Demand Forecasting Scheme for a Smart Grid

7.1 Overview

The increased use of renewable energy and smart grid technologies has raised the importance of accurately forecasting energy usage demand. Accurate forecasts help utilities and microgrid operators manage resources more effectively, implement demand-response programs, and maintain a stable energy supply. Typically, energy forecasts are performed by a central authority, such as a grid operator, requiring consumers or microgrids to share detailed energy consumption data. However, this centralized approach creates privacy issues, as energy usage patterns can reveal sensitive details about users' daily habits and lifestyles.

Privacy regulations, such as GDPR, require solutions that balance the need for accurate energy demand forecasting with strong privacy protection. FL has become a promising solution that enables collaborative training of a machine-learning model and does not require participants to share the actual data with other parties [115]. In this approach, each participant trains a local model using his private dataset, and only updates to these models (such as model parameters or gradients) are shared with a central system.

Despite its advantages, conventional FL faces specific challenges in smart grid applications. Firstly, there is still a risk that sensitive information could be leaked from the shared model updates during aggregation due to data reconstruction and membership attacks [116–118]. Secondly, differences between the datasets held by individual microgrids (known as statistical heterogeneity) can negatively affect the accuracy and speed of model training.

This chapter addresses the aforementioned privacy concerns by proposing a Privacy-preserving Energy Usage Forecasting Scheme for smart grids. The scheme integrates privacy protection methods, such as DP, into the federated learning process. This approach ensures strong protection of users' sensitive energy data, reduces the risk of information leakage during training, and maintains accurate forecasting results.

The key innovation of this work lies in formulating a time-dependent target accuracy function that guides the evolution of the model's performance expectations and drives the configuration of differential privacy parameters throughout training. Choosing an appropriate privacy budget is a long-standing challenge in the differential privacy domain, particularly because the relationship between privacy and utility (e.g., model accuracy) is often unclear in advance. This challenge becomes even more complex in distributed settings such as federated learning, where different clients may follow distinct privacy preferences. In this context, this work introduces a novel mechanism in which the central server defines a dynamic privacy budget based on a function that reflects the desired accuracy and the time horizon for achieving it. Rather than allowing clients to select their privacy budgets independently, the central server invites clients to adhere to the proposed privacy budget configuration. This mechanism ensures alignment between local behavior and global goals. The privacy budget is not static: it adapts throughout training, allowing for more relaxed privacy (and hence better utility) when high accuracy must be achieved early, and tightening privacy guarantees as the training progresses. Additionally, the server sets boundaries and rules for how the privacy budget evolves, reducing the decision-making burden on clients and promoting coordinated and privacy-aware learning.

The main design goal of this scheme is dynamic adaptability, where the privacy configuration adjusts over time to balance early-stage learning requirements with later-stage privacy preservation. The evaluation framework assesses this goal by training a central machine learning model using federated learning on distributed energy datasets. Accuracy is measured using the Root Mean Square Error (RMSE) metric, comparing the proposed dynamic approach against fixed-privacy baselines and non-private training.

7.2 Proposed Model

In this section, a novel privacy-preserving energy usage forecasting scheme for smart grids is proposed.

7.2.1 System Model

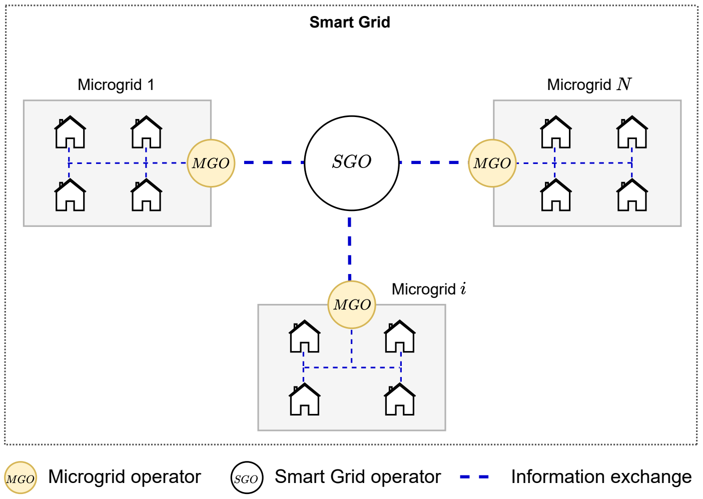

Figure 7.1: System model of a smart grid of microgrids

Fig. 7.1 shows the structure of a smart grid consisting of N microgrids. Each microgrid represents a group of consumers who can produce and consume energy. In addition, each microgrid includes a microgrid operator, who is responsible for aggregating energy usage data within the microgrid. The collected data can then be used to predict future energy demand within the microgrid. The difference between the datasets collected by different MGO lies in the levels of energy usage that reflect prosumers' behavior within a particular microgrid. To preserve prosumers' privacy, MGOs want to keep these datasets private while using them to collaboratively train an energy demand forecasting model to better predict future energy demand. Smart Grid Operator (SGO) operates as an aggregator or a central authority to enable collaborative training among microgrids. Each MGO trains a local machine learning model using its private dataset and only reports model updates (weights) to the central server (SGO). SGO is responsible for combining all received updates from MGOs and calculating new model weights, which are then distributed to MGOs. Thus, MGOs can collaboratively train one machine learning model to better predict future energy demand while keeping their datasets private.

The main aim of this work is to design a privacy-preserving energy demand forecasting scheme for a smart grid consisting of several microgrids.

In this work, a smart grid consists of N microgrids. Let $\mathcal{N} = \{1, 2, 3, ..., N\}$ denote the set of microgrids in the smart grid, where n is the microgrid index and $n \in \mathcal{N}$, whereas, the total number of microgrids is given by $N \triangleq |\mathcal{N}|$. Each microgrid represents a group of prosumers, whereas the number of prosumers in a microgrid may be different. Let $\mathcal{M}_n = \{1, 2, 3, ..., M_n\}$ denote the set of prosumers in the microgrid n, where m_n is the prosumer index and $m_n \in \mathcal{M}_n$, whereas, the total number of prosumers in the microgrid n is given by $M_n \triangleq |\mathcal{M}_n|$. Next, to simplify the mathematical equations, the index m will be used instead of m_n, meaning the prosumer m in a microgrid n.

MGO aggregates energy usage data of all prosumers in a microgrid at the end of each time slot, where each time slot is of one hour. Let $\mathcal{H} = \{1, 2, 3, ..., H\}$ denote the set of all time slots, where h is the time slot index and $h \in \mathcal{H}$, whereas, the total number of time slots is given by $H \triangleq |\mathcal{H}| = 24$.

Let QD_{m_h} denote the actual level of energy consumption for a prosumer m in a microgrid n in a time slot h and QS_{m_h} denote the actual level of energy production (supply) for a prosumer m in a microgrid n in a time slot h.

The first priority for prosumers is self-consumption; thus, each prosumer consumes energy it generates (PV) to cover its demand. If there is not enough energy (PV) produced on the prosumer's side, the rest is bought from the microgrid, which contributes to the total energy demand of a microgrid. On the other hand, if a prosumer can cover its demand by consuming its own energy (PV) and has excess energy to share, the surplus is sold to the microgrid, which can be consumed (bought) by other prosumers within that microgrid.

Let NP_{m_h} denote the net power for a prosumer m in a microgrid n in a time slot h, and it is defined as follows:

$$NP_{m_h} = QS_{m_h} - QD_{m_h} \tag{7.1}$$

At the end of each time slot, each prosumer reports its net power to the MGO, which then aggregates all received measurements and calculates total amount of energy consumed in a reporting time slot. Let NP_h^T denote the total net power (total demand) of a microgrid n in a time slot h, which is calculated as follows:

$$NP_h^T = \sum_{m=1}^{M_n} NP_{m_h} \tag{7.2}$$

7.2.2 Privacy-preserving energy demand forecasting model

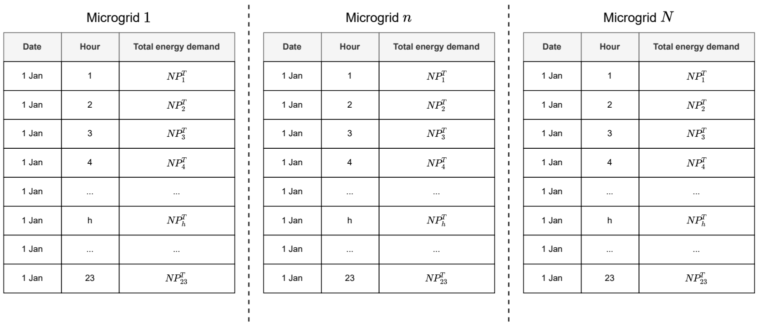

Microgrid 1

Date	Hour	Total energy demand
1 Jan	1	NP_1^T
1 Jan	2	NP_2^T
1 Jan	3	NP_3^T
1 Jan	4	NP_4^T
1 Jan	...	...
1 Jan	h	NP_h^T
1 Jan	...	...
1 Jan	23	NP_{23}^T

Microgrid n

Date	Hour	Total energy demand
1 Jan	1	NP_1^T
1 Jan	2	NP_2^T
1 Jan	3	NP_3^T
1 Jan	4	NP_4^T
1 Jan	...	...
1 Jan	h	NP_h^T
1 Jan	...	...
1 Jan	23	NP_{23}^T

Microgrid N

Date	Hour	Total energy demand
1 Jan	1	NP_1^T
1 Jan	2	NP_2^T
1 Jan	3	NP_3^T
1 Jan	4	NP_4^T
1 Jan	...	...
1 Jan	h	NP_h^T
1 Jan	...	...
1 Jan	23	NP_{23}^T

Figure 7.2: Private datasets in different microgrids

At the end of each time slot, prosumers within a microgrid report their net power energy data to the MGO, which aggregates received measurements and adds aggregated data to a local dataset. Fig. 7.2 shows different datasets stored privately within each microgrid.

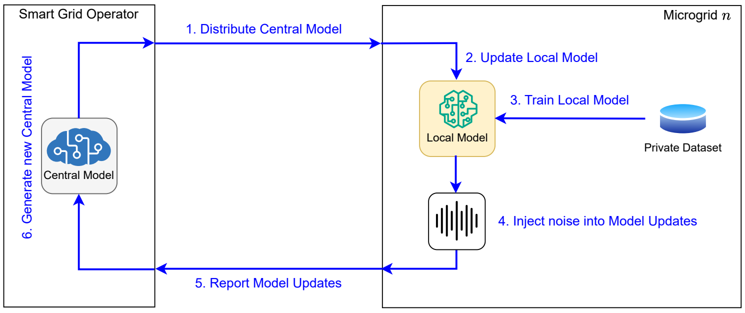

Figure 7.3: Lifecycle of collaborative model training

Fig. 7.3 shows the lifecycle of a collaborative machine learning model training by microgrids. At the first iteration, SGO initializes an initial central model and distributes the model configuration to all microgrid operators. Each MGO initializes a local model based on the received model configuration. Then, the training process repeats until a certain level of accuracy or a maximum number of training rounds is reached.

Each training round includes the following steps. SGO distributes the latest central model configuration to all microgrid operators. Each MGO updates its local model parameters according to the new central model configuration. Then, each MGO uses its private dataset (Fig. 7.2) to train the local model. Since reporting raw model updates (model weights) to SGO may cause data reconstruction and membership attacks [116–118], MGO injects noise into the model updates before sharing them with SGO. After SGO receives model updates from all microgrid operators, it applies a specific algorithm to generate new weights for a new central model. In this work, the FedAvg algorithm [119] is used to iteratively average the local model updates. At this step, if the number of training rounds has not been reached, SGO initiates a new training round by distributing a new central model to all microgrid operators.

Let $\mathcal{T} = \{1, 2, 3, ..., T^{max}\}$ denote the set of all training rounds, where t is the training round number and $t \in \mathcal{T}$, whereas, the total number of training rounds is defined by the central server and $T \triangleq |\mathcal{T}| = T^{max}$. It should be noted that the central server defines T^{max}, thus aiming to train the central model for only T^{max} rounds. There are a number of reasons why the central server defines the maximum number of rounds, including the prevention of model overfitting and eliminating unnecessary use of clients' computational resources.

Initially, the central server initializes the central model with randomly set or zero weights. Let w_0 ($w_{t=0}$) denote the initial weights of the central model. At the beginning of a new training round t, the central server distributes the model configuration (weights) to all clients participating in the training process. Each client updates its local model with a new configuration received from the central server and uses its private dataset to train a new local model. Let w^m_{t+1} denote the weights of the client's local model after training the local model in a round t, which can be expressed as follows:

$$w^m_{t+1} = TrainClientModel(w_t, PD_m) \tag{7.3}$$

where PD_m is the private dataset containing energy usage data within the microgrid m, which is stored privately on the MGO_m side.

After a client m updates its local model, it reports its model configuration (weights) w^m_{t+1} to the central server. When SGO (central server) receives new model parameters from all clients, it aggregates (averages) all clients' model updates and generates a new configuration for the central model.

