

Imperial College London

Department of Electrical and Electronic Engineering

Information-Theoretic Privacy in Smart Meters

Giulio Giaconi

July 2018

Supervised by Dr. Deniz Gündüz

A thesis submitted in fulfilment of the requirements for the degree of
Doctor of Philosophy in Electrical and Electronic Engineering of Imperial College London
and the Diploma of Imperial College London

Abstract

Smart grids promise to enhance drastically the efficiency of today's power networks. One of the key components of smart grids is the smart meter, which allows to monitor a user's electricity consumption with much higher resolution compared to conventional energy meters. However, the high resolution of smart meter measurements also entails serious privacy implications for the users, as it makes easier to distinguish the power signature of single appliances from the aggregate household consumption. This would lead an attacker, which may be a thief, a surveillance agency, or the utility provider itself, to gain insights into users' activities and behaviors.

In this dissertation we analyze several privacy-preserving techniques that protect users' privacy without diminishing the utility of smart grids. We adopt physical resources that are increasingly available at the users' premises, such as renewable energy sources and rechargeable batteries, and use them to minimize the information leaked about a user's electricity consumption, as well as the cost of energy. We deploy information-theoretic tools to characterize the fundamental limits of smart meter privacy, measuring privacy via mutual information, and characterizing single-letter expressions for the information leaked when considering infinite and zero-capacity rechargeable batteries. These scenarios represent lower and upper bounds on the privacy performance of more realistic settings with finite-capacity rechargeable batteries. When considering a finite-capacity battery, we express the information leakage as an additive quantity, and cast the problem as a stochastic control problem, which can be solved by dynamic programming. We also propose more empirical privacy-preserving strategies, testing their performance against real smart meter measurements and time of use pricing tariffs. In particular, we measure privacy as the squared difference between the smart meter measurements and a target profile, which we consider as a completely private power profile, and characterize the optimal trade-off between privacy and cost of energy.

Declaration of Originality

I hereby certify that this thesis is the result of my own work under the guidance of my Ph.D. advisor, Dr. Deniz Gündüz. Any ideas or quotations from the work of other people are appropriately referenced.

Imperial College London
London, United Kingdom
1 July 2018

Giulio Giaconi

Copyright Declaration

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

Acknowledgements

First of all, I would like to express my deepest gratitude to my supervisor, Dr. Deniz Gündüz, for giving me constant guidance and precious advices during my doctorate. Thanks to his vast knowledge and kind help, he has been an incredible source of motivation and inspiration for me, both from a professional and a human point of view. I have to thank him for passing his passion for research on to me, and, most importantly, for being able to motivate me every single time I stepped into his office. Moreover, I would like to acknowledge the financial support provided by the Engineering and Physical Sciences Research Council of the United Kingdom for funding my Ph.D. studies, and for enabling me to study in such an excellent research environment. I would also like to thank my examiners, Prof. Yiannis Demiris and Prof. Tobias Oechtering, for thoroughly reading this thesis, and for the interesting and useful discussion we had during the viva examination, which greatly improved the quality of this work.

I am also very thankful to the Intelligent Systems and Networks group's administrators, Patrick and Joan, who have always helped me without hesitation, always willing to go one step further to help me. I am also very grateful to my managers at BT Technology, Applied Research, Security Futures Practice, for supporting me in completing my thesis.

I have to thank all the people of the Information Processing and Communications Lab, who shared with me the ups and downs of conducting research, as well as uncountable coffee breaks, dinners, and interesting conversations: Iñaki, Pol, Morteza, Qianqian, Samuel, David, Mohammad, Sreejith, Junlin, Joan, Elif, Emre, Borzoo, Eirina, Nan; as well as the numerous visiting students. For the same reasons I have to thank Zaid, Kiarash, Yasin, Martina, Aikaterini, Rigas, Panagiotis, Gokce, Samuel M., Tobias, Theo, Yixing, Hyung Jin, Miguel, Ayse, Enrico, Giovanni, Carmine, Patricio, Milica, Huibo, Mihajlo. Thank you also to all of my new and old friends, with whom I have shared the best experiences of these years.

Vorrei ringraziare la mia Famiglia per il continuo supporto ed incoraggiamento, e per l'affetto incondizionato che mi ha offerto in tutti questi anni. Anche a migliaia di chilometri di distanza, ho potuto avvertire la loro costante vicinanza, senza la quale nulla di tutto ciò sarebbe stato possibile. Devo a loro ogni traguardo abbia raggiunto finora nella mia vita, e per questa ragione dedico a loro questa tesi. In particolare, desidero ringraziare mia madre per avermi trasmesso energia, positività ed ottimismo, che mi hanno aiutato in ogni aspetto della mia vita. Vorrei ringraziare i miei zii Marzio e Maria Palma, che sono semplicemente i migliori zii che si possano avere, e che mi hanno aiutato così tanto che ogni ringraziamento è insufficiente e riduttivo. Desidero ringraziare inoltre mia sorella Laura per la sua vicinanza, fisica oltre che emotiva, e per il bei momenti trascorsi insieme a Londra.

Infine, vorrei ringraziare Piera per il suo supporto impagabile, il suo amore e la sua pazienza infinita. In tutti questi anni sei sempre stata dalla mia parte, e hai avuto fiducia in me anche quando ho dubitato di me stesso. Grazie, questa tesi è dedicata anche a te.

London, United Kingdom

1 July 2018

Giulio

*Alla mia Famiglia,
la mia più grande fonte
di ispirazione.*

Contents

List of Figures	12
List of Tables	15
List of Acronyms	16
Notation	17
1 Introduction	18
1.1 Objectives	19
1.2 Outline and Contributions	20
2 Smart Meter Privacy by Means of Physical Resources	24
2.1 Introduction	24
2.2 Smart Meters for a Smart Grid	25
2.2.1 Smart Meter Privacy Risks	28
2.3 Privacy-Enabling Techniques for Smart Meters	30
2.3.1 Current Household Batteries, Typical Power Demands and Renewable Power Generation	34
2.4 A Signal-Processing Perspective on Smart Meter Privacy	35
2.4.1 The Energy Management Policy	38
2.5 A Heuristic Privacy Measure: Variations in the Grid Load Profile	40
2.5.1 Grid Load Variance as a Privacy Measure	40
2.5.2 Temporal and Spatial Similarities in the Grid Load as a Privacy Measure	43
2.5.3 Heuristic Algorithms	44
2.6 Theoretical Guarantees on Smart Meter Privacy	46
2.6.1 The Significance of Single-Letter Expressions	47
2.6.2 Markov Decision Process Formulation	47
2.6.3 Mutual Information as a Privacy Measure	49
2.6.3.1 Privacy with a Renewable Energy Source	50
2.6.3.2 Privacy with a Rechargeable Battery	52
2.6.3.3 Privacy with a Renewable Energy Source and a Rechargeable Battery	54
2.6.4 Detection Error Probability as a Privacy Measure	55
2.6.5 Fisher Information as a Privacy Measure	57
2.6.6 Empirical Mutual Information as a Privacy Measure	58
2.7 Conclusions	60
3 Fundamental Limits of Smart Meter Privacy	61
3.1 Overview	61

3.2	Introduction	61
3.3	System Model	63
3.4	Infinite Battery Capacity	68
3.4.1	Generated Renewable Energy not Known by the Utility Provider	69
3.4.2	Optimal Energy Management Policy for $B_{\max} = \infty$	69
3.4.3	Store-and-Hide Energy Management Policy	72
3.4.4	Generated Renewable Energy Known by the Utility Provider . .	73
3.5	Smart Meter System Without Energy Storage	75
3.5.1	Generated Renewable Energy not Known by the Utility Provider	76
3.5.2	Generated Renewable Energy Known by the Utility Provider . .	77
3.6	Binary Scenario	78
3.7	Finite Battery Capacity	80
3.7.1	Binary Alphabet: $\mathcal{X} = \mathcal{E} = \mathcal{Y} = \{0, 1\}$	80
3.7.1.1	Battery-Independent Policy	81
3.7.1.2	Battery-Conditioned Policy	83
3.7.2	Larger Alphabets: $ \mathcal{X} = \mathcal{Y} = \mathcal{E} > 2$	84
3.8	Continuous User Loads	86
3.8.1	No Battery - Renewable Energy not Known by the Utility Provider	88
3.8.2	No Battery - Renewable Energy Known by the Utility Provider .	88
3.9	Conclusions	88
4	Smart Meter Privacy with a Finite-Capacity Battery	90
4.1	Overview	90
4.2	Introduction	90
4.3	System Model	91
4.4	Renewable Energy not Known by the Utility Provider	93
4.4.1	Additive Formulation for the Information Leakage Rate	93
4.4.2	Markov Decision Process Formulation	95
4.5	Renewable Energy Known by the Utility Provider	99
4.5.1	Additive Formulation for the Information Leakage Rate	100
4.5.2	Markov Decision Process Formulation	101
4.6	Conclusions	103
5	Grid Load Variance as a Privacy Measure	104
5.1	Overview	104
5.2	Introduction	104
5.3	System Model	106
5.4	Target Load as a Piecewise Constant Function	109
5.4.1	Illustration of the Water-filling Solution	112
5.4.2	Selling Energy to the Grid	114
5.5	Numerical Results	114
5.5.1	Simulation Settings	115
5.6	Conclusions	119
6	Short Horizon Joint Privacy-Cost Optimization	120

6.1	Overview	120
6.2	Introduction	120
6.3	System Model	121
6.4	Target Load as a Constant Value	123
6.4.1	Solution to the Optimization Problem (6.5)	126
6.4.2	Illustration of the Water-filling Solution	128
6.4.3	Impact of the Duration of Prediction and Past Horizons	130
6.4.4	Alternative Privacy Measures	134
6.5	Target Load as Filtered User Load	136
6.6	A More Practical Energy Management Policy	137
6.7	Conclusions	140
7	Conclusions	141
	Research Challenges	144
	Bibliography	146
	Appendix A Proofs for Chapter 3	156
A.1	Proof of Lemma 3.5	156
A.2	Proof of Lemma 3.6	158
	Appendix B Proofs for Chapter 4	159
B.1	Proof of Theorem 4.1	159
B.2	Proof of Corollary 4.1	161
B.3	Proof of Theorem 4.2	163
B.4	Proof of Corollary 4.3	166

List of Figures

2.1	An example of a household electricity consumption profile with some appliances highlighted within the aggregated total consumption (data retrieved from the Dataport database [1]).	28
2.2	Some of the questions an attacker may be able to answer by having access to SM data.	29
2.3	Overview of the privacy-enabling approaches for SMs.	31
2.4	Representation of the system model. X_t , Y_t , E_t and B_t denote the consumer's power demand, i.e., the user load, the SM readings, i.e., the grid load, the energy produced by the RES, and the battery state of charge at TS t , respectively. The dashed line represents the meter readings being reported to the UP.	36
2.5	Example of user load, grid load, and constant target load profiles, where the "distance" $Y_t - W_t$ is highlighted. The aim of the algorithms presented in this section is to minimize the average squared distance. . . .	41
2.6	Examples of user load, grid load, and a piecewise constant target load profiles over the course of one day. The price periods are highlighted by arrows of different colors. Note that the target assumes a different constant value for each price period.	43
2.7	Representation of the system model where a single EMU and RES are shared among N users, each equipped with an SM. The EMU decides how much power each user can retrieve from the RES and from the SG.	52
3.1	Representation of the system model. X_t , Y_t , E_t and B_t denote the consumer's power demand, the SM readings, the energy produced by the RES, and the level of energy in the battery at time t , respectively. The dashed line represents the meter readings being reported to the UP.	64
3.2	Representation of the system model where the UP has perfect knowledge about the realizations of the renewable energy generation process, E^n , in addition to the power used from the grid that is reported through the SM readings, Y^n	74
3.3	Finite state diagram for the evolution of the level of energy in the RB with $\mathcal{B} = \{0, 1, \dots, B_{\max}\}$ and $\mathcal{X} = \mathcal{E} = \mathcal{Y} = \{0, 1\}$ for the battery-independent policy. The 4-tuple (x, e, v, y) represents, for every TS t , the values of the user load, the renewable energy produced, the power drawn from the battery by the EMU, and the grid load, respectively. . . .	81
3.4	Optimal p_v for the binary scenario and various battery capacities, when $q_x = 0.5$, for the battery-independent policy of Section 3.7.1.1.	82
3.5	Minimum information leakage with respect to the renewable energy generation rate p_e with $\mathcal{X} = \mathcal{E} = \mathcal{Y} = \{0, 1\}$, for the battery-conditioned policy of Section 3.7.1.2. As $B_{\max}$ increases, the performance rapidly approaches that of $B_{\max} = \infty$ and $\hat{P}_d = 1$	84

3.6	Minimum information leakage with respect to the renewable energy generation rate p_e with $\mathcal{X} = \mathcal{E} = \mathcal{Y} = \{0, 1, 2, 3, 4\}$. The leakage for $B_{\max} = \infty$ has been characterized by setting $\hat{P}_d = 4$	85
4.1	Representation of the system model where information about the realizations of the renewable energy process is available to the UP.	99
5.1	Representation of the system model. X_t , Y_t and $Y_t - X_t$ are the user load, the grid load, and the power exchanged with the RB at time t , respectively. The dashed line represents the SM readings being reported to the UP.	106
5.2	ToU tariff convention used in this chapter. $t_{c(i)}$, for $i = 1, \dots, M - 1$, are the TSs at which the price of energy changes, while $t_{c(0)} = 0$ and $t_{c(M)} = t_N$ are the extremes of the time interval considered.	110
5.3	Optimal grid and target load profiles for $\alpha = 0.5$, when energy selling is not allowed.	113
5.4	User, target and grid load profiles for a Powervault G200-LI-4KWH [2].	116
5.5	Comparison between allowing to sell energy to the UP and not allowing it, when using the piecewise constant target load, a Tesla Powerwall 2 [3] and $\alpha = 0.5$	117
5.6	Privacy-cost trade-off by using a Powervault G200-LI-4KWH RB [2]. . .	117
5.7	Privacy and cost versus battery capacity.	118
6.1	User, grid and target load profiles for the SHM and the LHM, for $\alpha = 0.5$ and $H_F = H_P = 2h$	125
6.2	Optimal grid and target load profiles, for $B_{\max} = 4$, $\hat{P}_c = \hat{P}_d = 2$, when energy selling is not allowed.	128
6.3	Optimal grid and target load profiles, for $B_{\max} = 20$, $\hat{P}_c = \hat{P}_d = 10$, when energy selling is allowed.	130
6.4	Comparison between different past and prediction horizons for $\alpha = 0.5$, when energy selling is not allowed.	131
6.5	Comparison between various past and prediction horizons for $\alpha = 0.5$, when energy selling is allowed.	132
6.6	Impact of the prediction horizon H_F on information leakage, cost and target load variance for $H_P = 2h$	132
6.7	Impact of the past horizon H_P on information leakage, cost and target load variance for $H_F = 2h$	133
6.8	Number of features versus H_F and H_P , when energy cannot be sold. . .	134
6.9	Power spectra versus H_F and H_P	135
6.10	Filtered target load scenario with cut-off frequency of 0.1mHz, $\alpha = 1$, and $H_F = H_P = 2h$	138
6.11	Practical energy management policy for $\alpha = 0.5$, $H_F = H_P = 2h$, $T_S = 1h$, and cut-off frequency set to 0.1mHz.	138
6.12	Filtered target load scenario with cut-off frequency of 0.1mHz, $\alpha = 1$, and no energy selling. For the SHM and the practical energy management policy we set $H_P = 2h$ and $T_S = 1h$	139

6.13	The privacy-cost trade-off, for $H_P = H_F = 2h$, $T_S = 1h$, energy selling allowed, and constant target load.	139
A.1	The evolution of the level of energy in the battery is represented by a random walk that starts at the beginning of the hiding phase and has a drift towards the positive direction of the axis representing the level of energy in the battery. We want to guarantee that the threshold α is never crossed by the random walk.	157

List of Tables

2.1	Time resolutions of some currently used SMs.	26
2.2	Specifications of some currently available residential RBs.	33
2.3	Distribution of average household power consumption (resolution refers to the measurement frequency). Values in each column indicate the percentage of time the average consumption falls into the corresponding interval.	33
2.4	Distribution of average power generated by residential photovoltaic systems. Values in each column indicate the percentage of time the average generation falls into the corresponding interval.	33
2.5	Specifications of the solar panels studied in the Microgen [4] database. Values in each column indicate the percentage of solar panels that satisfy the corresponding property.	34
3.1	Tuples and transition probabilities for the battery-independent policy when $\mathcal{X} = \mathcal{E} = \mathcal{Y} = \{0, 1\}$	81
5.1	Specifications of the RBs considered in the simulations in this chapter. .	115

List of Acronyms

BA	Blahut-Arimoto
BE	Best effort
DP	Dynamic program
DSO	Distribution system operator
EMU	Energy management unit
FI	Fisher information
i.i.d.	Independent and identically distributed
KKT	Karush Kuhn Tucker
KL	Kullback-Leibler
LHM	Long horizon model
MDP	Markov decision process
NILL	Non-intrusive load levelling
NILM	Non-intrusive load monitoring
RB	Rechargeable battery
RES	Renewable energy source
SG	Smart grid
SHM	Short horizon model
SLB	Shannon lower bound
SM	Smart meter
SMDM	Smart meter data manipulation
ToU	Time of use
TS	Time slot
TTP	Trusted third party
UDS	User demand shaping
UP	Utility provider

Notation

Throughout this dissertation we will use the following notation. We denote random variables by capital letters, e.g., X and Y , their realizations, i.e., the value they assume, by lower-case letters, e.g., x and y , and the corresponding alphabets, i.e., the set of values they can take on, by calligraphic letters, e.g., $\mathcal{X}$ and $\mathcal{Y}$. The probability distribution of a random variable X taking values in $\mathcal{X}$ is denoted by p_X . For integers $0 < a < b$, X_a^b denotes the sequence $X_a, X_{a+1}, \dots, X_b$, while $X^b \triangleq X_1^b$. All logarithms are in base 2, unless specified otherwise. When solving optimization problems, we denote the optimal value of a variable with a star, e.g., Y^* if the optimization is over Y .

Chapter 1

Introduction

The next-generation energy network, the so-called *smart grid* (SG), promises a tremendous increase in efficiency, safety and flexibility of managing the electricity grid as compared to the legacy energy network. This is needed today more than ever as the global energy consumption is growing at an unprecedented rate, and renewable energy sources (RESs) have to be seamlessly integrated into the grid to assure a sustainable human development. *Smart meters* (SMs) are among the crucial enablers of the SG concept; they supply accurate high-frequency information about users' household energy consumption to utility providers (UPs), that is, the entities that sell energy to customers, and to distribution system operators (DSOs), that is, the entities that operate the grid. The increased resolution and accuracy of SM data are essential for services such as time of use (ToU) pricing, rapid fault detection, energy theft prevention, and also provide user with more flexibility and near real-time control over their consumption. However, highly accurate and granular SM data also poses a threat to consumer privacy as non-intrusive load monitoring techniques (NILM) enable a malicious attacker to infer many details of a user's private life. The possibility of large-scale data collection from residential users has been a topic of heated public and political debate. If not properly addressed, the issue of privacy in smart meters would hinder the full development of smart grids; hence, privacy of SM data has recently become a topic of considerable research interest. In this regard, we expect that universities and research institutes will stress the issue of privacy in smart meters even further in order to push regulators to impose tighter rules about privacy.

1.1 Objectives

This dissertation analyzes privacy-enhancing energy management techniques that provide accurate energy consumption information to the grid operator, without sacrificing consumer privacy. In particular, we focus on techniques that shape and modify the actual user energy consumption, which we call the *user load*, by means of physical resources, such as rechargeable batteries (RBs), RESs, or demand shaping. The aim of these techniques is to make the SM measurements, which we call the *grid load*, as different as possible from the user load. Compared to techniques that are aimed at injecting noise or perturbations in the SM measurements, the techniques studied in this dissertation do not diminish the accuracy of the information provided to UPs, while at the same time guaranteeing privacy for the SG users. Moreover, these techniques have been receiving growing attention from the research community in the recent years, since the physical resources these techniques rely on are becoming increasingly available, thanks to government incentives and decreasing cost of solar panels and household RBs, as well as batteries for electric vehicles.

The main contribution of this dissertation is two-fold. Firstly, we aim at shedding further light on the fundamental limits of privacy achievable in limit scenarios. These limits represent the fundamental bounds on the performance of more practical SM settings, and help us understand the role played by the various components of the SM setting in terms of privacy performance. Moreover, the characterization of a rigorous mathematical framework and the use of information-theoretic tools enables to compare various privacy-preserving techniques and to establish theoretical guarantees on the level of privacy achieved, regardless of any monitoring technique and of the computational capability of the attacker. This is of significant importance, as monitoring techniques that are currently thought not to be feasible, may become available to attackers in the future, if computational capabilities improve, for example by the development of quantum computing, or if new methods are developed. However, since the presence of an RB introduces memory effects in the system, an absolute mathematical characterization for the level of privacy achieved in an SM system has been characterized only for specific settings under various simplifications, for example by

considering an independent and identically distributed (i.i.d.) or Markov user load or RES generation. For this reason, the second aim of this dissertation is to analyze more practical privacy-preserving strategies and test their performance against real SM measurements and ToU pricing tariffs.

1.2 Outline and Contributions

This dissertation presents first an extensive overview of the SM privacy problem and the main privacy-preserving approaches, which is conducted in Chapter 2, followed by several technical chapters dedicated to various privacy-preserving techniques. In the following, we outline the content and results of each chapter, as well as the corresponding publications.

Chapter 2

In Chapter 2 we present an overview of the main privacy-preserving techniques that make use of physical resources, that is, RESs or RBs, to guarantee privacy for the consumers. The content of this chapter has been accepted for publication in:

- Giulio Giaconi, Deniz Gündüz and H. Vincent Poor, “Privacy-Aware Smart Metering: Progress and Challenges”, *arXiv:1802.01166 [cs.IT]*, to appear in *IEEE Signal Processing Magazine*, May 2018.

Chapter 3

In Chapter 3 we study SM privacy by considering the presence of both an RES and an RB as a means of masking the user’s electricity consumption. The RB, which is considered to be constrained by instantaneous power constraints that limit the level of privacy attainable, is utilized to store the energy obtained from the RES. Privacy is measured by the *information leakage rate*, which denotes the average mutual information between the user’s real energy consumption, the user load, and the energy requested from the grid, the grid load, which the SM reads and reports to the UP. The impact on the privacy performance when assuming that the UP has knowledge of

the amount of renewable energy generated is also analyzed. The minimum information leakage rate is characterized as a computable information-theoretic single-letter expression when the battery capacity is either infinite or zero. Numerical results are presented for the finite battery capacity scenario to illustrate the potential privacy gains from the existence of an energy storage device. It is shown that, while the information leakage rate decreases with increasing availability of renewable energy, a larger storage capacity is needed to fully exploit the available energy to improve the privacy. The results in this chapter have been published in:

- Giulio Giaconi, Deniz Gündüz and H. Vincent Poor, “Smart Meter Privacy with an Energy Harvesting Device and Instantaneous Power Constraints”, in *Proceedings of the IEEE International Conference on Communications (ICC)*, London, UK, 8-12 June 2015, pp. 7216-7221.
- Giulio Giaconi, Deniz Gündüz and H. Vincent Poor, “Smart Meter Privacy With Renewable Energy and an Energy Storage Device”, in *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 1, pp. 129-142, January 2018.

Chapter 4

In Chapter 4 we analyze the SM privacy problem by considering the availability of an RES and an RB. We measure privacy by the average mutual information between the consumer’s power consumption and the renewable energy generation processes, and the power drawn from the grid, where the latter is known by the UP through the SM readings, and the former two are to be kept private. Differently from Chapter 3, in this chapter we allow the RB to store energy also from the grid. By expressing the information leakage as an additive quantity, we cast the problem as a stochastic control problem, and formulate the corresponding Bellman equations. The results in this chapter have been partially published in:

- Giulio Giaconi and Deniz Gündüz, “Smart meter privacy with renewable energy and a finite capacity battery”, in *Proceedings of the IEEE International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, Edinburgh, UK, 3-6 July 2016, pp. 1-5.

Chapter 5

In Chapter 5 we jointly optimize user privacy and cost of electricity by considering the presence of a finite-capacity RB, which can be utilized both to achieve privacy and to reduce the energy cost for the user. Privacy is measured via the mean squared distance between the SM readings and a target load, while ToU pricing is considered to compute the electricity cost. As target load, we consider a piecewise constant function, which sets a constant power value for each price period. The SM user has the possibility to sell electricity back to the UP to further improve the trade-off between the achievable privacy and the energy cost. The results in this chapter have been published in:

- Giulio Giaconi, Deniz Gündüz, and H. Vincent Poor, “Optimal Demand-Side Management for Joint Privacy-Cost Optimization with Energy Storage”, in *Proceedings of the IEEE International Conference on Smart Grid Communications (SmartGridComm)*, Dresden, Germany, 23-26 October 2017, pp. 265-270.

Chapter 6

In Chapter 6 we extend the analysis of Chapter 5, by allowing the user to know her future electricity consumption only for a short period of time. Moreover, we propose a privacy-preserving energy management policy, which characterizes its target load as a low-pass filtered version of the user load. Additionally, a simplified and more practical energy management policy is considered, which optimizes the energy management less frequently. Extensive numerical results are presented to compare the performances of various energy management policies in terms of the privacy-cost trade-off they achieve, considering a number of privacy indicators. The results in this chapter have been submitted for publication in:

- Giulio Giaconi, Deniz Gündüz, and H. Vincent Poor, *under revision in IEEE Transactions on Smart Grid*, June 2018.

Chapter 7

Finally, in Chapter 7 we provide the conclusions of the research presented in this dissertation, and discuss potential research directions that can be considered in the

future, as well as open questions and challenges that need to be addressed to pave the way to an effective protection of users' privacy in future SGs.

Other Publications

The following paper has also been published as a result of the work carried out during the Ph.D. studies, but its content has not been included in this dissertation:

- Jun-Xing Chin, Giulio Giaconi, Tomas Tinoco De Rubira, Deniz Gündüz, and Gabriela Hug, “Considering Time Correlation in the Estimation of Privacy Loss for Consumers with Smart Meters”, *Proceedings of the Power Systems Computation Conference (PSCC)*, Dublin, Ireland, 11-15 June 2018.

Chapter 2

Smart Meter Privacy by Means of Physical Resources

2.1 Introduction

The SG promises an exceptional increase in efficiency, safety and flexibility of managing the electricity grid as compared to the legacy energy network. This improvement is largely due to the utilization of SMs, which supply accurate high-frequency information about users' household energy consumption to a UP. However, the very characteristic that enables a massive improvements of the grid's performance, that is, the availability of high-resolution and accurate SM data, also poses a significant privacy threat to the consumers, as highly accurate and granular SM data may be used by a malicious attacker to infer many details of a user's private life. For this reason, many privacy-preserving techniques have been proposed in the literature. This chapter provides an overview of privacy-enhancing energy management techniques that provide accurate energy consumption information to the grid operator, without sacrificing consumer privacy. In particular, we focus on techniques that shape and modify the actual user power consumption by means of physical resources, such as RBs, RESs or demand shaping. A rigorous mathematical analysis of privacy is presented under various physical constraints on the available physical resources.

2.2 Smart Meters for a Smart Grid

The current energy grid is one of the engineering marvels of the 20th century. However, it has become inadequate to satisfy the steadily growing global electricity demand. In fact, the world energy consumption is predicted to increase 48% from 2012 to 2040 [5], driven by factors such as the growth of world's economy, the rise of the gross domestic product per person, the increase of world's population, the surge in penetration of electric vehicles, and the broader mobility revolution [6]. Other issues that need to be addressed are the effective integration of RESs and storage capabilities into the grid, the improvement of the grid's environmental sustainability, and the promotion of plug-in hybrid electric vehicles. To address these challenges, a new generation of electricity grid is being engineered, the SG. SGs are intended to substantially improve energy generation, transmission, distribution, consumption and security, providing improved reliability and quality of the electricity supply, quicker detection of energy outages and theft, better matching of energy supply with demand, and a better environmental sustainability by enabling an easier integration of distributed generation and storage capabilities [7]. The “smartness” of a SG resides in the advanced metering infrastructure, which enables two-way communication between the utility and its customers, and whose pivotal element in the distribution network is the SM, the device that monitors a user's electricity consumption in almost real-time.

In contrast to legacy grids, in which billing data is gathered at the end of a billing period, SMs send electricity consumption measurements automatically and at a much higher resolution. SMs enable two-way communication with the UP, transmitting a great amount of detailed information. SMs collect and send bidirectional readings of active, reactive and apparent power and energy, that is, the so-called 4-quadrant metering, that is purchased from the grid, or sold to the grid, if the user produces energy, for example by means of a photovoltaic panel. In the latter case, the user is referred to as a “prosumer”, that is, a producer and consumer of electricity at the same time, and can be financially rewarded for the energy sold to the grid. SMs also keep track of historical consumption data over the previous days, weeks and months, and provide high-resolution consumption data analytics to the customers to enable them to monitor

TABLE 2.1: Time resolutions of some currently used SMs.

Smart Meter Model	Time Resolution
Itron Centron [8]	1 min
REX2 [9]	5 min
Kamstrup Omnipower [10]	5 min
Enel Open Meter [11]	15 min

their energy consumption via an in-home display, web portal or smartphone application, in near real-time. SMs also send alerts about voltage quality measurements, helping UPs fulfill their obligations towards customers concerning energy, power and voltage quality, e.g., in accordance with the European standard EN 50160. Examples of these measurements include the root mean square voltage variations, e.g., voltage dropout, sags and swells, and the total harmonic distortion. Data used for billing, such as current ToU tariff, balance and debts, credit and prepayment modes, credit alerts and topping up, is also sent to the UP. SMs can detect if a tampering takes place, and send relevant data about it, along with security credentials for enabling the correct functioning of cryptographic protocols, e.g., hashing, digital signature, and cyclic redundancy check. Finally, SM firmware information and updates are also communicated.

The increased data resolution is crucial for enabling SG functionalities. Table 2.1 shows the smallest time resolution of some SMs currently in use, which is on the order of few minutes. The European Union recommends a time resolution of at least 15 minutes to allow the new SG functionalities [12]. For example, the current SM specifications in the UK impose that an SM should send integrated energy readings every 30 minutes to the UP, while the data sent to a user's in-home display can have a resolution of up to 10 seconds [13]. In Texas, SMs send measurements every 15 minutes [14]. It should be noted that, with the increased adoption of renewable energy generation by the prosumers, the increased penetration of electric vehicles and energy storage technologies, and the diversification of the energy market, it is expected that the SGs will become more volatile, requiring meter readings at a much higher rate in the near future [15].

SMs provide a wide range of benefits to all the parties in an SG. Thanks to SMs,

UPs can gain a better knowledge of their customers' needs, while reducing the cost of meter readings. SMs allow UPs to determine the electricity cost dynamically, as well as to produce more accurate bills, thus reducing customers' complaints and back-office rebilling. The implementation of ToU pricing can incentivize demand response and control customer behavior, while improved demand forecasts and load-shaping techniques can reduce peak electricity demands. Finally, energy theft can be detected more easily and quickly.

DSOs benefit from SMs as well, being able to better monitor and manage the grid. SMs allow DSOs to reduce operational costs and energy losses, and improve grid efficiency and system design, as well as distributed system state estimation and Volt and Var control. Moreover, DSOs are able to better match distributed resources with the ongoing electricity demand and the grid's power delivery capability, thus reducing the need to build new power plants.

Consumers themselves take advantage of SMs to monitor their consumption in near real-time, leading to better consumption awareness and energy usage management. Moreover, consumers receive accurate and timely billing services, with no more estimated bills, and benefit from ToU pricing, by shifting non urgent loads to off-peak price periods. Microgeneration and energy storage devices can be integrated more easily, and profits from selling the generated excess energy can be collected automatically. Failing or inefficient home appliances, unexpected activity or inactivity, and wasted energy are detected faster and more accurately, and switching between UPs is easier by requesting on-demand readings, which in turn increases the competition among UPs and reduces costs for consumers.

For the above reasons, the installation of SMs is proceeding rapidly, and is attracting massive investments globally. For example, all European Union countries are required to have 80% SM adoption by 2020 and 100% by 2022 [16], whereas China represents the current leading global SM market with more than 348 million SMs installed, accounting for 67.1% of tracked global installations. The SM market is expected to grow from an estimated \$12.79 billion in 2017 to \$19.98 billion by 2022, registering a compound annual growth rate of 9.34% [17]. Moreover, the global SM data analytics market, which includes demand response analytics and grid optimization tools, is expected to




FIGURE 2.1: An example of a household electricity consumption profile with some appliances highlighted within the aggregated total consumption (data retrieved from the Dataport database [1]).

reach \$4.6 billion by 2022 [18], while the global penetration of SMs is expected to climb from approximately 30% at the end of 2016 to 53% by the end of 2025 [19]. These figures show how timely and crucial the research in this field is, and highlight the need to quickly resolve potential obstructions that can threaten the future benefits from this critical technology.

2.2.1 Smart Meter Privacy Risks

The SM's ability to monitor a user's electricity consumption in almost real-time entails serious implications about consumer privacy. In fact, by employing NILM techniques, it is possible to identify the power signatures of specific appliances from the aggregated household SM measurements. NILM techniques date back to the work of George Hart in the 80s, who first proposed a prototype of a NILM device [20]. Since then, NILM techniques have improved in different directions, e.g., by assuming either high or low-frequency measurements, by considering known or learned signatures [21], or even by using off-the-shelf statistical methods without any a-priori knowledge of household activities [22]. An example of a typical power consumption profile along with some detected appliances is illustrated in Figure 2.1. As shown in Figure 2.2, the UP, a third party that has access to SM data, for example, by buying it from the UP, or




FIGURE 2.2: Some of the questions an attacker may be able to answer by having access to SM data.

a malicious eavesdropper, may gain insights into users' activities and behaviors, and determine, for example, a user's presence at home, her religious beliefs, disabilities, illnesses, and even the TV channel she is watching [23–25]. Apart from residential users, SM privacy is particularly critical for businesses, e.g., factories and data centers, as their power consumption profile may reveal sensitive information about the state of their businesses to their competitors. SM privacy has attracted significant public attention, and continues to be a topic of heated public and political debate, and even stopped the mandatory SM roll-out plan in the Netherlands in 2009, after a court decided that the forced installation of SMs would violate consumers' right to privacy, and would be in breach of the European Convention of Human Rights [26]. Indeed, concerns about consumer privacy threaten the widespread adoption of SMs and can be a major roadblock for this multi-billion dollar industry.

It is worth pointing out that the privacy problem in SMs is different from the SM data security problem [27]. In the latter there is a sharp distinction between legitimate users and malicious attackers, whereas in the privacy problem in the SM context, any legitimate receiver of data can also be considered malicious. To benefit from the

advantages provided by the SG, users need to share some information about their electricity consumption with the UP and DSO. However, by sharing accurate and real-time information about their energy consumption, consumers also expose their private lives and behaviors to the UP, which is a fully legitimate user and a potential malicious attacker at the same time. This renders traditional encryption techniques for data privacy ineffective in achieving privacy against the UP, and calls for novel privacy measures and privacy protection techniques.

2.3 Privacy-Enabling Techniques for Smart Meters

There is a growing literature on privacy-preserving methods for SMs, which can be classified into two main families. The first family, which we call the *SM data manipulation* (SMDM) approach, consists of techniques that process the SM data before reporting it to the UP, while the techniques in the second family, called *user demand shaping* (UDS), aim at modifying the user's actual power consumption. Considered within the first family are methods such as *data obfuscation*, *data aggregation*, *data anonymization*, and *down-sampling*. Data obfuscation, i.e., the perturbation of metering data by adding noise, is a classical privacy protection method, and has been adapted to SGs in [28] and [29]. In [30], *differential privacy*, a well-established concept in the data mining literature, is applied to SMs where noise is added not only to the user's energy consumption, via the RB, but also to the energy used for charging the RB itself to provide differential privacy guarantees. Along these lines, the authors in [31] introduce an information-theoretic framework to study the trade-off between the privacy obtained by altering the SM data and the utility of data for various SG functionalities. Note that, the more noise added to the data, the higher the privacy, but the less relevant and less useful the data is for monitoring and controlling the grid. In [31] an additive distortion measure is considered to model the utility, which allows the characterization of the optimal privacy-utility trade-off in an information-theoretic single-letter form. The data aggregation approach, proposed in [29], [32] and [33], considers sending the aggregate power measurements for a group of households so that the UP is prevented from distinguishing individual consumption patterns. The aggregation




FIGURE 2.3: Overview of the privacy-enabling approaches for SMs.

can be performed with or without the help of a *trusted third party* (TTP). The data anonymization approach, on the other hand, mainly considers utilizing pseudonyms rather than the real identities of consumers [34,35]. Another method proposed in [36], reduces the SMs sampling rate to a level that does not pose any privacy threat. However, the SMDM family suffers from the following shortcomings:

- Adding noise to the SM readings causes a mismatch between the reported values and the real energy consumption, which prevents DSOs and UPs from accurately monitoring the grid state, rapidly reacting to outages, energy theft or other problems, and producing accurate and timely billing services. These would significantly limit the benefits of SMs;
- DSOs, UPs, or more generally any eavesdropper can embed additional sensors right outside a household or a business (street level measurements are already available to the DSOs and UPs) to monitor the energy consumption, without fully relying on SM readings;
- The anonymization and aggregation techniques that include the presence of a TTP only shift the problem of trust from one entity (UP) to another (TTP).

