



UNIVERSITY OF THESSALY

DOCTORAL THESIS

Energy Management and Consumer Modeling in Smart Grid Systems

Author:

Vassiliki Hatzi

Supervisor:

Aspassia Daskalopulu

*A thesis submitted in partial fulfilment of the requirements
for the degree of Doctor of Philosophy*

in the

Department of Electrical and Computer Engineering

February 2017

Declaration of Authorship

I, Vassiliki Hatzi, declare that this thesis titled, 'Energy Management and Consumer Modeling in Smart Grid Systems' and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed:

Date:

The dissertation of Vassiliki Hatzi is approved by:

- Assistant Prof. Aspasia Daskalopulu
- Associate Prof. Iordanis Koutsopoulos
- Assistant Prof. Athanasios Korakis
- Prof. Lefteris Tsoukalas
- Prof. Emmanouil Vavalis
- Assistant Prof. Dimitrios Katsaros
- Assistant Prof. Antonios Argyriou

Abstract

Department of Electrical and Computer Engineering

Doctor of Philosophy

Energy Management and Consumer Modeling in Smart Grid Systems

by Vassiliki Hatzil

In the last decades, energy efficiency has turned into a major research issue, since the energy requirements of modern societies are growing continuously. Researchers have focused on optimizing the efficiency of the modernized power grid, i.e., the smart grid, which has evolved into a complex ecosystem with different actors such as consumers, operators and generators having different active roles in the system. They have also focused on improving the energy efficiency of ICT which is a prerequisite for efficient smart grids. Moreover, the advance of key technologies such as energy storage, renewable energy sources (RESs), communication and control has opened the way to new research directions. In this thesis, we present some key research problems in the context of this area, and we explore the use of control and optimization methods toward approaching them. Specifically, our goal is to address i) system challenges pertaining to the integration of energy storage and RESs into the smart grid and ii) challenges related to the energy consumer aspect.

Energy storage devices, like uninterruptible power supply (UPS) or batteries, and Plug-in Hybrid Electric Vehicles (PHEVs), are prime resources for smart-grid efficiency improvement which need to be appropriately managed and incorporated into smart grid systems. In this context, we study two fundamental problems in energy storage management and dimensioning in smart grids. Specifically, first, we introduce the optimal energy storage control problem faced by an energy supplier, which amounts to deciding when and how much to charge and discharge a storage device in order to achieve a certain optimization objective in terms of energy generation cost. We address the problem above, first for a single storage device, and then for multiple storage devices that are shared among multiple micro-grids. Our optimization objective leads us to policies which attempt to keep balanced grid power consumption at all times. Next, we study a joint energy storage placement, dimensioning and management problem, given an available storage budget, where the goal is to minimize the power generation cost. We are interested in the way storage capacity placement and control impact the overall cost of energy generation.

The solution policy for this problem involves various parameters such as the demand profiles of consumers, and power flow and balance constraints.

RESs rely mainly on energy that flows naturally through the environment on a continual but time-varying basis and have received major attention in the last years due to their significant potential in reducing the carbon footprint. In this thesis, we study the use of renewable energy sources and scheduling techniques toward reducing the carbon footprint induced by an ICT system, namely, the web crawling component of a web search engine. This component discovers and downloads new pages on the Web as well as refreshes previously downloaded pages in the web repository. We introduce the problem of green web crawling, where the objective is to devise a page refresh policy that keeps the web pages as much as possible fresh and the carbon emissions that the web crawling process incurs on remote web servers low enough. We devise an optimal policy, which can be implemented in an online fashion, based only on the type of energy consumed by the servers and the staleness of the pages in the web repository. We also devise heuristics along the lines of the optimal policy and study their performance through experiments with real data.

Demand-side management (DSM) activities for the smart grid aim to reduce or smoothen energy consumption by providing dynamic prices or incentives in the form of monetary or non-monetary rewards to consumers. Different DSM techniques have been proposed, however, a fundamental issue in their design and operation is the recruitment of users. Serious-games design is an emerging area that can address precisely the issue of maximizing user engagement in various contexts. In this thesis, we discuss the use of serious games for demand-side management in smart grids. We introduce the problem of optimal serious-games design for the purpose of enforcing prudent energy consumption. We present a mathematical model of a simple gamification mechanism through which a serious-game designer (e.g., a demand-side management entity) aims to motivate consumers to reduce their energy consumption at peak hours by setting up a contest and by providing them incentives in the context of a serious game. The game designer optimally selects the game parameters, so as the utility-maximizing choices of consumers to minimize the energy generation cost of the energy supplier. We demonstrate that even such a simple serious-game design can provide adequate incentives to users for engaging in demand-side management.

Finally, most proposed DSM schemes assume that consumers are rational decision makers which select their actions by solving complex optimization problems. However, consumers are humans and their decisions, which are driven by different factors, are far from rational. In this direction, we discuss the role of data in building behavior-based models for profiling energy consumers and predicting their behavior in DSM programs and

energy consumption curtailment campaigns. Our ultimate goal is to discover through these models the different factors that determine consumer actions and the different importance placed on them. We present two different data-driven approaches for consumer modeling which are based on a popular machine-learning tool and a cognitive heuristic. We show that both approaches succeed in capturing the different importance that each consumer places on different factors and the uncertainty on consumer actions. We also introduce the optimal load-reduction task and incentive allocation problem faced by the designer of an energy consumption reduction campaign. The aim of the designer is to target tasks and incentives appropriately based on the different consumer profiles, so as to best fulfill the purpose of the campaign.

Περίληψη

Τμήμα Ηλεκτρολόγων Μηχανικών και Μηχανικών Υπολογιστών

Διδακτορικό Δίπλωμα

Διαχείριση Ενέργειας και Μοντελοποίηση Καταναλωτών σε Έξυπνα Δίκτυα Ηλεκτρικής Ενέργειας

Βασιλική Χατζή

Τις τελευταίες δεκαετίες, η ενεργειακή αποδοτικότητα αποτελεί μείζον ερευνητικό θέμα, καθώς οι ενεργειακές απαιτήσεις των σύγχρονων κοινωνιών αυξάνονται συνεχώς. Οι ερευνητές έχουν επικεντρωθεί στην βελτιστοποίηση της αποδοτικότητας του εκμοντερνισμένου δικτύου ισχύος, δηλαδή του έξυπνου δικτύου ηλεκτρικής ενέργειας, το οποίο έχει εξελιχθεί σε ένα πολυσύνθετο οικοσύστημα με διάφορους παράγοντες/φορείς όπως καταναλωτές, φορείς εκμετάλλευσης (διαχειριστές) και παραγωγούς ηλεκτρικής ενέργειας οι οποίοι έχουν διαφορετικούς ενεργούς ρόλους στο σύστημα. Έχουν επίσης επικεντρωθεί στην βελτιστοποίηση της ενεργειακής αποδοτικότητας των Τεχνολογιών Πληροφορίας και Επικοινωνιών (ΤΠΕ), οι οποίες αποτελούν απαραίτητη προϋπόθεση για την ανάπτυξη αποδοτικών έξυπνων δικτύων ηλεκτρικής ενέργειας. Επιπλέον, η εξέλιξη βασικών τεχνολογιών, όπως η αποθήκευση ενέργειας, οι ανανεώσιμες πηγές ενέργειας (ΑΠΕ), η επικοινωνία και ο έλεγχος, άνοιξε τον δρόμο σε νέες ερευνητικές κατευθύνσεις. Στην παρούσα διδακτορική διατριβή, παρουσιάζουμε μερικά καιρία ερευνητικά προβλήματα στο πλαίσιο αυτής της ερευνητικής περιοχής, και διερευνούμε τη χρήση μεθόδων ελέγχου και βελτιστοποίησης για την προσέγγιση τους. Συγκεκριμένα, στόχος μας είναι να αντιμετωπίσουμε *α)* προκλήσεις σε επίπεδο συστήματος που αφορούν την ενσωμάτωση αποθηκών ενέργειας και ΑΠΕ στο έξυπνο δίκτυο ηλεκτρικής ενέργειας και *β)* προκλήσεις που σχετίζονται με τον ίδιο καταναλωτή.

Οι συσκευές αποθήκευσης ενέργειας, όπως η αδιάλειπτη παροχή ηλεκτρικού ρεύματος (UPS) ή οι μπαταρίες, και τα επαναφορτιζόμενα υβριδικά ηλεκτρικά αυτοκίνητα (PHEVs), είναι πρωταρχικοί πόροι για την βελτιστοποίηση της αποτελεσματικότητας του έξυπνου δικτύου ηλεκτρικής ενέργειας, οι οποίοι πρέπει να ενσωματωθούν και να διαχειρισθούν κατάλληλα σε συστήματα έξυπνων δικτύων ηλεκτρικής ενέργειας. Σε αυτό το πλαίσιο, μελετούμε δύο θεμελιώδη προβλήματα που αφορούν την διαχείριση και διαστασιολόγηση αποθηκών ενέργειας σε έξυπνα δίκτυα ηλεκτρικής ενέργειας. Συγκεκριμένα, αρχικά, παρουσιάζουμε το πρόβλημα βέλτιστου ελέγχου αποθήκευσης ενέργειας που αντιμετωπίζει ένας προμηθευτής ενέργειας, το οποίο αφορά την λήψη αποφάσεων σχετικά

με το πότε και πόσο να φορτίσει και να αποφορτίσει μια συσκευή αποθήκευσης έτσι ώστε να επιτύχει έναν συγκεκριμένο αντικειμενικό σκοπό βελτιστοποίησης σε όρους κόστους παραγωγής ενέργειας. Επιλύουμε το παραπάνω πρόβλημα, αρχικά για μία μοναδική συσκευή αποθήκευσης και έπειτα για πολλαπλές συσκευές αποθήκευσης οι οποίες χρησιμοποιούνται από κοινού από πολλαπλά μικρο-δικτύα. Ο αντικειμενικός στόχος του προβλήματος μας οδηγεί σε πολιτικές που προσπαθούν να διατηρήσουν μια ισορροπημένη κατανάλωση ισχύος στο δίκτυο ανά πάσα στιγμή. Έπειτα, μελετούμε ένα πρόβλημα τοποθέτησης, διαστασιολόγησης και διαχείρισης αποθηκών ενέργειας με δεδομένο ένα διαθέσιμο μπάτζετ αποθήκευσης, όπου ο στόχος είναι να ελαχιστοποιηθεί το κόστος παραγωγής ενέργειας. Μας ενδιαφέρει ο τρόπος με τον οποίο η τοποθέτηση και ο έλεγχος της χωρητικότητας αποθήκευσης επηρεάζουν το συνολικό κόστος παραγωγής ενέργειας. Η πολιτική που προκύπτει από την επίλυση του προβλήματος εμπλέκει διάφορες παραμέτρους όπως τα προφίλ ζήτησης των καταναλωτών, και περιορισμούς ροής και ισορροπίας ισχύος.

Οι ΑΠΕ βασίζονται κυρίως σε ενέργεια που ρέει φυσικά μέσω του περιβάλλοντος σε συνεχή αλλά χρονικά μεταβαλλόμενη βάση, και τα τελευταία χρόνια έχουν προσελκύσει το ενδιαφέρον λόγω της σημαντικής τους προοπτικής να μειώσουν το αποτύπωμα άνθρακα. Στην παρούσα διατριβή, μελετούμε τη χρήση ΑΠΕ και τεχνικών χρονοπρογραμματισμού με στόχο την μείωση του αποτυπώματος άνθρακα που προκαλείται από ένα σύστημα ΤΠΕ, συγκεκριμένα, από το πρόγραμμα ανίχνευσης παγκόσμιου ιστού μιας μηχανής διαδικτυακής αναζήτησης. Αυτό το πρόγραμμα ανακαλύπτει και κατεβάζει νέες ιστοσελίδες του παγκόσμιου ιστού καθώς επίσης, ανανεώνει σελίδες που έχουν κατεβεί προηγουμένως στον αποθηκευτικό χώρο του ιστού (web repository). Εισάγουμε το πρόβλημα της πράσινης (οικολογικής) ανίχνευσης παγκόσμιου ιστού, όπου αντικειμενικός σκοπός είναι να αναπτύξουμε μια πολιτική ανανέωσης σελίδων η οποία να διατηρεί τις ιστοσελίδες όσο το δυνατόν περισσότερο ανανεωμένες και τις εκπομπές διοξειδίου του άνθρακα, που η διαδικασία ανίχνευσης ιστού προκαλεί σε απομακρυσμένους διακομιστές διαδικτύου, αρκετά χαμηλές. Εξάγουμε μια βέλτιστη πολιτική η οποία μπορεί να εφαρμοστεί σε πραγματικό χρόνο βασιζόμενη μόνο στον τύπο της ενέργειας που καταναλώνεται από τους διακομιστές διαδικτύου και στην παλαιότητα των σελίδων που βρίσκονται στον αποθηκευτικό χώρο του ιστού. Αναπτύσσουμε επίσης ευριστικές πολιτικές έχοντας ως πρότυπο την βέλτιστη πολιτική και μελετούμε την απόδοση τους μέσω πειραμάτων με πραγματικά δεδομένα.

Οι δραστηριότητες διαχείρισης ζήτησης για το έξυπνο δίκτυο ηλεκτρικής ενέργειας έχουν ως στόχο να μειώσουν ή να εξομαλύνουν την κατανάλωση ενέργειας παρέχοντας στους καταναλωτές δυναμικές τιμές ή κίνητρα υπό την μορφή χρηματικών ή μη χρηματικών ανταμοιβών. Διάφορες τεχνικές διαχείρισης ζήτησης έχουν προταθεί, ωστόσο, ένα βασικό θέμα στον σχεδιασμό και στην λειτουργία τους είναι η στρατολόγηση των χρηστών. Ο σχεδιασμός σοβαρών παιγνίων είναι μια αναδυόμενη περιοχή που μπορεί να αντιμετωπίσει

με ακρίβεια το ζήτημα της μεγιστοποίησης της συμμετοχής των χρηστών σε διάφορα πλαίσια. Στην παρούσα διδακτορική διατριβή, συζητούμε την χρήση των σοβαρών παιγνίων στο ζήτημα της διαχείρισης της ζήτησης στα έξυπνα ηλεκτρικά δίκτυα. Εισάγουμε το πρόβλημα του βέλτιστου σχεδιασμού σοβαρών παιγνίων έχοντας ως στόχο την επιβολή συνετής κατανάλωσης ενέργειας. Παρουσιάζουμε ένα μαθηματικό μοντέλο ενός απλού μηχανισμού παιγνιοποίησης μέσω του οποίου ένας σχεδιαστής σοβαρού παιγνίου (π.χ. μια οντότητα διαχείρισης ζήτησης) επιδιώκει να παρακινήσει τους καταναλωτές να μειώσουν την ενεργειακή τους κατανάλωση σε ώρες αιχμής (υψηλής ζήτησης) στήνοντας έναν διαγωνισμό και παρέχοντας τους κίνητρα στο πλαίσιο ενός σοβαρού παιγνίου. Ο σχεδιαστής επιλέγει βέλτιστα τις παραμέτρους του παιγνίου έτσι ώστε οι επιλογές των καταναλωτών, οι οποίες μεγιστοποιούν την χρησιμότητα/όφελος των καταναλωτών, να ελαχιστοποιούν το κόστος παραγωγής ενέργειας του προμηθευτή ενέργειας.

Τέλος, οι περισσότερες προτεινόμενες στρατηγικές διαχείρισης ζήτησης υποθέτουν ότι οι καταναλωτές λαμβάνουν αποφάσεις ορθολογικά και ότι επιλέγουν τις δράσεις τους λύνοντας πολύπλοκα προβλήματα βελτιστοποίησης. Ωστόσο, οι καταναλωτές είναι άνθρωποι και οι αποφάσεις τους, οι οποίες οδηγούνται από διάφορους παράγοντες, απέχουν πολύ από τη λογική. Προς αυτή την κατεύθυνση, συζητούμε τον ρόλο των δεδομένων στο χτίσιμο μοντέλων που βασίζονται στην ανθρώπινη συμπεριφορά, τα οποία χρησιμοποιούνται για την δημιουργία προφίλ των καταναλωτών ενέργειας και την πρόβλεψη της συμπεριφοράς τους σε προγράμματα διαχείρισης ζήτησης και σε καμπάνιες/εκστρατείες μείωσης της κατανάλωσης ενέργειας. Στόχος μας είναι να ανακαλύψουμε μέσω αυτών των μοντέλων τους διαφορετικούς παράγοντες που καθορίζουν τις πράξεις/δράσεις των καταναλωτών και την διαφορετική βαρύτητα που δίνεται σε αυτούς. Παρουσιάζουμε δύο διαφορετικές προσεγγίσεις για την μοντελοποίηση των καταναλωτών, οι οποίες έχουν ως γνώμονα τα δεδομένα και βασίζονται σε ένα δημοφιλές εργαλείο μηχανικής μάθησης και σε ένα γνωσιακό ευριστικό μοντέλο. Δείχνουμε ότι και οι δύο προσεγγίσεις επιτυγχάνουν στο να συλλάβουν την διαφορετική βαρύτητα/σημασία που ο κάθε καταναλωτής δίνει στους διαφορετικούς παράγοντες και την αβεβαιότητα στις πράξεις των καταναλωτών. Εισάγουμε επίσης το πρόβλημα της βέλτιστης κατανομής κινήτρων και εργασιών μείωσης φορτίου που αντιμετωπίζει ο σχεδιαστής μιας καμπάνιας μείωσης της κατανάλωσης ενέργειας. Στόχος του σχεδιαστή είναι να κατευθύνει εργασίες και κίνητρα κατάλληλα βάσει των διαφορετικών προφίλ των καταναλωτών, έτσι ώστε να εκπληρώσει βέλτιστα τον σκοπό της καμπάνιας.

Acknowledgements

This thesis represents the culmination of research that was conducted towards my PhD degree from the Department of Electrical and Computer Engineering, University of Thessaly, Greece.

First and foremost, I would like to thank Associate Prof. Iordanis Koutsopoulos for not only giving me the opportunity to work with him but also being a great mentor. I am deeply indebted to him for first taking me as an undergraduate and Master's student, and for offering constructive criticism and advice throughout my time as his student. His patient guidance and understanding have been a great support in this long journey.

I would like to express my thanks and appreciation to the members of the thesis committee; Assistant Prof. Aspasia Daskalopulu, Associate Prof. Iordanis Koutsopoulos, Assistant Prof. Athanasios Korakis, Prof. Lefteris Tsoukalas, Prof. Emmanouil Vavalis, Assistant Prof. Dimitrios Katsaros and Assistant Prof. Antonios Argyriou.

I would like to thank Prof. Leandros Tassiulas for co-supervising the energy storage management problem presented in Chapter 2. One of the problems, the green web-crawling one presented in Chapter 3 was studied in collaboration with Yahoo! Labs Barcelona. I would like to thank Dr. Berkant Barla Cambazoglu (Yahoo! Labs) for the fruitful cooperation in the context of this work and for providing real data (sampled from a large web crawl performed by Yahoo!).

Finally, special thanks go to my parents, my sister and my husband for their love and patience all these years. Their continuous support made the accomplishment of this thesis possible. The least I can do in return is to dedicate this thesis to them.

Publications

The results of this thesis are included in the following publications:

International Conferences

- [C.1] V. Hatzi, B. B. Cambazoglu and I. Koutsopoulos, “Web Page Download Scheduling Policies for Green Web Crawling”, *Proc. 22nd International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, September 2014.
- [C.2] T. G. Papaioannou, V. Hatzi and I. Koutsopoulos, “Optimal Design of Serious Games for Demand Side Management”, *Proc. IEEE International Conference on Smart Grid Communications (SmartGridComm)*, Venice, Italy, November 2014.
- [C.3] I. Koutsopoulos, V. Hatzi, and L. Tassiulas, “Optimal Energy Storage Control Policies for the Smart Power Grid”, *Proc. IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pp. 475-480, Brussels, 2011.

International Journals and Magazines

- [J.1] V. Hatzi, B. B. Cambazoglu, and I. Koutsopoulos, “Optimal Web Page Download Scheduling Policies for Green Web Crawling”, *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 5, pp. 1378-1388, May 2016.
- [J.2] T. Papaioannou, V. Hatzi, and I. Koutsopoulos, “Optimal Design of Serious Games for Consumer Engagement in the Smart Grid”, *IEEE Transactions on Smart Grid*, June 2016.

Participation in Books

- [B.1] I. Koutsopoulos, T. Papaioannou, and V. Hatzi, “Modeling and Optimization of the Smart Grid Ecosystem”, *Foundations and Trends in Networking*, vol. 10, no. 2-3, pp. 115-316, June 2016.

In Preparation for Submission

[S.1] V. Hatzi, I. Koutsopoulos, and L. Tassiulas, “Optimal Energy Storage Management Policies for Smart Grid Systems”.

[S.2] V. Hatzi, and I. Koutsopoulos, “Optimal Design of Energy Consumption Reduction Campaigns Through Behavioral Data-Driven Consumer Profiling”.

Contents

Declaration of Authorship	i
Abstract	iii
Greek Abstract	vi
Acknowledgements	ix
Publications	x
Contents	xii
List of Figures	xvi
List of Tables	xviii
1 Introduction	1
1.1 Motivation	1
1.2 Contributions of this thesis	5
2 Optimal Energy Storage Management and Dimensioning in Smart Grid Systems	8
2.1 Introduction to the Employment of Energy Storage Devices in the Smart Grid	9
2.1.1 Our contribution	10
2.2 Optimal Energy Storage Control Policies for the Energy Supplier	11
2.2.1 System model	11
2.2.2 Problem formulation	15
2.2.3 An Asymptotically Optimal Control Policy	16
2.2.4 Extension to the model: Renewable source	18
2.2.5 Multiple Storage Devices and Grid Entities	19
2.2.5.1 Case A: Minimization of cost of aggregate demand	22
2.2.5.2 Case B: Minimization of sum of individual GE costs	23
2.2.5.3 Asymptotically Optimal Control Policy	24
2.2.6 Extensions to the model	25
2.2.6.1 Non-Stationarity of Renewable Energy Generation	25

2.2.6.2	Storage Device Modeling	26
2.2.6.3	Transmission Losses	26
2.2.7	Numerical Results	27
2.2.7.1	One Storage Device	27
2.2.7.2	Two storage devices and GEs	29
2.3	Storage Placement and Power Flow	30
2.3.1	Power flow analysis	31
2.3.2	Storage and power flow model	32
2.3.3	Problem formulation	34
2.3.4	Numerical Example 1: Storage placement for N=2 buses	35
2.3.5	Numerical Example 2: Storage dimensioning for N=2 buses	38
2.4	Related Work	39
2.5	Conclusion	42
3	Green Web Crawling	43
3.1	Introduction to the Energy Efficiency of Web Search Engines	44
3.2	System Model	46
3.2.1	Web Crawler	46
3.2.2	Web Servers	47
3.2.3	Greenness of Server Energy Consumption	48
3.2.4	Web Page Staleness	48
3.3	Problem Formulation	49
3.3.1	Single Web Server, Single Thread Scheduling Problem	49
3.3.2	Optimal Web Page Download Scheduling Policy	51
3.3.3	Extensions to the Model	52
3.3.3.1	Many Web Servers	52
3.3.3.2	Many Web Servers, Multiple Crawling Threads	54
3.3.3.3	Web Pages with Variable Freshness Requirements	54
3.4	Performance Evaluation	55
3.4.1	Data Set	55
3.4.2	Greenness and Staleness Computation	55
3.4.3	Heuristic Policies	57
3.4.4	Experimental Results	58
3.4.4.1	Performance of the Optimal Policy	59
3.4.4.2	Performance of the Heuristic Policies	60
3.4.4.3	Performance Comparison	62
3.5	Related Work	63
3.5.1	Refreshing Web Repositories	63
3.5.2	Job and Packet Processing	63
3.5.3	Energy Efficiency and Greenness of Data Centers	64
3.6	Conclusion	65
4	Optimal Design of Serious Games for Smart Grid Consumer Engagement	66
4.1	Introduction to the Concept of Serious Games for Demand-side Management	67
4.2	Serious Game Setup	69
4.3	System Model	70

4.3.1	Consumer's problem	72
4.3.1.1	Dissatisfaction	72
4.3.1.2	Social recognition	72
4.3.1.3	Feedback to consumers	73
4.3.1.4	Probability of inclusion in the top-K and the bottom-M lists	74
4.3.1.5	User utility function	76
4.3.2	Game Designer's Problem	76
4.3.2.1	Full information of consumer utility functions	77
4.3.2.2	Historical information about consumer actions	79
4.3.3	Equilibrium of the Stackelberg game arising from serious-game interactions	80
4.4	Simulation Results and Analysis	81
4.4.1	Experimental Setup	81
4.4.2	Steady-State Convergence and Wear-off Effects	81
4.4.3	Experimental Results	83
4.5	Related Work	85
4.6	Conclusion	86
5	Optimal Design of Energy Consumption Reduction Campaigns Through Behavioral Data-Driven Consumer Profiling	88
5.1	Introduction to Data-driven Energy Consumer Profiling and Behavior Prediction	89
5.1.1	Our contribution	91
5.2	Energy Consumption Reduction Campaign: A simple scenario	92
5.3	System Model	93
5.4	Incentive and load-reduction task allocation: Problem formulation	94
5.5	Consumer profiling	95
5.5.1	Logistic Regression model, and Optimal Task and Incentive Allocation Policy	96
5.5.1.1	Optimal allocation of incentives and load reduction tasks in the case of LR: A sigmoid optimization problem	96
5.5.2	Fast-and-Frugal Tree model, and Optimal Task and Incentive Allocation Policy	97
5.5.2.1	Optimal allocation of incentives and load reduction tasks in the case of FFTs: An integer linear programming (ILP) problem	101
5.6	Numerical Example	103
5.7	Related work	107
5.8	Conclusion	110
6	Concluding Remarks and Future Challenges	111
6.1	Summary of Contributions	111
6.2	Future challenges	113
A	Training the Logistic Regression models	116

B Training the Fast-and-Frugal Trees	118
C Building procedure details and performance analysis of the fitted LR and FFT models of section 5.6	121
Bibliography	126

List of Figures

1.1	Example electricity demand curve showing the effect of DR actions. Image source: [4].	2
1.2	Electricity consumption forecasts of ICT equipment during use [5].	3
1.3	Main components of electricity consumption for the ICT sector. Image source: [8].	4
1.4	Estimated distribution of global CO_2 emissions from ICT. Image source: [9].	4
2.1	Overview of system model with the energy storage device, the charging and discharging process, and the interaction with grid consumption. . . .	12
2.2	Power generation cost as a piece-wise linear convex function of demand load.	15
2.3	Extension to the model, with a renewable source feeding the battery. . . .	17
2.4	System with multiple storage devices and appended renewable energy sources.	19
2.5	Performance of the proposed energy storage control policy as a function of available storage capacity.	27
2.6	Amount of stored energy for A) $E_{max} = 10$ kWh and B) $E_{max} = 24$ kWh.	28
2.7	Total instantaneous grid load for A) $E_{max} = 10$ kWh and B) $E_{max} = 24$ kWh.	28
2.8	Minimum required capacity E_{max} for which the policy is optimal, versus the demand load.	29
2.9	Performance of the proposed policy for the case of multiple batteries as a function of available storage capacities with $\lambda_1 \neq \lambda_2$, $E_{max}^1 = E_{max}^2$	30
2.10	Power balance at node k	34
2.11	A power network consisting of two buses. It must hold $E_{max}^1 + E_{max}^2 \leq S_B$	35
2.12	The optimal system performance for different values of the available storage S_B and line capacity f_{12}	38
3.1	Our system model for the crawler: m crawling threads concurrently retrieve pages from N web servers at time slot t	47
3.2	The variation in (a) the solar irradiance [52] and (b) the $g_i(t)$ values during the day.	56
3.3	The performance of the optimal policy in comparison with that of the EDD-like policy in terms of a) total staleness, and b) carbon emissions as a function of λ	59
3.4	a) The average amount of carbon emissions (measured in grams (g)) generated by the proposed heuristic policies, and b) the performance of all heuristics in terms of staleness reduction.	59

3.5	The average page staleness of a) the first four heuristic policies and b) all heuristics at the end of the simulation.	61
4.1	The serious-game interactions among consumers, the serious-game designer and the utility company.	69
4.2	The Stackelberg game structure arising from serious-game interactions.	71
4.3	Sample performance histograms of players at a time slot; $\delta=0.125$	73
4.4	(A) Probability of consumer i to be included in the top- K or in the bottom- M lists with respect to r_i . (B) Expected utilities of consumer i based on equation (4.3) (“probability-based”) and on equation (4.6) (“indicator function-based”) with respect to r_i . Parameters values: $X \sim U[0, 1]$, $N = 100$, $K = M = 10$, $a_i = 1.5$, $h_i = 1.01$, $p_0 = 6$ kWh, $q = 0.124$	75
4.5	The steady (set of) states $\langle (r_l, r_u) \rangle$ for different $K=M$	80
4.6	The percentage of operational-cost reduction achieved by our approach for different percentages of consumers leaving the game.	81
4.7	Utility functions of two players for $K = 4$ and two different states.	82
4.8	Operational cost for different values of $K=M$	82
4.9	Social welfare for different values of $K=M$	83
4.10	Distribution of r_i for different pairs (r_l, r_u) in the cases of a) $K = 4$ and b) $K = 350$	84
5.1	(a) An example FFT and (b) its alternative form for a consumer i that prioritizes the cue d'_i of the suggestion \mathbf{r}_i over the cue p_i . $AP_{i,l}/N_{i,l}$ is the estimate of the probability that an instance (i.e., suggestion) assigned to leaf l belongs to class C_1 . $AP_{i,l}$ is the number of training instances of (actual) class C_1 at leaf l , and $N_{i,l}$ is the total number of training instances at that leaf.	98
5.2	An example FFT that prioritizes the cue p_i of the suggestion \mathbf{r}_i over the cue d'_i . Its cue ranking and exit location are different from those of the FFT in Fig. 5.1a.	98
5.3	The resulting FFT of consumer 1: priority to the payment attribute p_1	104
5.4	The resulting FFT of consumer 2: exclusive reliance on the percentage load reduction cue \tilde{d}_2	104
C.1	Classification tables for (a) the LR model and (b) the regularized LR model of consumer 1.	123
C.2	Classification tables for (a) the LR model and (b) the regularized LR model of consumer 2.	124
C.3	Classification tables for (a) the FFT of consumer 1, and (b) the FFT of consumer 2.	124
C.4	Performance comparison of the FFT and LR models of (a) consumer 1 and (b) consumer 2.	125

List of Tables

2.1	Average cost of the proposed policy for $M = 2$ storage devices	30
2.2	Optimal storage dimensioning, charging/discharging and power flow for $S_B = 5$ and $f_{12} = 0.5, 1$	39
3.1	Performance results as the impact of cold-start is reduced	61
3.2	Performance comparison of the proposed policies	62
B.1	Contingency table over an abstract attribute x	118
C.1	Example questionnaire and consumers' responses.	121
C.2	Contingency table over cue p_1	122
C.3	Contingency table over cue \tilde{d}_1	122
C.4	Contingency table over cue p_2	122
C.5	Contingency table over cue \tilde{d}_2	122
C.6	Performance statistics of the fitted LR and FFT models.	123

*Dedicated to my husband, Dimitris and to my family, Aristeidis,
Evaggelia, Marianthi.*

Chapter 1

Introduction

Contents

1.1 Motivation	1
1.2 Contributions of this thesis	5

1.1 Motivation

Electricity is the most widely used form of energy and global demand is constantly increasing. World electricity demand is projected to double between 2000 and 2030, growing at an annual rate of 2.4%. Also, electricity's share of total final energy consumption rises from 18% in 2000 to 22% in 2030 [1]. Our modern societies rely on the power grid system to satisfy their energy needs. The power grid must respond quickly to the increasing demand and continuously generate and route electricity to where it's needed the most. However, the generation of electrical energy makes a significant contribution to climate change since it relies on fossil fuels and is currently the largest source of carbon dioxide emissions. To satisfy both the increasing demand for energy and the need to reduce carbon dioxide emissions, we need a power grid system that can handle these challenges in a reliable and economic way [2].

According to [2], "a smart grid is an evolved power grid system that manages electricity demand in a sustainable, reliable and economic manner, built on advanced infrastructure and tuned to facilitate the integration of all involved". Smart power grids harness ICT and smart metering in order to enhance their reliability, enforce sensible use of energy, reduce CO_2 emissions, and efficiently incorporate and control components such as renewable energy sources (RESs), distributed micro-generators and energy storage entities (e.g. plug-in electric vehicles, batteries, etc.) [3]. These modernized electricity




FIGURE 1.1: Example electricity demand curve showing the effect of DR actions. Image source: [4].

networks are complex ecosystems which are composed of new entering actors such as prosumers, and traditional ones like consumers, operators and generators each having a fundamentally different, active role in the system. In this thesis, we focus on developing efficient mechanisms that optimally incorporate and manage energy storage devices in smart grid systems so as to exploit their potential in reducing the energy generation cost of energy supplier/s (utility operator).

Smart grids possess demand-side management (DSM) capability, also known as Demand Response (DR), to better match the demand for energy with the supply and avoid energy overconsumption. The goal of DR is to encourage consumers, through various methods such as financial incentives and behavioral change through education, to use less energy during peak hours, or to move non-emergency power demands to off-peak times such as nighttime and weekends (Fig. 1.1). By smoothing out the system power demand profile across time, grid reliability is enhanced as instabilities are reduced, power outages due to sudden increases of demand are avoided, and the need for activating supplementary power generation sources so as to satisfy high demand is eliminated. This supplementary power may be generated from expensive sources, or it may be imported at high prices from other countries. Therefore, in DR programs, end-users enjoy reduced electricity bills due to lower real-time electricity prices, while the operator enjoys reduced operating cost.

Current DR schemes provide incentives to the users, usually in the form of dynamic prices and monetary or non-monetary rewards. They rely on strong rationality assumptions about consumers, and they assume that the response of consumers to incentives is governed by the mathematics of optimization theory. However, it is known from behavioral




FIGURE 1.2: Electricity consumption forecasts of ICT equipment during use [5].

science that human behavior is far from rational due to the fact that humans often make decisions under certain cognitive biases, various natural predispositions, social norms, various sentiments and prejudices, or even in the presence of limited information and limited information processing capacity. Various DR schemes may often see hesitation and even negativity of consumers, primarily because they are not presented to consumers in an appropriate context. This means that the efficiency of the above programs mainly depends on the *participation* of consumers and their *behavior*. The interfaces with which such DR schemes are delivered to consumers should respect the limited time and attention of consumers and should make the interaction worthwhile and entertaining. In this thesis, we focus on designing efficient mechanisms that exploit the concept of serious games in order to encourage consumers to save energy by means of competition, and increase *consumer engagement* in DR programs by creating a more enjoyable experience for consumers. We also focus on providing *behavior-based* data-driven models for building personalized decision-making profiles for end-users, predicting consumer behavior, and optimally providing incentives in DR programs.

Toward the aforementioned goal of energy consumption and carbon emissions reduction, we also need to focus on the energy efficiency of large energy-consuming entities. Information and communication technologies (ICT) and networked systems have been proved to be big energy consumers with a large carbon footprint [5], [6] and thus, they need to undergo significant adjustments. ICT is a prerequisite for efficient smart grids but it is also responsible for a quick increase in worldwide energy consumption. The authors of [5] argue that the relative contribution of ICT in energy consumption is expected to grow from 8% in 2008 to more than 14% in 2020 (Fig. 1.2). Networks and data

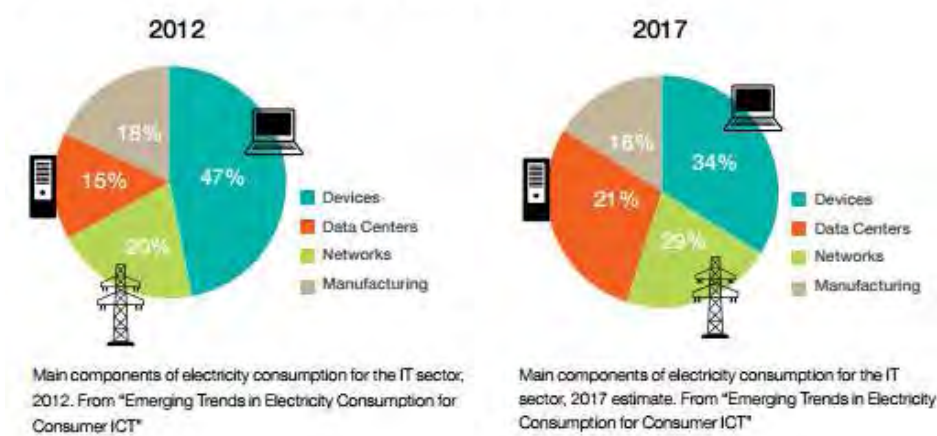


FIGURE 1.3: Main components of electricity consumption for the ICT sector. Image source: [8].




FIGURE 1.4: Estimated distribution of global CO₂ emissions from ICT. Image source: [9].

centers have been proved to be major energy consumers and their energy consumption is estimated to reach 29% and 21%, respectively, of the ICT sector's electricity usage in 2017 [7]. (Fig. 1.3)

Moreover, the increased use of ICT contributes to global warming. According to [6], the ICT sector itself (telecommunications, computing and the Internet, but excluding broadcasting transmitters and receivers) contributes around 2.5% of greenhouse gas emissions, at just under 1 Gigatonne of CO₂ equivalent, and this share may well increase over time. A portion 23% of this share is caused by data centers, while communication networks (fixed and mobile) contribute a 24% of the total (Fig. 1.4). Therefore, it is clear that the need for green energy-efficient ICT will be more and more evident in the next years. This thesis investigates energy efficiency and greenness at the level of the Internet, data centers and web servers. Specifically, we study the use of RESs and scheduling techniques toward reducing the carbon footprint induced by an ICT system,

namely, the web crawling component of a web search engine. We focus on the reduction of the carbon emissions that this component incurs on remote web servers, that do not belong to the search engine, during its operations. Our approach exploits the time-varying renewable energy availability (i.e., the dynamic RES generation patterns) at remote servers in order to schedule the search engine workload in a greener way.

1.2 Contributions of this thesis

We briefly describe the main contributions of this thesis which are analytically presented in the following chapters:

In **Chapter 2**, we study two fundamental problems in energy storage management and dimensioning. First, we address the optimal energy storage control problem from the side of an energy supplier with the aim to minimize the long-term average cost of generated power. The energy supplier controller receives power demand requests which are immediately activated and controls one or more energy storage devices, each of which is attached to a renewable energy source (RES). The supplier uses these storage devices to mitigate time variations in demand load and RESs. For a single storage device, we derive a threshold-based control policy that maintains balanced grid power consumption at all times, and it is shown to be asymptotically optimal as the storage device capacity becomes large. The optimal policy is then extended to the case of multiple storage devices that are shared among multiple micro-grid entities each with its own demand load. Our results provide evidence about the potential of our approaches in terms of operational cost efficiency, in that they approach a cost global lower bound. Next, we study the joint energy storage placement, dimensioning and management problem. We are interested in the way storage capacity placement and control impacts the overall cost of energy generation. It turns out that various aspects of power flow need to be taken into account in the determination of the optimal policy. We discuss some simple special cases of the problem by presenting two numerical examples. Our results show that the solution policy entails various parameters such as the demand profiles of prosumers and power flow constraints.