Let w_{t+1} denote the new weights of the central model after the training round t, which is calculated as follows:

$$w_{t+1} = \sum_{m=1}^{M} \frac{pd_m}{pd} w^m_{t+1} \tag{7.4}$$

where d_m is the number of records in the private dataset PD_m on the client m, and d is the total number of records in all private datasets on all clients.

The ratio $\frac{d_m}{d}$ determines the contribution of each client m to the update of the central model. Thus, clients with more records in their private datasets should have more influence on the update of the global model. However, in order to preserve the privacy of the clients, the proposed model eliminates the need to share any details regarding the private datasets with the central server. Hence, all clients contribute equally to the update of the central model, and 7.4 may be rewritten as follows:

$$w_{t+1} = \sum_{m=1}^{M} \frac{1}{M} w_{t+1}^{m} \tag{7.5}$$

where M is the total number of clients (microgrids) participating in the training.

To further enhance the privacy of clients, the proposed model integrates the concept of Local Differential Privacy, thus enabling clients (microgrid operators) to determine the required level of privacy. To preserve its privacy, each client injects controllable noise into the model updates before sharing them with the central server (Fig. 7.3). Thus, instead of receiving raw model parameters from clients, the central server receives noisy versions of clients' model weights. This approach is considered more lightweight compared to encryption-based approaches. However, to prevent eavesdropping attacks during the data transmission between clients and the central server, noisy model updates may be encrypted.

Let $w_{t+1}^{m^*}$ denote the model updates that the client shares with the central server, which is defined as follows:

$$w_{t+1}^{m^*} = w_{t+1}^{m} + Z(\cdot) \tag{7.6}$$

where w_{t+1}^{m} is the weights of the client m local model after a training round t, and $Z(\cdot)$ is the noise injection mechanism.

For example, clients may utilize the Laplace mechanism (6.4) to inject noise into the model updates, thus (7.6) may be rewritten as follows:

$$w_{t+1}^{m^*} = w_{t+1}^{m} + Lap\left(\frac{\Delta f}{\epsilon_m}\right) \tag{7.7}$$

where ϵ_m is the level of privacy defined by the client m. It has to be noted that a high level of privacy will lead to more noise injected into the data, which will degrade the accuracy of the final model.

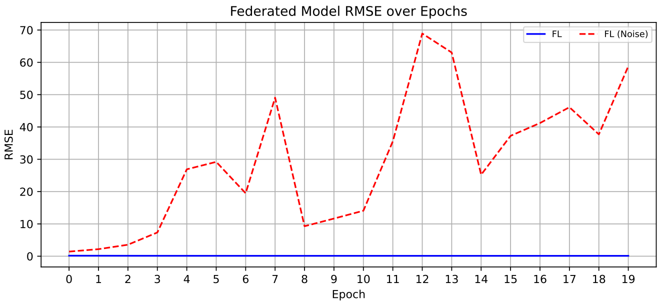

Figure 7.4: Federated Learning of the central model by two clients with a fixed privacy setting ($\epsilon_m = 0.5$)

In this work, Root Mean Square Error (RMSE) is used to evaluate the performance of the federated learning system. RMSE has been chosen because it emphasizes larger errors more strongly, which is critical when even small prediction inaccuracies can accumulate across distributed nodes. Unlike Mean Absolute Error (MAE), RMSE captures the variance of errors, making it more sensitive to outliers that could signal poor model generalization. RMSE is determined as follows:

$$RMSE = \sqrt{\sum_{i=1}^{n} \frac{(y_i^* - y_i)^2}{n}} \tag{7.8}$$

where y_i^* is predicted value, y_i is observed value, and n is the number of observations.

Fig. 7.4 shows the RMSE of the central model trained by two clients over 20 epochs. After each training round, the central model is evaluated using random parts of clients' test datasets to reflect the most realistic scenario. The blue line shows the RMSE of the central model when clients do not inject noise, whereas the red dashed line shows the RMSE when each client injects controllable noise into the model parameters using the Laplace mechanism ($\epsilon_m = 0.5$) before sharing them with the central server for averaging. It can be clearly observed that the noise injected on the clients' side completely degrades the accuracy of the central model while providing a high level of privacy to the clients.

In order to make sure that the accuracy of the central model does not degrade with the increasing number of training rounds, the system should motivate clients not to inject a lot of noise into their local models' updates that are reported to the central server (SGO) for averaging. The introduction of the motivation component into the training process means that parties participating in the training of the central model are either aiming at achieving the same outcome (high accuracy) beneficial for all parties or receiving some reward for their contribution.

The exact reward mechanism depends on the business case. For example, a central bank may have an initiative to develop a fraud prevention model that analyses financial transactions and marks malicious ones as fraud with high accuracy. In this case, the central authority (a central bank) will reward other banks for collaboratively training the central model without disclosing their private datasets. On the other hand, in the energy domain, microgrid operators in a bigger network (smart grid) are interested in better predicting future energy demand, which can be achieved by training a central model using datasets of other microgrids. In this case, smart grid operators do not receive incentives, while a highly accurate demand forecasting model directly affects the efficiency of managing an electricity network.

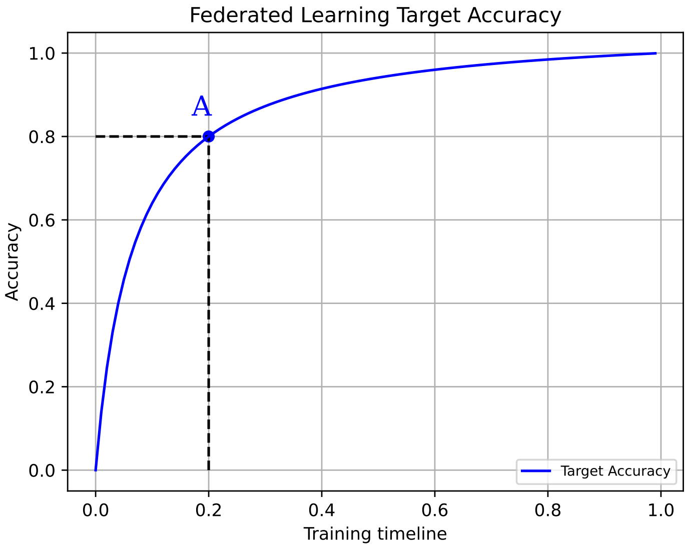

Figure 7.5: Federated Learning Target Accuracy

This work proposes a mechanism for determining rewards in a federated learning system, where participants are rewarded for contributing to the targeted accuracy that has to be achieved within a predefined number of training rounds. Fig. 7.5 shows the dependency of the targeted model accuracy and training timeline. In this particular example, the targeted accuracy of the central machine learning model is set to 80% after completing 20% of training. In other words, the consortium (participants) has a target of achieving an accuracy of 80% after completing 20% of the training.

Let ψ denote the function that defines the dependency between the targeted accuracy and the training time (number of training rounds), which is defined as follows:

$$\psi = 1 - \frac{T - x}{\omega x + T} \tag{7.9}$$

where T is the number of training rounds, and ω is a parameter that reflects the concavity of a curve or the growth rate of accuracy. By expressing ω from 7.9, it is defined as follows:

$$\omega = \frac{\xi - t}{t(1 - \xi)} \tag{7.10}$$

where ξ is the desired accuracy of the central model that has to be reached at the training round t.

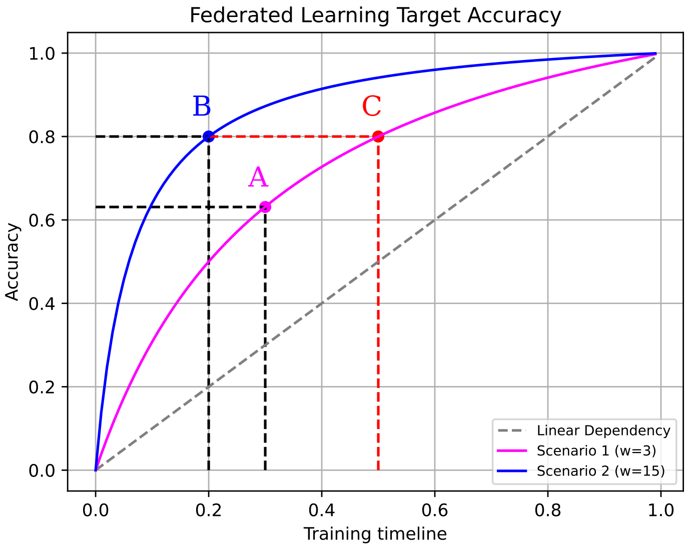

Figure 7.6: Different scenarios of federated learning

Fig. 7.6 shows different scenarios of federated learning depending on the desired level of accuracy ξ that has to be reached at the desired training round t. In scenario 1, the target for the accuracy of the central model is set to 63% by completing 30% of training. In scenario 2, the target for the accuracy is set to 80% by completing 20% of training. It can be observed that in scenario 1, the growth rate of accuracy is lower compared to scenario 2. For example, the accuracy of 80% is reached after completing 20% of training in scenario 2, while the same accuracy in scenario 1 will be reached after completing 50% of training. In addition, the chart shows the linear dependency between the targeted accuracy and the training round (dashed gray line). In this work, it is assumed that the real dependency between the accuracy and the training round can only be represented by a curve defined in 7.9 that lies above the linear dependency line. This assumption is based on the way the learning rate in machine learning algorithms decreases with the increase in training rounds.

The axes' values on Fig. 7.5 and Fig. 7.6 are set to an interval from 0 to 1, representing 100% accuracy for the y axis and 100% of training time for x axis. 100% of training time means the maximum number of training rounds, which can vary. These figures can be represented for a particular number of training rounds by scaling the axes, whereas the key aspect here is the growth rate of accuracy, which will be the same and does not depend on scale.

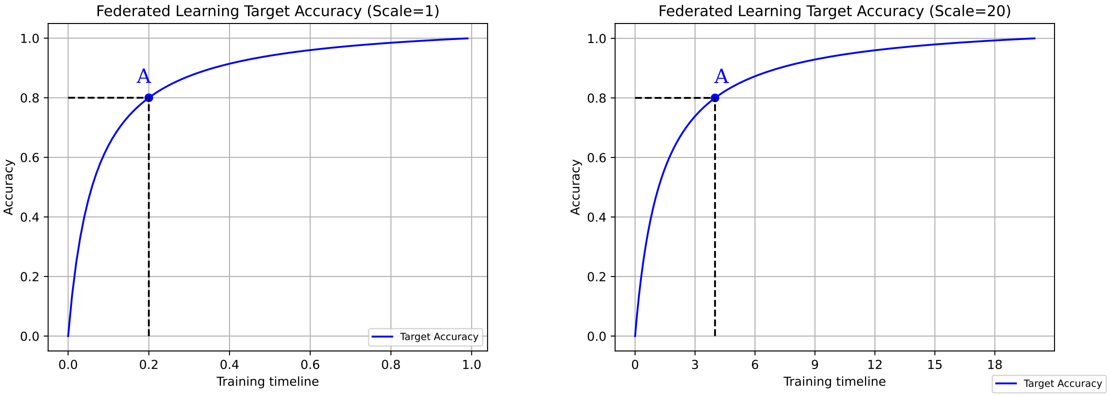

(a) Scale = 1

(b) Scale = 20

Figure 7.7: Representation of Federated Learning Target Accuracy using different scales

Fig. 7.7 shows the Federated Learning target accuracy using different scales, where the target for the accuracy is set to 80% after completing 20% of training. Fig. 7.7 (a) shows the dependency between the targeted accuracy and the training timeline using a scale equal to 1. It can be observed that the curve passes through the point A (0.2; 0.8), which means that the accuracy of 80% is reached at 20% of training time. Fig. 7.7 (b) shows the dependency between the targeted accuracy and the training timeline using a scale equal to the maximum number of training rounds ($T = 20$). It can be observed that the curve passes through the point A (4; 0.8), which means that the accuracy of 80% is reached at the 4th training round. Since $x = 4$ on the right figure is exactly equal to 20% of 20 training rounds, Fig. 7.7 (a) and Fig. 7.7 (b) represent the same dependency pattern between the targeted accuracy and the training time.

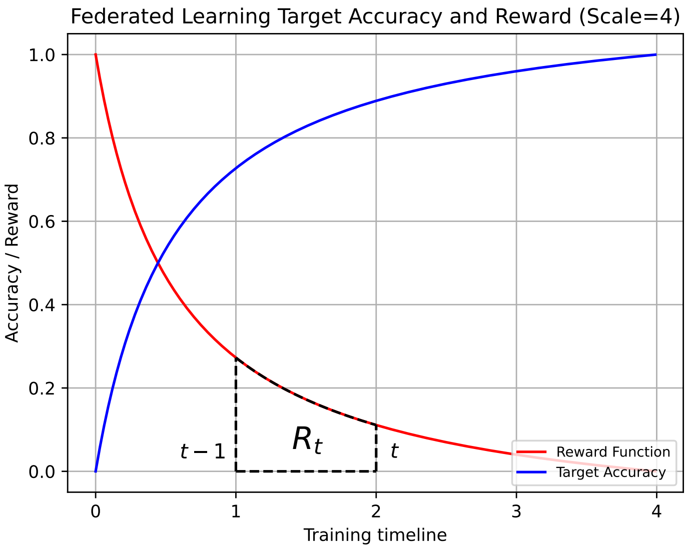

Figure 7.8: Federated Learning Target Accuracy and Reward Function

In business scenarios when the learning process implies rewarding participants for their contributions, the reward amount should correlate with participants' contributions. Fig. 7.8 shows how overall reward correlates with the training time. In the proposed approach, the reward function is inversely proportional to the growth rate of targeted accuracy. The rationale behind this is that at the beginning of the training, the growth rate of accuracy is higher compared to later rounds, which leads to a higher reward distributed among contributing participants. At the end of the training process, the increase in accuracy becomes small, which is reflected in the decreased reward.