These issues are avoided by the approaches in the UDS family, which directly modify the actual power consumption profile of the user, called the user load, rather than modifying the data sent to the UP, called the grid load. In this family the SM reports accurately the energy taken from the grid without any modification; however, this

is not the energy that is actually consumed by the appliances. This is achieved by filtering the user's actual electricity consumption via a rechargeable energy storage device, i.e., an RB, or by exploiting an RES, which can be used to partially hide the consumer's power consumption. Examples of RESs include solar panels or micro wind farms. Another technique is to partially shift user's demand. Thus, the idea is to physically differentiate the grid load from the user load. Note that the effect of using an RB or an RES can also be considered as adding noise to the household consumption, but the noise in this case corresponds to a physical variation in the power received from the grid. Moreover, differently from approaches in the SMDM family, the SM measurements provided by the UDS methods are exact and there is no issue of data mismatch between the SM data and the effective user demand from the grid. Thus, when UDS methods are deployed, the utility of SMs for the SG is not diminished since the users' energy consumption is neither misreported nor distorted. As a result, while the privacy-utility trade-off is of particular concern for the SMDM techniques, with the UDS techniques smart grid utility is never diminished and other trade-offs are considered instead, such as the *privacy-cost*, or the *privacy-wasted energy*. Figure 2.3 shows an overview of the privacy-enabling approaches.

The focus of this dissertation is on UDS techniques, which have been receiving growing attention from the research community in the recent years. The physical resources these techniques rely on, such as RBs or RESs at consumer premises, are already becoming increasingly available thanks to government incentives and decreasing cost of solar panels and household RBs, as well as the RBs for electric cars. Moreover, shaping and filtering users' actual power consumption by means of physical resources renders any data misreporting or distortion unnecessary, which, thus, do not undermine the utility of the SG concept itself. In this chapter we will present a signal processing perspective on SM privacy by treating the user load as a stochastic time series, which can be filtered and distorted by using an RB, an RES or demand shaping or scheduling. The available energy generated by the RES can also be modeled as a random sequence, whose statistics depend on the energy source (e.g., solar, wind) and the specifications of the renewable energy generator. Additionally, the finite-capacity battery imposes

TABLE 2.2: Specifications of some currently available residential RBs.

Residential RB	RB Capacity (kWh)	RB Charging Peak Power (kW)	RB Discharging Peak Power (kW)
Sunverge SIS-6848 [37]	7.7, 11.6, 15.5, 19.4	6.4	6
SonnenBatterie eco [38]	4 – 16	3 – 8	3 – 8
Tesla Powerwall 2 [3]	13.5	5	5
LG RESU 48V [39]	2.9, 5.9, 8.8	3, 4.2, 5	3, 4.2, 5
Panasonic Battery System LJ-SK56A [40]	5.3	2	2
Powervault G200-LI-2/4/6KWH [2]	2, 4, 6	0.8, 1.2	0.7, 1.4
Orison Panel [41]	2.2	1.8	1.8
Simpliphi PHI 3.4-48V [42]	3.4	1.5	1.5

TABLE 2.3: Distribution of average household power consumption (resolution refers to the measurement frequency). Values in each column indicate the percentage of time the average consumption falls into the corresponding interval.

Source	Location	Resolution	Time Frame	# of Houses	[0, 0.5] kW	(0.5, 1] kW	(1, 2] kW	(2, 3] kW	(3, 4] kW	(4, +∞) kW
Dataport [1]	Texas	60 mins	01/01/2016 - 31/05/2016	512	38	30	20	7	3	2
			01/01/2015 - 31/12/2015	703	36	26	20	9	5	4
			01/01/2014 - 31/12/2014	720	39	25	20	8	4	4
			01/01/2013 - 31/12/2013	419	35	25	21	9	5	5
			01/01/2012 - 31/12/2012	182	31	26	24	10	5	5
Intertek [43]	UK	2 mins	01/05/2010 - 31/07/2011	251	18	24	47	11	0	0
Dred [44]	Netherlands	1 sec	05/07/2015 - 05/12/2015	1	98	1.8	0.4	0	0	0
Uci [45]	France	1 min	16/12/2006 - 26/11/2010	1	47	9	28	8	4	2

TABLE 2.4: Distribution of average power generated by residential photovoltaic systems. Values in each column indicate the percentage of time the average generation falls into the corresponding interval.

Source	Location	Resolution	Time Frame	# of Houses	0 kW	(0, 0.5] kW	(0.5, 1] kW	(1, 2] kW	(2, 3] kW	(3, 4] kW	(4, +∞) kW
Dataport [1]	Texas	60 min	01/01/2012 - 31/05/2016	351	49	17	7	9	7	6	5
Microgen [4]	UK	30 min	01/01/2015 - 31/12/2015	100	51.7	36.4	9.8	2	0.1	0	0

instantaneous limitations on the available energy. We also note that such physical resources can also be used for cost minimization purposes by the users, e.g., by acquiring and storing energy over low-cost periods, and utilizing the stored energy in the RB and the energy generated by an RES over peak-cost periods. Accordingly, we also study the trade-off between privacy and cost, and the minimization of the wasted renewable energy. In the following, we describe and summarize the progress made in recent years towards quantifying the privacy leakage of SMs in a rigorous manner, report the most significant results, and highlight a number of future research directions.

TABLE 2.5: Specifications of the solar panels studied in the Microgen [4] database. Values in each column indicate the percentage of solar panels that satisfy the corresponding property.

Solar Panel Area (m^2)					Solar Panel Cell Type		Nominal Installed Capacity (kWp)			
(0, 15]	(15, 20]	(20, 25]	(25, 30]	(30, $+\infty$)	Monocrystalline	Polycrystalline	(0, 2]	(2, 3]	(3, 4]	(4, $+\infty$)
5	35	44	15	1	93	7	4	36	59	1

2.3.1 Current Household Batteries, Typical Power Demands and Renewable Power Generation

Table 2.2 lists the storage capacity and peak power of some of the currently available RBs for residential use. As it can be seen, the capacities are in the range of few kWhs. The typical household average power consumption also lies within the range of few kW, as shown in Table 2.3, where the distribution of the average user power consumption values over different years obtained from various databases is reported with various time resolutions. Analyzing the Dataport [1] database we observe that, independently of the period considered, the average user energy demand is less than 2 kWh for 80 – 90% of the time. Current batteries charged at full capacity would then be able to satisfy the demand continuously only for a few hours. However, completely covering the consumption over a few hours may come at the expense of revealing the energy consumption fully at future time periods (to be able to recharge the battery). We should also remark that the residential electricity consumption is forecast to increase significantly in the coming years [46], emphasizing the need to intelligently exploit limited capacity storage devices to hide energy consumption behavior. Moreover, as the typical electricity consumption is very spiky (see Figure 2.1 for an example), current batteries cannot fully hide the spikes in the consumption, because of their charging/discharging peak power constraints. We also would like to emphasize that the privacy leakage is caused mostly by these spikes, which are typically more informative (e.g., oven, microwave, heater) compared to more regular consumption (e.g., fridge). Moreover, due to electricity price variations users may prefer charging/discharging the battery in certain time periods, which limits the available energy that can be used for privacy. Finally, it is expected that the increasingly wider adoption of electrical vehicles and the mass production and adoption of energy-hungry “smart devices” will inevitably increase the typical household electricity consumption, limiting

further the rechargeable batteries' capability in fully hiding the user load.

Table 2.4 shows the average power generated via a solar panel, which is the most common residential RES. Locations, technology, as well as the inclination and size of the solar panel affect the generated power, as shown in Table 2.5 for one of the databases considered, where kWp denotes the kilowatt peak, i.e., the output power achieved by a panel under full solar radiation. As expected, around 50% of the time, i.e., at night, no energy is generated at all, while there are differences in the distribution of the average values for the two databases considered, due to the different locations. Comparing these values with those in Table 2.2, we note that the battery capacities are large enough to store many hours of average solar energy generated by the solar panels most of the time.

2.4 A Signal-Processing Perspective on Smart Meter Privacy

The generic discrete-time SM system model is depicted in Figure 2.4. In this discrete-time model, each time slot (TS) duration is normalized to unit time; therefore, power and energy values within a TS are used interchangeably. $X_t \in \mathcal{X}$ denotes the total power demanded by the appliances in the household in TS t , i.e., the user load, where $\mathcal{X}$ denotes the user load alphabet, i.e., the set of values that X_t can assume. X_t represents the user's private information, which needs to be protected. $Y_t \in \mathcal{Y}$ is the power received from the grid, i.e., the grid load, which is measured and reported to the UP through the SM, while $\mathcal{Y}$ denotes the grid load alphabet. We assume that the user and grid load power remain constant within a TS. This can be considered as a discrete-time linear approximation of a continuous load profile in practice. This approximation can be made arbitrarily accurate by reducing the TS duration.

In current systems, where no power manipulation is employed, $Y_t = X_t, \forall t$; that is, the actual power consumption of the appliances is reported to the UP. Instead, we assume that an RB and an RES are available to the user to physically distort the user's power consumption; so that what the user receives from the grid, Y_t , does not




FIGURE 2.4: Representation of the system model. X_t , Y_t , E_t and B_t denote the consumer's power demand, i.e., the user load, the SM readings, i.e., the grid load, the energy produced by the RES, and the battery state of charge at TS t , respectively. The dashed line represents the meter readings being reported to the UP.

reveal too much information about the power used by the appliances, X_t . We remark here that the TSs in our model correspond to time instants when the electricity is actually requested by the user and drawn from the grid, rather than the typically longer sampling interval used for sending SM measurements to the UP. In fact, we assume that the SM measures and records the output power values at each time slot; this is because our aim is to protect consumers' privacy not only from the UP, but also from the DSO or any other attacker that may deploy a sensor on the consumer's power line, recording the electricity consumption in almost real-time. The amount of energy stored in the RB at time t is $B_t \in [0, B_{\max}]$, where $B_{\max}$ denotes the maximum battery capacity. $X_t - Y_t$ denotes the power taken from the RB, and the battery charging and discharging processes are often constrained by the so-called charging and discharging power constraints $\hat{P}_c$ and $\hat{P}_d$, respectively, i.e., $-\hat{P}_c \leq X_t - Y_t \leq \hat{P}_d$, $\forall t$. There is typically a constraint on the average power that can be retrieved from an RB as well, imposed by an average power constraint $\bar{P}$, i.e., $\mathbb{E}[\frac{1}{n} \sum_{t=1}^n (X_t - Y_t)] \leq \bar{P}$. Losses

in the battery charging and discharging processes may also be taken into account to model a more realistic energy management system. The renewable energy generated at time t by the RES is denoted by $E_t \in \mathcal{E}$, where $\mathcal{E} = [0, E_{\max}]$. RBs and RESs are expensive facilities, and installation and operation costs can be reduced if they are shared by multiple users, e.g., users within the same neighborhood or block of flats. Moreover, sharing these resources allows a centralized management of the energy system for the RES and/or RB for multiple parties, which also leads to a more efficient use of the available resources. The renewable energy can be stored in the RB, or used immediately, so that a user can:

- *increase privacy*, by avoiding to report her actual power consumption to the UP;
- *decrease electricity costs*, by purchasing and storing electricity from the grid when it is cheaper, and use it to satisfy future demand, or even sell it back to the UP when the price increases;
- *increase energy efficiency*, by reducing the waste of generated renewable energy when it is not needed, and when it is not profitable to sell it to the UP.

We note that the system model represented in Figure 2.4 can be extended to consider multiple RBs or RESs. One example of this generalization is given in [47], where two rechargeable batteries, connected in series, are considered. However, we highlight that the optimal privacy-preserving solution for more complicated scenarios would be derived by extending the techniques developed for the system model of Figure 2.4. Hence, in the following we focus on characterizing the level of privacy achievable for the model depicted in Figure 2.4.

The random processes X and E are often modeled as a Markov process, or as a sequence of i.i.d. random variables. Although the UP does not know typically the instantaneous realizations of the stochastic processes X and E , it may well know their statistics. In some cases, the UP may know the realizations of the renewable energy process E , for example if it has access to additional information from sensors deployed near the household that measure different parameters, e.g., the solar or the wind power, and if it knows the specifications of the user's renewable energy generator, e.g., model and size of the solar panel.

Given the above definitions, the evolution of the level of energy in the RB is typically expressed via the following equation:

$$B_{t+1} = \min \left\{ B_t + E_t - (X_t - Y_t), B_{\max} \right\}, \quad (2.1)$$

where we remind that we can power and energy values interchangeably because each TS duration is normalized to unit time.

Sometimes the user load does not need to be satisfied immediately in its entirety. In fact, the user load can be further classified into demand that must be met immediately, e.g., lighting or cooking, and demand that can be satisfied at a later time, the so-called *elastic demand*, e.g., electric vehicle charging, dishwasher or clothes washer-dryer. For the latter demand, the user's only concern is that a certain task needs to be finished by a certain deadline, e.g., the electric car must be fully charged by 8 a.m., and it is not of interest at what exact time the consumption takes place. This flexibility allows the consumer to employ *demand response* to increase her privacy as well as to lower costs.

The electricity unit cost at time t , denoted by C_t , can be modeled as a random variable, or in accordance with a specific ToU tariff. The cost incurred by a user to purchase Y_t units of power over a time interval of τ_t at the price of C_t is thus given by $\tau_t Y_t C_t$. When the presence of an RES is considered, the prosumer may be able to sell part of the energy generated to the grid to further improve her privacy and to minimize the energy cost. If this occurs, the *net metering* approach is typically considered, i.e., the utilities purchase consumer-generated electricity at the current retail electricity rate. The battery wear and tear due to charging and discharging the RB can also be taken into account and modeled as an additional cost [48].

2.4.1 The Energy Management Policy

The energy management unit (EMU) is the intelligence of the system, located at the user's premises, where the SM privacy-preservation and cost-optimization algorithms are physically implemented. The energy management policy, implemented by

the EMU, determines at any time t the amount of power that should be drawn from the grid and the RB, given the previous values of the user load X^t , renewable energy E^t , level of energy in the battery B^t , and grid load Y^{t-1} , i.e.,

$$f_t : \mathcal{X}^t \times \mathcal{E}^t \times \mathcal{B}^t \times \mathcal{Y}^{t-1} \rightarrow \mathcal{Y}, \quad \forall t, \quad (2.2)$$

where $f \in \mathcal{F}$, and $\mathcal{F}$ denotes the set of feasible policies, i.e., policies that produce grid load values that satisfy the RB and RES constraints at any time, as well as the battery update equation in (2.1). The optimal policy is chosen to minimize the long-term information leakage about a consumer's electricity consumption, possibly along with other criteria, such as the minimization of electricity cost or wasted energy. The energy management policy prevents outages, and typically it is not allowed to draw more energy from the grid to be wasted simply for the sake of increased privacy.

The policy f_t in (2.2) corresponds to an *online energy management policy*, i.e., a policy in which the action taken by the EMU at any TS depends only on the information available causally right up to that time. Alternatively, in an *offline optimization* framework, the policy takes actions by knowing future information about the system state, i.e., user load and RES energy generation, in a non-causal fashion. In the SM privacy literature, both offline and online SM privacy-preserving algorithms have been considered. Online algorithms are more realistic and relevant for real-world applications; however, offline algorithms may lead to interesting intuitions or bounds on the performance. Moreover, non-causal knowledge of the electricity price process is a realistic assumption in today's energy networks; and even the non-causal knowledge of power consumption may be valid for certain appliances, such as refrigerators, boilers, heating and electric vehicles, whose energy consumption can be accurately predicted over certain finite time frames.

In the following, we will first present an overview of some heuristic solutions to the SM privacy problem. Next, we will describe the more rigorous and mathematically involved techniques, which are aimed at finding the relevant fundamental bounds and trade-offs.

2.5 A Heuristic Privacy Measure: Variations in the Grid Load Profile

As in many other problems involving privacy, a wide consensus over the best privacy measure for SMs has not been reached yet, and a number of privacy measures have been proposed in the literature, each one with its own benefits and limitations. Although it is clear that privacy is achieved when the UP cannot infer a user's behavior on the basis of SM measurements, it is challenging to define a corresponding mathematical measure that is independent of the particular detection technique employed by the attacker.

2.5.1 Grid Load Variance as a Privacy Measure

One can argue that privacy in SMs can be ensured by opportunely charging and discharging the RB so that the grid load is always constant. In fact, the differences in consecutive load measurements, $y_t - y_{t-1}$, are indicative of the appliances' switch-on/off events, the so-called *features*, and are typically exploited by the existing NILM algorithms [49]. Ideally, a completely "flat" grid load profile would not reveal any feature, and would only leak a user's long term average power consumption. However, this would require a very large battery capacity or a powerful RES. Alternatively, the level of privacy can be measured by the "distance" of the grid load from a completely flat *target load profile*, based on the intuition that the smaller the distance, the higher the level of privacy achieved [50]. Accordingly, privacy can be defined as the grid load variance around a prefixed target load profile W , i.e.,

$$\mathcal{V}_n \triangleq \frac{1}{n} \sum_{t=1}^n \mathbb{E}[(Y_t - W)^2], \quad (2.3)$$

where the expectation is over X_t and Y_t , and $W = \mathbb{E}[X]$ typically.

An important concern for consumers is to minimize their energy cost. With the integration of unreliable RESs into the grid, it is expected that the unit cost of energy from different UPs will fluctuate in time. Accordingly, RBs for residential use provide




FIGURE 2.5: Example of user load, grid load, and constant target load profiles, where the “distance” $Y_t - W_t$ is highlighted. The aim of the algorithms presented in this section is to minimize the average squared distance.

a significant flexibility for the consumers as they can buy and store energy during low-cost periods to be used during peak-price periods. The impact of RBs in reducing the cost of energy to the consumers have been extensively studied in the literature [51]. Note, however, that the operation of the EMU in order to minimize the energy cost does not necessarily align with the goal of minimizing privacy leakage. Therefore, it is essential to jointly optimize the electricity cost and user privacy. If the cost of energy and battery wear and tear are considered, the overall optimization problem becomes:

$$\min \frac{1}{n} \sum_{t=1}^n \mathbb{E} \left[C_t Y_t + \mathbb{1}_B(t) C_B + \alpha (Y_t - W)^2 \right], \quad (2.4)$$

where $\mathbb{1}_B(t) = 1$ if the battery is charging/discharging at time t , and 0 otherwise; C_B is the battery operating cost due to the battery wear and tear caused by charging and discharging the RB; and α strikes the trade-off between privacy and cost. The expectation in (2.4) is over the probability distributions of all the involved random variables, i.e., X_t , Y_t , and C_t .

If $W_t = \mathbb{E}[X]$, $\forall t$, the EMU tries to achieve a flat grid load profile around the average user power consumption with as little deviations as possible. This scenario is illustrated in Figure 2.5, where the straight blue line is the fixed target consumption

profile W_t , and the red line indicates the achieved grid load profile Y_t . For i.i.d. X and C processes, an online energy management policy can be obtained using Lyapunov optimization [48]. The online control algorithm can be formulated as a Lyapunov function with a perturbed weight and the *drift-plus-penalty* framework is adopted, which is typically used for stabilizing a queuing network, by minimizing the so-called *drift*, while minimizing at the same time a *penalty* function. Here, the penalty is represented by the optimization target, while the *Lyapunov drift* is defined as the difference of the level of energy in the RB at successive TSs. Authors in [48] show that this approach leads to a mixed-integer nonlinear program, which they solve by decomposing it into multiple cases, and by finding a closed-form solution to each of them.

This problem can also be studied in an offline framework, by assuming the future user demand profile can be accurately estimated for a certain time horizon, and the energy cost is known in advance. When privacy and cost of energy are jointly optimized over a certain time horizon, one can characterize the points on the Pareto boundary of the convex region formed by all the cost and privacy leakage pairs, by solving the following convex optimization problem [50]:

$$\min_{Y_t \geq 0} \sum_{t=1}^n \left[(1 - \alpha) Y_t C_t + \alpha (Y_t - W)^2 \right]. \quad (2.5)$$

It is shown in [50] that the optimal offline solution has a *water-filling* interpretation. However, differently from the classical water-filling algorithm, which appears as the solution of the power allocation problem across parallel Gaussian channels under a total power constraint, here the *water level* is not constant, and changes across time because of the instantaneous power constraints.

However, a completely flat consumption profile may not be feasible, or even desirable, for example if the cost varies greatly during the system operation due to ToU tariffs. Thus, it is reasonable to assume that a user requests more energy during the off-peak price periods as compared to the peak price [52]. An example of the application of this strategy is shown in Figure 2.6, applied to real power consumption data from the UK-Dale dataset [53]. This approach will be analyzed in detail in Chapters 5 and 6.




FIGURE 2.6: Examples of user load, grid load, and a piecewise constant target load profiles over the course of one day. The price periods are highlighted by arrows of different colors. Note that the target assumes a different constant value for each price period.

2.5.2 Temporal and Spatial Similarities in the Grid Load as a Privacy Measure

Variations in the grid load profile can be captured by considering power traces of single appliances and by computing differences in power consumption both in the time domain, i.e., consumption deviation over time of a specific appliance, and in the “space” domain, i.e., consumption profiles of different appliances. As these variations are computed over a certain time horizon, when an online algorithm is considered, future user electricity consumption is estimated by forecasting the future electricity prices and running Monte Carlo simulations. The optimal decision at any time is characterized by considering both the current inputs and the forecasts through a rolling online stochastic optimization process. *Load shifting*, i.e., the scheduling of the user’s flexible electricity demand in accordance with privacy as well as cost concerns, can also be considered. Load shifting is analyzed in [54] and in [55], where privacy, cost of energy, and battery wear and tear are jointly optimized, and an online algorithm is formulated. The objective is to minimize the sum of current and expected electricity

and charging/discharging costs together with the weighted power profile differences measured through the similarity parameters for an entire day. The effectiveness of three similarity measures are examined in [55] separately, as well as jointly, where only four typical appliances, an oven, a clothes dryer, a dishwasher and an electric vehicle, are considered for the sake of simplicity.

2.5.3 Heuristic Algorithms

While the grid load can be flattened by minimizing the variation of the grid load around a constant consumption target, several works in the literature propose heuristic battery charging and discharging algorithms, which keep the grid load variations limited. An intuitive approach is to try to keep the grid load equal to its most recent value by discharging (charging) the RB when the current user load is larger (smaller) than the previous one. This approach, called the *best-effort* (BE) *algorithm* in [56], tends to eliminate the higher-frequency components of the user load, while still revealing the lower-frequency components. In [56], the similarity between the two probability distributions of the user and grid load is quantified via the *empirical relative entropy* (*i.e.*, *Kullback-Leibler divergence*) [57]. In the same work, the authors also consider *cluster classification*, whereby data is clustered according to power levels, and *cross-correlation* and *regression* procedures, according to which the grid load is shifted in time at the point of maximum cross-correlation with the user load, and regression methods are then used to compare the two aligned signals.

A slightly more sophisticated approach is considered in [58], called the non-intrusive load levelling (NIL) algorithm, in which more than one grid load target value, namely a *steady state target* and *low* and *high recovery state targets* are allowed, and the EMU tries to maintain the grid load at one of these values across time. If the steady state load cannot be maintained, the EMU switches to a high (low) recovery state in case of persistent light (heavy) user demand. When one of the recovery states is reached, the target load is adapted accordingly to permit the battery to charge or discharge, similarly to the empirical strategies outlined in [59]. The value of the steady state target load can be updated whenever a recovery state is reached, in order to reduce

the occurrences of recovery states, which is achieved by using an exponential weighted moving average of the demand. To assess the performance of their proposed approach, the authors in [58] count the *number of features*, i.e., the number of times a device is recognized as being on or off, from the grid load, as compared to the user load, and they further consider the *empirical entropy*.

As also pointed out in [60], these heuristic algorithms suffer from precise load change recovery attacks that can identify peaks of user demand. We note that the NIL algorithm is essentially quantizing the user load to three values with the help of the RB. This idea is generalized in [60] by considering an arbitrary number of quantization levels. Since quantization is a “many-to-few mapping”, converting the grid load to a step function is inherently a non-linear and irreversible process, which can be used to provide privacy by maximizing the quantization error under battery limitations. More specifically, the grid load is forced to be a multiple of a quantity β , i.e., $y_t = h_t \beta$, where h_t is an integer value, and β is the largest value that satisfies battery’s maximum capacity and power constraints. At any TS, given the user load, the grid load is chosen between the two adjacent levels to the user load, namely $\left\lceil \frac{x_t}{\beta} \right\rceil$ and $\left\lfloor \frac{x_t}{\beta} \right\rfloor$, where $\left\lceil \cdot \right\rceil$ and $\left\lfloor \cdot \right\rfloor$ denote the ceiling and floor functions, respectively. Three *stepping algorithms* are proposed in [60], which have different quantization levels: 1) the *lazy stepping* algorithm, which tries to maintain the external load constant for as long as possible; 2) the *lazy charging* algorithm, which keeps charging (discharging) the battery until it is full (empty); and 3) the *random charging* algorithm that chooses its actions at random. While the simulation results show that these algorithms outperform the BE and NIL algorithms, with the lazy stepping algorithm typically performing the best, it is hard to make general claims due to the heuristic nature of these algorithms. In fact, these approaches do not provide theoretical guarantees on the level of privacy achieved; thus, they are not able to make any general claim about the strength of the proposed privacy-preserving approaches and their absolute performance. This is an important limitation as consumers would like to know the level of privacy they can achieve, even if it is in statistical terms. Also, because such heuristics are often based on deterministic schemes, they are prone to be easily reverse-engineered.

2.6 Theoretical Guarantees on Smart Meter Privacy

One of the challenges in SM privacy is to provide theoretical assurances and fundamental limits on the information leaked by an SM system, independent of any assumption on the capability of an attacker, or of the particular NILM algorithm employed. This is essential in privacy research as privacy-preserving techniques may perform extremely well against some NILM algorithms and very poorly against others. Moreover, the privacy assurances should not be based on the complexity limitations of a potential attacker, as techniques that are currently thought not to be feasible, may become available to attackers in the future, if computational capabilities improve, or if new methods are developed. Last but not least, establishing a coherent mathematical framework would allow us to compare various SM scenarios and the use of different physical resources, e.g., RBs of various capacities, RESs of various nature, etc., in a rigorous manner. Accordingly, signal processing and information-theoretic tools have been employed in the literature to provide theoretical privacy assurances. We will overview various statistical measures for privacy, in particular, the *conditional entropy* [59], *Fisher information* (FI) [61], or *type II error probability* in detecting user activity [62].

In this statistical framework, it is commonly assumed that the statistics of the user load and the RES are stationary over the period of interest, and are known to the EMU. This assumption is reasonable especially if the period of stationarity is sufficiently long for the EMU to observe and learn these statistics [63–65]. On the other hand, an online learning theoretic framework can also be considered to account for the convergence time of the learning algorithm. Alternatively, most of the works in the literature that carry out a theoretical analysis also propose suboptimal policies that can be applied on real power traces, thus allowing the reader to gain intuition about the practical application and performance of these theoretically-motivated techniques. We take a worst-case approach, and assume that the statistics governing the involved random processes are known also by the attacker. Note that this can only empower the attacker, and strengthen the stated privacy guarantees.

2.6.1 The Significance of Single-Letter Expressions

It is expected that a meaningful privacy measure should consider the leakage of a user's information over a certain time period of reasonable length, because of the memory effects introduced by the RB and the RES. The energy consumption over a short period of time can be easily covered by satisfying all the demand from the RB or the RES over this period, but this may come at the expense of revealing the energy consumption fully at future time periods. Therefore, the information-theoretic analysis typically considers an average information rate measured over a given finite time period, and often studies its infinite-horizon asymptotics as well. However, increasing the time horizon also increases the problem complexity, and one of the challenges of the information-theoretic analysis is to obtain a so-called “single-letter” expression for the optimal solution, which would reduce the problem complexity significantly, particularly when the involved random variables are defined over finite alphabets. Unfortunately, to date, closed-form or single-letter expressions for the information leaked in an SM system have been characterized only for specific settings under various simplifications, e.g., considering an i.i.d. or Markov user load or RES generation.

2.6.2 Markov Decision Process Formulation

In the online optimization framework, where the user load and the energy generated by the RES can be modeled as Markov processes (or, as an i.i.d. sequence as a special case), the SM privacy problem can be cast as a Markov decision process (MDP). An MDP is a discrete-time state-transition system, which is formally characterized by a *state space*, which includes the possible states for the SM system, an *action space*, which includes the possible actions that can be taken by the decision maker at each state, the *transition probabilities* from the current state to the next state, which describe the dynamics of the system, and the *reward* (or, inversely, the *cost*) process, which indicates the reward received (or, cost incurred) by the decision maker by taking a particular action in a particular state. The goal of an MDP is to find the optimal policy that minimizes the average (or, discounted) cost either by a specified time in the future, i.e., by considering the so-called *finite horizon* setting, or over an indefinite time period, by

considering the so-called *infinite horizon* setting. To solve the corresponding MDP the optimal Bellman optimality equations should be formulated [66], which are then solved to obtain the optimal policy at each state and time instant. The problem can be solved numerically for the finite horizon setting, while the value iteration algorithm can be employed to obtain the optimal stationary policy in the infinite horizon scenario.

In the SM problem, the *state* at any time t is typically represented by a combination of the current level of energy in the battery B_t , user demand X_t , and renewable energy E_t , since the values taken by these variables represent the current situation of the SM setting, in terms of both energy available for preserving privacy and power demanded by the user. The *action*, performed by the EMU, is represented by the current grid load and the power used from the RB and the RES. *State transitions* are modelled by the battery update equation, which is typically assumed to be deterministic, as well as transitions in the user demand and renewable energy generation states, which typically do not depend on user's actions. The *cost* function is the privacy loss that is experienced by the user when the SM system moves from one state to another by following a certain action. However, to consider the privacy as the cost function in an MDP, it is necessary to formulate the privacy leakage in an additive form across time, so that the total loss of privacy over multiple TSs is given by the summation of the privacy leakage at various TSs. This may be challenging depending on the privacy measure employed. For example, measuring privacy via the squared distance of the grid load from a constant target profile has a straightforward additive formulation, while the same does not hold when privacy is measured by the mutual information between the user and grid load sequences. This is because the mutual information takes into account the dependence between the realization of the user load at time t , X_t , and the current, past and future realizations of Y , $Y_1, \dots, Y_t, Y_{t+1} \dots$.

Moreover, when the state and action spaces are continuous, it is necessary to discretize them to solve the problem numerically. The accuracy of the numerical solution can be improved by decreasing the discretization step size, at the expense of significantly higher computational complexity. When the dimensions of the state and action spaces render numerical evaluation of the optimal policy unfeasible, one can resort to suboptimal solutions that are easier to optimize and compute numerically,

yet may provide near optimal performance or interesting intuitions. Also, when the information-theoretic privacy measures are used, it may be possible to simplify the infinite-horizon optimization problem, and write it in a single-letter form. We will provide further insights into this below.

2.6.3 Mutual Information as a Privacy Measure

The entropy of a random variable X , $H(X)$, is a measure of the uncertainty of its realization. The mutual information between random variables X and Y , $I(X; Y)$, measures the amount of information shared between the two random variables [57]. Mutual information can also be considered as a measure of dependence between the random variables X and Y , and it is equal to zero if and only if they are independent. Rewriting the mutual information as $I(X; Y) = H(X) - H(X|Y)$, where $H(X|Y)$ is the conditional entropy, we can also interpret mutual information as the average reduction in the uncertainty of X from the knowledge of Y . Therefore, we can measure the privacy leakage about the user load sequence X^n through the SM readings Y^n by the mutual information between the two sequences, $I(X^n; Y^n)$. This will measure the reduction in the uncertainty of the UP about the real power consumption of the appliances, X^n , after receiving the SM measurements, Y^n . Thus, minimizing the mutual information can be interpreted as a way of improving privacy for SM users. Mutual information as a measure of privacy leakage has also been considered in other domains, see for example [67–69].

For an SM system with only an RB (no RES) and a given energy management policy f in (2.2), running over n TSSs, the average *information leakage rate* $\mathcal{I}_f^n(B_{\max}, \hat{P}_d)$ is defined as

$$\mathcal{I}_f^n(B_{\max}, \hat{P}_d) \triangleq \frac{1}{n} I(X^n; Y^n) = \frac{1}{n} [H(X^n) - H(X^n|Y^n)], \quad (2.6)$$

where $0 \leq X_t - Y_t \leq \hat{P}_d$. The parameters $B_{\max}$ and $\hat{P}_d$ emphasize the dependence of the energy management policy, and therefore, of the achievable information leakage rate, on the battery capacity and the discharging peak power constraint. The optimal energy management policy and the corresponding minimum information leakage rate is obtained by minimizing (2.6) over all feasible policies $f \in \mathcal{F}$ to obtain $\mathcal{I}^n(B_{\max}, \hat{P}_d)$.

2.6.3.1 Privacy with a Renewable Energy Source

Alternatively, one can also consider the SM system of Figure 2.4 with an RES, but no RB. Assume that the renewable energy that can be used over the operation period is constrained by an average and a peak power constraint. We do not allow selling the generated renewable energy to the UP, as the goal is to understand the impact of the RES on providing privacy to the user. The minimum information leakage rate achieved under these assumptions and for an i.i.d. user load can be characterized by the so-called *privacy-power function* $\mathcal{I}(\bar{P}, \hat{P})$, and can be formulated in the following single-letter form:

$$\mathcal{I}(\bar{P}, \hat{P}) = \inf_{p_{Y|X} \in \mathcal{P}} I(X; Y), \quad (2.7)$$

where $\mathcal{P} \triangleq \{p_{Y|X} : y \in \mathcal{Y}, \mathbb{E}[(X - Y)] \leq \bar{P}, 0 \leq X - Y \leq \hat{P}\}$. This formulation is presented in [70] for a discrete user load alphabet, i.e., X can only assume values that are multiples of a fixed quantum, and in [71], for a continuous user load alphabet, i.e., X can assume any real value within the limits specified by the peak power constraints of the appliances. The optimal energy management policy that minimizes (2.7) is stochastic and memoryless, that is, the optimal grid load at each TS is generated randomly via the optimal conditional probability that minimizes (2.7) by only considering the current user load. Another interesting observation is that (2.7) is in a similar form to the well-known *rate-distortion function* in information theory, which characterizes the minimum compression rate R of data, in bits per sample, that is required for the receiver to reconstruct the source sequence within a specified average distortion level D [57]. Formally, the rate-distortion function $R(D)$ for an i.i.d. source $X \in \mathcal{X}$ with distribution p_X , reconstruction alphabet $\hat{\mathcal{X}}$, and distortion function $d(\hat{x}, x)$, where the average distortion between sequences X^n and $\hat{X}^n$ is given by $\frac{1}{n} \sum_{i=1}^n d(x_i, \hat{x}_i)$, characterizes the minimum rate with which an average distortion of D is achievable. The compression rate specifies the size of the codebook 2^{nR} , required to compress the source sequence of length n , X^n . Shannon showed that the rate-distortion function can be obtained in the following single-letter form:

$$R(D) = \min_{p_{\hat{X}|X} : \sum_{(x, \hat{x})} p_X p_{\hat{X}|X} d(x, \hat{x}) \leq D} I(\hat{X}; X). \quad (2.8)$$

The analogy between (2.7) and (2.8) becomes clear considering the following distortion measure

$$d(x, y) = \begin{cases} x - y, & \text{if } 0 \leq x - y \leq \hat{P}, \\ \infty, & \text{otherwise,} \end{cases} \quad (2.9)$$

and such analogy enables to use tools from rate-distortion theory to evaluate the privacy-power function for an SM system. However, it is important to highlight that despite the functional similarity, there are major conceptual differences between the two problems, namely: *i*) in the SM privacy problem Y^n is the direct output of the encoder rather than the reconstruction at the decoder side; *ii*) unlike the lossy source encoder, the EMU does not operate over blocks of user load realizations; instead, it operates symbol by symbol, acting instantaneously after receiving the appliance load at each TS. For discrete user load alphabets, the grid load alphabet can be constrained to the user load alphabet without loss of optimality [72], and since mutual information is a convex function of the conditional probability $p_{Y|X} \in \mathcal{P}$, the privacy-power function can be written as a convex optimization problem with linear constraints. Algorithms such as the Blahut-Arimoto (BA) algorithm can be used to numerically compute the optimal conditional distribution [57]. For continuous user load distributions, the Shannon lower bound is derived in [72], which is a computable lower bound to the rate-distortion function widely used in the literature, and is shown to be tight for exponential user load distributions.

These results can be generalized to a multi-user scenario in which N users, each equipped with a single SM, share the same RES [72]. This scenario is represented in Figure 2.7, where the objective is to minimize the total privacy loss of N consumers (or, devices) considered jointly, rather than minimizing the privacy loss for each of them separately. This requires the EMU to allocate the shared RES among all the users in the most effective manner. The average information leakage rate can still be written as in (2.6), by replacing X_t and Y_t with $\mathbf{X}_t = [X_{1,t}, \dots, X_{N,t}]$ and $\mathbf{Y}_t = [Y_{1,t}, \dots, Y_{N,t}]$, where the boldface characters denote the vectors representing the N power measurements. The privacy power function has the same expression in (2.7), and, for the case of independent, but not necessarily identically distributed user loads, the optimization




FIGURE 2.7: Representation of the system model where a single EMU and RES are shared among N users, each equipped with an SM. The EMU decides how much power each user can retrieve from the RES and from the SG.

problem (ignoring the peak power constraint) can be cast as

$$\mathcal{I}(\bar{P}) = \inf_{\sum_{i=1}^N P_i \leq \bar{P}} \sum_{i=1}^N \mathcal{I}_{X_i}(P_i), \quad (2.10)$$

where $\mathcal{I}_{X_i}(\cdot)$ denotes the privacy-power function for the i -th user having user load distribution $p_{X_i}(x_i)$. For continuous and exponential user loads, the optimal allocation of the energy generated by an RES can be obtained by the *reverse water-filling* algorithm, according to which only energy from the RES is used to satisfy the users with a low average load, while users with higher average load need to request energy from the grid as well.