In **Chapter 3**, motivated by a real-world problem, we consider energy efficiency in the context of a concrete application. Namely, we study the carbon footprint induced by the web crawling component of a large-scale web search engine. This component, during its page refresh operations, issues a large number of HTTP requests to remote web servers that do not belong to the search engine. These requests increase the energy consumption and carbon footprint of the web servers since computational resources are used while serving the requests. We focus on reducing the induced carbon footprint by exploiting

the renewable energy availability at remote servers. We model the dynamics of renewable energy generation by means of an appropriately defined time-varying (greenness) index. We introduce the problem of green web crawling, where the objective is to devise a page refresh policy that minimizes the total staleness of pages in the repository of a web crawler, subject to a constraint on the amount of carbon emissions due to the processing on web servers. For the case of one web server and one crawling thread, the optimal policy turns out to be a greedy one. At each iteration, the page to be refreshed is selected based on a metric that considers the page's staleness, its size, and the greenness of the energy consumed at the web server premises. We then extend the optimal policy to the cases of (i) many servers, (ii) multiple threads, and (iii) pages with variable freshness requirements. We present experimental results for the optimal page refresh policy as well as for various heuristics, in an effort to study the effect of different factors on performance. This work was done in collaboration with Yahoo! Labs Barcelona which provided us with a large, real-life dataset in order to evaluate the performance of the proposed policies.

Next, in **Chapter 4**, we theoretically model and study a simple gamification mechanism through which a demand-response entity aims to motivate consumers to adopt prudent energy consumption patterns by setting up a contest and by providing them incentives in the context of a serious game. Serious games are a promising approach for demand-side management, and they aim at increasing user engagement and active participation. We mathematically formulate the problem of optimal serious-game design for energy consumption reduction. The serious game designer presents publicly to consumers a list of top- K consumers and a list of bottom- M consumers according to their respective energy-consumption reduction at peak hours. The consumer parameters that determine user consumption behavior are the user discomfort due to demand load reduction, the user desire for social approval, and the user sensitivity to social outcasting. The consumer aims at maximizing her net expected utility function, while the game designer aims at minimizing the total power generation cost. Central in the selection of the game designer is the choice of parameters K , M as well as the statistical feedback provided to the consumers. We experimentally show how the choices of K , M affect the energy consumption reduction for different types of customers.

In **Chapter 5**, we migrate from conventional rational models for the consumer, and we bring into the foreground behavior-based data-driven models for modeling and predicting consumer behavior. We aim at understanding the different criteria under which consumers reach decisions regarding their participation in an energy consumption reduction campaign, and the different importance placed on them. Such criteria may be monetary incentives, social or community recognition, environment-friendly behavior, perceived dissatisfaction from consumption reduction, and willingness for cooperation.

They are expressed by means of attributes that characterize the consumption reduction suggestions made by the campaign designer to consumers. We consider that a consumer profile and her/his decision on whether to participate in the campaign, are determined by two controllable attributes of the suggestion made to her/him. We use a logistic regression model from machine learning and a fast-and-frugal tree model inspired from cognitive psychology and behavioral science to model the consumer decisions. We also introduce the optimal task and incentive allocation problem faced by the designer of the campaign which amounts to appropriately controlling the suggestion attributes in order to achieve a certain expected load saving at minimum economic cost. We show that in the case of LR the optimization problem turns out to be a sigmoid optimization one, while in the case of FFTs it is an integer linear programming one.

We conclude with **Chapter 6**, where we summarize our findings and discuss various future directions that warrant further investigation.

Chapter 2

Optimal Energy Storage Management and Dimensioning in Smart Grid Systems

Contents

2.1	Introduction to the Employment of Energy Storage Devices in the Smart Grid	9
2.1.1	Our contribution	10
2.2	Optimal Energy Storage Control Policies for the Energy Supplier	11
2.2.1	System model	11
2.2.2	Problem formulation	15
2.2.3	An Asymptotically Optimal Control Policy	16
2.2.4	Extension to the model: Renewable source	18
2.2.5	Multiple Storage Devices and Grid Entities	19
2.2.6	Extensions to the model	25
2.2.7	Numerical Results	27
2.3	Storage Placement and Power Flow	30
2.3.1	Power flow analysis	31
2.3.2	Storage and power flow model	32
2.3.3	Problem formulation	34
2.3.4	Numerical Example 1: Storage placement for N=2 buses	35
2.3.5	Numerical Example 2: Storage dimensioning for N=2 buses	38
2.4	Related Work	39
2.5	Conclusion	42

2.1 Introduction to the Employment of Energy Storage Devices in the Smart Grid

The smart power grid relies on information and communication technologies and advanced control methods to manage the dynamic demand load and to ensure efficient use of electric energy [10]. Smart metering and bidirectional communication enable real-time interconnection of the consumer and operator premises through IP addressable components over the Internet. These technologies allow power-consumption monitoring, automated control of consumption of customer appliances through messages from the operator Command and Control center, real-time electricity price signaling, and fault diagnosis.

Demand-load management is primarily employed by utility operators so as to reduce the grid operational costs. The rationale of demand-load control is to alleviate high demand load at peak times. This can be achieved for instance by using the time slack of delay-tolerant demands so as to temporally shift part of the peak load in time when it is feasible to do so [3]. Thus, the risk of a potential grid failure is reduced, while the operational cost is lowered by avoiding the use of more expensive or less efficient power generation sources.

Recent advances in electric energy storage technologies have rendered backup devices like uninterruptible power supply (UPS) or batteries, and Plug-in Hybrid Electric Vehicles (PHEVs) prime candidates for demand-load management. With appropriate storage management policies, these devices can be quite advantageous for electric utility operators and consumers. If stored energy control is delegated to the *grid operator*, a valid objective is to minimize the grid operational cost. Batteries can be charged at off-peak-load times, and this stored energy can be used to satisfy increased demand load at peak times. If energy storage management is performed at the *consumer* level (e.g. through PHEVs), the goal is to minimize the cost of power consumption (i.e., the electricity bill), assuming that an instantaneous time-of-use price per unit of consumed power is fed back. Thus, energy can be stored when the price per unit of consumed power is low, and it can be used to satisfy part of the demand when the price is high.

In future smart grid architectures, we anticipate the existence of multiple consumer entities, oftentimes organized as micro-grids, each with its own power demand, and its own power generation sources. In addition, energy storage devices with attached RESs are also expected to be part of the architecture. RESs rely mainly on energy that flows naturally through the environment on a continual but time-varying basis and have received major attention in the last years due to their significant potential in reducing

the carbon footprint. The potential of RESs is amplified when used in synergy with energy storage devices.

In this chapter, first, we address the problem of *optimal energy storage control* faced by the utility operator (energy supplier), which has a number of storage devices under its control. The objective is to find a policy for managing the charging and discharging processes of storage devices such that the long-term average cost of generated power is minimized. This cost is modeled as a convex function of instantaneous total power consumption. We study the problem above, first for a single storage device, and then for multiple storage devices that are shared among multiple micro-grids. Next, we review a joint energy storage placement, dimensioning and control problem, given an available storage budget [11], where the goal is to minimize the energy generation cost.

2.1.1 Our contribution

Our work contributes on several levels to the literature.

- For the case of a single energy storage device, we formulate the online storage control problem by devising a stochastic model for continually generated demands and completions, and we consider minimizing long-term average cost. A threshold-based control policy is derived that attempts to maintain balanced power consumption from the grid by adaptively managing the storage device charge/discharge processes. We prove that this policy is asymptotically optimal as the battery capacity becomes large, and we show that it performs quite well even for finite capacity.
- We also extend the model and structure of the optimal policy to account for the case that one renewable energy source feeds the battery.
- The approach is then extended to the case of multiple storage devices that are shared among multiple micro-grids each with its own demand load. First, we study the problem of minimizing the grid operational cost in the case that micro-grids act cooperatively in the sense that the total consumed load of micro-grids from the main grid is aggregated, and the total cost is due to this aggregate load. Then, we study the same problem in the case that each micro-grid is treated as a separate grid with its own consumed load and operational cost, and the total system cost is the sum of individual micro-grid operational costs. The proposed policy which is common for both problems is again of threshold type and is asymptotically optimal as storage capacities of batteries become large. These policies for the different cases we consider provide a systematic way of handling storage facilities.

- We study the problem of how much storage capacity should be placed on each node of a power network given an available storage capacity budget. It turns out that various aspects of power flow need to be taken into account in order to determine the optimal policy. Our objective is to minimize the total average generation cost through optimal storage dimensioning, storage management, and power flow. We provide two simple numerical examples in order to present the optimal solutions for some special cases of the problem.

The rest of the chapter is organized as follows. In section 2.2, first, we study the online control problem for a single storage device, we show the optimality of the proposed threshold policy, and we present an extension to the model that includes a RES which feeds the storage device. Then, we study the problem with multiple storage devices and multiple grid entities. Subsection 2.2.6 provides hints for extensions of the proposed models. In subsection 2.2.7, we evaluate the performance of our policies and present numerical results. In section 2.3, the joint energy storage placement, dimensioning and control problem is studied. Finally, section 2.4 presents the related literature and section 2.5 concludes our study. This chapter contains material from works [12], [13] and [14]. The terms “battery” and “storage device” are used interchangeably with the same meaning in the chapter.

2.2 Optimal Energy Storage Control Policies for the Energy Supplier

2.2.1 System model

Power demand arrival and service processes. We consider an energy supplier (energy generation facility) that needs to serve the power demand tasks generated by a pool of consumers. In the online version of the problem, power demand requests are generated continually by consumers and arrive at a central controller at the supplier premises according to a Poisson process, with rate λ .² The time duration s_n of each request n is a random variable that is exponentially distributed with parameter s , i.e.,

$$\Pr(s_n \leq x) = 1 - e^{-sx}, x \geq 0. \quad (2.1)$$

²A non-homogeneous Poisson process with time-varying request arrival rate $\lambda(t)$ could model time-variation of the demand load and affect load variation and peaks at different times. Here, we choose to adhere to a homogeneous Poisson process assumption with average rate λ . Since we study the system in an infinite time-horizon, our approach encompasses the case of time-varying $\lambda(t)$ if $\lim_{T \rightarrow +\infty} \frac{1}{T} \sum_{t=0}^{T-1} \lambda(t)$ exists and is equal to λ .

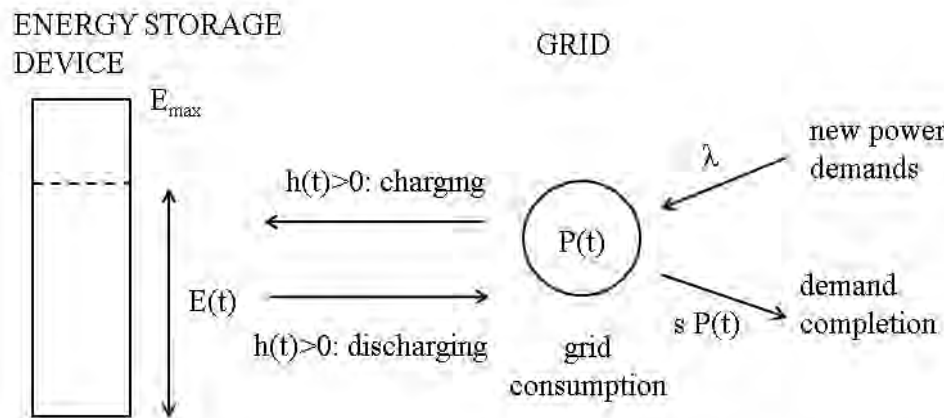


FIGURE 2.1: Overview of system model with the energy storage device, the charging and discharging process, and the interaction with grid consumption.

Equivalently, the mean duration of demand requests is $1/s$ time units, and s may be seen as the average service rate of power demand tasks by the supplier. These assumptions are motivated for mathematical tractability as they facilitate the derivation of the structure of the optimal policy. However, they are quite close to reality since they capture (i) the burst of arriving demand requests, (ii) the different durations of requests and (iii) the fact that the chances of having large durations decrease fast. Further, let p_n denote the power requirement (in Watts) of demand n . Denote by $P(t)$ the demand load on the system at time t as a result of the demand arrival and completion processes above.

Consider the system above, first without the storage device. Let \mathcal{A}_t denote the set of active demands at time t . If the power requirement p_n of each task n , is fixed and equal to 1 unit of power, then the instantaneous demand load $P(t)$ equals $N(t) = |\mathcal{A}_t|$, i.e. the number of active demand tasks at t , and it is a Markov chain. Since each power demand task is activated upon arrival and there exists no waiting time or loss, $P(t)$ is the *occupation process of an $M/M/\infty$ service system*. From state $P(t)$, there are transitions to state:

- $P(t) + 1$ with rate λ , when a new demand request arrives.
- $P(t) - 1$ with rate $sP(t)$, when one of the current $P(t)$ active demands is completed.

The steady-state probabilities of $P(t)$, $q_i = \lim_{t \rightarrow \infty} \Pr(P(t) = i)$, for $i = 1, 2, \dots$, are derived from equilibrium equations [15, Sec.3.4.2] as

$$q_i = \left(\frac{\lambda}{s}\right)^i \frac{e^{-\lambda/s}}{i!}. \quad (2.2)$$

Thus, $P(t)$ is Poisson distributed with parameter $\frac{\lambda}{s}$, and the expected number of active requests at steady state is $\mathbb{E}[P(t)] = \frac{\lambda}{s}$, where the expectation is with respect to the stationary Poisson distribution of $P(t)$.

The extension to different power requirements goes as follows. Suppose that the power requirement \hat{P}_n of each demand n is a random variable with a discrete probability distribution on set $\{p_1, \dots, p_L\}$, with associated probabilities w_1, \dots, w_L (the case of continuous distribution of \hat{P}_n is tackled similarly). Random variables \hat{P}_n , are independent from process $N(t)$. Let \hat{P} denote the random power requirement of a demand. Also, let $\mathbb{E}[\hat{P}] = \mathbb{E}[\hat{P}_n] = \sum_{k=1}^L w_k p_k$ be the expected power requirement of a demand n . The power consumption at slot t is $P(t) = \sum_{n \in \mathcal{A}_t} \hat{P}_n$, and the average power consumption at steady state is $\mathbb{E}[P(t)] = \mathbb{E}[N(t)]\mathbb{E}[\hat{P}] = \frac{\lambda}{s}\mathbb{E}[\hat{P}]$.

Time-slotted model: We turn the continuous-time model above to a time-slotted one. Let $A(t)$ be the process that describes the total number of request arrivals from time 0 to time t . Then, the number of arrivals at time interval $(t_1, t_2]$ is $A(t_2) - A(t_1)$. Assume a slotted system with slot length $\tau = 1$. The number of arrivals is Poisson distributed with parameter λ , that is

$$\Pr[A(t+1) - A(t) = n] = e^{-\lambda} \frac{\lambda^n}{n!}. \quad (2.3)$$

Similarly, let $D(t)$ be the process that describes the total number of request completions from time 0 to time t . The number of completions is Poisson distributed with parameter s , that is

$$\Pr[D(t+1) - D(t) = n] = e^{-s} \frac{s^n}{n!}. \quad (2.4)$$

Energy Storage Device. There exists an energy storage device (battery) of storage capacity E_{max} kWh that is located at the supplier premises and is controllable by the energy supplier. At each time slot t , the battery may be charging or discharging. Let $E(t)$ be the stored amount of energy at the battery at the beginning of slot t . Define the decision variable $h(t)$ as the *rate* at which the battery is charged/discharged at slot t , with the following convention. If $h(t) > 0$, the battery charges, and energy flows into it from the energy generation facility with rate $h(t)$. This in turn implies that the total power demand load to be served from the grid is $h(t)$ plus the power demand $P(t)$, i.e. $P(t) + h(t)$. If $h(t) < 0$, the battery is discharged, namely energy flows out of the battery at rate $|h(t)|$, and it is used to serve part of the demand load $P(t)$. Hence, the amount of demand load that is actually served through the energy generation facility is $P(t) + h(t) < P(t)$, since an amount $|h(t)|$ of the power demand is served directly from the battery. The system model is depicted in Fig. 2.1.

The level of stored energy at the battery evolves with time as:

$$E(t + 1) = E(t) + h(t). \quad (2.5)$$

Since it is $0 \leq E(t) \leq E_{max}$ at all slots t , and time slots are of unit size, it becomes evident that $h(t)$ must satisfy:

$$-E(t) \leq h(t) \leq E_{max} - E(t), \forall t. \quad (2.6)$$

Remark: In order to demonstrate the structure of our approach, we assume there are no constraints on the maximum charge and discharge rate other than those implied by (2.6). We also neglect the power losses during the charge and discharge process as these could be easily absorbed in the charge and discharge rate. Finally, we assume that there exists no switching delay for transferring power demands from the grid to the battery and vice versa. The former transfers take place when we decide to discharge the battery, and the latter occur when the battery empties while some tasks are served. We discuss later how these and other modeling constraints can be incorporated in the problem and affect the policy.

Cost Model. Let $X(t) = P(t) + h(t)$ denote the total demand load to be served from the energy supplier at slot t . We denote the instantaneous *operational/generation cost* associated with load $X(t)$ as $C(X(t))$, where $C(\cdot)$ is an increasing, differentiable convex function. Convexity of $C(\cdot)$ reflects the fact that the *differential* or *marginal* cost of power generation for the supplier increases as the demand load increases. That is, each unit of additional power that is needed to satisfy the increasing demand becomes more expensive to generate/obtain and make available to the consumers.

For example, the following average energy generation marginal costs are reported in United States: Geothermal, 47 \$/MWh; wind on-shore, 73 \$/MWh; natural gas with combined cycle, 75 \$/MWh; hydro, 83 \$/MWh; nuclear, 95 \$/MWh; natural gas with combustion turbine, 141 \$/MWh; solar photo-voltaic, 125 \$/MWh; wind offshore 197 \$/MWh. Electricity suppliers have several different generation units at their disposal. In order to cover the demand, suppliers activate units in increasing order of marginal energy generation cost. Namely, when the energy generation capacity of one unit is reached, they activate the next more expensive one. Thus, supplementary power for serving high demand may be generated from expensive sources, or it may be imported at high prices from other countries.

The form of the convex cost function may be considered as a piecewise linear one (differentiable at every point), in which each linear segment represents one alternative energy




FIGURE 2.2: Power generation cost as a piece-wise linear convex function of demand load.

generation source, and the slopes of linear segments are the respective marginal generation costs (Fig. 2.2).

2.2.2 Problem formulation

At the beginning of each slot t , the controller observes the total grid consumption level $X(t)$ (and thus, also $P(t)$), as well as the current energy level $E(t)$ in order to decide whether it will charge or discharge the battery and how much.

Denote the system state at slot t as $\mathbf{x}(t) = (P(t), E(t))$. For now, assume that all quantities are restricted to positive integer values. Let the initial state be $(P(0), E(0))$. A horizon of T time slots is assumed. The objective is to find a battery charge-discharge control policy that minimizes the long-term average operational cost,

$$\lim_{T \rightarrow +\infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}_t [C(X(t))] = \mathbb{E}[C(X(t))], \quad (2.7)$$

where the first expectation above is with respect to the randomness of $P(t)$, and the second one is with respect to the stationary distribution of $\{P(t)\}$. A policy π selects control variables $\{h(t)\}_{t=0,1,\dots}$, subject to the evolution equation (2.5) and constraint (2.6). A policy π^* is optimal if it minimizes the long-term average cost (2.7) over all policies that satisfy the constraints above.

Under the assumptions above on arrival and completion processes, the problem of may be cast as a Markov Decision Process (MDP) one. Instead of tackling the problem through solving an MDP formulation, our focus here will be on simple control policies

that operate based on $P(t)$ as the system state. We obtain a simple policy and show that it is asymptotically optimal for large storage capacity values.

2.2.3 An Asymptotically Optimal Control Policy

Consider the following battery charge-discharge control policy, which we refer to as Policy (P). There exists a demand threshold, P_0 . At the beginning of each slot t , the controller checks the amount of current demand to be fulfilled, $P(t)$. Recall that the current demand consists of all demand requests that are active, i.e., they are not yet completed.

- If $P(t) \leq P_0$, then all current demand load is served with the available energy from the grid, and a decision to *charge* the battery is taken, at a rate

$$h(t) = P_0 - P(t) \geq 0. \quad (2.8)$$

- If $P(t) > P_0$, a decision to *discharge* the battery is taken, at a rate

$$h(t) = P(t) - P_0. \quad (2.9)$$

Hence, a portion P_0 of the demand load is served by the available energy from the grid, while the rest, $P(t) - P_0$ is satisfied by the battery.

Whenever the battery energy level becomes zero during discharging, the controller serves from the grid the rest of the load that the battery was serving at that slot t , and it holds $E(t+1) = 0$. Since $P_0 - P(t)$ can take positive or negative values, the policy above can be succinctly described as:

$$h(t) = \max\{-E(t), \min\{P_0 - P(t), E_{max} - E(t)\}\}. \quad (2.10)$$

Theorem 2.1. *Policy (P) is asymptotically optimal, in the sense that its performance converges to the lower bound (2.11) as $E_{max} \rightarrow \infty$, and therefore it minimizes the long-term average cost (2.7).*

Proof. We provide an intuitive sketch of the proof. We need to realize that, regardless of the charging/discharging decisions, it is $\mathbb{E}[X(t)] = \mathbb{E}[P(t)] = \frac{\lambda}{s}$. That is, charging/discharging decisions may change the instantaneous values and steady-state probabilities of process $X(t)$, however the average total active demand will remain fixed and equal to $\frac{\lambda}{s}$. This is because there are no charging or discharging losses; such losses would exist if the battery was empty during discharge or it was full during charge.

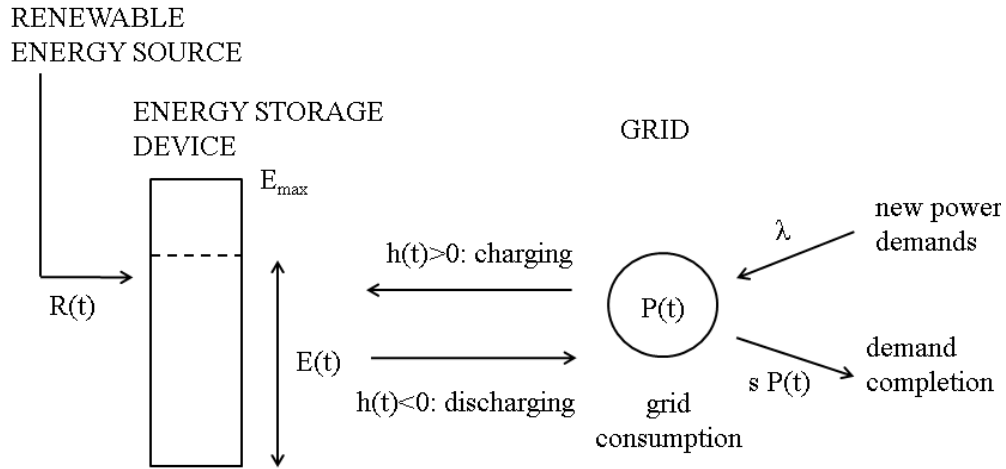


FIGURE 2.3: Extension to the model, with a renewable source feeding the battery.

To make this more intuitive, note that as the available storage capacity grows larger, more energy can be stored when there is opportunity to do so (i.e. the values of process $\{E(t)\}$ increase in general as well), which in turn implies that there will almost always exist enough stored energy to discharge. Thus, as E_{max} increases, the process $h(t)$ takes values in the respective set of feasible controls such that $-E(t) \leq h(t) \leq E_{max} - E(t)$ with probability that approaches 1. Hence, the events $\{P(t) - P_0 \geq E(t)\}$ and $\{P_0 - P(t) \geq E_{max} - E(t)\}$ have diminishing probability as $E_{max} \rightarrow \infty$. Thus, $h(t) < E_{max} - E(t)$ and $h(t) > -E(t)$ almost surely as $E_{max} \rightarrow \infty$. These imply that the battery will almost never go empty during discharge and it will almost never get full during charge.

From Jensen's inequality, for a random variable X and a convex function $C(\cdot)$, it is $\mathbb{E}[C(X)] \geq C(\mathbb{E}[X])$. Equality holds if and only if $X = \mathbb{E}[X]$, i.e., when random variable X is constant. In our setup, Jensen's inequality implies:

$$\mathbb{E}[C(X(t))] \geq C(\mathbb{E}[X(t)]). \quad (2.11)$$

The policy above maintains $\mathbb{E}[X(t)] = P_0$. If we combine this with the rationale above, we get that the optimal threshold must be $P_0 = \mathbb{E}[P(t)] = \frac{\lambda}{s}$.

□

As a byproduct of the above, we have that $\lim_{E_{max} \rightarrow \infty} \mathbb{E}[h(t)] = 0$, since the battery charges with rate $P_0 - P(t)$ whenever $P(t) < P_0$, and it discharges with rate $P(t) - P_0$ whenever $P(t) > P_0$. For $P_0 = \frac{\lambda}{s}$, the charging and discharging events take place for half the amount of time each on average in the long-run, and thus $\mathbb{E}[h(t)] \rightarrow 0$.

2.2.4 Extension to the model: Renewable source

We now consider the following extension to the model. There exists a RES which feeds the battery. Denote by $R(t)$ the renewable energy generation process, which is assumed to be an arbitrary (wide-sense) stationary process with expectation $\bar{R} = \mathbb{E}[R(t)]$. The rest of the model is the same as the one presented in subsection 2.2.1, and the system is depicted in Fig. 2.3.

Now, the battery energy level evolution equation is,

$$E(t+1) = E(t) + R(t) + h(t), \quad (2.12)$$

and the feasible control set at slot t is such that,

$$-R(t) - E(t) \leq h(t) \leq E_{max} - E(t) - R(t). \quad (2.13)$$

Note that this system is equivalent to one without the RES, where $h(t)$ is substituted by $h(t) - R(t)$.

Again, our objective is to minimize the long-term average energy generation cost $\mathbb{E}[C(X(t))]$. From Jensen's inequality, it is $\mathbb{E}[C(P(t) + h(t))] \geq C(\bar{P} + \bar{h})$, where $\bar{h} = \mathbb{E}[h(t)]$ and $\bar{P} = \mathbb{E}[P(t)]$. From (2.12), we have:

$$\begin{aligned} \lim_{T \rightarrow +\infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}_t[E(t+1)] - \lim_{T \rightarrow +\infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}_t[E(t)] &= \bar{R} + \bar{h} \Rightarrow \\ \lim_{T \rightarrow +\infty} \frac{1}{T} (\mathbb{E}[E(T)] - \mathbb{E}[E(0)]) &= \bar{R} + \bar{h} \Rightarrow \bar{h} = -\bar{R}. \end{aligned} \quad (2.14)$$

which means that in the long run, the average amount (dis)charged from the battery is equal to the average amount of energy generated from the renewable source. Note also that for $\bar{R} = 0$, equation (2.14) is a rigorous way of showing $\bar{h} = 0$ in the Theorem above. Thus,

$$\mathbb{E}[C(P(t) + h(t))] \geq C(\bar{P} - \bar{R}). \quad (2.15)$$

This holds with equality if and only if:

$$P(t) + h(t) = \bar{P} - \bar{R} \Rightarrow h(t) = \bar{P} - \bar{R} - P(t).$$

Therefore, the policy that minimizes long-term average cost for the setup with a RES is as follows:

- If $P(t) \leq \bar{P} - \bar{R}$, charge the battery at rate $\bar{P} - \bar{R} - P(t)$.

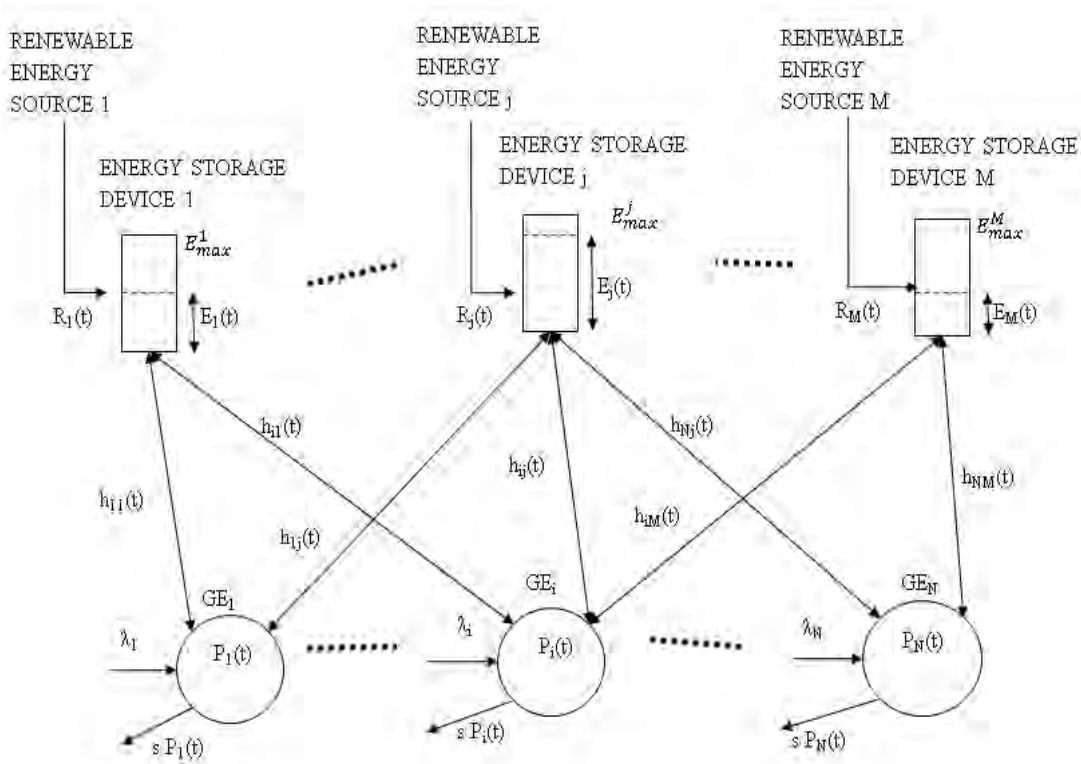


FIGURE 2.4: System with multiple storage devices and appended renewable energy sources.

- If $P(t) > \bar{P} - \bar{R}$, discharge the battery at rate $P(t) - \bar{P} + \bar{R}$.

Observe that if $\bar{P} - \bar{R} < 0$, battery (dis)charge does not take place since in that case the RES output fully covers the power demand load requirements.

2.2.5 Multiple Storage Devices and Grid Entities

System Architecture. Consider now N grid entities (GEs). A first instance of GE may be an aggregator that aggregates demand load requests of consumers and acts as an intermediary between the utility operator and consumers. The set of all GEs results in a total consumed load from the main grid (and hence the need to generate this load), and the total cost is due to the needed amount of energy in order to satisfy this total load. A second instance of GE may be that of a micro-grid with its own power demand and autonomous power generation, and hence its own operational cost. In that case, each micro-grid is treated as a separate grid with its own demand load and operational cost, and the cost of total generated power for the entire system is the sum of individual micro-grid (GE) operational costs. There also exist M energy storage devices, each with an attached RES. These storage devices are shared among the N GEs, but they

are owned and controlled by the central operator. That is, they can be charged from power generated by the GEs, or they can be discharged to satisfy demand load in GEs. Continual arrival and completion of power demands takes place at each GE location. The system is depicted in Fig. 2.4.

Power Demand Model. Power demand requests arrive at each GE i according to a Poisson process, with average arrival rate λ_i requests per unit of time. The time duration s_n of each power demand task n is exponentially distributed with parameter s and $p_n = 1, \forall n$. With no loss of generality, the average service rate s is taken to be the same for demand tasks of all GEs. Since $p_n = 1, \forall n$, let $P_i(t)$ be the number of active requests on GE i at time t and it is a continuous-time Markov chain. Therefore, the expected number of active requests on GE i at steady state is $\bar{P}_i = \mathbb{E}[P_i(t)] = \frac{\lambda_i}{s}$. The extension to different power requirements is similar to the one described in section 2.2. In this case, the average power consumption on GE i at steady state is $\bar{P}_i = \mathbb{E}[P_i(t)] = \mathbb{E}[N_i(t)]\mathbb{E}[\hat{P}] = \frac{\lambda_i}{s}\mathbb{E}[\hat{P}]$, where $N_i(t)$ is the number of active requests on GE i at t . The process to turn the continuous-time model to a time-slotted one is similar to that described in section 2.2.

Energy Storage Devices. Each storage device $j, j = 1, \dots, M$ has storage capacity E_{max}^j kWh, $j = 1, \dots, M$, and it is accessible by some or all GEs. Denote by \mathcal{K}_i the subset of storage devices that are accessible to GE i . Let $R_j(t)$ be the time-varying amount of generated energy at RES j which is appended to storage device j . This is assumed to be an arbitrary (wide-sense) stationary process with expectation $\bar{R}_j = \mathbb{E}[R_j(t)]$. Renewable source j is appended to storage device j in the sense that its energy feeds storage device j . At the beginning of each slot t , the central controller decides whether each GE i will charge or discharge each storage device j , and by how much. The charging and discharging processes for each battery are determined based on the GE power demand levels and the renewable energy generation levels. Specifically, when the demand load of a GE i is low, it may charge one or more batteries, since the energy generation cost is low. On the other hand, at peak consumption, when the cost of power generation is high, GE i may satisfy part of its increased demand load using part of the stored energy of its accessible batteries.

Let $E_j(t)$ be the amount of stored energy at storage device j at the beginning of slot t . Define the decision variable $h_{ij}(t)$ as the charging/discharging rate for storage device j and GE i at slot t . If $h_{ij}(t) > 0$, battery j is charged, and energy flows into it from GE i with rate $h_{ij}(t)$. If $h_{ij}(t) < 0$, battery j discharges, namely energy flows out of the battery with rate $|h_{ij}(t)|$, and it is used to serve part of the demand load of GE i .

Define $h_i(t)$ as the total charging rate of all batteries that are accessible by GE i ,

$$h_i(t) = \sum_{j \in \mathcal{K}_i} h_{ij}(t), \quad (2.16)$$

and expectation $\bar{h}_i = \mathbb{E}[h_i(t)]$. Hence, for GE i , the amount of demand load that is actually served by the grid (or GE i) at slot t is $X_i(t) = P_i(t) + h_i(t)$, with expectation $\bar{X}_i = \mathbb{E}[X_i(t)]$.

Further, let $h_{tot,j}(t)$ be the net rate at which energy flows into the storage device j at slot t ,

$$h_{tot,j}(t) = \sum_{i=1}^N h_{ij}(t). \quad (2.17)$$

Note that a battery j may be charged from some GEs and discharged from some others at the same time, which means that it can be used in a different way by each GE at each slot t . For instance, for battery j it can be $h_{kj}(t) < 0$ and $h_{lj}(t) > 0$ for two different GEs k and l . The energy level of storage device j evolves as,

$$E_j(t+1) = E_j(t) + R_j(t) + h_{tot,j}(t) \quad (2.18)$$

and the feasible control set at slot t is such that,

$$-R_j(t) - E_j(t) \leq h_{tot,j}(t) \leq E_{max}^j - E_j(t) - R_j(t), \forall t. \quad (2.19)$$

By using a similar argument as for the case of a single storage device, we get

$$\sum_{i=1}^N \bar{h}_i = - \sum_{j=1}^M \bar{R}_j. \quad (2.20)$$

We distinguish two different objectives that make sense in future smart grid architectures. In case A, the different micro-grids receive energy from one energy generator entity. In this case, the total energy generation cost amounts to the generation cost of the aggregate demand load of all micro-grids. In case B, each micro-grid is treated as a separate grid with its own energy generation facility. Now a plausible objective is to minimize the total cost of energy generation for the system, which is the sum of individual generation costs for each micro-grid (GE). In the sequel, we formulate and discuss the two corresponding optimization problems.

2.2.5.1 Case A: Minimization of cost of aggregate demand

In case A, the demand loads of all GEs are aggregated, and the controller aims to minimize the energy generation cost so as to satisfy the total demand, by determining appropriate charging and discharging rates for each battery.

The generation cost associated with total power consumption $X(t) = \sum_{i=1}^N X_i(t)$, is given as:

$$C(X(t)) = C\left(\sum_{i=1}^N X_i(t)\right) = C\left(\sum_{i=1}^N (P_i(t) + h_i(t))\right). \quad (2.21)$$

Moreover, for the long-term average cost, we have the following lower bound from Jensen's inequality:

$$\begin{aligned} \mathbb{E}\left[C\left(\sum_{i=1}^N X_i(t)\right)\right] &\geq C\left(\sum_{i=1}^N (\mathbb{E}[P_i(t)] + \mathbb{E}[h_i(t)])\right) = \\ &\stackrel{(2.20)}{=} C\left(\sum_{i=1}^N \bar{P}_i - \sum_{j=1}^M \bar{R}_j\right). \end{aligned} \quad (2.22)$$

Optimal policy for $N = 2$ GEs and $M = 2$ storage devices:

From (2.22),

$$\mathbb{E}[C(X_1(t) + X_2(t))] \geq C(\bar{P}_1 + \bar{P}_2 - \bar{R}_1 - \bar{R}_2), \quad (2.23)$$

and equality holds if

$$X_1(t) + X_2(t) = \bar{P}_1 + \bar{P}_2 - \bar{R}_1 - \bar{R}_2. \quad (2.24)$$

The conditions that make the inequality above hold with equality are:

$$\begin{aligned} P_1(t) + h_1(t) &= \bar{P}_1 - \alpha(\bar{R}_1 + \bar{R}_2) \Rightarrow \\ h_1(t) &= \bar{P}_1 - \alpha(\bar{R}_1 + \bar{R}_2) - P_1(t) \end{aligned} \quad (2.25)$$

and

$$\begin{aligned} P_2(t) + h_2(t) &= \bar{P}_2 - (1 - \alpha)(\bar{R}_1 + \bar{R}_2) \Rightarrow \\ h_2(t) &= \bar{P}_2 - (1 - \alpha)(\bar{R}_1 + \bar{R}_2) - P_2(t), \end{aligned} \quad (2.26)$$

where $\alpha \in [0, 1]$. The quantities

$$P_0^1 = \bar{P}_1 - \alpha(\bar{R}_1 + \bar{R}_2) \quad (2.27)$$

and

$$P_0^2 = \bar{P}_2 - (1 - \alpha)(\bar{R}_1 + \bar{R}_2) \quad (2.28)$$

show one possibility for setting the thresholds for the optimal policy for GE 1 and GE 2, respectively for charging and discharging. For $\alpha = 1/2$, the thresholds are:

$$P_0^1 = \bar{P}_1 - \frac{\bar{R}_1 + \bar{R}_2}{2} \quad (2.29)$$

and

$$P_0^2 = \bar{P}_2 - \frac{\bar{R}_1 + \bar{R}_2}{2} \quad (2.30)$$

2.2.5.2 Case B: Minimization of sum of individual GE costs

We assume that each GE represents an separate entity with its own power generation capabilities and power demand, and thus its own operational cost. Hence, the demand loads of the different GEs cannot be aggregated, and the meaningful objective for the controller is to minimize the sum of individual GE costs. The cost model is the same as the one above, but now each GE i has its own operational cost function, $C_i(\cdot)$.

The instantaneous generation cost for GE i associated with power consumption $X_i(t)$ is $C_i(X_i(t))$, and the total long-run average total cost is $\lim_{T \rightarrow +\infty} \mathbb{E}[\frac{1}{T} \sum_{t=0}^{T-1} \sum_{i=1}^N C_i(X_i(t))]$.

From Jensen's inequality, we get:

$$\begin{aligned} \mathbb{E}[\sum_{i=1}^N C_i(X_i(t))] &= \sum_{i=1}^N \mathbb{E}[C_i(P_i(t) + h_i(t))] \geq \\ &\geq \sum_{i=1}^N C_i(\bar{P}_i + \bar{h}_i) \geq \\ &\geq \min_{\bar{h}_i} \left\{ \sum_{i=1}^N C_i(\bar{P}_i + \bar{h}_i) \right\}, \end{aligned} \quad (2.31)$$

where the last inequality follows from the selection of values for \bar{h}_i , $i = 1, \dots, N$ that minimize $\sum_{i=1}^N C_i(\bar{P}_i + \bar{h}_i)$.