Let φ denote the reward function. The growth rate of the reward function is inversely proportional to the growth rate of targeted accuracy. Thus, the reward function is defined as follows:

$$\varphi = \frac{T - x}{\omega x + T} \tag{7.11}$$

where T is the number of training rounds, and ω is a parameter that reflects the concavity of the growth rate of the targeted accuracy in the original function (7.9). Let R^V denote the value of a total reward allocated for a whole training process. Let R^T denote the overall reward represented as an integral of the reward function φ for a whole training process, which is defined as follows:

$$R^T = \int_0^T \varphi \, dx = \int_0^T \frac{T - x}{\omega x + T} \, dx \tag{7.12}$$

Let R_t denote the value of the reward that is distributed among participants in the training round t and represented as an integral of the reward function φ on the interval $[t - 1; t]$. Thus, R_t is defined as follows:

$$R_t = \int_{t-1}^t \varphi \, dx = \int_{t-1}^t \frac{T - x}{\omega x + T} \, dx \tag{7.13}$$

where $t - 1$ is the number of the previous training round, and t is the current training round.

Since the area under the reward function curve does not equal to the value of the total reward (R^V), the value of the integral R_t cannot be directly used for the reward allocation. Instead, the actual value of the reward in a training round t depends on the ratio of the integral of the reward function on the interval $[t - 1; t]$ and the integral of the reward function on the interval $[0; T]$. This ratio represents the portion of the total reward (R^V) that is allocated for the training round t.

First of all, the integral of the reward function φ (7.11) can be evaluated as follows:

$$F(x) = \int \varphi \, dx = \int \frac{T - x}{\omega x + T} \, dx$$

$$= -\frac{(\omega x + T) + (-\omega T - T) ln(\omega x + T)}{\omega^2} \tag{7.14}$$

The overall reward (R^T) is the area under the reward function curve φ on the interval $[0; T]$ and can be evaluated using the result of (7.14) as follows:

$$R^T = \int_0^T \varphi \, dx = \int_0^T \frac{T - x}{\omega x + T} \, dx = F(x)\Big|_0^T$$

$$= -\frac{(\omega T + T) + (-\omega T - T) ln(\omega T + T)}{\omega^2}$$

$$+ \frac{T + (-\omega T - T) ln(T)}{\omega^2}$$

$$= \frac{(-\omega T - T)\Big(ln(T) - ln(\omega T + T)\Big) - \omega T}{\omega^2} \tag{7.15}$$

The reward that is allocated for the training round t is the portion of the overall reward R^T. R_t is the area under the reward function curve φ on the interval $[t - 1; t]$ (Fig. 7.8) and can be evaluated using the result of (7.14) as follows:

$$R_t = \int_{t-1}^t \varphi \, dx = \int_{t-1}^t \frac{T - x}{\omega x + T} \, dx = F(x)\Big|_{t-1}^t$$

$$= -\frac{(\omega t + T) + (-\omega T - T) ln(\omega t + T)}{\omega^2}$$

$$+\frac{(\omega(t-1)+T)+(-\omega T-T)ln(\omega(t-1)+T)}{\omega^2}$$

$$=\frac{(-\omega T-T)\Big(ln(\omega(t-1)+T)-ln(\omega t+T)\Big)-\omega}{\omega^2} \tag{7.16}$$

Since the area under the reward function curve φ does not equal to the value of the reward (R^V), the actual value of the reward for the training round t can be determined based on the ratio of the reward for the training round t and the overall reward R^T.

Let ζ_t denote the ratio of the reward for the training round t and the overall reward R^T, which can be expressed using (7.15) and (7.16) and is defined as follows:

$$\zeta_t=\frac{R_t}{R^T}=\frac{\dfrac{(-\omega T-T)\Big(ln(\omega(t-1)+T)-ln(\omega t+T)\Big)-\omega}{\omega^2}}{\dfrac{(-\omega T-T)\Big(ln(T)-ln(\omega T+T)\Big)-\omega T}{\omega^2}}$$

$$=\frac{(-\omega T-T)\Big(ln(\omega(t-1)+T)-ln(\omega t+T)\Big)-\omega}{(-\omega T-T)\Big(ln(T)-ln(\omega T+T)\Big)-\omega T} \tag{7.17}$$

The ratio ζ_t represents the portion of the overall reward R^T that is allocated for the training round t and is distributed among contributing participants. In other words, ζ_t quantifies what part of the overall reward should be spent in the training round t, which is based on the targeted accuracy in this training round.

Let Δ_t denote the actual portion of the value of the overall reward R^V that is allocated for the training round t and is determined based on the ratio ζ_t and the value of the overall reward R^V as follows:

$$\Delta_t=R^V\frac{R_t}{R^T}=R^V\zeta_t$$

$$= R^V \frac{(-\omega T - T)\Big(ln(\omega(t-1)+T) - ln(\omega t + T)\Big) - \omega}{(-\omega T - T)\Big(ln(T) - ln(\omega T + T)\Big) - \omega T} \tag{7.18}$$

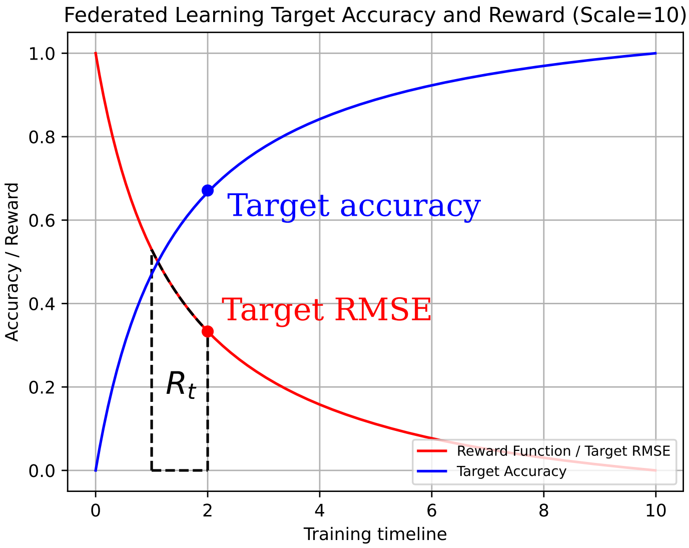

Figure 7.9: Federated Learning Target Accuracy, Target RMSE, and Reward Function

Fig. 7.9 shows the dependency between the target accuracy and reward in the training round t. The number of training rounds is set equal to 10 ($T = 10$), whereas the parameter ω that controls the concavity of the targeted accuracy (ψ) and reward function (φ) curves is set equal to 7 ($\omega = 7$). The reward will be based on the portion of the reward ($R_{t=2}$) and calculated according to (7.18). In this particular example, the value of the reward that will be distributed among participants in the training round $t = 2$ will be calculated as follows:

$$\Delta_2 = R^V \frac{R_2}{R^{10}} = R^V \zeta_2$$

$$= R^V \frac{(-7*10-10)\Big(ln(7(2-1)+10) - ln(7*2+10)\Big) - 7}{(-7*10-10)\Big(ln(10) - ln(7*10+10)\Big) - 7*10}$$

$$= R^V \frac{(-80)\Big(ln(17) - ln(24)\Big) - 7}{(-80)\Big(ln(10) - ln(80)\Big) - 70}$$

$$= R^V * \frac{20.587}{96.355} = 0.213 * R^V$$

Depending on the type of machine learning problem, the targeted accuracy function ψ can be changed to the targeted RMSE function, which is calculated in the same way as the reward function φ. For example, the targeted accuracy function can be used for classification problems, whereas the targeted RMSE function can be used for predicting continuous numeric values, such as energy consumption, temperature, and so on.

To receive a reward after a training round t, clients must contribute to the overall objective of the federated learning task, which is to achieve a desired level of accuracy of the central model. At the same time, the privacy of clients should be guaranteed by the scheme, enabling collaborative training of a machine learning model without disclosing sensitive information. In the proposed scheme, each client injects controllable noise into the machine learning model updates (model weights) using an LDP mechanism, before sharing the model updates with the central server.

The privacy budget ϵ can vary from 0 to 1, where $\epsilon = 0$ corresponds to the highest level of privacy, and $\epsilon = 1$ to the lowest level of privacy. The value of $\epsilon = 1$ cannot be used because it would open doors for adversaries to conduct data reconstruction attacks, as no noise would be injected into the model updates. Similarly, the value of $\epsilon = 0$ cannot be used either, because an LDP mechanism will inject too much noise into the data, thereby degrading the accuracy of the central model.

To address this issue, the proposed scheme involves the central server determining the interval for the privacy budget, within which clients can select their custom preferred privacy interval. Let ϵ^{min} denote the lower bound for the privacy budget interval, and ϵ^{max} denote the upper bound for the privacy budget interval determined by the central server.

Let ϵ_m^{min} denote the lower bound for the privacy budget interval, and ϵ_m^{max} denote the upper bound for the privacy budget interval determined by the client m, where

$$\epsilon_m^{max} > \epsilon_m^{min} \quad and \quad \epsilon_m^{min} \geq \epsilon^{min} \quad and \quad \epsilon_m^{max} \leq \epsilon^{max}$$

Let ϵ_m^t denote the privacy budget for a client m in a training round t, which follows the reward function curve φ pattern and is defined as follows:

$$\epsilon_m^t = \epsilon_m^{max} * \varphi + (1 - \varphi) * \epsilon_m^{min} \tag{7.19}$$

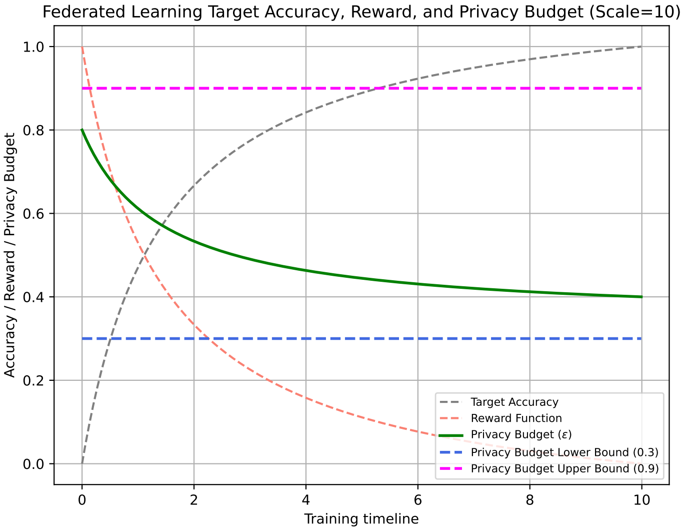

Figure 7.10: Federated Learning Target Accuracy, Reward, and Privacy Budget

Fig. 7.10 shows the lower and upper bounds for the privacy budget, which are determined by the central server, specifically $\epsilon^{min} = 0.3$ and $\epsilon^{max} = 0.9$. In this example, a client m selected $\epsilon_m^{min} = 0.4$ and $\epsilon_m^{max} = 0.8$, where the green curve reflects how the client's privacy budget changes over time, following the pattern of the reward curve. Thus, at the beginning of the training, the reward will be higher due to the lower noise injected into the model updates, whereas at the end of the training, the privacy budget is smaller, as is the reward.

7.3 Threat Model

The proposed federated learning framework for energy demand forecasting involves a Smart Grid Operator (SGO) acting as a central server and multiple Microgrid Operators (MGOs) acting as clients. In this subsection, the assumed behaviour of these entities, the adversarial capabilities, and the specific misbehaviour scenarios considered in our evaluation are described.

Participating entities and behavioural assumptions

- **Smart Grid Operator (SGO).** The SGO is assumed to be honest-but-curious. It faithfully executes the federated learning protocol, including distributing the global model, computing the target RMSE and reward functions, and defining the dynamic privacy budget interval. However, it may attempt to infer information about individual MGOs' underlying energy consumption patterns from the model updates it receives.
- **Microgrid Operators (MGOs / clients).** MGOs are semi-honest: they locally train on their own historical data and apply the prescribed LDP mechanism before sending model updates to the SGO. However, they may be curious about other MGOs' data.
- **Communication channels.** It is assumed that MGOs and the SGO communicate over authenticated and encrypted channels, and that each participating MGO has a unique, authenticated identity registered with the SGO. Under this assumption, Sybil attacks are not considered.