2.6.3.2 Privacy with a Rechargeable Battery

We can also consider the presence of only an RB in the system, which is thus charged only via the grid (no RES is available to the EMU). Including an RB complicates the problem significantly, and the amount of energy stored in the battery, B_t , plays an important role when designing a feasible energy management policy.

This problem can be solved by putting it in the form of an MDP and by finding a suitable additive formulation for the privacy cost function [73]. The optimization

problem is formulated as

$$L^* = \min_f \frac{1}{n} I(B_1, X^n; Y^n), \quad (2.11)$$

where f can be any feasible policy, as specified in (2.2) (without including the renewable energy process). Eq. (2.11) has been cast in an additive formulation in [73] by noting that there is no loss of optimality in restricting the focus to charging strategies f' that decide on the grid load only based on the current values of the user load X_t and level of energy in the battery B_t , and on the past values of the grid load Y^{t-1} , i.e., the general strategy f in (2.2) is specialized as $f'_t : \mathcal{X} \times \mathcal{B} \times \mathcal{Y}^{t-1} \rightarrow \mathcal{Y}, \forall t$, because of the following inequality:

$$\frac{1}{n} I(X^n, B_1; Y^n) \geq \frac{1}{n} \sum_{t=1}^n I(X_t, B_t; Y_t | Y^{t-1}). \quad (2.12)$$

The conditional distributions in (2.12) grow exponentially with time because of the term Y^{t-1} , so that the problem becomes computationally infeasible very quickly. To overcome this problem, the knowledge of Y^{t-1} is summarized into a *belief* state, defined as $p(X_t, B_t | Y^{t-1})$, which can be computed recursively and interpreted as the belief that the UP has about (X_t, B_t) at time t , given its past observations, Y^{t-1} . This way, the optimal Bellman equations can be formulated, and the optimal policy can be identified numerically (with a discretization of the belief state).

For an i.i.d. user load, the single-letter characterization of the minimum information leakage rate is given by [73] as

$$J^* = \min_{\theta \in \mathcal{P}_B} I(B - X; X), \quad (2.13)$$

where θ is the probability distribution over B given the past grid load and actions, i.e., $\theta_t \triangleq p(b_t | y^{t-1}, a^{t-1})$, and the action a_t is defined as the transition probability from the current belief, user load and level of energy in the battery to the current grid load. This result is obtained by considering a belief on $W_t \triangleq B_t - X_t$, rather than (B_t, X_t) , and by further restricting to policies of the type $f''_t : \mathcal{W} \times \mathcal{Y}^{t-1} \rightarrow \mathcal{Y}, \forall t$. Since (2.13) is convex in θ , the optimal θ^* may be obtained by using the BA algorithm. The optimal grid load turns out to be i.i.d. and indistinguishable from the demand, while

the optimal policy is memoryless and the distribution of Y_t depends only on W_t . Such a characterization is provided in [74] for a binary i.i.d. user load, while the authors extend it to an i.i.d. user load of generic alphabet size in [73, 75, 76].

Apart from only maximizing privacy, it is of interest to also minimize the cost. Differently from privacy, cost of energy has an immediate additive formulation and can be easily incorporate into the MDP formulation. Considering the random price vector $C^t = (C_1, \dots, C_t)$, where C_t denotes the unit cost of energy at TS t , privacy can be defined in the long time horizon as

$$\mathcal{P} \triangleq \lim_{t \rightarrow \infty} \frac{H(X^t | Y^t, C^t)}{t}. \quad (2.14)$$

This formulation is presented in [59], where the corresponding MDP is formulated, and two suboptimal algorithms are proposed. The first is a *greedy algorithm*, which maximizes at any time the current instantaneous reward, while the second is a *battery centering approach* that is aimed at keeping the battery at a medium level of charge so that the EMU is less constrained by the battery or the demand in determining the grid load. In the latter approach, if the grid load depends not on the current user load or the battery level, but only on the current electricity price, the system is said to be in a *hidden state*, while it is said to be in a *revealing state*, otherwise. The latter strategy is analyzed for an i.i.d. user load by considering the system as a recurrent Markov chain and by adopting random walk theory.

2.6.3.3 Privacy with a Renewable Energy Source and a Rechargeable Battery

When both an RES and an RB are present, the information-theoretic privacy analysis becomes more challenging. As an initial step, we can consider infinite and zero battery capacities, which represent, respectively, lower and upper bounds on the privacy leakage achievable for a practical SM system with a finite-capacity battery [77, 78]. This scenario, along with a suboptimal energy management policy, will be analyzed in detail in Chapter 3.

The presence of a finite-capacity battery increases the problem complexity dramatically due to the memory effects induced by the finite battery, and single-letter expressions are still lacking for this scenario. A possible approach to find a theoretical solution to this problem is by extending the MDP formulation, as investigated in [79]. The latter approach will be analyzed in detail in Chapter 4.

2.6.4 Detection Error Probability as a Privacy Measure

So far we have considered approaches that try to hide the complete user energy demand from the UP. However, rather than hiding the entire energy consumption profile, in some cases it may be more meaningful to keep specific particular user activities, such as “is there anybody at home?”, “has the alarm been activated?” or “is the fridge working well?”. In order to keep the answer to such questions private, the goal of the EMU is to maximize the attacker’s probability of making errors when attempting to answer them.

Let the consumer’s behavior that needs to be kept private belong to a set of M possible activities. Thus, we can treat the attacker’s decision and the user’s action as an M -ary hypothesis, i.e., $H \in \mathcal{H} = \{h_0, h_1, \dots, h_{M-1}\}$. When $M = 2$, the hypothesis test is said to be binary and, by convention, the hypothesis h_0 , called the *null hypothesis*, represents the absence of some factor or condition, while the hypothesis h_1 , called the *alternative hypothesis*, is the complementary condition. For example, answering the question “is somebody at home?” corresponds to a binary hypothesis test, where h_0 is the hypothesis “somebody is not at home” and h_1 is the hypothesis “somebody is at home”. It is reasonable to assume that the user load will have different statistics under these two hypotheses; accordingly, we assume that under hypothesis h_0 (h_1), the power demand at TS t is i.i.d. with $p_{X|h_0}$ ($p_{X|h_1}$). Based on the SM readings, the attacker aims at determining the best *decision rule* $\hat{H}(\cdot)$, i.e., the optimal map between the SM readings and the underlying hypothesis. The space of all possible SM readings, $\mathcal{Y}^n$, is partitioned into the two disjoint decision regions $\mathcal{A}_0$ and $\mathcal{A}_1$, defined as follows:

$$\mathcal{A}_0 \triangleq \{y^n | \hat{H}(y^n) = h_0\}, \quad (2.15)$$

$$\mathcal{A}_1 \triangleq \{y^n | \hat{H}(y^n) = h_1\}, \quad (2.16)$$

which correspond to the SM readings subsets for which the attacker decides for one of the two hypotheses. The attacker's binary hypothesis test can incur two types of errors:

- Type I error probability: make a decision h_1 when h_0 is the true hypothesis (*false positive* or *false alarm*), i.e., $p_I = p_{Y^n|h_1}(\mathcal{A}_0)$;
- Type II error probability: make a decision h_0 when h_1 is the true hypothesis (*false negative* or *miss*), i.e., $p_{II} = p_{Y^n|h_0}(\mathcal{A}_1)$.

The Neyman-Pearson test minimizes the type II error probability for a fixed maximum type I error probability and makes decisions by thresholding the likelihood ratio $\frac{p_{Y^n|h_0}(y^n|h_0)}{p_{Y^n|h_1}(y^n|h_1)}$. Consider the worst case of an all-powerful attacker, which has the perfect knowledge of the energy management employed, in the asymptotic regime $n \rightarrow \infty$, and denote by $p_{II}^{\min}$ the minimal type II probability of error subject to a constraint on the type I probability of error. Assuming that a memoryless energy management is employed by the EMU, that is, the grid load at any TS t depends only on the user load at the same TS, then the attacker runs a Neyman-Pearson detection test on the grid load. Chernoff-Stein Lemma [57] links the minimal type II error probability $p_{II}^{\min}$ to the Kullback-Leibler (KL) divergence $D(\cdot||\cdot)$ between the grid load distributions conditioned on the two hypotheses in the limit of the number of observations going to infinity:

$$\lim_{n \rightarrow \infty} -\frac{\log p_{II}^{\min}}{n} = D(p_{Y|h_0}||p_{Y|h_1}), \quad (2.17)$$

while the KL divergence between two probability distribution functions on X , p_X and q_X , is defined as [57]

$$D(p_X||q_X) \triangleq \sum_{x \in \mathcal{X}} p_X(x) \log \frac{p_X(x)}{q_X(x)}. \quad (2.18)$$

Not surprisingly, to maximize the privacy the goal of the EMU is to find the optimal grid load distributions, which, given the user load X and the true hypothesis H , minimizes the KL divergence in (2.17), or equivalently, minimizes the asymptotic exponential decay rate of $p_{II}^{\min}$. However, the EMU is constrained by the available resources in making the two user load distributions produce similar grid load distributions. In

particular, we impose a constraint on the average RES it can use. Thus, the objective is to solve the following minimization problem:

$$\min_{p_{Y|H} \in \mathcal{P}_{Y|H}} D(p_{Y|h_0} || p_{Y|h_1}), \quad (2.19)$$

where $\mathcal{P}_{Y|H}$ is the set of feasible energy management policies, i.e., those that satisfy the average RES generation rate $\bar{P}$, so that $\frac{1}{n}\mathbb{E}[\sum_{i=1}^n X_i - Y_i | h_j] \leq \bar{P}$, $j = 0, 1$. This setting is studied in [80], where the asymptotic single-letter expressions of two privacy-preserving energy management policies in the worst case scenario are considered, i.e., when the probability of type I error is close to 1. The first policy is a memoryless hypothesis-aware policy that decides on Y_t based only on the current X_t and H , while the second policy is unaware of the correct hypothesis H but takes into account all the previous realizations of X and Y . It is noteworthy that even if the hypothesis-unaware policy with memory does not have access to the current hypothesis, it performs at least as well as the memoryless hypothesis-aware policy. This is because the hypothesis-unaware policy is able to learn the hypothesis with negligible error probability after observing the power demand process for a sufficiently long period. Additionally, the power supply alphabet can be constrained to the power demand alphabet without loss of optimality, which greatly simplifies the numerical solution to the problem.

2.6.5 Fisher Information as a Privacy Measure

FI is another statistical measure that can be employed as a measure of SM privacy [61]. Let some sample data x be drawn according to a distribution depending on an underlying parameter θ . Then, FI is a measure of the amount of information that x contains about θ . In the SM setting, Y^n is the sample data available to the attacker, while X^n is the parameter underlying the sample data that is to be estimated by the attacker. Let $\hat{X}^n$ denote the estimate of the attacker. The FI can be generalized to the multivariate case by the FI matrix, defined as

$$\mathcal{FI}(x^n) = \int_{y^n \in \mathcal{Y}^n} p(y^n | x^n) \left[\frac{\partial \log(p(y^n | x^n))}{\partial x^n} \right] \left[\frac{\partial \log(p(y^n | x^n))}{\partial x^n} \right]^T dy^n. \quad (2.20)$$

Assuming an *unbiased estimator* at the attacker, i.e., the difference between the estimator's expected value and the true average value of the parameter being estimated is zero, the variance of the estimation error can be bounded via the Cramér-Rao bound as follows:

$$\mathbb{E}[\|x^n - \hat{x}^n(y^n)\|_2^2] \geq \text{Tr}(\mathcal{FI}(x^n)^{-1}), \quad (2.21)$$

where $\|x^n - \hat{x}^n(y^n)\|_2^2$ denotes the squared Euclidean norm, and $\text{Tr}(A)$ denotes the trace of matrix A . To maximize the privacy it is then necessary to maximize the trace of the inverse of the FI matrix. In [61], two SM settings with RB are studied, specifically when the battery charging policy is independent of the user load, and when it is dependent non-causally on the entire user load sequence. For both cases single-letter expressions are obtained for the maximum privacy. Moreover, the case of biased estimators, wear and tear of the batteries, and peak power charging and discharging constraints are also briefly analyzed in [61].

2.6.6 Empirical Mutual Information as a Privacy Measure

Approaches aimed at determining theoretical privacy limits provide important insights and intuitions for the optimal energy management policy in order to limit the privacy leakage. However, they are often difficult to optimize or even evaluate numerically, and the relatively simplified formulation obtained in various special cases rely on restrictive assumptions, e.g., i.i.d. user load, infinite RB capacity, etc. An alternative is to follow a suboptimal or heuristic energy management policy. Although such a policy does not provide theoretical privacy guarantees, one can evaluate the corresponding privacy leakage numerically using empirical mutual information.

One way to compute the empirical mutual information is by simulating a discrete-time system for a “large enough” time interval and sampling the resulting X^n and Y^n sequences [81]. The mutual information between two sequences x^n and y^n can be approximated as

$$I(X; Y) \approx -\frac{1}{n} \log p(y^n) - \frac{1}{n} \log p(x^n) + \frac{1}{n} \log p(x^n, y^n), \quad (2.22)$$

where $p(y^n)$, $p(x^n)$ and $p(x^n, y^n)$ are calculated recursively through a sum-product computation. When using this method, the RB is modeled as a finite state machine, and the level of energy in the battery evolves in time through a Markov chain with transition probabilities depending on the specific policy implemented. The support space for the parameters is discretized, and the optimal combination of parameters is found, which minimizes the empirical mutual information. This approach is followed in [82], where only a binary RB is present and for an i.i.d. Bernoulli distributed user demand, and in [83] where an RES is also considered. The latter work also analyzes the wasted energy and characterize the privacy-energy efficiency trade-off for the binary scenario and equiprobable user load and renewable energy generation processes. For larger battery capacities and for an equiprobable user load, the authors note that there is a symmetry and complementarity in the optimal transition probabilities in the finite state machine model, which simplifies the numerical analysis. This model is also employed in [78] by considering an RES and designing a suboptimal policy, which, at each TS, decides among using all of the available energy, half of it, or no energy at all, according to a probability chosen to minimize the overall information leakage. Further details about the latter approach can be found in Chapter 3.

Another technique for approximating mutual information is to assume X and Y to be i.i.d. over a time interval, and approximate the mutual information via the relative frequency of events (X_t, Y_t) during the same time window. In [84] this approach is enriched by additive smoothing, i.e., avoiding zero probability estimates by adding a positive scalar, and it is employed together with a *model-distribution predictive controller*, such that, at each TS t , the EMU decides its actions for a prediction horizon of length T , i.e., up to time $t + T$. Privacy and cost are jointly optimized by considering non-causal knowledge of the renewable energy generation process, user load and energy cost, while EMU's actions, i.e., the power that is requested from the grid and the battery, are forecast over the prediction horizon. The user and grid load processes are assumed to be i.i.d. within a time window $N \gg T$, which also includes the prediction horizon T , and finite alphabets $\mathcal{X}$ and $\mathcal{Y}$ are considered. As $N \gg T$, first-order Taylor approximation of the logarithm function is used, and the corresponding mixed integer quadratic program is formulated, which is of manageable size and can be

solved recursively whenever new SM readings are available. Results show that considering a relatively small prediction horizon T prevents the EMU from fully utilizing the RB capacity, as the user load that is considered by the algorithm is generally smaller than the RB capacity. Allowing a longer prediction interval dramatically improves the performance in terms of both privacy and cost, at the expense of a much higher computational complexity. The work also shows that by increasing the alphabet sizes of $\mathcal{X}$ and $\mathcal{Y}$ a better privacy performance can be achieved.

Empirical mutual information normalized by the empirical entropy of the user load is considered in [85], where an RB is used to minimize the energy cost subject to privacy constraints. Here two cost tariffs are considered, a low-price and a high-price, and a dynamic programming approach is developed to maximize the energy stored in the battery at the end of the low-price period, and minimize it at the end of the high-price period. At every TS, the optimal probability distribution of the grid load is computed, which is forced to be independent of the user load distribution.

2.7 Conclusions

Privacy in SMs is a very debated topic and many solutions have been proposed so far, which differ in many aspects, most notably in the actual definition of privacy. In this chapter we have presented an overview of privacy-preserving approaches within the UDS family, that is, which make use of physical resources to achieve privacy for consumers. These techniques all share the significant advantage of providing the UP and the DSO with accurate SM data, without adding any modification on it. Thus, these techniques provide privacy to consumers without undermining the usefulness of the SG concept itself. Nonetheless, these techniques are still very different in the way they define privacy, and in the mathematical tools they employ to solve the privacy problem. For all of them, we have presented the high-level system model and the main constraints, and we have explained the most significant results they achieve. Some of the techniques mentioned in this chapter will be analyzed more in depth in the following chapters.

Chapter 3

Fundamental Limits of Smart Meter Privacy

3.1 Overview

In this chapter we study SM privacy by considering the presence of an RES and an RB at the user's premises. Privacy is measured by the information leakage rate, which measures the average mutual information between the user load and the grid load processes. The loss of privacy arising from the fact that the UP knows the amount of power generated by the RES is also characterized. The minimum information leakage rate is characterized as a computable information-theoretic single-letter expression in the two limit scenarios, that is, when the battery capacity is either infinite or zero. Numerical results are presented for the finite battery capacity case to illustrate the potential privacy gains from the existence of an RB. It is shown that, while the information leakage rate decreases with increasing availability of an RES, a larger storage capacity is needed to fully exploit the available energy to improve the privacy.

3.2 Introduction

In this chapter we draw from Section 2.6.3, and study privacy from an information-theoretic point of view. For some scenarios, we provide closed-form single-letter expressions for the best privacy performance achievable. A similar model, studied in [72], imposes only average and peak power constraints on the RES, which can be a micro-grid, capable of providing any amount of energy at each time instant. However, the energy produced by an RES at each time instant is typically random, and its statistics

depend on the energy source (e.g., solar, wind) and the energy generator specifications. In addition, the finite-capacity battery imposes further limitations on the available energy. Thus, in this chapter we study the minimum amount of user's power consumption information leaked to the UP by taking into account instantaneous power constraints. We study the two limit scenarios of zero and infinite battery capacity with a discrete-alphabet user load, as well as the more practical scenario with a finite-capacity storage device, and the case of a continuous-alphabet user load.

Following up on [82], [83] and [72], we model user's power consumption profile as a randomly generated time series whose statistics are known by the UP, and measure the user's information leakage by the average mutual information between the user and grid load vectors, i.e., between the real power consumption profile of the appliances and the SM readings, which is called the information leakage rate and defined in (2.6). More details about the use of mutual information as a privacy measure can be found in Section 2.6.3.

The main contributions of this chapter can be summarized as follows:

- We provide computable closed-form single-letter expressions for the minimum information leakage rate when the battery capacity is either zero or infinite, providing detailed proofs for these results. These two asymptotic limit scenarios can be considered as upper and lower bounds, respectively, on the achievable privacy performance for a more practical SM system with a finite-capacity battery.
- When $B_{\max} = \infty$, the problem can be shown to be equivalent to the average and peak power-constrained scenario, and, interestingly, the privacy performance does not deteriorate even if the UP knows the exact amount of renewable energy generated. This shows that, keeping the renewable energy generation process private is more critical when the RB has a limited capacity. Two distinct energy management policies are shown to achieve this lower bound.
- When $B_{\max} = 0$, the renewable energy that can be used at any TS is limited by the amount of energy generated within that TS, and, as expected, we show that assuming the knowledge of the renewable energy process at the UP significantly degrades the privacy performance in this scenario.

- For a finite-capacity battery scenario, we propose a suboptimal parameterized energy management policy, and optimize the policy parameters using a policy search technique that exploits stochastic gradient descent. We show numerically that the performance of the proposed energy management policy approaches the one with an infinite battery even with a relatively small battery size. This shows the efficacy of the proposed privacy preservation scheme.
- We show that the information leakage rate decreases with the rate of the available RES, and that a larger RB is needed to fully exploit the available energy to improve the privacy.

The remainder of this chapter is organized as follows. In Section 3.3 the system model is introduced. In Section 3.4 the limit scenario with an infinite-capacity battery is studied, while in Section 3.5 another limit scenario with no energy storage is considered. For both scenarios, we also study the case in which the UP knows the realizations of the renewable energy process. In Section 3.6 we study the binary scenario, while in Section 3.7 we propose achievable schemes for the generic finite battery capacity scenario, and present the corresponding numerical results. In Section 3.8 a continuous user load is considered, while conclusions are drawn in Section 3.9.

3.3 System Model

The system model considered in this chapter shares many similarities with the general scenario introduced in the previous chapter. However, we recall the most important quantities and assumptions for an easier accessibility to the reader. We adopt the discrete-time system model of Figure 3.1. $X_t \in \mathcal{X}$ is the total amount of power demanded by a user in TS t , i.e., the user load, where $\mathcal{X} = [0, \dots, X_{\max}]$, while $Y_t \in \mathcal{Y}$ is the power received from the UP in TS t , i.e., the grid load, where $\mathcal{Y} = [0, \dots, Y_{\max}]$. In this model, each TS duration is normalized to unit time; therefore, power and energy values within a TS are used interchangeably and there is no need to explicitly mention the TS duration in the following equations. For simplicity, we assume that the entries of the user load sequence $\{X_t\}_{t=1}^{\infty}$ are i.i.d. with distribution p_X . In TS t , $E_t \in \mathcal{E}$




FIGURE 3.1: Representation of the system model. X_t , Y_t , E_t and B_t denote the consumer's power demand, the SM readings, the energy produced by the RES, and the level of energy in the battery at time t , respectively. The dashed line represents the meter readings being reported to the UP.

units of energy are generated by the RES, which becomes available to the EMU at the beginning of TS t . The entries of the renewable energy sequence $\{E_t\}_{t=1}^{\infty}$ are also i.i.d. with distribution p_E and alphabet $\mathcal{E} = [0, \dots, E_{\max}]$, while the average renewable energy rate is denoted by $\bar{P}_E \triangleq \mathbb{E}[E]$. We further consider the presence of an RB in which the renewable energy can be stored for future use. The level of energy in the battery at time t is $B_t \in [0, \dots, B_{\max}]$, and its capacity is $B_{\max}$. We assume no losses in the battery charging and discharging processes.

The EMU always satisfies user's power demands by drawing power from either the UP or the RB; that is, outages or demand shifting are not allowed. As a consequence, we have $X_{\max} \geq Y_{\max} \geq X_{\max} - B_{\max}$. We do not allow extra power to be drawn from the grid and then wasted. This could provide additional privacy, albeit at a significantly higher energy cost. Also, the battery is exclusively for storing the generated renewable energy, and it cannot be recharged with grid energy. While storing grid energy in the battery to be supplied later to the appliances can provide additional privacy [82], here we limit the use of the battery to renewable energy storage to isolate and understand the privacy benefits of RESs. Hence, we impose

$$0 \leq Y_t \leq X_t, \quad \forall t, \quad (3.1)$$

while $X_t - Y_t$ is the amount of power obtained from the RB in TS t . The energy retrieved from the battery must be smaller than the energy available in it, i.e.,

$$X_t - Y_t \leq B_t + E_t, \quad \forall t, \quad (3.2)$$

where we remind that we can compare energy with power because of the assumption of having unit time TSs.

We also consider a peak power discharging constraint $\hat{P}_d$ on the amount of power that can be requested at any time from the RB, i.e.,

$$0 \leq X_t - Y_t \leq \hat{P}_d, \quad \forall t, \quad (3.3)$$

and for the rest of the chapter we assume that $\bar{P}_E \leq \hat{P}_d$.

Given $(X_t, E_t, B_t) = (x_t, e_t, b_t)$ and the constraints (3.1), (3.2), and (3.3), the set of feasible power requests at time t is

$$\bar{\mathcal{Y}}(x_t, e_t, b_t) \triangleq \left\{ y_t \in \mathcal{Y} : [x_t - \min\{b_t + e_t, \hat{P}_d\}]^+ \leq y_t \leq x_t \right\}, \quad (3.4)$$

where $[a]^+ = a$ if $a > 0$, and 0 otherwise.

The battery update equation can be written as in the following

$$B_{t+1} = \min \left\{ B_t + E_t - (X_t - Y_t), B_{\max} \right\}, \quad \forall t. \quad (3.5)$$

We aim at designing *energy management policies* $f = (f_1, f_2, \dots)$ that decide on the amount of power to request from the UP at each TS t , given the previous values of user load X^t , renewable energy E^t , level of energy in the battery B^t , and grid load Y^{t-1} :

$$f_t : \mathcal{X}^t \times \mathcal{E}^t \times \mathcal{B}^t \times \mathcal{Y}^{t-1} \rightarrow \mathcal{Y}, \quad \forall t, \quad (3.6)$$

where $f \in \mathcal{F}$, and $\mathcal{F}$ denotes the set of feasible policies, i.e., policies that produce grid load values that satisfy the constraint given in (3.4) at any time, as well as the battery update equation in (3.5).

We measure privacy via the information leakage rate, i.e.,

$$\mathcal{I}_f^i(B_{\max}, \hat{P}_d) \triangleq \limsup_{n \rightarrow \infty} \frac{1}{n} I(X^n; Y^n), \quad (3.7)$$

where the subscript f denotes the specific energy management policy employed, and the superscript i stresses the fact that we are considering instantaneous power constraints. Eq. (3.7) characterizes the amount of information about the user's actual consumption that is leaked through the SM measurements. We note that in Eq. (3.7) we use the $\limsup$ since the limit of $I(X^n; Y^n)$ as $n \rightarrow \infty$ may not necessarily exist, and because we want to consider a worst-case scenario for evaluating the privacy loss. Thus, the optimization problem can be written as the minimization of (3.7) over all feasible policies $f \in \mathcal{F}$, i.e.,

$$\mathcal{I}^i(B_{\max}, \hat{P}_d) \triangleq \inf_{f \in \mathcal{F}} \limsup_{n \rightarrow \infty} \frac{1}{n} I(X^n; Y^n). \quad (3.8)$$

A single-letter expression for the information leakage rate is provided in [70–72] when the EMU is constrained only by the average and peak power constraints. In general, because of the memory effects introduced by the RB and the RES, satisfying the user load from the RB or the RES at some time period may come at the expense of revealing more information about the energy consumption at future time periods. For this reason, the information-theoretic analysis typically focuses on the average performance, measured over a period of n TSs, and aims at understanding the fundamental performance bounds by letting this time period go to infinity, i.e., $n \rightarrow \infty$, as in (3.7). However, the definition of the information leakage rate in (3.7) involves n -length sequences X^n and Y^n , and the asymptotic performance limit corresponds to an infinite-dimensional optimization problem, which cannot be solved numerically. On the contrary, characterizing a single-letter expression allows the optimal solution to be described as an optimization problem in terms of the single-letter random variables, which can be a finite-dimensional optimization problem when the involved random variables are defined over finite alphabets. Therefore, a single-letter characterization of the information-theoretic privacy is desirable to be able to evaluate the minimum information leakage rate.

In [71] the *privacy-power function* $\mathcal{I}(\bar{P}, \hat{P}_d)$ is defined as the minimum information leakage rate that can be achieved when the energy management policy satisfies the average power constraint $\mathbb{E}[\frac{1}{n} \sum_{t=1}^n (X_t - Y_t)] \leq \bar{P}$, as well as the peak power constraint $0 \leq X_t - Y_t \leq \hat{P}_d, \forall t$, on the amount of energy provided by an RES. The privacy-power function has the single-letter characterization provided by the following theorem.

Theorem 3.1. [71, Theorem 1] *The privacy-power function $\mathcal{I}(\bar{P}, \hat{P}_d)$ for an i.i.d. user load vector X with distribution $p_X(x)$ and grid load vector Y , when the average and peak values of the power provided by the RB are limited by $\bar{P}$ and $\hat{P}_d$, respectively, is given by*

$$\mathcal{I}(\bar{P}, \hat{P}_d) \triangleq \inf_{p_{Y|X} \in \mathcal{P}} I(X; Y), \quad (3.9)$$

where $\mathcal{P} \triangleq \{p_{Y|X} : y \in \mathcal{Y}, \mathbb{E}[(X - Y)] \leq \bar{P}, 0 \leq X - Y \leq \hat{P}_d\}$.

From Theorem 3.1 we can see that the optimal energy management policy is memoryless, i.e., the optimal grid load is produced through the optimal conditional probability distribution by simply considering the current input load.

Lemma 3.1. [71, Lemma 1] *The privacy-power function $\mathcal{I}(\bar{P}, \hat{P}_d)$, given above, is a non-increasing convex function of $\bar{P}$ for a given $\hat{P}_d$.*

Authors in [72] show that, when the user load alphabet is discrete, i.e., $\mathcal{X} = \{0, 1, \dots, X_{\max}\}$, the grid load alphabet $\mathcal{Y}$, which is not necessarily discrete, can be restricted to the user load alphabet, i.e., $\mathcal{Y} = \mathcal{X}$, without loss of optimality. Given this restriction and the convexity of the privacy-power function, $\mathcal{I}(\bar{P}, \hat{P}_d)$ can be numerically evaluated, e.g., by the efficient BA algorithm [86]. The following lemma states that this property holds also in our setting for the various battery capacities we analyze in the following. Thus, in the discrete case, we can assume that all the involved random processes are defined over finite alphabets and that there is a minimum quantum of power such that all the aforementioned quantities are integer multiples of this quantum.

Lemma 3.2. *If the user load alphabet $\mathcal{X}$ is discrete, the grid load alphabet $\mathcal{Y}$ can be constrained to the user load alphabet without loss of optimality.*

Proof. The proof is similar to that of [72, Theorem 2]. Let $\mathcal{X}$ be the discrete user load alphabet and let $X(y) = \min_{x \in \mathcal{X}} \{x \geq y\}$. Then, for any given energy management policy, and the resultant grid load Y^n , we define a new grid load as $\hat{Y}(t) = X(Y(t))$, that is, $\hat{Y}$ is a post-processed version of Y , and $\hat{\mathcal{Y}} = \mathcal{X}$. By construction, we have that $X(t) \geq \hat{Y}(t) \geq Y(t), \forall t$, i.e., the power demanded by the EMU from the RB cannot have a larger peak value than the original demanded power. Similarly, the new grid load satisfies all the instantaneous power constraints as well. This proves that the policy is feasible. Also, the information leakage rate is not increased as $\hat{Y}$ is a deterministic function of Y , and thus $X \rightarrow Y \rightarrow \hat{Y}$ forms a Markov chain, and $I(X, Y) \geq I(X, \hat{Y})$ by the data processing inequality [57]. $\square$

Here we introduce a generic energy management policy, which we later specialize to the different scenarios we consider. This is a stationary and memoryless policy that generates Y_t randomly using a conditional probability distribution that is based only on the current user load X_t and the available total renewable energy $B_t + E_t$, i.e.,

$$\tilde{p}_{Y|X, B+E} : \mathcal{X} \times (\mathcal{B} + \mathcal{E}) \rightarrow \mathcal{Y}. \quad (3.10)$$

Note that, in the presence of an RB, in which the generated renewable energy is stored and used for privacy, a memoryless energy management policy is suboptimal in general, as it ignores the history. However, in the following we show that a memoryless policy is able to achieve the minimum information leakage rate in the two limit scenarios of $B_{\max} = \infty$ and $B_{\max} = 0$.

3.4 Infinite Battery Capacity

In this section we relax the constraint on the battery capacity and consider $B_{\max} = \infty$. This is an extreme situation that may model a battery with a relatively large capacity compared to the average generation rate of renewable energy, $\bar{P}_E$, and the average user load. This scenario provides useful insights on the best achievable privacy performance, and also serves as a bound on the performance achievable with a finite-capacity RB.

In each TS, the EMU is limited by both the peak power constraint (3.3) and the energy available in the RB, which is the difference between the total renewable energy generated and the total energy that has been requested from the battery up to that TS, i.e.,

$$\sum_{t=1}^n (X_t - Y_t) \leq \sum_{t=1}^n E_t, \quad \forall n. \quad (3.11)$$

3.4.1 Generated Renewable Energy not Known by the Utility Provider

In this section E^n is treated as a random sequence whose realization is known only to the consumer in a causal manner. This scenario may occur if the renewable energy originates from sources which could be extremely difficult, if not impossible, for the UP to track.

The following Theorem 3.2 states that the minimum information leakage rate when $B_{\max} = \infty$ is equivalent to the average and peak power-constrained scenario, as in [71]; that is, the cumulative constraints on the EMU policy do not reduce the achievable privacy if the battery capacity is sufficiently large.

Theorem 3.2. *If $B_{\max} = \infty$ and the peak power constraint on the amount of power taken from the RB is $\hat{P}_d$, then the minimum information leakage rate for an i.i.d. user load X and an i.i.d. renewable energy generation process with average power $\bar{P}_E$ is*

$$\mathcal{I}^i(\infty, \hat{P}_d) = \mathcal{I}(\bar{P}_E, \hat{P}_d). \quad (3.12)$$

$\mathcal{I}(\bar{P}_E, \hat{P}_d)$ is a trivial lower bound on $\mathcal{I}^i(\infty, \hat{P}_d)$, as the former quantity is not constrained by any battery. In the following section an energy management policy that achieves $\mathcal{I}^i(\infty, \hat{P}_d)$ is presented. The proposed policy is a specialization of the generalized memoryless policy introduced in (3.10).

3.4.2 Optimal Energy Management Policy for $B_{\max} = \infty$

Consider the following energy management policy. In each TS t , the EMU, based on the instantaneous user load X_t , decides on the optimal portion of the user load to be

Algorithm 1 Best-Effort Privacy Policy for $B_{\max} = \infty$.

```

1: Initial level of energy in RB:  $B_0$ .
2: Find  $p_{Y|X}^*$  that minimizes (3.9) for given  $\bar{P}_E$  and  $\hat{P}_d$ .
3: for  $t = 1, \dots, n$  do
4:   Input:  $X_t, B_t, E_t$ .
5:   Generate  $Y_t^*$  according to  $p_{Y|X}^*$ .
6:   if  $B_t + E_t \geq X_t - Y_t^*$  then
7:     Optimal policy is followed:  $Y_t = Y_t^*$  and  $X_t - Y_t^*$  taken from the battery.
8:   else
9:     Full leakage occurs:  $Y_t = X_t$ .
10:  end if
11:  Next battery state:  $B_{t+1} = \min\{B_t + E_t - (X_t - Y_t), B_{\max}\}$ .
12: end for

```

received from the grid, Y_t^* , by using the optimal conditional probability distribution $p_{Y|X}^*$ that minimizes (3.9). If there is enough energy available to fully satisfy the EMU requests, i.e., $B_t + E_t \geq X_t - Y_t^*$, the EMU uses $X_t - Y_t^*$ units of renewable energy and Y_t^* units of power from the grid, i.e., $Y_t = Y_t^*$; otherwise, all the user load is satisfied directly from the grid, i.e., $Y_t = X_t$, thus leading to the maximum information leakage for that TS, that is, the UP learns X_t perfectly. The TSs at which such leakage occurs cannot be determined beforehand, since they depend on the realizations of the renewable energy process, user and grid loads. Given the nature of this policy, which tries to follow the optimal policy generated by ignoring the current amount of energy in RB, we name it the *best-effort energy management policy*. Algorithm 1 summarizes this policy.

The following equation specializes policy (3.10) to the best-effort policy:

$$\begin{aligned}
& \tilde{p}_{Y|X, B+E}(y|x, b+e) \\
&= \begin{cases} p_{Y|X}^*(y|x), & \text{if } x - y^* \leq b + e \text{ and } y^* \neq x, \\ p_{Y|X}^*(y|x) + \sum_{\{y' \in \mathcal{Y}: x - y' > b + e\}} p_{Y|X}^*(y'|x), & \text{if } y^* = x, \\ 0, & \text{if } x - y^* > b + e. \end{cases} \quad (3.13)
\end{aligned}$$

The second case in (3.13) includes all the instances for which $p_{Y|X}^*$ outputs either $y^* = x$ (first term), or an infeasible output, i.e, for which $x - y^* > b + e$ (second term). In the latter case, all of the user load is satisfied from the grid load, i.e., $y = x$.

Since the energy arrival is stochastic, it may seem that very little can be said about the information leakage rate. However, if the condition $\mathbb{E}[X - Y^*] < \bar{P}_E$ holds, then it is possible to show that the number of TSs full leakage of information occurs due to unavailability of energy is relatively small compared to the operating time of the system n . This is proved in the following lemma.

Lemma 3.3. *As $n \rightarrow \infty$, $\mathbb{E}[X - Y^*] < \bar{P}_E$, and the EMU follows the best-effort energy management policy, then almost surely the condition $B_t + E_t < X_t - Y_t^*$ holds with probability zero for all large enough n .*

Proof. We remind that by Y_t^* we denote the solution given by the optimal privacy-preserving policy $p_{Y|X}^*(y|x)$, which is independent of B_t and E_t , whereas by Y_t we denote the actual value of power demanded from the grid, which depends on $B_t + E_t$ as well as X_t . We note that the sequence $E_t - (X_t - Y_t^*)$ is i.i.d, since E_t and X_t are i.i.d., and Y_t^* depends on X_t through the fixed memoryless policy $p_{Y|X}^*(y|x)$.