Optimal policy for $N = 2$ GEs with the same cost function and $M = 2$ storage devices:

From (2.20), it is

$$\bar{h}_1 + \bar{h}_2 = -\bar{R}_1 - \bar{R}_2. \quad (2.32)$$

Substituting from (2.32), the right side of (2.31) becomes:

$$\min_{\bar{h}_1} \{ C(\bar{P}_1 + \bar{h}_1) + C(\bar{P}_2 - \bar{R}_1 - \bar{R}_2 - \bar{h}_1) \}. \quad (2.33)$$

Setting the derivative of the term above with respect to h_i equal to zero and assuming that $C'(\cdot)$ is a one-to-one function, we get

$$\bar{h}_1^* = \frac{\bar{P}_2 - \bar{P}_1 - \bar{R}_1 - \bar{R}_2}{2} \quad (2.34)$$

and

$$\bar{h}_2^* = \frac{\bar{P}_1 - \bar{P}_2 - \bar{R}_1 - \bar{R}_2}{2}. \quad (2.35)$$

From (2.31), we have

$$\sum_{i=1}^2 \mathbb{E}[C(P_i(t) + h_i(t))] \geq \sum_{i=1}^2 C(\bar{P}_i + \bar{h}_i^*) \quad (2.36)$$

and equality holds if

$$P_1(t) + h_1(t) = \bar{P}_1 + \bar{h}_1^* \quad (2.37)$$

and

$$P_2(t) + h_2(t) = \bar{P}_2 + \bar{h}_2^*. \quad (2.38)$$

Therefore, using (2.34) and (2.35), we have the charging / discharging decision control for the two micro-grids given by

$$h_i(t) = \frac{\bar{P}_1 + \bar{P}_2 - \bar{R}_1 - \bar{R}_2}{2} - P_i(t), \quad i = 1, 2. \quad (2.39)$$

The quantity $P_0^i = \frac{\bar{P}_1 + \bar{P}_2 - \bar{R}_1 - \bar{R}_2}{2}$, which is the same for both GEs, will be the threshold for the asymptotically optimal policy.

2.2.5.3 Asymptotically Optimal Control Policy

We now describe some meaningful policies for a system with multiple GEs. There exist N thresholds, P_0^i , one for each GE i . At the beginning of each slot t , the controller checks each $P_i(t)$. If $P_i(t) \leq P_0^i$, then all active demand requests on GE i are served by the grid (or GE i), and a decision to charge one or more batteries is taken, with total charging rate $h_i(t) = P_0^i - P_i(t)$. On the other hand, if $P_i(t) > P_0^i$, a decision to discharge one or more batteries is taken, with total rate $h_i(t) = P_i(t) - P_0^i$. Whenever the battery j energy level becomes zero during discharging, the controller sets $h_{ij}(t) = 0$ for all GEs i that were discharging battery j at that slot t (namely, it serves from these GEs the rest of their loads that battery j was serving at slot t), and it holds $E_j(t+1) = 0$.

The key difference between the policy presented in section 2.2 and the one presented here is that each charging (or discharging) rate $h_i(t)$ must now be allocated among

the different batteries based on certain criteria pertaining to the storage devices. For instance, each charging or discharging rate may be split equally among the batteries. In this case, it is $h_{ij}(t) = \frac{h_i(t)}{|\mathcal{K}_i|}$, $\forall j \in \mathcal{K}_i$. As a second case, we may use the amount of stored energy $E_j(t)$ as a criterion. We may also use the intuition that batteries with smaller amount of stored energy and larger capacity need to be charged with higher charging rate, and that batteries with higher battery energy level should be discharged with a higher rate so as to avoid the case of empty or fully charged batteries that cannot be further discharged or charged, respectively. Based on the criterion above, each GE i could assign time-varying weights $w_{ij}^c(t)$ and $w_{ij}^d(t)$ to each storage device $j \in \mathcal{K}_i$ in case of charging and discharging, respectively that take values between 0 and 1. For one GE i , the controller would then set $h_{ij}(t) = h_i(t)w_{ij}^c(t)$ or $h_{ij}(t) = h_i(t)w_{ij}^d(t)$, $\forall j \in \mathcal{K}_i$, taking also into account constraint (2.19). Observe that $\sum_{j \in \mathcal{K}_i} h_{ij}(t) = h_i(t)$ since $\sum_{j \in \mathcal{K}_i} w_{ij}^c(t) = 1$ and $\sum_{j \in \mathcal{K}_i} w_{ij}^d(t) = 1$.

The expressions for the thresholds that we have provided in subsections 2.2.5.1 and 2.2.5.2 can be generalized for $N, M > 2$, and they are:

$$P_0^i = \bar{P}_i - \frac{\sum_{j=1}^M \bar{R}_j}{N}, \quad \forall i \quad (2.40)$$

and

$$P_0^i = \frac{\sum_{i=1}^N \bar{P}_i - \sum_{j=1}^M \bar{R}_j}{N}, \quad \forall i \quad (2.41)$$

for the first and second problem, respectively.

Theorem 2.2. *For both the cases above, the proposed policy with the above-mentioned thresholds is asymptotically optimal, in the sense that its performance converges to the lower bound as $E_{max}^j \rightarrow \infty$, $\forall j$, and therefore it minimizes the long-term average total cost.*

2.2.6 Extensions to the model

2.2.6.1 Non-Stationarity of Renewable Energy Generation

In subsection 2.2.4, we assumed that the RES that feeds the battery, is characterized by a (wide-sense) stationary energy generation process. This assumption does not affect the generality of the model, since for time-varying average generated power $\bar{R}(t)$ we assume that it holds $\bar{R} = \lim_{T \rightarrow +\infty} \frac{1}{T} \sum_{t=0}^{T-1} \bar{R}(t)$ and that the limit exists.

However, some RESs like solar and wind are characterized by variations of energy generation and might exhibit non-stationary behavior. In order to tackle this non-stationary

behavior, we propose an approach similar to that presented in [16]. Namely, we derive an empirical estimate for the time-varying average of generated power based on the most recent historical data, by using a moving time window of length L .

Let us denote by \tilde{P} and \tilde{R} the empirical estimate of the average of power demand $P(t)$ and the empirical estimate of the average of generated energy $R(t)$, respectively. These quantities are estimated from the data gathered over the L most recent observation periods, i.e. the actual average power demands $\bar{P}_{d_1}, \dots, \bar{P}_{d_L}$, and the actual average amounts of generated energy $\bar{R}_{d_1}, \dots, \bar{R}_{d_L}$, respectively. We have: $\tilde{P} = (\bar{P}_{d_1} + \dots + \bar{P}_{d_L})/L$ and $\tilde{R} = (\bar{R}_{d_1} + \dots + \bar{R}_{d_L})/L$.

Based on the above, a *heuristic* extension of the optimal energy storage control policy can be provided by keeping the structure of the optimal policy, and by setting $\bar{P} = \tilde{P}$ and $\bar{R} = \tilde{R}$. In the same sense, a *heuristic* policy can be devised for multiple GEs and batteries.

2.2.6.2 Storage Device Modeling

Charge/discharge rate bounds and losses: In this work, we assumed that there are no constraints on charge/discharge rates other than those implied by (2.6), (2.13) or (2.19), and we neglected losses during the charge and discharge process. In reality, a storage device has bounded charge and discharge rate. One could derive various extensions to the basic policy presented above that are expected to perform close to the optimal. For instance, if $h_B > 0$ denotes the absolute upper bound on charge or discharge rate, i.e. $|h(t)| < h_B$, then one could consider the following policy:

$$h(t) = \max\{-E(t), P_0 - P(t), h_B\}, \quad (2.42)$$

whose performance approaches the optimal one as h_B increases. Furthermore, a storage device may have charge and discharge power losses. Battery performance is affected by the charging and discharging efficiencies $\eta_c, \eta_d \in (0, 1]$ of the storage technology used. Namely the power flowing in or out of the battery at slot t is $\eta_c h(t)$ or $\frac{1}{\eta_d} h(t)$, respectively.

2.2.6.3 Transmission Losses

One could incorporate to the model possible transmission losses as the power is transported from storage devices to GEs. The power loss function for transported amount of power x is $P^{loss}(x) = \beta x^2$, where $\beta = \frac{Rd}{V^2}$. In this expression, R is the ohmic resistance

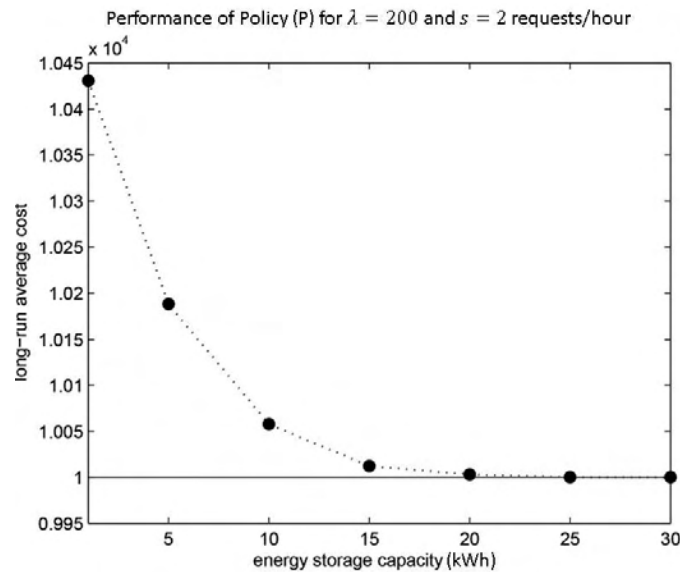


FIGURE 2.5: Performance of the proposed energy storage control policy as a function of available storage capacity.

per unit distance of the transmission line, V is the operating voltage of the transmission line that connects the two end-points of the transport, i.e the storage device and the GE, and d is the distance between them. Typical values of R range from 0.1 to 0.5 Ohm/km, while V could be 220V or tens of kV. Depending on the distance considered, the losses could be either ignored or incorporated in the model. The structure of the optimal policy will not change.

2.2.7 Numerical Results

2.2.7.1 One Storage Device

In order to evaluate the performance of the proposed policy (P) for a single storage device, we first compute the long-term average energy generation cost and compare it to the lower bound (2.11). Our simulation scenario ran for a horizon of 240 hours. The power demand arrival and completion processes are Poisson with $\lambda = 200$ requests / hour and $s = 2$ requests / hour, i.e. the average demand duration is $1/s = 1/2$ hour. Also, the average power requirement per demand is 1 kW. Thus, the average demand load is $E[P(t)] = \frac{\lambda}{s} = 100$ kW. We consider here the cost function $C(x) = x^2$ for simplicity.

Fig. 2.5 shows the average generation cost resulting from policy (P) for different values of storage capacity E_{max} . As anticipated, the average cost decreases as E_{max} increases, and it ultimately converges to the lower bound $(\lambda/s)^2$. It can be observed that the

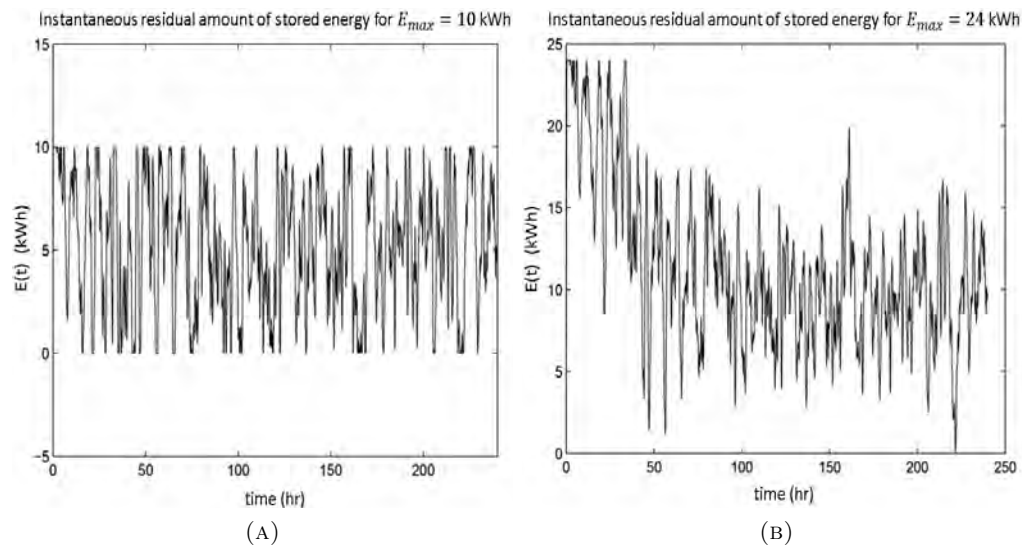


FIGURE 2.6: Amount of stored energy for A) $E_{max} = 10$ kWh and B) $E_{max} = 24$ kWh.

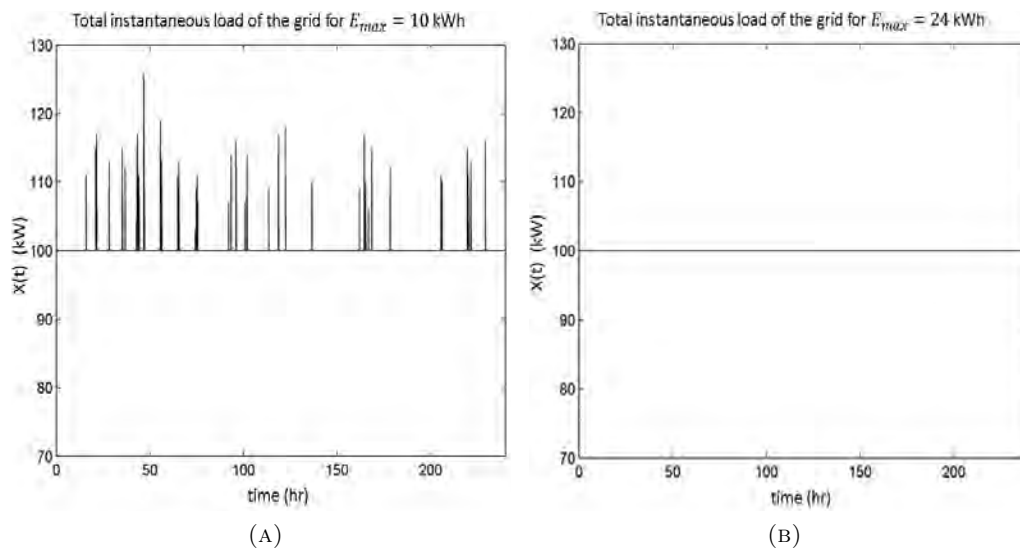


FIGURE 2.7: Total instantaneous grid load for A) $E_{max} = 10$ kWh and B) $E_{max} = 24$ kWh.

performance of the policy is optimal at $E_{max} \geq 24$ kWh, when the resulting average generation cost reaches the lower bound.

Next, we compare two scenarios, one with storage capacity $E_{max} = 10$ kWh and one with $E_{max} = 24$ kWh. In Fig. 2.6a and 2.6b, we show the instantaneous residual stored energy across the time horizon for the two scenarios. When $E_{max} = 10$ kWh, the battery empties ($E(t) = 0$) for a total of 69 times, whereas for $E_{max} = 24$ kWh, the battery is *never* fully discharged. In Fig. 2.7a and 2.7b we depict the total instantaneous load

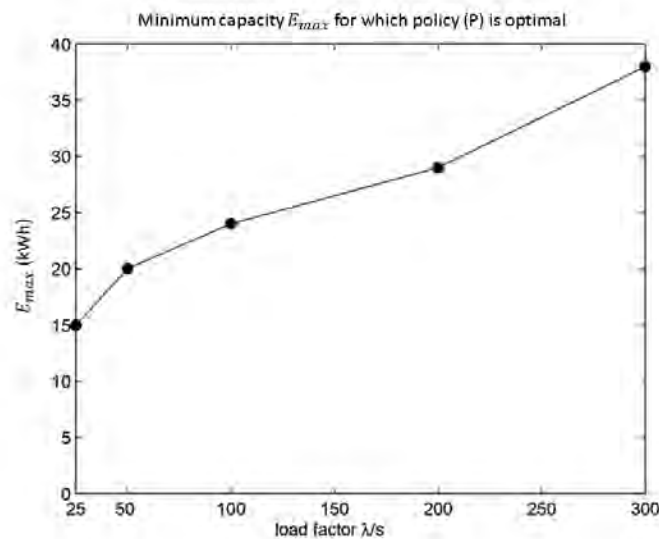


FIGURE 2.8: Minimum required capacity E_{max} for which the policy is optimal, versus the demand load.

on the grid, $X(t)$, for the two cases above. Observe that if $E_{max} = 10$ kWh, the total load exceeds $\bar{P} = 100$ for a total of 69 times as well. Clearly, these are the times when the battery empties, and the load that the battery was serving is moved to the grid. On the contrary, for $E_{max} = 24$ kWh, $X(t)$ remains constant and equals $\bar{P} = 100$, since the battery is never fully discharged. These findings verify the earlier theoretical investigation.

We also seek the value of the minimum required storage capacity for which our policy reaches the lower bound in (2.11), as a function of demand load λ/s . Fig. 2.8 presents our findings. The minimum required capacity can be seen to increase in a concave-like fashion in four out of the five test points as the demand load increases. This seems to imply that the additional amount of capacity needed to fulfill an additional unit of load decreases as the load increases, which is in accordance to the existence of a limiting value of storage capacity derived above.

2.2.7.2 Two storage devices and GEs

Now, we study the performance of the policy presented in section 2.2.5 for the case of minimization of the sum of individual GE costs, by evaluating the long-term average cost for the optimal policy and comparing it to the lower bound. We consider $N = 2$ GEs and $M = 2$ storage devices with no renewable resources attached, and a cost function $C(x) = x^2$ for both GEs. We also assume $\lambda_1 = 200$ requests/hour, $\lambda_2 = 100$ requests/hour, and $s = 2$ requests/hour. The average power requirement per demand

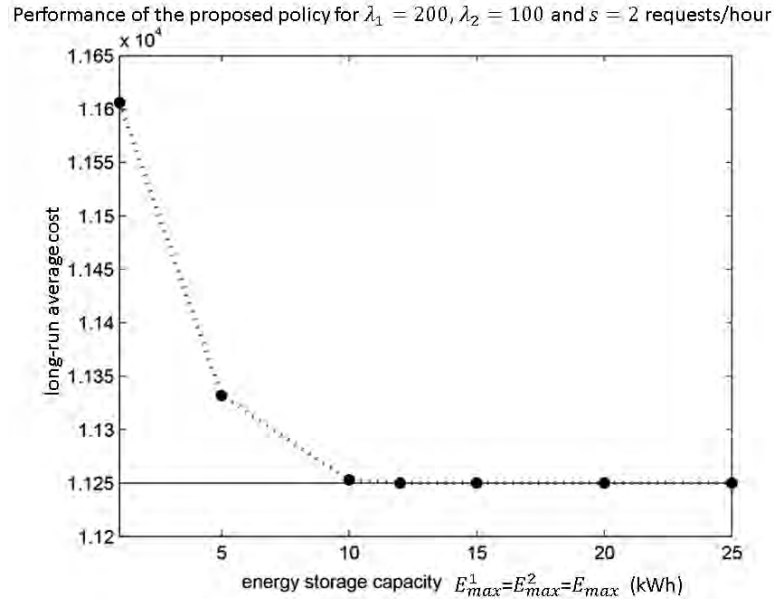


FIGURE 2.9: Performance of the proposed policy for the case of multiple batteries as a function of available storage capacities with $\lambda_1 \neq \lambda_2$, $E_{max}^1 = E_{max}^2$.

TABLE 2.1: Average cost of the proposed policy for $M = 2$ storage devices

E_{max}^1 (kWh)	E_{max}^2 (kWh)	average cost
1	4	11479
4	7	11323
7	10	11267
10	13	11251
13	16	11250

is 1 kW and the batteries have equal storage capacities, $E_{max}^1 = E_{max}^2 = E_{max}$. The threshold values are computed to be $P_0^1 = P_0^2 = 75$.

In Fig. 2.9, we show the values of the total long-term average cost for different values of the common storage capacity E_{max} . We observe that, as storage capacities increase, the average total cost converges to the lower bound which is 11250. Specifically, for $E_{max} \geq 12$ kWh the cost reaches the lower bound. Finally, we ran the same experiment for $E_{max}^1 \neq E_{max}^2$. Some of the resulting average costs are depicted in Table 2.1. Again, it can be observed that as storage capacities increase, the average total cost converges to the lower bound (11250).

2.3 Storage Placement and Power Flow

In this section, we study the problem of how much storage capacity should be placed on each node –prosumer entity or bus– of a power network given an available storage capacity budget. We are interested in the way storage capacity placement impacts the

overall cost of energy generation. In determining the optimal policy, it turns out that various aspects of power flow need to be taken into account. The solution policy for this energy storage dimensioning problem entails various parameters such as the demand profiles of prosumers and the power flow constraints. This section is based on work [11]. The references for the power flow analysis are [17], [18] and [19].

2.3.1 Power flow analysis

We consider a power network which consists of a set $\mathcal{N} = 1, 2, \dots, N$ of N prosumer nodes. Each node may be connected to other nodes through power distribution links. We also use the equivalent, smart-grid term “bus” to refer to a node. Let (k, ℓ) denote the power distribution link between nodes $k, \ell \in \mathcal{N}$. For each link (k, ℓ) , let $Y_{k\ell}$ be its admittance, which is a complex number. Each node can be a power supplier (generator) through some renewable energy source (RES) (e.g., a solar panel or a wind-turbine), or a power consumer, both or neither.

Let V_k be the complex voltage at bus $k \in \mathcal{N}$ and $|V_k|$ denote its magnitude. Let I_k denote the current injected at bus k . The node equation at bus k is,

$$I_k = \sum_{l=1}^N Y_{kl} V_l. \quad (2.43)$$

Define S_k to be the net complex power injection (generation minus load) at bus $k \in \mathcal{N}$. This is written as,

$$S_k = P_k + jQ_k = V_k I_k^*, \quad (2.44)$$

where I_k^* is the complex conjugate of the current at bus k , and P_k, Q_k are the active (real) and reactive powers. We can write,

$$I_k^* = \frac{P_k + jQ_k}{V_k} \Rightarrow I_k = \frac{P_k - jQ_k}{V_k^*} \quad (2.45)$$

and thus

$$P_k - jQ_k = V_k^* \sum_{l=1}^N Y_{kl} V_l = \sum_{l=1}^N Y_{kl} V_k^* V_l. \quad (2.46)$$

Now let $Y_{kl} = |Y_{kl}| e^{j\delta_{kl}}$ and $V_k = |V_k| e^{j\theta_k}$. Since

$$P_k - jQ_k = \sum_{l=1}^N |Y_{kl}| |V_k| |V_l| e^{j(\delta_{kl} + \theta_l - \theta_k)} \quad (2.47)$$

we have

$$P_k = \sum_{l=1}^N |Y_{kl}| |V_k| |V_l| \cos(\delta_{kl} + \theta_l - \theta_k) \quad (2.48)$$

and

$$Q_k = - \sum_{l=1}^N |Y_{kl}| |V_k| |V_l| \sin(\delta_{kl} + \theta_l - \theta_k) \quad (2.49)$$

Here, we use the linearized DC approximation. In this approximation, the network is assumed to be lossless, the voltage magnitudes $|V_k|$ are assumed to be at their nominal values at all nodes $k \in \mathcal{N}$, and the voltage phase angle differences $(\theta_k - \theta_l)$ between end-points of a link (k, l) are small.

The first step to this approximation is to neglect the reactive power, i.e., $Q_k \approx 0$. Through standard trigonometry, we obtain

$$P_k = \sum_{l=1}^N |Y_{kl}| |V_k| |V_l| (\cos \delta_{kl} \cos(\theta_l - \theta_k) - \sin \delta_{kl} \sin(\theta_l - \theta_k)). \quad (2.50)$$

The second step is to neglect the real part of the admittance, i.e., $\delta_{kl} \approx \pi/2$. Then

$$P_k = \sum_{l=1}^N |Y_{kl}| |V_k| |V_l| \sin \delta_{kl} \sin(\theta_k - \theta_l). \quad (2.51)$$

Finally, we assume that $|V_k| = |V_l| = 1$, and that the voltage angle differences $(\theta_k - \theta_l)$ between end-points of a link (k, l) are very small, so that $\sin(\theta_k - \theta_l) \approx \theta_k - \theta_l$. We also set $B_{kl} = |Y_{kl}| \sin \delta_{kl}$ for the imaginary part of the admittance, which is also called *reactance*, and we get the power flow equation

$$P_k = \sum_{l=1}^N B_{kl} (\theta_k - \theta_l). \quad (2.52)$$

This power-flow equation is called direct-current (DC) one because of its analogies with the flow equation in a DC circuit.

If we denote with $P_{kl} \geq 0$ the real power flowing from node k to node l through line (k, l) , and $k, l \in \mathcal{N}$, it is

$$P_{kl} = B_{kl} (\theta_k - \theta_l). \quad (2.53)$$

2.3.2 Storage and power flow model

For the power flow model, we use the linearized DC approximation described above. For our model, we assume a time-slotted system with slots of unit length, and study

its operation for an horizon of T time slots. The value of power flow $P_{kl}(t)$, for nodes $k, l \in \mathcal{N}$, at different times t is limited by stability constraints and thermal effects, and we have:

$$|P_{kl}(t)| \leq f_{kl}, \quad (2.54)$$

where f_{kl} is the capacity of the line connecting nodes k and l .

Let $L_k(t)$ be the instantaneous demand load at node $k \in \mathcal{N}$ at time slot t and let $G_k(t)$ be the amount of generated power at node k at slot t . We denote the instantaneous operational cost associated with power generation $G_k(t)$ as $C_k(G_k(t))$, where $C_k(\cdot)$ is an increasing, differentiable convex function.

Typically, the energy demand of prosumers follows specific patterns such as periodic ones [20], [21]. Thus, the energy demand patterns of buses can be characterized. In order to extract these patterns, estimation techniques from historical data regarding the power demand of buses may be used. Here, we will assume that the demand profiles for a finite horizon T , which we denote here as vectors $\mathbf{L}_k = (L_k(t) : t = 1, \dots, T)$, for $k \in \mathcal{N}$, are obtained and known in advance.

The decision variable $h_k(t)$ is the charging/discharging rate of the battery of bus k at slot t , respectively. Also, let $E_k(t)$ be the stored amount of energy at the battery of node k at the beginning of slot t . Define $E_{max}^k \geq 0$ as the storage capacity allocated at node k . Clearly,

$$\sum_{k \in \mathcal{N}} E_{max}^k \leq S_B \quad (2.55)$$

where S_B is the total available storage budget, and $0 \leq E_k(t) \leq E_{max}^k$ for $t = 1, \dots, T$. The level of stored energy $E_k(t)$ at the battery of node k evolves with time as:

$$E_k(t+1) = E_k(t) + h_k(t). \quad (2.56)$$

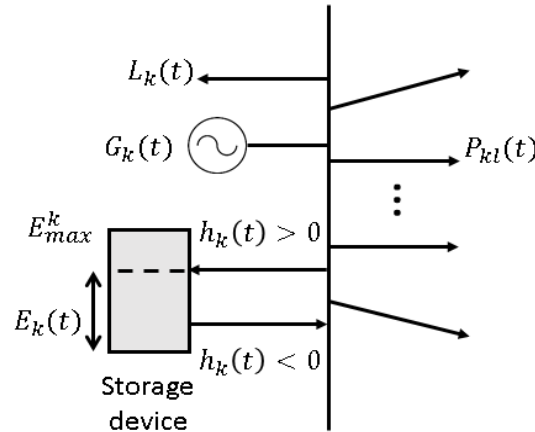
Since $0 \leq E_k(t) \leq E_{max}^k$ at all slots t and time slots are of unit size, the quantity $h_k(t)$ must satisfy,

$$-E_k(t) \leq h_k(t) \leq E_{max}^k - E_k(t) \quad (2.57)$$

Here, we assume that each battery is empty at installation time $t = 0$, i.e., $E_k(0) = 0$. We have the following succinct time-evolution formula:

$$E_k(t+1) = \min\{E_{max}^k, \max\{E_k(t) + h_k(t), 0\}\}. \quad (2.58)$$

Clearly, if $G_k(t) > L_k(t)$ there exists a power surplus and the excess energy $G_k(t) - L_k(t)$ of prosumer k can be used to charge her installed battery with rate $h_k(t) > 0$, and/or


 FIGURE 2.10: Power balance at node k .

satisfy any energy needs of another prosumer $l \neq k$ through transferring power $P_{kl}(t)$ to her. On the other hand, if $G_k(t) < L_k(t)$, then there exists a power deficit and thus prosumer k may satisfy her residual demand load $L_k(t) - G_k(t)$ by discharging her battery at rate $|h_k(t)|$, and/or by receiving power $|P_{kl}(t)|$ from a prosumer l .

At each time t , the power that flows in and out of a bus k should be balanced, i.e., it must hold:

$$G_k(t) = L_k(t) + h_k(t) + \sum_{l \neq k} P_{kl}(t) \quad (2.59)$$

This power balance constraint can be observed in Fig. 2.10.

2.3.3 Problem formulation

In the system design phase, there exists an amount S_B of available storage budget/capacity which should be allocated among buses. The charging/discharging and energy exchange decisions of the nodes would need to be considered as well.

Given the energy demand pattern \mathbf{L}_k for each prosumer k , we need to decide at each time slot: (i) whether and how much the battery will be charged or discharged, and (ii) whether power $|P_{kl}(t)|$ will flow from or to some other prosumer l . Further, partitioning the finite storage capacity across buses and placing corresponding storage facilities there to optimize the generation cost is needed. The jointly optimal policy will encompass the variables above and can be derived from the solution of the following optimization problem:




FIGURE 2.11: A power network consisting of two buses. It must hold $E_{max}^1 + E_{max}^2 \leq S_B$

$$\min \sum_{t=1}^T \sum_{k \in \mathcal{N}} C_k(G_k(t)) \quad (2.60)$$

over $(h_k(t), P_{kl}(t), E_{max}^k), t = 1, \dots, T, k, l \in \mathcal{N}$

subject to: (2.54), (2.55), (2.58), (2.57), (2.59),

where constraint (2.54) should be satisfied by a feasible power flow, constraint (2.55) stands for the limited storage capacity over all storage devices, constraint (2.58) represents the constraint on the stored amount of energy at each battery, constraint (2.57) is about the charging/discharging rates of the energy storage devices, and constraint (2.59) represents the power balance constraint at each node of the network.

In the sequel, we present and discuss some simple special cases of the problem.

2.3.4 Numerical Example 1: Storage placement for N=2 buses

Consider a toy power network with two nodes, which is depicted in Fig. 2.11. Both buses $k = 1, 2$ have generator units and loads, and let \bar{G}_k, \bar{L}_k denote the average power generation and consumption load respectively at bus k . To simplify the scenario, we assume that storage capacity and power link capacity are unlimited, and we focus on the impact of placing the entire storage at bus 1 or 2. If storage is placed at bus 1, the power conservation equations at the two buses are

$$\bar{G}_1 - P_{12} - h_1 = \bar{L}_1 \quad (2.61)$$

and

$$\bar{G}_2 + P_{12} = \bar{L}_2. \quad (2.62)$$

Note that the average charging rate cannot exceed the total injected power in the bus, and it cannot also exceed the load, i.e.

$$h_1 \leq \bar{G}_1 - P_{12}, \text{ and } h_1 \leq \bar{L}_1. \quad (2.63)$$

The objective is to minimize the total generation cost, $C(\bar{G}_1) + C(\bar{G}_2)$, which is written as,

$$\min_{P_{12}, h_1} C(\bar{L}_1 + h_1 + P_{12}) + C(\bar{L}_2 - P_{12}). \quad (2.64)$$

We can take the partial derivatives with respect to P_{12} and h_1 . We get

$$P_{12} = \frac{1}{2}(\bar{L}_2 - \bar{L}_1 - h_1), \quad C'(\bar{L}_1 + h_1 + P_{12}) = 0. \quad (2.65)$$

If for example $C(x) = x^2$, we have $h_1 = -(\bar{L}_1 + P_{12})$. Finally,

$$h_1 = \min\{-(\bar{L}_1 + P_{12}), \bar{L}_1, \bar{G}_1 - P_{12}\}. \quad (2.66)$$

Depending on the arithmetic values of the load and generation parameters involved, we obtain the expressions. We can make similar computations for storage placed in bus 2, and we then decide where to place storage. For a numerical illustration, let the cost of generation at each bus k be $C(\bar{G}_k) = \bar{G}_k^2$, and the average demand loads at nodes 1 and 2 be $\bar{L}_1 = 10$ and $\bar{L}_2 = 20$. Also, suppose that the line capacity f_{12} is infinite. If there is an infinite storage budget, then:

- if the storage capacity is placed in bus 1, then $h_1 = -30$, $P_{12} = 20$, $\bar{G}_1 = 0$, $\bar{G}_2 = 0$ and thus, $C(\bar{G}_1) + C(\bar{G}_2) = 0$
- if the storage capacity is placed in bus 2, then $h_2 = -30$, $P_{12} = -10$, $\bar{G}_1 = 0$, $\bar{G}_2 = 0$ and thus, $C(\bar{G}_1) + C(\bar{G}_2) = 0$.

Thus, the average cost remains the same regardless of the bus at which storage is placed, and this is true for other values of the loads as well. Even for a different set of values of the average demand loads, the result is the same. Namely, for $\bar{L}_1 = 15$ and $\bar{L}_2 = 5$:

- if the storage capacity is placed in bus 1, then $h_1 = -20$, $P_{12} = 5$, $\bar{G}_1 = 0$, $\bar{G}_2 = 0$ and $C(\bar{G}_1) + C(\bar{G}_2) = 0$
- if the storage capacity is placed in bus 2, then $h_2 = -20$, $P_{12} = -15$, $\bar{G}_1 = 0$, $\bar{G}_2 = 0$ and $C(\bar{G}_1) + C(\bar{G}_2) = 0$.

This cost invariance holds even if the storage capacity available is finite, e.g. $S_B = 5$, and for example $\bar{L}_1 = 10$ and $\bar{L}_2 = 20$:

- if the storage capacity is placed in bus 1, then $h_1 = -5$, $P_{12} = 7.5$, $\bar{G}_1 = 12.5$, $\bar{G}_2 = 12.5$ and thus, $C(\bar{G}_1) + C(\bar{G}_2) = 312.5$
- if the storage capacity is placed in bus 2, then $h_2 = -5$, $P_{12} = 2.5$, $\bar{G}_1 = 12.5$, $\bar{G}_2 = 12.5$ and thus, $C(\bar{G}_1) + C(\bar{G}_2) = 312.5$,

and for $\bar{L}_1 = 15$ and $\bar{L}_2 = 5$:

- if the storage capacity is placed in bus 1, then $h_1 = -5$, $P_{12} = -2.5$, $\bar{G}_1 = 7.5$, $\bar{G}_2 = 7.5$ and thus, $C(\bar{G}_1) + C(\bar{G}_2) = 112.5$
- if the storage capacity is placed in bus 2, then $h_2 = -5$, $P_{12} = -7.5$, $\bar{G}_1 = 7.5$, $\bar{G}_2 = 7.5$ and thus, $C(\bar{G}_1) + C(\bar{G}_2) = 112.5$.

Again, we observe that the total cost is independent of the placement location of the storage capacity. It can be checked that the same holds for every possible value of S_B .

However, for $\bar{L}_1 = 10$, $\bar{L}_2 = 20$, $S_B = 5$ and finite line capacity $f_{12} = 1$, we calculate that:

- if the storage capacity is placed in bus 1, then $h_1 = -5$, $P_{12} = 1$, $\bar{G}_1 = 6$, $\bar{G}_2 = 19$ and thus, $C(\bar{G}_1) + C(\bar{G}_2) = 397$
- if the storage capacity is placed in bus 2, then $h_2 = -5$, $P_{12} = 1$, $\bar{G}_1 = 11$, $\bar{G}_2 = 14$ and thus, $C(\bar{G}_1) + C(\bar{G}_2) = 317$,

which means that the storage capacity should be placed in bus 2.

On the other hand, for $\bar{L}_1 = 15$, $\bar{L}_2 = 5$, $S_B = 5$ and $f_{12} = 1$:

- if the storage capacity is placed in bus 1, then $h_1 = -5$, $P_{12} = -1$, $\bar{G}_1 = 9$, $\bar{G}_2 = 6$ and thus, $C(\bar{G}_1) + C(\bar{G}_2) = 117$
- if the storage capacity is placed in bus 2, then $h_2 = -5$, $P_{12} = -1$, $\bar{G}_1 = 14$, $\bar{G}_2 = 1$ and thus, $C(\bar{G}_1) + C(\bar{G}_2) = 197$,

which means that the storage capacity should be placed in bus 1.

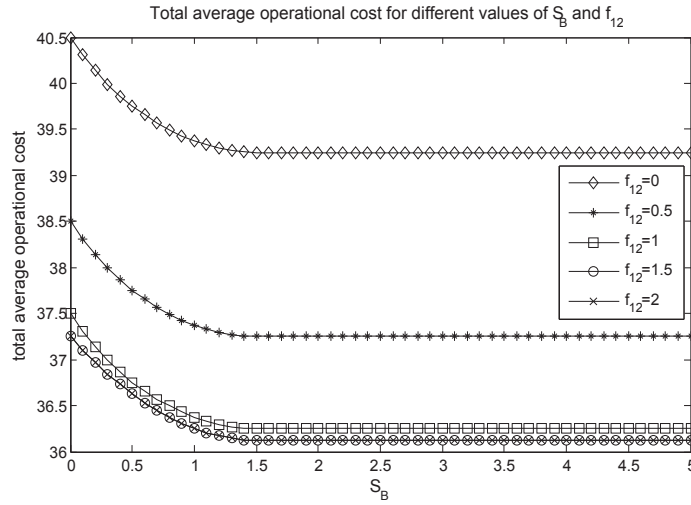


FIGURE 2.12: The optimal system performance for different values of the available storage S_B and line capacity f_{12} .

2.3.5 Numerical Example 2: Storage dimensioning for N=2 buses

Consider the same toy power network of Fig. 2.11. We assume that the demand load profiles $L_k(t)$ of nodes $k = 1, 2$ and the power flow capacity f_{12} of the line connecting the two nodes are known. Further, we assume that there is an available storage budget S_B which can be split into amounts E_{max}^1, E_{max}^2 and allocated at the two nodes, so that $E_{max}^1 + E_{max}^2 = S_B$. We study our system for a time horizon of $T = 2$ slots.

The objective is to minimize the total average generation cost through optimal storage dimensioning, storage management, and power flow. Namely, we need to optimally share the available storage budget S_B among the two nodes, and find the charging/discharging rates $h_k(t)$ and power flow $P_{12}(t)$ at each slot t , so as to minimize the total average generation cost, $\frac{1}{2} \sum_{t=1}^2 \sum_{k=1}^2 C_k(G_k(t))$, which can be written as

$$\frac{1}{2} \sum_{t=1}^2 C_1(L_1(t) + h_1(t) + P_{12}(t)) + C_2(L_2(t) + h_2(t) - P_{12}(t)). \quad (2.67)$$

We conduct a numerical study for different values of demand loads $L_1(t), L_2(t)$, power flow capacity f_{12} and storage budget S_B . Here, we assume $C_1(\cdot) = C_2(\cdot) = C(\cdot)$ and $C(G_k(t)) = (G_k(t))^2$. Our results are depicted in Fig. 2.12 and Table 2.2.