Adversarial goals and capabilities.

The main security objective is to preserve the privacy of each MGO's underlying energy consumption data while training an accurate global forecasting model. Adversaries' goals are as follows:

- to increase their own privacy (by injecting more noise into model updates) without following the SGO's prescribed privacy configuration; or

- to infer information about other MGOs' data from the shared model updates.

This chapter focuses on misconfiguration attacks on privacy parameters and their impact on the accuracy of the global model. Other types of attacks commonly studied in federated learning are outside the scope of this work.

Misbehaviour scenarios.
In this work, two scenarios of clients' malicious behaviour are considered. Both scenarios imply that some clients in the federated learning setting do not follow the algorithm. The objective of malicious users is to benefit from an enhanced level of privacy by not following the privacy configuration defined by the central server. More precisely, malicious users are users who choose smaller values of the privacy budget compared to the values defined by the central server. Smaller values of the privacy budget result in higher noise being injected into the user's model updates, thereby enhancing the user's level of privacy. These scenarios help to understand the impact of clients' misbehaviour on the final model RMSE. The scenarios are as follows:

- The central server (SGO) determines the privacy configuration to be used during the training of the machine learning model. A number of malicious clients may not follow the algorithm by selecting a custom privacy configuration within the boundaries defined by the central server. In an ideal scenario, clients should use the privacy configuration defined by the central server. For example, if the central server defines the upper bound $\epsilon^{max} = 0.9$ and the lower bound $\epsilon^{min} = 0.3$, clients should follow this configuration, whereas the privacy budget for each client would be decreasing from 0.9 to 0.3. However, a malicious client may decide to set an initial privacy budget ϵ_m^{max} equal to 0.5 and keep $\epsilon_m^{min} = \epsilon^{min}$, aiming at enhanced level of privacy.
- The central server (SGO) determines the privacy configuration to be used during the training of the machine learning model. A number of malicious clients may choose a privacy budget outside the boundaries defined by the central server. Clients

aim to maximize their level of privacy, which means they often choose relatively small values for the privacy budget. To understand the effect of misbehavior in this scenario, the values of privacy budget below the lower bound defined by the central server are considered.

Although various attacks may occur in distributed systems, this work examines the effect of malicious behaviour only in terms of users not following the privacy configuration defined by the central server. All other types of attacks are beyond the scope of this work.

7.4 Evaluation and Discussion

This section presents the simulation results to evaluate the proposed federated learning scheme. In this work, a smart grid environment is studied as a specific use case, whereas the framework can be applied to other scenarios, such as the detection of fraudulent transactions, where banks collaboratively train a central model without sharing their data while applying custom privacy settings.

A smart grid is modelled with a Smart Grid Operator (SGO) and multiple microgrids, each represented by a Microgrid Operator (MGO). Real electricity consumption data from 100 London households [93] are randomly partitioned into multiple datasets, serving as the private historical datasets of individual MGOs. Each MGO trains a local forecasting model using its own data and shares noisy model updates with the SGO via an ϵ-LDP mechanism (Laplace noise). The central model is a feed-forward neural network implemented using Keras: two hidden layers with 32 and 16 ReLU units and a single linear output neuron for regression. Training and simulations are executed in Python on an Apple M1 CPU with 8 GB RAM. The Root Mean Square Error (RMSE) is used as the main evaluation metric, computed on client-side test sets after each training round. The experiments investigate: (i) the impact of injecting noise into local updates, (ii) the effect of different centrally defined privacy-budget intervals, and (iii) the effect of clients deviating from the recommended privacy configuration.

In this particular case study, the central server (SGO) has its own dataset on the energy consumption in the smart grid, which is utilized for validation of the central model, namely to perform energy demand forecasting. The validation process performed by the SGO results in RMSE metric. SGO randomly selects data from its dataset when validating the central model, thus making the process more realistic.

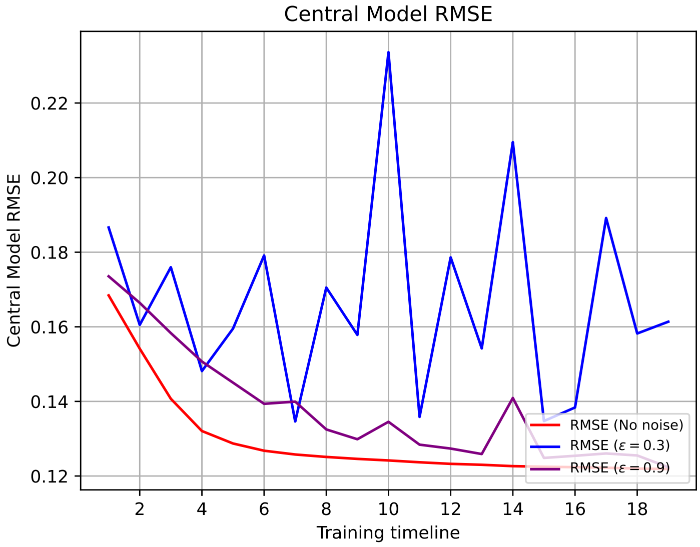

Figure 7.11: Central Model RMSE for different FL scenarios

First of all, to demonstrate the impact of the noise injected into the model updates on the central model RMSE, simulations were conducted for three different scenarios, namely without injecting noise, injecting a small amount of noise ($\epsilon = 0.9$), and injecting a high amount of noise ($\epsilon = 0.3$). Fig. 7.11 shows RMSE of the central machine learning model for different configurations of federated learning. In the first scenario (red curve), clients perform training on a local machine learning model and report their model updates to the central server without injecting noise into the model updates, which can be considered a classical federated learning approach. It can be seen that without injecting noise, RMSE of the central model is the lowest compared to other scenarios. In the second scenario (blue curve), clients use a fixed privacy budget $\epsilon = 0.3$, which provides a better privacy guarantee by injecting more noise into the model updates before sharing them with the central server. Although this configuration guarantees a high level of privacy for clients, RMSE of the central model is very high compared to the first scenario. In the third scenario (purple curve), clients use a fixed privacy budget $\epsilon = 0.9$, which provides some privacy guarantee by injecting a lower amount of noise compared to the second scenario. It leads to the reduced RMSE of the central model, as well as the reduced privacy guarantee. It can be clearly observed that the more noise is injected into the model updates, the higher RMSE of the central model will be. Although the first scenario yields the best results in terms of RMSE of the central model compared to other scenarios involving noise injection, it does not guarantee any privacy for the clients.

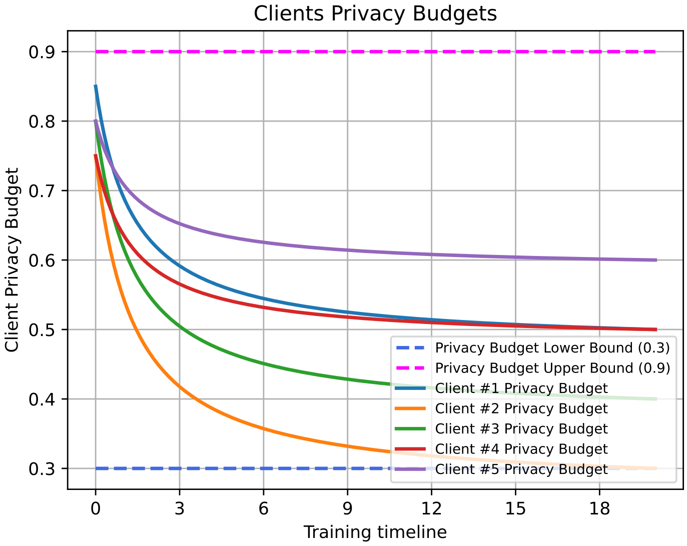

Figure 7.12: Clients Privacy Budgets

To explore the effect of different privacy configurations chosen by clients on the central model's RMSE, various privacy settings were configured within the boundaries defined by the central server for five clients (Fig. 7.12). It should be noted that ideally, clients should use the privacy configuration defined by the central server, namely the starting ϵ should be equal to ϵ^{max}, and ϵ in the final training round should be equal to ϵ^{min}. The purpose of randomly choosing an initial ϵ_m^{max} and a final ϵ_m^{min} ($\epsilon_m^{max} > \epsilon_m^{min}$) for each client is to examine how clients' misbehaviour in terms of not using the central server's privacy configuration affects the RMSE of the central model. Clients are considered semi-honest when they do not strictly follow the algorithm by choosing $\epsilon_m^{max} < \epsilon^{max}$ and/or $\epsilon_m^{min} > \epsilon^{min}$, while the privacy budget remains within the boundaries at any particular training round. In this particular example, SGO defined the number of training rounds equal to 20, the lower bound for the privacy budget as 0.3, and the upper bound as 0.9. As mentioned before, the privacy budget $\epsilon = 1$ cannot be used because the amount of noise injected into the model updates will be too small, which may lead to a data reconstruction attack. Similarly, $\epsilon = 0$, which guarantees the highest level of privacy, cannot be used because it would lead to an enormous amount of noise being injected into the model updates, drastically degrading the central model's accuracy. Taking into account the lower and upper bounds for the privacy budget defined by the central servers, clients set their privacy configurations within this interval ($[0.9; 0.3]$). An important aspect of the proposed framework is that the client's privacy budget is not fixed and changes over time according to the target RMSE equation. Thus, each client m randomly selects ϵ_m^{max} and ϵ_m^{min} ($\epsilon_m^{max} > \epsilon_m^{min}$), whereas the proposed algorithm adjusts the epsilon value throughout the training. At the beginning of the training, clients use a high privacy budget, which means that the amount of noise injected will be lower compared to the later training rounds. It allows the model to achieve the desired level of accuracy in a smaller number of rounds. For using a high privacy budget at the beginning, clients receive a bigger reward compared to the reward in later rounds.

To demonstrate the effect of a client not following the algorithm by using a fixed privacy configuration, a corresponding scenario was simulated. Fig. 7.13 shows the target RMSE (dashed magenta curve), average RMSE (dashed brown curve), and RMSE for five different clients over 20 training rounds when one of the clients (client 1) decided to use a fixed privacy budget parameter $\epsilon_m^{max} = \epsilon_m^{min} = 0.4$. The privacy configurations ($[\epsilon_m^{max};\epsilon_m^{min}]$) for clients 2-4 were set to $[0.85; 0.5]$, $[0.75; 0.3]$, $[0.8; 0.4]$, $[0.75; 0.5]$ respectively. RMSE is calculated for each client based on the evaluation of the central model using the client's private test dataset at the end of each training round. It can be seen that RMSE is decreasing, following the pattern of the target accuracy function, even though one of the clients did not follow the algorithm. It should be noted that even if a client uses a fixed privacy budget within the boundaries ($\epsilon^{max} < \epsilon_m > \epsilon^{min}$), the RMSE

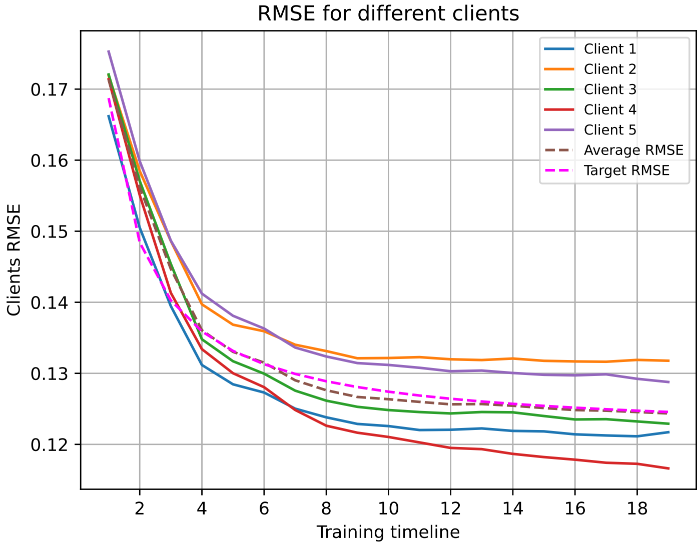

Figure 7.13: RMSE for different clients

for this client may show better results (as in this simulation) compared to honest clients who strictly follow the algorithm. This is happening due to the nature of the federated learning mechanism, where several clients collaboratively train the central model, which is then evaluated using each client's private test dataset.