Let $\mathbb{E}[X - Y^*] = \bar{P}_E - \epsilon$, for some $\epsilon > 0$. The sequence $E_t - (X_t - Y_t^*) - \epsilon$ has zero mean. By the strong law of large numbers, the sample average of the sequence converges almost surely to its expected value, i.e., the sequence of events $\{\frac{1}{n} \sum_{t=1}^n [E_t - (X_t - Y_t^*) - \epsilon] < -\epsilon\}_{n=1}^\infty$, and thus the event $\{\frac{1}{n} \sum_{t=1}^n [E_t - (X_t - Y_t^*)] < 0\}_{n=1}^\infty$ stops occurring for all large enough n . Multiplying both sides of the inequality by n and adding the level of energy at the beginning, B_1 , which is a nonnegative quantity, the previous condition implies that, with Y_t^* generated according to the best-effort policy, the unavailability of energy $\{B_1 + \sum_{t=1}^n E_t < \sum_{t=1}^n X_t - Y_t^*\}_{n=1}^\infty$, which can be rewritten as $\{B_n + E_n < X_n - Y_n^*\}_{n=1}^\infty$, stops occurring for any large enough time.

□

Lemma 3.4. *If $\mathbb{E}[X - Y^*] < \bar{P}_E$, then the minimum information leakage rate of the best-effort policy tends to $\mathcal{I}^i(\infty, \hat{P}_d)$, as $n \rightarrow \infty$.*

Proof. Divide the sequence of user and grid loads according to the TSs in which a private SM operation is achieved, i.e., the TSs the EMU can fully emulate $p_{Y|X}^*$, and the TSs in which full leakage occurs. From Lemma 3.3 we know that as $n \rightarrow \infty$, the probability of not achieving the level of privacy induced by $p_{Y|X}^*$ is 0, i.e., the number

of times, say m , for which the condition $B_t + E_t < X_t - Y_t^*$ holds is $m = o(n)$. This is because after a large enough n , there is always enough energy in the battery to satisfy user needs. Then, we can write the following:

$$\frac{1}{n}I(X^n; Y^n) = \frac{1}{n} \left[H(X^n) - H(X^n | Y^n) \right] \quad (3.14a)$$

$$= \frac{1}{n} \sum_{t=1}^n \left[H(X_t) - H(X_t | X^{t-1}, Y^n) \right] \quad (3.14b)$$

$$\geq \frac{1}{n} \sum_{t=1}^n \left[H(X_t) - H(X_t | Y_t) \right] \quad (3.14c)$$

$$= \frac{1}{n} \left[\sum_{t \in \mathcal{T}^C} I(X_t; Y_t = Y_t^*) + \sum_{t \in \mathcal{T}} I(X_t; Y_t = X_t) \right] \quad (3.14d)$$

$$\geq \frac{n-m}{n} \mathcal{I}^i(\infty, \hat{P}_d) + \frac{m}{n} H(X) \quad (3.14e)$$

$$\xrightarrow{n \rightarrow \infty} \mathcal{I}^i(\infty, \hat{P}_d), \quad (3.14f)$$

where $\mathcal{T}$ is the set of instants when full leakage of information takes place, i.e., for which $Y_t = X_t$, and $\mathcal{T}^C$ is the set of TSs in which the output is generated through $p_{Y|X}^*$, i.e., $Y_t = Y_t^*$; (3.14c) follows since conditioning reduces entropy; (3.14f) follows since $m = o(n)$. $\square$

3.4.3 Store-and-Hide Energy Management Policy

Here we provide an alternative energy management policy in the case of an infinite-capacity battery. The *store-and-hide energy management policy* consists of an initial *storage* phase, during which all the energy requests of the user are satisfied from the grid while all the generated renewable energy is stored in the battery, and a second *hiding* phase, during which the EMU deploys the optimal policy $p_{Y|X}^*$.

More formally, consider n TSs. In the first $s(n)$ TSs, the so-called *storage phase*, no privacy is achieved because we have $Y_t = X_t$, for $t = 1, 2, \dots, s(n)$. In the remaining $n - s(n)$ TSs, the so-called *hiding phase*, the user demand is satisfied by taking energy from both the grid and the battery according to the optimal policy $p_{Y|X}^*$. We assume that $s(n) = o(n)$, with $\lim_{n \rightarrow \infty} s(n) = \infty$, and $\lim_{n \rightarrow \infty} n - s(n) = \infty$. The initial waiting time $s(n)$ enables the battery to store on average $s(n)\bar{P}_E$ units of energy. In

the following lemma we show that the energy stored in the initial storage phase is sufficient to let the EMU follow the optimal energy management policy $p_{Y|X}^*$ during the hiding phase, without energy outages almost surely. After $s(n)$ units of time, thanks to the energy already stored in the RB, the system is able to overcome the uncertainty in the energy arrival, and is able to adopt the optimal privacy-preserving energy management policy for the remaining time.

Remark. It is noteworthy that no information about the recharge process of the battery is required, and all the EMU needs to know is the average power generated by the renewable energy process, $\bar{P}_E$.

Lemma 3.5. *With a storage phase of length $s(n) = o(n)$, where $\lim_{n \rightarrow \infty} s(n) = \infty$, and $\lim_{n \rightarrow \infty} n - s(n) = \infty$, the store-and-hide policy satisfies the energy constraints in (3.11) almost surely provided that $\mathbb{E}[X - Y^*] < \bar{P}_E$.*

The proof of Lemma 3.5 can be found in Appendix A.1.

By means of Lemma 3.5 it is possible to show that the minimum information leakage rate of the store-and-hide policy approaches $\mathcal{I}^i(\infty, \hat{P}_d)$ as $n \rightarrow \infty$, as shown in the following lemma, whose proof can be found in Appendix A.2.

Lemma 3.6. *If $\mathbb{E}[X - Y^*] < \bar{P}_E$, then the information leakage rate of the store-and-hide policy with $s(n)$ as specified in Lemma 3.5 approaches $\mathcal{I}^i(\infty, \hat{P}_d)$ as $n \rightarrow \infty$.*

Remark. Even though the two schemes described above achieve the same privacy performance as $n \rightarrow \infty$, they do have some conceptual differences. During the initial phase of energy saving, the store-and-hide policy satisfies all the user electricity demands from the grid leaking full information. Therefore, the SM readings reveal user's activity completely in this period. While the impact of this on the information leakage rate vanishes as $n \rightarrow \infty$, this might not be preferable in practice. Therefore, we believe that the best-effort policy is more appropriate for practical applications.

3.4.4 Generated Renewable Energy Known by the Utility Provider

Here we assume that the UP knows the realization of the renewable energy process E^n , as highlighted in Figure 3.2. This scenario can occur if, for example, we consider




FIGURE 3.2: Representation of the system model where the UP has perfect knowledge about the realizations of the renewable energy generation process, E^n , in addition to the power used from the grid that is reported through the SM readings, Y^n .

solar energy as the RES, and the UP can accurately estimate the renewable energy produced from its own observations in nearby locations, weather forecast of the area, and the specifications of the solar panel. This is a worst-case situation and we expect the amount of leaked information in this case to be greater than or equal to that of the previous scenario, in which only the EMU knows the current state of the renewable energy produced. In this setting, the information leakage rate is defined as

$$\bar{\mathcal{I}}^i(\infty, \hat{P}_d) \triangleq \inf_{f \in \mathcal{F}} \limsup_{n \rightarrow \infty} \frac{1}{n} I(X^n; Y^n | E^n). \quad (3.15)$$

The following theorem states that E^n does not necessarily provide more information to the UP compared to the scenario where the UP does not have access to this information.

Theorem 3.3. *If $B_{\max} = \infty$, the minimum information leakage rates for the cases in which E^n is either known or not known to the UP are the same, i.e., $\bar{\mathcal{I}}^i(\infty, \hat{P}_d) = \mathcal{I}^i(\infty, \hat{P}_d)$.*

Proof. We have the following chain of inequalities:

$$\limsup_{n \rightarrow \infty} \frac{1}{n} I(X^n; Y^n | E^n) = \limsup_{n \rightarrow \infty} \frac{1}{n} I(X^n; Y^n, E^n) \quad (3.16a)$$

$$= \limsup_{n \rightarrow \infty} \frac{1}{n} [I(X^n; Y^n) + I(E^n; X^n | Y^n)] \quad (3.16b)$$

$$\geq \limsup_{n \rightarrow \infty} \frac{1}{n} I(X^n; Y^n), \quad (3.16c)$$

where (3.16a) follows from the chain rule of mutual information and because X and E are independent from each other, and (3.16c) is due to the non negativity of mutual information. Thus, we have $\bar{\mathcal{I}}^i(\infty, \hat{P}_d) \geq \mathcal{I}^i(\infty, \hat{P}_d)$.

The inequality in (3.16c) becomes an equality if $I(E^n; X^n | Y^n) = 0$. This condition can be achieved by the store-and-hide policy. In fact, at the end of the storage phase the battery is filled up with an infinite amount of energy, and, as a consequence, the optimal policy during the hiding phase $p_{Y|X}^*$ does not need to take the information about the RES into account. This implies that $\limsup_{n \rightarrow \infty} I(E^n; X^n | Y^n) = 0$; and therefore, $\limsup_{n \rightarrow \infty} \frac{1}{n} I(X^n; Y^n | E^n) = \limsup_{n \rightarrow \infty} \frac{1}{n} I(X^n; Y^n)$, and that $\bar{\mathcal{I}}^i(\infty, \hat{P}_d) = \mathcal{I}^i(\infty, \hat{P}_d)$. $\square$

3.5 Smart Meter System Without Energy Storage

In this section we focus on another limit scenario in which there is no RB for storing extra renewable energy, i.e., $B_{\max} = 0$. The renewable energy available at TS t , E_t , can be considered as an i.i.d. state information, and could be known, or not, to the UP. Given E_t and X_t , the EMU decides on the amount of energy to use from the grid and from the RES. In each TS $t = 1, \dots, n$ the energy that can be obtained from the RES, $X_t - Y_t$, is limited by the energy generated in the same TS t , E_t , i.e., $0 \leq X_t - Y_t \leq E_t$. Thus, this is an SM system with a stochastic peak power constraint on the energy that the EMU can obtain from the RES. Therefore, this section can be considered as a generalization of [72], where the authors consider a fixed peak power constraint.

Remark. We note that a peak power constraint other than E_t can be easily incorporated into the model, as this would simply correspond to a new instantaneous power constraint of $X_t - Y_t \leq \min\{E_t, \hat{P}_d\}$. Therefore, for the brevity of the presentation we do not consider a peak power constraint in this section.

Note that, as opposed to the infinite-capacity battery scenario, here the past has no influence on the energy constraint, since there is no battery, and thus, no memory, in

the system.

To analyze this scenario, we first consider the minimum information leakage rate when the generated renewable energy is constant in every TS, i.e., $\mathcal{E} = \{e\}$, which is known by both the EMU and the UP. The privacy-power function is obtained by considering only a peak power constraint, which can be obtained as a special case of Theorem 3.1.

Lemma 3.7. *If $B_{\max} = 0$ and $\mathcal{E} = \{e\}$, the privacy-power function for an i.i.d. user load X is given by $\mathcal{I}(e, e)$.*

3.5.1 Generated Renewable Energy not Known by the Utility Provider

As in Section 3.4.1, here the realization of the renewable energy process is assumed to be known only by the EMU, while the UP only knows the probability distribution p_E .

Theorem 3.4. *If $B_{\max} = 0$, the user load is i.i.d. with distribution p_X , and the renewable energy produced by the RES is i.i.d. with distribution p_E , then the minimum information leakage rate, denoted by $\mathcal{I}^i(0)$, is given by*

$$\mathcal{I}^i(0) \triangleq \inf_{\substack{p_{Y|X} : p_{Y|X} = \sum_{e \in \mathcal{E}} p_{Y|X,E}(y|x,e)p_E(e); \\ p_{Y|X,E} \in \mathcal{P}^i}} I(X; Y), \quad (3.17)$$

where $\mathcal{P}^i \triangleq \{p_{Y|X,E} : p_{Y|X,E}(y|x,e) = 0 \text{ if } y > x \text{ or } y < x - e\}$.

Proof. Achievability. We consider a conditional probability distribution $p_{Y|X,E}(y|x,e)$ that satisfies the conditions of Theorem 3.4. At each TS, for given x_t and e_t , y_t is generated independently using the conditional distribution $p_{Y|X,E}(y_t|x_t, E_t = e_t)$. Since the user and grid load sequences are generated i.i.d. with the induced joint distribution $p_X(x)p_{Y|X}(y|x)$, the information leakage rate is given by $I(X; Y)$, whereas the instantaneous peak power constraint is satisfied for all conditional distributions in $\mathcal{P}^i$.

Converse. We assume that there is an energy management policy that satisfies the instantaneous peak power constraints, i.e., $x_t - y_t \leq e_t, \forall t$. Then, the information

leakage rate satisfies the following chain of inequalities:

$$\frac{1}{n}I(X^n; Y^n) = \frac{1}{n} \left[H(X^n) - H(X^n|Y^n) \right] \quad (3.18a)$$

$$= \frac{1}{n} \left[\sum_{t=1}^n H(X_t) - \sum_{t=1}^n H(X_t|X^{t-1}, Y^n) \right] \quad (3.18b)$$

$$\geq \frac{1}{n} \sum_{t=1}^n \left[H(X_t) - H(X_t|Y_t) \right] \quad (3.18c)$$

$$= \frac{1}{n} \sum_{t=1}^n I(X_t; Y_t) \quad (3.18d)$$

$$\geq \frac{1}{n} \sum_{t=1}^n \mathcal{I}^i(0) = \mathcal{I}^i(0), \quad (3.18e)$$

where (3.18b) follows since X is i.i.d.; (3.18c) follows since conditioning reduces entropy; and (3.18e) follows from the definition of $\mathcal{I}^i(0)$ in (3.17). $\square$

3.5.2 Generated Renewable Energy Known by the Utility Provider

Here we assume the UP also knows the realization of the renewable energy $E_t, \forall t$.

Theorem 3.5. *If $B_{\max} = 0$, the user load is i.i.d. with distribution p_X , the renewable energy produced by the RES is i.i.d. with distribution p_E , and the amount of generated renewable energy is also known by the UP at each time t , then the minimum information leakage rate $\bar{\mathcal{I}}^i(0)$ is given by*

$$\bar{\mathcal{I}}^i(0) \triangleq \inf_{p_{Y|X,E} \in \mathcal{P}^i} I(X; Y|E) = \mathbb{E}_E[\mathcal{I}(E, E)], \quad (3.19)$$

where $\mathcal{P}^i \triangleq \{p_{Y|X,E} : p_{Y|X,E}(y|x, e) = 0 \text{ if } y > x \text{ or } y < x - e\}$.

Proof. Achievability of (3.19) follows trivially by employing the optimal $p_{Y|X,E}$ that minimizes (3.19) at each TS. To prove the converse, we show that any energy management policy that satisfies the stochastic peak power constraint at each TS satisfies the

following chain of inequalities:

$$\frac{1}{n}I(X^n; Y^n | E^n) = \frac{1}{n} \left[H(X^n | E^n) - H(X^n | Y^n, E^n) \right] \quad (3.20a)$$

$$= \frac{1}{n} \sum_{t=1}^n \left[H(X_t | X^{t-1}, E^n) - H(X_t | X^{t-1}, Y^n, E^n) \right] \quad (3.20b)$$

$$\geq \frac{1}{n} \sum_{t=1}^n \left[H(X_t | E_t) - H(X_t | Y_t, E_t) \right] \quad (3.20c)$$

$$= \frac{1}{n} \sum_{t=1}^n \sum_{k=1}^{|\mathcal{E}|} p_E(E = e_k) I(X_t; Y_t | E_t = e_k) \quad (3.20d)$$

$$\geq \frac{1}{n} \sum_{t=1}^n \sum_{k=1}^{|\mathcal{E}|} p_E(E = e_k) \mathcal{I}(e_k, e_k) \quad (3.20e)$$

$$= \sum_{k=1}^{|\mathcal{E}|} p_E(E = e_k) \mathcal{I}(e_k, e_k) \quad (3.20f)$$

$$= \mathbb{E}_E [\mathcal{I}(E, E)], \quad (3.20g)$$

where (3.20c) follows because X and E are independent of each other and across time, and conditioning reduces entropy; (3.20d) follows by explicitly considering all the possible realizations of E_t ; and (3.20e) follows from Lemma 3.7. $\square$

From the chain rule of mutual information, we have

$$I(X; Y, E) = I(X; E) + I(X; Y | E) = I(X; Y | E), \quad (3.21a)$$

$$I(X; Y, E) = I(X; Y) + I(X; E | Y), \quad (3.21b)$$

where (3.21a) follows since X and E are independent of each other. From (3.21a) and (3.21b), we get $I(X; Y) \leq I(X; Y | E)$. Hence, from Theorems 3.4 and 3.5, we have $\mathcal{I}^i(0) \leq \bar{\mathcal{I}}^i(0)$, as expected.

3.6 Binary Scenario

In order to provide further insights into the behavior of the information leakage rate, here we consider a simple scenario with binary power demands, binary energy

generation and binary grid load, i.e., $\mathcal{X} = \mathcal{E} = \mathcal{Y} = \{0, 1\}$. This scenario may represent appliances that are either on or off/standby. X and E follow independent Bernoulli distributions with $\Pr\{X = 1\} = q_x$ and $\Pr\{E = 1\} = p_e$, respectively. We compare the minimum information leakage rates for the infinite and zero battery capacities.

If $B_{\max} = \infty$, the minimum information leakage rate can be characterized explicitly, by deriving the joint probability $p_{X,Y}$ as in the following. When $X = 0$, Y cannot be equal to 1 because of the constraint (3.1), i.e., $p_{X,Y}(X = 0, Y = 1) = 0$, whereas $p_{X,Y}(X = 0, Y = 0) = p_X(X = 0) = 1 - q_x$. When $X = 1$, $Y = 0$ with probability p_e , i.e., $p_{X,Y}(X = 1, Y = 0) = p_e$, whereas the remaining $p_{X,Y}(X = 1, Y = 1)$ is given by subtracting the previous probabilities from 1, i.e., $p_{X,Y}(X = 1, Y = 1) = q_x - p_e$. Summarizing, we obtain the following joint distribution:

$$p_{X,Y} = \begin{bmatrix} 1 - q_x & 0 \\ p_e & q_x - p_e \end{bmatrix}. \quad (3.22)$$

From (3.22) we can write the minimum information leakage rate for $B_{\max} = \infty$ as in the following:

$$\mathcal{I}^i(\infty, 1) = \mathcal{I}(p_e, 1) = \begin{cases} p_e \log p_e - q_x \log q_x \\ \quad - (1 - q_x + p_e) \log(1 - q_x + p_e), & \text{if } p_e \leq q_x, \\ 0, & \text{otherwise,} \end{cases} \quad (3.23)$$

where we set the peak power constraint to $\hat{P}_d = 1$.

When $B_{\max} = 0$, there are two scenarios. If the generated renewable energy is known only by the EMU, the minimum information leakage rate for this scenario is given by

$$\mathcal{I}^i(0; p_e, p_v, q_x) = h(1 - q_x + q_x p_e p_v) - q_x h(p_e p_v), \quad (3.24)$$

where $h(\cdot)$ is the binary entropy function defined as $h(p) \triangleq -p \log p - (1-p) \log(1-p)$; q_x is fixed, and p_v is the probability of using the energy available in the battery whenever $X = 1$ and $E = 1$; $H(Y) = h(1 - q_x + q_x p_e p_v)$, since $p(Y = 0) = 1 - q_x + q_x p_e p_v$; $H(Y|X) = H(Y|X = 1) = q_x h(p_e p_v)$, since $H(Y|X = 0) = 0$.

Proposition 3.1. *For every p_e and q_x , the information leakage rate $\mathcal{I}^i(0; p_e, p_v, q_x)$ is minimized with $p_v = 1$.*

Proof. The proof follows from observing that $\frac{d\mathcal{I}^i(0, p_v)}{dp_v} \leq 0, \forall p_e, q_x$. Thus, the minimum of $\mathcal{I}^i(0; p_e, p_v, q_x)$ is reached when p_v takes its maximum value, i.e., $p_v = 1$. $\square$

When E_t is known also by the UP, if the peak power constraint is $e = 1$, no information is leaked, whereas if $e = 0$, the user load is known perfectly by the UP, leading to a leakage of $H(X)$. Hence, the minimum information leakage rate when the state information is known by the UP is

$$\bar{\mathcal{I}}^i(0; p_e, q_x) = (1 - p_e)h(q_x). \quad (3.25)$$

A numerical comparison of the information leakage rate for zero and infinite battery capacities in the binary scenario will be presented in the next section together with the results corresponding to a finite battery capacity.

3.7 Finite Battery Capacity

A closed-form expression for the finite-capacity battery scenario is elusive as the presence of a finite battery brings memory into the system, and the future energy usage depends on how much renewable energy has been generated in the previous TSs, how much of that energy has already been used by the EMU, and how much is available in the RB. Instead, we propose a low-complexity energy management policy and compare it to the two previous scenarios, which represent upper and lower bounds on the system performance for the finite battery scenario.

3.7.1 Binary Alphabet: $\mathcal{X} = \mathcal{E} = \mathcal{Y} = \{0, 1\}$

In this setting X , E and Y have binary alphabets and we consider a discrete-time system, modeled via a finite state machine. As in Section 3.6, we set $\Pr\{X = 1\} = q_x$




FIGURE 3.3: Finite state diagram for the evolution of the level of energy in the RB with $\mathcal{B} = \{0, 1, \dots, B_{\max}\}$ and $\mathcal{X} = \mathcal{E} = \mathcal{Y} = \{0, 1\}$ for the battery-independent policy. The 4-tuple (x, e, v, y) represents, for every TS t , the values of the user load, the renewable energy produced, the power drawn from the battery by the EMU, and the grid load, respectively.

TABLE 3.1: Tuples and transition probabilities for the battery-independent policy when $\mathcal{X} = \mathcal{E} = \mathcal{Y} = \{0, 1\}$.

$\mathbf{B}_t$	$\mathbf{X}_t$	$\mathbf{E}_t$	$\mathbf{V}_t$	$\mathbf{Y}_t$	$\mathbf{B}_{t+1}$	Transition Probability
$B_t = 0$	0	0	0	0	0	$(1 - q_x)(1 - p_e)$
	0	1	0	0	1	$(1 - q_x)p_e$
	1	0	0	1	0	$q_x(1 - p_e)$
	1	1	0	1	1	$q_x p_e(1 - p_v)$
	1	1	1	0	0	$q_x p_e p_v$
$0 < B_t \leq B_{\max}$	0	0	0	0	B_t	$(1 - q_x)(1 - p_e)$
	0	1	0	0	$\min\{B_t + 1, B_{\max}\}$	$(1 - q_x)p_e$
	1	0	0	1	B_t	$q_x(1 - p_e)(1 - p_v)$
	1	0	1	0	$B_t - 1$	$q_x(1 - p_e)p_v$
	1	1	0	1	$\min\{B_t + 1, B_{\max}\}$	$q_x p_e(1 - p_v)$
	1	1	1	0	B_t	$q_x p_e p_v$

and $\Pr\{E = 1\} = p_e$, while $V^n \triangleq X^n - Y^n$ represents the power taken by the EMU from the battery, with $\mathcal{V} = \{0, 1\}$.

3.7.1.1 Battery-Independent Policy

Here we consider a time-invariant policy according to which the evolution of the battery state can be modeled as the Markov chain of Figure 3.3, where the 4-tuples (x, e, v, y) represent the realizations at TS t of the user load X , the renewable energy E , the power drawn from the battery by the EMU V , and the grid load Y , respectively. At every TS, the RB can be charged, discharged or remain with the same level of




FIGURE 3.4: Optimal p_v for the binary scenario and various battery capacities, when $q_x = 0.5$, for the battery-independent policy of Section 3.7.1.1.

energy, depending on the transition probabilities. We note that a similar model has been adopted in [83], with the difference that in [83] the RB can also store energy from the grid. We define p_v as the probability that the power is drawn from the battery provided that the user is asking for energy and that there is energy available for use, i.e., $p_v \triangleq \Pr\{V = 1 | X = 1, E + B \geq 1\}$. Since the value of p_v does not change according to the current level of energy in the RB, we name this policy *battery-independent policy*. Table 3.1 lists all the possible states and transition probabilities for this scenario. In particular, the table shows for each transition from B_t to B_{t+1} and each combination of the tuple (X_t, E_t, V_t, Y_t) the corresponding transition probability.

To compute the information leakage rate, all the distributions are considered to be Bernoulli. For $B_{\max} = \infty$ and $B_{\max} = 0$ we use the single-letter expressions derived in Section 3.6, and set $\hat{P}_d = 1$ for $B_{\max} = \infty$. For a finite-capacity battery, we implement the achievable scheme described above, and by means of the algorithm in [81] we simulate the system for very long sequences and evaluate the information leakage between the user and the grid loads numerically and for different battery capacities. Moreover, for each p_e , we find the value of p_v that achieves the minimum information leakage rate by searching over a discretized set of p_v values. As an example,

Figure 3.4 represents the optimal p_v values for each p_e , when the user load is uniformly distributed and $B_{\max} = \{1, 2, 5, 10\}$. In the figure, $p_e = 0$ is not represented because, regardless of p_v , the leakage when $p_e = 0$ is always equal to the entropy of the user load. Also, the figure shows that for higher p_e values, the minimum leakage is achieved for $p_v = 1$, i.e., it is better to always use the energy when available.

3.7.1.2 Battery-Conditioned Policy

Here we consider a policy, in which p_v , as defined before, can be different for various levels of energy in the battery, i.e., the policy is characterized by a specific p_{v_i} for each level of energy contained in the battery $B_t = i$, for $i = \{0, \dots, B_{\max}\}$. Thus, we now have the vector

$$\bar{p}_v = [p_{v_0}, p_{v_1}, \dots, p_{v_{B_{\max}}}] \quad (3.26)$$

To find the optimal $\bar{p}_v$ for each p_e and $B_{\max}$ we deploy a stochastic gradient descent algorithm, specifically we use the least square-based finite difference method to approximate the gradient [87]. Briefly, the algorithm works as follows. At any step, small perturbations are applied to each p_{v_i} according to a uniform distribution over a predefined interval, and the leakage corresponding to the resulting perturbed vector $\bar{p}_v$ is computed. The gradient of the leakage function can thus be approximated numerically by employing the leakage corresponding to a number of different perturbations. A new $\bar{p}_v$ is finally computed using the gradient estimate and a predefined learning rate, and its corresponding leakage is determined and compared with that of the previous step. If the difference between the two leakage rates is below a certain threshold, the algorithm stops. Otherwise, the algorithm keeps on iterating.

Figure 3.5 shows the information leakage rate with respect to the renewable energy generation rate p_e , for different battery capacities. For $B_{\max} = \{1, 2, 5, 10\}$, we adopt the battery-conditioned policy, which has only a small gain with respect to the battery-independent policy. In particular, this gain is focused around smaller p_e values. As expected, the least information leakage rate is achieved when $B_{\max} = \infty$ and $\hat{P}_d = 1$, while the maximum leakage occurs when $B_{\max} = 0$ and the UP knows the realizations of the renewable energy process. When $B_{\max} = 0$ the information leakage rate reduces




FIGURE 3.5: Minimum information leakage with respect to the renewable energy generation rate p_e with $\mathcal{X} = \mathcal{E} = \mathcal{Y} = \{0, 1\}$, for the battery-conditioned policy of Section 3.7.1.2. As $B_{\max}$ increases, the performance rapidly approaches that of $B_{\max} = \infty$ and $\hat{P}_d = 1$.

significantly if the realizations of the renewable energy process are not known by the UP and, more interestingly, we observe that the performance of the proposed suboptimal memoryless scheme approaches that of the infinite-capacity battery with relatively small battery sizes. In addition, we can see that the gain from the battery is much higher when the renewable generation rate is higher, i.e., when p_e is high. This is expected because when p_e is low, there is less energy to be stored for future TSs, hence the importance of having a large RB is diminished.

3.7.2 Larger Alphabets: $|\mathcal{X}| = |\mathcal{Y}| = |\mathcal{E}| > 2$

Here we consider larger alphabets for X , E and Y . As the alphabet sizes grow, so does the complexity of searching for the optimal policy. Instead, we consider the following suboptimal policy. At each time instant, the policy chooses among using all of the available energy, half of it, or no energy at all and we model the probability p_v




FIGURE 3.6: Minimum information leakage with respect to the renewable energy generation rate p_e with $\mathcal{X} = \mathcal{E} = \mathcal{Y} = \{0, 1, 2, 3, 4\}$. The leakage for $B_{\max} = \infty$ has been characterized by setting $\hat{P}_d = 4$.

as in the following:

$$p_v(B_t + E_t, X_t) = \begin{cases} (p_1, p_4), & \text{if } B_t + E_t < X_t, \\ (p_2, p_5), & \text{if } B_t + E_t = X_t, \\ (p_3, p_6), & \text{if } B_t + E_t > X_t. \end{cases} \quad (3.27)$$

The probability pairs in (3.27) refer to the probability of using all the available energy and the probability of using half of it. Therefore, we have $0 \leq p_i \leq 1$, for $i = 1, \dots, 6$, and $p_i + p_{i+3} \leq 1$, for $i = 1, 2, 3$. For example, if $B_t + E_t < X_t$, all of the available energy is used with probability p_1 , half of it, or the nearest integer value lower than that, is used with probability p_4 , and none of it is used with probability $1 - p_1 - p_4$.

Figure 3.6 shows the results for the scenario for $|\mathcal{X}| = |\mathcal{E}| = |\mathcal{Y}| = 5$ when $B_{\max} = \{0, 1, 2, \infty\}$. The user load is uniformly distributed over the alphabet $\mathcal{X}$, while the renewable energy generation follows a binomial distribution with parameters $|\mathcal{X}|$ (number of possible units of energy that can be generated) and p_e (probability of generating a single unit of renewable energy). The information leakage rate for the infinite and zero

battery capacities is computed by using the single-letter expressions which are evaluated by efficient numerical algorithms, specifically the BA algorithm [86] and the CVX package [88]. In particular, for $B_{\max} = \infty$ we set $\hat{P}_d = X_{\max}$. For the finite battery scenario, we adopt the aforementioned policy (3.27) and optimize the performance by trying different combinations of the probabilities p_i , $1 \leq i \leq 6$. Similar considerations to that of Figure 3.5 can be drawn for Figure 3.6 as well.

Remark. We remark here that, in order to isolate the privacy benefits of RESs, we do not allow charging the battery directly from the grid, which can potentially reduce the information leakage. It is known that modulating grid energy intake by employing a storage device provides privacy even in the absence of an RES [73, 82], or jointly with an RES [79]. The additional privacy benefits of allowing charging of the RB from the grid will depend on the battery capacity. When $B_{\max} = \infty$, perfect privacy can be achieved by charging the battery initially, and using the battery throughout the operation. In the other extreme scenario, that is, when $B_{\max} = 0$, obviously it is not possible to charge a non-existing battery from the grid. We leave a more detailed study of a finite-capacity storage device that can be charged by both the RES and the grid in the following chapters.

3.8 Continuous User Loads

In the simulation results presented above, we have considered discrete alphabets for all the involved random variables. A set of fixed discrete values for the energy demands may not be an accurate model for all the appliances in the real world. However, as discussed in Section 3.3, such hypothesis enables to constrain the grid load alphabet to the user load alphabet without loss of optimality, and to apply efficient algorithms to find the minimum amount of information leakage.

For continuous user loads, the optimal alphabet is also continuous. Thus, low-complexity numerical algorithms, such as the BA algorithm, cannot be applied. However, one can provide a lower bound on the privacy-power function by using the Shannon lower bound (SLB) [57, 89], which has been introduced by Shannon, and widely

used in the literature to provide a computable lower bound to the rate-distortion function. Although it is not always a tight bound, it is shown in [72] that the SLB provides a tight bound for the information leakage rate for an exponentially distributed user load. The SLB for the rate-distortion function $R(D)$ is defined as $H(X) - \phi(D)$ where $\phi(D) = \max_{p: \sum_{i=1}^m p_i d_i \leq D} H(p)$. The truncated exponential distribution maximizes the entropy for a given mean value $\bar{P}$ and a peak power constraint $0 \leq X \leq \hat{P}$ [57] and has the form given in [71] and reported in the following Eq. (3.28):

$$f_X(x) = \begin{cases} \frac{1}{\lambda_0} e^{-\frac{x}{\lambda_1}}, & \text{if } 0 \leq x \leq \hat{P}, \\ 0, & \text{otherwise,} \end{cases} \quad (3.28)$$

where $\lambda_0 \geq 0$ and $\lambda_1 \geq 0$ are chosen to satisfy the constraints on the moments:

$$\int_0^\infty x f_X(x) dx = \bar{P} \quad \rightarrow \quad -\frac{\lambda_1}{\lambda_0} \left(\hat{P} e^{-\frac{\hat{P}}{\lambda_1}} + \lambda_1 e^{-\frac{\hat{P}}{\lambda_1}} - \lambda_1 \right) = \bar{P}, \quad (3.29)$$

$$\int_0^\infty f_X(x) dx = 1 \quad \rightarrow \quad \frac{\lambda_1}{\lambda_0} (1 - e^{-\frac{\hat{P}}{\lambda_1}}) = 1. \quad (3.30)$$

The differential entropy of this truncated exponential distribution is given by

$$h(X) = - \int_0^{\hat{P}} \frac{1}{\lambda_0} e^{-\frac{x}{\lambda_1}} \log \left(\frac{1}{\lambda_0} e^{-\frac{x}{\lambda_1}} \right) dx \quad (3.31a)$$

$$= \frac{1}{\ln 2} \left(\log(\lambda_0) + \frac{\bar{P}}{\lambda_1} \right). \quad (3.31b)$$

Thus, the SLB for the privacy-power function introduced in Theorem 3.1 is given by (considering $\hat{P}_d$ instead of $\hat{P}$):

$$\mathcal{I}_{SLB}(\bar{P}, \hat{P}_d) = h(X) - \frac{1}{\ln 2} \left(\log(\lambda_0) + \frac{\bar{P}}{\lambda_1} \right). \quad (3.32)$$

Authors in [71] show that the SLB is indeed achievable for peak and average power constraints, by finding the conditional distribution $f_{Y|X}(y|x)$ that satisfies the SLB with equality, provided that the power drawn from the battery, $X - Y$, is distributed according to a truncated exponential distribution with mean $\bar{P}$ and peak $\hat{P}_d$.

Authors in [72] provide the SLB for the average power constraint, which, as we have shown, is equivalent to the infinite-capacity battery scenario.

3.8.1 No Battery - Renewable Energy not Known by the Utility Provider

Here only a peak power constraint is considered, i.e., $X - Y$ is constrained by $0 \leq X - Y \leq \hat{P}_d$. The distribution that maximizes the entropy over an interval is the uniform distribution

$$f_X(x) = \begin{cases} \frac{1}{\hat{P}_d}, & \text{if } 0 \leq x \leq \hat{P}_d, \\ 0, & \text{otherwise.} \end{cases} \quad (3.33)$$

For a fixed $\hat{P}_d$, the differential entropy of this distribution is $\log(\hat{P}_d)$. Then, the SLB in the case of zero capacity battery is

$$\mathcal{I}_{SLB}(\hat{P}_d) = h(X) - \log(\hat{P}_d). \quad (3.34)$$

3.8.2 No Battery - Renewable Energy Known by the Utility Provider

As in the previous scenario, only peak power constraints are considered and thus the entropy maximizing distribution is still the uniform distribution given in (3.33). The privacy-power function is given by the expected value of the privacy-power function related to every realization of the renewable energy process, over the distribution of all the realizations of the renewable energy process. Hence, the SLB is given by:

$$\mathcal{I}_{SLB}(\hat{P}_d) = \mathbb{E}_{\hat{P}_d}[h(X) - \log(\hat{P}_d)]. \quad (3.35)$$

3.9 Conclusions

In this chapter we have studied the information leakage in an SM system by considering an RES along with an RB. For infinite and zero battery capacities, we have

provided single-letter information-theoretic expressions for the minimum information leakage rate, which can be efficiently evaluated when the user load has a discrete alphabet. For these scenarios, we have also studied the information leakage rate when the UP knows the exact amount of renewable energy generated in each TS. In addition, for the finite-capacity battery scenario, we have proposed a suboptimal low-complexity energy management policy, and evaluated the corresponding privacy performance using a stochastic gradient descent algorithm. Our results show that the privacy achieved by the proposed low-complexity policy approaches the theoretical lower bound obtained by assuming an infinite-capacity battery with a relatively small battery capacity, especially when the generation rate of the RES is low or high.

Chapter 4

Smart Meter Privacy with a Finite-Capacity Battery

4.1 Overview

In this chapter we address the SM privacy problem by considering the availability of an RES and a finite-capacity RB. The RB can be charged by using both renewable energy and energy from the grid, and we adopt the average mutual information as a privacy measure. By expressing the information leakage as an additive quantity, we cast the problem as a stochastic control problem, and formulate the corresponding Bellman equations.

4.2 Introduction

In this chapter we focus on the analysis of the privacy performance when considering a finite-capacity RB and an RES at the user's premises. We adopt an information-theoretic approach, and measure privacy by the mutual information rate between the consumer's power consumption and the renewable energy generation process, and the power drawn from the grid, where the latter is known by the UP through the SM readings, and the former two are to be kept private.