In Fig. 2.12, the optimal total average generation cost (2.67) for $S_B \in [0, 5]$ and line capacity $f_{12} = 0, 0.5, 1, 1.5, 2$ is depicted. For all values of f_{12} , we observe that as the available storage capacity S_B increases, the optimal generation cost decreases, and for storage larger than a certain value, i.e., for $S_B \geq 1.4$, it becomes constant. Also, as we

	Values for $f_{12} = 0.5$				Values for $f_{12} = 1$			
$L_1(1)$	5	1	7	3	5	1	7	3
$L_1(2)$	7	7	1	7	7	7	1	7
$L_2(1)$	1	5	5	5	1	5	5	5
$L_2(2)$	3	3	3	1	3	3	3	1
$h_1(1)$	1	2.5	0	1.67	1	2	0	1.33
$h_1(2)$	-1	-2.5	0	-1.67	-1	-2	0	-1.33
$h_2(1)$	1	0	0	0	1	0.001	0	0
$h_2(2)$	-1	0	0	0	-1	-0.001	0	0
$P_{12}(1)$	-0.5	0.5	-0.5	0.17	-1	1	-1	0.33
$P_{12}(2)$	-0.5	-0.5	0.5	-0.5	-1	-1	1	-1
E_{max}^1	2.67	3.54	2.77	3.22	2.66	3.33	2.78	3.11
E_{max}^2	2.33	1.46	2.23	1.78	2.34	1.67	2.22	1.89

TABLE 2.2: Optimal storage dimensioning, charging/discharging and power flow for $S_B = 5$ and $f_{12} = 0.5, 1$.

allow larger line capacity, the optimal cost decreases, but it does not reduce beyond line capacity $f_{12} = 1.5$.

Table 2.2 shows the optimal battery capacities, charging/discharging rates and power flows of the two buses for different values of the demand profiles $L_1(t)$, $L_2(t)$ and for line capacity $f_{12} = 0.5, 1$. It can be seen that for both values of f_{12} , more storage capacity is placed in buses which need greater charging/discharging rates. When the two buses need to charge or discharge with the same rate, the available capacity is almost equally shared among the two buses.

2.4 Related Work

There exists significant amount of work on leveraging stored energy in various contexts.

From a wireless networking perspective, the work [22] uses the stochastic dynamic programming framework to obtain an optimal battery discharge policy for maximizing the lifetime of power-limited wireless nodes. The problem of energy-aware routing is studied in [23]. The designed algorithm computes an energy-weighted shortest path from the source to the destination node of each service request. The wireless nodes are assumed to have knowledge of the future short-term recharge process. However, this is an assumption that the authors of [24] dispense with in their work. They present a modified adaptive backpressure policy that maximizes throughput and whose asymptotic optimality is shown using the notion of Lyapunov drift. A similar problem is considered in [25] using a similar Lyapunov optimization approach. The proposed algorithms jointly manage the stored energy and make power allocation decisions for packet transmissions

without requiring any knowledge of the harvestable energy process and the channel qualities.

One classic problem in the literature is the optimal stored energy control problem from the consumer point of view, in the presence of time-varying prices. In [26], [27], for instance, the proposed cost-minimizing storage control policy is shown to be a simple threshold-based one, while [28] proposes SmartCharge, a charging/discharging system that determines when and how much to store low-cost energy for use during high-cost periods. SmartCharge's algorithm leverages next-day electricity prices and a prediction model which forecasts future demand. Based on the Lyapunov optimization technique, the authors of [29] develop energy management schemes for load-serving and demand-response. Each power consuming entity's objective is to find a control policy for determining the load consuming, purchasing/selling, and charging/discharging actions, so as to minimize her average cost. In [30], the authors proposed a threshold-based control policy for a single battery and extended their policy to the multiple-battery case in order to reduce the cost paid by users for energy supply.

In the context of game theory, two games are discussed in [31]. The first is a non-cooperative one played between storage unit owners, who schedule their energy use to minimize their energy cost. The second is a Stackelberg game played between the utility provider and the energy consumers, in which the users minimize their cost, while the utility maximizes its profit. Work [32], instead, introduces a prospect theory-based framework which explains how real-life user decisions can deviate from those predicted by conventional game theory. A non-cooperative game is formulated between storage unit owners, which aim to maximize their utility functions that capture benefit-cost tradeoffs.

Energy storage has also received significant attention as a way to efficiently integrate renewables into the electricity grid. In particular, in [33] storage has been proposed as a strategy for increasing the average energy sold and reducing curtailment of wind energy in a wind-based system that trades electricity in electricity markets. The corresponding problem is formulated as a Markov decision process and a triple-threshold policy is proposed. The stochastic network calculus framework is extended in [34] to investigate the effects of energy storage on the power supply reliability in configurations with different levels of renewable generation.

In the same context, the impact of wind prediction quality on the performance of a system equipped with a storage system and a set of wind sources is studied in [35]. The authors present two production scheduling policies: a deterministic one, and a heuristic one which is obtained using the statistics of wind-forecast errors. Two techniques to mitigate the intermittent nature of renewables are considered in [36]: the use of storage

and the concept of distributed generation combined with cooperation among distributed sources. The objective is to minimize the time average cost of energy exchange among the grid. An algorithm is developed using the technique of Lyapunov optimization.

The works [37] and [38] present ways to reduce power costs in data centers. Using the technique of Lyapunov optimization, an online control algorithm is developed in [37] that exploits UPS storage devices to minimize the time average electricity bill in a data center. The proposed algorithm operates without any knowledge of the statistics of the workload or electricity cost processes. The authors of [38] apply the same optimization technique to solve the problem of optimal traffic distribution and battery charging/discharging management in Internet data centers under location-varying and time-varying electricity prices.

The problem of optimal dimensioning and placing of energy storage systems within the distribution grid is studied in [11], [39]. In [11], the placement of storage is shown to affect the power generation costs, and the objective is to minimize the power generation cost, yet assuming that there is no cost for storage devices. The goal is to optimally place and partition a given amount of storage capacity and to control the storage unit across a power network so as to minimize the long-term average generation cost. The work [39] proposes a model for collaborative prosumption of energy in a community of prosumers that collectively use energy storage systems. The objective of the prosumers is to minimize the total cost of battery deployment and that of power drawn from the grid. The decisions amount to battery charging and discharging at each time slot as well as partitioning of a storage budget among a set of candidate locations. Nash bargaining theory is used to determine how the total cost should be shared in a fair fashion among prosumers.

The benefits of cooperation are studied in [40] for the case of micro-grids that transfer power among themselves. The proposed algorithm, based on coalitional game theory, allows the microgrids to form coalitions so as to minimize the costs due to losses of power over the distribution lines. However, the potential of energy storage is not addressed. An other work that deals with micro-grid challenges is [41]. Leveraging the dual decomposition method, a distributed power scheduling approach is developed, tailored for a micro-grid with RESs. The objective is to minimize the microgrid net cost, but neither cooperation among micro-grids nor sharing of storage resources is considered.

The majority of those works study the optimal energy storage management problem from the perspective of the consumer. In the first part of this chapter, we study the optimal energy storage control problem from the perspective of the utility operator, with the goal of minimizing the long-term average grid operational cost. It turns out that the convexity of cost gives rise to simple intuitive control policies.

2.5 Conclusion

In this chapter, we studied the optimal stored energy control problem for the case of a single storage device, and for multiple storage devices that are shared among multiple micro-grid entities. Our objective of minimizing the average grid operational cost (or, the sum of GE operational costs), led us to policies which attempt to keep balanced grid power consumption at all times. Our numerical results demonstrate that the proposed asymptotically optimal policies are excellent approximations even for *finite* storage capacity values, as small as a few tens of kWh. We also studied the joint energy storage placement, dimensioning and control problem given an available storage budget, where the goal is to minimize the power generation cost. We provided two simple numerical examples in order to present the optimal solutions for some special cases of this problem.

Chapter 3

Green Web Crawling

Contents

3.1	Introduction to the Energy Efficiency of Web Search Engines	44
3.2	System Model	46
3.2.1	Web Crawler	46
3.2.2	Web Servers	47
3.2.3	Greenness of Server Energy Consumption	48
3.2.4	Web Page Staleness	48
3.3	Problem Formulation	49
3.3.1	Single Web Server, Single Thread Scheduling Problem	49
3.3.2	Optimal Web Page Download Scheduling Policy	51
3.3.3	Extensions to the Model	52
3.4	Performance Evaluation	55
3.4.1	Data Set	55
3.4.2	Greenness and Staleness Computation	55
3.4.3	Heuristic Policies	57
3.4.4	Experimental Results	58
3.5	Related Work	63
3.5.1	Refreshing Web Repositories	63
3.5.2	Job and Packet Processing	63
3.5.3	Energy Efficiency and Greenness of Data Centers	64
3.6	Conclusion	65

3.1 Introduction to the Energy Efficiency of Web Search Engines

The operations of a web search engine can be grouped under three main components: web crawling, indexing, and query processing [42]. Web crawling is responsible for traversing the hyperlink structure among the web pages to discover and download the content in the Web, as well as for refreshing already downloaded pages in the web repository. The indexing component converts downloaded content into compact data structures that can be easily searched. Finally, the query processing component evaluates user queries by processing these data structures and matches each query to a set of pages deemed to be relevant to the query.

Motivated by a real-world problem, in this chapter, we consider energy efficiency in the context of a concrete application that comes from web crawling. Namely, we focus on the greenness of the web crawling process. In search data centers, thousands of computers are allocated to crawl the Web. Maintaining an infrastructure of this scale results in certain implications in terms of energy consumption and carbon footprint. In general, web crawlers lead to carbon emissions in two different ways, due to (i) local operations performed on the crawling nodes in the data centers (e.g., parsing web pages) and (ii) remote operations performed on the web servers while serving HTTP requests (e.g., retrieving a page from disk). In this chapter, we are interested in the latter case, i.e., the carbon emissions that the web crawler incurs on remote computers that do not belong to the search engine.

In practice, a web crawler may lead to significant energy consumption on web servers while the HTTP requests issued by the crawler are processed on the servers (e.g., during disk accesses, processing in the CPU, and network operations). We motivate this by some back-of-the-envelope calculations: Let us assume that there are five billion pages in the Web [43]. According to a conservative estimate, we can assume an average of 200 J (0.055 Wh) of energy consumption per HTTP request [44]. Let us assume that each page in the web repository is refreshed once per minute, on average. Now, refreshing only one-tenth of the repository requires about 40 GWh of energy per day. Web crawling is also a costly operation in terms of the carbon emissions of web servers. In our example, if the carbon footprint of the fuel used to generate electricity is 0.85 kg/KWh [45], on average, the carbon emissions due to web crawling can be estimated as 34,000 tons per day.

Unfortunately, there is little a web crawler can do to reduce the energy consumption it incurs to web servers without sacrificing the coverage or freshness of its web repository. This is because the amount of energy consumed on web servers depends only on factors

related to the hardware and software resources that are not managed by the search engine company. Nevertheless, certain optimizations can be employed to reduce the carbon emissions that a crawler incurs to web servers. Our main observation is that the carbon footprint of a web server depends on the type (greenness) of the consumed energy, which varies depending on the time of the day if the server is supported by a RES (e.g., a solar panel). For example, a web server is more likely to consume green energy in daytime, while it is more likely to consume brown energy during the night. This intra-day variation creates an opportunity to reduce the carbon footprint of web servers as HTTP requests may be scheduled such that pages are downloaded from web servers that are more likely to be consuming green energy.

Motivated by the observation above, we devise a web repository refreshing technique that takes into account both the greenness and staleness concepts when scheduling the download of web pages. This technique aims to reduce the total staleness of pages in the web repository while constraining web servers' total carbon footprint resulting from the activities of the crawler. The goal of our approach is to reduce/constraint the carbon footprint of web servers by exploiting the dynamics of renewable energy generation, which are here modeled by means of an appropriately defined time-varying (greenness) index. Beyond creating an environment-friendly crawler, our work has implications for large-scale web search engines, which should comply with regulations about carbon footprint reduction.

The concept of staleness is unique to our green web crawling problem and it arises only in the context of keeping the downloaded pages fresh enough. Our objective of minimizing the total staleness of a web repository is related to the problem of minimizing the total waiting time of scheduled jobs [46, 47]. However, our problem differs in that the crawler needs to consider both the staleness and greenness aspects while making its scheduling decisions. The concept of freshness could be related to a hard deadline constraint on the refreshing of pages. However, the nature of the crawling process necessitates a softer version of the deadline, that of staleness.

Due to the constraints on the amount of carbon emissions, the opportunities for scheduling each page weigh differently, and the concept of greenness needs to be considered explicitly in scheduling. To accommodate the average greenness constraints in our problem, we weigh them with the Lagrange multiplier λ and include them in the objective function. In that respect, our work could be considered, in abstract terms, as a scheduling problem under both deadline and average energy constraints. This same problem could also be studied from the point of view of a smart grid system; continually generated power demand requests have to be scheduled under both timeliness and average

greenness constraints taking into account the dynamic patterns of renewables. In contrast to works that present policies for time- or energy-efficient offline job scheduling under constraints [48], our work proposes online policies, where the crawler makes its decisions as the page download requests arrive. Finally, the concepts of page size and page freshness requirement differentiate our work further from previous works.

Our main contributions are the following:

- We introduce the green web crawling problem, where we study a page refresh policy that minimizes the total staleness of pages in the web repository of a crawler, while keeping the amount of carbon emission on remote web servers low enough.
- For one web server and one crawling thread, we show that the optimal page refresh policy is a greedy one. At each time slot, the page to be refreshed is selected based on a metric that considers the staleness of the page, its size, and the greenness of the energy consumed by the web server.
- We extend the optimal policy to the cases of (i) many web servers, (ii) multiple crawling threads, and (iii) web pages with variable freshness requirements. We also propose various heuristics along the lines of the optimal policy.
- We conduct simulations with a large-scale, real data set obtained from Yahoo.

The rest of this chapter is organized as follows. In section 3.2, we describe our system model. In section 3.3, we formulate our core optimization problem, devise an optimal policy, and provide some extensions. In section 3.4, we describe our experimental setup, present some heuristics, and set forth numerical results obtained via simulations. Finally, section 3.5 provides an overview of related work and section 3.6 concludes our study. This chapter is based on works [49] and [50].

3.2 System Model

3.2.1 Web Crawler

In large-scale search engines, web crawling is performed by clusters of computers, where each computer runs multiple crawling threads. A crawling thread either downloads a newly discovered URL or re-downloads a previously stored page in the repository by issuing HTTP requests to web servers. These two operations are known as discovery and refresh, respectively. The discovery operations increase the size of the web repository of the crawler. The refresh operations help to decrease the staleness of pages in the

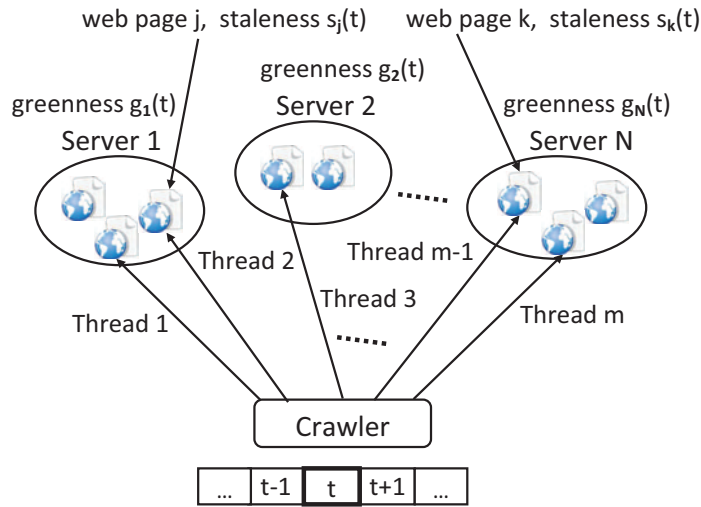


FIGURE 3.1: Our system model for the crawler: m crawling threads concurrently retrieve pages from N web servers at time slot t .

repository. The focus of this work is on the page refresh operations only. Decreasing the staleness of a page repository is important since this has a direct impact on the quality of the search results presented to the users, thus affecting the monetization of the search engine.

3.2.2 Web Servers

The system model is depicted in Fig. 3.1. We assume that there exists a set \mathcal{K} of N web servers. Also, let \mathcal{W}_i be the set of web pages that are hosted in server $i \in \mathcal{K}$ and thus, $\mathcal{W} = \bigcup_{i \in \mathcal{K}} \mathcal{W}_i$ be the set of all pages in the system. For each page $j \in \mathcal{W}_i$, let p_j be the size of its content, in bytes. The hardware devices located in the remote servers consume energy for serving the page download requests issued by the crawler. Fetching the requested pages from the disk, processing them in the CPU, and transmitting them through the server transmit circuit, all consume energy. To account for the factors above, we assume that the amount of energy consumed to download a page j is a linear function of p_j , i.e., $e_j = \alpha p_j + \beta$, with $\alpha, \beta > 0$, known constants. These assumptions are made to better expose the main characteristics of our approach, namely server greenness and web page staleness, to be presented below.

We also assume that each page may have its own freshness requirement which is mainly determined by two factors: (i) its PageRank and (ii) its likelihood of change. The PageRank of a page measures its relative importance within the set of web pages [51]. On the other hand, the likelihood of change of a page is determined by the estimated frequency of content change of the page. Pages with high PageRank value and high

likelihood of change may have higher freshness requirements. We define weights γ_j , $\forall j \in \mathcal{W}_i$ with $i \in \mathcal{K}$, which indicate the different freshness requirements of pages. Web pages with high γ_j should be refreshed more often. These weights can be estimated based on crawler statistics and past crawling history.

3.2.3 Greenness of Server Energy Consumption

For each server i , we define a time-varying value, $g_i(t)$, which indicates (on a given scale) the amount of produced carbon emissions per unit of consumed energy (Wh). This value denotes the “greenness” of the energy consumed to run the server. For example, in the $[0, 1]$ scale, the carbon emissions (or greenness) of a very green energy source is close to zero (no carbon emissions, maximum greenness) and, as the value increases, the source becomes more “brown”. For instance, a server may be powered entirely by clean nuclear energy or wind power, or entirely by brown energy generated from coal or lignite, or by a mixture of these. The $g_i(t)$ values may vary in time due to the time-varying power output generated by renewable energy sources or because the local electricity company increases the brownness due to high demand. We assume that there is no a priori knowledge of the $g_i(t)$ values, and they can be communicated from the web servers to the crawler only at the time that a decision needs to be made.

3.2.4 Web Page Staleness

We assume that time is divided into slots and that the slot size is large enough to cover the download of the largest page. Hence, page downloads occur on a time slot basis. This assumption is made to simplify the subsequent analysis and better expose the structure of the optimal policy. Nevertheless, it is not restrictive since it still captures different energy consumption incurred to remote web servers.

At the beginning of each slot, m threads are directed by the crawler toward m web pages that are selected for download. For each page j , we need to define a measure of its staleness. Let $s_j(t)$ be the staleness of a page j at the beginning of slot t . If this page is selected for download, then at most by the end of the slot the page download will finish, and at the end of the slot its staleness will be 0. However, if this page is not selected, then its staleness will increase and at the end of the slot it will be $s_j(t) + 1$. We observe that as long as page j is not selected, it becomes more and more stale, which means that the staleness of the page depends on the time elapsed since the last time it was downloaded. It should be noted that the slot assumption above leads to a somewhat conservative consideration in terms of the computed staleness, in the sense that it leads

to larger staleness increase for pages that are not scheduled for download. However, even if this assumption is relaxed, the structure of our analysis is not expected to change.

3.3 Problem Formulation

We are interested in scheduling the page download threads from the crawler to the web servers with the objective to keep the web pages as much as possible fresh and the carbon emissions due to page download requests low enough. The decision at each time t is to pick a server $i \in \mathcal{K}$ and a page $j \in \mathcal{W}_i$ to download in such a way that the total staleness of the pages is minimized and the amount of carbon emissions is kept below a given threshold. On the one hand, we would like to choose pages with large staleness. On the other hand, we would like to schedule downloads of pages from servers that have low $g_i(t)$ values. Also, out of all pages it is not clear whether we should schedule for download the ones with smaller size or the ones with larger size. Small-size pages consume less energy. However, larger pages should also be downloaded at some point in order to reduce their staleness. There may also be pages with high freshness requirements γ_j , which should be given high priority in the download process.

The joint consideration of all parameters above and the conflicting objectives of keeping the pages fresh and the carbon emissions at the remote servers low, make the thread scheduling problem non-trivial. First, we formulate and solve the basic single-server, single-thread problem presenting the various intuitions behind it. We then extend it to the cases of many servers, multiple crawling threads, and pages with different freshness requirements.

3.3.1 Single Web Server, Single Thread Scheduling Problem

First, we consider the simple case of one server that hosts a set \mathcal{W} of web pages. A horizon of $(T + 1)$ time slots is assumed. As we mentioned above, the page download time is one slot, and we assume that $g(t)$ remains stable over the time slot. At each time t , $m = 1$ thread is sent to fetch one web page. We define the vector variable $\mathbf{x}(t) = (x_j(t) : j \in \mathcal{W})$, where $x_j(t) = 1$ if at time slot t the web page $j \in \mathcal{W}$ is downloaded, else $x_j(t) = 0$. Clearly, $\sum_{j \in \mathcal{W}} x_j(t) = 1, \forall t = 0, \dots, T - 1$, since at each slot t , only one thread to a web page is allowed to be active to the server. We note that the staleness of a web page j at the beginning of slot $t + 1$ is:

$$s_j(t + 1) = \begin{cases} 0 & \text{if } x_j(t) = 1 \\ s_j(t) + 1 & \text{if } x_j(t) = 0. \end{cases} \quad (3.1)$$

The time evolution of the staleness of a page j can be written as $s_j(t+1) = (s_j(t) + 1)(1 - x_j(t))$. Our goal is to find a policy $\mathbf{x}^* = (\mathbf{x}(t) : t = 0, \dots, T-1)$ that minimizes the total (over time and over pages) staleness of the server, i.e.,

$$\min_{\mathbf{x}^*} \sum_{t=0}^T \sum_{j \in \mathcal{W}} s_j(t) \quad (3.2)$$

subject to the constraints that only one thread to a web page is allowed to be active to the server at each time t , i.e.,

$$s.t. \quad \sum_{j \in \mathcal{W}} x_j(t) = 1, \quad \forall t = 0, \dots, T-1, \quad (3.3)$$

and that the total amount of carbon emissions due to page download requests does not exceed a given threshold \bar{G} , i.e.,

$$s.t. \quad \sum_{t=0}^{T-1} \sum_{j \in \mathcal{W}} x_j(t) e_j g(t) \leq \bar{G}, \quad (3.4)$$

where \bar{G} is set by the agreement between the crawler and the remote server. Since we estimate the total staleness (over pages) at the beginning of each slot t , we assume that the last scheduling decision is made at the beginning of slot $T-1$, the last page download finishes by the end of slot $T-1$, the last total staleness estimation is made at the beginning of slot T , and no action is taken during slot T . We observe that the time evolution of the total (over pages) staleness is:

$$\sum_{j \in \mathcal{W}} s_j(t+1) = \sum_{j \in \mathcal{W}} s_j(t) + (|\mathcal{W}| - 1) - s_{j^*(t)}(t), \quad (3.5)$$

where $j^*(t)$ is the selected web page at time slot t with staleness $s_{j^*(t)}(t)$ (which becomes zero at the beginning of slot $t+1$). We observe that for $t = T-1$, and after a number of calculations, equation (3.5) can be written as:

$$\sum_{j \in \mathcal{W}} s_j(T) = \sum_{j \in \mathcal{W}} s_j(0) + T(|\mathcal{W}| - 1) - \sum_{t'=0}^{T-1} s_{j^*(t')}(t').$$

Therefore, for any $t = 0, \dots, T-1$, we have:

$$\sum_{j \in \mathcal{W}} s_j(t) = \sum_{j \in \mathcal{W}} s_j(0) + t(|\mathcal{W}| - 1) - \sum_{t'=0}^{t-1} s_{j^*(t')}(t'), \quad (3.6)$$

and using the decision variable $\mathbf{x}(t)$, (3.6) can be written as:

$$\sum_{j \in \mathcal{W}} s_j(t) = \sum_{j \in \mathcal{W}} s_j(0) + t(|\mathcal{W}| - 1) - \sum_{t'=0}^{t-1} \sum_{j \in \mathcal{W}} x_j(t') s_j(t'). \quad (3.7)$$

Using (3.7), the total staleness $\sum_{t=0}^T \sum_{j \in \mathcal{W}} s_j(t)$ can now be written as:

$$(T+1) \sum_{j \in \mathcal{W}} s_j(0) + \frac{T(T+1)}{2} (|\mathcal{W}| - 1) - \sum_{t=0}^T \sum_{t'=0}^{t-1} \sum_{j \in \mathcal{W}} x_j(t') s_j(t'). \quad (3.8)$$

In order to minimize (3.8) we just have to maximize the quantity

$\sum_{t=0}^T \sum_{t'=0}^{t-1} \sum_{j \in \mathcal{W}} x_j(t') s_j(t')$. Moreover, in order to accommodate constraint (3.4) in our problem, we include it in the objective function parameterized by Lagrange multiplier $\lambda \in \mathbb{R}_+$. Thus, we end up with the following objective:

$$\max_{\mathbf{x}^*} \sum_{t=0}^T \sum_{t'=0}^{t-1} \sum_{j \in \mathcal{W}} x_j(t') s_j(t') - \lambda \left(\sum_{t=0}^{T-1} \sum_{j \in \mathcal{W}} x_j(t) e_j g(t) - \bar{G} \right). \quad (3.9)$$

Parameter λ denotes the significance of greenness for the server. Its value, which is set in consultation with the crawler, depends on the capability of using green energy at the local server premises. If there is no such capability, then $\lambda = 0$ and the crawler's goal is to just minimize the total staleness. On the other hand, if there is high potential for energy from renewable sources, the value of λ is high. In that case, besides minimizing staleness, the crawler also wishes to keep the carbon footprint as low as possible, as suggested by the values of λ and \bar{G} . We observe that as λ increases, the amount of carbon emissions becomes more and more important for the server and the page download scheduling is largely determined by its value.

3.3.2 Optimal Web Page Download Scheduling Policy

Now, we will try to understand the structure of objective (3.9). We write (3.9) for $T = 2$ and after some algebra we get:

$$\max_{\mathbf{x}^*} \sum_{j \in \mathcal{W}} (2s_j(0) - \lambda e_j g(0)) x_j(0) + \sum_{j \in \mathcal{W}} (s_j(1) - \lambda e_j g(1)) x_j(1) + \lambda \bar{G}. \quad (3.10)$$

For $T > 2$, the objective is written as:

$$\begin{aligned} \max_{\mathbf{x}^*} \sum_{j \in \mathcal{W}} (Ts_j(0) - \lambda e_j g(0))x_j(0) + \sum_{j \in \mathcal{W}} ((T-1)s_j(1) - \lambda e_j g(1))x_j(1) + \dots + \\ + \sum_{j \in \mathcal{W}} (s_j(T-1) - \lambda e_j g(T-1))x_j(T-1) + \lambda \bar{G}. \end{aligned} \quad (3.11)$$

Since $g(t)$ is not known a priori, we study the online version of the problem. Objective (3.11) can be decomposed to separate terms to be optimized with respect to the scheduling decision only at slot t , which means that the crawler's decision at t is independent of the decisions made at the other slots. Thus, the optimal policy involves greedy decision making.

Optimal policy: At the beginning of each slot t , the crawler chooses $\mathbf{x}(t)$ that maximizes

$$\sum_{j \in \mathcal{W}} ((T-t)s_j(t) - \lambda e_j g(t))x_j(t)$$

such that $\sum_{j \in \mathcal{W}} x_j(t) = 1$. For the selected web page $j^*(t)$, it holds:

$$j^*(t) = \arg \max_{j \in \mathcal{W}} ((T-t)s_j(t) - \lambda e_j g(t)),$$

which means that at each time t , given $g(t)$, the crawler decides based on the values of $s_j(t)$ and e_j , for $j \in \mathcal{W}$.

In the special case where all pages need the same energy to be downloaded, i.e., $e_j = e$, $\forall j \in \mathcal{W}$, the crawler chooses the page j with the maximum staleness $s_j(t)$. On the other hand, if all pages have the same staleness at t , i.e., $s_j(t) = s$, $\forall j \in \mathcal{W}$, then the crawler chooses the page j with the minimum required energy e_j , i.e., the one with the minimum size p_j .

3.3.3 Extensions to the Model

3.3.3.1 Many Web Servers

Consider a set \mathcal{K} of $N > 1$ web servers where each server is characterized by its own $g_i(t)$ value. Again, we assume that at each slot t , only one thread to a web page is allowed to be active to any server ($m = 1$).

Our goal is to find a policy $\mathbf{x}^* = (\mathbf{x}(t) : t = 0, \dots, T-1)$ that minimizes total staleness (over time and all pages) of all servers, i.e.,

$$\min_{\mathbf{x}^*} \sum_{t=0}^T \sum_{i=1}^N \sum_{j \in \mathcal{W}_i} s_j(t), \quad (3.12)$$

subject to the constraints that only one thread to a web page is allowed to be active to any server at each slot t , i.e.,

$$s.t. \quad \sum_{i \in \mathcal{K}} \sum_{j \in \mathcal{W}_i} x_j(t) = 1, \quad \forall t = 0, \dots, T-1, \quad (3.13)$$

and that the total amount of carbon emissions of each remote server i does not exceed a given threshold \bar{G}_i , i.e.,

$$s.t. \quad \sum_{t=0}^{T-1} \sum_{j \in \mathcal{W}_i} x_j(t) e_j g_i(t) \leq \bar{G}_i, \quad \forall i \in \mathcal{K}. \quad (3.14)$$

where \bar{G}_i is set by the agreement between the crawler and the remote server i . Based on (3.13), we observe that a policy \mathbf{x}^* is a sequence of vectors $\mathbf{x}(t)$, where for given t , $x_j(t) = 1$ for only one i and only one $j \in \mathcal{W}_i$.

Similar to Section 3.3.1, objective (3.12) is converted to:

$$\max_{\mathbf{x}^*} \sum_{t=0}^T \sum_{t'=0}^{t-1} \sum_{i=1}^N \sum_{j \in \mathcal{W}_i} x_j(t') s_j(t'). \quad (3.15)$$

By bringing (3.14) in the objective with Lagrange multipliers $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \dots, \lambda_N) \in \mathbb{R}_+^N$, our problem becomes:

$$\max_{\mathbf{x}^*} \left[\sum_{t=0}^T \sum_{t'=0}^{t-1} \sum_{i=1}^N \sum_{j \in \mathcal{W}_i} x_j(t') s_j(t') - \sum_{i=1}^N \lambda_i \left(\sum_{t=0}^{T-1} \sum_{j \in \mathcal{W}_i} x_j(t) e_j g_i(t) - \bar{G}_i \right) \right], \quad (3.16)$$

subject to (3.13). As in Section 3.3.2, the objective can be decomposed to separate objectives, where each objective needs to be optimized with respect to the scheduling decision only at time slot t . Similarly, the optimal policy in the case of many web servers involves greedy decision making.

Optimal policy for many web servers: At the beginning of each slot t , the crawler chooses the server i and the web page $j \in \mathcal{W}_i$ with the maximum value of $((T-t)s_j(t) - \lambda_i e_j g_i(t))$.

We observe that at each slot t , the crawler makes its decisions based on the values of $s_j(t)$, e_j , λ_i , and $g_i(t)$, for $i = 1, \dots, N$ and $j \in \mathcal{W}_i$. We identify some special cases here:

- If at slot t , $s_j(t) = s$, $\forall i \in \mathcal{K}$, $\forall j \in \mathcal{W}_i$, then the crawler chooses the server $i \in \mathcal{K}$ and the web page $j \in \mathcal{W}_i$ with the minimum product $\lambda_i e_j g_i(t)$. If it is also $\lambda_i = \lambda$ and $e_j = e$, $\forall i \in \mathcal{K}$, $\forall j \in \mathcal{W}_i$, then the crawler chooses randomly a page from the server with minimum $g_i(t)$.
- If $\lambda_i = \lambda$ and $e_j = e$, $\forall i \in \mathcal{K}$, $\forall j \in \mathcal{W}_i$, then at slot t the crawler chooses the server $i \in \mathcal{K}$ and the web page $j \in \mathcal{W}_i$ with the maximum $s_j(t) - g_i(t)$. If also at slot t , $g_i(t) = g(t)$, $\forall i \in \mathcal{K}$, then the crawler chooses out of all pages the web page j with maximum staleness $s_j(t)$.

3.3.3.2 Many Web Servers, Multiple Crawling Threads

Now, we assume that the crawler uses $m > 1$ threads at each slot. The problem is the same as the one above but now, at each slot t , the crawler selects m pages for fetching. This means that in the above problem, constraint (3.13) is transformed to:

$$\sum_{i \in \mathcal{K}} \sum_{j \in \mathcal{W}_i} x_j(t) = m, \quad \forall t = 0, \dots, T-1. \quad (3.17)$$

Again, the optimal policy turns out to involve greedy decision making.

Optimal policy for many web servers and multiple simultaneous crawling threads: At the beginning of each time slot t , the crawler selects the m server-page pairs (i, j) with largest values of $(T-t)s_j(t) - \lambda_i e_j g_i(t)$.

3.3.3.3 Web Pages with Variable Freshness Requirements

Here, we incorporate into our problem the weights γ_j defined in Section 3.2, which indicate the variable freshness requirements of web pages. Our optimization problem becomes:

$$\min_{\mathbf{x}^*} \sum_{t=0}^T \sum_{i=1}^N \sum_{j \in \mathcal{W}_i} \gamma_j s_j(t), \quad (3.18)$$

subject to constraints (3.14) and (3.17).

Optimal policy for many web servers, multiple threads and pages with variable freshness requirements: At the beginning of each slot t the crawler chooses the m server-page pairs (i, j) with largest values of $(T-t)\gamma_j s_j(t) - \lambda_i e_j g_i(t)$.

The following simple example for $m = 1$ shows the influence of weights γ_j on the crawler's decisions: Assume that at slot t the crawler has to decide between two pages k, l hosted in the same server (i.e., the same $g_i(t)$, λ_i). If $s_k(t) = s_l(t)$, $e_k = e_l$ and $\gamma_k > \gamma_l$, the

crawler gives priority to the one with the biggest freshness requirement, i.e., it selects page k .

3.4 Performance Evaluation

3.4.1 Data Set

As a web collection, we used web pages sampled from a large web crawl performed by Yahoo. The web crawling was a continuous process. Therefore, our data represents only a snapshot of this crawl (the snapshot was obtained in November 2011). The data contains the most important web servers (about 500,000 servers) and the pages hosted on those servers (about two billion pages). The importance of a web server was estimated by a proprietary link analysis metric. This metric is similar to PageRank in that the importance of a server is computed based on the number of its inbound links and the estimated importance of servers that provide those links. Hence, our collection is large and also represents high-quality content that is of importance to a web search engine. The scale of the collected dataset makes us confident that the numerical investigation of our problem (Section 3.4.4) is realistic and gives a good sense of the performance in a realistic system.

In our simulations, we estimate the greenness value of a web server based on the time-zone of the country in which it is physically located and the time an HTTP request is issued to the web server. We estimate the country information for each server by means of a proprietary classifier. The classifier assigns each server to a country based on a number of features, including the IP address of the server, some link features (e.g., the country information associated with the servers providing inbound links) and some content features (e.g., the language of pages hosted on the server). For each server, we also compute some information about the number of pages hosted on it as well as its average page size $\bar{p}_i = (\sum_{j \in \mathcal{W}_i} p_j) / |\mathcal{W}_i|$.

3.4.2 Greenness and Staleness Computation

Estimation of Carbon Emissions: First, we assume that each server can be powered by a mixture of solar (green) energy during the day and energy provided by the grid (brown) during the night. The actual solar insolation/irradiance reaching a solar array is strongly dependent on the solar array's position on the Earth and on local weather conditions [52]. It varies throughout the day from 0 kW/m^2 at night to a maximum of about 1 kW/m^2 . Fig. 3.2a shows an example in which the solar irradiance reaches

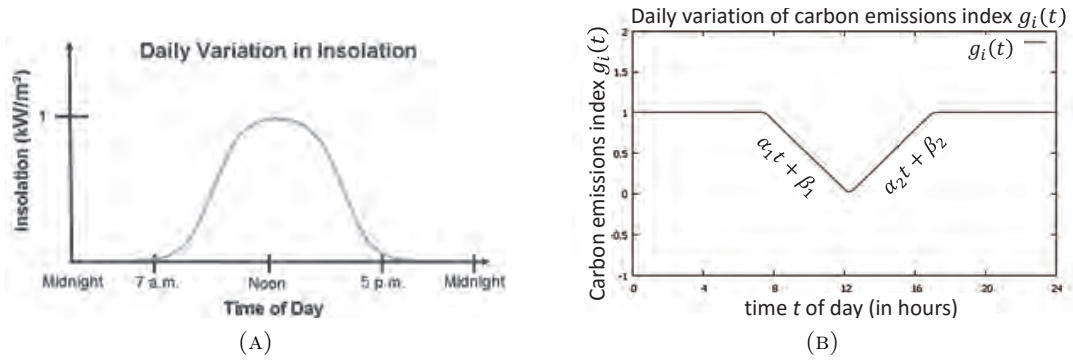


FIGURE 3.2: The variation in (a) the solar irradiance [52] and (b) the $g_i(t)$ values during the day.

its maximum value at noon when the sun is at its highest point in the sky. The actual solar irradiance is usually below this value because it depends on the angle of incidence of the sun's rays with the ground. Here, for simplicity, we ignore the case of clouds on the horizon.

Based on the day of the year, the time-zone, latitude and longitude of each server, we estimate the day length as well as the sunrise and sunset times at the server location. We define $g_i(t) = 1 - \text{normalized solar irradiance}$ (using the maximum solar insolation as a measure of scale). This definition captures the fact that the amount of solar insolation determines the amount of generated solar energy and thus, the amount of produced carbon emissions. We observe that $g_i(t)$ takes its minimum value in the middle of the day, whereas during the night it takes its maximum value 1.

For example, if for a server i located in a country in the northern hemisphere, the sun in winter rises at 07:30 a.m. and sets at 17:00 p.m., the $g_i(t)$ value is 1 at 07:30 a.m., around 0 at 12:15 p.m., again 1 at 17:00 p.m. and 1 during the night. From Fig. 3.2a, we observe that during the intermediate hours the values of $g_i(t)$ will range between 0 and 1 and that they can be approximately determined by a linear function $\alpha_1 t + \beta_1$ between sunrise and noon, and by a function $\alpha_2 t + \beta_2$ between noon and sunset. For the pair of points (07:30,1), (12:15,0), we get $\alpha_1 = -0.21$ and $\beta_1 = 2.57$, while for the pair (12:15,0), (17:00,1) we get $\alpha_2 = 0.21$ and $\beta_2 = -2.57$. Fig. 3.2b shows the daily variation of $g_i(t)$ for the above example.

Note that our approach is transparent to such derivations and may be used in conjunction with various regimes of computing the time evolution of $g_i(t)$. For example, in the case of wind energy, $g_i(t)$ could be computed using wind velocity and direction [53]. Also, $g_i(t)$ could be estimated using the expected energy generation pattern of each renewable

source installed at server i premises, which can be approximated by an average time sequence based on historical data [21].