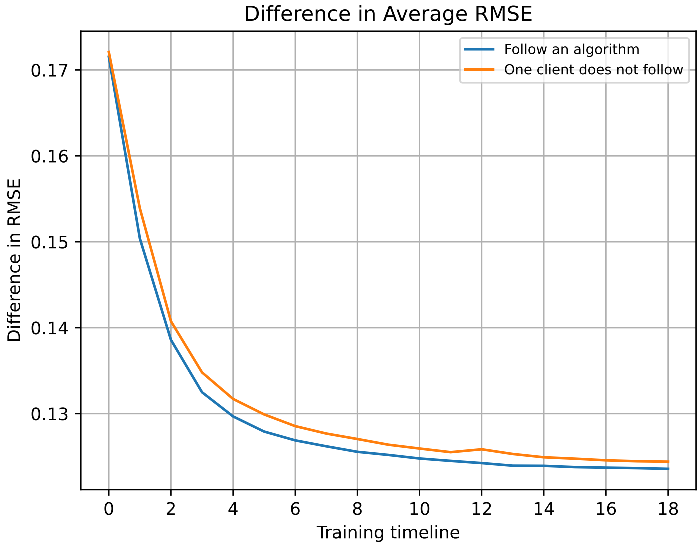

Figure 7.14: Difference in average RMSE

To examine the impact of a client using a fixed value for the privacy budget on the average RMSE among all clients, the following scenario was simulated. (Fig. 7.14) shows the RMSE for two simulation scenarios, namely when all clients follow the algorithm (blue curve) and when one of the clients does not follow the algorithm by using a fixed privacy budget configuration (orange curve). Similar to Fig. 7.13, the average RMSE is calculated based on the RMSE of all clients. It can be observed that in the scenario where one of the clients does not follow the algorithm, the average RMSE of the central model is slightly higher compared to the scenario where all clients follow the algorithm. The highest difference of 0.0023 was observed in the training round at $t = 4$, after which it stabilized and the difference was reduced in the subsequent rounds. At the final training round, the difference in average RMSE was around 0.00083. Although a slight deviation in RMSE of the final central model can be observed when some clients do not follow the algorithm, the RMSE of the central model stabilizes with time and follows the pattern of the target accuracy function.

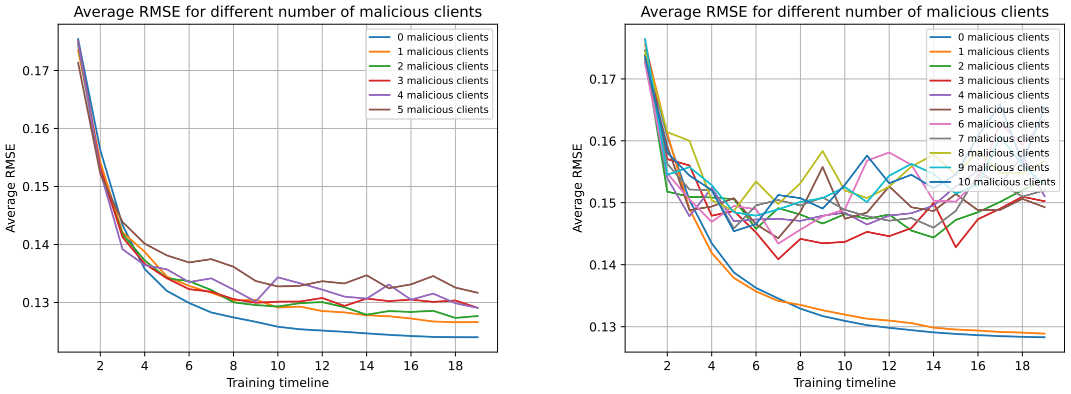

(a) $\epsilon_m \in [\epsilon^{min}; \epsilon^{max}]$

(b) $\epsilon_m \notin [\epsilon^{min}; \epsilon^{max}]$

Figure 7.15: Average RMSE depending on the number of clients not following the algorithm

To examine the impact of clients not following the algorithm by choosing a fixed privacy budget within the boundaries and outside the boundaries on the RMSE of the central model, the corresponding simulations were conducted. Fig. 7.15 shows the average RMSE for 10 clients over 20 training rounds, depending on the number of clients not following the algorithm by using a fixed privacy budget. Fig. 7.15a refers to the scenario when clients choose a fixed value of the privacy budget within the boundaries defined by the central server $[\epsilon^{max}; \epsilon^{min}]$, whereas Fig. 7.15b represents the scenario when clients choose a fixed value for the privacy budget outside the boundaries. The simulation results represent an average value of the RMSE based on five repetitions. It can be clearly observed that when clients do not follow the algorithm and choose a fixed privacy budget within the boundaries defined by the central server, the average RMSE of the central model increases with the increasing number of clients who don't follow the algorithm. More precisely, the smallest RMSE can be observed when all clients follow the algorithm, namely their privacy configuration ϵ_m changes over time. The highest RMSE is observed when all clients use a fixed privacy budget, which results in more noise being injected at later training rounds. When clients decide to choose the privacy budget outside the boundaries $[0.9; 0.3]$ ($\epsilon_m < \epsilon^{min}$), the model performs good when the number of malicious clients is small (in this scenario - one client). It can be observed that the average RMSE increases with the increasing number of clients choosing a small privacy budget ($\epsilon_m \in [0.3; 0.01]$), which causes an increased amount of noise injected into the model updates.

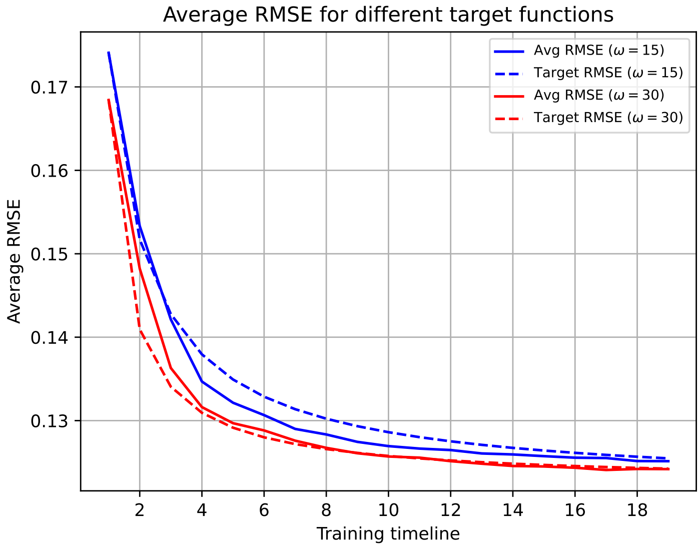

Figure 7.16: Difference in average RMSE

To evaluate the proposed algorithm's performance under various target functions, simulations were conducted for two different scenarios, with each scenario run five times for

10 clients over 20 training rounds. The privacy configuration ($[\epsilon^{max};\epsilon^{min}]$) defined by the central server is set to $[0.8;0.4]$, where clients randomly choose ϵ_m^{max} and ϵ_m^{min} within the defined interval. Fig. 7.16 shows the average RMSE for different target functions with $\omega = 15$ and $\omega = 30$. It can be observed that each curve representing the average RMSE follows its target RMSE function, and not another one. In the following simulations, RMSE of the central model refers to the average RMSE of all clients participating in the model training.

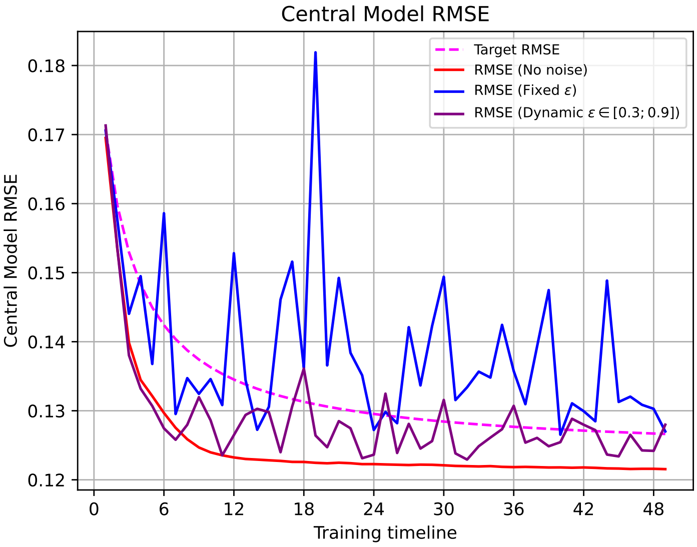

Figure 7.17: Central Model RMSE for different privacy configurations

The following simulations were conducted to demonstrate how RMSE of the central model is affected by different privacy settings, including a classical FL scenario without injecting noise into model updates, clients using a dynamic privacy budget, and clients misbehaving by using a fixed value of the privacy budget. Fig. 7.17 shows RMSE of the central machine learning model for three different scenarios over 50 training rounds. The interval for the privacy budget is set to $[0.9;0.3]$, within which clients determine their custom interval for the privacy budget ($[\epsilon_m^{max};\epsilon_m^{min}]$). An example of the clients' privacy budget configurations is shown in Fig. 7.12. The red curve (classical federated learning) shows RMSE of the central machine learning model in the scenario when no noise was injected into the model updates on the clients' side. It can be seen that in this scenario, RMSE is less compared to other scenarios. The dashed magenta curve shows the target RMSE (reward function φ) of the proposed federated learning scheme with $\omega = 15$. The blue curve shows RMSE of the central model in the scenario when clients use a fixed privacy budget ($\epsilon = 0.7$) when injecting noise into the model updates. It can be seen that RMSE in this case is very high compared to the classical federated learning approach. Although the privacy budget is relatively high ($\epsilon = 0.7$) compared to clients' configurations in Fig. 7.12, this approach does not perform well in terms of RMSE of the central model. As shown in Fig. 7.7, this curve shows the way the parameters ($epsilon$) of the proposed model change over training time, whereas it can be represented using various scales. In this particular example, the curve was scaled to the value of RMSE of the central model (the purple curve) after the first training round to illustrate the correlation between the target RMSE and the actual RMSE. The purple curve shows RMSE of the central model when clients follow the proposed algorithm and use a dynamic privacy budget that changes according to the target RMSE function φ. It can be seen that RMSE of the central model in this case is closer to the RMSE of the classical approach compared to the scenario when clients use a fixed privacy budget (the blue curve). Moreover, in this particular example, the actual RMSE is less than the target RMSE (magenta dashed curve) in most of the training rounds.

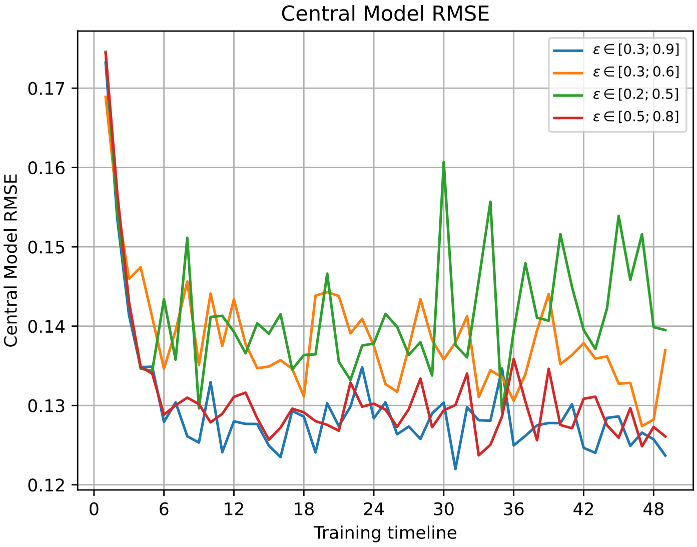

Figure 7.18: Central Model RMSE for different privacy budget intervals

To demonstrate the effect of different privacy configurations defined by the central server on the RMSE of the central model, the following simulations were conducted. Fig. 7.18 shows the RMSE of the central model for different privacy budget intervals defined by the central server. The training was performed over 50 training rounds, whereas the number of clients (MGOs) was set to 5, and clients' privacy budget intervals were randomly selected within the intervals defined by the central server. It can be seen that the RMSE of the central model in the scenario when $\epsilon \in [0.5; 0.2]$ (green curve) is higher compared to other scenarios because of the smaller values of ϵ used by the noise injection mechanism. This means that more noise was injected into the model updates, which has led to a higher error. In the second scenario, when $\epsilon \in [0.6; 0.3]$ (orange curve), the RMSE of the central model is lower compared to the previous scenario due to the fact that clients were able to choose higher ϵ, which means that less noise was injected into the model updates. In the scenario, when $\epsilon \in [0.9; 0.3]$ (blue curve), the RMSE of the central model is much better compared to other scenarios because the privacy budget interval defined by the central server empowers clients to choose higher ϵ. It should be noted that the privacy budget intervals on the clients' side are generated randomly, which means that not all clients choose a low ϵ. In contrast, when the central server forces clients to use a low privacy budget, the RMSE of the central model becomes higher (green curve) due to the fact that clients could not choose a higher ϵ.

7.5 Summary

Recognizing the growing privacy concerns associated with centralized data collection in energy demand forecasting, this chapter presented a novel privacy-preserving energy demand forecasting scheme for smart grids composed of multiple microgrids. The proposed approach integrates DP mechanisms into a federated learning framework, allowing microgrid operators to collaboratively train an accurate machine learning model for energy demand forecasting without exposing raw energy usage data.

The chapter introduced two key contributions. First, a dynamic privacy configuration mechanism was proposed, in which the central server defines a time-dependent target accuracy function to drive the evolution of the privacy budget during training. This approach enables a more coordinated and utility-aware privacy mechanism, ensuring a balance between privacy and model performance throughout the training process. The system also allows individual clients to define their own privacy configurations within the centrally specified range, thus preserving customization while maintaining consistency with global goals.