The main contributions of this chapter can be summarized as follows:

- Differently from Chapter 3, in this chapter we allow the RB to be charged from the grid to further increase user's privacy, and we focus specifically on the finite capacity battery.

- We characterize the information leakage as an additive quantity over time, which enable us to cast the problem as an MDP. We note that a similar approach has been followed in [73], where a battery has been considered. Here, we also consider the presence of an RES to further hide user's consumption.
- We formulate the corresponding dynamic program (DP), which can be solved numerically in order to identify the optimal energy management policy.

The remainder of this chapter is organized as follows. In Section 4.3 the system model is introduced. In Section 4.4 it is assumed that the UP does not know the realizations of the renewable energy process, whereas in Section 4.5 it does. For both cases, the information leakage minimization problem can be cast as an MDP. Conclusions are drawn in Section 4.6.

4.3 System Model

The discrete-time system model adopted in this chapter is analogous to the general model introduced in Chapter 2 and illustrated in Figure 2.4. In this chapter we assume the user load $\{X_t\}_{t=1}^{\infty}$ to be a first-order time-homogeneous Markov chain with transition probability q_X , whose initial state X_1 is distributed according to p_{X_1} . Moreover, we assume the renewable energy process $\{E_t\}_{t=1}^{\infty}$ to be a first-order time-homogeneous Markov chain with transition probability q_E and initial state E_1 , distributed according to p_{E_1} . The initial state of the battery B_1 has distribution p_{B_1} . As in the previous chapter, each TS duration is normalized to unit time; therefore, power and energy values within a TS are used interchangeably and there is no need to explicitly mention the TS duration in the following equations. We would like to highlight that, compared to Chapter 3, where the X and E processes are i.i.d., here the processes are considered to be Markov, which is a more realistic assumption. We assume that all the involved random processes are defined over finite alphabets and that there is a minimum unit of energy such that all the aforementioned quantities are integer multiples of this unit. Moreover, we assume $Y_{\max} \geq X_{\max} + B_{\max}$. We would like to highlight that, differently from Chapter 3, in this chapter we allow the RB to be charged via the grid.

As usual, we assume that user demands have to be satisfied at all times, i.e., the constraint in Eq. (3.2) holds. Moreover, we do not allow intentional energy waste, or selling energy back to the grid; that is, we impose

$$0 \leq Y_t \leq [B_{\max} + X_t - B_t - E_t]^+, \quad \forall t. \quad (4.1)$$

Similarly, we do not allow wasting of renewable energy when the battery is not full. These actions could provide additional privacy to the user, albeit at a significantly higher energy cost. However, energy can still be wasted inevitably, for example, when the battery is full and the user load is smaller than the generated energy. The level of energy in the battery is updated as in Eq. (3.5).

Let $W_t \triangleq B_t + E_t - X_t$. The feasible set for Y_t , given $W_t = w_t$, is given by

$$\bar{\mathcal{Y}}(w_t) \triangleq \left\{ y_t \in \mathcal{Y} : [-w_t]^+ \leq y_t \leq [B_{\max} - w_t]^+ \right\}. \quad (4.2)$$

The energy management policy $f = (f_1, f_2, \dots)$ determines the amount of energy to request from the UP at each TS t , given the previous values of user load X^t , renewable energy E^t , battery state B^t , and grid load Y^{t-1} . We consider randomized policies, that is, each f_t is a conditional probability $f_t(y_t | x^t, e^t, b^t, y^{t-1})$, with $f \in \mathcal{F}$, i.e.,

$$f_t : \mathcal{X}^t \times \mathcal{E}^t \times \mathcal{B}^t \times \mathcal{Y}^{t-1} \rightarrow \mathcal{Y}, \quad \forall t. \quad (4.3)$$

Our goal is to minimize the leakage to the UP of information about user's energy consumption, as well as the amount of energy generated by the RES. Accordingly, the information leakage rate induced by policy f over n TSs is

$$\mathcal{I}(n, f) \triangleq \frac{1}{n} I^f(X^n, E^n, B_1; Y^n). \quad (4.4)$$

We note that including B_1 and E^n in Eq. (4.4) is necessary to allow for an easier mathematical formulation, as it will be clear from the following.

Thus, our optimization problem can be written as

$$\mathcal{I}(n) \triangleq \inf_{f \in \mathcal{F}} \frac{1}{n} I^f(X^n, E^n, B_1; Y^n), \quad (4.5)$$

where the minimization is over all the feasible policies $f \in \mathcal{F}$. Note that, in the above form, this is an extremely complex optimization problem as we need to specify $p(Y_t|X^t, E^t, B_1, Y^{t-1})$ for every t , and every possible realization $(X^t, E^t, B_1, Y^{t-1}) = (x^t, e^t, b_1, y^{t-1})$.

In the following, we identify two different scenarios based on the information available at the UP regarding the renewable energy process. For both scenarios, we will rewrite Eq. (4.5) in an additive form, which consequently can be formulated as a DP. Given q_X , q_E , p_{X_1} , p_{E_1} and p_{B_1} , our goal is to find an optimal battery charging policy $f = (f_1, f_2, \dots) \in \mathcal{F}$ that solves (4.5), for both finite and infinite horizon settings.

4.4 Renewable Energy not Known by the Utility Provider

In this section E^n is treated as a random sequence whose realizations are known causally only to the user. This scenario may occur if, for example, E^n originates from light, vibration, thermal, or biological sources, which could be extremely difficult, if not impossible, for the UP to track.

4.4.1 Additive Formulation for the Information Leakage Rate

To formulate the problem as an MDP we need to write the cost function, i.e., the information leakage rate, in an additive form. An introduction to MDPs can be found in Section 2.6.2. We aim at a formulation in which, conditioned on Y^{t-1} , the grid load at time t , Y_t , leaks information only on X_t , E_t and B_t , but not on their previous values. This will be achieved by restricting the set of possible policies to a smaller set $\mathcal{F}' \subseteq \mathcal{F}$, where each $f' \in \mathcal{F}'$ is a conditional probability of the form $f'_t(y_t|x_t, e_t, b_t, y^{t-1})$, i.e.,

$$f'_t : \mathcal{X} \times \mathcal{E} \times \mathcal{B} \times \mathcal{Y}^{t-1} \rightarrow \mathcal{Y}, \quad \forall t. \quad (4.6)$$

The joint measure induced by f' is

$$\begin{aligned} & p^{f'}(X^n = x^n, E^n = e^n, B^n = b^n, Y^n = y^n) \\ &= p_{X_1}(x_1)p_{E_1}(e_1)p_{B_1}(b_1)f'_1(y_1|x_1, e_1, b_1) \prod_{t=2}^n \left[\mathbb{1}_{b_t} \left\{ \min\{b_{t-1} + e_{t-1} - x_{t-1}, B_{\max}\} + y_{t-1} \right\} \right. \\ & \quad \left. \times q_X(x_t|x_{t-1})q_E(e_t|e_{t-1})f'_t(y_t|x_t, e_t, b_t, y^{t-1}) \right], \quad (4.7) \end{aligned}$$

where $\mathbb{1}$ is the indicator function, i.e., $\mathbb{1}_{b_t}\{a\} = 1$ if $b_t = a$, and 0 otherwise. $b_t \in \{0, \dots, B_{\max}\}$ holds since $Y_t \in \bar{\mathcal{Y}}(w_t)$.

The following theorem states that this restriction is without loss of optimality. Its proof, which follows steps similar to [73, Appendix II], can be found in Appendix B.1.

Theorem 4.1. *There is no loss of optimality in focusing only on charging strategies $f' \in \mathcal{F}'$, where $\mathcal{F}' \subseteq \mathcal{F}$. Moreover, the minimum information leakage rate can be expressed in the following additive form*

$$\mathcal{I}(n) = \inf_{f' \in \mathcal{F}'} \frac{1}{n} \sum_{t=1}^n I^{f'}(X_t, E_t, B_t; Y_t | Y^{t-1}). \quad (4.8)$$

It is possible to generalize Theorem 4.1 to the scenario in which X and E are Markov chains with order higher than 1, i.e., $p_{X_t|X^{t-1}} = p_{X_t|X_{t-m}^{t-1}}$ and $p_{E_t|E^{t-1}} = p_{E_t|E_{t-l}^{t-1}}$. If new policies $f'' \in \mathcal{F}''$ are defined such that

$$f''_t : \mathcal{X}^m \times \mathcal{E}^l \times \mathcal{B} \times \mathcal{Y}^{t-1} \rightarrow \mathcal{Y}, \quad \forall t, \quad (4.9)$$

where $\mathcal{F}'' \subseteq \mathcal{F}$, then the following corollary holds.

Corollary 4.1. *Let X and E be Markov chains of order m and l , respectively. There is no loss of optimality in focusing only on charging strategies $f'' \in \mathcal{F}''$, and for those strategies the minimum information leakage rate can be expressed in the following additive form*

$$\mathcal{I}(n) = \inf_{f'' \in \mathcal{F}''} \frac{1}{n} \sum_{t=1}^n I^{f''}(X_{t-m+1}^t, E_{t-l+1}^t, B_t; Y_t | Y^{t-1}). \quad (4.10)$$

The proof of Corollary 4.1, which follows steps similar to that of Theorem 4.1, can be found in Appendix B.2.

4.4.2 Markov Decision Process Formulation

Our next goal is to cast the problem as a stochastic control problem, which can be formulated as a DP. For this, we need to specify the state space, the control actions and the instantaneous cost corresponding to state-action pairs. The per-step cost in Eq. (4.8) depends on the past observations y^{t-1} , which could be considered as the state of the DP at TS t . However, this would mean a state space growing with time. To avoid this, we follow the approach of [90] and [76], and introduce a belief state, which can be shown to replace the y^{t-1} sequence.

The state of the DP at TS t is considered to be the causal posterior probability distribution over the triplet $S_t \triangleq (X_t, E_t, B_t)$, with $S_t \in \mathcal{S}$, given the knowledge of the previous realizations of the grid load, y^{t-1} :

$$\beta_t(s_t) \triangleq \begin{cases} p(s_1), & \text{if } t = 1, \\ p(s_t|y^{t-1}), & \text{otherwise.} \end{cases} \quad (4.11)$$

$\beta_t(s_t)$ can be considered as the *belief* that the UP has about s_t at TS t , given the SM measurements it has received in the previous TSs, Y^{t-1} .

The control action $U_t \in \mathcal{U}$ is the conditional probability $u_t(y_t|s_t)$. A *randomized history-dependent policy* $\pi = \{\pi_1, \pi_2, \dots\}$ chooses control action at time t via $u_t = \pi_t(h)$, where h represents the history available to the controller. Thus, the time ordering of the events is $S_1, U_1, Y_1, S_2, U_2, Y_2, \dots$. For the finite horizon problem, if the control action space and the initial state distribution spaces are countable, then without loss of optimality, we can focus on randomized *Markov policies* that depend only on the current state β_t , i.e., $\pi_t(\beta_t) = \pi_t(h)$. This holds also in the infinite horizon setting since our problem has nonnegative cost per stage [66]. Furthermore, for the finite horizon setting and the infinite horizon setting with finite state and control spaces, we can restrict ourselves to stationary deterministic Markov policies [66]. Policy π induces

the following joint measure:

$$\begin{aligned}
& p^\pi(X^n = x^n, E^n = e^n, B^n = b^n, Y^n = y^n) \\
&= p_{X_1}(x_1)p_{E_1}(e_1)p_{B_1}(b_1)u_1(y_1|x_1, e_1, b_1) \prod_{t=2}^n \left[\mathbb{1}_{b_t} \left\{ \min\{b_{t-1}+e_{t-1}-x_{t-1}, B_{\max}\} + y_{t-1} \right\} \right. \\
&\quad \left. \times q_X(x_t|x_{t-1})q_E(e_t|e_{t-1})u_t(y_t|x_t, e_t, b_t) \right]. \quad (4.12)
\end{aligned}$$

The state can be updated recursively, i.e., $\beta_{t+1} = \phi(\beta_t, u_t)$, as shown in the following:

$$\beta_{t+1}(s_{t+1}) = p(s_{t+1}|y^t) = \sum_{s_t} p(s_t, s_{t+1}|y^t) \quad (4.13a)$$

$$= \sum_{s_t} \frac{p(s_t, s_{t+1}, y_t|y^{t-1})}{p(y_t|y^{t-1})} \quad (4.13b)$$

$$= \frac{\sum_{s_t} p(s_t|y^{t-1})p(y_t|s_t, y^{t-1})p(s_{t+1}|s_t, y_t)}{\sum_{s_t, s_{t+1}} p(s_t|y^{t-1})p(y_t|s_t, y^{t-1})p(s_{t+1}|s_t, y_t)} \quad (4.13c)$$

$$\begin{aligned}
&= \frac{\sum_{s_t} \beta_t(s_t)u_t(y_t|s_t)q_X(x_{t+1}|x_t)q_E(e_{t+1}|e_t)\mathbb{1}_{b_{t+1}} \left\{ \min\{b_t + e_t - x_t, B_{\max}\} + y_t \right\}}{\sum_{s_t, s_{t+1}} \beta_t(s_t)u_t(y_t|s_t)q_X(x_{t+1}|x_t)q_E(e_{t+1}|e_t)\mathbb{1}_{b_{t+1}} \left\{ \min\{b_t + e_t - x_t, B_{\max}\} + y_t \right\}}, \\
&\quad (4.13d)
\end{aligned}$$

where (4.13c) follows from Bayes rule and the Markov chain $Y^{t-1} \rightarrow (S_t, Y_t) \rightarrow S_{t+1}$; and (4.13d) is due to the definitions of β_t and u_t .

Given $Y^{t-1} = y^{t-1}$, the per-step cost of taking action u_t is

$$g_t(x_t, e_t, b_t, u_t, y^t) \triangleq \log \frac{u_t(y_t|x_t, e_t, b_t)}{p(y_t|y^{t-1})}, \quad (4.14)$$

which is obtained from $I(X_t, E_t, B_t; Y_t|Y^{t-1})$ through the following steps:

$$\frac{p(x_t, e_t, b_t, y_t|y^{t-1})}{p(x_t, e_t, b_t|y^{t-1})p(y_t|y^{t-1})} = \frac{p(x_t, e_t, b_t|y^{t-1})p(y_t|x_t, e_t, b_t, y^{t-1})}{p(x_t, e_t, b_t|y^{t-1})p(y_t|y^{t-1})} \quad (4.15a)$$

$$= \frac{p(y_t|x_t, e_t, b_t, y^{t-1})}{p(y_t|y^{t-1})} \quad (4.15b)$$

$$= \frac{u_t(y_t|x_t, e_t, b_t)}{p(y_t|y^{t-1})}. \quad (4.15c)$$

It is possible to show that this new formulation is equivalent to the original problem, by considering the average n -horizon cost and the knowledge of y^{t-1} as in $\mathcal{I}(n, f')$

$$\mathcal{I}(n, \pi) = \frac{1}{n} \mathbb{E}^\pi \left[\sum_{t=1}^n g_t(x_t, e_t, b_t, u_t, y^t) \right] \quad (4.16a)$$

$$= \frac{1}{n} \sum_{t=1}^n \sum_{s_t \in \mathcal{S}, y^t \in \mathcal{Y}^t} p(s_t, y_t | y^{t-1}) \log \frac{u_t(y_t | s_t)}{p(y_t | y^{t-1})} \quad (4.16b)$$

$$= \mathcal{I}(n, f'), \quad (4.16c)$$

where we remind that u_t is also a function of y^{t-1} , since $u_t = \pi_t(\beta_t)$. Given a policy π , β_t and u_t are determined by y^{t-1} . The average information leakage at time t can be rewritten as in the following:

$$\mathbb{E}^\pi \left[g_t(X_t, E_t, B_t, U_t, Y^t) \right] = I(X_t, E_t, B_t; Y_t | Y^{t-1} = y^{t-1}) \quad (4.17a)$$

$$= \sum_{\substack{x_t \in \mathcal{X}, e_t \in \mathcal{E} \\ b_t \in \mathcal{B}, y^t \in \mathcal{Y}^t}} p(x_t, e_t, b_t | y^{t-1}) p(y_t | x_t, e_t, b_t, y^{t-1}) \\ \times \log \frac{u_t(y_t | x_t, e_t, b_t)}{p(y_t | y^{t-1})} \quad (4.17b)$$

$$= \sum_{s_t \in \mathcal{S}, y_t \in \mathcal{Y}} \beta_t(s_t) u_t(y_t | s_t) \log \frac{u_t(y_t | s_t)}{\sum_{\tilde{s}_t \in \mathcal{S}} \beta_t(\tilde{s}_t) u_t(y_t | \tilde{s}_t)} \quad (4.17c)$$

$$= I(X_t, E_t, B_t; Y_t | \beta_t, u_t), \quad (4.17d)$$

where the last step confirms that $Y^{t-1} \rightarrow (\beta_t, u_t) \rightarrow (X_t, E_t, B_t, Y_t)$ is a Markov chain.

The following lemma summarizes the results of this section.

Lemma 4.1. *Without loss of optimality, the SM privacy problem (4.5) for the scenario in which the UP does not know the realizations of the renewable energy process can be modeled as an MDP, such that*

1. the state at time t is given by (4.11),
2. the action at time t is specified by $u_t(y_t | x_t, b_t, e_t)$,
3. and the instantaneous cost is given by (4.17d).

In order to formulate the Bellman equations, it is convenient to first define an operator T for the DP as follows

$$(TJ)(\beta) = g(s, \pi(\beta), \beta) + \sum_{s \in \mathcal{S}, y \in \mathcal{Y}} \beta(s) u(y|s) J(\phi(\beta, u)), \quad (4.18)$$

for $\beta \in \bar{\mathcal{B}}$, where $J : \bar{\mathcal{B}} \rightarrow \mathbb{R}$ is the value function.

For the finite horizon setting, let J_t denote the value function at time $t \leq n$, with $J_{n+1} = 0$. For $t \leq n$, we have

$$J_t(\beta) = \inf_{u \in \mathcal{U}} [TJ_{t+1}](\beta). \quad (4.19)$$

The minimization problem is solved by going backwards in time from $t = n$ to $t = 1$ in order to find the optimal policy $\pi^* = (\pi_1^*, \pi_2^*, \dots, \pi_n^*)$ that minimizes (4.19) for every t .

In the infinite horizon scenario, since the total information leakage over an infinite number of stages is generally infinite, we minimize the average information leakage per stage, i.e.,

$$J(\beta_0) = \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}^\pi \left[\sum_{t=0}^{n-1} g_t(s_t, \pi_t(\beta_t), \beta_t) \right]. \quad (4.20)$$

The solution for the infinite horizon problem can be determined as the solution to the following Bellman equation

$$\lambda + J(\beta) = \inf_{u \in \mathcal{U}} [TJ](\beta), \quad (4.21)$$

where $\lambda \in \mathbb{R}$ is the optimal average information leakage, and the vector $J(\beta)$ is the relative or differential privacy leakage, i.e., the difference of the expected leakage to reach a conventional state and the cost that would be incurred if the cost per stage was equal to λ for all states. Eq. (4.21) can be solved via efficient dynamic programming algorithms, e.g., value iteration and policy iteration [66], and an optimal stationary policy $\pi^* = (\pi^*, \pi^*, \dots, \pi^*)$ can be found.




FIGURE 4.1: Representation of the system model where information about the realizations of the renewable energy process is available to the UP.

The following corollary generalizes our result to a user load and a renewable energy process as Markov chains of higher orders.

Corollary 4.2. *Let X and E be Markov chains of order m and l , respectively. Let $S_t \triangleq (X_{t-m+1}^t, E_{t-l+1}^t, B_t)$. The previous steps follow also for this scenario, where*

1. *the state is $\beta_t(s_t) \triangleq p(s_t|y^{t-1})$,*
2. *the action is $u_t(y_t|s_t)$,*
3. *and the cost is $g_t(s_t, u_t, y^t) \triangleq \log \frac{u_t(y_t|s_t)}{p(y_t|y^{t-1})}$.*

4.5 Renewable Energy Known by the Utility Provider

Here we assume that the UP knows the realizations of the renewable energy process E^n , as shown in Figure 4.1. This scenario can occur if we consider solar energy as the RES, and the UP can accurately estimate the renewable energy produced from its own observations in nearby locations, weather forecast of the area, and the specifications of the solar panel. The goal is to find a battery charging policy $f \in \mathcal{F}$ that minimizes the following information leakage rate

$$\mathcal{I}(n, f) = \frac{1}{n} I(X^n, E^n, B_1; Y^n | E^n) = \frac{1}{n} I(X^n, B_1; Y^n | E^n), \quad (4.22)$$

where the charging policy at TS t is defined in (4.3).

4.5.1 Additive Formulation for the Information Leakage Rate

In the same way of Section 4.4.1, we want to express the problem in an additive form. We define policies $\bar{f}'_t(y_t|x_t, e^t, b_t, y^{t-1})$ as

$$\bar{f}'_t : \mathcal{X} \times \mathcal{E}^t \times \mathcal{B} \times \mathcal{Y}^{t-1} \rightarrow \mathcal{Y}, \quad \forall t, \quad (4.23)$$

and state the following theorem.

Theorem 4.2. *There is no loss of optimality in focusing only on charging strategies $\bar{f}' \in \bar{\mathcal{F}}'$, where $\bar{\mathcal{F}}' \subseteq \mathcal{F}$. Moreover, the minimum information leakage rate can be written in the following additive form*

$$\mathcal{I}(n) = \inf_{\bar{f}' \in \bar{\mathcal{F}}'} \frac{1}{n} \sum_{t=1}^n I^{\bar{f}'}(X_t, B_t; E_t, Y_t | E^{t-1}, Y^{t-1}). \quad (4.24)$$

The proof, which follows similar steps to the proof of Theorem 4.1 and [73, Appendix II], is reported in Appendix B.3 for completeness.

For the setting in which X and E are Markov processes of order m and l , respectively, we define policies $\bar{f}''$ such that

$$\bar{f}''_t : \mathcal{X}^m \times \mathcal{E}^t \times \mathcal{B} \times \mathcal{Y}^{t-1} \rightarrow \mathcal{Y}, \quad \forall t. \quad (4.25)$$

Then the following Corollary 4.3 holds, whose proof follows similar steps to that of Theorem 4.2 and can be found in Appendix B.4.

Corollary 4.3. *Let X and E be Markov chains of order m and l , respectively. There is no loss of optimality in focusing only on charging strategies $\bar{f}'' \in \bar{\mathcal{F}}''$, and for those strategies the minimum information leakage rate can be written in the following additive form*

$$\mathcal{I}(n) = \inf_{\bar{f}'' \in \bar{\mathcal{F}}''} \frac{1}{n} \sum_{t=1}^n I^{\bar{f}''}(X_{t-m+1}^t, B_t; E_t, Y_t | Y^{t-1}, E^{t-1}). \quad (4.26)$$

4.5.2 Markov Decision Process Formulation

As in Section 4.4.2 we specify the state space, control actions and the instantaneous cost. The state of the DP at time t is considered to be the causal posterior probability distribution over $S_t \triangleq (X_t, B_t)$, given the knowledge of Y^{t-1} and E^{t-1} :

$$\beta_t(s_t) \triangleq \begin{cases} p(s_1), & \text{if } t = 1, \\ p(s_t|y^{t-1}, e^{t-1}), & \text{otherwise.} \end{cases} \quad (4.27)$$

$\beta_t(s_t)$ has the interpretation of *belief* that the UP has about s_t at TS t , given (Y^{t-1}, E^{t-1}) .

The action $U_t \in \mathcal{U}$ is the conditional probability $u_t(y_t|s_t, e_t)$ given by $u_t = \pi_t(h)$. As before, we consider without loss of optimality Markov policies $\pi = \{\pi_1, \pi_2, \dots\}$ that depend only on the current state β_t , i.e., $\pi_t(\beta_t) = \pi_t(h)$. As in (4.13d), β is updated recursively, i.e., $\beta_{t+1} = \phi'(\beta_t, u_t)$. e_t is not in the belief as the UP has perfect knowledge about it.

We follow steps similar to those of Section 4.4.2, and define the cost of taking action u_t when $s_t = (x_t, b_t)$ as

$$g_t(x_t, e_t, b_t, u_t, y^t) \triangleq \log \frac{q_E(e_t|e_{t-1})u_t(y_t|x_t, e_t, b_t)}{p(y_t, e_t|y^{t-1}, e^{t-1})}, \quad (4.28)$$

which is obtained from $I(X_t, B_t; E_t, Y_t|Y^{t-1}, E^{t-1})$ through the following steps

$$\frac{p(x_t, e_t, b_t, y_t|y^{t-1}, e^{t-1})}{p(x_t, b_t|y^{t-1}, e^{t-1})p(y_t, e_t|y^{t-1}, e^{t-1})} = \frac{p(x_t, e_t, b_t|y^{t-1}, e^{t-1})p(y_t|x_t, e_t, b_t, y^{t-1}, e^{t-1})}{p(x_t, b_t|y^{t-1}, e^{t-1})p(y_t, e_t|y^{t-1}, e^{t-1})} \quad (4.29a)$$

$$= \frac{q_E(e_t|e_{t-1})p(x_t, b_t|y^{t-1}, e^{t-1})u_t(y_t|x_t, e_t, b_t)}{p(x_t, b_t|y^{t-1}, e^{t-1})p(y_t, e_t|y^{t-1}, e^{t-1})} \quad (4.29b)$$

$$= \frac{q_E(e_t|e_{t-1})u_t(y_t|x_t, e_t, b_t)}{p(y_t, e_t|y^{t-1}, e^{t-1})}. \quad (4.29c)$$

By considering the average n -horizon cost, it is possible to show that this formulation is equivalent to the original problem

$$\mathcal{I}(n, \pi) = \frac{1}{n} \mathbb{E}^\pi \left[\sum_{t=1}^n g_t(x_t, e^t, b_t, u_t, y^t) \right] \quad (4.30a)$$

$$= \frac{1}{n} \sum_{t=1}^n \sum_{\substack{s_t \in \mathcal{S} \\ y^t \in \mathcal{Y}^t, e^t \in \mathcal{E}^t}} p(s_t, e_t, y_t | y^{t-1}, e^{t-1}) g_t(s_t, e^t, u_t, y^t) \quad (4.30b)$$

$$= \mathcal{I}(n, \bar{f}'), \quad (4.30c)$$

where u_t is also a function of (y^{t-1}, e^{t-1}) since $u_t = \pi_t(\beta_t)$. Given a policy π , β_t and u_t are determined by (y^{t-1}, e^{t-1}) . Then, it is possible to write

$$\mathbb{E}^\pi \left[g_t(X_t, E^t, B_t, U_t, Y^t) \right] = I(X_t, B_t; E_t, Y_t | Y^{t-1} = y^{t-1}, E^{t-1} = e^{t-1}) \quad (4.31a)$$

$$= \sum_{\substack{x_t \in \mathcal{X}, e^t \in \mathcal{E}^t \\ b_t \in \mathcal{B}, y^t \in \mathcal{Y}^t}} p(x_t, e_t, b_t | y^{t-1}, e^{t-1}) p(y_t | x_t, b_t, e^t, y^{t-1}) \\ \times \log \frac{q_E(e_t | e_{t-1}) u_t(y_t | x_t, e_t, b_t)}{p(y_t, e_t | y^{t-1}, e^{t-1})} \quad (4.31b)$$

$$= \sum_{\substack{s_t \in \mathcal{S}, y_t \in \mathcal{Y} \\ e_t \in \mathcal{E}}} \beta_t(s_t) q_E(e_t | e_{t-1}) u_t(y_t | s_t, e_t) \\ \times \log \frac{q_E(e_t | e_{t-1}) u_t(y_t | s_t, e_t)}{\sum_{\tilde{s}_t \in \mathcal{S}} \beta_t(\tilde{s}_t) q_E(e_t | e_{t-1}) u_t(y_t | \tilde{s}_t, e_t)} \quad (4.31c)$$

$$= I(X_t, B_t; E_t, Y_t | \beta_t, q_E, u_t), \quad (4.31d)$$

where the last step confirms that $(Y^{t-1}, E^{t-1}) \rightarrow (\beta_t, q_E, u_t) \rightarrow (X_t, E_t, B_t, Y_t)$ is a Markov chain.

The following lemma summarizes the results of this section.

Lemma 4.2. *Without loss of optimality, the SM privacy problem (4.5) for the scenario in which the UP knows the realizations of the process E^n can be modeled as an MDP, such that*

1. the state at time t is given by (4.27),
2. the action at time t is specified by $u_t(y_t | x_t, e_t, b_t)$,

3. and the instantaneous cost is given by (4.31d).

Bellman equations for the finite and infinite horizon problems can be obtained as in Section 4.4, with the consequent changes in the formulations of β , u and π .

The following corollary, counterpart of Corollary 4.2, holds.

Corollary 4.4. *Let X and E be Markov chains of order m and l , respectively, $S_t \triangleq (X_{t-m+1}^t, B_t)$, and q_E^l the l -th order transition probability. The previous steps follow also for this scenario, where*

1. the state is $\beta_t(s_t) \triangleq p(s_t|y^{t-1}, e^{t-1})$,
2. the action is $u_t(y_t|s_t, e_t)$,
3. the cost is $g_t(s_t, u_t, y^t, e^t) \triangleq \log \frac{q_E^l(e_t|e_{t-l}^{t-1})u_t(y_t|s_t, e_t)}{p(y_t, e_t|y^{t-1}, e^{t-1})}$.

4.6 Conclusions

In this chapter we have studied the information leakage rate in an SM system by considering the availability of an RES and a finite-capacity battery at the consumer's premises. The minimum information leakage rate has been characterized for both scenarios in which the UP knows and does not know the realizations of the renewable energy process. For both scenarios, we have formulated the minimum information leakage rate as an additive cost function, and cast the problem as an MDP, thereby finding the expressions for the corresponding Bellman equations. The optimal leakage rate for a given scenario can be obtained by discretizing the continuous belief state and applying dynamic programming techniques.

Chapter 5

Grid Load Variance as a Privacy Measure

5.1 Overview

In this chapter the privacy maximization problem is addressed together with the energy cost minimization problem for a SM user. It is assumed that an RB is available for the consumer, which can be utilized both to achieve privacy and to reduce the energy cost by modifying the energy consumption profile. Privacy is measured via the mean squared-error between the SM readings and a target load, while ToU pricing is considered for energy cost calculation. The optimal trade-off between the achievable privacy and the energy cost is characterized by taking into account the limited capacity of the battery as well as the capability to sell energy to the UP. Extensive numerical simulations are presented to evaluate the performance of the proposed strategy for different system settings.

5.2 Introduction

In this chapter we combine the objective of privacy maximization with that of minimizing the cost of energy for the user. While a widely accepted definition of privacy is elusive, one would expect that privacy is achieved when it is not possible to distinguish a specific appliance load from an aggregated household power consumption [56]. Statistical privacy measures have been studied in [50, 59, 72, 73, 78, 84], measuring the leaked information by the mutual information between the user and the grid loads; however, this requires the knowledge of the underlying statistics, and the results are

typically valid under various simplifications, such as assuming i.i.d. user load profiles, and over sufficiently long time horizons. As mentioned in Section 2.5.1, an alternative approach is based on the idea that a high degree of privacy can be achieved by flattening the power consumption around a constant target value, as considered in [48, 50, 91]. Accordingly, in this chapter we measure privacy as the mean squared error between the grid load and a *target load profile*, that is, the distance from a completely private profile. On the other hand, a completely constant consumption may not be practically viable or desirable, since the energy cost may vary greatly during the system operation due to ToU tariffs. For this reason, in this chapter the EMU is allowed to target a different fixed power value for each price period. In this model it is assumed that the EMU knows, or, accurately predicts, the load profile for the time period of interest, and obtains the optimal energy management policy by solving an optimization problem. We consider the availability of an RB at the user's premises, and characterize the optimal trade-off between the achievable privacy and the energy cost by taking into account the limited capacity of the battery, as well as the capability for the user to sell energy to the UP.

The main contributions of this chapter can be summarized as follows:

- We extend the privacy analysis carried out in the previous chapters by also considering the minimization of the cost of energy for the user. Moreover, we also allow the user to sell energy to the UP, to further enhance her privacy-cost trade-off.
- We generalize the target load profile introduced in [48, 50, 91], which we call the *constant target load*, by allowing the EMU to fix a constant target value for each price period, which we call the *piecewise constant target load*.
- We characterize the optimal solution, and highlight its water-filling interpretation. We show detailed numerical results to demonstrate the performance improvement of the proposed piecewise target load as compared to the constant target load

The remainder of this chapter is organized as follows. In Section 5.3 we present the system model, while in Section 5.4 we analyze the proposed target load analytically.




FIGURE 5.1: Representation of the system model. X_t , Y_t and $Y_t - X_t$ are the user load, the grid load, and the power exchanged with the RB at time t , respectively. The dashed line represents the SM readings being reported to the UP.

Numerical results are presented in Section 5.5, while conclusions are drawn in Section 5.6.

5.3 System Model

We consider the discrete-time system model depicted in Figure 5.1, in which each discrete time t represents a TS of duration d seconds, for $1 \leq t \leq N$, where N is the time horizon of interest. For TS t , the user load is denoted by $X_t \in \mathcal{X}$, while the grid load is $Y_t \in \mathcal{Y}$. We remark that the TSs in our model correspond to time instants when the electricity is actually requested by the user and drawn from the grid, rather than the typically longer sampling interval used for sending SM measurements to the UP. We assume that the SM measures and records the output power values at each TS; this is because our aim is to protect consumers' privacy not only from the UP, but also from the DSO or any other attacker that may deploy a sensor on the consumer's power line, recording the electricity consumption in almost real-time. We assume the presence of an RB at the user's premises, of capacity $B_{\max}$, which is used both to filter the user load in order to provide privacy, and to shift energy intake from the grid to

minimize the cost. The EMU computes the amount of power to draw from the grid, Y_t , and to exchange with the RB, $Y_t - X_t$. Let $B_t \in [0, B_{\max}]$ denote the amount of energy in the RB at the end of TS t , and we set $B_0 = 0$. The RB is charging if $Y_t - X_t > 0$, discharging if $Y_t - X_t < 0$, or remains with the same level of energy if $Y_t - X_t = 0$. We do not allow wasting grid energy; that is, there are no battery overflows, i.e.,

$$\sum_{\tau=1}^t (Y_\tau - X_\tau)d \leq B_{\max}, \quad t = 1, 2, \dots, N. \quad (5.1)$$

While additional energy can be stored in the battery for future use, we do not allow energy rescheduling or demand shifting, i.e., user's energy requests are always satisfied as they occur:

$$\sum_{\tau=1}^t Y_\tau \geq \sum_{\tau=1}^t X_\tau, \quad t = 1, 2, \dots, N. \quad (5.2)$$

In particular, Eq. (5.2) is also implicitly verified by the equation expressing the evolution of the level of energy in the RB:

$$0 \leq B_{t+1} = B_t + Y_{t+1}d - X_{t+1}d. \quad (5.3)$$

The power the RB can be charged or discharged at is constrained by $\hat{P}_c$ and $\hat{P}_d$, respectively. Thus, $\forall t$ we have:

$$Y_t - X_t \leq \hat{P}_c, \quad (5.4)$$

$$X_t - Y_t \leq \hat{P}_d. \quad (5.5)$$

Here we assume that the battery has perfect efficiency, that is, no leakages, to understand the optimal privacy-cost trade-off in an ideal setting. The effects of battery inefficiencies and degradations may be taken into account by introducing a constant rate leakage for the energy stored in the battery, or by imposing a further cost whenever the battery is used, within the framework studied in this chapter.

At first we do not consider selling energy to the UP, i.e., the power drawn from the grid is always non-negative

$$Y_t \geq 0, \quad \forall t. \quad (5.6)$$

Later we will lift this constraint. Given $(X_t, B_t) = (x_t, b_t)$, $B_{\max}$ and the constraints (5.4)-(5.6), the set of feasible power requests at time t is

$$\bar{\mathcal{Y}}_t(x_t, b_t) \triangleq \left\{ y_t \in \mathcal{Y} : \left[x_t - \min \left\{ \frac{b_t}{d}, \hat{P}_d \right\} \right]^+ \leq y_t \leq x_t + \min \left\{ \hat{P}_c, \frac{B_{\max} - b_t}{d} \right\} \right\}, \quad (5.7)$$

where $[a]^+ = a$ if $a > 0$, and 0 otherwise.

When $t = N$ we impose the RB to reach the same level of energy of the beginning of the time horizon of interest, i.e., we impose the RB to be emptied. Hence, we have the following constraint:

$$\sum_{t=1}^N Y_t = \sum_{t=1}^N X_t. \quad (5.8)$$

We aim at designing an *energy management policy* that decides on the amount of power to request from the UP at each TS t while satisfying (5.2)-(5.8). We consider an offline setting, in which the user load X^N and the cost of energy C^N are known beforehand and our goal is to jointly minimize the information leaked about the user's energy consumption as well as the cost the user incurs to purchase energy from the UP. While non-causal knowledge of the electricity price is a realistic assumption in today's energy networks, non-causal knowledge of power consumption is appropriate for appliances whose energy consumption can be accurately predicted over certain finite time frames, such as refrigerators, boilers, heaters and electric vehicles. Chapter 6 will analyze the case of limited information about the future power consumption. We also remark that the solution of the offline problem yields meaningful lower bounds on the performance of the more general online setting. In fact, in the offline setting, the knowledge of future energy consumption enables the user to better match the grid load with the target load, thus leading to the minimum amount of information leaked to the UP.