Staleness Computation: For the experimental scenario, we assume that the scheduling of a thread towards a web page j results in a download time of k_j slots, which is proportional to size p_j . We assume that $k_j = \ell + \frac{p_j}{b}$, where b is the bandwidth and ℓ is the network latency between the crawler machine and the server, which we take for simplicity to be the same across servers. Since the available data does not allow us to estimate the staleness of each page separately, we use a more practical method to compute the total staleness $S_i(t)$ of a server i . For example, let us consider a server n with $|\mathcal{W}_n|$ pages. We assume that at time t_0 a page j from that server is selected to be crawled. After k_j time units, the total staleness of each server i increases by $|\mathcal{W}_i|k_j$, i.e., $S_i(t_0 + k_j) = S_i(t_0) + |\mathcal{W}_i|k_j$, since each page gets k_j time units older. Now, at time $t_0 + k_j + 1$, we assume:

$$S_n(t_0 + k_j + 1) = \frac{S_n(t_0 + k_j + 1)(|\mathcal{W}_n| - 1)}{|\mathcal{W}_n|}, \quad (3.19)$$

i.e., after crawling a page from server n , we simply decrease its current total staleness by its average staleness.

3.4.3 Heuristic Policies

Here, we devise some heuristic page download scheduling policies, along the lines of the optimal policy, in order to evaluate the relative importance of the different system parameters and study the tradeoff between staleness and greenness. Due to the nature of the collected data described in Section 3.4.1, the proposed policies work at the granularity of servers. Each policy achieves a different objective and relies on different parameters to make its decisions, using less information than the optimal policy. We define two more metrics: the average page staleness (over all pages) at slot t as,

$$\bar{s}(t) = \frac{\sum_{i \in \mathcal{K}} \sum_{j \in \mathcal{W}_i} s_j(t)}{\sum_{i \in \mathcal{K}} |\mathcal{W}_i|}, \quad (3.20)$$

as well as the average staleness of a server i at slot t as,

$$\bar{S}_i(t) = \frac{\sum_{j \in \mathcal{W}_i} s_j(t)}{|\mathcal{W}_i|}. \quad (3.21)$$

The proposed heuristic policies are the following:

- *Random Server Selection (RS)*: The crawler picks a server uniformly at random out of all servers in the system. Then, it selects a random page in that server to crawl. We use this policy for benchmarking.
- *Maximum Greenness Server Selection (MG)*: The crawler picks the server with the minimum $g_i(t)$ value (maximum greenness). Then, it selects a random page in that server to crawl. This policy tries to minimize the amount of carbon emissions produced by the crawling process.
- *Maximum Average Staleness Server Selection (MS)*: The crawler selects the server with the maximum average staleness $\bar{S}_i(t)$. Then, it selects a random page in that server to crawl. This policy does not take greenness into account at all. Its goal is to minimize the average staleness of the servers (and thus, that of the pages).
- *Maximum Product of Average Server Staleness and (1-carbon emissions) (MPSC)*: The crawler selects the server with the maximum product $\bar{S}_i(t)(1 - g_i(t))$. Then, it selects a random page in that server to crawl. Among the servers with $g_i(t) = 0$, the crawler chooses the one with the maximum average staleness. The goal is to maintain the amount of carbon emissions and the average staleness of the servers at low levels.
- *Maximum Product of Average Server Staleness and (1/average page size) (MPSS)*: The crawler selects the server with the maximum product $\bar{S}_i(t)(1/\bar{p}_i)$. Then, it selects a random page in that server to crawl.
- *Maximum Product of Average Server Staleness, (1/average page size), and (1-carbon emissions) (MPSSC)*: The crawler selects the server with the maximum product $\bar{S}_i(t)(1/\bar{p}_i)(1 - g_i(t))$. Then, it selects a random page in that server to crawl. Among the servers with $g_i(t) = 0$ (maximum greenness), the crawler prefers servers with large average staleness and small average page size.

By experimenting with the last two policies we try to understand the impact of average page size on system performance in terms of reduction in carbon emissions and staleness.

3.4.4 Experimental Results

Here, we present our experimental results in order to (i) show the performance of the optimal policy in terms of total staleness and total carbon emissions as a function of parameter λ , and compare it with the performance of an EDD-like policy, (ii) show the performance of our heuristics in terms of average amount of carbon emissions and average staleness reduction, and (iii) compare the performance of all proposed policies.

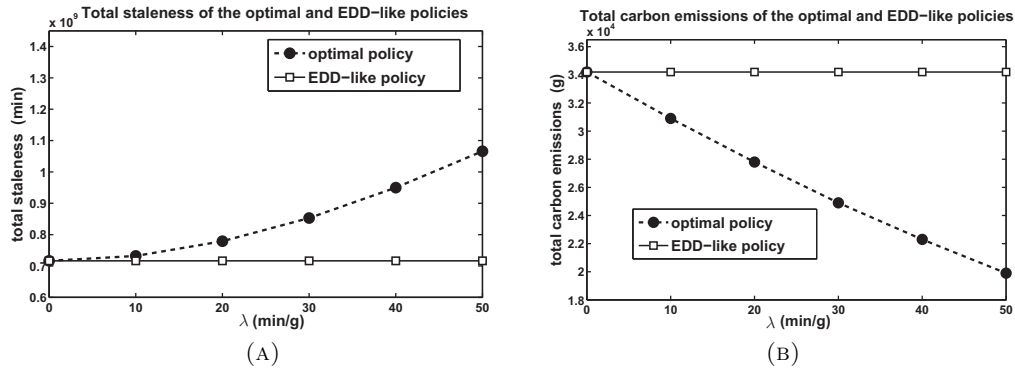


FIGURE 3.3: The performance of the optimal policy in comparison with that of the EDD-like policy in terms of a) total staleness, and b) carbon emissions as a function of λ .

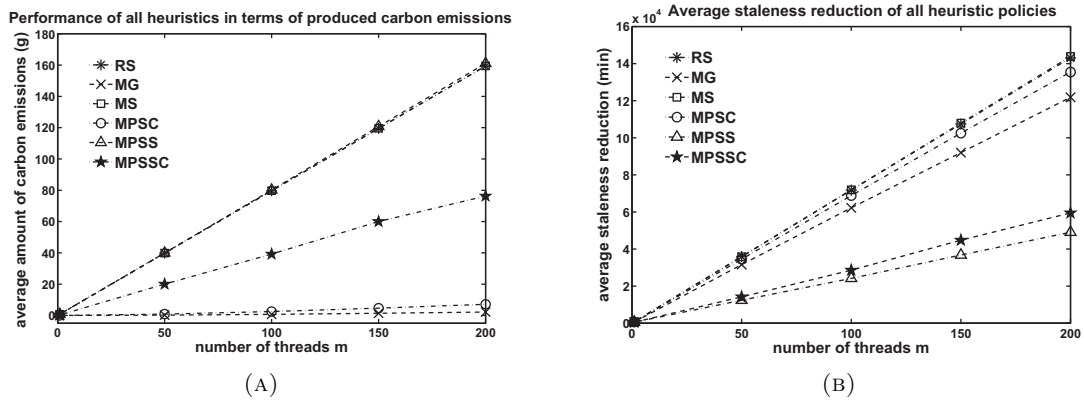


FIGURE 3.4: a) The average amount of carbon emissions (measured in grams (g)) generated by the proposed heuristic policies, and b) the performance of all heuristics in terms of staleness reduction.

We run our experiments for $T=24$ hrs, $\ell=100$ ms, $b=1$ Mbps, and $m=1, 50, 100, 150, 200$.

3.4.4.1 Performance of the Optimal Policy

First, we study the performance of the optimal policy for the case of a single server and a single thread as a function of λ . Here, we use synthetic data as a way to study the performance at the level of web pages. We use a set of 1000 pages and normalized values between 0 and 1 for the energy required for page download. As mentioned in Section 3.3.1, the page download time is one slot, and during this slot $g(t)$ remains stable. The page staleness values are calculated using the method described in Section 3.2 and 3.3.1. The $g(t)$ values are computed according to the method described in Section 3.4.2. We assume that staleness is measured in minutes (min) and that the

amount of carbon emissions at each slot t , $e_j g(t)$, is measured in grams (g). Since the crawler decides based on the value of $(T - t)s_j(t) - \lambda e_j g(t)$, the measurement unit of λ is minutes/grams (min/g).

Figs. 3.3a and 3.3b show the performance of the optimal policy in terms of total (over time and over pages) staleness and total carbon emissions, respectively. As λ increases, the total staleness increases in a non-linear convex manner, whereas the total carbon emissions decrease almost linearly. It is obvious that there is a tradeoff between page staleness and server greenness, and that λ can be set to quantify this tradeoff. These results stem from the fact that as λ increases, the potential of using green energy increases as well and the server wants more and more to exploit this potential. Thus, although the crawler wants to minimize the staleness of pages, it is hindered by the server's desire to keep its carbon footprint low.

Figs. 3.3a and 3.3b also depict the performance of a carbon-unaware web crawling technique that places emphasis only on the freshness of the downloaded web pages. This policy is reminiscent of the earliest due date (EDD) policy [46] since at each slot it selects the web page with the maximum staleness. As expected, this EDD-like policy outperforms our optimal policy in terms of total staleness for all values of λ , whereas the reverse occurs in terms of total carbon emissions. The performance values of the two policies coincide for $\lambda = 0$.

3.4.4.2 Performance of the Heuristic Policies

Here, we study the performance of our heuristics using the collected real data described in Section 3.4.1. Figs. 3.4a and 3.4b show their performance in terms of average (over time) amount of carbon emissions and average (over time) staleness reduction, respectively. As the number of threads increases, the crawler potential increases, and the differences between the policies are more pronounced. We observe that both carbon emissions and staleness reduction constantly grow due to the increasing number of downloaded pages per unit of time.

In Fig. 3.4a, MG outperforms the other five policies in terms of reduced carbon emissions as expected. Compared to the case where there is no possibility of using green energy, i.e., $g_i(t) = 1, \forall i$ and $\forall t$, it achieves on average a 99.35% reduction in the amount of carbon emissions. The other two policies that also take the value of $g_i(t)$ into account, MPSC and MPSSC, are characterized by an average carbon emissions reduction of 97.74% and 63.42%, respectively. Their gain is lower since they also count in server staleness and/or web page size. The tradeoff between staleness and greenness prevents the amount of carbon emissions from being as small as possible.

TABLE 3.1: Performance results as the impact of cold-start is reduced

	Day 1	Day 2	Day 3	Day 4
Avg. staleness reduction of MS (in min)	71.9	215.7	359.1	502.1
Avg. staleness reduction of RS (in min)	71.8	214.2	355.4	495.1
difference (in min)	0.1	1.5	3.7	7.0
difference (%)	0.14	0.67	1.03	1.39

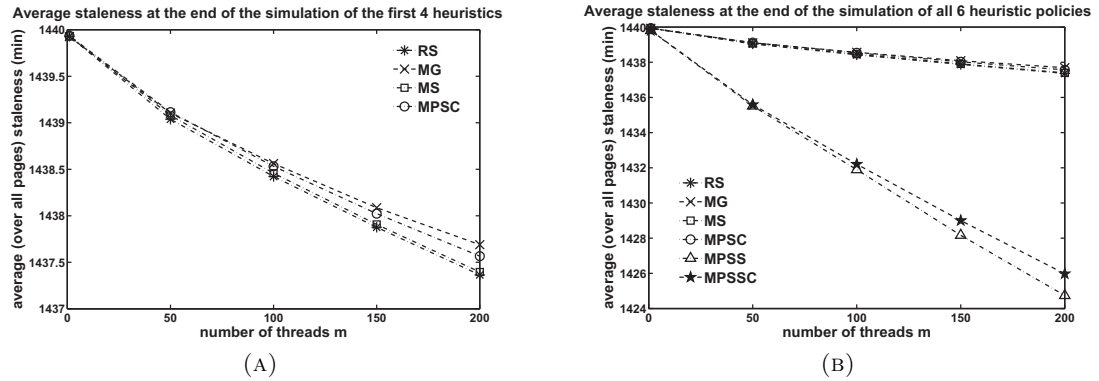


FIGURE 3.5: The average page staleness of a) the first four heuristic policies and b) all heuristics at the end of the simulation.

In Fig. 3.4b, we observe that MS, which can be considered as an EDD-like heuristic policy, outperforms four of the other five policies in terms of staleness reduction. The fifth one (i.e., RS) performs almost the same as MS. This happens due to a “cold start” phenomenon: in the first hours of the simulation, the great majority of servers have the same average staleness since the number of downloaded pages (and hence, the number of staleness reduction events) during this period is too small compared to the number of the pages that are not downloaded. Thus, each time a decision must be made, RS is more likely to select a page from a server whose staleness is at maximum. In order to eliminate the impact of this phenomenon on our results, we ran both RS and MS for three more days ($m = 1$). Table 3.1 shows the results. We observe that as the algorithm runs for more days, the difference between the two policies in terms of staleness reduction builds up.

In Figs. 3.4a and 3.4b, we observe that MPSC tries to strike a balance between staleness and greenness. Besides its relatively high gain in terms of reduced carbon emissions, its average staleness reduction is only 4.34% lower than that of MS.

Figs. 3.5a and 3.5b depict the average page staleness after running the proposed policies for one day. Considering only the first four policies, we observe that RS and MS, as expected, achieve the lowest average page staleness. On the other hand, the performance of MPSS and MPSSC, as shown in Fig. 3.5b, reveals the impact of average page size on

TABLE 3.2: Performance comparison of the proposed policies

	Avg. page staleness (end of day) (min)	Total staleness reduction ($\times 10^6$) (min)	Total carbon emissions ($\times 10^4$) (g)
Optimal 1	125.3	19.72	3.13
Optimal 2	126.2	19.7	3.11
RS	250.1	17.84	3.3
MG	673.0	11.5	0.41
MS	248.6	17.87	3.29
MPSC	431.5	15.12	1.72
MPSS	248.7	17.86	3.28
MPSSC	431.1	15.13	1.7

page freshness. Although the average staleness reduction of both policies remains low (Fig. 3.4b), the average page staleness at the end of the day is much lower than that of the first four policies. Their ability to keep pages more fresh stems from the following fact: Since both policies give priority to pages with relatively small size, the number of downloaded pages during the horizon of the $T + 1$ slots is much greater (88.98% and 87.44% on average for the MPSS and the MPSSC, respectively) than that of the other policies. This yields to an increased total staleness reduction, which offsets the low average staleness reduction per unit of time, and which in turn yields to a relatively low average page staleness at the end of the day.

In Figs. 3.4a and 3.5b, we observe that MPSSC could also be used by the crawler according to our objective. Although the carbon emission reduction of MPSSC is lower than that of MPSC, its average page staleness at the end of the simulation is lower than that of MPSC.

3.4.4.3 Performance Comparison

In order to compare the performance of all policies, we use synthetic data. Our system consists of 15 servers and each server hosts 1000 pages. Again, we use normalized values between 0 and 1 for the energy required for page download. For the optimal policy, we consider the same value of λ for all servers, i.e., $\lambda_i = \lambda$ for $i = 1, \dots, 15$. As we mentioned above, the page download time is one slot, and we assume that during this slot the $g_i(t)$ values, for $i = 1, \dots, 15$, remain stable. At each slot t , $m = 1$ thread is sent to fetch one web page. The decision at each time slot t is to pick 1 out of the 15 servers and a page from the selected server to download.

Table 3.2 shows the performance of all proposed policies. We ran our optimal policy both for $\lambda = 100$ (Optimal 1) and for $\lambda = 500$ (Optimal 2). As we can see, both versions of the optimal policy outperform all the heuristic policies in terms of total staleness reduction and average page staleness at the end of the simulation. This stems from the fact that the optimal policy works at the granularity of web pages. Specifically, for relatively low values of λ , it applies more weight to the staleness factor. Thus, in contrast to all heuristics, which select a web page from the selected server at random, it has the ability to examine in more detail the characteristics of all web pages and pick a web page with relatively high staleness. As λ increases, the servers' desire to keep their carbon footprint low increases as well and thus, the crawler is obstructed from minimizing the staleness of pages. As we mentioned above, λ is a parameter that can be set to quantify the tradeoff between staleness and greenness. On the other hand, for the specific values of λ , the policy that outperforms all the other ones in terms of reduced carbon emissions is MG. However, the produced carbon emissions of the optimal policy (in both cases) are less than those of the RS, MS (EDD-like heuristic) and MPSS policies which do not consider the greenness of the servers at all.

3.5 Related Work

3.5.1 Refreshing Web Repositories

The pages in the web repository of a crawler need to be refreshed to prevent the search engine from presenting stale results to its users. A research problem here is to find an order in which the pages will be refreshed such that some freshness metric is optimized over the entire repository. Early papers on the topic suggest refreshing frequently updated pages more often, relying on the observed update frequency of pages as a proxy [54], [55]. Certain works use the update history of related pages to capture the actual update likelihood of a page better [56], [57]. Works in another line devise page refresh strategies that aim to minimize the negative impact on users due to stale search results [58], [59]. Finally, the work in [60] suggests avoiding to refresh fast-changing content as much as possible. To the best of our knowledge, so far, no prior work has investigated this research problem taking into account the greenness aspect of web servers as a constraint.

3.5.2 Job and Packet Processing

A line of work that is relevant to ours is that of scheduling jobs under deadlines. The work in [46] presents the earliest due date (EDD) scheduling policy for wire-line networks. In this policy, at each time, the job with the earliest due date is scheduled, hence minimizing

the expected lateness. From a wireless networking perspective, the work in [47] presents a wireless channel-aware version of EDD, the feasible earliest due date (FEDD) policy. The authors of [48] study the problem of scheduling constant-bit-rate traffic over wireless channels and devise a policy that minimizes the packet loss rate due to packet delivery deadline expiration. The work in [61] studies the design of a downlink packet scheduler for real-time multimedia applications and proposes a channel- and QoS-aware version of EDD, named CA-EDD. We note that none of the above works considers energy efficiency.

A number of works aim to minimize the total energy consumption. In [62], a CPU scheduler is presented for mobile devices. This scheduler integrates dynamic voltage scaling into soft real-time scheduling and decides when, how long, and how fast to execute multimedia applications based on their demand distribution. In [63], the authors propose offline and online packet scheduling algorithms for uplink and downlink in wireless networks with constraints imposed by packet deadlines and finite buffers. The reader may refer to [64] for a survey of studies on energy-efficient scheduling without deadlines. In contrast to these works, our work tries to minimize the total staleness of a web page repository while keeping the amount of carbon emissions of web servers below a given threshold.

3.5.3 Energy Efficiency and Greenness of Data Centers

There is significant amount of work on reducing the energy consumption of data centers. The work in [65] aims to optimize the workload, power, and cooling system management of a data center with the objective of saving energy. In order to emphasize the role of communication fabric in energy consumption, the work in [66] presents a scheduling technique which makes decisions based on a run-time feedback from data center switches and links. The work in [67] uses delay-tolerant jobs to fill the extra capacity of data centers and designs energy-efficient mechanisms that achieve good delay performance. The survey in [68] analyzes software- and hardware-based techniques and architectures as well as mechanisms to control data center resources for energy-efficient operations.

A different line of works focuses on increasing renewable energy utilization and reducing the carbon footprint of data centers. Specifically, the authors of [53] propose an adaptive data center job scheduler that leverages green energy prediction to reduce the number of canceled jobs due to lack of available green energy. The work in [69] focuses on renewable energy capacity planning and proposes an optimization-based framework to achieve specified carbon footprint goals at minimal cost. The work in [70] proposes a policy for request distribution across data centers with the objective of data center cost minimization while enabling Internet services to leverage green energy and respect their

SLAs. The work in [71] investigates three issues related to the feasibility of powering an Internet-scale system exclusively with renewable energy: the impact of geographical load balancing, the optimal mix of renewable sources, and the role of storage.

Data centers hosting cloud applications consume huge amounts of energy. The works in [72], [73], and [74] use virtualization as a power management and resource allocation technique for energy efficiency in cloud computing environments. In particular, the work in [72] assumes deterministic virtual machine (VM) demands, whereas stochastic VM demands are considered in [73]. The key idea of the approach presented in work [74] is to match the VM load with the renewable energy source (RES) provisioned power with the goal of minimizing the total cost of power consumption for the cloud provider. The work in [75] proposes a graph-based approach which utilizes Voronoi partitions to control the operation of a cloud system with the goal of minimizing a combination of average request time, electricity cost, and carbon emissions.

The works surveyed above (subsections 3.5.1, 3.5.2, 3.5.3) are related either to refreshing the repositories of crawlers, effective processing of jobs, or energy efficiency and greenness of data centers. Our work introduces the green web crawling problem, which combines these research threads. It brings together the scheduling concepts with the problem of reducing the staleness of a web repository while limiting the carbon emissions incurred to web servers. We believe that this combination, along with the observation that the type of the energy consumed by a web server varies in time, are unique to our work.

3.6 Conclusion

In this chapter, we introduced the problem of green web crawling: minimizing the total staleness of pages in the repository of a web crawler while keeping the amount of carbon emissions on web servers, due to HTTP requests issued by the crawler, low enough. We devised an optimal policy, which can be implemented in an online fashion, based only on the instantaneous values of page staleness and greenness indicators of the servers. We also devised some heuristics along the lines of the optimal policy and studied their performance through experiments with real data.

Chapter 4

Optimal Design of Serious Games for Smart Grid Consumer Engagement

Contents

4.1	Introduction to the Concept of Serious Games for Demand-side Management	67
4.2	Serious Game Setup	69
4.3	System Model	70
4.3.1	Consumer's problem	72
4.3.2	Game Designer's Problem	76
4.3.3	Equilibrium of the Stackelberg game arising from serious-game interactions	80
4.4	Simulation Results and Analysis	81
4.4.1	Experimental Setup	81
4.4.2	Steady-State Convergence and Wear-off Effects	81
4.4.3	Experimental Results	83
4.5	Related Work	85
4.6	Conclusion	86

4.1 Introduction to the Concept of Serious Games for Demand-side Management

The reduction of carbon footprint is the holy grail of our times and is to be realized primarily through prudent energy consumption. Different techniques have been proposed, spanning the entire chain of energy generation, transmission, distribution and consumption in the context of realizing a smart energy grid toward reaching the goal above. However, the weakest link in the chain above remains the end-consumer. No matter how sophisticated these techniques become for the rest of the chain, it is the end-consumer that determines to a large extent the mode of energy consumption in the end.

Demand-side management (DSM), which includes demand response (DR)¹, is an active research area that aims to reduce or smoothen energy consumption. Pricing-based and incentive-based DR schemes have been proposed; e.g., by having different prices per unit of energy presented to consumers for different times of the day (i.e., TOU pricing), rational consumers are forced to shift part of their demand-load from peak-times to off-peak times. Incentives may be provided to the consumer in the form of monetary or non-monetary rewards. For example, in the Critical Peak Rebate incentive scheme, participants are paid for the amounts of power by which they reduce consumption below their predicted consumption levels during critical peak hours. In the presence of a DR scheme, consumers optimize some form of utility functions that factor the monetary gains from load shifting/reduction and the inconvenience cost induced by the shift.

The main shortcomings of pricing-based and incentive-based schemes are the strong rationality assumptions about consumers. They assume and that consumers are inherently driven by financial rationality and that the response of consumers to incentives is governed by the mathematics of optimization theory. Also, more often than not, consumers are not the bill-payers (e.g., they are employees in an office building or younger/elderly household members), while their energy literacy level (i.e., awareness and capability to act on energy savings) varies.

Various incentives schemes may often see hesitation and even negativity of consumers, primarily because they are not presented to consumers in an appropriate context. Consumers are humans and more often than not, they are driven in their decisions by

¹DR mechanisms involve behavioral changes in energy consumption as a response to certain signals to the consumers, e.g., price changes, rewards, etc.

factors other than financial motives, such as various biases and sentiments, or the desire to outperform others. The interface with which the incentive scheme is delivered to the consumer is important, in the sense that should make the interaction worthwhile and entertaining, and it should respect the limited time and attention budget of consumers.

Serious games design is an emerging area that aims at addressing precisely the issues of educating and maximizing user engagement in various contexts [76]. A serious game is a game that is designed for a purpose that goes beyond that of offering pure entertainment. There have been some initial attempts to use serious games for demand-side management [77–80] with great success in realistic case studies [81]. To the best of our knowledge, building a foundational theory on modeling and understanding serious games with the purpose of extracting guidelines for serious games design in this context has not hitherto been explored.

In this chapter, we introduce the problem of *optimal serious-games design for the purpose of enforcing prudent energy consumption*, and we make a first attempt to mathematically model it. We define a simple serious-game scenario that does not employ direct monetary incentives for the consumers and a generic game-theoretic mathematical framework for the optimization of the parameters of the serious-game. We assume that a serious-game designer entity (e.g., the energy supplier, a private entity or an energy-efficiency minded administrative authority) aims to design a serious game for smoothening the energy consumption behavior of consumers across time. The serious-game designer runs daily contests on energy consumption reduction on behalf of the players during the peak hours.

We consider a simple class of serious games, where the serious-game designer publicly announces a list of top- K consumers and a list of bottom- M consumers according to their respective energy-consumption reduction at the contest of the previous day. Here, the driving forces are the user discomfort due to consumption-load reduction, the user desire to enter to top- K list and the user sensitivity to social outcasting if she/he enters the bottom- M list.

We formulate the problem faced by the serious-game designer as an operational-cost minimization problem for the utility company and that of the consumer as a utility-maximization one. The serious-game-design problem is to decide on K, M and on the feedback provided to the consumers, while the consumer-side problem amounts to selecting the consumption-load reduction that maximizes the user utility under the serious-game rules. These two problems constitute the two stages of an overall *Stackelberg-game* setting, where the choice of the parameters K, M affects the optimal responses of consumers in terms of their energy-consumption reduction. Most importantly, the serious-game designer does not necessarily need any information about the user-utility parameters that characterize the user discomfort due to the consumption-load reduction

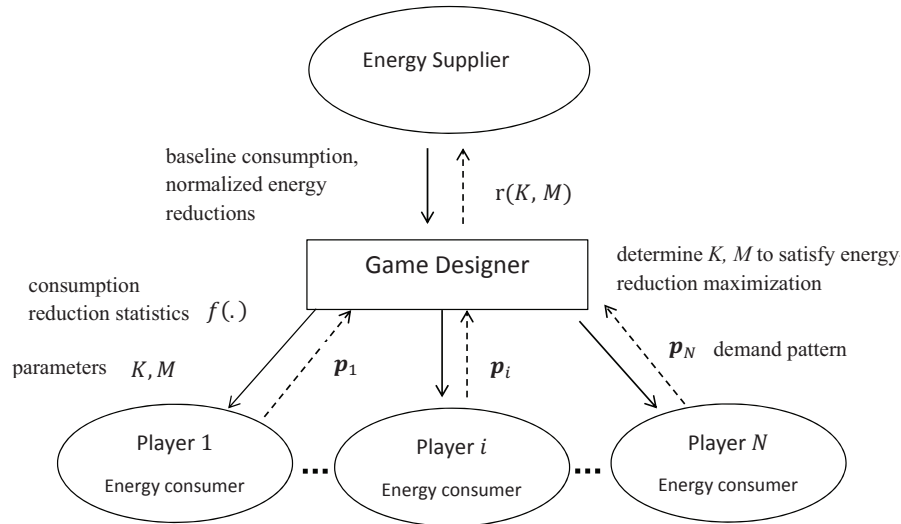


FIGURE 4.1: The serious-game interactions among consumers, the serious-game designer and the utility company.

or the importance for each user to be in either list. Through analysis and a series of numerical simulations, we show how the various serious-game-design choices affect the energy consumption reduction and investigate the driving forces in this serious game. Evidently, even such a simple serious-game design can provide adequate incentives to users for engaging in demand-side management.

The remainder of this chapter is structured as follows: In Section 4.2, we describe the setup of our serious game. In Section 4.3, we analyze our system model and formulate its constituent problems, namely the consumer problem and the serious-game-designer problem. In Section 4.4, we experimentally analyze the effect of various serious-game-design parameters. Section 4.5 reviews the related work. Finally, Section 4.6 concludes our work. This chapter is based on works [82] and [83].

4.2 Serious Game Setup

We consider a energy supplier that aims to reduce the total energy generation (or operational) cost. If the energy generation cost is a convex function of demand, this problem is equivalent to that of smoothening the total energy consumption across time. The energy supplier has instrumented its customer premises with smart-meters. Based on consumption baselines [84] that are dynamically calculated on a sliding window of historical data, it identifies the time slots where peak overall consumption is observed. We also consider a serious-game designer to which the energy supplier outsources the task

of designing a gamification platform for reducing the total energy generation cost. The serious-game designer organizes contests about energy consumption reduction among consumers for each peak time slot of the day. On a daily basis, the customers play the serious game for a specific peak time slot by adjusting their energy demand in that time slot of the next day. Their actions are monitored by smart meters and are communicated to the serious-game designer by the energy supplier.

The next day of each serious-game instance at a peak time slot, the serious-game designer ranks customers according to their relative energy-consumption reduction in the corresponding time slot of the previous day, and it selects the top- K ones to be the “winners” and the bottom- M to be the “losers” of this day-round of the serious game at the particular time slot. K, M are serious-game parameters determined by the designer. The outcome of each competition and the serious-game parameters K, M are announced to customers via the mobile app and may also be shared with customer friends in various social applications. The serious-game designer keeps statistics on customer performance. Based on these statistics, it communicates to the customers useful historical information (e.g., in the form of histograms) on the performance effort that, if exerted in the next day, is expected to include them in the upper list of top- K consumers or in the bottom list of bottom- M consumers of the serious game the day after the next one. The overall serious-game scenario is depicted in Fig. 4.1.

This approach for “winner” and “loser” determination aims to exploit the desire of customers for social approval [85]. On the one hand, a consumer that is included in the upper list may feel proud for her accomplishment relative to others. On the other hand, the inclusion of a player in the lower list is intended to make her feel embarrassed and perhaps socially pinpointed for her bad consumption behavior, so as to force her to be more active in the future if she is sensitive on that issue. However, the extent to which the consumption behavior of each player is affected depends on her sensitivity to the social approval or outcasting.

4.3 System Model

We consider an energy supplier, N consumers (residential users) and a serious-game designer entity. The serious-game designer sets the parameters K and M that determine the sizes of the upper list \mathcal{U} of “prudent” consumers and of the lower list \mathcal{L} of “non-prudent” consumers, respectively. We assume a fixed unit price for the electricity denoted by q \$/kWh. Time is divided into slots. We focus on a serious-game instance that concerns energy-consumption reduction of consumers at a specific time slot. Let p_i be the energy demand of consumer/player i , $i = 1, \dots, N$, at the target time slot with

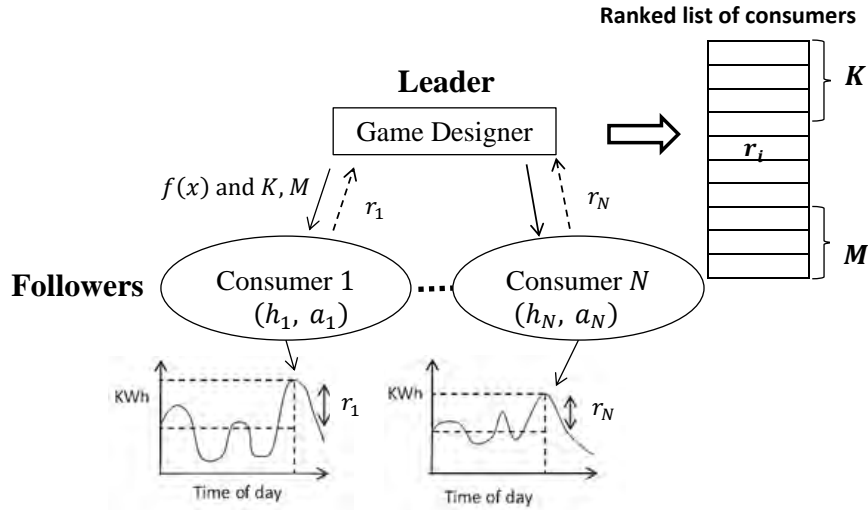


FIGURE 4.2: The Stackelberg game structure arising from serious-game interactions.

$0 \leq p_i \leq p_{i,max}$, where $p_{i,max}$ is an upper bound on the demand of consumer i at each time slot (given fixed power capacity of residential circuit). Without loss of generality, we assume that $p_{i,max} = p_{max}$, for $i = 1, \dots, N$. Also, let $P = \sum_{i=1}^N p_i$ be the total energy demand of all consumers at that time slot.

We denote by $r_i = (p_i^0 - p_i)/p_i^0$ the *normalized energy reduction* of player i , where p_i^0 is the baseline (nominal) energy demand of consumer i which is assumed to be known to the supplier and the game designer. Note that p_i is the modified energy demand as a result of the serious game. We refer to r_i as the *consumption reduction performance* of consumer i . Consumers are ranked by the serious-game designer according to their performance. Each customer i adjusts her demand p_i , so as to maximize her net utility, to be defined in the sequel, given the values of the parameters K, M set by the serious-game designer. The game designer, on the other hand, aims to optimally select the values of K, M , so as to minimize the operational cost of the energy supplier. It also chooses the type of statistical feedback to provide to consumers; this will be taken into account by consumers when deciding on their strategy.

The overall serious-game design problem resembles the two levels of a *Stackelberg* game, where the serious-game designer is the leader and the residential customers are the followers (see Fig. 4.2). The serious-game designer aims to optimally select the parameters K, M , so that the consumption strategies of consumers, in their effort to enter the upper list and/or avoid the bottom list, result in such energy consumption that minimizes energy generation cost. First, we formulate the consumer problem and then we proceed to the game designer's problem.

4.3.1 Consumer's problem

The utility of a consumer i is quantified through a utility function that factors her dissatisfaction caused by energy-consumption reduction, her desire to enter the top- K list \mathcal{U} and her sensitivity in being included in the bottom- M list \mathcal{L} .

4.3.1.1 Dissatisfaction

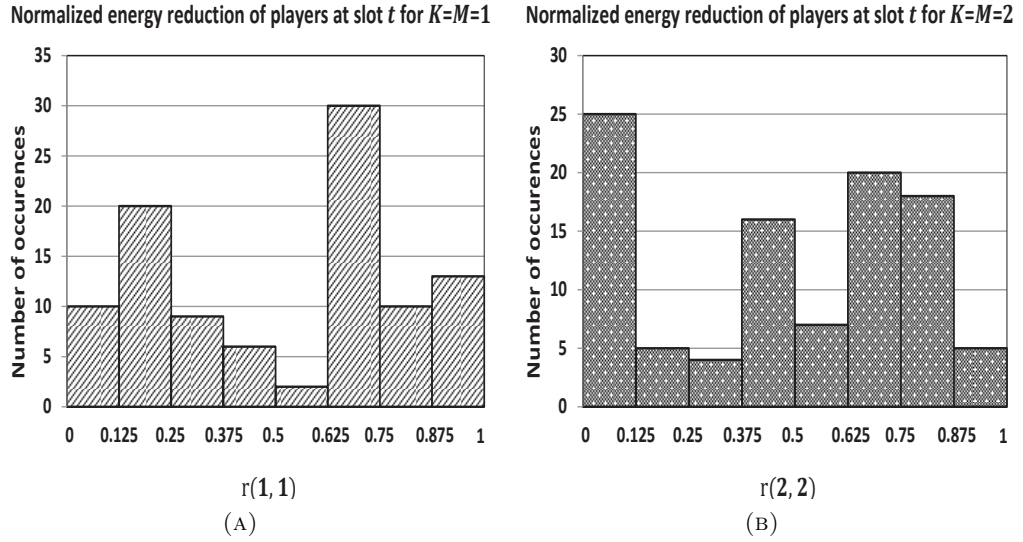
In order to capture the dissatisfaction caused due to deviation of the modified demand p_i from the nominal demand p_i^0 of consumer i , we define a function $d_i(r)$. We assume that $d_i(\cdot)$ is increasing and convex so as to model the increasing marginal dissatisfaction for the consumer as a function of deviation from the nominal consumption pattern, as in [86]. Here, we assume

$$d_i(r) = a_i r^2, r \in [0, 1], \quad (4.1)$$

while $d_i(r) = 0$ for $r < 0$. Parameter $a_i \geq 0$ is referred to as the *inelasticity parameter* of consumer i and it models her tolerance or sensitivity on the deviation from the nominal consumption profile. The intuition of such a simple generic dissatisfaction model, which is chosen here for algebraic tractability, is the following: Assume a set of energy consumption activities of a consumer in a time slot. Each activity involves a number of appliances. Reducing the overall energy consumption at the time slot is done by shifting less important activities at a different time slot or by canceling them. As the total shed energy increases, more important consumer activities need to be shifted/canceled and thus, the marginal dissatisfaction of the consumer increases more per shed energy unit. Considering more sophisticated dissatisfaction models, e.g., as in [87], is left for future work.

4.3.1.2 Social recognition

Social recognition can be defined as the appreciation an observer holds for the person she observes. It is a three-mode phenomenon, since one may receive positive, neutral, or negative recognition. Here, we consider that each consumer elicits one of the three levels of social recognition based on her performance strategy, which leads to one of three possible outcomes: (a) she is included in the upper list \mathcal{U} of the top- K consumers, (b) she is included in the bottom list \mathcal{L} of the bottom- M consumers, (c) neither of the above. Social recognition for consumer i can be modeled as a reward h_i or a penalty $-h_i$, according to her position and the value she places on social recognition. For each consumer i , we define a social recognition function $S_i(r_i, \mathbf{r}_{-i})$ as function of consumer

FIGURE 4.3: Sample performance histograms of players at a time slot; $\delta=0.125$.

i 's action r_i and others' action vector \mathbf{r}_{-i} ,

$$S_i(r_i, \mathbf{r}_{-i}) = \begin{cases} h_i, & \text{if } i \in \mathcal{U} \\ -h_i, & \text{if } i \in \mathcal{L} \\ 0, & \text{otherwise.} \end{cases} \quad (4.2)$$

The dependence on r_i and \mathbf{r}_{-i} is implied in the inclusion in the lists. The sensitivity $h_i \geq 0$ of a consumer quantifies the importance for the consumer of being included in the \mathcal{U} or \mathcal{L} list. For reasons of symmetry, we assume that for each consumer, the emotional reward extracted from pride for being in the list of best consumers is equal to the emotional cost due to embarrassment for being in the list of worst consumers. For example, if a player i with a great value of h_i is included in \mathcal{U} , then she receives a social recognition, which is considered rewarding for her. On the other hand, if she is included in \mathcal{L} , she receives a negative recognition that makes her feel embarrassed. Being somewhere in the middle in terms of consumption is assumed to be a neutral situation with no reward or penalty inflicted to the consumer.

4.3.1.3 Feedback to consumers

In order to aid consumers in taking energy consumption decisions, the game designer may provide cumulative statistics about consumer actions (e.g., in the form of the histograms (Fig. 4.3)) as feedback to them. This feedback may take the form of the empirical cumulative distribution function (CDF) $F(x)$ and empirical probability density function (PDF) $f(x)$, where x stands for consumer performance concerning energy reduction. All

consumers are treated as one virtual consumer with a PDF $f(x)$ and CDF $F(x)$, namely these show the empirical probability distribution of energy reduction of *any* consumer in the past. For example, $F(x)$ may be interpreted as the percentage of times that the energy reduction performance of a consumer is below x . The absence of such feedback may be modeled by taking $F(x), f(x)$ as corresponding to a uniform distribution over a certain range of values. Other types of statistical information may be available.