Secondly, the chapter introduced a reward allocation model aligned with the target accuracy function. A mathematical formulation was introduced to calculate reward distribution over training rounds, incentivizing participants to contribute honestly to the model training. The reward function is inversely related to the growth rate of the targeted model accuracy, encouraging meaningful early contributions and discouraging behavior that could degrade model performance.

In conclusion, the research in this chapter has presented a new privacy-aware forecasting framework that combines federated learning with dynamic differential privacy and reward mechanisms, enabling accurate and robust energy demand prediction while preserving user confidentiality in smart grid environments.

Chapter 8

Conclusion and Future Work

8.1 Overview

Modern energy systems are undergoing a transformative shift driven by Smart Grid technologies that integrate renewable energy sources and distributed energy sources, as well as advanced data-driven technologies. The increasing reliance on distributed energy resources, including photovoltaic panels and wind turbines, requires new mechanisms for energy trading, data aggregation and sharing, as well as energy demand forecasting that prioritizes the efficiency and privacy of end-users. At the same time, regulatory frameworks, such as the GDPR, require that end-users retain control over their personal data (including energy consumption data), which creates the need for redesigning traditional energy management approaches. The overarching aim of this thesis is to explore the aforementioned challenges in the Smart Grid ecosystem, including inefficient energy trading, data aggregation, privacy concerns, and energy demand forecasting. Due to a wide range of challenges and the large scope of the research domain, this thesis focuses on addressing the research objectives in four steps.

First, in Chapter 4, the research focuses on designing a novel dynamic pricing model to facilitate efficient energy trading within a microgrid of prosumers with photovoltaic panels. Although various pricing models are present in the literature, most approaches assume consistent user engagement or rely on complex bidding mechanisms. Unlike existing solutions, the proposed scheme introduces a deviation-sensitive pricing mechanism that ensures fair distributions of costs and profits for energy community participants based on their actual versus predicted energy consumption and generation patterns. The developed pricing scheme adapts to real-time conditions and integrates energy transfer losses into cost calculations. The inclusion of penalty and reward mechanisms based on deviation from forecasted behavior ensures accountability and enhances the reliability of energy predictions, thereby maintaining fairness among participants, while offering more favorable buying and selling prices compared to the main grid. According to the simulation results, the total energy usage cost of all prosumers may be reduced by 61-73%, compared to the maximum of 23.77% achieved by existing approaches.

Secondly, Chapter 5 focuses on the development of a blockchain-based data aggregation scheme that integrates on-chain and off-chain storage, ensuring both data integrity and scalability. Most existing approaches often compromise on user privacy, add computational overhead (e.g., homomorphic encryption), and contribute to unmanageable blockchain growth. These challenges are particularly critical in modern energy systems, where edge devices, such as smart meters, must handle sensitive data of prosumers in a secure fashion. The proposed scheme allows the preservation of end-users' privacy against cyber threats while maintaining high aggregation accuracy. The system eliminates the risks associated with centralized architectures by leveraging a dual-reporting mechanism, where each smart meter transmits obfuscated energy readings to two randomly selected edge nodes. This approach ensures that even in the presence of eavesdropping or man-in-the-middle attacks, the original consumption data cannot be reconstructed. The use of blockchain and smart contracts ensures tamper-proof storage and verifiable data integrity without incurring unbounded blockchain growth by using temporary off-chain storage. Simulation results indicate that the evaluation metrics remain relatively stable with an increasing number of prosumers. For example, the average time for aggregating the energy usage data and writing the aggregated data to the blockchain for a different number of prosumers stays within the interval from 0.7s to 1.2s. Notably, the number of blocks in the blockchain scales with the number of edge nodes, rather than the number of prosumers, demonstrating scalability. Critically, the total aggregated energy consumption retrieved from the blockchain equals the actual sum of prosumers' original readings.

Next, Chapter 6 focuses on the development of a novel user-centric privacy-preserving data-sharing framework for microgrids. First, the chapter introduces a novel bimodal probability distribution (the Boiarkin distribution) that provides control over the ranges of random variables from which the noise is drawn with high probability, and enables noise compensation. Second, a novel privacy-preserving mechanism that satisfies ϵLDP- has been introduced, which allows for adding more noise compared to the Laplace mechanism, while maintaining the same relative error in utility over time. Finally, a novel user-centric data-sharing scheme has been designed that allows the end-users to specify the boundaries for utility, whereas the maximum possible level of privacy (privacy budget ϵ) is provided by the scheme. Although various data-sharing mechanisms utilizing differential and local differential privacy are present in the literature, most of them heavily rely on trusted aggregators and provide end-users with limited control over the use of their personal information, which makes existing solutions suboptimal for user-centric applications. In the proposed framework, the amount of noise injected into the data can be fine-tuned by adjusting the parameter p of the Boiarkin mechanism, which affects the relative error for one iteration of the mechanism. According to the mutual information score analysis, injecting more noise using the Boiarkin mechanism ($p \geq 0.2$) makes the noisy and original data more independent compared to the Laplace mechanism. Thus, the Boiarkin mechanism can provide better privacy compared to the Laplace mechanism, whereas the relative error in utility is always within the boundaries defined by the end-user.

Finally, Chapter 7 focuses on the development of a novel privacy-preserving framework for energy demand forecasting in smart grids by integrating federated learning and local differential privacy. First, an approach is proposed to determine an equation for the target accuracy over time, enabling the central server to determine how quickly the desired level of accuracy should be achieved during model training. Secondly, an approach to formulate a dependency between the target accuracy of the central machine learning model and the privacy budget configuration is proposed. Most of the existing solutions propose the use of a specific value for the privacy budget ϵ determined based on the number of simulations performed for a particular use case, and cannot guarantee the same outcome for another application scenario. In contrast, in the proposed scheme, the privacy budget used by clients changes over time, which is mathematically linked to the target accuracy. To reward clients for their contribution, a dependency between the target accuracy and the reward is mathematically formulated. In the simulations, the Root Mean Square Error (RMSE) metric is used to evaluate the accuracy of the federated learning system. The results demonstrate that the average RMSE of the central model follows its target function, establishing a theoretical foundation for linking privacy configuration, machine learning model training algorithms, and reward mechanisms. Notably, simulations show that when a few clients (1 out of 5 or 10) deviate from the protocol by using fixed or even out-of-bound privacy budgets, the impact on the central model's RMSE is negligible. When all clients adhere to the algorithm and use dynamically adjusted privacy budgets, the model consistently approaches the target RMSE with reduced error over time. The results show that the actual RMSE of the proposed model is lower when the privacy budget used by clients changes over time compared to a fixed value used in most existing schemes.

8.2 Future Work

After researching areas of energy trading, privacy, and demand forecasting in modern smart grids and energy systems, this thesis presents a number of approaches aimed at enhancing user privacy, energy trading efficiency, and privacy-preserving data sharing within smart grid environments. This research began with the design of a dynamic pricing model for a microgrid of prosumers that incentivizes local energy trading while reducing costs for prosumers. Next, a blockchain-based data aggregation scheme was proposed by combining blockchain, smart contracts, and off-chain storage to ensure the robust management of energy consumption data within edge-enabled microgrids. To further safeguard the personal data of end-users, this work introduces a privacy-preserving user-centric data-sharing scheme for a microgrid of prosumers, utilizing local differential privacy. This enables end-users to control data utility while automatically providing the maximum possible level of privacy based on utility thresholds. Finally, a privacy-preserving energy demand forecasting scheme was designed by integrating federated learning and local differential privacy. In the proposed scheme, dependencies between accuracy, privacy, and rewards have been mathematically formulated. Unlike existing models that use static privacy configurations, the proposed scheme involves changing privacy configurations over time within the boundaries defined by the central server.

While this thesis successfully addresses critical gaps in the smart grid environment, including end-users' privacy, privacy-preserving data sharing, and energy trading, it also opens doors for future research. Specifically, this work focused on four primary domains: energy trading, blockchain-based data aggregation, privacy-preserving data sharing, and privacy-preserving energy demand forecasting in smart grids. These contributions are essential for enabling robust and efficient energy management within modern energy systems.

The following outlines potential future research directions that can further enrich the foundational work presented in this thesis.

- **Adaptive and Personalized Privacy Mechanisms**

  One potential future research direction is to explore personalized privacy configurations that adapt to users' preferences, behaviors, and contextual factors. While this thesis explored mechanisms to tune privacy settings based on the utility preferences, further work is needed to:

  - Develop user profiling techniques to recommend and adopt privacy settings (e.g., privacy budget values) tailored to individual behavior, risk sensitivity, or context, such as location, time, or activity.
  - Investigate reverse privacy engineering, allowing users to specify measurable goals (e.g., acceptable accuracy degradation or monetary cost) and derive appropriate privacy budgets from them. The future research could be based on the model proposed in Chapter 6.
  - Develop a mechanism to protect a system from malicious clients by designing a model that adjusts its parameters based on the impact caused by misbehaving

clients.

- **Privacy for Time-Aware and Sequential Data**

  Another possible future research direction is to address privacy risks over time, particularly given the time-series nature of energy data and the risks associated with sequential data releases. Further work can focus on the following aspects:

  - Creating dynamic privacy budgets that adapt over time or based on release history, minimizing cumulative privacy loss.
  - Proposing epsilon scheduling strategies or algorithms that optimize the privacy-utility trade-off across multiple data releases.
  - Studying sensitivity re-evaluation in light of temporal patterns or changes in user behavior, possibly recalibrating privacy guarantees accordingly.

- **Generalization Beyond Smart Grids and Toward Federated Utility Incentives**

  Finally, the methods developed could be extended beyond energy systems and enhanced within federated learning settings. Some ideas for further research include:

  - Application of privacy-preserving models to other IoT domains (e.g., vehicular telemetry for insurance), where the requirements of privacy and utility vary across applications.
  - In-depth investigation of reward mechanisms in federated learning to incentivize participants who contribute high-utility models, while still respecting heterogeneous privacy settings. It is worth considering the incorporation of reputation mechanisms.
  - Development of new models for multi-objective optimization in federated settings, balancing user-specific privacy budgets, model accuracy, and resource usage.

References

[1] Enerdata, *Electricity Forecast Consumption — Long-term electricity*, 2022 (accessed Oct 6, 2022), https://eneroutlook.enerdata.net/forecast-world-electricity-consumption.html.

[2] IEA, *Global Energy Review 2020*, 2022 (accessed Oct 6, 2022), https://www.iea.org/reports/global-energy-review-2020/electricity.

[3] Le Monde, *Global coal consumption still at a record high*, 2024 (accessed Nov 12, 2024), https://www.lemonde.fr/en/economy/article/2024/10/02/coal-global-consumption-still-at-a-record-high_6727918_19.html.

[4] IEA, *Pathways for the energy mix*, 2024 (accessed Nov 10, 2024), https://www.iea.org/reports/world-energy-outlook-2024/pathways-for-the-energy-mix.

[5] European Commission, *Renewable Energy Directive*, 2023 (accessed Oct 1, 2024), https://energy.ec.europa.eu/topics/renewable-energy/renewable-energy-directive-targets-and-rules/renewable-energy-directive_en.

[6] Department for Energy Security and Net Zero, *Powering up Britain*, 2023 (accessed Sep 10, 2024), https://www.gov.uk/government/publications/powering-up-britain.

[7] IEA *Net Zero by 2050*, 2021 (accessed Oct 6, 2022), https://www.iea.org/reports/net-zero-by-2050.

[8] Reuters, *Europe's economic woes may worsen as key power prices rise*, 2024 (accessed Nov 20, 2024), https://www.reuters.com/business/energy/europes-economic-woes-may-worsen-key-power-prices-rise-maguire-2024-11-20.

[9] Dow Jones & Company, Inc, *12 Stocks to Play Growing Energy Demand, From Oil and Gas to Nuclear and Solar*, 2024 (accessed Aug 11, 2024), https://www.barrons.com/articles/energy-roundtable-oil-gas-nuclear-solar-stock-picks-da4948af.

[10] IEA, *Electricity 2024*, 2024 (accessed Apr 17, 2024), https://www.iea.org/reports/electricity-2024.

[11] F. Alfaverh, M. Denai, and Y. Sun, "A dynamic peer-to-peer electricity market model for a community microgrid with price-based demand response," *IEEE Transactions on Smart Grid*, vol. 14, pp. 3976–3991, 2023.

[12] W. Lou, S. Zhu, B. Xu, T. Zhu, L. Sun, M. Wang, and X. Wang, "Integrated interactive control of distribution systems with multi-building microgrids based on game theory," *Buildings*, 2024.

[13] European Data Protection Supervisor, *TechDispatch #2: Smart Meters in Smart Homes*, 2019 (accessed Jun 11, 2023), https://www.edps.europa.eu/data-protection/our-work/publications/techdispatch/techdispatch-2-smart-meters-smart-homes_en.

[14] T. A. Alghamdi, R. Khalid, and N. Javaid, "A survey of blockchain based systems: Scalability issues and solutions, applications and future challenges," *IEEE Access*, vol. 12, pp. 79 626–79 651, 2024.