We measure the privacy leakage as the average variance of the grid load Y^N from the target load profile W^N over N TSs, i.e.,

$$\mathcal{P} \triangleq \frac{1}{N} \sum_{t=1}^N (Y_t - W_t)^2, \quad (5.9)$$

according to which, perfect privacy is achieved when $Y_t = W_t, \forall t$. We adopt squared distance in (5.9) so that we do not differentiate between negative and positive deviations of the grid load with respect to the target load. When energy cannot be sold to the UP, W_t is non-negative, i.e.,

$$W_t \geq 0, \quad t = 1, 2, \dots, N. \quad (5.10)$$

The average cost incurred by the user during time N is

$$\mathcal{C} \triangleq \frac{1}{N} \sum_{t=1}^N C_t Y_t, \quad (5.11)$$

where C_t is the cost of energy at time t , which is determined by the specific ToU tariff employed by the UP.

Compared to [50] where the target load is set to be a constant function over time, in the following section we consider a more general target load, specifically a function which assumes a constant value for each price period.

5.4 Target Load as a Piecewise Constant Function

In this section the objective of the optimization problem is to determine both the optimal target $\{W^*\}_{t=1}^N$ and the optimal grid load $\{Y^*\}_{t=1}^N$, so that the average squared-distance between them is minimized. The target load is set to be piecewise constant over price periods. We argue that such a grid load trades off privacy with cost. Multiple target values for power intake from the grid intuitively reveal more information about the user's real energy consumption behavior compared to a single constant target value. However, they also provide more flexibility to reduce the cost and to match




FIGURE 5.2: ToU tariff convention used in this chapter. $t_{c(i)}$, for $i = 1, \dots, M-1$, are the TSs at which the price of energy changes, while $t_{c(0)} = 0$ and $t_{c(M)} = t_N$ are the extremes of the time interval considered.

the target load, representing a trade-off between cost and privacy.

Let $C^{(i)}$ be the cost of the energy purchased from the UP during the i -th price period, where $1 \leq i \leq M$, and M is the total number of price periods within the time interval N . The i -th price period spans from TS $t_{c(i-1)}$ to $t_{c(i)} - 1$, and $N^{(i)} \triangleq (t_{c(i)} - t_{c(i-1)})/d$ is the number of TSs contained within price period i . Figure 5.2 depicts the timing convention considered in this scenario. Let $W^{(i)}$ be the constant target load for the i -th time period. Given the nature of the objective functions and the constraints, pairs of $(\mathcal{P}, \mathcal{C})$ form a convex region and the optimal points can be characterized by the Pareto boundary of this region. Then, our problem can be cast as a weighted combination of privacy (5.9) and cost (5.11), i.e.,

$$\min_{Y_t, W^{(i)}} \sum_{i=1}^M \sum_{t=t_{c(i-1)}}^{t_{c(i)}-1} \left[\alpha (Y_t - W^{(i)})^2 + (1 - \alpha) Y_t C^{(i)} \right], \quad (5.12)$$

where $0 \leq \alpha \leq 1$ is the weighting parameter between privacy and cost, according to which if $\alpha = 0$ only cost of energy is minimized, whereas if $\alpha = 1$ only information leakage is minimized.

Based on (5.12) and the constraints (5.1)-(5.2), (5.4)-(5.6), (5.8) and (5.10), we define the Lagrangian function as in the following

$$\begin{aligned}
\mathcal{L}(x, \boldsymbol{\lambda}, \nu) = & \sum_{i=1}^M \sum_{t=t_{c(i-1)}}^{t_{c(i)}-1} \left[\alpha(Y_t - W^{(i)})^2 + (1 - \alpha)Y_t C^{(i)} \right] \\
& + \sum_{t=1}^N \lambda_t^{(1)} \sum_{\tau=1}^t (X_\tau - Y_\tau) + \sum_{t=1}^N \lambda_t^{(2)} \left\{ \left[\sum_{\tau=1}^t (Y_\tau - X_\tau) d \right] - B_{\max} \right\} \\
& + \sum_{t=1}^N \lambda_t^{(3)} (Y_t - X_t - \hat{P}_c) + \sum_{t=1}^N \lambda_t^{(4)} (X_t - Y_t - \hat{P}_d) - \sum_{t=1}^N \lambda_t^{(5)} Y_t \\
& - \sum_{i=1}^{C_M} \lambda_i^{(6)} W^{(i)} + \nu \sum_{t=1}^N (Y_t - X_t), \tag{5.13}
\end{aligned}$$

where $\lambda_t^{(j)} \geq 0$ for $1 \leq j \leq 5$ and $1 \leq t \leq N$, $\lambda_i^{(6)} \geq 0$ for $1 \leq i \leq C_M$, and $\nu \geq 0$ are the Lagrange multipliers; and, denoting the vectors in bold, we have $\boldsymbol{\lambda} = [\boldsymbol{\lambda}^{(1)}, \boldsymbol{\lambda}^{(2)}, \boldsymbol{\lambda}^{(3)}, \boldsymbol{\lambda}^{(4)}, \boldsymbol{\lambda}^{(5)}, \boldsymbol{\lambda}^{(6)}]$. The corresponding slackness conditions are imposed on the inequality constraints:

$$\lambda_t^{(1)} \sum_{\tau=1}^t (X_\tau - Y_\tau) = 0, \quad t = 1, \dots, N, \tag{5.14}$$

$$\lambda_t^{(2)} \left\{ \left[\sum_{\tau=1}^t (Y_\tau - X_\tau) d \right] - B_{\max} \right\} = 0, \quad t = 1, \dots, N, \tag{5.15}$$

$$\lambda_t^{(3)} (Y_t - X_t - \hat{P}_c) = 0, \quad t = 1, \dots, N, \tag{5.16}$$

$$\lambda_t^{(4)} (X_t - Y_t - \hat{P}_d) = 0, \quad t = 1, \dots, N, \tag{5.17}$$

$$\lambda_t^{(5)} Y_t = 0, \quad t = 1, \dots, N, \tag{5.18}$$

$$\lambda_i^{(6)} W^{(i)} = 0, \quad i = 1, \dots, C_M. \tag{5.19}$$

We then apply the Karush Kuhn Tucker (KKT) conditions and set the gradient of the Lagrangian to zero, i.e.,

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial Y_t} = & 2\alpha(Y_t - W^{(i)}) + (1 - \alpha)C^{(i)} + \sum_{\tau=t}^N (\lambda_\tau^{(2)} d - \lambda_\tau^{(1)}) \\
& + \lambda_t^{(3)} - \lambda_t^{(4)} - \lambda_t^{(5)} + \nu^{(1)} = 0, \tag{5.20}
\end{aligned}$$

$$\frac{\partial \mathcal{L}}{\partial W^{(i)}} = - \sum_{t=t_{c(i-1)}}^{t_{c(i)}-1} 2\alpha(Y_t - W^{(i)}) - \lambda_i^{(6)} = 0. \quad (5.21)$$

From (5.20), we obtain the optimal value for Y_t as

$$Y_t^* = \left[\frac{a_t}{2\alpha} + W^{(i)*} - \tilde{C}^{(i)} \right]^+, \quad t = 1, \dots, N, \quad (5.22)$$

where $a_t = \sum_{\tau=t}^N (\lambda_\tau^{(1)} - \lambda_\tau^{(2)}d) - \lambda_t^{(3)} + \lambda_t^{(4)} + \lambda_t^{(5)} - \nu^{(1)}$ and $\tilde{C}^{(i)} = \frac{(1-\alpha)C^{(i)}}{2\alpha}$. From (5.21), we obtain the optimal value for $W^{(i)}$ as

$$W^{(i)*} = \left[\frac{\lambda_i^{(6)}}{2\alpha N^{(i)}} + \frac{\sum_{t=t_{c(i-1)}}^{t_{c(i)}-1} Y_t^*}{N^{(i)}} \right]^+, \quad i = 1, \dots, M. \quad (5.23)$$

The optimal solution (5.22) resembles the classical water-filling algorithm that determines the optimal power allocation for Gaussian parallel channels under a total power constraint [57]. However, differently from the classical water-filling formulation, here the water level, $Y_t^* + \tilde{C}^{(i)} = \frac{a_t}{2\alpha} + W^{(i)*}$, is not constant, but varies over time due to the instantaneous constraints. The optimal target for price period i , given in (5.23), is a function of the Lagrange multipliers and of the mean of the optimal grid loads in the same price period. We note that, since privacy and cost in (5.12) may have significantly different magnitudes, they need to be further normalized to get the Pareto optimal solution consistent with α . Hence, we implement the normalization approach described in [92].

5.4.1 Illustration of the Water-filling Solution

The simple scenario illustrated in Figure 5.3 offers some intuition about the optimal solution and its water-filling interpretation. We consider two price periods, $C^{(1)} = 1$ and $C^{(2)} = 3$, two batteries of parameters $B_{\max,1} = 4$, $\hat{P}_{c,1} = \hat{P}_{d,1} = 2$ and $B_{\max,2} = 8$, $\hat{P}_{c,2} = \hat{P}_{d,2} = 4$, respectively, $N = 4$ and $\alpha = 0.5$. The user load is $X^4 = [1, 4, 2, 5]$. For the sake of simplicity, we consider $d = 1$, so that power and energy can be used interchangeably, with no specific units of measure adopted. Figure 5.3a, shows the optimal




FIGURE 5.3: Optimal grid and target load profiles for $\alpha = 0.5$, when energy selling is not allowed.

solution and water levels for $B_{\max,1}$. Since $C^{(1)} < C^{(2)}$, more energy is requested from the grid during the first price period and stored in the battery to partially satisfy the user's energy consumption during the second price period. However, the energy that can be stored is limited not only by the battery capacity $B_{\max}$ but also by the peak power constraint $\hat{P}_{c,1}$ (5.4). In the first TS it is indeed $\hat{P}_{c,1}$ that limits the grid load. We have $Y_1^* = X_1 + \hat{P}_{c,1} = 3$ and the amount of energy in the battery at the end of the TS is $B_1 = \hat{P}_{c,1} = 2$. The water level for $t = 1$ is given by $Y_1^* + \tilde{C}^{(1)} = \frac{4a_1}{2\alpha} + W^{(1)*}$, where $a_1 = \sum_{\tau=1}^4 (\lambda_{\tau}^{(1)} - \lambda_{\tau}^{(2)}d) - \lambda_1^{(3)} + \lambda_1^{(4)} + \lambda_1^{(5)} - \nu^{(1)}$, and $W^{(1)*}$ is computed jointly with Y_1^* and Y_2^* (5.23). In the second TS the grid load is limited by both the charging peak power constraint and the battery capacity, i.e., $Y_2^* = X_2 + \hat{P}_{c,1} = 6$, and $B_2 = B_{\max} = 4$. In the third TS $Y_3^* = 0$, which does not violate the discharging peak power constraint as $X_3 = \hat{P}_{d,1} = 2$, and $B_3 = 2$. In the last TS, $Y_4^* = X_4 - \hat{P}_{d,1} = 3$ and the battery becomes empty. Figure 5.3b shows that by increasing $B_{\max}$, $\hat{P}_c$ and $\hat{P}_d$, it is possible to further reduce the variance of the grid load from the target and the cost of energy, as a larger amount of cheaper electricity is stored in the battery during the first 2 TSs, for use in the later TSs.

5.4.2 Selling Energy to the Grid

Here we consider the possibility of selling energy to the UP in order to further improve the user's privacy and save costs. By allowing this, the UP can practically utilize user devices as distributed energy storage. Thus, we lift constraints (5.6) and (5.10), and we have

$$Y_t : \begin{cases} \geq 0, & \text{if energy is purchased from the UP,} \\ < 0, & \text{if energy is sold to the UP.} \end{cases} \quad (5.24)$$

The price of energy sold to the grid is set equal to the price of energy bought from it, i.e., we implement the *net metering* approach, in which a single SM can measure bi-directional energy flows [93].

The set of feasible power requests at time t is given by

$$\tilde{\mathcal{Y}}_t(x_t, b_t) \triangleq \left\{ y_t \in \mathcal{Y} : x_t - \min \left\{ \frac{b_t}{d}, \hat{P}_d \right\} \leq y_t \leq x_t + \min \left\{ \hat{P}_c, \frac{B_{\max} - b_t}{d} \right\} \right\}. \quad (5.25)$$

The optimization problem is expressed as in Eq. (5.12) and the Lagrangian formulation is the same as in Eq. (5.13) without including the constraints corresponding to $\lambda_t^{(5)}$ and $\lambda_i^{(6)}$. The slackness conditions follow similarly, and by applying the KKT conditions, we obtain the following optimal values for Y_t and $W^{(i)}$:

$$Y_t^* = \frac{a_t}{2\alpha} + W^{(i)*} - \tilde{C}^{(i)}, \quad t = 1, \dots, N, \quad (5.26)$$

$$W^{(i)*} = \frac{\sum_{t=t_c^{(i-1)}}^{t_c^{(i)}-1} Y_t^*}{N^{(i)}}, \quad i = 1, \dots, M, \quad (5.27)$$

where $a_t = \sum_{\tau=t}^N (\lambda_\tau^{(1)} - \lambda_\tau^{(2)} d) - \lambda_t^{(3)} + \lambda_t^{(4)} - \nu^{(1)}$.

5.5 Numerical Results

In this section we evaluate numerically the performance achieved when using the piecewise constant target and compare it with the constant power consumption target

TABLE 5.1: Specifications of the RBs considered in the simulations in this chapter.

Residential RB	RB Capacity $B_{\max}$ (kWh)	RB Charging Peak Power $\hat{P}_c$ (kW)	RB Discharging Peak Power $\hat{P}_d$ (kW)
Powervault G200-LI-4KWH [2]	4	1.2	1.4
Tesla Powerwall 2 [3]	13.5	5	5

in [50]. We note that we also include the peak power constraints (5.4)-(5.5) and the total energy constraint (5.8).

5.5.1 Simulation Settings

For numerical simulations, we use real SM consumption traces from the UK Dale dataset [53], which has a time resolution of 6 seconds. For our simulations we convert the original readings to a time resolution of 1 minutes to reduce the computational complexity. We adopt two commercial batteries, listed in Table 5.1, in which the charging (discharging) peak power denotes the amount of power that the battery is able to input (output) sustainably over time. We consider a ToU tariff currently being offered in the UK [94], in which the off-peak price is 4.99p/kWh during 23:00 to 6:00, the medium price is 11.99p/kWh during 6:00 to 16:00 and during 19:00 to 23:00, and the peak price is 24.99p/kWh during 16:00 to 19:00.

Figure 5.4 shows the user, grid and target load profiles for $\alpha = 0.001$, when using a constant and piecewise constant target load in Figure 5.4a and Figure 5.4b, respectively; and the load profiles for $\alpha = 0.999$, when using a constant and a piecewise constant target profile in Figure 5.4c and Figure 5.4d, respectively. These weights are chosen to simulate users that are mainly interested in minimizing cost of energy ($\alpha = 0.001$) or maximizing privacy ($\alpha = 0.999$). For these figures we adopt a Powervault G200-LI-4KWH battery [2], and we assume that energy cannot be sold to the UP. As the figures show, adopting a piecewise constant target reduces the variance of the grid load around the target, especially for $\alpha = 0.001$, thanks to the flexibility introduced by the piecewise constant target function. However, because of the rather small battery capacity, the peaks in the user load are mostly revealed in the grid load for both strategies. When $\alpha = 0.001$, since cost is more important than privacy, the battery is charged more during the off-peak price period, and the grid loads for the




FIGURE 5.4: User, target and grid load profiles for a Powervault G200-LI-4KWH [2].

two scenarios are more similar, since reducing the variance from the target assumes less importance.

Figure 5.5 shows the grid and target loads when a piecewise constant target is considered for both cases when energy selling is not allowed (see Figure 5.5a) and it is allowed (see Figure 5.5b). We consider a Tesla Powerwall 2 [3] as RB, and $\alpha = 0.5$. The results show that the target in this case can also have negative values, which allow the user to sell energy to the grid during the peak price period. Selling energy to the grid further minimizes the costs and improves the overall privacy-cost trade-off.

Figure 5.6 shows the trade-off between privacy and costs for the two strategies by considering the Powervault G200-LI-4KWH [2], and including the setting in which energy can be sold. The figure confirms the results of Figure 5.5 by showing that,




FIGURE 5.5: Comparison between allowing to sell energy to the UP and not allowing it, when using the piecewise constant target load, a Tesla Powerwall 2 [3] and $\alpha = 0.5$.




FIGURE 5.6: Privacy-cost trade-off by using a Powervault G200-LI-4KWH RB [2].

when energy can be sold, much lower costs are achieved for the same level of privacy. As expected, adopting a piecewise constant target leads to a much better privacy-cost trade-off compared to the constant target, both when selling and not selling energy.

Figure 5.7a shows the degree of privacy that can be achieved for different battery capacities for $\alpha = 1$, while Figure 5.7b shows the hourly average energy cost with respect to $B_{\max}$ for $\alpha = 0$. For these plots, we set $\hat{P}_c = \hat{P}_d = B_{\max}/2$ which is a good




FIGURE 5.7: Privacy and cost versus battery capacity.

approximation of the specifications of currently available residential batteries. Figure 5.7a shows that the information leakage drops greatly even for relatively small battery sizes, and it saturates at around $B_{\max} = 8$ kWh for all the strategies considered. This shows that, if a user is only interested in her privacy and has a medium-sized battery (≥ 6 kWh) available, she can achieve almost zero grid load variance by using any of the target loads considered. Also, for smaller battery capacities the piecewise constant target achieves better performance. However, we remark here that having multiple power levels has an inherent information leakage which is not captured solely by the load variance from the target. It is noteworthy that selling energy does not improve privacy when $\alpha = 1$, since selling introduces negative target values which are more difficult to match with the positive grid load unless a large battery is used. In the setting of Figure 5.7b, as privacy is not relevant, the only difference between the two targets is due to allowing or not to sell energy to the UP. As one would expect, the flexibility to sell energy to the grid can significantly reduce the energy cost, and in principle the user can even make a profit (corresponding to a negative cost in the figure) when using RBs with a large capacity.

5.6 Conclusions

In this chapter we have studied the joint optimization of privacy and cost for an SM system with an RB. Privacy is measured via the mean squared-error between the SM measurements and a target load profile, which is set to be a piecewise constant function over price periods. The user is allowed to sell excess energy to the UP. The optimal solution of the corresponding optimization problem has been characterized, highlighting its water-filling interpretation. Detailed numerical simulations have been presented, showing the performance improvement of the proposed piecewise constant target load as compared to a constant target load.

Chapter 6

Short Horizon Joint Privacy-Cost Optimization

6.1 Overview

In this chapter we draw upon Chapter 5 and study the joint privacy and cost optimization for a SM user, which is achievable by means of an RB. As in Chapter 5 we measure privacy via the mean squared distance between the SM readings and a target load profile, while ToU pricing is considered to compute the electricity cost. We allow the possibility for the consumer to sell electricity to the UP to further improve the privacy-cost trade-off. Two privacy-preserving energy management policies are proposed, which differ in the way the target load is characterized. Additionally, a simplified and more practical energy management policy, which optimizes the energy management less frequently, is considered. Numerical results are presented to compare the performances of various energy management policies in terms of the privacy-cost trade-off they achieve, considering a number of privacy indicators.

6.2 Introduction

In this chapter we extend the joint privacy-cost analysis carried out in Chapter 5, and following up on [50] and [52], we model SM privacy as the squared distance between grid and target load profiles. However, in this chapter we consider a more general target load profile, and assume that the consumer has only a partial knowledge of her future energy consumption and cost.

The main contributions of this chapter can be summarized as follows:

1. While full information on the future electricity consumption is assumed to be available at the EMU in [50] and in Chapter 5, which we call the *long horizon* model (LHM), in this chapter we consider a more realistic scenario whereby the consumer's future consumption profile is only partially known to the EMU, in a *moving* or *receding horizon* manner, which we call the *short horizon* model (SHM). The optimal solution at any time is computed only based on the currently available information within the *prediction horizon* by adopting a model predictive controller, recently implemented in a SM setting in [84]. In this chapter we present a detailed comparison of the results for SHM and LHM.
2. We introduce a target load profile computed as a low-pass filtered version of the user load, since higher-frequency components of a user's consumption profile leak more information about her behavior, compared to lower-frequency components.
3. We propose a more practical energy management policy which performs the optimization less frequently. The optimal solution is computed in batch, reducing the algorithm's computational load at the expense of the privacy-cost trade-off. Finally, we compare the privacy-cost trade-offs for all the scenarios considered.

The remainder of this chapter is organized as follows. In Section 6.3 we present the system model, while in Sections 6.4 and 6.5 we consider the SHM for a constant and a filtered target load profile, respectively. A more practical energy management policy with less regular policy updates is analyzed in Section 6.6, while conclusions are drawn in Section 6.7.

6.3 System Model

We consider the same discrete-time system depicted in Figure 5.1. We assume that the user's electricity consumption and electricity price are known for a horizon of H_F TSs beyond the current TS, naming H_F as the *prediction horizon*. Additionally, we assume that the EMU has memory about the past H_P TSs, which we call the *past horizon*. At each TS t , the EMU computes an energy management policy for the following H_F TSs, using its knowledge of the user load within the prediction horizon,

and the user and grid loads, and the RB level of energy within the past horizon. Let $\overline{t + H_F} \triangleq \min\{t + H_F, N\}$. As in Chapter 5, we do not allow wasting grid energy; that is, there are no battery overflows, i.e.,

$$B_{t-1} + \sum_{\tau=t}^{\overline{t+H_F}} (Y_\tau - X_\tau)d \leq B_{\max}, \quad \forall t, \quad (6.1)$$

and we impose that user's energy demands are always satisfied at the time of request. This leads to the following constraint:

$$\sum_{\tau=t}^{\overline{t+H_F}} (X_\tau - Y_\tau)d \leq B_{t-1}, \quad \forall t, \quad (6.2)$$

which is also implicitly verified by the equation expressing the evolution of the energy level in the battery (5.3). The power the RB can be charged or discharged at is constrained by $\hat{P}_c$ and $\hat{P}_d$, respectively, that is, the power constraints (5.4)-(5.5) hold. The model could be made more accurate by introducing further constraints, e.g., battery charging and discharging efficiency parameters, which we leave for future research that will focus on the practical implications of the proposed UDS techniques. As in Chapter 5, we study and compare the two scenarios in which energy can or cannot be sold to the UP, i.e., for which the constraints (5.6) and (5.24) hold, respectively, and we implement the net metering approach. The set of feasible energy requests at time t is given by Eqs.(5.7) and (5.25) for the scenarios in which energy selling is not allowed and allowed, respectively.

We design an energy management policy that decides on the grid load at each TS t while satisfying the above constraints. We consider a model predictive control approach, whereby the user load and the cost of energy are known beforehand within the prediction horizon $[t + 1, \dots, t + H_F]$, and our goal is to jointly minimize the information leaked about a user's energy consumption as well as the cost the user incurs to purchase energy from the UP. We note that the setting studied in Chapter 5, which assumes all future energy consumption and cost information to be known beforehand, is a lower bound on the setting studied in this chapter, as more information leads to a better privacy-cost trade-off.

We adopt the same definitions of privacy leakage $\mathcal{P}$ and cost $\mathcal{C}$ adopted in Chapter 5, which we rewrite here for an easier accessibility to the the reader:

$$\mathcal{P} \triangleq \frac{1}{N} \sum_{t=1}^N (Y_t - W_t)^2, \quad (6.3)$$

$$\mathcal{C} \triangleq \frac{1}{N} \sum_{t=1}^N C_t Y_t, \quad (6.4)$$

where W_t and C_t are the target load and the cost of energy at time t , respectively.

For numerical simulations, we use real SM consumption traces from the UK Dale dataset [53], converting the original readings to a time resolution of 10 minutes to reduce the computational complexity. As RB, we adopt the Tesla Powerwall 2 [3], for which $B_{\max} = 13.5\text{kWh}$, and $\hat{P}_c = \hat{P}_d = 5\text{kW}$. We consider the same ToU tariff adopted in Chapter 5 [94], in which the off-peak price is 4.99p/kWh during 23:00 to 6:00, the medium price is 11.99p/kWh during 6:00 to 16:00 and during 19:00 to 23:00, and the peak price is 24.99p/kWh during 16:00 to 19:00. All the simulation results are obtained for a time interval spanning 14 consecutive days, to average over a considerably large amount of data.

6.4 Target Load as a Constant Value

In this section, following up on [48, 56] and [50], we assume that the goal of the EMU is to keep the grid load as constant as possible. In [50] it is assumed that all the future user load and energy cost values are known, and the EMU can fix a target value for the whole duration, e.g., one whole day. In our model the information available to the EMU on X_t and C_t is limited to the prediction horizon, and changes over time; thus, the target load cannot be constant, and its variability depends on the length of the past and prediction horizons. In this section, given the knowledge of the cost of energy and the user's power consumption, the aim is to characterize both the optimal target load W^* and the optimal grid load G^* so as to optimize the overall privacy-cost trade-off.

Given the nature of the objective functions and the constraints, pairs of $(\mathcal{P}, \mathcal{C})$ form a convex region and the optimal points can be characterized by the Pareto boundary of this region. Hence, the objective can be cast as a weighted sum of privacy (6.3) and cost (6.4):

$$\min_{Y_t^{\overline{t+H_F}}, W_t} \alpha \sum_{\tau=\overline{t-H_P}}^{\overline{t+H_F}} (Y_\tau - W_t)^2 + (1 - \alpha) \sum_{\tau=t}^{\overline{t+H_F}} Y_\tau C_\tau, \quad (6.5)$$

where $0 \leq \alpha \leq 1$ is the weighting parameter, i.e., if $\alpha = 0$ only cost of energy is minimized, whereas if $\alpha = 1$ only information leakage is minimized; and $\overline{t - H_P} \triangleq \max\{t - H_P, 0\}$. We remark that setting the value of α is up to the consumer, who is in charge of deciding whether to focus more on protecting her privacy or on saving costs. The result of the minimization in (6.5) is the grid load for the current TS and the entire duration of the prediction horizon $Y_t^{\overline{t+H_F}}$, and the target load W_t . Eq. (6.5) characterizes the target load value W_t for the finite prediction horizon, which leads to the optimal privacy-cost trade-off over this horizon, based on the available information. At TS $t + 1$, the minimization (6.5) is carried out again based on the additional information that becomes available, i.e., Y_{t+H_F+1} and C_{t+H_F+1} , and Y_{t+1} and W_{t+1} are determined. The past horizon $\sum_{\tau=\overline{t-H_P}}^{t-1} (Y_\tau - W_t)^2$ is considered when optimizing for the privacy objective, since it ensures smoother variations of the overall target load profile. We note that, since privacy and cost in (6.5) may have significantly different magnitudes, we implement the same normalization approach utilized in Chapter 5 [92].

When $\alpha = 0$, (6.5) reduces to the following linear program:

$$\min_{Y_t^{\overline{t+H_F}}} \sum_{\tau=t}^{\overline{t+H_F}} Y_\tau C_\tau. \quad (6.6)$$

Remark. Differently from Chapter 5, here we do not impose the RB to be emptied at the end of each time window $[t - H_P, t + H_F]$, since in this context the end of the prediction horizon does not typically coincide with the end of the time horizon of interest, and the energy remaining in the RB can be utilized in the following TSs. Since the algorithm jointly optimizes privacy and cost, the RB is normally emptied at the end of the time horizon of interest N , unless α is high. If $\alpha \rightarrow 1$ and the RB is large, a sustained demand of energy may take place in the short term, which is ultimately constrained by




FIGURE 6.1: User, grid and target load profiles for the SHM and the LHM, for $\alpha = 0.5$ and $H_F = H_P = 2h$.

$B_{\max}$ and $\hat{P}_c$.

Figure 6.1 compares the load profiles of the SHM (Figures 6.1a and 6.1c) and the LHM (Figures 6.1b and 6.1d) over the course of one day. As one would expect, the LHM provides much better performance, in the sense that the resultant profile is much flatter, hiding most of the spikes in consumption, which typically reveal more information about user's behavior. However, the SHM, despite relying on much less data, also leads to a reasonable suppression of the peaks in the profile. One can also argue that the SHM reveals more information about the low-frequency variation of user's energy consumption, which, however, is common across households, and thus provides only a limited amount of personal information. Moreover, Figures 6.1a and

6.1c show that the peaks of the grid load generated by the SHM are not necessarily aligned with those of the user load.

6.4.1 Solution to the Optimization Problem (6.5)

In the following we consider the optimization problem (6.5) for $\alpha \neq \{0, 1\}$. We analyze first the optimal solution when selling of energy is allowed, i.e., (5.6) and (5.10) do not hold. Based on (6.5) and the constraints (6.1)-(6.2) and (5.4)-(5.5), we define the following Lagrangian function:

$$\begin{aligned} \mathcal{L}(Y_t^{\overline{t+H_F}}, W_t, \boldsymbol{\lambda}) = & \alpha \sum_{\tau=t-H_P}^{\overline{t+H_F}} (Y_\tau - W_t)^2 + (1-\alpha) \sum_{\tau=t}^{\overline{t+H_F}} Y_\tau C_\tau \\ & + \sum_{\tau=t}^{\overline{t+H_F}} \lambda_\tau^{(1)} \left[d \sum_{s=t}^{\tau} (Y_s - X_s) - B_{\max} + B_{t-1} \right] \\ & + \sum_{\tau=t}^{\overline{t+H_F}} \lambda_\tau^{(2)} \left[d \sum_{s=t}^{\tau} (X_s - Y_s) - B_{t-1} \right] \\ & + \sum_{\tau=t}^{\overline{t+H_F}} \lambda_\tau^{(3)} (Y_\tau - X_\tau - \hat{P}_c) + \sum_{\tau=t}^{\overline{t+H_F}} \lambda_\tau^{(4)} (X_\tau - Y_\tau - \hat{P}_d), \end{aligned} \quad (6.7)$$

where $\lambda_\tau^{(j)} \geq 0$, for $1 \leq j \leq 4$, are the Lagrange multipliers, and $t \leq \tau \leq \overline{t+H_F}$. Denoting the vectors in bold, we have $\boldsymbol{\lambda} = [\boldsymbol{\lambda}^{(1)}, \boldsymbol{\lambda}^{(2)}, \boldsymbol{\lambda}^{(3)}, \boldsymbol{\lambda}^{(4)}]$. The slackness conditions are imposed on the inequality constraints, for $\tau = t, t+1, \dots, \overline{t+H_F}$:

$$\lambda_\tau^{(1)} \left[d \sum_{s=t}^{\tau} (Y_s - X_s) - B_{\max} + B_{t-1} \right] = 0, \quad (6.8)$$

$$\lambda_\tau^{(2)} \left[d \sum_{s=t}^{\tau} (X_s - Y_s) - B_{t-1} \right] = 0, \quad (6.9)$$

$$\lambda_\tau^{(3)} (Y_\tau - X_\tau - \hat{P}_c) = 0, \quad (6.10)$$

$$\lambda_\tau^{(4)} (X_\tau - Y_\tau - \hat{P}_d) = 0. \quad (6.11)$$

Let $a_\tau \triangleq d \sum_{s=\tau}^{\overline{t+H_F}} (\lambda_s^{(2)} - \lambda_s^{(1)}) - \lambda_t^{(3)} + \lambda_t^{(4)}$, and $\tilde{C}_\tau \triangleq \frac{(1-\alpha)}{2\alpha} C_\tau$. Applying the KKT conditions and setting the gradient of the Lagrangian in Eq. (6.7) to zero, we obtain

the following expressions for the optimal grid and target load values:

$$Y_\tau^* = \frac{a_\tau}{2\alpha} - \tilde{C}_\tau + W_t^*, \quad \text{for } \tau = t, \dots, \overline{t+H_F}, \quad (6.12)$$

$$W_t^* = \frac{\sum_{\tau=\overline{t-H_P}}^{\overline{t+H_F}} Y_\tau^*}{1 + \min\{H_P, t\} + \min\{H_F, N-t\}}. \quad (6.13)$$

The optimal solution for the grid load given in Eq. (6.12) resembles the classical water-filling algorithm [57]. However, differently from the classical water-filling formulation and similarly to Chapter 5, here the water level, $Y_\tau^* + \tilde{C}_\tau = \frac{a_\tau}{2\alpha} + W_t^*$, is not constant, but varies over time due to the instantaneous constraints. The optimal solutions in (6.12) and (6.13) depend on the values of the Lagrangian multipliers and can be determined numerically.

When $\alpha = 0$, the only objective is to minimize the cost, and (6.5) reduces to a linear program (LP,) which can be solved using standard LP solvers. On the other hand, when the user is not concerned about the cost, i.e., $\alpha = 1$, (6.5) leads to a quadratic program analogous to the general case.

When energy selling is not allowed, the constraints given by Eq. (5.6) and Eq. (5.10) hold, and the Lagrangian is given by the following expression:

$$\begin{aligned} \mathcal{L}(Y_t^{\overline{t+H_F}}, W_t, \tilde{\lambda}) = & \alpha \sum_{\tau=\overline{t-H_P}}^{\overline{t+H_F}} (Y_\tau - W_t)^2 + (1-\alpha) \sum_{\tau=t}^{\overline{t+H_F}} Y_\tau C_\tau \\ & + \sum_{\tau=t}^{\overline{t+H_F}} \lambda_\tau^{(1)} \left[d \sum_{s=t}^{\tau} (Y_s - X_s) - B_{\max} + B_{t-1} \right] \\ & + \sum_{\tau=t}^{\overline{t+H_F}} \lambda_\tau^{(2)} \left[d \sum_{s=t}^{\tau} (X_s - Y_s) - B_{t-1} \right] \\ & + \sum_{\tau=t}^{\overline{t+H_F}} \lambda_\tau^{(3)} (Y_\tau - X_\tau - \hat{P}_c) + \sum_{\tau=t}^{\overline{t+H_F}} \lambda_\tau^{(4)} (X_\tau - Y_\tau - \hat{P}_d) \\ & - \sum_{\tau=t}^{\overline{t+H_F}} \lambda_\tau^{(5)} Y_\tau - \lambda^{(6)} W_t, \end{aligned} \quad (6.14)$$




FIGURE 6.2: Optimal grid and target load profiles, for $B_{\max} = 4$, $\hat{P}_c = \hat{P}_d = 2$, when energy selling is not allowed.

where $\tilde{\lambda} = [\lambda^{(1)}, \lambda^{(2)}, \lambda^{(3)}, \lambda^{(4)}, \lambda^{(5)}, \lambda^{(6)}]$ are the Lagrange multipliers. The slackness conditions are given by Eqs. (6.8)-(6.11) with the addition of the following two:

$$\lambda_{\tau}^{(5)} Y_{\tau} = 0, \quad \tau = t, \dots, \overline{t + H_F}, \quad (6.15)$$

$$\lambda^{(6)} W_t = 0. \quad (6.16)$$

Let $\tilde{a}_t \triangleq d \sum_{s=\tau}^{\overline{t+H_F}} (\lambda_s^{(2)} - \lambda_s^{(1)}) - \lambda_{\tau}^{(3)} + \lambda_{\tau}^{(4)} + \lambda_{\tau}^{(5)}$. Then, by applying the KKT conditions and by setting the gradient of the Lagrangian in Eq. (6.14) to zero we obtain the following expressions, counterpart of (6.12) and (6.13):

$$Y_{\tau}^* = \left[\frac{\tilde{a}_{\tau}}{2\alpha} - \tilde{C}_{\tau} + W_t^* \right]^+, \quad \text{for } \tau = t, \dots, \overline{t + H_F}, \quad (6.17)$$

$$W_t^* = \left[\frac{\sum_{\tau=t-H_P}^{\overline{t+H_F}} Y_{\tau}^* + \lambda^{(6)}}{1 + \min\{H_P, t\} + \min\{H_F, N - t\}} \right]^+. \quad (6.18)$$

6.4.2 Illustration of the Water-filling Solution

Here we present the solution for some simple scenarios to acquire an intuition on the solution of the optimization problems (6.5) and (6.6), and on their water-filling interpretation. We assume that energy selling is not allowed, $N = 6$, and $d = 1$, so

that power and energy can be used interchangeably. Let the user load and cost vectors be $X^6 = [1, 2, 6, 5, 2, 4]$ and $C^6 = [1, 2, 5, 3, 1, 3]$, respectively, and $H_P = H_F = 2$ TSs. Consider an RB with $B_{\max} = 4$ and $\hat{P}_c = \hat{P}_d = 2$. Figure 6.2a shows the optimal solution for $\alpha = 0$, $Y^{*,6} = [3, 4, 4, 3, 4, 2]$. Since the electricity is cheaper in the first and second TSs, more energy is requested from the grid during these TSs and stored in the RB to partially satisfy the demand at later TSs. However, the energy that can be requested in advance, and stored in the battery, is limited by both $B_{\max}$ and $\hat{P}_c$. In the first TS $\hat{P}_c$ limits the grid load, since $Y_1^* = X_1 + \hat{P}_c = 3$, and the level of energy in the RB at the end of the first TS is $B_1 = \hat{P}_c = 2$. In the second TS, the grid load is limited by $\hat{P}_c$ and $B_{\max}$ simultaneously, and $Y_2^* = X_2 + \hat{P}_c = 4$, and $B_2 = B_{\max} = 4$. Note that, although the third TS is the most expensive, $Y_3^* = 4$ because the battery cannot be discharged by more than 2 units of energy ($\hat{P}_d = 2$). For the same reason, $Y_4^* = 3$, while the remaining energy was stored in the second TS. Similar considerations hold for the last two TSs. Figure 6.2b illustrates the optimal solution for $\alpha = 1$, $Y^{*,6} = [3, 3.5, 4, 3.75, 3.88, 3.81]$. In this scenario, the energy management policy tries to match Y_t to the target load W_t , even at the cost of asking more energy than needed. The energy demand in the first TS is the same as the case $\alpha = 0$, whereas in the second TS less energy is stored in the battery to be used in the fourth TS, so that the water level matches the target load in this TS (at the expense of a higher cost in the fourth TS). It is noteworthy that more energy than needed is requested in the fourth TS, and used to satisfy the demand during the fifth TS, which is cheaper. Finally, more energy than needed is requested in the last two TSs and depicted in yellow, which is consequently stored in the RB to be used in future TSs. This allows the privacy-preserving algorithm to match the target and the grid load during these TSs.