4.3.1.4 Probability of inclusion in the top-K and the bottom-M lists

Denote X_j the random variable of the energy consumption reduction of player j and $F_j(\cdot)$ the cumulative distribution function (CDF) of her distribution. Also, denote $F(\cdot)$ the CDF of the joint performance distribution for all consumers and X the random variable of the performance of any individual consumer. Given a normalized energy consumption reduction r_i of player i , a rational software agent residing at the consumer side, e.g., at the mobile app calculates the probability that consumer i is included in the upper list (\mathcal{U}) or in the bottom list \mathcal{L} . These probabilities are computed as:

$$\begin{aligned} \Pr(i \in \mathcal{U}|r_i) &= \sum_{l=0}^{K-1} \binom{N-1}{l} \prod_{j=1}^l (1 - F_j(r_i)) \prod_{j=1}^{N-l-1} F_j(r_i) \\ \Pr(i \in \mathcal{L}|r_i) &= \sum_{l=0}^{M-1} \binom{N-1}{l} \prod_{j=1, j \neq i}^l F_j(r_i) \prod_{j=1, j \neq i}^{N-l-1} (1 - F_j(r_i)) \end{aligned} \quad (4.3)$$

$\Pr(i \in \mathcal{U}|r_i)$ (resp. $\Pr(i \in \mathcal{L}|r_i)$) considers all combinations of other consumers ranked above (resp. below) consumer i , when consumer i takes any position in the top- K (resp. bottom- M) list. $F_j(\cdot)$ could be estimated based on a histogram on the performance of consumer j . However, there are practical difficulties toward this direction. First, this would require $O(N)$ storage space for the histograms of one K, M pair. Second, all these histograms would have to be communicated to all consumers prior to each serious-game round, thus creating a high communication overhead. Third, this approach would raise privacy concerns for the consumers. Fourth, the statistical significance of individual histograms would be low, since the number of individual choices is much lower than the total number of choices of all consumers. Given the statistical information $f(r), F(r)$ about the ensemble of consumers, a consumer may assume $F_j(\cdot) = F(\cdot), \forall j \neq i$ for calculating her probabilities to be included in the top or the bottom lists (according to equation (4.3)).

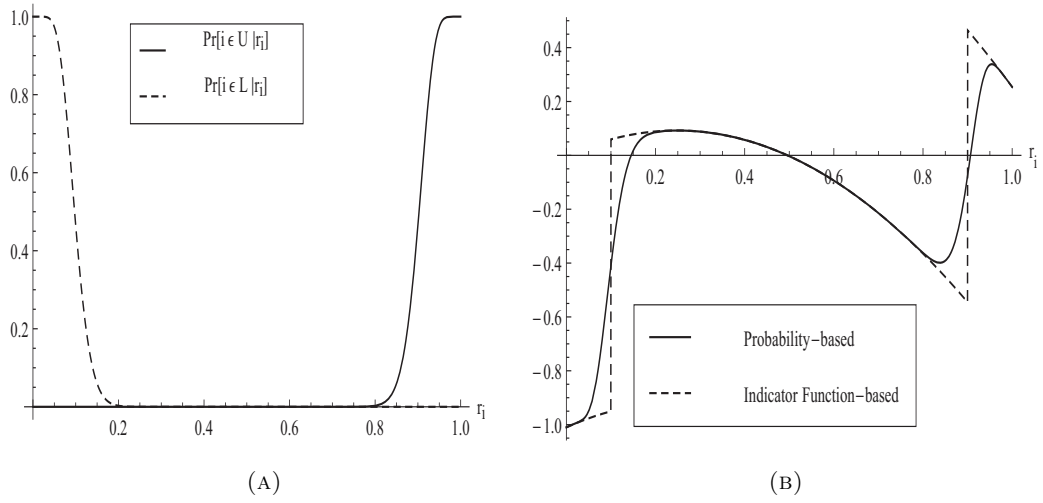


FIGURE 4.4: (A) Probability of consumer i to be included in the top- K or in the bottom- M lists with respect to r_i . (B) Expected utilities of consumer i based on equation (4.3) (“probability-based”) and on equation (4.6) (“indicator function-based”) with respect to r_i . Parameters values: $X \sim U[0, 1]$, $N = 100$, $K = M = 10$, $a_i = 1.5$, $h_i = 1.01$, $p_0 = 6$ kWh, $q = 0.124$.

The probabilities in equation (4.3) can be simplified according to the following observations. Observe that $\Pr(i \in \mathcal{U} | r_i)$ increases with r_i and $\Pr(i \in \mathcal{L} | r_i)$ increases as r_i decreases. Denote r_u (resp. r_l) the performance threshold above (resp. below) which the expected number of consumers is K (resp. M). r_u, r_l can be calculated as follows:

$$\begin{aligned} \sum_{j=0}^{(1-r_u)/\delta} f(X = r_u + j\delta) &= \frac{K}{N} \\ \sum_{j=0}^{(r_l-r_{min})/\delta} f(X = r_{min} + j\delta) &= \frac{M}{N}, \end{aligned} \quad (4.4)$$

where $r_{min} = (\min_i p_i^0 - p_{max}) / \min_i p_i^0$ is the minimum r_i . Interestingly, both the $\Pr(i \in \mathcal{U} | r_i)$ with $r_i \geq r_u$ and $\Pr(i \in \mathcal{L} | r_i)$ with $r_i \leq r_l$ approach 1 for various numerical instantiations of the problem (see Fig. 4.4a). More formally, it can be assumed that:

$$\begin{aligned} \sum_{j=0}^{(1-r_u)/\delta} \Pr(i \in \mathcal{U} | r_i = r_u + j\delta) &\approx 1 \\ \sum_{j=0}^{(r_l-r_{min})/\delta} \Pr(i \in \mathcal{L} | r_i = r_{min} + j\delta) &\approx 1 \end{aligned} \quad (4.5)$$

We can then simplify (4.3) via the following approximations:

$$\begin{aligned}\Pr(i \in \mathcal{U} | r_i \geq r_u) &= \mathbb{1}_{r_i \geq r_u} \\ \Pr(i \in \mathcal{L} | r_i \leq r_l) &= \mathbb{1}_{r_i \leq r_l},\end{aligned}\tag{4.6}$$

where $\mathbb{1}_A$ is an indicator function on event A with $\mathbb{1}_A = 1$ or 0 if event A is true or false respectively.

4.3.1.5 User utility function

We define the expected utility of each player i with strategy r_i as follows:

$$\mathbb{E}\{u_i(r_i)\} = \Pr(i \in \mathcal{U} | r_i)h_i - \Pr(i \in \mathcal{L} | r_i)h_i - d_i(r_i) + qp_i^0 r_i,\tag{4.7}$$

This definition captures the net payoff for consumer i from normalized load reduction r_i , which consists of the expected social reward/cost (if any) minus the incurred dissatisfaction, plus her energy consumption savings from load reduction. Note that the benefit from energy consumption savings becomes cost if $r_i < 0$. The expected utilities of the consumer i based on the formulas of equations (4.3) and (4.6) are depicted in Fig. 4.4b for certain numerical parameters. As becomes evident from this figure, the simplifying assumptions of equation (4.6) can be quite accurate approximations of equation (4.3).

The objective of each consumer i is to determine a normalized consumption reduction r_i that achieves optimal balance of the factors above in the sense of maximizing her expected utility function. Thus, based on (4.3) and (4.7) the consumer optimization problem becomes:

$$\max_{r_i} (\mathbb{1}_{r_i \geq r_u} - \mathbb{1}_{r_i < r_l})h_i - d_i(r_i) + qp_i^0 r_i\tag{4.8}$$

4.3.2 Game Designer's Problem

We denote the generation/operational cost for the energy supplier with an increasing convex function $C(P)$ of total demand load P at a specific time slot (i.e., marginal energy generation cost increases with energy demand). Under the current common practice of flat pricing, a fixed price q is charged per unit of energy consumption irrespectively of the consumption time. Therefore, the energy supplier maximizes its revenue by minimizing the operational expenses for satisfying the energy demand.

The supplier aims to minimize energy generation costs, and it delegates this task to the game designer, with the anticipation that the appropriate induced user consumption

behavior will achieve the goal. The serious-game designer selects the settings of the serious game which in our case are abstracted as (i) size K of the best-consumer list, (ii) size M of the worst-consumer list and (iii) feedback $f(\cdot)$ to be provided to consumers.

Assuming that consumers (or rather software agents that reside in their side) are rational self-interested entities that seek to optimize their expected utility in the sense specified above, different values of K and different modes of feedback $f(\cdot)$ incur different consumption strategies. Specifically, different values of K incentivize differently consumers to be included in the top- K list. On the other hand, M serves as disincentive to users for avoiding the bottom- M list. Of course the impact of a given set of (K, M) values to different consumers depends on consumer parameters h_i, a_i .

Given a serious game with parameters K, M , the total demand of consumers at a specific time slot is:

$$P(K, M) = \sum_{i=1}^N p_i^0 (1 - r_i(K, M)), \quad (4.9)$$

where $r_i(K, M)$ is the normalized consumption reduction of consumer i given serious-game parameters K, M . Thus, the problem faced by the serious-game designer is as follows:

$$\min_{K, M \in \{1, 2, \dots, N/2\}} C(P(K, M)). \quad (4.10)$$

The initial reference consumption p_i^0 of each player i is given to the serious-game designer by the utility company. Therefore, the game designer only needs to estimate the performance $r_i(K, M)$ of each player i for certain serious-game parameters K, M , in order to calculate $C(P(K, M))$.

4.3.2.1 Full information of consumer utility functions

If full information was available to the serious-game designer regarding consumer utility functions (e.g., $h_i, d_i(\cdot)$, etc.), then a closed-form solution for K, M could be found as the equilibrium of the two-stage Stackelberg game. Alternatively, a numerical solution to the equilibrium of the Stackelberg game can be found as follows: The serious-game designer finds optimal $r_i^*(K, M)$ values for each consumer i by solving (4.8) for a certain set of pairs of K, M parameters and replaces the $r_i^*(K, M)$ values to (4.10), in order to find the values of K, M that lead the system to Stackelberg equilibrium. Specifically, the expected utility function of the player in equation (4.7) can be re-written using equation (4.6) in a segmented form as follows:

$$\mathbb{E}\{u_i(r_i(K, M))\} = \begin{cases} h_i - d_i(r_i(K, M)) + qp_i^0 r_i(K, M), & r_i(K, M) \geq r_u(K) \\ -d_i(r_i(K, M)) + qp_i^0 r_i(K, M), & r_l(M) < r_i(K, M) < r_u(K) \\ -h_i - d_i(r_i(K, M)) + qp_i^0 r_i(K, M), & r_i(K, M) \leq r_l(M) \end{cases} \quad (4.11)$$

For convenience, denote the upper segment of equation (4.11) as $u_1(\cdot)$, the middle segment as $u_2(\cdot)$ and the lower segment as $u_3(\cdot)$. Toward maximizing the expected utility of the player, the first derivative of $\mathbb{E}\{u_i(r_i(K, M))\}$ with respect to $r_i(K, M)$ is given by:

$$(\mathbb{E}\{u_i(r_i(K, M))\})' = -2a_i r_i(K, M) + qp_i^0 \quad (4.12)$$

We identify three cases according to the sign of (4.12):

- (i) (4.12) is negative for all $r_i(K, M) \in [0, 1]$. Then, the expected utility of the player is maximized at $r_i^*(K, M) = \arg \max\{u_1(r_u(K)), u_2(r_l(M) + \varepsilon), u_3(0)\}$, where $0 < \varepsilon \ll 1$ an arbitrarily small positive number.
- (ii) (4.12) is positive for all $r_i(K, M) \in [0, 1]$. Then, the expected utility of the player is maximized at $r_i^*(K, M) = \arg \max\{u_1(1), u_2(r_u(K) - \varepsilon), u_3(r_l(M))\}$.
- (iii) (4.12) is 0 at $r_i(K, M) = x_0 := qp_i^0 / 2a_i$ and $x_0 \in [0, 1]$. Here, there are three subcases:
 - If $x_0 \leq r_l(M)$, then (4.12) is negative for $r_i(K, M) > r_l(M)$ and the expected utility of the player is maximized at $r_i^*(K, M) = \arg \max\{u_1(r_u(K)), u_2(r_l(M) + \varepsilon), u_3(x_0)\}$.
 - If $r_l(M) < x_0 < r_u(K)$, then (4.12) is positive for $r_i(K, M) \leq r_l(M)$ and negative for $r_i(K, M) \geq r_u(K)$. Therefore, the expected utility of the player is maximized at $r_i^*(K, M) = \arg \max\{u_1(r_u(K)), u_2(x_0), u_3(r_l(M))\}$.
 - If $x_0 \geq r_u(K)$, then (4.12) is positive for $r_i(K, M) < r_u(K)$ and the expected utility of the player is maximized at $r_i^*(K, M) = \arg \max\{u_1(x_0), u_2(r_u(K) - \varepsilon), u_3(r_l(M))\}$.

We replace the optimal $r_i^*(K, M)$ values for the N players in (4.9) and if the partial derivatives of $C(\cdot)$ in (4.10) with respect to K, M are defined, then closed-form solutions for optimal K^*, M^* can be analytically found. Observe, that the energy consumption behavior of players whose performance is maximized in values other than those involving $r_u(K), r_l(M)$ does not depend on K, M . According to the above, the complexity for finding optimal K^*, M^* parameters in the full-information case is either $O(N)$ when

closed-form solutions can be analytically derived for K^*, M^* or $O(N^3)$ when K^*, M^* are numerically approximated by evaluating $C(\cdot)$ over all (K, M) pairs.

4.3.2.2 Historical information about consumer actions

Since full information may not be readily available in some practical scenarios, the serious-game designer can aggregate historical data on the consumption performance of customers in the serious game for different values of K, M and employ them to extract personalized information about player behavior. Histograms like the ones depicted in Fig. 4.3 can be created by the game designer and the appropriate information can be provided to supplier so as to estimate $P(K, M)$ for the various (K, M) pairs.

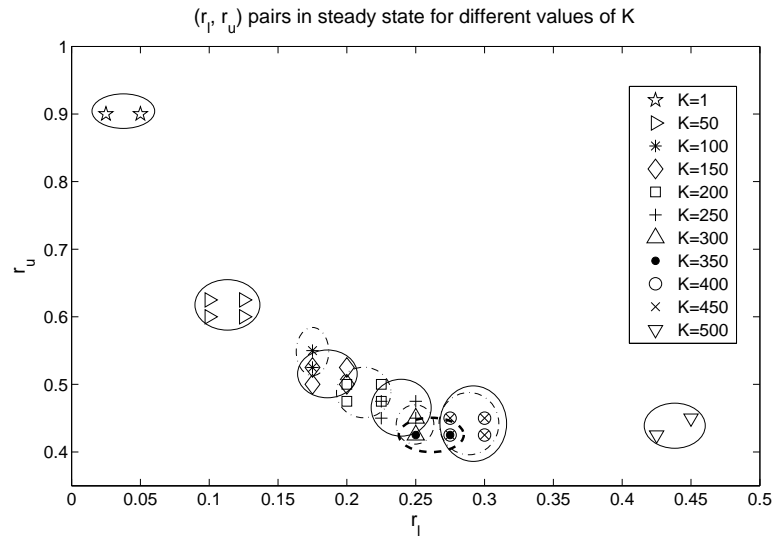
Given the 2D histogram for specific values of K, M , the operator can estimate the expected normalized consumption reduction $\mathbb{E}\{r(K, M)\}$ from any consumer as follows:

$$\mathbb{E}\{r(K, M)\} = \sum_{y \in \mathcal{Y}} y \Pr(r(K, M) = y), \quad (4.13)$$

where \mathcal{Y} is the set of values ranges (with width $\delta = 0.125$ in Fig. 4.3) of energy consumption reduction that are summarized in the histogram, and the probability $\Pr(r(K, M) = y)$ equals the relative number of consumers that had a normalized energy reduction $r(K, M) = y$ in the previous serious games with parameters K, M . Thereafter, the serious-game designer can calculate its expected total load $\mathbb{E}\{P(K, M)\}$ and the operational cost for the energy supplier as follows:

$$C(\mathbb{E}\{P(K, M)\}) = C\left(\sum_{i=1}^N p_i^0 (1 - \mathbb{E}\{r(K, M)\})\right). \quad (4.14)$$

In order to find the values of K, M that minimize the energy generation cost, the game designer should repetitively run serious games for multiple pairs K, M of parameters and build performance histograms similar to those of Fig. 4.3 for each pair of parameters. In theory, K, M can take values in $\{1, 2, \dots, N/2\}$. However, in practical settings, the maximum cardinalities of the upper and lower lists are expected to be much lower than $N/2$. Then, iterating among all pairs of K, M , the serious-game designer can find the one that gives the minimum operational cost according to equation (4.14). According to the above, the expected performance of any consumer is calculated in $O(1)$ and the minimum operational cost is found after $O(N^2)$ iterations. Thus, the complexity of the historical-information case is $O(N^2)$.

FIGURE 4.5: The steady (set of) states $\langle r_l, r_u \rangle$ for different $K=M$.

4.3.3 Equilibrium of the Stackelberg game arising from serious-game interactions

Given the empirical cumulative distribution function $F(\cdot)$ for each K, M pair, followers' game accepts one equilibrium and so does the overall Stackelberg game. This is because, according to the analysis of equation (4.11) (which is an approximation of equation (4.7) using (4.3)), for fixed r_u, r_l , each consumer has a unique utility maximizing strategy. Iterating through K, M pairs or by means of the partial derivatives of the operation cost with respect to K, M , the serious-game designer selects the pair (K^*, M^*) that minimizes the energy generation cost.

Regarding the followers' game, it is important to argue on why the cumulative performance statistics for each K, M pair will converge for repetitive instances of the serious game. Observe that players play according to equation (4.7) using (4.3), while the game designer uses equation (4.11) that results from equation (4.7) using the indicator function (4.6). Thus, the utility functions of the players are continuous, the strategy space is non-empty and compact, and the set of players is finite. So, the followers' stage of our Stackelberg game is a continuous game on its own. According to Glicksberg theorem [88], every continuous game has a mixed strategy Nash equilibrium. Therefore, the consumers' game will have a mixed strategy Nash equilibrium, in which each consumer i plays according to a mixed strategy regarding her performance r_i . It then follows that the empirical cumulative statistics will converge for every K, M pair. Note that until it converges, the empirical cumulative distribution function $F(\cdot)$ (and thus, $f(\cdot)$) of consumers is the sole feedback through which a consumer may adapt her best-response performance strategy toward the equilibrium point.

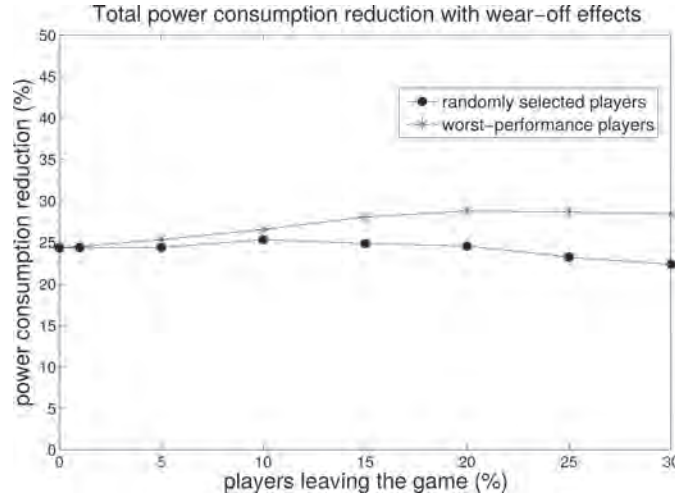


FIGURE 4.6: The percentage of operational-cost reduction achieved by our approach for different percentages of consumers leaving the game.

4.4 Simulation Results and Analysis

4.4.1 Experimental Setup

We consider $N=1000$ consumers, fixed power-unit price $q=0.124$ $\$/kWh$, fixed upper bound on the demand of each consumer $p_{max}=9$ kWh and $\varepsilon=10^{-4}$. For simplicity, we assume $K=M$. The baseline energy demand p_i^0 of each player i is assumed to follow the normal distribution $N(5, 1)$. Moreover, first, we assume that a_i follows $U(5, 40)$ and h_i follows $U(1, 4)$. At the beginning of each serious-game instance, each consumer is provided with the evolving probability distribution of consumer performance in the form of a histogram with bar width $\delta=0.025$. The performance thresholds r_u , r_l are calculated according to (4.4).

4.4.2 Steady-State Convergence and Wear-off Effects

For the aforementioned setup, we experimentally observed that for each pair of K, M our system passes through various transient states. Each state is characterized by a specific pair of thresholds (r_l, r_u) . After conducting a number of serious-game rounds for each pair of K, M , we noticed that our system reaches in a few tens of game rounds either a unique steady state or a steady set of states that alternate each other in a closed loop, as it was theoretically expected in Subsection 4.3.3. Fig. 4.5 shows the states, i.e., the pairs of thresholds, that appear for different (K, M) pairs. Employing a factor $\beta = 0.985$ to discount the significance of the past in the probability distribution of consumer performance, we experimentally found the convergence time to be around 40

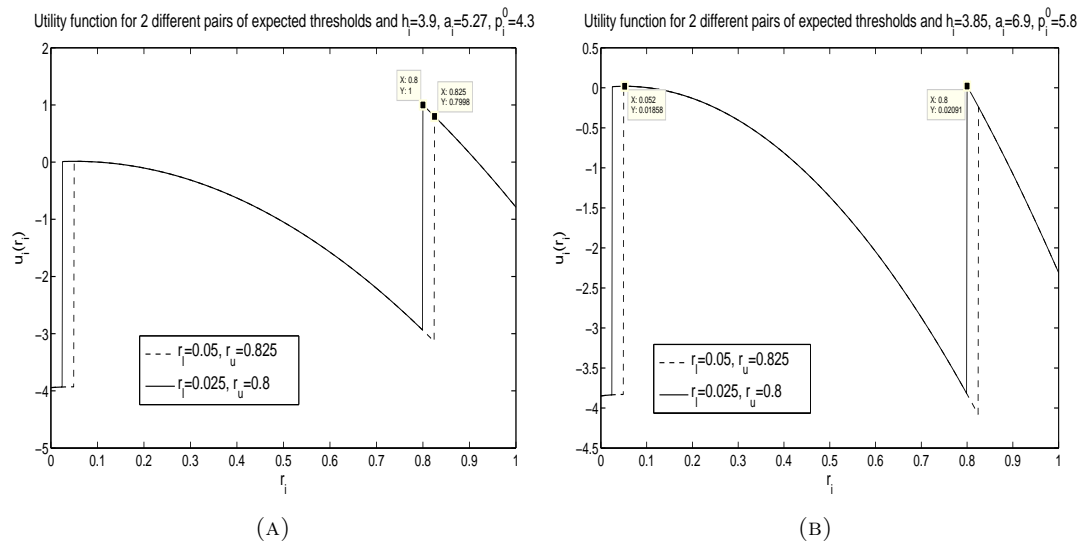


FIGURE 4.7: Utility functions of two players for $K = 4$ and two different states.

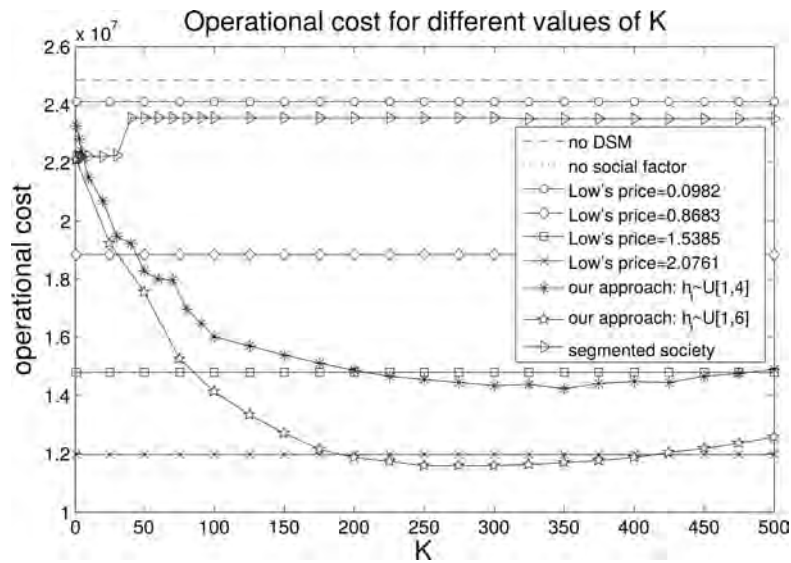
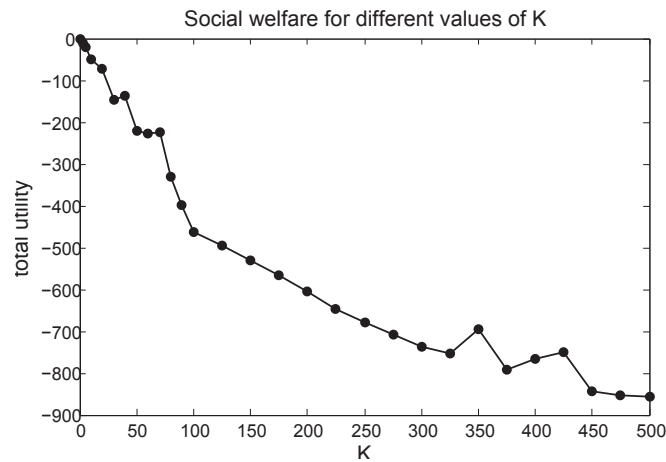


FIGURE 4.8: Operational cost for different values of $K=M$.

to 50 game rounds. For larger societies, the convergence time has been experimentally found to be shorter.

We also investigated the impact of gaming *wear-off*, i.e., players leaving the game, to the effectiveness of our approach. As depicted in Fig. 4.6 for $K=M=350$, the percentage reduction in the operational cost for remaining players achieved by our approach remains roughly the same when a percentage of random or worst-performing players leave the game.

FIGURE 4.9: Social welfare for different values of $K=M$.

4.4.3 Experimental Results

First, we show in Fig. 4.7 how the different values of h_i , a_i and p_i^0 affect the utility functions of two different consumers of the community and the values of their performance r_i that maximize these functions. We consider two different states $(r_{l_1}, r_{u_1})=(0.025, 0.8)$ and $(r_{l_2}, r_{u_2})=(0.05, 0.825)$. From Fig. 4.7a, we observe that for $h_i=3.9$, $a_i=5.27$ and $p_i^0=4.3$, the consumer's utility is maximized at $r_i=r_{u_1}=0.8$ and at $r_i=r_{u_2}=0.825$, respectively for the different states. Fig. 4.7b shows that for the consumer with parameters $h_i=3.85$, $a_i=6.9$ and $p_i^0=5.8$, the optimal utility results for $r_i=r_{u_1}=0.8$ and $r_i=x_0=0.052$, respectively for the different states.

Based on these results, we observe that when consumers desire to be included in \mathcal{U} due to their high sensitivity h_i to social impact, they choose the minimum value of r_i that enables them to be on \mathcal{U} , i.e., the threshold r_u . On the other hand, when they are relatively socially-indifferent (i.e., low h_i), then they choose the minimum r_i that enables them to stay out of \mathcal{L} , i.e., a value slightly larger than the threshold r_l . The selection of the minimum possible value stems from the desire of consumers to achieve social recognition with the minimum possible dissatisfaction cost. However, bigger values of a_i reduce the ability of consumers to be on \mathcal{U} or stay out of \mathcal{L} , since these a_i values discourage them from making substantial efforts. The baseline energy demand p_i^0 also influences each player's decision r_i , since it determines the benefit or the cost from energy consumption savings (equation (4.7)) and thus, each player's net payoff.

Fig. 4.8 depicts the average (over all states for each $K=M$) operational cost of our approach for different values of $K=M$ and for different communities of consumers, as compared to the dynamic-pricing approach of [89] ("Low") and to the no Demand Side Management approach ("no DSM"). For a community of consumers with $a_i \sim U(5, 40)$,

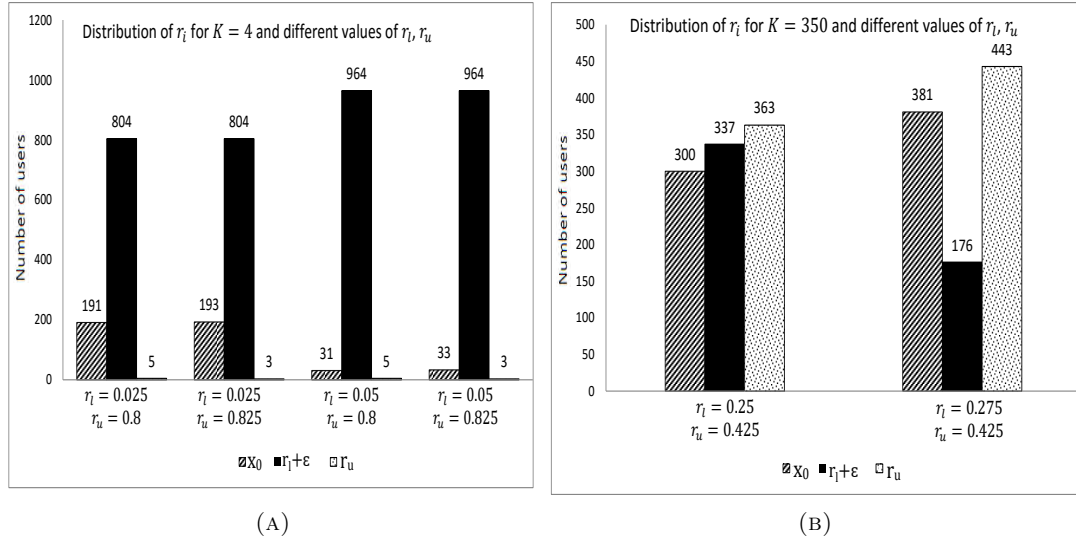


FIGURE 4.10: Distribution of r_i for different pairs (r_l, r_u) in the cases of a) $K = 4$ and b) $K = 350$.

$h_i \sim U(1, 4)$, the value of K that minimizes the operational cost is $K = M = 350$, and it becomes $K = M = 250$ when $h_i \sim U(1, 6)$. For a “segmented” society with two different communities of consumers, one of 30 consumers with $a_i \sim U(5, 40)$, $h_i \sim N(40, 2)$ and another of 970 consumers with $a_i \sim U(110, 115)$, $h_i \sim U(0.1, 0.3)$, the minimum generation cost arises for $K = 1$. Note that the first community is very socially-sensitive, while the second is expected to exert zero effort for any K . Therefore, K is selected in this case, so as to minimize the operational cost of the first community.

In general, the lower the mean sensitivity \bar{h}_i of consumers to social approval, the higher the value of K that must be selected by the serious-game designer in order to: a) considerably decrease the threshold r_u so as to be easier and more beneficial for these consumers to exert the effort to enter \mathcal{U} , and also, b) considerably increase the threshold r_l so as to force them make substantial effort to avoid entering \mathcal{L} .

When a community of consumers is socially-insensitive (“no social factor”), i.e., $h_i = 0$, then the choice of K makes no difference, while the operational cost is high. As compared to the dynamic-pricing DSM approach of [89], referred to as “Low”, we observe that the minimum operational cost achieved by our approach can be lower than that of [89] for low prices per energy unit. Thus, the preferable choice between our serious-game approach and that of [89] depends on the financial rationality, the elasticity of energy demand and the sensitivity of consumers to social approval. Fig. 4.9 shows the social welfare (total utility of players) for different K values for a consumer community with $a_i \sim U(5, 40)$, $h_i \sim U(1, 4)$.

Fig. 4.10 shows the distribution of performance r_i for the different pairs of thresholds that arise for $K = 4$ (Fig. 4.10a) and $K = 350$ (Fig. 4.10b). Only three of the seven possible values of r_i (described in Subsection 4.3.2) appear in our results: x_0 , $r_l + \varepsilon$ and r_u . For $K=4$, a slight increase in the value of r_u (for both $r_l = 0.025$ and $r_l = 0.05$) forces 2 out of the 5 players with performance $r_i = r_u$ make an effort smaller than r_u and namely, x_0 . Similarly, for $K = 350$, a slight increase in the value of r_l forces 161 out of the 337 players with performance $r_i = r_l + \varepsilon$ make an effort smaller (i.e., x_0) or greater (i.e., r_u) than $r_l + \varepsilon$.

Thus, we observe that for low values of K (Fig. 4.10a), a slight increase in the value of r_l increases the number of players that achieve $r_l + \varepsilon$ due to the fact that their net payoff for $r_i = r_l + \varepsilon$ is greater than the one for $r_i = x_0$. In contrast, for higher values of K (Fig. 4.10b), such an increase reduces the number of consumers with performance $r_l + \varepsilon$. Furthermore, we observe that the number of consumers that choose $r_i = r_u$ for $K = 4$ is much smaller than the one for $K = 350$. This stems from the fact that a low value of K makes it more difficult for consumers to achieve $r_i = r_u$.

4.5 Related Work

There have been some prior efforts to employ serious games for demand side management [77–80, 90, 91], albeit with no modeling or analysis on the serious-game design, as opposed to our work. In [79], a serious game for smart grids is organized as a virtual world with many user roles and actions, involving direct actions and training for sharing a Medium/Low Voltage transformer among prosumers. A serious game for energy conservation among students is described in [80]. The serious-game website and associated game mechanics are provided by the Makahiki system [92]. Similarly to our setting, no monetary rewards are included in the game; incentives are introduced through competition among consumers for points for energy conservation actions and for participation to online educational and real-world activities. According to [80], energy feedback systems should be actionable, include training and be time-persistent to have long-term effect into energy consumption behavior. Our serious-game model is time-persistent.

Also, the game “Energy Battle” [90], similarly to [80], aimed at encouraging occupants of student-households to save energy by means of competition. In [93], the authors review multiple energy competitions among university students and identify several pitfalls in their design. Specifically, the use of total energy consumption or (relative) energy-consumption reduction for winner determination is deemed as not adequate when static baseline calculation methods are employed and may be unfair for already “green” consumers.

An online game for improving home energy behavior, named Power House, is proposed in [78]. Its objective is to track activities and assist each member of a virtual family to save energy, while real-world energy behaviors produce particular in-game advantages and disadvantages. An online serious game (“EnerCities”) is presented in [77] to increase the environmental and energy-related awareness of secondary school students, and to influence their energy-related behaviors. Also, a virtual pet game designed for energy use reduction in a commercial office setting is presented in [91], where device-specific energy consumption is reflected in the fitness of virtual pets. There are also a number of studies on gamification in general [94], [95], which verify that specific serious-game design elements, such as leaderboards, points and levels, positively influence user participation, engagement and behavioral change.

In a different class of work, a number of game-theoretic dynamic-pricing schemes that involve interaction between the utility company and the consumers for energy-consumption smoothing have been proposed [89], [96]. However, [97] shows that dynamic pricing mechanisms can lead to peak-shifting when consumers rationally respond to price signals, unless specific strategies of bounded rationality are employed. In our paper, consumers take decisions based on social influence, as opposed to financial incentives.

Finally, prospect theory is employed in [32] for studying the problem of customer-owned energy storage management in the smart grid in a less rational manner, as opposed to the von Neumann-Morgenstern utility theorem employed here. In [32], a human player subjectively observes and makes her charging/discharging decisions based on the potential value of the benefit from selling energy and of the penalty from power regulation rather than the final outcome.

4.6 Conclusion

In this chapter, we introduced the problem of optimal serious-game design for reducing energy consumption and increasing user engagement. We considered a serious game, where a serious-game designer entity presents publicly to all consumers a list of top- K consumers and a list of bottom- M consumers according to their respective energy consumption reduction at peak hours. The driving forces of this serious game are the user discomfort due to demand load reduction, the user desire for social approval and the user sensitivity to social outcasting. We formulated the problems of the serious-game designer as an operational-cost minimization one for the utility company and that of each consumer as a utility maximization one. The serious-game design problem faced by the utility company is to decide on K, M and on the feedback provided to the consumers, while the consumer-side problem amounts to selecting the behavioral change to energy

consumption that maximizes the expected user utility. By a series of simulations, we showed how the choices of K, M affect the energy consumption reduction for different types of customers.

Chapter 5

Optimal Design of Energy Consumption Reduction Campaigns Through Behavioral Data-Driven Consumer Profiling

Contents

5.1	Introduction to Data-driven Energy Consumer Profiling and Behavior Prediction	89
5.1.1	Our contribution	91
5.2	Energy Consumption Reduction Campaign: A simple scenario	92
5.3	System Model	93
5.4	Incentive and load-reduction task allocation: Problem formulation	94
5.5	Consumer profiling	95
5.5.1	Logistic Regression model, and Optimal Task and Incentive Allocation Policy	96
5.5.2	Fast-and-Frugal Tree model, and Optimal Task and Incentive Allocation Policy	97
5.6	Numerical Example	103
5.7	Related work	107
5.8	Conclusion	110

5.1 Introduction to Data-driven Energy Consumer Profiling and Behavior Prediction

The penetration of new components in the traditional energy management system, such as smart meters, energy storage systems, bidirectional communication links between consumers and utility operators, and renewable energy sources, has led to a fundamental change in the energy consumers' role. In this new and smart energy grid, consumers can now influence their energy consumption profile, i.e., curtail or shift in time their energy demand loads, and they can also generate and store power. Moreover, the proliferation of smartphones has facilitated the participation of consumers in demand-response programs and energy consumption reduction campaigns. Such programs/campaigns aim at optimally matching energy demand and supply, and at reducing consumption in peak periods by making suggestions/offers to users on tasks to be fulfilled, along with offered incentives. A fundamental issue in their design and operation is the recruitment of users. Consumer participation is subject to various types of costs she/he may experience, such as discomfort from shifting energy demand tasks in time or from not performing them at all, or expenditure of time and effort on using a dedicated mobile app and performing a suggested task.

In literature, there is a great number of works on incentive mechanisms that try to motivate users to participate in demand-response programs and energy consumption curtailment efforts [83, 86, 98–102]. Most of them assume that consumers are rational decision-makers which aim at maximizing some notion of expected net utility and which select their actions by solving complex optimization problems. However, consumers are humans and their decisions are far from rational. Decisions regarding their participation in a campaign may depend on various *attributes* that characterize the demand load shifting/reduction suggestions, which are prioritized or weighed differently by different users (different consumer *preferences*).

In this work, instead of relying on conventional rational models for the consumer, we study the role of data in building behavior-based data-driven models for predicting consumer behavior. First, we build two different models to learn consumer individual preferences through historical data consisting of past load reduction suggestions and user decisions. The first model is based on a popular machine-learning technique while the second one on a cognitive psychological heuristic. The derived consumer preference models (profiles) are then used to predict the probability that a consumer will carry out specific consumption reduction tasks. Each load reduction suggestion made to a consumer comprises attributes that are present in the user decision-making model and which are either beyond or under the control of the campaign designer (CD). Then, the

ultimate goal of the CD is to determine appropriately the controllable attributes so as to accomplish the purpose of the campaign.

Our first data-driven approach to model consumer preferences and predict consumer behavior is based on Logistic Regression (LR). LR [103, 104] is a popular machine-learning framework for probabilistic classification. It is the most celebrated instance of generalized linear models. The binary logistic regression model is used to estimate the probability of a binary response based on one or more predictor variables (attributes). Important properties of LR are: (i) the objective for learning the model parameters is convex, so there are no local optima in optimization problems that involve this objective, and (ii) the model is probabilistic, hence it comes with well-calibrated estimates of uncertainty in the classification decision [103, 104].

It has been proved that there are small differences in performance (predictive accuracy) between simple mathematical models, known as simple decision heuristics, and computationally more complex methods developed in statistics, such as logistic regression, neural networks, categorization and regression trees (CART) and naive Bayes [105–107]. These heuristic models aim to a) capture the *bounded rationality* of decision-makers and the fact that decision-makers try to *satisfice* rather than optimize [108], and b) give insight into the actual decision-making mechanism. They need *less* time, information, and effort compared to more complex methods.