[15] M. Krichen, M. Ammi, A. Mihoub, and M. Almutiq, "Blockchain for modern applications: A survey," *Sensors*, vol. 22, no. 14, 2022. [Online]. Available: https://www.mdpi.com/1424-8220/22/14/5274

[16] R. Cummings, D. Desfontaines, D. Evans, R. Geambasu, Y. Huang, M. Jagielski, P. Kairouz, G. Kamath, S. Oh, O. Ohrimenko, N. Papernot, R. Rogers, M. Shen, S. Song, W. Su, A. Terzis, A. Thakurta, S. Vassilvitskii, Y.-X. Wang, L. Xiong, S. Yekhanin, D. Yu, H. Zhang, and W. Zhang, "Advancing Differential Privacy: Where We Are Now and Future Directions for Real-World Deployment," *Harvard Data Science Review*, vol. 6, no. 1, jan 16 2024, https://hdsr.mitpress.mit.edu/pub/sl9we8gh.

[17] M. Yang, T. Guo, T. Zhu, I. Tjuawinata, J. Zhao, and K.-Y. Lam, "Local differential privacy and its applications: A comprehensive survey," *Computer*

Standards & Interfaces, vol. 89, p. 103827, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0920548923001083

[18] B. Yurdem, M. Kuzlu, M. K. Gullu, F. O. Catak, and M. Tabassum, “Federated learning: Overview, strategies, applications, tools and future directions,” *Heliyon*, vol. 10, no. 19, p. e38137, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2405844024141680

[19] S. Banabilah, M. Aloqaily, E. Alsayed, N. Malik, and Y. Jararweh, “Federated learning review: Fundamentals, enabling technologies, and future applications,” *Information Processing & Management*, vol. 59, no. 6, p. 103061, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306457322001649

[20] V. Boiarkin, M. Rajarajan, J. Al-Zaili, and W. Asif, “A novel dynamic pricing model for a microgrid of prosumers with photovoltaic systems,” *Applied Energy*, vol. 342, July 2023, © 2023 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/). [Online]. Available: https://openaccess.city.ac.uk/id/eprint/30328/

[21] V. Boiarkin and M. Rajarajan, “A novel blockchain-based data-aggregation scheme for edge-enabled microgrid of prosumers,” in *2022 Fourth International Conference on Blockchain Computing and Applications (BCCA)*, 2022, pp. 63–68.

[22] V. Boiarkin, B. Bogaz Zarpelao, J. Al-Zaili, and M. Rajarajan, “A privacy-preserving user-centric data-sharing scheme,” *IEEE Access*, vol. 12, pp. 149 977–149 987, October 2024, 2024 The Authors. This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0. [Online]. Available: https://openaccess.city.ac.uk/id/eprint/33969/

[23] W. Tushar, T. K. Saha, C. Yuen, D. Smith, and H. V. Poor, “Peer-to-peer trading in electricity networks: An overview,” *IEEE Transactions on Smart Grid*, vol. 11, no. 4, pp. 3185–3200, 2020.

[24] L. Ali, S. Muyeen, H. Bizhani, and A. Ghosh, “A peer-to-peer energy trading for a clustered microgrid – game theoretical approach,” *International Journal of Electrical Power & Energy Systems*, vol. 133, p. 107307, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0142061521005469

[25] B. N. Alhasnawi, B. H. Jasim, B. E. Sedhom, and J. M. Guerrero, “Consensus algorithm-based coalition game theory for demand management scheme in smart microgrid,” *Sustainable Cities and Society*, vol. 74, p. 103248, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2210670721005254

[26] M. H. Ullah and J.-D. Park, “Dlmp integrated p2p2g energy trading in distribution-level grid-interactive transactive energy systems,” *Applied Energy*, vol. 312, p. 118592, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261922000721

[27] W. Tushar, T. K. Saha, C. Yuen, T. Morstyn, M. D. McCulloch, H. V. Poor, and K. L. Wood, “A motivational game-theoretic approach for peer-to-peer energy trading in the smart grid,” *Applied Energy*, vol. 243, pp. 10–20, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261919305185

[28] A. Paudel, K. Chaudhari, C. Long, and H. B. Gooi, “Peer-to-peer energy trading in a prosumer-based community microgrid: A game-theoretic model,” *IEEE Transactions on Industrial Electronics*, vol. 66, no. 8, pp. 6087–6097, 2019.

[29] Z. Zhang, H. Tang, P. Wang, Q. Huang, and W.-J. Lee, “Two-stage bidding strategy for peer-to-peer energy trading of nanogrid,” *IEEE Transactions on Industry Applications*, vol. 56, no. 2, pp. 1000–1009, 2020.

[30] M. R. Alam, M. St-Hilaire, and T. Kunz, “Peer-to-peer energy trading among smart homes,” *Applied Energy*, vol. 238, pp. 1434–1443, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261919300935

[31] L. He, Y. Liu, and J. Zhang, “Peer-to-peer energy sharing with battery storage: Energy pawn in the smart grid,” *Applied Energy*, vol. 297, p. 117129, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261921005699

[32] M. Mehdinejad, H. Shayanfar, and B. Mohammadi-Ivatloo, “Peer-to-peer decentralized energy trading framework for retailers and prosumers,” *Applied Energy*, vol. 308, p. 118310, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S030626192101566X

[33] A. Paudel, L. P. M. I. Sampath, J. Yang, and H. B. Gooi, “Peer-to-peer energy trading in smart grid considering power losses and network fees,” *IEEE Transactions on Smart Grid*, vol. 11, no. 6, pp. 4727–4737, 2020.

[34] T. AlSkaif, J. L. Crespo-Vazquez, M. Sekuloski, G. van Leeuwen, and J. P. S. Catalão, “Blockchain-based fully peer-to-peer energy trading strategies for residential energy systems,” *IEEE Transactions on Industrial Informatics*, vol. 18, no. 1, pp. 231–241, 2022.

[35] W. Liu, D. Qi, and F. Wen, “Intraday residential demand response scheme based on peer-to-peer energy trading,” *IEEE Transactions on Industrial Informatics*, vol. 16, no. 3, pp. 1823–1835, 2020.

[36] L. Li and S. Zhang, “Peer-to-peer multi-energy sharing for home microgrids: An integration of data-driven and model-driven approaches,” *International Journal of Electrical Power & Energy Systems*, vol. 133, p. 107243, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0142061521004828

[37] S. Xu, Y. Zhao, Y. Li, and Y. Zhou, “An iterative uniform-price auction mechanism for peer-to-peer energy trading in a community microgrid,” *Applied Energy*, vol. 298, p. 117088, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261921005390

[38] H. Haggi and W. Sun, "Multi-round double auction-enabled peer-to-peer energy exchange in active distribution networks," *IEEE Transactions on Smart Grid*, vol. 12, no. 5, pp. 4403–4414, 2021.

[39] L. He and J. Zhang, "A community sharing market with pv and energy storage: An adaptive bidding-based double-side auction mechanism," *IEEE Transactions on Smart Grid*, vol. 12, no. 3, pp. 2450–2461, 2021.

[40] A. Yu, X. Tang, Y. J. Zhang, and J. Huang, "Continuous group-wise double auction for prosumers in distribution-level markets," *IEEE Transactions on Smart Grid*, vol. 12, no. 6, pp. 4822–4833, 2021.

[41] Z. Zheng, T. Wang, A. K. Bashir, M. Alazab, S. Mumtaz, and X. Wang, "A decentralized mechanism based on differential privacy for privacy-preserving computation in smart grid," *IEEE Transactions on Computers*, vol. 71, no. 11, pp. 2915–2926, 2022.

[42] M. B. Gough, S. F. Santos, T. AlSkaif, M. S. Javadi, R. Castro, and J. P. S. Catalão, "Preserving privacy of smart meter data in a smart grid environment," *IEEE Transactions on Industrial Informatics*, vol. 18, no. 1, pp. 707–718, 2022.

[43] K. Kapusta, G. Memmi, and H. Noura, "Poster: A keyless efficient algorithm for data protection by means of fragmentation," 10 2016, pp. 1745–1747.

[44] Z. Guan, X. Zhou, P. Liu, L. Wu, and W. Yang, "A blockchain-based dual-side privacy-preserving multiparty computation scheme for edge-enabled smart grid," *IEEE Internet of Things Journal*, vol. 9, no. 16, pp. 14 287–14 299, 2022.

[45] S. Garg, K. Kaur, G. Kaddoum, J. J. P. C. Rodrigues, and M. Guizani, "Secure and lightweight authentication scheme for smart metering infrastructure in smart grid," *IEEE Transactions on Industrial Informatics*, vol. 16, no. 5, pp. 3548–3557, 2020.

[46] J. Srinivas, A. K. Das, X. Li, M. K. Khan, and M. Jo, "Designing anonymous signature-based authenticated key exchange scheme for internet of things-enabled

smart grid systems," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 7, pp. 4425–4436, 2021.

[47] J. Qian, Z. Cao, X. Dong, J. Shen, Z. Liu, and Y. Ye, "Two secure and efficient lightweight data aggregation schemes for smart grid," *IEEE Transactions on Smart Grid*, vol. 12, no. 3, pp. 2625–2637, 2021.

[48] X. Luo, K. Xue, J. Xu, Q. Sun, and Y. Zhang, "Blockchain based secure data aggregation and distributed power dispatching for microgrids," *IEEE Transactions on Smart Grid*, vol. 12, no. 6, pp. 5268–5279, 2021.

[49] Q. Yang and H. Wang, "Privacy-preserving transactive energy management for iot-aided smart homes via blockchain," *IEEE Internet of Things Journal*, vol. 8, no. 14, pp. 11 463–11 475, 2021.

[50] Y. Su, Y. Li, J. Li, and K. Zhang, "Lceda: Lightweight and communication-efficient data aggregation scheme for smart grid," *IEEE Internet of Things Journal*, vol. 8, no. 20, pp. 15 639–15 648, 2021.

[51] W. Lu, Z. Ren, J. Xu, and S. Chen, "Edge blockchain assisted lightweight privacy-preserving data aggregation for smart grid," *IEEE Transactions on Network and Service Management*, vol. 18, no. 2, pp. 1246–1259, 2021.

[52] Z. Wang, Y. Liu, Z. Ma, X. Liu, and J. Ma, "Lipsg: Lightweight privacy-preserving q-learning-based energy management for the iot-enabled smart grid," *IEEE Internet of Things Journal*, vol. 7, no. 5, pp. 3935–3947, 2020.

[53] T. Zhu, D. Ye, W. Wang, W. Zhou, and P. S. Yu, "More than privacy: Applying differential privacy in key areas of artificial intelligence," *IEEE Transactions on Knowledge and Data Engineering*, vol. 34, no. 6, pp. 2824–2843, 2022.

[54] B. Jiang, J. Li, G. Yue, and H. Song, "Differential privacy for industrial internet of things: Opportunities, applications, and challenges," *IEEE Internet of Things Journal*, vol. 8, no. 13, pp. 10 430–10 451, 2021.

[55] A. Krall, D. Finke, and H. Yang, “Mosaic privacy-preserving mechanisms for healthcare analytics,” *IEEE Journal of Biomedical and Health Informatics*, vol. 25, no. 6, pp. 2184–2192, 2021.

[56] J. Wei, Y. Lin, X. Yao, J. Zhang, and X. Liu, “Differential privacy-based genetic matching in personalized medicine,” *IEEE Transactions on Emerging Topics in Computing*, vol. 9, no. 3, pp. 1109–1125, 2021.

[57] J. Wei, Y. Lin, X. Yao, and J. Zhang, “Differential privacy-based location protection in spatial crowdsourcing,” *IEEE Transactions on Services Computing*, vol. 15, no. 1, pp. 45–58, 2022.

[58] Y.-E. Sun, H. Huang, W. Yang, S. Chen, and Y. Du, “Toward differential privacy for traffic measurement in vehicular cyber-physical systems,” *IEEE Transactions on Industrial Informatics*, vol. 18, no. 6, pp. 4078–4087, 2022.

[59] X. Zheng, M. Guan, X. Jia, L. Guo, and Y. Luo, “A matrix factorization recommendation system-based local differential privacy for protecting users’ sensitive data,” *IEEE Transactions on Computational Social Systems*, vol. 10, no. 3, pp. 1189–1198, 2023.

[60] L. Zhang, T. Zhu, P. Xiong, W. Zhou, and P. S. Yu, “A robust game-theoretical federated learning framework with joint differential privacy,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 4, pp. 3333–3346, 2023.

[61] B. Jiang, J. Li, H. Wang, and H. Song, “Privacy-preserving federated learning for industrial edge computing via hybrid differential privacy and adaptive compression,” *IEEE Transactions on Industrial Informatics*, vol. 19, no. 2, pp. 1136–1144, 2023.

[62] W. Lin, B. Li, and C. Wang, “Towards private learning on decentralized graphs with local differential privacy,” *IEEE Transactions on Information Forensics and Security*, vol. 17, pp. 2936–2946, 2022.