Figure 6.3 shows the same scenario with $B_{\max} = 20$, $\hat{P}_c = \hat{P}_d = 10$, and when energy can be sold. When $\alpha = 0$, Figure 6.3a shows that energy is bought when it is cheaper and sold back to the grid when it is more expensive, while $\tilde{C}_3$ and $\tilde{C}_4$ are plotted as negative since energy is sold in these TSs. In fact, the RB is emptied of the energy stored during the first two TSs at the end of the fourth TS, and it is emptied again of the energy stored during the fifth TS at the end of the sixth TS. This is done to maximize the user's profit from selling energy to the grid when it is more expensive.




FIGURE 6.3: Optimal grid and target load profiles, for $B_{\max} = 20$, $\hat{P}_c = \hat{P}_d = 10$, when energy selling is allowed.

When $\alpha = 1$, Figure 6.3b shows that a larger RB permits greater flexibility but also boosts the amount of energy requested. This is not necessarily a disadvantage, as such energy can be used during later TSs.

6.4.3 Impact of the Duration of Prediction and Past Horizons

Figure 6.4 shows the load profiles for various combinations of prediction and past horizons H_F and H_P , respectively. As expected, a larger H_F produces flatter target and grid loads (see Figures 6.4c and 6.4d), compared to a smaller H_F (see Figure 6.4a). However, when H_F is larger the resulting grid load is more distant from the target load, thus resulting in a higher information leakage according to our definition of privacy leakage in Eq. (6.3). In fact, when H_F is small, the grid load values that are compared to the target load are few, and the energy management policy is able to determine a target load that is close to the grid load within the analyzed time window. On the other hand, when H_F is large, the energy management policy needs to find a single target load that matches a longer interval of predicted grid load values; as a result, the target load may be less representative for some periods. A longer past horizon leads to a flatter target load (see Figure 6.4b), however, the improvement for the grid load is not as evident as in the case of a longer prediction horizon (see Figure 6.4c). When the prediction horizon is long, i.e., $H_F = 12\text{h}$, the improvement obtained




FIGURE 6.4: Comparison between different past and prediction horizons for $\alpha = 0.5$, when energy selling is not allowed.

by considering a longer past horizon is limited, as the comparison of Figures 6.4c and 6.4d show.

Figure 6.5 shows the power profiles for various combinations of prediction and past horizons H_F and H_P , respectively, when energy selling is allowed. Considerations similar to those made for Figure 6.4 can be made here as well. In particular, it is clear that allowing a longer prediction horizon (see Figures 6.5c and 6.5d) enables to reduce the energy cost as more energy is drawn from the grid when it is cheaper, and sold back to the UP when it is more expensive. Additionally, a longer prediction horizon also enables a better suppression of the peaks in the user demand.

Figure 6.6 shows the average information leakage $\mathcal{P}$, average cost $\mathcal{C}$, and average




FIGURE 6.5: Comparison between various past and prediction horizons for $\alpha = 0.5$, when energy selling is allowed.




FIGURE 6.6: Impact of the prediction horizon H_F on information leakage, cost and target load variance for $H_P = 2\text{h}$.

target load variance with respect to H_F . The average target load variance, which can





FIGURE 6.7: Impact of the past horizon H_P on information leakage, cost and target load variance for $H_F = 2h$.

be considered as another privacy indicator, is defined as

$$\mathcal{V} \triangleq \frac{1}{N} \sum_{t=1}^N (W_t - \mu_W)^2, \quad (6.19)$$

where μ_W is the mean of W across time. When α is small, i.e., the main focus is to minimize the cost, a larger H_F reduces the average cost up to a certain extent, beyond which it cannot be further reduced (see Figure 6.6b); whereas for $\alpha = 1$ the cost does not change considerably with H_F . Opposite considerations hold for the information leakage, which even slightly increases when $\alpha \neq 1$ (see Figure 6.6a). This is due to the fact that a longer prediction horizon generates a grid load that is more distant from the target load, except for $\alpha = 1$, when the focus is on privacy only. Figure 6.6c shows that increasing H_F induces a smaller variance of the target load.







Figure 6.7, which shows $\mathcal{P}$, $\mathcal{C}$ and $\mathcal{V}$ with respect to H_P , exhibits similar behaviors to those in Figure 6.6, with some notable differences. The y -axis ranges are more limited here, confirming that the knowledge of past consumption is less significant for the energy management policy compared to the knowledge of future consumption and costs. This explains the far smaller reduction in cost in Figure 6.7b, as compared to Figure 6.6b, and the increase in the information leakage when $\alpha = 1$ in Figure 6.7a. Figure 6.7c shows that the target variance is higher and more variable when $\alpha = 1$; however, this corresponds to the case in which the grid load is closer to the target load, i.e., the most private scenario according to our original privacy measure. This contradiction of the two privacy indicators shows that evaluating the variance of the

FIGURE 6.8: Number of features versus H_F and H_P , when energy cannot be sold.

target load does not fully reflect the level of privacy achieved, as defined in Eq. (6.3).

6.4.4 Alternative Privacy Measures

As opposed to the LHM studied in [50], where W_t is fixed throughout the operation time, in the model considered in this chapter W_t is allowed to vary over time. Therefore, the squared distance between Y_t and W_t may not be sufficient as a privacy measure on its own. Accordingly, we consider alternative measures of privacy to see the impact of the proposed model predictive control framework on those measures. One of the objectives of privacy-preserving algorithms for SM data is to mask the difference between consecutive power measurements, called *features*, which NILM algorithms exploit to identify appliances' switch-on/off events [23]. Thus, it is possible to evaluate the performance of an energy management policy against such algorithms by computing the number of features present in the grid load [48]. We classify as features those differences that are larger or equal to 50 W, which represent a typical household electricity consumption of lights. We plot the number of features with respect to H_F and H_P in Figure 6.8. As expected, a larger value of H_F leads to a reduction in the number of features in the grid load (see Figure 6.8a); however, this does not hold with increasing H_P (see Figure 6.8b), which does not seem to have a big impact on the number of features.

(A) $\alpha = 0$, $H_P = 2h$, no energy selling.(B) $\alpha = 0$, $H_P = 2h$, energy selling.(C) $\alpha = 1$, $H_P = 2h$, no energy selling.(D) $\alpha = 1$, $H_P = 2h$, energy selling.(E) $\alpha = 1$, $H_F = 2h$, no energy selling.(F) $\alpha = 1$, $H_F = 2h$, energy selling.FIGURE 6.9: Power spectra versus H_F and H_P .

Another way of assessing the performance of privacy-preserving algorithms is by analyzing the power spectrum of the resulting grid load. In fact, the higher-frequency components of the grid load spectrum correspond typically to more sensitive information about a user's power consumption compared to the lower-frequency components [95]. Figure 6.9 shows the grid load spectra corresponding to using different values of H_F and H_P when energy selling is allowed or not. Larger values of H_F lead to better suppression of higher-frequency components when $\alpha = 1$ (see Figures 6.9c and 6.9d), whereas for $\alpha = 0$ even additional high-frequency components are introduced (see Figures 6.9a and 6.9b). When $\alpha = 1$, increasing H_P also attenuates the higher-frequency components (see Figures 6.9e and 6.9f), but less markedly compared to increasing H_F . As the spectral analysis of the grid load captures better the information leaked, in the following section we consider a privacy-preserving approach whose aim is to filter out directly the higher-frequency components of the user load.

6.5 Target Load as Filtered User Load

When the target load is set to a constant value, one can consider this as the DC component of the Fourier transform of the user load profile. If the grid load can be maintained at the average value of the user load at all times, this is equivalent to filtering out all the positive frequency components of the user load profile. However, as shown in Section 6.4, this is not always possible due to the RB capacity and power constraints, and the information leakage is measured as average squared error distance from this constant DC component. In this section we consider a more general target load profile, obtained by low-pass filtering the user load, which is equivalent to removing only the high-frequency variations. The motivation for this is two-fold. Firstly, the EMU is able to better approximate the target load profile by keeping the low-frequency components; and secondly, the high-frequency components are the ones that leak more information about user behavior. Low-frequency devices are those that typically have continuous periodic operation, e.g., the fridge, and are not particularly privacy sensitive. We would like to remark that, differently from the previous section, here W is not an optimization variable but it is determined based only on the user load, and

prior to solving the optimization problem. The optimization problem is expressed as

$$\min_{Y_t^{\overline{t+H_F}}} \alpha \sum_{\tau=t}^{\overline{t+H_F}} (Y_\tau - W_\tau)^2 + (1 - \alpha) \sum_{\tau=t}^{\overline{t+H_F}} Y_\tau C_\tau, \quad (6.20)$$

where $W_t, W_{t+1}, \dots, W_{\overline{t+H_F}}$ are obtained as low-pass filtered versions of the user load, subject to the same constraints of the constant target scenario, i.e., (6.1)-(6.2) and (5.4)-(5.5). More specifically, the target load at time t , W_t , is computed as follows. The EMU considers only the user load within the time window $[\overline{t-H_P}, \overline{t+H_F}]$, which, by hypothesis, is assumed to be known. It computes the spectral representation of the user load in this time window by means of discrete Fourier transform, which is then low-pass filtered according to a predefined cut-off frequency. Finally, the inverse transform provides the target load profile W_t . We note that, although H_P does not appear explicitly in Eq. (6.20), the target load computed at time t is determined by low-pass filtering the user load within the window $[t-H_P, t+H_F]$ to prevent the target load from varying dramatically over different TSs. When $\alpha = 0$, Eq. (6.20) reduces to the same linear programming problem in (6.6). The optimal solutions to Eq. (6.20) can be characterized by following the same steps of Section 6.4, apart from W_t , which here is not an optimization variable. The optimal solutions are given in (6.12) and (6.17) for the scenarios where selling energy is allowed and not allowed, respectively. Figure 6.10 compares the SHM and LHM, showing that the SHM generates profiles that are smooth and similar to that of the LHM, despite relying only the knowledge of 2 hours of future consumption.

6.6 A More Practical Energy Management Policy

In the previous sections it is assumed that the EMU solves the optimization problem at each TS. However, in practice it may not be feasible to obtain the future predictions at each TS, and may be impractical to compute the target profile so often. Thus, in this section we consider a more practical energy management policy where the optimization problem is solved once every T_S TSs. The optimization problems at time t are still given by (6.5) and (6.20) for the constant and filtered target load scenarios,




FIGURE 6.10: Filtered target load scenario with cut-off frequency of 0.1mHz, $\alpha = 1$, and $H_F = H_P = 2h$.




FIGURE 6.11: Practical energy management policy for $\alpha = 0.5$, $H_F = H_P = 2h$, $T_S = 1h$, and cut-off frequency set to 0.1mHz.

respectively, such that the sequences $[Y_t^{t+T_S}]$ and $[W_t^{t+T_S}]$ are obtained at time t on the basis of the available information for TSs $[t - H_P, t + H_F]$, where $H_F \geq T_S$.

Figure 6.11 shows the practical energy management policy for a constant and a filtered target load when energy selling is not allowed. The practical energy management policy creates a piecewise target and grid load, resembling the piecewise target load profile approach studied in [52] and presented in Chapter 5. Due to the discontinuities induced in the grid load profile, spikes at high frequencies may appear in the spectrum




FIGURE 6.12: Filtered target load scenario with cut-off frequency of 0.1mHz, $\alpha = 1$, and no energy selling. For the SHM and the practical energy management policy we set $H_P = 2h$ and $T_S = 1h$.




FIGURE 6.13: The privacy-cost trade-off, for $H_P = H_F = 2h$, $T_S = 1h$, energy selling allowed, and constant target load.

of the grid load produced by this strategy, leading to a higher privacy loss. When $\alpha = 1$, Figure 6.12 shows that the practical energy management policy reaches virtually the same performance of the SHM, despite computing the grid load six times less often than the LHM.

Finally, in Figure 6.13 we present the privacy-cost trade-offs for the various scenarios when energy selling is allowed and for the constant target load scenario, which highlights the loss in performance due to the limited amount of information available to

the EMU.

6.7 Conclusions

In this chapter we have studied the joint optimization of privacy and cost for an SM system equipped with an RB. Privacy is measured via the mean squared-error between the SM measurements and a target load profile, which is set to be either a constant function or a low-pass filtered version of the user load. We assume that only partial information about the user's future electricity consumption and electricity cost is known to the EMU, and we cast the joint privacy and cost optimization as a model predictive control problem. The scenario in which the user is allowed to sell excess energy to the UP is also studied, which is shown to achieve a better privacy-cost trade-off. The optimal solutions for the constant and filtered target load profiles have been characterized, highlighting their water-filling interpretation. Detailed numerical simulations have been presented, alternative privacy measures have been discussed, and the privacy-cost trade-off for the various scenarios has been characterized.

Chapter 7

Conclusions

Privacy and “the right to be let alone” are an individual’s inalienable fundamental rights, which are safeguarded in many national constitutions worldwide. Moreover, in Europe, the General Data Protection Regulation (GDPR) [96], which will be enforced from 25 May 2018, will set even more stringent requirements for every technology or device that collects and processes customer data, hence including SMs. For these reasons, addressing the SM privacy problem is necessary for the full development of the SG concept. More generally, considering the growing privacy concerns of the consumers for SMs, as well as for many other emerging technologies [97], a critical growth in SM adoption and other SG technologies will occur only when consumers are given full control of their privacy, and they feel they have clear and honest information on how their data is being used. Only then, consumer resistance can be overcome and their trust will be assured, thus paving the way to a more fertile and fair ground for new products and increased innovation in this domain.

UPs and their partners, including governments, may be too keen on collecting users’ data indiscriminately, and less incentivized to develop privacy-enhancing technologies. Therefore, legislators, public commissions, consumer advocacy groups and researchers have an important role in tackling the SM privacy problem, and preventing the SM data from being gathered indiscriminately and sold to third parties without an explicit user consent, or even passed to government intelligence agencies for mass surveillance scopes. An example of such initiatives is the enforcement of the GDPR in Europe. However, given that such a legal framework is still mostly lacking or not yet fully developed, it becomes imperative to push forward the concept of privacy by design, according to which privacy should be “designed-in” to new products and services, rather than considered only after user complaints and regulatory impositions. This is because a wider range of options are available during the design stage as compared

to modifying the product following a privacy incident or a user complaint. Achieving privacy by design is the ultimate goal of the techniques proposed in this dissertation.

In this dissertation we have focused exclusively on techniques that adopt physical resources, such as RESs and RBs, to provide privacy to users. The main motivation and benefit of these techniques is that they do not undermine the benefits of the SG concept. Each of the outlined techniques has its unique advantages and disadvantages, and focuses on a particular aspect of privacy.

In Chapter 2 we have presented an extensive overview of the main privacy preserving techniques within the UDS family, comparing their performance and analyzing their most significant results.

In Chapter 3 we have studied fundamental bounds to the SM privacy problem, by considering the availability of an RB and an RES, and by measuring privacy as the mutual information between the user and the grid load sequences. As described in Chapter 2, mutual information is only one of the many possible ways in which privacy can be defined. We have characterized single-letter closed-form expressions for the minimum information leakage rate in the limit scenarios of infinite and zero battery capacities. These asymptotic scenarios can be considered as a benchmark for comparing more practical privacy-preserving strategies which adopt a finite-capacity battery. In the same chapter we have also proposed a suboptimal parameterized energy management policy for a finite-capacity battery, and compared it to the limit scenarios.

In Chapter 4 we have targeted specifically the finite-capacity battery scenario, and measured privacy via the mutual information rate between the consumer's power consumption and the renewable energy generation process, and the power drawn from the grid, where the latter is known by the UP through the SM readings, and the former two are to be kept private. We have characterized the information leakage as an additive quantity over time, enabling us to cast the problem as an MDP, which can be solved by using dynamic programming techniques. The worst-case scenario where the UP knows the realizations of the renewable energy process is also considered.

In Chapter 5 we have jointly optimized privacy and cost of energy. As privacy measure we have adopted the mean squared error between the grid load and a target

load profile, and provided a water-filling interpretation for the optimal solution. As target load profile we have adopted a piecewise constant function, which sets a fixed value for each price period. Our proposed privacy-preserving method has been tested by using real SM data, and the optimal privacy-cost trade-off has been characterized.

In Chapter 6 we have extended the analysis of Chapter 5 by considering a more realistic scenario whereby the consumer's future consumption profile is only partially known to the EMU, in a moving horizon manner. We have introduced a new target load profile, computed as a low-pass filtered version of the user load, and proposed a practical energy management policy which performs the optimization less frequently. We have characterized the privacy-cost trade-off for the various techniques analyzed.

This thesis can be ideally divided into two parts. Firstly, in Chapters 3 and 4 we adopt a theoretical measure for privacy, i.e., mutual information, and, for some specific scenarios, we characterize the fundamental limits of privacy in terms of single-letter closed-form mathematical expressions. We begin our analysis in Chapter 3 by considering limit scenarios, i.e., infinite and zero capacity RBs, whereas in Chapter 4 we analyze the more practical case of a finite-capacity RB. Secondly, in Chapters 5 and 6 we adopt more empirical privacy measures, i.e., the squared distance between the grid load and a target load profile, and the number of features in the grid load, which we validate against real SM data and ToU tariffs. Although rather different in the way they approach the problem, these two parts are in fact complementary and part of the same effort to identify the best privacy measure. In fact, while the approaches studied in Chapters 3 and 4 may seem to be designed for very narrow settings, e.g., i.i.d. or Markovian user loads, they still provide an intuition and interesting upper and lower bounds on the performance of more practical scenarios. On the contrary, the approaches studied in Chapters 5 and 6 may seem to be more interesting because they are validated against real data; however, differently from the approaches in the earlier chapters, they lack a theoretically sound interpretation for privacy, and it is harder to derive general conclusions in terms of privacy assurances from these. We expect that future research will fill the gap between more empirical and more practical measures, with the aim of characterizing the most effective privacy measure that could be implemented in real-world systems.

In the context of SM privacy, UDS-based methods manipulate a physical quantity, energy, to ensure privacy for users. This entails that physical constraints, such as those related to an RB or an RES, play a crucial role in finding the optimal privacy-preserving strategy. We expect that the techniques developed for enhancing SM privacy can prove useful in other privacy-sensitive settings, in which physical quantities are involved, such as gas and water meters, or location privacy.

Research Challenges

In this thesis we have studied several problems related to SM privacy. However, there is still a number of open research questions that need to be properly addressed before privacy by design can become a reality in SM systems.

Firstly, a generic privacy measure, or a combination of various measures, must be determined and adopted in order to formally quantify the loss of privacy, in the same way a user's electricity bill is computed. Such a measure should be device-independent, and enable the comparison of various privacy-preserving strategies. It is also necessary to understand the implications of the various privacy measures on the grid load. From this point of view, theoretical measures may be preferable due to their abstract and fundamental nature, i.e., they are independent of any assumptions on the attacker's algorithms. However, their relevance in real-world scenarios must be assessed further, and, if necessary, valid suboptimal privacy measures or algorithms should be put forward and standardized as a proxy for more rigorous privacy assurances.

Another important goal is to give consumers as much flexibility as possible in setting their desired level of privacy, trading off privacy with the cost of electricity. We have considered this possibility in Chapters 5 and 6, where we studied the weighted joint optimization of privacy and cost. Along these lines, we think that future research should focus on allowing to set different privacy requirements for different devices, as users may value the information about the usage of certain devices more sensitive compared to others. This may happen because certain devices are naturally more correlated to the user's activities or presence at home, such as the use of a kettle, a microwave or an oven, or because a user may decide to hide the usage of a certain

appliance for personal reasons. Moreover, future research should characterize privacy in more complicated scenarios, in which multiple RBs and RESs are considered.

In the near future, a wider use of electric vehicles will also bring additional complications to the SM privacy problem, as mobility patterns may be inferred by analyzing the charging and discharging events. This problem can be tackled by load shifting, which is expected to play an important role in jointly optimizing electricity cost and privacy. Load shifting, as well as other privacy-preserving techniques introduced here, will be more accurate and relevant thanks to the development of reliable prediction techniques for future electricity consumption, e.g., by using machine learning techniques. The proliferation of various energy-hungry “smart devices” will complicate the problem further and overburden the RBs even more.

Finally, the use of shared physical resources should also be investigated in more depth, as cities are becoming more and more dense and users may want to team up to install storage devices or energy generators that are still rather costly. In cities, solar panels or mini wind turbines may be installed on the roof of blocks of flats and RBs may be put in communal areas, and these resources can be used jointly by all the users in a building. Such resource sharing models make the privacy problem even more complicated and challenging, and might call for a game theoretic formulation of the problem.

In conclusion, despite the considerable efforts put into the development of SM privacy-preserving techniques, the full extent of the privacy problem in SMs is far from being completely understood, and a unified and coherent vision for SM privacy, just like in many other domains, is still missing. However, we hope that the research work presented in this Ph.D. thesis has contributed towards the solution of the SM privacy problem, as well as to further encourage research and development in this area so that remaining open issues will be solved and the full potential of SMs will be unleashed.

Bibliography

- [1] Pecan Street Inc. Dataport. [Online]. Available: <https://dataport.cloud/>
- [2] Powervault. Technical specifications. [Online]. Available: <http://www.powervault.co.uk/technical/technical-specifications/>
- [3] Tesla Powerwall. [Online]. Available: https://www.tesla.com/en_GB/powerwall
- [4] The University of Sheffield - Sheffield Solar. Microgen database. [Online]. Available: <https://microgen-database.sheffield.ac.uk/>
- [5] U.S. Energy Information Administration, “International energy outlook,” May 2016.
- [6] BP, “BP energy outlook,” 2016.
- [7] Y. Mo, T.-H. Kim, K. Brancik, D. Dickinson, H. Lee, A. Perrig, and B. Sinopoli, “Cyber-physical security of a smart grid infrastructure,” *Proc. IEEE*, vol. 100, no. 1, pp. 195–209, Jan. 2012.
- [8] Itron. Centron C12.19, C1SD, C1ST, C1SL. [Online]. Available: <https://www.itron.com/na/technology/product-services-catalog/products/0/7/5/centron>
- [9] Elster. REX2 meter. [Online]. Available: https://www.elstersolutions.com/en/product-details-na/826/en/REX2_meter
- [10] Kamstrup. Omnipower. [Online]. Available: <https://www.kamstrup.com/en-us/products-and-solutions/smart-grid/electricity-meters>
- [11] Enel. Open meter. [Online]. Available: https://www.enel.com/content/dam/enel-com/pressrelease/porting_pressrelease/1666038-1_PDF-1.pdf
- [12] European Regulators Group for Electricity and Gas, “Final guidelines of good practice on regulatory aspects of smart metering for electricity and gas,” Feb. 2011.

- [13] Department of Energy and Climate Change, “Smart metering equipment technical specifications: Version 2,” Nov. 2014.
- [14] Smart Meter Texas. [Online]. Available: https://www.smartmetertexas.com/CAP/public/home/home_about_us.html
- [15] R. Segovia and M. Sánchez, “A joint contribution of DG ENER and DG INFSO towards the Digital Agenda, Action 73: Set of common functional requirements of the smart meter,” 2011. [Online]. Available: https://ec.europa.eu/energy/sites/ener/files/documents/2011_10_smart_meter_funtionalities_report.pdf
- [16] European Union, “Directive 2009/72/EC of the European parliament and of the council of 13 July 2009 concerning common rules for the internal market in electricity and repealing directive 2003/54/EC,” *Official J. European Union*, vol. 52, no. L211, pp. 55–93, Aug. 2009.
- [17] MarketsandMarkets, “Smart meters market by type (electric, water, and gas), application (commercial, residential, and industrial), technology (automatic meter reading and advanced metering infrastructure), and by region - global forecasts to 2022,” Apr. 2017.
- [18] Transparency Market Research, “Smart grid data analytics market - global industry analysis, size, share, growth, trends and forecast 2015 - 2022,” Nov. 2015.
- [19] Navigant Research, “Market data: Smart meters,” Sep. 2016.
- [20] G. Hart, “Prototype nonintrusive appliance load monitor,” *MIT Energy Laboratory Tech. Rep., and Electric Power Research Inst. Tech. Rep.*, Sep. 1985.
- [21] —, “Nonintrusive appliance load monitoring,” *Proc. IEEE*, vol. 80, no. 12, pp. 1870–1891, Dec. 1992.
- [22] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin, “Private memoirs of a smart meter,” in *Proc. ACM Workshop Embedded Sensing Syst. for Energy-Efficiency in Building*, Zurich, Switzerland, Nov. 2010, pp. 61–66.

- [23] A. Prudenzi, “A neuron nets based procedure for identifying domestic appliances pattern-of-use from energy recordings at meter panel,” in *Proc. IEEE Power Eng. Soc. Winter Meeting*, vol. 2, New York, NY, USA, Jan. 2002, pp. 941–946.
- [24] E. Quinn, “Privacy and the new energy infrastructure,” *Social Sci. Research Network*, Feb. 2009.
- [25] I. Rouf, H. Mustafa, M. Xu, W. Xu, R. Miller, and M. Gruteser, “Neighborhood watch: Security and privacy analysis of automatic meter reading systems,” in *Proc. ACM Conf. Comput. and Commun. Security*, Raleigh, NC, USA, Oct. 2012, pp. 462–473.
- [26] C. Cuijpers and B.-J. Koops, “Smart metering and privacy in Europe: Lessons from the Dutch case,” *European Data Protection: Coming of Age*, pp. 269–293, Feb. 2012.
- [27] J. Loeb, “Smart meters: what would it take to stop the national rollout juggernaut?” *IET, E&T*, vol. 12, no. 5, pp. 32–34, May 2017.
- [28] Y. Kim, E. Ngai, and M. Srivastava, “Cooperative state estimation for preserving privacy of user behaviors in smart grid,” in *Proc. IEEE Int. Conf. Smart Grid Commun.*, Brussels, Belgium, Oct. 2011, pp. 178–183.
- [29] J.-M. Bohli, C. Sorge, and O. Ugus, “A privacy model for smart metering,” in *Proc. IEEE Int. Conf. Commun.*, Cape Town, South Africa, May 2010, pp. 1–5.
- [30] M. Backes and S. Meiser, “Differentially private smart metering with battery recharging,” in *Proc. Int. Workshop Data Privacy Manage. and Autonomous Spontaneous Security*, Egham, UK, Sep. 2014, pp. 194–212.
- [31] L. Sankar, S. Rajagopalan, S. Mohajer, and H. V. Poor, “Smart meter privacy: A theoretical framework,” *IEEE Trans. Smart Grid*, vol. 4, no. 2, pp. 837–846, Jun. 2013.
- [32] F. D. Garcia and B. Jacobs, “Privacy-friendly energy-metering via homomorphic encryption,” in *Proc. Int. Conf. Security and Trust Manage.*, Athens, Greece, Sep. 2010, pp. 226–238.

- [33] F. Li, B. Luo, and P. Liu, "Secure and privacy-preserving information aggregation for smart grids," *Int. J. of Security and Networks*, vol. 6, no. 1, pp. 28–39, Apr. 2011.
- [34] R. Petric, "A privacy-preserving concept for smart grids," in *Sicherheit in vernetzten Systemen:18. DFN Workshop*. Books on Demand GmbH, 2010, pp. B1–B14.
- [35] C. Efthymiou and G. Kalogridis, "Smart grid privacy via anonymization of smart metering data," in *Proc. IEEE Int. Conf. Smart Grid Commun.*, Gaithersburg, MD, USA, Oct. 2010, pp. 238–243.
- [36] A. Cárdenas, S. Amin, and G. A. Schwartz, "Privacy-aware sampling for residential demand response programs," in *Proc. ACM Int. Conf. High Confidence Networked Syst.*, Beijing, China, Apr. 2012.
- [37] Sunverge. Ac coupled system spec sheet. [Online]. Available: <http://www.sunverge.com/energy-management/>
- [38] The sonnenBatterie. [Online]. Available: <https://sonnen-batterie.com/en-us/sonnenbatterie>
- [39] LG Chem. Ess battery. [Online]. Available: <http://www.lgchem.com/global/ess/ess/product-detail-PDEC0001>
- [40] Panasonic. Residential storage battery system lj-sk84a. [Online]. Available: <http://www.panasonic.com/au/consumer/energy-solutions/residential-storage-battery-system/lj-sk56a.html>
- [41] Orison. Meet the Orison family. [Online]. Available: <http://orison.energy/>
- [42] Simpliphi Power. Phi3.4TM smart-tech battery. [Online]. Available: <http://simpliphipower.com/product/phi3-4-smart-tech-battery/>
- [43] J.-P. Zimmermann, M. Evans, J. Griggs, N. King, L. Harding, P. Roberts, and C. Evans, "Intertek Report R66141, Household Electricity Survey, A study of domestic electrical product usage," Intertek Testing & Certification Ltd, Tech. Rep., May 2012.

- [44] S. N. A. U. Nambi, A. R. Lua, and R. V. Prasad, "LocED: Location-aware energy disaggregation framework," in *Proc. ACM Int. Conf. Embedded Syst. for Energy-Efficient Built Environments*, Nov. 2015, pp. 45–54.
- [45] D. Dheeru and E. Karra Taniskidou, "UCI machine learning repository," 2017. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [46] European Environment Agency. Total electricity consumption - outlook from IEA. [Online]. Available: <https://www.eea.europa.eu/data-and-maps/indicators/total-electricity-consumption-outlook-from-iea/total-electricity-consumption-outlook-from-1>
- [47] Y. H. Liu, S. H. Lee, and A. Khisti, "Information-theoretic privacy in smart metering systems using cascaded rechargeable batteries," *IEEE Signal Process. Lett.*, vol. 24, no. 3, pp. 314–318, Mar. 2017.
- [48] L. Yang, X. Chen, J. Zhang, and H. V. Poor, "Cost-effective and privacy-preserving energy management for smart meters," *IEEE Trans. Smart Grid*, vol. 6, no. 1, pp. 486–495, Jan. 2015.
- [49] A. Zoha, A. Gluhak, M. A. Imran, and S. Rajasegarar, "Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey," *Sensors*, vol. 12, no. 12, pp. 16 838–16 866, Dec. 2012.
- [50] O. Tan, J. Gómez-Vilardebó, and D. Gündüz, "Privacy-cost trade-offs in demand-side management with storage," *IEEE Trans. Inform. Forens. Security*, vol. 12, no. 6, pp. 1458–1469, Jun. 2017.
- [51] A. Lever, D. Sanders, N. Lehmann, M. Ravishankar, M. Ashcroft, G. Strbac, M. Aunedi, F. Teng, and D. Pudijanto, "Can storage help reduce the cost of a future UK electricity system?" *Carbon Trust*, Feb. 2016.
- [52] G. Giaconi, D. Gündüz, and H. V. Poor, "Optimal demand-side management for joint privacy-cost optimization with energy storage," in *Proc. IEEE Int. Conf. Smart Grid Commun.*, Dresden, Germany, Oct. 2017, pp. 265–270.

- [53] J. Kelly and W. Knottenbelt, “The UK-DALE dataset, domestic appliance-level electricity demand and whole-house demand from five UK homes,” *Scientific Data*, vol. 2, no. 150007, Mar. 2015.
- [54] J. Wu, J. Liu, X. S. Hu, and Y. Shi, “Privacy protection via appliance scheduling in smart homes,” in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design*, Austin, TX, USA, Nov. 2016, pp. 1–6.
- [55] Z. Chen and L. Wu, “Residential appliance DR energy management with electric privacy protection by online stochastic optimization,” *IEEE Trans. Smart Grid*, vol. 4, no. 4, pp. 1861–1869, Dec. 2013.
- [56] G. Kalogridis, C. Efthymiou, S. Denic, T. Lewis, and R. Cepeda, “Privacy for smart meters: Towards undetectable appliance load signatures,” in *Proc. IEEE Int. Conf. Smart Grid Commun.*, Gaithersburg, MD, USA, Oct. 2010, pp. 232–237.
- [57] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York, NY, USA: Wiley-Interscience, 1991.
- [58] S. McLaughlin, P. McDaniel, and W. Aiello, “Protecting consumer privacy from electric load monitoring,” in *Proc. ACM Conf. Comput. and Commun. Security*, Chicago, IL, USA, Oct. 2011, pp. 87–98.
- [59] J. Yao and P. Venkitasubramaniam, “The privacy analysis of battery control mechanisms in demand response: Revealing state approach and rate distortion bounds,” *IEEE Trans. Smart Grid*, vol. 6, no. 5, pp. 2417–2425, Sep. 2015.
- [60] W. Yang, N. Li, Y. Qi, W. Qardaji, S. McLaughlin, and P. McDaniel, “Minimizing private data disclosures in the smart grid,” in *Proc. ACM Conf. Comput. and Commun. Security*, Raleigh, NC, USA, Oct. 2012, pp. 415–427.
- [61] F. Farokhi and H. Sandberg, “Fisher information as a measure of privacy: Preserving privacy of households with smart meters using batteries,” *IEEE Trans. Smart Grid*, Feb. 2017.
- [62] Z. Li and T. J. Oechetering, “Privacy on hypothesis testing in smart grids,” in *Proc. IEEE Inform. Theory Workshop*, Jerusalem, Israel, Oct. 2015, pp. 337–341.

- [63] K. Qian, C. Zhou, M. Allan, and Y. Yuan, “Modeling of load demand due to EV battery charging in distribution systems,” *IEEE Trans. Power Syst.*, vol. 26, no. 2, pp. 802–810, May 2011.
- [64] P. A. Leicester, C. I. Goodier, and P. N. Rowley, “Probabilistic analysis of solar photovoltaic self-consumption using bayesian network models,” *IET Renewable Power Generation*, vol. 10, no. 4, pp. 448–455, Mar. 2016.
- [65] W. Labeeuw and G. Deconinck, “Residential electrical load model based on mixture model clustering and Markov models,” *IEEE Trans. Ind. Informat.*, vol. 9, no. 3, pp. 1561–1569, Aug. 2013.
- [66] D. P. Bertsekas, *Dynamic Programming and Optimal Control, Vol. II*, 3rd ed. Belmont, MA, USA: Athena Scientific, 2007.
- [67] K. Chatzikokolakis, C. Palamidessi, and P. Panangaden, “Anonymity protocols as noisy channels,” *Information and Computation*, vol. 206, no. 2, pp. 378 – 401, 2008.
- [68] B. Köpf and D. Basin, “An information-theoretic model for adaptive side-channel attacks,” in *Proc. ACM Conf. Comput. and Commun. Security*, Alexandria, VA, USA, Oct. 2007, pp. 286–296.
- [69] D. Clark, S. Hunt, and P. Malacaria, “Quantitative analysis of the leakage of confidential data,” *Electron. Notes in Theoretical Comput. Sci.*, vol. 59, no. 3, pp. 238–251, Nov. 2002.
- [70] D. Gündüz and J. Gómez-Vilardebó, “Smart meter privacy in the presence of an alternative energy source,” in *Proc. IEEE Int. Conf. Commun.*, Budapest, Hungary, Jun. 2013, pp. 2027–2031.
- [71] J. Gómez-Vilardebó and D. Gündüz, “Privacy of smart meter systems with an alternative energy source,” in *Proc. IEEE Int. Symp. Inform. Theory*, Istanbul, Turkey, Jul. 2013, pp. 2572–2576.
- [72] J. Gómez-Vilardebó and D. Gündüz, “Smart meter privacy for multiple users in the presence of an alternative energy source,” *IEEE Trans. Inform. Forens. Security*, vol. 10, no. 1, pp. 132–141, Jan. 2015.

- [73] S. Li, A. Khisti, and A. Mahajan, “Information-theoretic privacy for smart metering systems with a rechargeable battery,” *IEEE Trans. Inform. Theory*, Feb. 2018.
- [74] —, “Structure of optimal privacy-preserving policies in smart-metered systems with a rechargeable battery,” in *Proc. IEEE Int. Workshop Signal Process. Advances in Wireless Commun.*, Jun. 2015, pp. 375–379.
- [75] —, “Privacy preserving rechargeable battery policies for smart metering systems,” in *Proc. Int. Zurich Seminar Commun.*, Zurich, Switzerland, Mar. 2016, pp. 121–124.
- [76] —, “Privacy-optimal strategies for smart metering systems with a rechargeable battery,” in *Proc. American Control Conf.*, Boston, MA, USA, Jul. 2016, pp. 2080–2085.
- [77] G. Giaconi, D. Gündüz, and H. V. Poor, “Smart meter privacy with an energy harvesting device and instantaneous power constraints,” in *Proc. IEEE Int. Conf. Commun.*, London, UK, Jun. 2015, pp. 7216–7221.
- [78] —, “Smart meter privacy with renewable energy and an energy storage device,” *IEEE Trans. Inform. Forens. Security*, vol. 13, no. 1, pp. 129–142, Jan. 2018.
- [79] G. Giaconi and D. Gündüz, “Smart meter privacy with renewable energy and a finite capacity battery,” in *Proc. IEEE Int. Workshop Signal Process. Advances in Wireless Commun.*, Edinburgh, UK, Jul. 2016, pp. 1–5.
- [80] Z. Li, T. J. Oechtering, and D. Gündüz, “Smart meter privacy based on adversarial hypothesis testing,” in *Proc. IEEE Int. Symp. Inform. Theory*, Aachen, Germany, Jun. 2017, pp. 774–778.
- [81] D.-M. Arnold, H.-A. Loeliger, P. Vontobel, A. Kavcic, and W. Zeng, “Simulation-based computation of information rates for channels with memory,” *IEEE Trans. Inform. Theory*, vol. 52, no. 8, pp. 3498–3508, Aug. 2006.
- [82] D. Varodayan and A. Khisti, “Smart meter privacy using a rechargeable battery: Minimizing the rate of information leakage,” in *Proc. IEEE Int. Conf. Acoust., Speech and Signal Process.*, Prague, Czech Republic, May 2011, pp. 1932–1935.