Our second approach is based on such a decision heuristic, namely the Fast-and-Frugal Tree (FFT) model [109]. FFTs are cognitive heuristic models which are suitable for classification tasks [105, 110]. They are decision trees with at least one exit leaf at every level and are composed of sequentially ordered attributes. This means that for every checked attribute, at least one of its branches can lead to a decision. It has been shown that they operate as lexicographic heuristics for categorization [110]. FFTs are constructed with binary attributes and a binary criterion but are generalized to other cases. They possess three advantages over other decision models: 1) fastness and frugality, 2) simple decision rules, and 3) potential robustness (less susceptible to overfitting). However, FFTs are not naturally probabilistic; they only predict the class that a given sample should belong to. Nevertheless, there exist a number of methods [111–113] that can generate calibrated probability estimates from decision trees and thus, from FFTs.

5.1.1 Our contribution

This work studies how different behavior-based data-driven models can be applied to model and predict consumer behavior. It also investigates the impact of consumer behavior uncertainty on the determination of the load reduction suggestions to be made to consumers, and on the resulting system performance. We rely on logistic regression from machine learning, and on fast-and-frugal trees inspired from cognitive psychology and behavioral science to build probabilistic models of consumer participation. Our main goal is to discover through these models the different factors that determine consumer actions and the different importance placed on them, and to predict consumer reaction to load reduction suggestions. We consider that a consumer views a suggestion made to her/him as a pair of two controllable attributes; namely, the offered monetary payment/reward and the suggested amount of load reduction. The CD aims to appropriately control/set these attributes in order to achieve a certain expected load saving at minimum economic cost.

Specifically, our main contributions are the following:

- We formulate the problem of minimizing the maximum potential total cost (i.e., the sum of offered payments) of the campaign designer that results from the allocation of load reduction tasks and monetary rewards (incentives) to consumers in order to achieve at least a certain expected load reduction.
- We build personalized decision-making profiles for consumers using a logistic-regression model and a fast-and-frugal tree model, based on past load reduction suggestions and the corresponding consumer responses. We show that in the case of LR the optimization problem turns out to be a sigmoid optimization one, while in the case of FFTs it is an integer linear programming one.
- We present a numerical example for a system of two consumers that highlights the need for optimal allocation of tasks and incentives (i.e., optimal control of suggestion attributes), and that shows how the selection of the attribute values affects the consumer offer acceptance probabilities and thus, the expected total energy consumption reduction.

The rest of the chapter is organized as follows. In section 5.2, we describe the simple scenario studied in this work. In section 5.3, we analyze our system model, while in Section 5.4, the CD optimization problem is formulated. In section 5.5, we present the two behavior-based data-driven approaches for consumer profiling, and the form of the optimization problem in each approach. Section 5.6 includes a representative numerical

example that concerns a simple special case of the problem. In section 5.7, we present the related literature. Finally, section 5.8 concludes the work. This chapter is based on work [114].

5.2 Energy Consumption Reduction Campaign: A simple scenario

We consider a peak-load reduction campaign that asks consumers to reduce their baseline energy demand in exchange for some payment. We assume that there is an expected overload at time t that the campaign designer (CD) aims to avoid. The campaign mobile app makes suggestions of the following form: (a) “Please reduce your demand load X at time slot t by d kWh for a payment $\$p$ ”, or (b) “If you reduce your demand load X at time slot t by $d\%$, you will save $\$p$ in your bill and you will reduce your carbon footprint by $z\%$ ”.

Let us assume that a consumer decides whether to curtail a demand load X based on the following parameters (attributes) that characterize a suggestion:

- The payment, $\$p$
- The amount of load reduction, d (in kWh)

in the case of example recommendation (a), and

- The payment, $\$p$
- The percentage load reduction, d (%)
- The percentage reduction of the consumer’s carbon footprint, z (%)

in the case of example suggestion (b). Here, we focus on suggestions of the form (a). The above attributes are prioritized or weighted differently by different consumers. Also, each suggestion may comprise attributes that are beyond the control of the CD and attributes that are under its control, such as the incentives to be designed.

The first goal of the CD is to learn consumer individual preferences, i.e., derive user profiles that capture the different importance that each user places on the different attributes, through records of past load reduction suggestions and user decisions. Then, its second goal is to determine appropriately the controllable attributes of each suggestion, based on the derived decision-making profiles, so as to achieve an expected total load reduction of *at least* Q units of energy with the minimum possible total offered compensation).

5.3 System Model

Denote by \mathcal{I} the set of energy consumers that participate in the peak-load reduction campaign. Time is divided into slots. Here, as we mentioned above, we focus on a *single* time slot of the day at which the CD needs to achieve an *expected* total load reduction greater than or equal to Q . The CD interacts with consumers through a mobile app. It suggests a specific task for load curtailment to each consumer offering at the same time a monetary reward as an incentive to carry out the task, i.e., it makes a single load reduction suggestion \mathbf{r}_i to each consumer $i \in \mathcal{I}$. Each suggestion \mathbf{r} can be viewed as a set of n attributes (interchangeably called cues hereafter) $\mathbf{r} = (c_1, c_2, \dots, c_n)$. These attributes include the monetary reward $c_1 = p$, the amount of load to be curtailed $c_2 = d$, and maybe the description of the corresponding energy demand task(s) to be canceled c_4, c_5, c_6, \dots .

The attributes p, d can be manipulated by the campaign designer in order to incentivize consumers to undertake specific load reduction tasks. The attributes that are related to the energy demand tasks to be cancelled, such as the time when the demand task(s) is(are) scheduled to start by the consumer (during the target peak slot), the appliance(s) involved, and the power consumption level of the appliance(s), are inherent ones to the energy demand task(s). We assume that the time of the day or the type of load does not determine whether the user will do the reduction or not (we are talking about homogeneous load for all users). Therefore, it is safe to assume that each consumer $i \in \mathcal{I}$ views a suggestion \mathbf{r}_i only as a pair of attributes p_i, d_i , i.e., $\mathbf{r}_i = (p_i, d_i)$. There may exist other attributes besides the two above that determine the decision of the consumer. Some might be related to the specific load reduction task, while some others might be related to the consumer per se, e.g., the social recognition consumers get about their prudent energy consumption or the intrinsic motivation to be environment friendly.

Without loss of generality, attributes p, d are taken to be continuous variables. Let $p_{min} \geq 0$ and $p_{max} > 0$ be the minimum and maximum allowed payment that can be offered to a consumer i , respectively, and $d_i^{max} \geq 0$ be the maximum possible load reduction of each consumer i . If the baseline energy demand of consumer i at the target slot t is e_i^0 , with $0 \leq e_i^0 \leq e_{max}^0$ then $d_i^{max} = e_i^0$. We denote as $d'_i = d_i/e_i^0$ the normalized demand load reduction, with $0 \leq d'_i \leq 1$. We observe that d'_i can also be expressed as a percentage, i.e., $\tilde{d}_i = 100 d'_i \%$.

Consumers decide to undertake a demand load curtailment, and they do so probabilistically, depending on how well suggestions compare with their own individual preferences. This probabilistic reaction of consumers depends on the suggested load reduction tasks and the corresponding offered payments, and on the consumer profiles that define their

preferences. Consumer preferences are expressed over the attributes p, d , which are controllable by the campaign designer. Hence, each consumer is faced with a two-attribute decision problem.

We approach consumer reaction to a load reduction suggestion as a two-class probabilistic classification problem. The possible outcomes of the consumer decision-making process determine the two classes. Specifically, when a suggestion \mathbf{r}_i is made to a consumer i , class 1 (C_1) corresponds to accepting it and performing the suggested load curtailment, while class 0 (C_0) corresponds to rejecting it. Thus, an instance $\mathbf{r}_i = (p_i, d_i)$ is assumed to belong either to class C_1 or to class C_0 . For a given vector of values $\mathbf{r} = (p, d)$ for the different suggestion attributes, the probability that consumer i performs the suggested load reduction task, i.e., accepts the offer \mathbf{r} , is $P_a^i(\mathbf{r}) = P_i(C_1|\mathbf{r}) = P_i(C_1|(p, d))$. This offer acceptance probability is not assumed to be a priori known, but it will emerge from our subsequent analysis.

In our scenario, each consumer decides whether to undertake a load reduction task with a specific fixed reward and if s/he does, s/he deterministically receives the offered payment (decision-making under certainty). We assume for now that whether a consumer will perform a suggested task is independent of all other consumers' decisions. Also, the outcomes (i.e., the received reward) for a decision-maker depend only on her/his own choices/actions.

5.4 Incentive and load-reduction task allocation: Problem formulation

Given a set of users \mathcal{I} , each user $i \in \mathcal{I}$ with a given preference/decision-making profile, the CD has to find the most economically efficient payment and load-reduction task allocation scheme that achieves an expected total load reduction of *at least* Q units of energy. I.e., the CD seeks to optimally control the suggestion attributes p_i, d_i so as to achieve an expected load saving greater than or equal to Q with the minimum possible total offered reward (i.e., sum of offered payments). Namely, the optimization problem faced by the CD is:

$$\min_{\mathbf{p}, \mathbf{d}} \sum_{i \in \mathcal{I}} p_i \quad (5.1)$$

$$s.t. \quad \sum_{i \in \mathcal{I}} d_i P_a^i(p_i, d_i) \geq Q \quad (5.2)$$

$$p_{min} < p_i \leq p_{max}, \quad \forall i \in \mathcal{I} \quad (5.3)$$

$$0 \leq d_i \leq d_i^{max}, \quad \forall i \in \mathcal{I}, \quad (5.4)$$

where $\mathbf{p} = (p_1, p_2, \dots, p_{|\mathcal{I}|})$ is the vector of the offered payments and $\mathbf{d} = (d_1, d_2, \dots, d_{|\mathcal{I}|})$ is the vector that contains the suggested amounts of load reduction. The objective function denotes the sum of offered payments, namely the maximum potential cost of the CD which is the cost that results from the acceptance of all recommended offers. Introducing the concept of normalized load reduction into the above formulation, the optimization problem becomes:

$$\min_{\mathbf{p}, \mathbf{d}'} \sum_{i \in \mathcal{I}} p_i \quad (5.5)$$

$$s.t. \quad \sum_{i \in \mathcal{I}} d'_i P_a^i(p_i, d'_i) \geq Q' \quad (5.6)$$

$$p_{min} < p_i \leq p_{max}, \quad \forall i \in \mathcal{I} \quad (5.7)$$

$$0 \leq d'_i \leq 1, \quad \forall i \in \mathcal{I}. \quad (5.8)$$

In order to solve the optimization problem (5.1)-(5.4) or the modified problem (5.5)-(5.8), we need to predict the probabilities $P_a^i(p_i, d_i)$ or $P_a^i(p_i, d'_i)$, $\forall i \in \mathcal{I}$, respectively.

5.5 Consumer profiling

As we mentioned above, our main goal is to predict the probability $P_a^i(\mathbf{r}_i)$ that a consumer i will accept an offer \mathbf{r}_i (i.e., will perform a suggested load reduction task), so as to solve the optimization problem (5.1)-(5.4) or (5.5)-(5.8) defined in section 5.4. In order to achieve our goal, we need to profile consumers. Namely, we need to build a model for each consumer i , using historical data, that describes the way s/he assesses a suggestion \mathbf{r}_i to reach a decision regarding her/his participation/action in the campaign. Such a decision depends on the suggestion attributes p_i, d_i (or d'_i or \tilde{d}_i), which are prioritized or weighed differently by different consumers.

Here, we employ two different techniques to learn consumer individual preferences (profiles) through records of past load reduction suggestions and user decisions: a) the Logistic Regression (LR) model and b) the Fast-and-Frugal Tree (FFT) model. The derived consumer profiles capture the different importance that each consumer places on the different attributes that characterize the suggestions. We then use these personalized consumer preference models to predict the probability $P_a^i(\mathbf{r}_i)$, $\forall i \in \mathcal{I}$, and subsequently solve the optimization problem (5.5)-(5.8).

5.5.1 Logistic Regression model, and Optimal Task and Incentive Allocation Policy

In this section, we model the way each consumer i assesses a suggestion \mathbf{r}_i to reach a decision (accept (C_1), or reject (C_0)) using the LR model. Given a vector of values $\mathbf{r} = (p, d')$ for the different attributes, LR predicts the probability that consumer i will make the positive choice (accept (C_1)) as follows:

$$P_a^i(\mathbf{r}) = P_i(C_1|\mathbf{r}) = \frac{1}{1 + e^{-\mathbf{w}_i \cdot \mathbf{r}}} = \sigma(\mathbf{w}_i \cdot \mathbf{r}), \quad (5.9)$$

where $\mathbf{w}_i = (w_{i,0}, w_{i,p}, w_{i,d'})$ is the vector of the attribute weights (logistic regression coefficients) for consumer i and $\sigma(x) = \frac{e^x}{e^x + 1} = (1 + e^{-x})^{-1}$ (with $\sigma(x) \in (0, 1), \forall x$) is the standard logistic (or sigmoid) function. $w_{i,0}$ is the intercept coefficient, while $w_{i,p}$ and $w_{i,d'}$ are the weights assigned by consumer i to attributes p_i and d'_i , respectively. An additional pseudo-attribute $c_0 = 1$, corresponding to the intercept term $w_{i,0}$, is added to the suggestion \mathbf{r} , i.e., $\mathbf{r} = (c_0, p, d') = (1, p, d')$. Also, $\mathbf{w}_i \cdot \mathbf{r}$ is the dot product of vectors \mathbf{w}_i and \mathbf{r} . We observe that the LR model predicts the above probability based on a weighted integration of the available attribute values p, d' . On the other hand, the alternative choice (reject (C_0)) is selected with probability $P_i(C_0|\mathbf{r}) = 1 - P_i(C_1|\mathbf{r}) = 1 - P_a^i(\mathbf{r})$.

For the prediction of the probability $P_a^i(\mathbf{r})$, the CD needs to learn the weight vector (*profile*) $\{\mathbf{w}_i\}, i \in \mathcal{I}$, which captures the importance that consumer i places on the different attributes in reaching a decision. These weights are learned from historical data (i.e., past load reduction suggestions and the corresponding responses) through a supervised learning process, whereby the individual consumer LR models are trained using these data (training datasets). The supervised learning process followed in the case of LR models is analytically described in Appendix A.

5.5.1.1 Optimal allocation of incentives and load reduction tasks in the case of LR: A sigmoid optimization problem

In section 5.4, we mentioned that the goal of the CD is to optimally control the suggestion attributes p_i, d'_i so as to achieve an expected load reduction greater than or equal to Q (units of energy) with the minimum possible total offered reward. Also, we showed above (Eq. (5.9)) that in the case of LR, the probability $P_a^i(\mathbf{r}_i)$ that a consumer i will undertake a suggested load reduction task is a sigmoid function of attributes p_i, d'_i . So, using Eq. (5.9), the optimization problem (5.5)-(5.8) of section 5.4 becomes:

$$\min_{\mathbf{p}, \mathbf{d}'} \sum_{i \in \mathcal{I}} p_i \quad (5.10)$$

$$s.t. \quad \sum_{i \in \mathcal{I}} d'_i \sigma(w_{i,0} + w_{i,p} p_i + w_{i,d'} d'_i) \geq Q' \quad (5.11)$$

$$p_{min} < p_i \leq p_{max}, \quad \forall i \in \mathcal{I} \quad (5.12)$$

$$0 \leq d'_i \leq 1, \quad \forall i \in \mathcal{I}. \quad (5.13)$$

Since both the objective function (5.10) and the function of constraint (5.11) have continuous first partial derivatives, the method of Lagrange multipliers can be used to solve the above problem. Namely, using a Lagrange multiplier $\lambda \geq 0$ for the constraint (5.11) (the other two are discarded and are introduced later through a constraint on the corresponding payment and task allocation policy of the CD), the Lagrange function is defined by

$$L(\mathbf{p}, \mathbf{d}', \lambda) = \sum_{i \in \mathcal{I}} p_i - \lambda \left(\sum_{i \in \mathcal{I}} d'_i \sigma(w_{i,0} + w_{i,p} p_i + w_{i,d'} d'_i) - Q' \right). \quad (5.14)$$

The attribute weights \mathbf{w}_i , $\forall i \in \mathcal{I}$ are obtained through the training process described in Appendix A. The Lagrangian (5.14) is the new objective function that needs to be minimized. We observe that the optimization problem turns out to be a *sigmoid optimization* one. The goal is to find some $(\mathbf{p}^*, \mathbf{d}'^*) \in \mathcal{C}$ and $\lambda^* \geq 0$ such that $L(\mathbf{p}^*, \mathbf{d}'^*, \lambda^*) = \min\{L(\mathbf{p}, \mathbf{d}', \lambda) : (\mathbf{p}, \mathbf{d}') \in \mathcal{C}, \lambda \in \mathbb{R}^+\}$, where $\mathcal{C} = \{(\mathbf{p}, \mathbf{d}') \in \mathbb{R}^{2|\mathcal{I}|} \mid p_{min} \leq p_i \leq p_{max}, 0 \leq d'_i \leq 1, \forall i \in \mathcal{I}\}$ is a closed convex set which determines the constrained feasible solution space for all controllable attributes p_i, d'_i , for $i \in \mathcal{I}$. The solution $(\mathbf{p}^*, \mathbf{d}'^*)$ of the above problem is the optimal payment and task allocation scheme in the case of LR models.

The problem (5.10)-(5.13) is a nonlinear programming (NLP) one. There are commercially available NLP solvers that employ Lagrange multiplier algorithms, such as the *interior-point* method or the *sequential quadratic programming* method [115], for solving constrained nonlinear programming problems. In section 5.6, we use such as a solver in order to obtain our numerical results.

5.5.2 Fast-and-Frugal Tree model, and Optimal Task and Incentive Allocation Policy

In this section, we build a behavioral decision-making profile for each consumer i using the FFT model. A tree such as the one in Fig. 5.1a is called a fast-and-frugal tree, according to the definition we presented above (Section 5.1). This is a typical FFT that models the following decision-making process: If $d'_i \leq T_{d'_i}$ then consumer i decides to accept the offer \mathbf{r}_i (i.e., \mathbf{r}_i belongs to class C_1). Else, if $d'_i > T_{d'_i}$ then the consumer considers p_i ; if $p_i \geq T_{p_i}$ then the consumer decides to accept the offer \mathbf{r}_i , else if $p_i < T_{p_i}$ then the consumer rejects \mathbf{r}_i (i.e., \mathbf{r}_i belongs to class C_0). T_{p_i} and $T_{d'_i}$ are appropriately

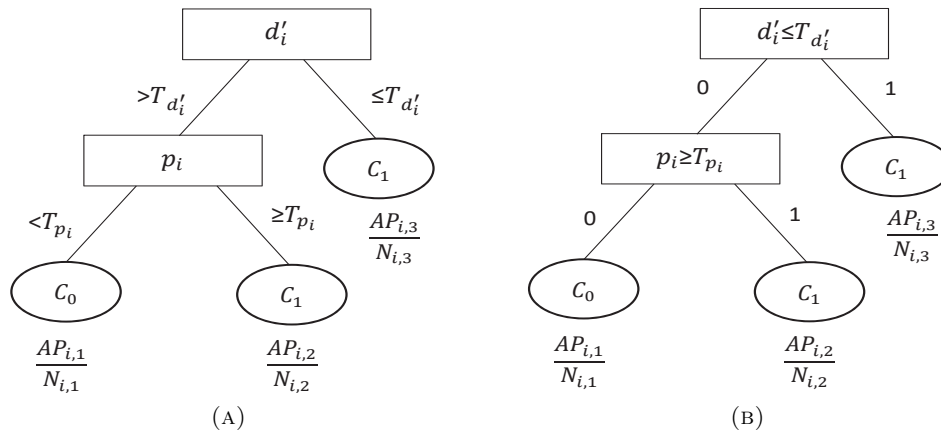


FIGURE 5.1: (a) An example FFT and (b) its alternative form for a consumer i that prioritizes the cue d'_i of the suggestion \mathbf{r}_i over the cue p_i . $AP_{i,l}/N_{i,l}$ is the estimate of the probability that an instance (i.e., suggestion) assigned to leaf l belongs to class C_1 . $AP_{i,l}$ is the number of training instances of (actual) class C_1 at leaf l , and $N_{i,l}$ is the total number of training instances at that leaf.

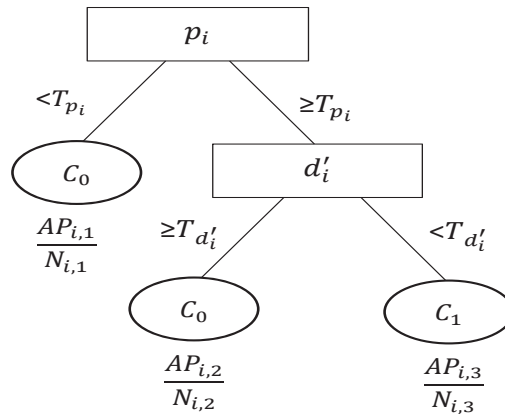


FIGURE 5.2: An example FFT that prioritizes the cue p_i of the suggestion \mathbf{r}_i over the cue d'_i . Its cue ranking and exit location are different from those of the FFT in Fig. 5.1a.

selected decision/classification thresholds. The structure of the tree implies that consumer i ranks the attribute d_i first and the attribute p_i second, i.e., s/he prioritizes d'_i over p_i . A different ranking would lead to a different tree. Fig. 5.2 depicts an example FFT with different cue order and different exit structure that may arise in our scenario.

Fig. 5.1b depicts the alternative form of the FFT in Fig. 5.1a; we apply the following convention here: each right branch is associated with the fulfillment of the condition at the top the branch and is labeled with 1, whereas each left branch is labeled with 0. Also, if a leaf stems from a branch labeled 1, then the decision at this leaf will be positive (i.e., its associated class is C_1), otherwise the decision will be negative (C_0).

In order to build (i.e., train) a FFT model for each consumer i , the CD needs to (i) find appropriate decision thresholds T_{p_i} , $T_{d'_i}$ for the available cues p_i , d'_i (or p_i , \tilde{d}_i), respectively, (ii) order the cues, and (iii) find an appropriate tree shape (exit structure). These three problems are solved using historical information (i.e., a training dataset for each consumer i consisting of past load-reduction suggestions and consumer responses) and a set of simple rules that are based on the optimization of a particular model evaluation measure. The training process followed in the case of FFTs is analytically described in Appendix B. After the tree has been constructed, its use for classification involves only traversing the tree by answering one simple question (i.e., checking a condition) at each level, and assigning the class label associated with the final node (leaf).

Although decision-tree models are not naturally probabilistic, they can predict probabilities, using information extracted from the classification of the training instances [110–112], as follows: The full training dataset sits at the top of the constructed tree and each training observation/instance traverses the tree (checking the condition at each node) until it reaches a leaf node. Each leaf comprises a subset of the full set of training instances. Then, each new observation that needs to be classified goes through the tree to reach a leaf node and a class-membership probability is obtained for this observation from the (training) relative frequency of the class in question at that leaf.

For example, if a new observation reaches a leaf node l that comprises 7 training instances of (actual) class C_1 and 3 training instances of (actual) class C_0 , then a probability $7/10 = 0.7$ would be assigned to be from class C_1 . Namely, the probability that an instance at leaf l belongs to class C_1 is $P(C_1|l) = 7/10$. Let AP_l be the number of training instances of (actual) class C_1 at leaf l and N_l be the total number of training instances at that leaf. Then, $P(C_1|l) = AP_l/N_l$. Since each leaf of the decision tree corresponds to a region of the two-dimensional feature space, the above probability estimation procedure assigns the same probability estimate to all points in the region.

It has been proved that this methodology provides poor probability estimates, since the decision tree was constructed using these same training instances and so the estimates tend to be too extreme, i.e., close to 0 or 1. One way of improving them is to make them smoother, i.e., to adjust them to be less extreme [111–113]. The Laplace correction method presented in [111], which has become a de facto standard for practitioners, can be used for smoothing these frequency-based probability estimates. This method replaces the conditional probability estimate $P(C_1|l) = AP_l/N_l$ by $P'(C_1|l) = (AP_l+1)/(N_l+C)$, where C is the number of classes. This simple, common smoothing method has been used for the production of improved probability estimates from decision-tree models like

CART and C4.5 [111, 113]. Since these models are more complex versions of FFTs ¹, the Laplace correction can also be used to improve the class-membership probability estimates given by FFTs ².

Applying the Laplace correction method to the frequency-based probability estimates derived from the FFT model built for a consumer i , the probability $P_a^i(\mathbf{r}_i)$ that the consumer will accept an offer \mathbf{r}_i (i.e., that instance \mathbf{r}_i belongs to class C_1) is given by:

$$P_a^i(\mathbf{r}_i) = P_i(C_1|\mathbf{r}_i) = \sum_{l=1}^{\Lambda_i} P_i(l|\mathbf{r}_i)P'_i(C_1|l) = \sum_{l=1}^{\Lambda_i} P_i(l|\mathbf{r}_i) \frac{AP_{i,l} + 1}{N_{i,l} + 2}, \quad (5.15)$$

where Λ_i is the number of leaves of the FFT, $P_i(l|\mathbf{r}_i)$ is the probability that leaf l is reached by instance (suggestion) \mathbf{r}_i , and $P'_i(C_1|l)$ is the smoothed/improved estimate of the probability that an instance assigned to l belongs to class C_1 . $AP_{i,l}$ is the number of training instances of (actual) class C_1 at leaf l and $N_{i,l}$ is the total number of training instances at that leaf. In our two-class problem, the Laplace correction method replaces the probability estimates $P_i(C_1|l) = AP_{i,l}/N_{i,l}$, for $l = 1, \dots, \Lambda_i$, obtained from the FFT of consumer i by $P'_i(C_1|l) = (AP_{i,l} + 1)/(N_{i,l} + 2)$.

In order to predict the probability $P_a^i(\mathbf{r}_i)$ of Eq. (5.15), the CD needs to find a) the probabilities $P_i(l|\mathbf{r}_i)$ and b) the values of $AP_{i,l}$, $N_{i,l}$, $\forall l = 1, \dots, \Lambda_i$. Based on the constructed FFT of a consumer i , a) each probability $P_i(l|\mathbf{r}_i)$ can be expressed as an indicator function of the conditions (tests) leading to leaf l , and b) the values of $AP_{i,l}$, $N_{i,l}$ can be estimated after the classification of all training instances using the tree (as described above). For example, from the FFT of Fig. 5.1a, it is:

$$P_a^i(\mathbf{r}_i) = \mathbf{1}_{\{d'_i > T_{d'_i}\}} \mathbf{1}_{\{p_i < T_{p_i}\}} \left(\frac{AP_{i,1} + 1}{N_{i,1} + 2} \right) + \mathbf{1}_{\{d'_i > T_{d'_i}\}} \mathbf{1}_{\{p_i \geq T_{p_i}\}} \left(\frac{AP_{i,2} + 1}{N_{i,2} + 2} \right) + \mathbf{1}_{\{d'_i \leq T_{d'_i}\}} \left(\frac{AP_{i,3} + 1}{N_{i,3} + 2} \right), \quad (5.16)$$

where $\mathbf{1}_A$ is an indicator function of the event A with $\mathbf{1}_A=1$ or 0 if event A is true or not, respectively. Here, we assume that leaves are numbered from left to right. We observe that in the case of FFTs the probability $P_a^i(\mathbf{r}_i)$ is a step (piecewise constant) function of attributes p_i , d'_i . Also, for a given suggestion \mathbf{r} , only one of the Λ_i terms of equation (5.15) is nonzero.

¹In the case of CART and C4.5, a) cues are looked up sequentially but not necessarily one at a time, b) the decision on which cues to use at each level and how to order them is made by statistical tests, and c) it is not necessarily the case that a decision can be made at each level of the tree [105, 116].

²Here, we assume that the available datasets for the consumers are not highly imbalanced. In the case of highly imbalanced datasets (i.e., the two classes are far from equiprobable), the m -estimation method presented in [112] might be more appropriate for the production of well-calibrated probability estimates.

5.5.2.1 Optimal allocation of incentives and load reduction tasks in the case of FFTs: An integer linear programming (ILP) problem

As we mentioned in section 5.4, the objective of the CD is to optimally control the attributes p_i , d'_i so as to achieve an expected total load reduction of at least Q' units of energy with the minimum possible total offered reward. Also, from the structure of a general FFT for a consumer i , we observe that each path from the root to a leaf l is characterized by *a*) a set of conditions/tests (one at each branch), *b*) the sets of values for the decision variables p_i , d'_i that satisfy the conditions at its branches, and *c*) a *fixed* probability estimate $(AP_{i,l} + 1)/(N_{i,l} + 2)$ that denotes the probability that consumer i will accept an offer \mathbf{r}_i that corresponds to this path. According to this observation and the objective of the CD, we identify the following cases:

- A path passes through a branch labeled $p_i < T_{p_i}$ or $p_i \leq T_{p_i}$. Then, the payment to be offered to consumer i should equal p_{min} (i.e., the element of the set of values $[p_{min}, T_{p_i})$ or $[p_{min}, T_{p_i}]$, respectively, that minimizes the total offered reward (i.e., the maximum potential cost of the CD).
- A path passes through a branch labeled $p_i > T_{p_i}$. Then, the offered payment to consumer i should be $T_{p_i} + \epsilon_p$ (i.e., one of the smallest elements of the interval $(T_{p_i}, p_{max}]$), where $0 < \epsilon_p < 1$ is an arbitrarily small positive number.
- A path passes through a branch labeled $p_i \geq T_{p_i}$. Then, the offered payment to consumer i should equal T_{p_i} (i.e., the minimum element of the interval $[T_{p_i}, p_{max}]$).
- A path passes through a branch labeled $d'_i > T_{d'_i}$ or $d'_i \geq T_{d'_i}$. Then, the normalized load reduction to be suggested to consumer i should be $d_i'^{max} = 1$ (i.e., the element of the set of values $(T_{d'_i}, 1]$ or $[T_{d'_i}, 1]$, respectively, that maximizes the expected load reduction of consumer i).
- A path passes through a branch labeled $d'_i < T_{d'_i}$. Then, the normalized load reduction to be suggested to consumer i should equal $T_{d'_i} - \epsilon_d$ (i.e., one of the largest elements of the interval $[0, T_{d'_i})$), where $0 < \epsilon_d < 1$.
- A path passes through a branch labeled $d'_i \leq T_{d'_i}$. Then, the normalized load reduction to be suggested to consumer i should be $T_{d'_i}$ (i.e., the maximum element of the interval $[0, T_{d'_i}]$).
- A path passes *only* through a *single* branch labeled $p_i < T_{p_i}$ or $p_i > T_{p_i}$ or $p_i \leq T_{p_i}$ or $p_i \geq T_{p_i}$. Then, the payment to be offered is determined as detailed above and the normalized load reduction to be suggested should be $d_i'^{max} = 1$ (i.e., the maximum element of the interval $[0, 1]$).

- A path passes *only* through a *single* branch labeled $d'_i < T_{d'_i}$ or $d'_i > T_{d'_i}$ or $d'_i \leq T_{d'_i}$ or $d'_i \geq T_{d'_i}$. Then, the normalized load reduction to be suggested is determined as detailed above and the payment to be offered to consumer i should be p^{min} (i.e., the minimum element of the interval $[p_{min}, p_{max}]$).

For example, there are three distinct paths in the FFT of Fig. 5.1a. For the path that passes through branches $d'_i > T_{d'_i}$ and $p_i < T_{p_i}$, the offer \mathbf{r}_i should be $(p_{min}, d_i^{max}) = (p_{min}, 1)$ and the estimated probability that consumer i will accept it is $(AP_{i,1} + 1)/(N_{i,1} + 2)$. For the path that passes through branches $d'_i > T_{d'_i}$ and $p_i \geq T_{p_i}$, the offer should be $\mathbf{r}_i = (T_{p_i}, 1)$ and it would be accepted with probability $(AP_{i,2} + 1)/(N_{i,2} + 2)$, whereas for the last path, it should be $\mathbf{r}_i = (p_{min}, T_{d'_i})$ with acceptance probability $(AP_{i,3} + 1)/(N_{i,3} + 2)$.

Therefore, the CD has a finite (countable) number of options (extracted from the FFT of a consumer i) for a suggestion \mathbf{r}_i , each of which corresponds to a path and has given p, d' values. Also, each path is associated with a probability estimate that shows the likelihood that consumer i will accept the option/offer that corresponds to this path. So, each of the Λ_i options is a predetermined optimal offer with a fixed acceptance probability, corresponding to a certain path l (the number of paths is equal to the number of leaves) of the FFT of consumer i , and is denoted as $\mathbf{r}_{il}^* = (p_{il}^*, d_{il}^*)$, for $l = 1, \dots, \Lambda_i$.

From the above observations, the optimization problem (5.5)-(5.8) of section 5.4 in the case of FFTs turns out to be a discrete optimization one and it can be reformulated as a variant of the standard knapsack problem [117]. In the following reformulated problem, the objective function represents the total offered reward (by the CD) that results from the selection of only one option (from the group of options) for each consumer i , and the first constraint shows that the expected total load reduction that results from the aforementioned selection of options needs to be greater than or equal to Q' (units of energy).

$$\min_{\mathbf{x}} \sum_{i=1}^{|\mathcal{I}|} \sum_{l \in \mathcal{M}_i} v_{il} x_{il} \quad (5.17)$$

$$s.t. \quad \sum_{i=1}^{|\mathcal{I}|} \sum_{l \in \mathcal{M}_i} z_{il} x_{il} \geq Q' \quad (5.18)$$

$$\sum_{l \in \mathcal{M}_i} x_{il} = 1, \quad \forall 1 \leq i \leq |\mathcal{I}| \quad (5.19)$$

$$x_{il} \in \{0, 1\}, \quad \forall 1 \leq i \leq |\mathcal{I}|, \forall l \in \mathcal{M}_i. \quad (5.20)$$

In this formulation, $\mathcal{M}_i = \{(v_{il}, z_{il}) | l = 1, \dots, \Lambda_i\}$ for $i \in \mathcal{I}$ is a set of items numbered from 1 up to Λ_i , each with a value $v_{il} = p_{il}^*$ and a weight $z_{il} = d_{il}^* P_a^i(\mathbf{r}_{il}^*)$. Each of these

items corresponds to a (predetermined) optimal offer $\mathbf{r}_{il}^* = (p_{il}^*, d_{il}^*)$ associated with a certain path l of the FFT of a consumer i . The value v_{il} depicts the payment p_{il}^* to be offered to consumer i and the weight z_{il} represents the expected normalized load reduction $d_{il}^* P_a^i(\mathbf{r}_{il}^*)$ of consumer i due to suggestion \mathbf{r}_{il}^* . Here, the decision variable x_{il} represents the selection ($x_{il} = 1$) or not ($x_{il} = 0$) of the l -th item of set \mathcal{M}_i , i.e., of the l -th option \mathbf{r}_{il}^* for the consumer i , for $i \in \mathcal{I}$ and $l = 1, \dots, \Lambda_i$. Constraint (5.19) denotes that exactly one item must be selected from each set \mathcal{M}_i , i.e., that only one of the available Λ_i options must be selected for each consumer i .

In terms of the knapsack problem, the problem for the CD is to decide which items (options) to include in a collection so that the total weight (expected total load reduction) is greater than or equal to a given constant Q' and the total value (total offered reward) is as small as possible. In other words, the optimization problem amounts to accounting over all possible options and selecting those ones that minimize the maximum potential cost of the CD and lead to an expected total load reduction greater than or equal to Q' . From the FFT of Fig. 5.1a, we have

$$\begin{aligned} \mathcal{M}_i &= \{(v_{i1}, z_{i1}), (v_{i2}, z_{i2}), (v_{i3}, z_{i3})\} \\ &= \left\{ \left(p_{min}, \frac{AP_{i,1} + 1}{N_{i,1} + 2} \right), \left(T_{p_i}, \frac{AP_{i,2} + 1}{N_{i,2} + 2} \right), \left(p_{min}, T_{d'_i} \frac{AP_{i,3} + 1}{N_{i,3} + 2} \right) \right\}. \end{aligned} \quad (5.21)$$

We observe that the problem (5.17)-(5.20) is an integer linear programming (ILP) one. Several advanced algorithms for solving integer linear programs are available, e.g., the *branch-and-bound* or the *cutting plane* method [118, 119]. There are a number of commercial solvers that implement such algorithms for the solution of ILPs [120]. We use such a solver for the derivation of our numerical results in section 5.6.

5.6 Numerical Example

We consider a system consisting of two consumers, i.e., $\mathcal{I} = \{1, 2\}$. We assume that a simple questionnaire (sequence of offers) is provided to each one of them which invites consumers to consider that they receive suggestions about load reduction tasks on their smartphones and have to reply to them (accept or reject). Each offer consists of a monetary payment p awarded to the consumer if s/he carries out the suggested task, and a percentage \tilde{d} by which a consumer's load needs to be reduced. The provided questionnaire and the corresponding responses of a consumer i , for $i = 1, 2$, form the training dataset for this consumer.

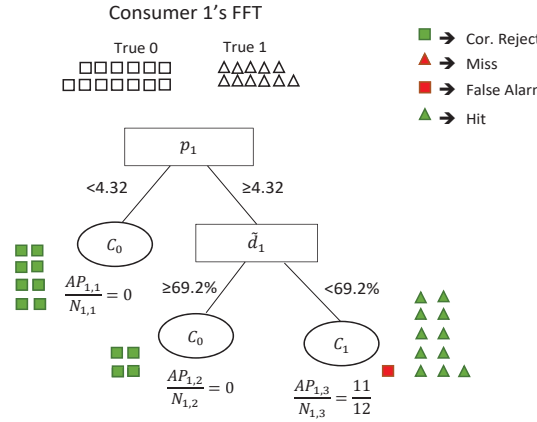


FIGURE 5.3: The resulting FFT of consumer 1: priority to the payment attribute p_1 .

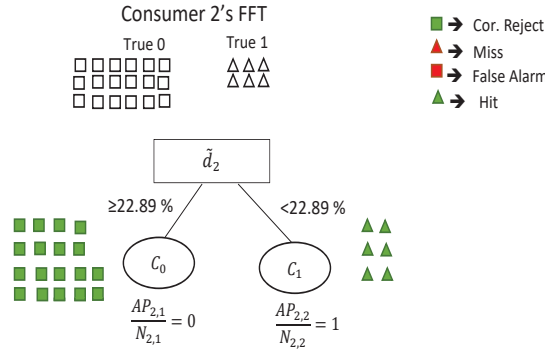


FIGURE 5.4: The resulting FFT of consumer 2: exclusive reliance on the percentage load reduction cue \tilde{d}_2 .

Table C.1 in Appendix C depicts the data we used for the creation of the synthetic training datasets $\mathcal{S}_{tr}^1, \mathcal{S}_{tr}^2$ for the two consumers. We generated 24 random pairs (p, \tilde{d}) (i.e., offers) and labeled them as 1 (accept (C_1)) or 0 (reject (C_0)) based on the different consumer preferences. Specifically, we assume that consumer 1 prefers offers with medium to high compensation (p) and low to medium effort (\tilde{d}), whereas consumer 2 has a preference for low effort tasks, irrespective of reward. Initially, we assume that $p_{min} = \$0.5$ and $p_{max} = \$15$.