[63] N. Gai, K. Xue, B. Zhu, J. Yang, J. Liu, and D. He, “An efficient data aggregation scheme with local differential privacy in smart grid,” *Digital*

Communications and Networks, vol. 8, no. 3, pp. 333–342, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2352864822000049

[64] H. Gupta, P. Agarwal, K. Gupta, S. Baliarsingh, O. P. Vyas, and A. Puliafito, “Fedgrid: A secure framework with federated learning for energy optimization in the smart grid,” *Energies*, vol. 16, no. 24, 2023. [Online]. Available: https://www.mdpi.com/1996-1073/16/24/8097

[65] A. Iqbal, P. Gope, and B. Sikdar, “Privacy-preserving collaborative split learning framework for smart grid load forecasting,” 2024. [Online]. Available: https://arxiv.org/abs/2403.01438

[66] Y.-H. Lin and J.-C. Ciou, “A privacy-preserving distributed energy management framework based on vertical federated learning-based smart data cleaning for smart home electricity data,” *Internet of Things*, vol. 26, p. 101222, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S254266052400163X

[67] M. A. A. Sarker, B. Shanmugam, S. Azam, and S. Thennadil, “Enhancing smart grid load forecasting: An attention-based deep learning model integrated with federated learning and xai for security and interpretability,” *Intelligent Systems with Applications*, vol. 23, p. 200422, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2667305324000966

[68] R. Rahman, N. Kumar, and D. C. Nguyen, “Electrical load forecasting in smart grid: A personalized federated learning approach,” 2024. [Online]. Available: https://arxiv.org/abs/2411.10619

[69] M. Akin, Y. Canbay, and P. Canbay, “Privacy-preserving prediction of electricity consumption in smart grids,” in *2024 4th International Conference on Smart Grid and Renewable Energy (SGRE)*, 2024, pp. 1–6.

[70] P. Riedel, K. Belkilani, M. Reichert, G. Heilscher, and R. von Schwerin, “Enhancing pv feed-in power forecasting through federated learning with differential privacy

using lstm and gru," *Energy and AI*, vol. 18, p. 100452, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2666546824001186

[71] Y. Chang, J. Li, and W. Li, "2d2ps: A demand-driven privacy-preserving scheme for anonymous data sharing in smart grids," *Journal of Information Security and Applications*, vol. 74, p. 103466, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2214212623000509

[72] S. Lee, L. Xie, and D.-H. Choi, "Privacy-preserving energy management of a shared energy storage system for smart buildings: A federated deep reinforcement learning approach," *Sensors*, vol. 21, no. 14, 2021. [Online]. Available: https://www.mdpi.com/1424-8220/21/14/4898

[73] F. E. Abrahamsen, Y. Ai, and M. Cheffena, "Communication technologies for smart grid: A comprehensive survey," *Sensors*, vol. 21, no. 23, 2021. [Online]. Available: https://www.mdpi.com/1424-8220/21/23/8087

[74] J. Powell, A. McCafferty-Leroux, W. Hilal, and S. A. Gadsden, "Smart grids: A comprehensive survey of challenges, industry applications, and future trends," *Energy Reports*, vol. 11, pp. 5760–5785, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2352484724003299

[75] A. Gopstein, C. Nguyen, C. O'Fallon, N. Hastings, and D. A. Wollman, "Nist framework and roadmap for smart grid interoperability standards, release 4.0," 2021-02-18 00:02:00 2021. [Online]. Available: https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=931882

[76] "Ieee guide for smart grid interoperability of energy technology and information technology operation with the electric power system (eps), end-use applications, and loads," *IEEE Std 2030-2011*, pp. 1–126, 2011.

[77] E. Kabalci and Y. Kabalci, *Smart Grids and Their Communication Systems*, 1st ed. Springer Publishing Company, Incorporated, 2018.

[78] N. Moosavi, H. Taherdoost, N. Mohamed, M. Madanchian, Y. Farhaoui, and I. U. Khan, "Blockchain technology, structure, and applications: A survey," *Procedia Computer Science*, vol. 237, pp. 645–658, 2024, international Conference on Industry Sciences and Computer Science Innovation. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1877050924011669

[79] A. S. Rajasekaran, M. Azees, and F. Al-Turjman, "A comprehensive survey on blockchain technology," *Sustainable Energy Technologies and Assessments*, vol. 52, p. 102039, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2213138822000911

[80] Bitcoin, *Block Chain*, 2020 (accessed Oct 10, 2023), https://developer.bitcoin.org/devguide/block_chain.html.

[81] Ethereum, *Smart contract languages*, 2024 (accessed Aug 3, 2024), https://ethereum.org/en/developers/docs/smart-contracts/languages/.

[82] H. Bibi, M. Abolhasan, J. Lipman, M. Abdollahi, and W. Ni, "A comprehensive survey on privacy-preserving technologies for smart grids," *Computers and Electrical Engineering*, vol. 124, p. 110371, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0045790625003143

[83] M. Yang, L. Lyu, J. Zhao, T. Zhu, and K.-Y. Lam, "Local differential privacy and its applications: A comprehensive survey," 2020. [Online]. Available: https://arxiv.org/abs/2008.03686

[84] C. Dwork, "Differential privacy," in *Automata, Languages and Programming*, M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 1–12.

[85] T. Zhu, G. Li, W. Zhou, and P. S. Yu, "Differential privacy and applications," 1 2017. [Online]. Available: https://dro.deakin.edu.au/articles/book/Differential_privacy_and_applications/20822704

[86] P. Zhao, G. Zhang, S. Wan, G. Liu, and T. Umer, “A survey of local differential privacy for securing internet of vehicles,” *The Journal of Supercomputing*, vol. 76, 11 2020.

[87] Q. Yang, Y. Liu, Y. Cheng, Y. Kang, T. Chen, and H. Yu, “Federated learning,” *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 2019. [Online]. Available: https://api.semanticscholar.org/CorpusID:233176472

[88] IEA, *World Energy Outlook 2019, IEA, Paris*, 2019 (accessed Dec 8, 2021), https://www.iea.org/reports/world-energy-outlook-2019.

[89] IEA *Global Energy Review 2020, IEA, Paris*, 2020 (accessed Dec 8, 2021), https://www.iea.org/reports/global-energy-review-2020.

[90] K. Anoh, S. Maharjan, A. Ikpehai, Y. Zhang, and B. Adebisi, “Energy peer-to-peer trading in virtual microgrids in smart grids: A game-theoretic approach,” *IEEE Transactions on Smart Grid*, vol. 11, no. 2, pp. 1264–1275, 2020.

[91] Official Journal of the European Union, *Directive (EU) 2019/944 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on common rules for the internal market for electricity*, 2019 (accessed Mar 06, 2023), https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32019L0944&from=EN.

[92] Official Journal of the European Union , *Directive (EU) 2018/2001 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on the promotion of the use of energy from renewable sources*, 2018 (accessed Mar 06, 2023), https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32018L2001&from=EN.

[93] U. P. Networks, *SmartMeter Energy Consumption Data in London Households*, 2022 (accessed Apr 19, 2022), https://data.london.gov.uk/dataset/smartmeter-energy-use-data-in-london-households.

[94] U. P. Networks *Photovoltaic (PV) Solar Panel Energy Generation data*, 2022 (accessed Apr 19, 2022), https://data.london.gov.uk/dataset/photovoltaic--pv--solar-panel-energy-generation-data.

[95] N. Liu, X. Yu, C. Wang, C. Li, L. Ma, and J. Lei, “Energy-sharing model with price-based demand response for microgrids of peer-to-peer prosumers,” *IEEE Transactions on Power Systems*, vol. 32, no. 5, pp. 3569–3583, 2017.

[96] A. S. Yahaya, N. Javaid, F. A. Alzahrani, A. Rehman, I. Ullah, A. Shahid, and M. Shafiq, “Blockchain based sustainable local energy trading considering home energy management and demurrage mechanism,” *Sustainability*, vol. 12, no. 8, 2020.

[97] UKPower.co.uk Limited, *Compare energy prices per kWh*, 2021 (accessed Dec 8, 2021), https://www.ukpower.co.uk/home_energy/tariffs-per-unit-kwh.

[98] UK government, *Feed-in tariffs: get money for generating your own electricity*, 2021 (accessed Dec 8, 2021), https://www.gov.uk/feed-in-tariffs.

[99] S. Energy, *Time of Use Tariffs: Saving the UK Grid?*, 2021 (accessed Apr 11, 2022), https://blog.spiritenergy.co.uk/homeowner/time-of-use-tariffs-uk.

[100] GreenMatch, *What Is the Smart Export Guarantee?*, 2021 (accessed Apr 11, 2022), https://www.greenmatch.co.uk/green-energy/grants/smart-export-guarantee.

[101] M. B. Rasheed, M. A. Qureshi, N. Javaid, and T. Alquthami, “Dynamic pricing mechanism with the integration of renewable energy source in smart grid,” *IEEE Access*, vol. 8, pp. 16 876–16 892, 2020.

[102] H. T. Javed, M. O. Beg, H. Mujtaba, H. Majeed, and M. Asim, “Fairness in Real-Time Energy Pricing for Smart Grid Using Unsupervised Learning,” *The Computer Journal*, vol. 62, no. 3, pp. 414–429, 07 2018.

[103] K. Aurangzeb, S. Aslam, S. M. Mohsin, and M. Alhussein, “A fair pricing mechanism in smart grids for low energy consumption users,” *IEEE Access*, vol. 9, pp. 22 035–22 044, 2021.

[104] Department for Energy Security & Net Zero (DESNZ), *Energy white paper: Powering our net zero future*, 2020 (accessed Sep 3, 2024), https://www.gov.uk/government/publications/energy-white-paper-powering-our-net-zero-future.

[105] European Data Protection Supervisor (EDPS), *TechDispatch #2: Smart Meters in Smart Homes*, 2019 (accessed Aug 9, 2023), https://edps.europa.eu/data-protection/our-work/publications/techdispatch/techdispatch-2-smart-meters-smart-homes_en.

[106] Official Journal of the European Union, *REGULATION (EU) 2016/679 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)*, 2016 (accessed Aug 09, 2023), https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016R0679&from=EN.

[107] U.S. Department of Health & Human Services, *Health Insurance Portability and Accountability Act of 1996 (HIPAA)*, 2022 (accessed Aug 20, 2023), https://www.cdc.gov/phlp/publications/topic/hipaa.html.

[108] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Found. Trends Theor. Comput. Sci.*, vol. 9, no. 3–4, p. 211–407, aug 2014. [Online]. Available: https://doi.org/10.1561/0400000042

[109] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," *Journal of Privacy and Confidentiality*, vol. 7, no. 3, p. 17–51, May 2017. [Online]. Available: https://journalprivacyconfidentiality.org/index.php/jpc/article/view/405

[110] L. Lyu, K. Nandakumar, B. Rubinstein, J. Jin, J. Bedo, and M. Palaniswami, "Ppfa: Privacy preserving fog-enabled aggregation in smart grid," *IEEE Transactions on Industrial Informatics*, vol. 14, no. 8, pp. 3733–3744, 2018.

[111] M. Li, Y. Tian, J. Zhang, D. Fan, and D. Zhao, "The trade-off between privacy and utility in local differential privacy," in *2021 International Conference on Networking and Network Applications (NaNA)*, 2021, pp. 373–378.

[112] M. A. Mustafa, S. Cleemput, A. Aly, and A. Abidin, "A secure and privacy-preserving protocol for smart metering operational data collection," *IEEE Transactions on Smart Grid*, vol. 10, no. 6, pp. 6481–6490, 2019.

[113] B. Rassouli and D. Gündüz, "Optimal utility-privacy trade-off with total variation distance as a privacy measure," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 594–603, 2020.

[114] Department for Energy Security & Net Zero (DESNZ), *Smart Metering Implementation Programme: review of the Data Access and Privacy Framework*, 2018 (accessed Sep 10, 2023), https://www.gov.uk/government/publications/smart-metering-implementation-programme-review-of-the-data-access-and-privacy-framework.

[115] A. Grataloup, S. Jonas, and A. Meyer, "A review of federated learning in renewable energy applications: Potential, challenges, and future directions," *Energy and AI*, vol. 17, p. 100375, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2666546824000417

[116] H. Yang, M. Ge, K. Xiang, and J. Li, "Using highly compressed gradients in federated learning for data reconstruction attacks," *IEEE Transactions on Information Forensics and Security*, vol. 18, pp. 818–830, 2023.

[117] H. Hu, X. Zhang, Z. Salcic, L. Sun, K.-K. R. Choo, and G. Dobbie, "Source inference attacks: Beyond membership inference attacks in federated learning," *IEEE Transactions on Dependable and Secure Computing*, vol. 21, no. 4, pp. 3012–3029, 2024.

[118] E. Nowroozi, I. Haider, R. Taheri, and M. Conti, "Federated learning under attack: Exposing vulnerabilities through data poisoning attacks in computer networks," *IEEE Transactions on Network and Service Management*, vol. 22, no. 1, pp. 822–831, 2025.

[119] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-efficient learning of deep networks from decentralized data,” 2023. [Online]. Available: https://arxiv.org/abs/1602.05629