- [83] O. Tan, D. Gündüz, and H. V. Poor, “Increasing smart meter privacy through energy harvesting and storage devices,” *IEEE J. Sel. Areas in Commun.*, vol. 31, no. 7, pp. 1331–1341, Jul. 2013.
- [84] J. X. Chin, T. T. D. Rubira, and G. Hug, “Privacy-protecting energy management unit through model-distribution predictive control,” *IEEE Trans. Smart Grid*, vol. 8, no. 6, pp. 3084–3093, Nov. 2017.
- [85] J. Koo, X. Lin, and S. Bagchi, “Privatus: Wallet-friendly privacy protection for smart meters,” in *Proc. European Symp. Research in Comput. Security*, Pisa, Italy, Sep. 2012, pp. 343–360.
- [86] R. Blahut, “Computation of channel capacity and rate-distortion functions,” *IEEE Trans. Inform. Theory*, vol. 18, no. 4, pp. 460–473, Jul. 1972.
- [87] M. P. Deisenroth, G. Neumann, and J. Peters, “A survey on policy search for robotics,” *Foundations and Trends in Robotics*, vol. 2, pp. 1–142, Aug. 2013.
- [88] M. Grant and S. Boyd, “CVX: Matlab software for disciplined convex programming, version 2.1,” Mar. 2014. [Online]. Available: <http://cvxr.com/cvx>
- [89] T. Berger, *Rate Distortion Theory: A Mathematical Basis for Data Compression*. Englewood Cliffs, NJ: Prentice-Hall, 1971.
- [90] S. Tatikonda and S. Mitter, “The capacity of channels with feedback,” *IEEE Trans. Inform. Theory*, vol. 55, no. 1, pp. 323–349, Jan. 2009.
- [91] L. Yang, X. Chen, J. Zhang, and H. V. Poor, “Optimal privacy-preserving energy management for smart meters,” in *Proc. IEEE Int. Conf. Comput. Commun.*, Toronto, ON, Canada, Apr. 2014, pp. 513–521.
- [92] G. Narzisi. (2008, Jan.) Multi-objective optimization, a quick introduction. [Online]. Available: <https://cims.nyu.edu/~gn387/glp/lec1.pdf>
- [93] A. M. Payne, R. D. Duke, and R. H. Williams, “The impact of net metering on the residential rooftop PV market,” in *Proc. IEEE Photovoltaic Specialists Conf.*, Anchorage, AK, USA, Sep. 2000, pp. 1391–1394.

-
- [94] Green Energy UK. Tide. [Online]. Available: <http://www.greenenergyuk.com/Tide>
- [95] D. Engel and G. Eibl, “Wavelet-based multiresolution smart meter privacy,” *IEEE Trans. Smart Grid*, vol. 8, no. 4, pp. 1710–1721, Jul. 2017.
- [96] The European Parliament and the Council of the European Union, “Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 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),” *Official J. of the European Union*, May 2016.
- [97] PwC, “Consumer intelligence series: Protect.me, an in-depth look at what consumers want, what worries them, and how companies can earn their trust - and their business,” Sep. 2017.
- [98] R. G. Gallager, *Discrete Stochastic Processes*. Springer US, 1996.

Appendix A

Proofs for Chapter 3

A.1 Proof of Lemma 3.5

Proof. During the hiding phase, the random variable $Q = E - X + Y^*$ is i.i.d., as E and X are i.i.d and Y^* is generated from X through a memoryless policy. Q can assume both positive and negative values with positive probability. The stochastic process

$$S_t = Q_1 + Q_2 + \dots + Q_t, \quad \forall t, \quad (\text{A.1})$$

is a random walk based on Q that moves along the axis representing the level of energy in the battery. Since by hypothesis $\mathbb{E}[E] = \bar{P}_E > \mathbb{E}[X - Y^*]$, then $\mathbb{E}[Q] = \mathbb{E}[E - X + Y^*] > 0$, meaning that the random walk S_t has a positive drift, i.e., as $t \rightarrow \infty$, S_t drift towards the positive values of the axis representing the level of energy in the battery.

By the law of large numbers, when $s(n) \rightarrow \infty$ the amount of energy stored in the battery at the end of the *storage phase* is $s(n)\bar{P}_E$, almost surely. Let $\alpha \triangleq -s(n)\bar{P}_E$. When $s(n) \rightarrow \infty$, $\alpha \rightarrow -\infty$. At $s(n) + 1$, when the *hiding phase* begins, the energy in the battery is used according to the optimal privacy-preserving policy $p_{Y|X}^*$ and the random walk state is $S_1 = Q_1 = E_1 - X_1 + Y_1^*$. For any t , $s(n)\bar{P}_E + S_t$ represents the level of energy in the battery at time t . Our objective is to prove that the battery is never emptied, i.e., that the probability of crossing the threshold α for any time t is zero:

$$\Pr\{S_t \leq \alpha\} = 0, \quad \forall t. \quad (\text{A.2})$$

This scenario is represented in Figure A.1. We recall a corollary of Wald's Identity [98, Chapter 7.5, Corollary 2], which is applied to find exponential bounds on the




FIGURE A.1: The evolution of the level of energy in the battery is represented by a random walk that starts at the beginning of the hiding phase and has a drift towards the positive direction of the axis representing the level of energy in the battery. We want to guarantee that the threshold α is never crossed by the random walk.

probability of threshold crossing. In particular, the corollary states that if we consider Q as having a finite moment-generating function $\gamma(r) = \ln\{\mathbb{E}[\exp(rQ)]\}$ over an interval (r_-, r_+) , a negative drift $\mathbb{E}[Q] < 0$ and r^* being the positive root of $\gamma(r)$, then the probability of crossing threshold $\alpha > 0$ by the random walk $S_t = Q_1 + Q_2 + \dots + Q_t$ is

$$\Pr\{S_\tau \geq \alpha\} \leq \exp(-r^*\alpha), \quad (\text{A.3})$$

where τ is the minimum t for which the threshold α is crossed. Having a finite moment generating function means that Q must have moments of all orders and the tails of its distribution function must decay at least exponentially in q as $q \rightarrow \infty$ and $q \rightarrow -\infty$. In our specific setting, $\mathbb{E}[Q] > 0$, $\alpha < 0$, and $r^* < 0$. We can still apply Wald's identity by changing the signs of r^* and α and by considering the probability of crossing a negative threshold. Thus, we have

$$\Pr\{S_\tau \leq \alpha\} \leq \exp(-r^*\alpha), \quad (\text{A.4})$$

where $\alpha < 0$ and $r^* < 0$. When $\lim_{n \rightarrow \infty} n - s(n) = \infty$ and $\lim_{n \rightarrow \infty} s(n) = \infty$, $\alpha \rightarrow -\infty$ and $\exp(-r^*\alpha) \rightarrow 0$. Thus, we obtain

$$\lim_{n \rightarrow \infty} \Pr\{S_\tau \leq \alpha\} = 0. \quad (\text{A.5})$$

□

A.2 Proof of Lemma 3.6

Proof. Split the sequence of user and grid load symbols into the storage and hiding phases of duration $s(n)$ and $n - s(n)$, respectively and let $s(n) = o(n)$. Then, it is possible to write

$$\frac{1}{n}I(X^n; Y^n) = \frac{1}{n} \sum_{i=1}^n \left[H(X_i | X^{i-1}) - H(X_i | X^{i-1}, Y^n) \right] \quad (\text{A.6a})$$

$$\geq \frac{1}{n} \sum_{i=1}^n \left[H(X_i) - H(X_i | Y_i) \right] \quad (\text{A.6b})$$

$$= \frac{1}{n} \left[\sum_{i=1}^{s(n)} I(X_i; Y_i) + \sum_{i=s(n)+1}^n I(X_i; Y_i) \right] \quad (\text{A.6c})$$

$$= \frac{1}{n} \left\{ s(n)H(X) + [n - s(n)]\mathcal{I}^i(\infty, \hat{P}_d) \right\} \quad (\text{A.6d})$$

$$= \frac{s(n)H(X)}{n} + \frac{[n - s(n)]\mathcal{I}^i(\infty, \hat{P}_d)}{n}, \quad (\text{A.6e})$$

where (A.6b) follows because X is i.i.d. and conditioning reduces entropy; (A.6d) follows since in the first $s(n)$ time instants leakage of full information $H(X)$ takes place, while in the following $n - s(n)$ TSs private operation is assured via the optimal strategy of Theorem 3.2.

If we take the limit $n \rightarrow \infty$, since $s(n) = o(n)$ and $H(X)$ is finite, we obtain the leakage rate

$$\lim_{n \rightarrow \infty} \frac{s(n)}{n} H(X) + \frac{n - s(n)}{n} \mathcal{I}^i(\infty, \hat{P}_d) = \mathcal{I}^i(\infty, \hat{P}_d). \quad (\text{A.7})$$

□

Appendix B

Proofs for Chapter 4

B.1 Proof of Theorem 4.1

We want to prove that restricting the set of feasible policies to $\mathcal{F}' \subseteq \mathcal{F}$ does not lose optimality. We first prove the following lemma.

Lemma B.1. *For any strategy $f \in \mathcal{F}$*

$$\mathcal{I}(n, f) = \frac{1}{n} I(X^n, E^n, B_1; Y^n) \geq \frac{1}{n} \sum_{t=1}^n I(X_t, E_t, B_t; Y_t | Y^{t-1}), \quad (\text{B.1})$$

with equality if and only if $f \in \mathcal{F}'$.

Proof. We have:

$$\mathcal{I}(n, f) = \frac{1}{n} I(X^n, E^n, B_1; Y^n) \quad (\text{B.2a})$$

$$= \frac{1}{n} \sum_{t=1}^n H(Y_t | Y^{t-1}) - H(Y_t | Y^{t-1}, X^n, E^n, B_1) \quad (\text{B.2b})$$

$$= \frac{1}{n} \sum_{t=1}^n I(X^t, E^t, B_1; Y_t | Y^{t-1}) \quad (\text{B.2c})$$

$$= \frac{1}{n} \sum_{t=1}^n I(X^t, E^t, B^t; Y_t | Y^{t-1}) \quad (\text{B.2d})$$

$$\geq \frac{1}{n} \sum_{t=1}^n I(X_t, E_t, B_t; Y_t | Y^{t-1}), \quad (\text{B.2e})$$

where (B.2c) follows because $(X_{t+1}^n, E_{t+1}^n) \rightarrow (X^t, E^t, B_1, Y^{t-1}) \rightarrow Y_t$ form a Markov chain; (B.2d) follows because B^t can be deterministically obtained from the knowledge of X^t, E^t, B_1 and Y^{t-1} ; and (B.2e) follows because mutual information is decreased if terms are removed. It is easy to see that the equality in (B.2e) holds if policies $f' \in \mathcal{F}'$

are considered, since f' decide on Y_t based only on the values of the variables at TS t . $\square$

Next, we show that policies $f' \in \mathcal{F}'$, defined in Eq. (4.6), achieve the minimum information leakage rate. For ease of notation, we define $Z_t \triangleq (X_t, E_t, B_t)$.

Lemma B.2. *For any $f \in \mathcal{F}$, there exists a $f' \in \mathcal{F}'$ such that*

$$\sum_{t=1}^n I^f(Z_t; Y_t | Y^{t-1}) = \sum_{t=1}^n I^{f'}(Z_t; Y_t | Y^{t-1}). \quad (\text{B.3})$$

Proof. We prove this lemma by showing that, for any $f \in \mathcal{F}$, it is possible to produce a corresponding $f' \in \mathcal{F}'$ that achieves the same leakage. For each t and all possible (y^t, z_t) such that the output is in the feasible set, i.e., $y_t \in \bar{\mathcal{Y}}(w_t)$, let

$$f'_t(y_t | z_t, y^{t-1}) = p_{Y_t | Z_t, Y^{t-1}}^f(y_t | z_t, y^{t-1}). \quad (\text{B.4})$$

To prove the lemma we show that the joint probability distributions are equal, i.e., $p_{Z_t, Y^t}^f = p_{Z_t, Y^t}^{f'}, \forall t$. By construction of f'_t in Eq. (B.4) we have:

$$p_{Y_t | Z_t, Y^{t-1}}^f = p_{Y_t | Z_t, Y^{t-1}}^{f'}. \quad (\text{B.5})$$

Then, it only remains to prove the following relation:

$$p_{Z_t, Y^{t-1}}^f = p_{Z_t, Y^{t-1}}^{f'}. \quad (\text{B.6})$$

By using induction on t , for $t = 1$ we have $p_{Z_1}^f = p_{Z_1}^{f'}$, and we obtain

$$p_{Z_{t+1}, Y^t}^f(z_{t+1}, y^t) = \sum_{z_t} p_{Z_{t+1}, Z_t, Y^t}^f(z_{t+1}, z_t, y^t) \quad (\text{B.7a})$$

$$= \sum_{z_t} p(z_{t+1} | y^t, z_t) p_{Y_t | Z_t, Y^{t-1}}^f(y_t | z_t, y^{t-1}) p_{Z_t, Y^{t-1}}^f(z_t, y^{t-1}) \quad (\text{B.7b})$$

$$= \sum_{z_t} p(z_{t+1} | y^t, z_t) f'_t(y_t | z_t, y^{t-1}) p_{Z_t, Y^{t-1}}^{f'}(z_t, y^{t-1}) \quad (\text{B.7c})$$

$$= p_{Z_{t+1}, Y^t}^{f'}(z_{t+1}, y^t), \quad (\text{B.7d})$$

where $p(z_{t+1}|y^t, z_t) = \mathbb{1}_{b_{t+1}}\{\min\{b_t + e_t - x_t, B_{\max}\} + y_t\}q_X(x_{t+1}|x_t)q_E(e_{t+1}|e_t)$ represents the system evolution; (B.7c) follows from (B.4) and from the induction hypothesis at TS t . $\square$

By combining Lemmas B.1 and B.2, we show that for any strategy $f \in \mathcal{F}$, it is possible to find a strategy $f' \in \mathcal{F}'$, such that

$$I^f(X^n, E^n, B_1; Y^n) \geq I^{f'}(X^n, E^n, B_1; Y^n), \quad (\text{B.8})$$

and, for the strategies $f' \in \mathcal{F}'$, the right hand side of Eq. (B.8) actually coincides with the additive expression $\sum_{t=1}^n I^{f'}(X_t, E_t, B_t; Y_t|Y^{t-1})$ in Eq. (4.8).

B.2 Proof of Corollary 4.1

We first prove the following lemma.

Lemma B.3. *For any strategy $f \in \mathcal{F}$*

$$\mathcal{I}(n) = \frac{1}{n}I(X^n, E^n, B_1; Y^n) \geq \frac{1}{n} \sum_{t=1}^n I(X_{t-m+1}^t, E_{t-l+1}^t, B_t; Y_t|Y^{t-1}), \quad (\text{B.9})$$

with equality if and only if $f \in \mathcal{F}''$.

Proof. We have:

$$\mathcal{I}(n, f) = \frac{1}{n}I(X^n, E^n, B_1; Y^n) \quad (\text{B.10a})$$

$$= \frac{1}{n} \sum_{t=1}^n I(X^t, E^t, B_1; Y_t|Y^{t-1}) \quad (\text{B.10b})$$

$$= \frac{1}{n} \sum_{t=1}^n I(X^t, E^t, B^t; Y_t|Y^{t-1}) \quad (\text{B.10c})$$

$$\geq \frac{1}{n} \sum_{t=1}^n I(X_{t-m+1}^t, E_{t-l+1}^t, B_t; Y_t|Y^{t-1}), \quad (\text{B.10d})$$

where (B.10b) follows because $(X_{t+1}^n, E_{t+1}^n) \rightarrow (X^t, E^t, B_1, Y^{t-1}) \rightarrow Y_t$ form a Markov chain; (B.10c) follows because B^t can be deterministically obtained from the knowledge

of X^t , E^t , B_1 and Y^{t-1} ; and (B.10d) follows because mutual information is decreased if terms are removed. It is easy to see that the equality in (B.10d) holds if $f \in \mathcal{F}''$. $\square$

Next, we show that the policies $f'' \in \mathcal{F}''$, defined in Eq. (4.9), achieve the minimum information leakage rate. Let $R_t \triangleq (X_{t-m+1}^t, E_{t-l+1}^t, B_t)$.

Lemma B.4. *For any strategy $f \in \mathcal{F}$, there exists a strategy $f'' \in \mathcal{F}''$ that achieves the same information leakage, i.e.,*

$$\sum_{t=1}^n I^f(R_t; Y_t | Y^{t-1}) = \sum_{t=1}^n I^{f''}(R_t; Y_t | Y^{t-1}). \quad (\text{B.11})$$

Proof. We want to show that for any leakage that can be achieved by adopting the energy management policy f , the same leakage can also be achieved by adopting the energy management policy f'' . For any policy $f = (f_1, \dots, f_n) \in \mathcal{F}$, we construct a corresponding policy $f'' = (f''_1, \dots, f''_n) \in \mathcal{F}''$ as follows. For each t and all possible (y^t, r_t) such that the output is in the feasible set, i.e., $y_t \in \bar{\mathcal{Y}}(w_t)$, let

$$f''_t(y_t | r_t, y^{t-1}) = p_{Y_t | R_t, Y^{t-1}}^f(y_t | r_t, y^{t-1}). \quad (\text{B.12})$$

To prove the lemma we show that the joint probability distributions p_{R_t, Y^t}^f and $p_{R_t, Y^t}^{f''}$ are equal $\forall t$. By construction of f''_t in Eq. (B.12) we have:

$$p_{Y_t | R_t, Y^{t-1}}^f = p_{Y_t | R_t, Y^{t-1}}^{f''}. \quad (\text{B.13})$$

Then, it only remains to prove the following relation:

$$p_{R_t, Y^{t-1}}^f = p_{R_t, Y^{t-1}}^{f''}. \quad (\text{B.14})$$

By using induction on t , for $t = 1$ we have $p_{R_1}^f = p_{R_1}^{f''}$. Then, we obtain

$$\begin{aligned} p_{R_{t+1}, Y^t}^f(r_{t+1}, y^t) &= \sum_{x_{t-m}, e_{t-l}, b_t} p_{X_{t-m}^{t+1}, B_t^{t+1}, E_{t-l}^{t+1}, Y^t}^f(x_{t-m}^{t+1}, b_t^{t+1}, e_{t-l}^{t+1}, y^t) \\ &= \sum_{x_{t-m}, e_{t-l}, b_t} p(x_{t+1}, b_{t+1}, e_{t+1} | y^t, r_t) \end{aligned} \quad (\text{B.15a})$$

$$\begin{aligned}
& \times p_{Y_t|R_t, Y^{t-1}}^f(y_t|r_t, y^{t-1}) p_{R_t, Y^{t-1}}^f(r_t, y^{t-1}) \quad (\text{B.15b}) \\
& = \sum_{x_{t-m}, e_{t-l}, b_t} p(x_{t+1}, b_{t+1}, e_{t+1}|y^t, r_t) \\
& \quad \times f_t''(y_t|r_t, y^{t-1}) p_{R_t, Y^{t-1}}^{f''}(r_t, y^{t-1}) \quad (\text{B.15c})
\end{aligned}$$

$$= p_{R_{t+1}, Y^t}^{f''}(r_{t+1}, y^t), \quad (\text{B.15d})$$

where (B.15c) is due to the induction hypothesis and Eq. (B.12), and

$$\begin{aligned}
p(x_{t+1}, b_{t+1}, e_{t+1}|y^t, r_t) &= q_X^m(x_{t+1}|x_{t-m+1}^t) q_E^l(e_{t+1}|e_{t-l+1}^t) \\
&\quad \times \mathbb{1}_{b_{t+1}}\{\min\{b_t + e_t - x_t, B_{\max}\} + y_t\}, \quad (\text{B.16})
\end{aligned}$$

represents the system evolution, where q_X^m and q_E^l denote the transition probabilities of the higher order Markov chains for the user load and the renewable energy process, respectively. $\square$

By combining Lemmas B.3 and B.4, we show that for any strategy $f \in \mathcal{F}$, it is possible to characterize a strategy $f'' \in \mathcal{F}''$, such that

$$I^f(X^n, E^n, B_1; Y^n) \geq I^{f''}(X^n, E^n, B_1; Y^n), \quad (\text{B.17})$$

and, for the strategies $f'' \in \mathcal{F}''$, the right hand side of Eq. (B.17) is equivalent to the additive expression $\sum_{t=1}^n I^{f''}(X_{t-m+1}^t, E_{t-l+1}^t, B_t; Y_t|Y^{t-1})$ in Eq. (4.10).

B.3 Proof of Theorem 4.2

For the scenario in which the renewable energy process is known by the UP, we show in the following that expressing the original problem into an additive form does not lose optimality. First, we prove that the leakage in the original problem formulation is always greater or equal to the leakage in the additive problem, the leakage being equal only when strategies $\bar{f}' \in \bar{\mathcal{F}}'$ are deployed.

Lemma B.5. For any strategy $f \in \mathcal{F}$

$$I(X^n, B_1; Y^n | E^n) \geq \sum_{t=1}^n I(X_t, B_t; E_t, Y_t | E^{t-1}, Y^{t-1}) \quad (\text{B.18})$$

with equality if and only if $f \in \bar{\mathcal{F}}'$.

Proof. We have:

$$\mathcal{I}(n, f) = \frac{1}{n} I(X^n, B_1; Y^n | E^n) \quad (\text{B.19a})$$

$$= \frac{1}{n} \sum_{t=1}^n H(Y_t | Y^{t-1}, E^n) - H(Y_t | Y^{t-1}, X^n, B_1, E^n) \quad (\text{B.19b})$$

$$= \frac{1}{n} \sum_{t=1}^n H(Y_t | Y^{t-1}, E^t) - H(Y_t | Y^{t-1}, X^t, B_1, E^t) \quad (\text{B.19c})$$

$$= \frac{1}{n} \sum_{t=1}^n I(X^t, B_1; Y_t | Y^{t-1}, E^t) \quad (\text{B.19d})$$

$$= \frac{1}{n} \sum_{t=1}^n I(X^t, B^t; Y_t | Y^{t-1}, E^t) \quad (\text{B.19e})$$

$$\geq \frac{1}{n} \sum_{t=1}^n I(X_t, B_t; Y_t | Y^{t-1}, E^t) \quad (\text{B.19f})$$

$$= \frac{1}{n} \sum_{t=1}^n H(X_t, B_t | Y^{t-1}, E^t) - H(X_t, B_t | Y^t, E^t) \quad (\text{B.19g})$$

$$= \frac{1}{n} \sum_{t=1}^n H(X_t, B_t | Y^{t-1}, E^{t-1}) - H(X_t, B_t | Y^t, E^t) \quad (\text{B.19h})$$

$$= \frac{1}{n} \sum_{t=1}^n I(X_t, B_t; E_t, Y_t | E^{t-1}, Y^{t-1}), \quad (\text{B.19i})$$

where (B.19c) follows because of the Markov chains $(X_{t+1}^n, E_{t+1}^n) \rightarrow (X^t, E^t, B_1, Y^{t-1}) \rightarrow Y_t$ and $E_{t+1}^n \rightarrow (E^t, Y^{t-1}) \rightarrow Y_t$; (B.19e) follows because B^t can be deterministically obtained from X^t, E^t, B_1 and Y^{t-1} ; (B.19f) follows because removing the terms X_1^{t-1} and B_1^{t-1} results in a lower mutual information; (B.19h) follows because of the Markov chain $E_t \rightarrow (Y^{t-1}, E^{t-1}) \rightarrow (X_t, B_t)$, since knowing E_t does not give more information about (B_t, X_t) than the knowledge of (Y^{t-1}, E^{t-1}) (E_t and X_t are independent, while B_t depends on E_{t-1} instead). The lower bound can be achieved if we deploy strategies that neglect the past and only decide the grid load based on the current values of the variables, thus when strategies $\bar{f}' \in \bar{\mathcal{F}}'$ are deployed. $\square$

Secondly, we need to show that the strategies $\bar{f}' \in \bar{\mathcal{F}}'$ achieve the least amount of information leakage. Thus, we need to show that for every leakage that is achievable by using the more general strategies $f \in \mathcal{F}$, the same leakage can be also achieved by adopting strategies $\bar{f}' \in \bar{\mathcal{F}}'$. For ease of notation, we define $K_t \triangleq (X_t, B_t)$.

Lemma B.6. *For any strategy $f \in \mathcal{F}$, there exists a strategy $\bar{f}' \in \bar{\mathcal{F}}'$ such that*

$$\sum_{t=1}^n I^f(K_t; E_t, Y_t | E^{t-1}, Y^{t-1}) = \sum_{t=1}^n I^{\bar{f}'}(K_t; E_t, Y_t | E^{t-1}, Y^{t-1}). \quad (\text{B.20})$$

Proof. For any strategy $f = (f_1, f_2, \dots) \in \mathcal{F}$ it is possible to find a strategy $\bar{f}' = (\bar{f}'_1, \bar{f}'_2, \dots) \in \bar{\mathcal{F}}'$ that achieves the lower bound in Lemma B.5 in the following way. For each t and all possible (y^t, e^t, k_t) such that the output is in the feasible set, i.e., $y_t \in \bar{\mathcal{Y}}(w_t)$, let

$$\bar{f}'_t(y_t | k_t, y^{t-1}, e^t) = p_{Y_t | K_t, Y^{t-1}, E^t}^f(y_t | k_t, y^{t-1}, e^t). \quad (\text{B.21})$$

To prove this lemma we show that the joint probability distributions are equal, i.e., $p_{K_t, Y^t, E^t}^f = p_{K_t, Y^t, E^t}^{\bar{f}'}, \forall t$. By construction of $\bar{f}'_t$ in (B.21) we have:

$$p_{Y_t | K_t, Y^{t-1}, E^t}^f = p_{Y_t | K_t, Y^{t-1}, E^t}^{\bar{f}'}. \quad (\text{B.22})$$

Then, it only remains to prove the following relation:

$$p_{K_t, Y^{t-1}, E^t}^f = p_{K_t, Y^{t-1}, E^t}^{\bar{f}'}. \quad (\text{B.23})$$

By using induction on t , for $t = 1$ we have $p_{K_1, E_1}^f = p_{K_1, E_1}^{\bar{f}'}$. Then, we obtain

$$p_{K_{t+1}, Y^t, E^{t+1}}^f(k_{t+1}, y^t, e^{t+1}) = \sum_{k_t} p_{K_{t+1}, K_t, Y^t, E^{t+1}}^f(k_{t+1}, k_t, y^t, e^{t+1}) \quad (\text{B.24a})$$

$$\begin{aligned} &= \sum_{k_t} p(k_{t+1}, e_{t+1} | y_t, k_t, e_t) \\ &\quad \times p_{Y_t | K_t, Y^{t-1}, E^t}^f(y_t | k_t, y^{t-1}, e^t) p_{K_t, Y^{t-1}, E^t}^f(k_t, y^{t-1}, e^t) \end{aligned} \quad (\text{B.24b})$$

$$= \sum_{k_t} p(k_{t+1}, e_{t+1} | y_t, k_t, e_t)$$

$$\times \bar{f}_t'(y_t|k_t, y^{t-1}, e^t) p_{K_t, Y^{t-1}, E^t}^{\bar{f}'}(k_t, y^{t-1}, e^t) \quad (\text{B.24c})$$

$$= p_{K_{t+1}, Y^t, E^{t+1}}^{\bar{f}'}(k_{t+1}, y^t, e^{t+1}), \quad (\text{B.24d})$$

where

$$p(k_{t+1}, e_{t+1}|y_t, k_t, e_t) = p(x_{t+1}, b_{t+1}, e_{t+1}|y_t, x_t, b_t, e_t) \quad (\text{B.25a})$$

$$= q_X(x_{t+1}|x_t) q_E(e_{t+1}|e_t) \mathbb{1}_{b_{t+1}}\{\min\{b_t + e_t - x_t, B_{\max}\} + y_t\}, \quad (\text{B.25b})$$

represents the system evolution, and (B.24c) follows from (B.21) and from the induction hypothesis at time t . $\square$

By combining Lemmas B.5 and B.6, we show that for any strategy $f \in \mathcal{F}$, it is possible to find a strategy $\bar{f}' \in \bar{\mathcal{F}}'$, such that

$$I^f(X^n, B_1; Y^n|E^n) \geq I^{\bar{f}'}(X^n, B_1; Y^n|E^n), \quad (\text{B.26})$$

and, for the strategies $\bar{f}' \in \bar{\mathcal{F}}'$, the right hand side of (B.26) actually coincides with the additive expression $\sum_{t=1}^n I^{\bar{f}'}(X_t, B_t; E_t, Y_t|Y^{t-1}, E^{t-1})$ in Eq. (4.24).

B.4 Proof of Corollary 4.3

We first prove the following lemma.

Lemma B.7. *For any strategy $f \in \mathcal{F}$*

$$\mathcal{I}(n, f) = \frac{1}{n} I(X^n, B_1; Y^n|E^n) \geq \frac{1}{n} \sum_{t=1}^n I(X_{t-m+1}^t, B_t; E_t, Y_t|Y^{t-1}, E^{t-1}), \quad (\text{B.27})$$

with equality if and only if $f \in \bar{\mathcal{F}}''$.

Proof. We have:

$$\mathcal{I}(n, f) = \frac{1}{n} I(X^n, B_1; Y^n|E^n) \quad (\text{B.28a})$$

$$= \frac{1}{n} \sum_{t=1}^n H(Y_t|Y^{t-1}, E^n) - H(Y_t|Y^{t-1}, X^n, B_1, E^n) \quad (\text{B.28b})$$

$$= \frac{1}{n} \sum_{t=1}^n I(X^t, B_1; Y_t|Y^{t-1}, E^t) \quad (\text{B.28c})$$

$$= \frac{1}{n} \sum_{t=1}^n I(X^t, B^t; Y_t|Y^{t-1}, E^t) \quad (\text{B.28d})$$

$$\geq \frac{1}{n} \sum_{t=1}^n I(X_{t-m+1}^t, B_t; Y_t|Y^{t-1}, E^t) \quad (\text{B.28e})$$

$$= \frac{1}{n} \sum_{t=1}^n H(X_{t-m+1}^t, B_t|Y^{t-1}, E^t) - H(X_{t-m+1}^t, B_t|Y^t, E^t) \quad (\text{B.28f})$$

$$= \frac{1}{n} \sum_{t=1}^n H(X_{t-m+1}^t, B_t|Y^{t-1}, E^{t-1}) - H(X_{t-m+1}^t, B_t|Y^t, E^t) \quad (\text{B.28g})$$

$$= \frac{1}{n} \sum_{t=1}^n I(X_{t-m+1}^t, B_t; E_t, Y_t|E^{t-1}, Y^{t-1}), \quad (\text{B.28h})$$

where (B.28c) follows because $(X_{t+1}^n, E_{t+1}^n) \rightarrow (X^t, E^t, B_1, Y^{t-1}) \rightarrow Y_t$ is a Markov chain; (B.28d) follows because B^t can be deterministically obtained from the knowledge of X^t, E^t, B_1 and Y^{t-1} ; (B.28e) follows as removing the terms X_1^{t-m} and B_1^{t-1} reduces the mutual information, and from the fact that we need to carry on information about the previous m values for X , since X is a Markov chain of order m ; (B.28g) follows because of the Markov chain $E_t \rightarrow (Y^{t-1}, E^{t-1}) \rightarrow (X_{t-m+1}^t, B_t)$, since knowing E_t does not give more information about (B_t, X_{t-m+1}^t) compared to the knowledge of (Y^{t-1}, E^{t-1}) (E_t and X_t are independent between each other, while B_t depends on E_{t-1} instead). It is easy to see that the equality in (B.28e) holds if $f \in \bar{\mathcal{F}}''$. $\square$

Next, we show that the strategies $\bar{f}'' \in \bar{\mathcal{F}}''$ achieve the minimum information leakage rate. Let $V_t \triangleq (X_{t-m+1}^t, B_t)$.

Lemma B.8. *For any strategy $f \in \mathcal{F}$, there exists a strategy $\bar{f}'' \in \bar{\mathcal{F}}''$ that induces the same information leakage, i.e.,*

$$\sum_{t=1}^n I^f(V_t; E_t, Y_t|Y^{t-1}, E^{t-1}) = \sum_{t=1}^n I^{\bar{f}''}(V_t, E_t; Y_t|Y^{t-1}, E^{t-1}). \quad (\text{B.29})$$

Proof. We want to show that for any leakage that can be achieved by the policy f , the same leakage can be achieved also by the policy $\bar{f}''$. Thus, for any strategy

$f = (f_1, \dots, f_n) \in \mathcal{F}$, we construct a corresponding $\bar{f}'' = (\bar{f}_1'', \dots, \bar{f}_n'') \in \bar{\mathcal{F}}''$ as follows. For each t and all possible (y^t, v_t, e^t) such that the output is in the feasible set, i.e., $y_t \in \bar{\mathcal{Y}}(w_t)$, let

$$\bar{f}_t''(y_t|v_t, y^{t-1}, e^t) = p_{Y_t|V_t, Y^{t-1}, E^t}^f(y_t|v_t, y^{t-1}, e^t). \quad (\text{B.30})$$

To prove the lemma we show that the joint probability distributions are equal, i.e., $p_{V_t, Y^t, E^t}^f = p_{V_t, Y^t, E^t}^{\bar{f}''}, \forall t$. By construction of $\bar{f}_t''$ in Eq. (B.30), we have:

$$p_{Y_t|V_t, Y^{t-1}, E^t}^f = p_{Y_t|V_t, Y^{t-1}, E^t}^{\bar{f}''}. \quad (\text{B.31})$$

Then, it only remains to prove the following relation:

$$p_{V_t, Y^{t-1}, E^t}^f = p_{V_t, Y^{t-1}, E^t}^{\bar{f}''}. \quad (\text{B.32})$$

By using induction on t , for $t = 1$ we have $p_{V_1, E_1}^f = p_{V_1, E_1}^{\bar{f}''}$. Then, we obtain

$$p_{V_{t+1}, Y^t, E^{t+1}}^f(v_{t+1}, y^t, e^{t+1}) = \sum_{x_{t-m}, b_t} p_{X_{t-m}^{t+1}, B_t^{t+1}, E^{t+1}, Y^t}^f(x_{t-m}^{t+1}, b_t^{t+1}, e^{t+1}, y^t) \quad (\text{B.33a})$$

$$\begin{aligned} &= \sum_{x_{t-m}, b_t} p(x_{t+1}, b_{t+1}, e_{t+1}|y_t, v_t, e_{t-l+1}^t) \\ &\quad \times p_{Y_t|V_t, Y^{t-1}, E^t}^f(y_t|v_t, y^{t-1}, e^t) p_{V_t, Y^{t-1}, E^t}^f(v_t, y^{t-1}, e^t) \end{aligned} \quad (\text{B.33b})$$

$$\begin{aligned} &= \sum_{x_{t-m}, b_t} p(x_{t+1}, b_{t+1}, e_{t+1}|y_t, v_t, e_{t-l+1}^t) \\ &\quad \times \bar{f}_t''(y_t|v_t, y^{t-1}, e^t) p_{V_t, Y^{t-1}, E^t}^{\bar{f}''}(v_t, y^{t-1}, e^t) \end{aligned} \quad (\text{B.33c})$$

$$= p_{V_{t+1}, Y^t, E^{t+1}}^{\bar{f}''}(v_{t+1}, y^t, e^{t+1}), \quad (\text{B.33d})$$

where (B.33c) follows from (B.30) and from the induction hypothesis at TS t ; and

$$\begin{aligned} p(x_{t+1}, b_{t+1}, e_{t+1}|y_t, x_{t-m+1}^t, b_t, e_{t-l+1}^t) = \\ q_X^m(x_{t+1}|x_{t-m+1}^t) q_E^l(e_{t+1}|e_{t-l+1}^t) \mathbb{1}_{b_{t+1}} \{ \min\{b_t + e_t - x_t, B_{\max}\} + y_t \}, \end{aligned} \quad (\text{B.34})$$

represents the system evolution, where q_X^m and q_E^l denote the transition probabilities

of the higher order Markov chains for the user load and the renewable energy process, respectively. $\square$

By combining Lemmas B.7 and B.8, we show that for any strategy $f \in \mathcal{F}$, it is possible to find a strategy $\bar{f}'' \in \bar{\mathcal{F}}''$, such that

$$I^f(X^n, E^n, B_1; Y^n) \geq I^{\bar{f}''}(X^n, E^n, B_1; Y^n), \quad (\text{B.35})$$

and, for the strategies $\bar{f}'' \in \bar{\mathcal{F}}''$, the right hand side of Eq. (B.35) actually coincides with the additive expression $\sum_{t=1}^n I^{\bar{f}''}(X_{t-m+1}^t, B_t; E_t, Y_t | Y^{t-1}, E^{t-1})$ in Eq. (4.26).