First, we train the LR models for the two consumers using the available training datasets. Namely, we estimate the $\mathbf{w}_1, \mathbf{w}_2$ vectors of attribute weights that best fit the models using the questionnaire suggestions and the corresponding consumer responses. Following the training procedure described in Appendix A, we get $w_{1,0} = -0.0882$, $w_{1,p} = 0.3794$ and $w_{1,\tilde{d}} = -0.0608$ for the LR model that captures the preferences of consumer 1, while for the LR model built for consumer 2 we have $w_{2,0} = 1.4642$, $w_{2,p} = -0.0164$ and $w_{2,\tilde{d}} = -0.0696$. Thus, the LR model of consumer 1 predicts (using Eq. (5.9)) that

consumer 1 will accept an offer $r_1 = (p_1, \tilde{d}_1)$ with probability

$$P_a^1(p_1, \tilde{d}_1) = \sigma(w_{1,0} + w_{1,p} p_1 + w_{1,\tilde{d}} \tilde{d}_1) = \sigma(-0.0882 + 0.3794 p_1 - 0.0608 \tilde{d}_1), \quad (5.22)$$

and the LR model of consumer 2 predicts that consumer 2 will accept an offer $r_2 = (p_2, \tilde{d}_2)$ with probability

$$P_a^2(p_2, \tilde{d}_2) = \sigma(w_{2,0} + w_{2,p} p_2 + w_{2,\tilde{d}} \tilde{d}_2) = \sigma(1.4642 - 0.0164 p_2 - 0.0696 \tilde{d}_2), \quad (5.23)$$

Next, applying the simple FFT construction algorithm (described in Appendix B) on the available training datasets, we obtain the FFTs for the two consumers. The resulting decision trees, which are depicted in Fig. 5.3 and 5.4, provide ample evidence of the high heterogeneity in consumer preferences. Consumer 1 assigns more importance to attribute p_1 , since s/he inspects the offered payment p_1 first and then considers the suggested percentage of load reduction \tilde{d}_1 for making her/his decisions. Consumer 2 appears to rely exclusively on single cue \tilde{d}_2 for making her/his choices.

We classify the training offers using the constructed trees and we get $AP_{1,1} = 0$, $AP_{1,2} = 0$, $AP_{1,3} = 11$, $AP_{2,1} = 0$, $AP_{2,2} = 6$, $N_{1,1} = 8$, $N_{1,2} = 4$, $N_{1,3} = 12$, $N_{2,1} = 18$ and $N_{2,2} = 6$. Also, we have $\Lambda_1 = 3$, $\Lambda_2 = 2$, and $C = 2$. From the FFT of consumer 1 (Fig. 5.3) and Eq. (5.15), the probability that consumer 1 will accept an offer $\mathbf{r}_1 = (p_1, \tilde{d}_1)$ is

$$\begin{aligned} P_a^1(p_1, \tilde{d}_1) &= \mathbf{1}_{\{p_1 < T_{p_1}\}} \left(\frac{AP_{1,1} + 1}{N_{1,1} + 2} \right) + \mathbf{1}_{\{p_1 \geq T_{p_1}\}} \mathbf{1}_{\{\tilde{d}_1 \geq T_{\tilde{d}_1}\}} \left(\frac{AP_{1,2} + 1}{N_{1,2} + 2} \right) \\ &\quad + \mathbf{1}_{\{p_1 \geq T_{p_1}\}} \mathbf{1}_{\{\tilde{d}_1 < T_{\tilde{d}_1}\}} \left(\frac{AP_{1,3} + 1}{N_{1,3} + 2} \right) \\ &= \mathbf{1}_{\{p_1 < 4.32\}} \left(\frac{1}{10} \right) + \mathbf{1}_{\{p_1 \geq 4.32\}} \mathbf{1}_{\{\tilde{d}_1 \geq 69.2\%\}} \left(\frac{1}{6} \right) + \mathbf{1}_{\{p_1 \geq 4.32\}} \mathbf{1}_{\{\tilde{d}_1 < 69.2\%\}} \left(\frac{12}{14} \right), \end{aligned} \quad (5.24)$$

while the set of options (extracted from this tree) for the suggestion \mathbf{r}_1 , is $\{\mathbf{r}_{11}^*, \mathbf{r}_{12}^*, \mathbf{r}_{13}^*\} = \{(p_{min}, \tilde{d}_{max}), (T_{p_1}, \tilde{d}_{max}), (T_{p_1}, T_{\tilde{d}_1} - \epsilon_d)\} = \{(\$0.5, 100\%), (\$4.32, 100\%), (\$4.32, 68.7\%\}$ (for $\epsilon_d = 0.5\%$) with acceptance probabilities $1/10$, $1/6$ and $12/14$, respectively.

From the FFT of consumer 2 (Fig. 5.4), the probability that the consumer will accept an offer $\mathbf{r}_2 = (p_2, \tilde{d}_2)$ is

$$\begin{aligned} P_a^2(p_2, \tilde{d}_2) &= \mathbf{1}_{\{\tilde{d}_2 \geq T_{\tilde{d}_2}\}} \left(\frac{AP_{2,1} + 1}{N_{2,1} + 2} \right) + \mathbf{1}_{\{\tilde{d}_2 < T_{\tilde{d}_2}\}} \left(\frac{AP_{2,2} + 1}{N_{2,2} + 2} \right) \\ &= \mathbf{1}_{\{\tilde{d}_2 \geq 22.89\%\}} \left(\frac{1}{20} \right) + \mathbf{1}_{\{\tilde{d}_2 < 22.89\%\}} \left(\frac{7}{8} \right) \end{aligned} \quad (5.25)$$

and the set of options for the suggestion \mathbf{r}_2 is $\{\mathbf{r}_{21}^*, \mathbf{r}_{22}^*\} = \{(p_{min}, \tilde{d}_{max}), (p_{min}, T_{\tilde{d}_2} - \epsilon_d)\} = \{(\$0.5, 100\%), (\$0.5, 22.39\%)\}$ with acceptance probabilities $1/20$ and $7/8$, respectively.

In Appendix C, we provide details about the building process of the above LR models and FFTs. We also study the performance (fitting accuracy) of the resulting models, i.e., how well the model decisions fit the decisions (responses) of the two consumers, and we present a performance comparison between them.

Now, assume that the CD needs to incentivize the two consumers so as to achieve an expected total load reduction of at least $Q = 1.5$ kWh at a specific peak slot of the day. The baseline energy demand of each consumer at the target slot is $e_1^0 = 3$ kWh and $e_2^0 = 2$ kWh, respectively. Thus, $d_1^{max} = 3$ kWh and $d_2^{max} = 2$ kWh. The following question arises: Which are the most economically efficient suggestions that can induce an expected total load reduction greater than or equal to Q ? The answer can be given through the solution of the optimization problems (5.10)-(5.13) and (5.17)-(5.20) of subsections 5.5.1.1 and 5.5.2.1, respectively. We solve the two problems using MATLAB functions that implement the interior-point and branch-and-bound algorithms.

In the case of LR, the optimal suggestions are $\mathbf{r}_{LR,1}^* = (\$12.26, 59.3\%)$ and $\mathbf{r}_{LR,2}^* = (\$0.5, 25.1\%)$, which can induce a maximum potential cost of \$12.76 on the CD and an expected total load reduction of $1.5 \text{ kWh} \geq Q$. On the other hand, in the case of FFTs, the optimal suggestions are $\mathbf{r}_{FFT,1}^* = (\$4.32, 68.7\%)$, $\mathbf{r}_{FFT,2}^* = (\$0.5, 22.39\%)$, and they can achieve an expected total load reduction of $2.16 \text{ kWh} > Q$ at a total offered payment (maximum potential cost) of \$4.82. If the CD wants to achieve an expected load reduction greater than or equal to $Q = 0.8$ kWh, the optimal suggestions are $\mathbf{r}_{LR,1}^* = (\$6.44, 35.96\%)$, $\mathbf{r}_{LR,2}^* = (\$0.5, 25.1\%)$ and $\mathbf{r}_{FFT,1}^* = (\$4.32, 68.7\%)$, $\mathbf{r}_{FFT,2}^* = (\$0.5, 22.39\%)$. The total offered reward and the expected total load reduction are \$6.94 and $0.8 \text{ kWh} \geq Q$, respectively, in the case of LR, while in the case of FFTs they are \$4.82 and $2.16 \text{ kWh} > Q$.

From the above results, we observe that in both approaches and for both values of Q , the optimal offers for consumer 2 are approximately the same. These optimal offers imply that the CD should offer consumer 2 the minimum allowed reward (p_{min}) for the suggested percentage reduction ($\approx 23\%$) that causes the maximum expected load reduction. The attribute values show that both models for consumer 2 have successfully discovered the decision-making profile of the consumer (which prefers low effort tasks ($\tilde{d}_2 < 22.89\%$) irrespective of reward). The obtained optimal suggestions for consumer 1 aim to achieve the rest of the required expected reduction (which is smaller for $Q = 0.8$). In the case of LR and for both values of Q , the optimal payments to be offered to consumer 1 are proportional to the suggested percentage load reductions (due to the

weighted integration of p_1 , \tilde{d}_1 for the estimation of the offer acceptance probability). The FFT optimal suggestions, in contrast, are identical for both values of Q . This stems from the fact that the CD has a countable number of options ($\{\mathbf{r}_{11}^*, \mathbf{r}_{12}^*, \mathbf{r}_{13}^*\}$) for the suggestion \mathbf{r}_1 , and it selects the same option for both values of Q . Again, the optimal attribute values signify the successful identification of the consumer 1 preferences (medium to high reward for low to medium effort).

We observe that the CD provides suggestions to both consumers since $p_{min} = \$0.5$ (and so a percentage load reduction that is proportionate to the \$0.5 payment should at least be suggested to each consumer) and due to the equality constraint (5.19) (that restricts the CD to select exactly one option from the available set of options for a suggestion \mathbf{r}_i). However, if we set $p_{min} = 0$, then the solution of the optimization problem (5.10)-(5.13) implies that suggestions will be offered only to those consumers for whom $p_i > 0$. In the optimization problem (5.17)-(5.20), for $p_{min} = 0$, we can relax the constraint (5.19) such that $\sum_{l \in \mathcal{M}_i} x_{il} \leq 1$, $\forall 1 \leq i \leq |\mathcal{I}|$, and we can also assume that for a path that passes through a branch labeled $p_i < T_{p_i}$ or $p_i \leq T_{p_i}$ the offered payment should be $p_{min} + \epsilon_p$. Then, the CD designer can select *at most* one option from the set of options for a suggestion \mathbf{r}_i , and thus, only those consumers for whom $\sum_{l \in \mathcal{M}_i} x_{il} = 1$ will be offered a suggestion.

5.7 Related work

There is a line of works that relies on principles of optimization and game theory to model energy consumer behavior and interaction, and to propose appropriate incentive mechanisms for consumer engagement in demand-response programs. In order to formulate the participation of customers in DR programs, work [102] develops an economic load model which represents the changes of customer's demand with respect to electricity price changes, incentives as well as penalties imposed to the customers. Each customer's action is determined by maximizing a quadratic customer benefit function. In work [101], the consumer is characterized by a concave utility function of total consumption that also factors her personal valuation. A Vickrey-Clarke-Groves pricing mechanism is performed under the assumptions that user private consumption information is not available, and that users are asked to reveal their consumption. The work [121] studies a prosumer-centric approach for achieving social optimality in energy trading between prosumers and a central unit, and for encouraging energy prosumers to participate in energy trading with the central unit. Each prosumer decides on the amount of energy to supply so as to maximize a concave utility function, while the central unit aims to set the price so as to minimize the total cost of energy.

In work [99], the authors consider a dynamic pricing scheme which is convex to the total hourly energy demand. Each user, which is characterized by a payoff function, exchanges daily consumption schedules with others and updates them accordingly to minimize her/his energy cost for the next day. Work [122] considers two market models, a competitive one and an oligopolistic one, for demand response in power networks, where customers are price taking. In both market cases, the utility company's goal is to charge appropriate prices so as to match electricity supply or to shape electricity demand, and each customer's goal is to maximize her/his net utility which is a concave function of her/his shed load or power draw, respectively. The work [123] is a survey of works where the consumer aspect arises in problems of energy exchange, energy consumption scheduling, appliance scheduling and storage facility management.

In the same context, work [86] devises a hierarchical demand response market model where several competing aggregators, which act as intermediaries, provide monetary incentives to home users to modify their demand pattern. Each user tries to maximize her/his (concave) net payoff function which captures the received compensation and the dissatisfaction caused due to deviation from the reference consumption. In work [124], the authors design incentive offers to consumers with unknown characteristics. Each consumer is characterized by a cost per unit of load reduction that needs to be truthfully elicited, and by a probability of accepting the offer made, which needs to be learned. Work [98] introduces the uncertainty in load curtailment stemming from end-consumer non-engagement through the definition of a conformance probability that depends on consumer valuation and on the provisioned incentives. This probability as well as incentive mechanisms in the form of rewards and fines are inserted in appropriately defined utility functions for the consumers which capture the expected benefit from the curtailment.

In the context of cognitive psychology and behavioral economics, there is significant amount of work on modeling the bounded rationality of decision-makers. Two different points of view have arisen from researchers' attempts to construct models of the heuristics that people use to make decisions. Work [125] expresses the first one. The "heuristics and biases" program presented in that work documented deviations of human behavior from the classical economic model through a number of laboratory experiments. The second point of view is expressed in work [126]. The presented "fast and frugal heuristics" program developed a suite of simple mathematical models known as simple decision heuristics, and analyzed how well these models describe human behavior. Works [127, 128] also presented various such heuristics that the authors have identified and tested. Due to their simplicity, they have been used in many real-world decision tasks.

For example, in work [106] they have been used for making coronary-artery disease diagnoses, in [129] for prescribing antibiotics to children, in [130] for detecting depression, and in work [107] for making legal decisions.

There is a large amount of work that compares the performance of decision heuristics and computationally more complex methods developed in statistics and artificial intelligence. For example, works [105–107, 116, 131–134] show that simple heuristics can outperform (in inferential speed and accuracy) traditional more complex methods. Specifically, the authors of [116] showed, by computer simulation, that the predictive accuracy of fast-and-frugal trees compares well with that of logistic regression and of classification and regression trees (CARTs). Also, works [135, 136] made a performance comparison between heuristics and compensatory integration models (including regression models) using simulated judgement data, whereas work [107] made a similar comparison based on human judgement data. Works [134, 137] also identified the conditions/situations under which simple heuristics are more effective than traditional “rational” models. Similarly, in work [105], the author reviews studies in business, medicine, and psychology, where computer simulations and mathematical analyses reveal conditions under which heuristics make better inferences than optimization, and vice versa.

Another line of works outlines the importance of behavioral science in energy consumption [138, 139]. The work [138] presents an overview of behavioral-science foundations that may be used to improve performance of smart-grid methods such as demand-response programs, time-of-use pricing, energy feedback through smart meters, disaggregation at the appliance level, and smart automation through smart appliances. The work [140] highlights key principles from psychology and behavioral economics to predict and change household energy consumption and energy conservation behavior. It is argued that this behavior often fails to align with personal values or material interests of consumers. The authors in [141] design a large-scale behavioral intervention survey by developing a psychological model based on the theory of planned behavior (TPB) to investigate the influence of information on environmental behavior and green electricity purchase. Results show that price is not the only barrier to purchasing green electricity, and that behavioral, normative and control beliefs influence the decision to purchase green electricity.

Finally, works [142, 143] are closely related to our work. The authors of both works acknowledge that users that participate in mobile crowdsensing (MCS) campaigns exhibit high diversity in decision making because they assess differently attributes related to recommended/presented tasks. In work [142], the authors draw on logistic-regression to learn users’ individual preferences from past data and formulate non-linear optimization problems to determine the tasks and incentives that should be optimally offered to each

user. Work [143], in contrast, demonstrates how users can be profiled in light of cognitive heuristic models and draws on empirical data to compare different decision-making heuristics for MCS task selection. It is assumed that users receive offers about pairs of crowdsensing tasks on their smartphones, and that they are asked to choose one of them to carry out.

5.8 Conclusion

In this chapter, we introduced the optimal incentive and load reduction task allocation problem faced by the designer of an energy consumption reduction campaign. The aim of the designer is to target incentives and load reduction tasks appropriately based on the different consumer profiles, so as to best fulfill the purpose of the campaign. The individual consumer preferences (profiles) and the corresponding probabilities of performing the suggested tasks are discovered through two different behavior-based data-driven approaches, which employ the logistic regression model and the fast-and-frugal tree model. In the case of logistic regression the optimization problem turns out to be a sigmoid optimization one, while in the case of fast-and-frugal trees it is an integer linear programming one. We also provided a numerical example that shows the optimal solution for a special case of the problem.

Chapter 6

Concluding Remarks and Future Challenges

Contents

6.1 Summary of Contributions	111
6.2 Future challenges	113

6.1 Summary of Contributions

In this thesis, we studied the issues of energy storage management and dimensioning, renewable energy exploitation, consumer engagement and consumer modeling. Specifically,

- **Chapter 2:** In this chapter, we underlined the importance of energy storage management and dimensioning in smart grid systems. First, we studied the problem of optimal storage device management with the goal to achieve a certain optimization objective for the energy supplier. Our setup and methodology took into account the dynamics of demand load and renewable sources, and the optimal policy turned out to be a threshold-based one that is expressed through the statistics of these dynamics. Our results provided evidence about the potential of our approaches in reducing energy-supplier operational costs even with limited storage capacities. Next, we dealt with the joint energy storage placement, dimensioning and control problem. We explored the way storage capacity placement impacts the overall cost of energy generation. In determining the optimal policy, it turned out that various aspects of power flow need to be taken into account. It has been

shown that the solution policy for this problem involves various parameters such as the demand profiles of prosumers, and power flow and balance constraints. Our numerical results demonstrate the significant impact of energy storage placement, dimensioning and control on power generation costs.

- **Chapter 3:** Here, we emphasized that in the pursue of green web crawling (and by extension, of green smart grid operations), the dynamics of RESs should be considered. We studied the real-world problem of minimizing the total staleness of web pages in the repository of a web crawler while keeping the amount of carbon emissions on remote web servers, due to the page refresh operations of the crawler, low enough. The optimal page refresh policy, which turned out to be a greedy one, can be implemented in an online fashion taking into account the staleness of the pages in the web repository and the dynamics of the RESs at the remote servers. Each proposed heuristic policy achieves a different objective and relies on different parameters to make its scheduling decisions, using less information than the optimal policy. The real-life dataset provided by the Yahoo! Labs Barcelona, in order to evaluate the performance of the proposed policies, was large and also represented high-quality content. Our experimental results *a)* proved the existence of a tradeoff between staleness and greenness, *b)* indicated the relative importance of different system parameters, and *c)* gave as a good sense of the performance in a realistic system due to the scale of the collected dataset.
- **Chapter 4:** In this chapter, we made a first attempt to develop a theory from first principles on the design of a simple class of serious games for energy efficiency in smart grids. The game designer optimally selects the game parameters, so as the utility-maximizing choices of consumers to minimize the operational cost of the energy supplier for energy production. The sole game parameters utilized are the sizes of the upper list (i.e., winners) and of the lower list (i.e., losers) of consumers according to their energy-consumption reduction. Our simulation experiments showed that even such simple serious games can provide adequate incentives to the consumers, so that the energy supplier achieves specific demand-side management objectives. Our serious-game model can be deployed in practical settings and is privacy-friendly, since only normalized energy consumption increase/decrease needs to be shared by the users with the game designer.
- **Chapter 5:** Here, we presented two data-driven approaches to profile the energy consumer and to bring in the foreground behavior-based models for predicting consumer behavior. We used a machine-learning tool and a cognitive heuristic to build such models for the consumers. Namely, the models are based on logistic regression and fast-and-frugal trees. The ultimate goal was to target load-curtailment

suggestions/offers appropriately based on the different derived criteria that consumers employ to decide whether to accept such offers. In this work, we made an effort to draw links between the presented optimization problem and important modeling work in the areas of machine learning and cognitive sciences. The proposed approaches appear to be effective in capturing the different importance that each consumer places on different criteria and the uncertainty on consumer actions. Although this work could be implemented on several smart grid applications, our emphasis has been on an energy consumption reduction effort.

6.2 Future challenges

We conclude this thesis by discussing some research challenges in the area of smart grid that open interesting research directions.

A first research challenge is that of predicting uncertainties that significantly affect the smart grid system performance. The unpredictability of renewable energy generation is a first source of uncertainty. In this direction, novel mechanisms need to be developed in order to encompass such uncertainties in smart grid systems. These mechanisms should charge prosumers based on the impact of uncertainty on the system, and the cost of bearing such uncertainty. Moreover, new incentives need to be designed in order to manage uncertainty both at the prosumer side (reduction of uncertainty) and at the generation side (actions to mitigate its effects).

The great number of unpredictable and uncontrollable RESs creates another challenge; that of tackling their variable energy output. The real-time control of the renewable energy supply is costly and difficult. Nevertheless, a feasible and cost-efficient alternative is to exploit any possible flexibility in demand so as to absorb variations in the RESs' output. It is expected that consumers will have an active role in the provision of services for balancing supply and demand.

Another challenge arises from the uncertainty on demand. Demand uncertainty has significant implications in real-time control decisions at a system level. Demand forecasting, which might not be feasible to be performed at the level of end-users, is a difficult and costly operation. However, the potential of crowdsourcing could be leveraged in order to predict demand. In the unpublished manuscript [144], the authors made a first attempt to discuss an example scenario where consumers are asked to submit their forecasts about their demand in the next period, and are compensated through a pricing scheme based on the quality of prediction.

In order to better balance supply and demand, entities such as aggregators could act as mediators between consumers and energy suppliers. Aggregators could interact both with the energy suppliers by buying energy, and with the consumers or micro-grid entities by controlling energy storage devices and flexible demand loads. In such a setting, the aggregators would like to coordinate any consumer interactions so as to reach their balancing goals, while profit-maximizing consumers might compete with each other. Here, we identify the need to design a hierarchical DR market model that will capture the interactions of consumers (home users), energy suppliers, and aggregators. Such a hierarchical market needs to be designed with operating rules and protocols such as contracts or options to regulate the market. The work [86] is a first step toward studying a hierarchical 3-layer market comprising generator entities, aggregators and consumers.

Prosumers with co-located storage devices in the premises of their renewable energy sources may refrain from selling the generated renewable energy to the main grid and they may engage in a game-theoretic interaction so as to maximize their profit. The volatility of RESs and dynamic electricity prices create possible situations that may significantly impact the design of electricity markets and regulation policies for the smart grid. Game theory concepts can be applied to devise novel pricing mechanisms that could guide the system to a socially optimal equilibrium. Such schemes will need to be appropriately integrated in such a hybrid energy grid (with both distributed and centralized energy generation), since they are expected to have significant impact on the centralized energy generation.

Furthermore, concepts and models stemming from the sharing economy could be applied to smart grid systems. The sharing economy and the concept of collaborative consumption are based on the sharing of resources. RESs and storage devices are shared by multiple smart grid entities in various sharing regimes and architectures. In this direction, long-term planning decisions (e.g., storage and RES dimensioning and placement), and short-term dynamic decisions (e.g., storage management, power flow, dynamic incentives, dynamic interactions with the main grid) will need to be modeled. These decisions aim for real-time alignment of energy supply and demand at minimum cost for the system.

Finally, the use of behavioral science and behavioral economics in order to understand the consumer/prosumer behavior in the smart grid system is another challenge. Prospect theory, one of the various dominant theories in behavioral sciences, could be used to understand consumers' choices based on their actual behavior and their assessment of potential gains and losses versus assessed levels of risk. Various cognitive heuristics such as the one used in chapter 5 will also aid in understanding the role of consumers/prosumers in the smart grid. This knowledge is expected to reshape the role of utility operators,

renewable source integration, demand-response programs and incentive provisioning in an effort to guide prosumers to energy consumption behaviors that are beneficial to the system. Moreover, novel social-driven demand-response mechanisms could be designed, which would exploit the impact of the social factor, both in the effectiveness of demand-response mechanisms and in cooperative consumer/prosumer coalition formations by means of social norms and competition for social pressure. It is clear that the confluence of multiple disciplines such as algorithm design, game theory, power systems, economics, machine learning, and cognitive psychology is required.

Appendix A

Training the Logistic Regression models

For each consumer i , there is a training set \mathcal{S}_{tr}^i consisting of K past load-reduction suggestions and her/his responses to each,

$$\mathcal{S}_{tr}^i = (\mathbf{r}_{i1}, y_{i1}), (\mathbf{r}_{i2}, y_{i2}), \dots, (\mathbf{r}_{iK}, y_{iK}), \quad (\text{A.1})$$

where $\mathbf{r}_{ij} = (p_{ij}, d'_{ij})$ is the j -th suggestion to consumer i and y_{ij} is the corresponding response, 1 (accept) or 0 (reject), depending on her/his choice.

One way to estimate the consumer weight vector $\mathbf{w}_i = (w_{i,0}, w_{i,p}, w_{i,d'})$ is to follow the Maximum-Likelihood principle and maximize the logarithm of the following likelihood function with respect to \mathbf{w}_i ,

$$P(\mathbf{y}_i | \mathbf{w}_i) = \prod_{j=1}^K P_i(C_1 | \mathbf{r}_{ij})^{y_{ij}} (1 - P_i(C_1 | \mathbf{r}_{ij}))^{1-y_{ij}},$$

with $\mathbf{r}_{ij} = (1, p_{ij}, d'_{ij})$ (a pseudo-attribute $c_{ij,0} = 1$ is added), and $\mathbf{y}_i = (y_{i1}, \dots, y_{iK})$. However, in order to prevent overfitting and improve the generalization of the learned model, in practice we apply a regularization technique. Namely, we add a regularization penalty $\frac{\lambda_i}{2} \|\mathbf{w}_i\|^2$ on weights (L2-regularization), which places preference on smaller values. This technique corresponds to imposing a Gaussian prior distribution on model parameters \mathbf{w}_i [103], [104].

Now, the aim of the training process is to find the weights \mathbf{w}_i that minimize the overall error function

$$\begin{aligned} E(\mathbf{w}_i) &= -\ln P(\mathbf{y}_i|\mathbf{w}_i) + \frac{\lambda_i}{2} \|\mathbf{w}_i\|^2 \\ &= -\sum_{j=1}^K [y_{ij} \cdot \ln(P_i(C_1|\mathbf{r}_{ij})) \\ &\quad + (1 - y_{ij}) \cdot \ln(1 - P_i(C_1|\mathbf{r}_{ij}))] + \frac{\lambda_i}{2} \|\mathbf{w}_i\|^2, \end{aligned} \quad (\text{A.2})$$

with $P_i(C_1|\mathbf{r}_{ij}) = \sigma(\mathbf{w}_i \cdot \mathbf{r}_{ij})$.

The gradient of the above overall error function (deviance plus a cost on the size of attribute weights) (Eq. (A.2)) with respect to \mathbf{w}_i can be shown to be

$$\nabla E(\mathbf{w}_i) = \sum_{j=1}^K [P_i(C_1|\mathbf{r}_{ij}) - y_{ij}] \mathbf{r}_{ij} + \lambda_i \mathbf{w}_i. \quad (\text{A.3})$$

This gradient can be used by a gradient-descent algorithm¹ to iteratively converge to the optimum vector \mathbf{w}_i through

$$\mathbf{w}_i^{(\tau+1)} = \mathbf{w}_i^{(\tau)} - \eta \nabla E(\mathbf{w}_i), \quad (\text{A.4})$$

where the step size (learning rate parameter) η determines the aggressiveness with which the algorithm moves toward the minimum. Since the error function above is a convex function [103, 104], the minimum is a global one. The regularization parameter λ_i is selected so that the classification performance of the LR model in a validation dataset is maximized (cross-validation technique) [103, 104].

¹Gradient descent is a first-order iterative optimization algorithm. To find a local minimum of a function using gradient descent, one takes steps proportional to the negative of the gradient (or of the approximate gradient) of the function at the current point. If instead one takes steps proportional to the positive of the gradient, one approaches a local maximum of that function; the procedure is then known as gradient ascent [145].

Appendix B

Training the Fast-and-Frugal Trees

In order to create an FFT, one can, of course, test all possible orderings of attributes and shapes of trees on a provided training dataset and optimize fitting performance [110]. However, in the general case, this requires enormous computation if the number of attributes is large. They argue that naive decision makers will basically have a good feeling of how well each cue alone predicts the criterion (class label). Thus, they define and use Diagnosticity and Validity¹ of a cue as measures of goodness (or predictivity or accuracy) to order the cues. Another measure of global goodness of an attribute that could also be used for the ranking of attributes is Informedness or Youden’s J². All these metrics can be viewed as statistical measures of the performance (classification accuracy) of a binary classification test that uses a single attribute as an indicator (predictor) to determine the outcome (class). They are estimated using only counting and ratios. Once cue validities/accuracies have been estimated, tree construction involves only a few simple rules [110].

	actual C_1	actual C_0
$x \geq T_x$ (predicted C_1)	α_x	β_x
$x < T_x$ (predicted C_0)	γ_x	δ_x

TABLE B.1: Contingency table over an abstract attribute x .

¹Diagnosticity is the average of sensitivity and specificity, while Validity is the average of positive and negative validity (or predictive value) [110, 146, 147].

²Informedness or Youden’s J (J) is the sum of sensitivity and specificity minus 1. Youden’s J statistic (also called Youden’s index) captures the performance of a binary classification test. It is also known as Δ and generalizes from the binary/dichotomous to the multiclass case as Informedness. It quantifies how informed a predictor (i.e., attribute) is for the specified condition (i.e., outcome or criterion). It specifies the probability that a prediction/decision is informed in relation to the condition (versus chance (i.e., random guess)), and takes into account all predictions [147–149].

FFT construction: First, a four cell contingency table [150, 151] is constructed for the consumer response (1 (accept, C_1) or 0 (reject, C_0)) over each attribute (along with a classification threshold) using the training dataset \mathcal{S}_{tr}^i . These tables summarize the predictions (both correct and wrong) based on each cue for given classification thresholds (split points) and can be used for the estimation of the aforementioned measures of cue accuracy (i.e., diagnosticity, validity, informedness, etc.). They are expressed using raw counts of the number of times each predicted class label is associated with each real class. Each row of the table represents the instances in a predicted class while each column represents the instances in an actual class. Table B.1 shows an example contingency table over an abstract attribute x for given threshold T_x .

A simple rule for the selection of the classification threshold T_x is to pick that value (and the corresponding inequality sign) that maximizes some measure of goodness of the cue x . Here, the criterion we use for measuring the accuracy of a cue and judging whether one attribute x (with classification threshold T_x) is better (i.e., more predictive, accurate or informative) than another is Youden's J . Its magnitude gives the probability of an informed decision between the two classes C_1, C_0 based only on attribute x and given threshold T_x ; its value ranges from -1 to 1 (> 0 represents appropriate use of information, 0 represents chance-level performance, < 0 represents perverse use of information, and 1 indicates perfect prediction)[146, 149]. It is defined as $J_x = TPR_x - FPR_x$ [146–149]. From the contingency table B.1, the proportion $TPR_x = \alpha_x / (\alpha_x + \gamma_x)$ is the true positive rate (or sensitivity or recall or hit rate (HR)), i.e., the proportion of real positive (C_1) cases that are correctly predicted positive based only on cue x and threshold T_x . The proportion $FPR_x = \beta_x / (\beta_x + \delta_x)$, which defines an error rate, is the false positive rate (or fallout or false alarm rate (FAR)), i.e., the proportion of real negative (C_0) cases that are incorrectly predicted positive based only on x and T_x . Thus,

$$J_x = TPR_x - FPR_x = \frac{\alpha_x}{\alpha_x + \gamma_x} - \frac{\beta_x}{\beta_x + \delta_x}. \quad (\text{B.1})$$

The effectiveness (diagnostic power) of an FFT heuristic is determined by both the TPR and the FPR likelihood ratios. Specifically, a high TPR is desired, as it means accurate identification of positive (C_1) cases. However, this benefit might be offset by false positives that are triggered when negative (C_0) cases are also identified as positive.

Using the index of performance J , the following simple algorithm for FFT construction solves the problems of classification threshold (and inequality sign (of the corresponding condition)) selection, cue order and exit location. It is the same one used for the construction of FFTs in R (programming language and software environment) [152], and it can be summarized in four steps:

Fast-and-Frugal Tree construction algorithm

1. For each cue, select a classification threshold (and inequality sign) that maximizes the difference between the TPR and FPR (i.e., the index J) of classifications of all training instances based on that cue. If the attribute is numeric, the threshold is a number. If the attribute is a factor, the threshold is one or more factor levels. These thresholds are calculated completely independently of all other cues.
2. Rank cues in order of their highest TPR-FPR (J) value calculated in step 1.
3. Create all possible trees by varying the exit point at each level to a maximum of 5 levels.
4. Reduce the size of trees by removing lower levels containing less than 5% of the dataset.

From the example FFT of Fig. 5.1a (section 5.5.2) and the above algorithm we deduce that $J_{d'_i} > J_{p_i}$ for optimal thresholds $T_{d'_i}, T_{p_i}$ (that maximize the difference TPR-FPR), i.e., consumer i prioritizes the suggested (normalized) load reduction d'_i over the reward p_i given for carrying the suggested task out.

p (\$)	\tilde{d} (%)	Consumer 1's response (A_1)	Consumer 2's response (A_2)
10	9	1	1
12	46	1	0
13	58	1	0
14	96	0	0
5	12	1	1
6	33	1	0
8	63	1	0
9	92	0	0
1	21	0	1
2	25	0	0
3	71	0	0
4	80	0	0
10.5	10	1	1
11	40	1	0
13.5	59	1	0
15	97	0	0
5.7	13	1	1
7	35	1	0
8.5	64	0	0
9.5	93	0	0
0.5	18	0	1
1.6	23	0	0
2.3	27	0	0
3.5	78	0	0

TABLE C.1: Example questionnaire and consumers' responses.

	C_1	C_0
$p_1 \geq T_{p_1}$ (predicted C_1)	11	5
$p_1 < T_{p_1}$ (predicted C_0)	0	8

TABLE C.2: Contingency table over cue p_1 .

	C_1	C_0
$\tilde{d}_1 < T_{\tilde{d}_1}$ (predicted C_1)	11	6
$\tilde{d}_1 \geq T_{\tilde{d}_1}$ (predicted C_0)	0	7

TABLE C.3: Contingency table over cue \tilde{d}_1 .

	C_1	C_0
$p_2 \leq T_{p_2}$ (predicted C_1)	2	0
$p_2 > T_{p_2}$ (predicted C_0)	4	18

TABLE C.4: Contingency table over cue p_2 .

	C_1	C_0
$\tilde{d}_2 < T_{\tilde{d}_2}$ (predicted C_1)	6	0
$\tilde{d}_2 \geq T_{\tilde{d}_2}$ (predicted C_0)	0	18

TABLE C.5: Contingency table over cue \tilde{d}_2 .

Appendix C

Building procedure details and performance analysis of the fitted LR and FFT models of section 5.6

Table C.1 depicts the data (sequence of offers and consumer responses) we used for the creation of the synthetic training datasets \mathcal{S}_{tr}^1 , \mathcal{S}_{tr}^2 for the two consumers in the

	Specificity (%)	Sensitivity (%)	d-prime (D')	AUC
LR _{nr,1}	100	91	3.109	0.955
LR _{r,1}	100	91	3.109	0.95
FFT ₁	92	100	3.121	0.96
LR _{nr,2}	100	100	3.296	1
LR _{r,2}	100	83	2.877	0.91
FFT ₂	100	100	3.296	1

TABLE C.6: Performance statistics of the fitted LR and FFT models.




FIGURE C.1: Classification tables for (a) the LR model and (b) the regularized LR model of consumer 1.

numerical example of section 5.6. These datasets were used for the building of the individual consumer LR models and FFTs.

For the building of the two LR models of section 5.6 we followed the training process of Appendix A, while for the construction of the consumer FFTs we applied the simple algorithm of Appendix B). The regularization parameters that maximize the classification performance of the two LR models are $\lambda_1 = 0.0289$ and $\lambda_2 = 0.0321$. We got two FFTs for consumer 1 and one single FFT for consumer 2. Here, we kept only one tree for consumer 2; namely, the one that has the maximum difference TPR-FPR. The classification thresholds that maximize the difference between the TPR and FPR (i.e., the index J) of classifications of all data based on each cue $p_1, \tilde{d}_1, p_2,$ and \tilde{d}_2 are $T_{p_1} = 4.32, T_{\tilde{d}_1} = 69.2\%, T_{p_2} = 1.26$ and $T_{\tilde{d}_2} = 22.89\%$, respectively. The contingency tables over cues $p_1, \tilde{d}_1, p_2, \tilde{d}_2$ and for the above thresholds $T_{p_1}, T_{\tilde{d}_1}, T_{p_2}$ and $T_{\tilde{d}_2}$ are depicted in Tables C.2, C.3, C.4, C.5. Furthermore, for the same thresholds it is $J_{p_1}=0.62, J_{\tilde{d}_1}=0.54, J_{p_2}=0.33,$ and $J_{\tilde{d}_2}=1$.

After we have derived the LR and FFT models for the two consumers in section 5.6, we would like to know how well the model decisions fit the decisions made by the consumers.




FIGURE C.2: Classification tables for (a) the LR model and (b) the regularized LR model of consumer 2.

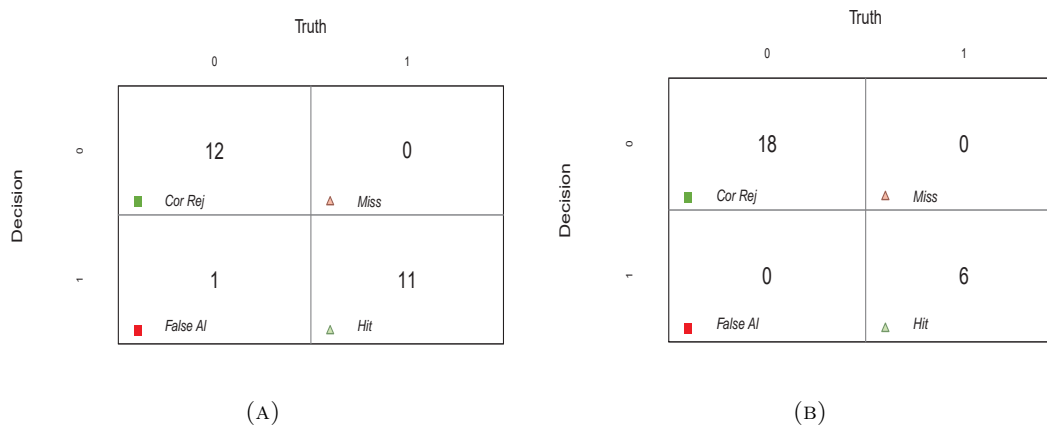


FIGURE C.3: Classification tables for (a) the FFT of consumer 1, and (b) the FFT of consumer 2.

There is a series of performance statistics that can be calculated and checked for that purpose. First, we construct a classification table for each LR (both regularized (LR_r) and non-regularized (LR_{nr})¹) and FFT model. These tables, which are depicted in Fig. C.1, C.2, C.3, show how each model's decisions compare to the truth. Entries on the main diagonal (Cor. Rej. and Hit) correspond to correct decisions, while the other entries (Miss and False Al.) correspond to incorrect decisions.

As we can see, the obtained models performed exceptionally well: the LR_{nr} , LR_r and FFT models of consumer 1, and the LR_r model of consumer 2 made correct decisions in 23 out of all 24 cases (96% correct). The FFT and LR_{nr} models of consumer 2 decided

¹Here, we also provide a performance analysis of the non-regularized LR model in order to emphasize the benefits of regularization (overfitting avoidance). The weights (regression coefficients) of the non-regularized LR models for the two consumers are $\mathbf{w}_1^{nr} = (1.286, 1.44, -0.243)$, $\mathbf{w}_2^{nr} = (330.05, -15.91, -14.06)$.

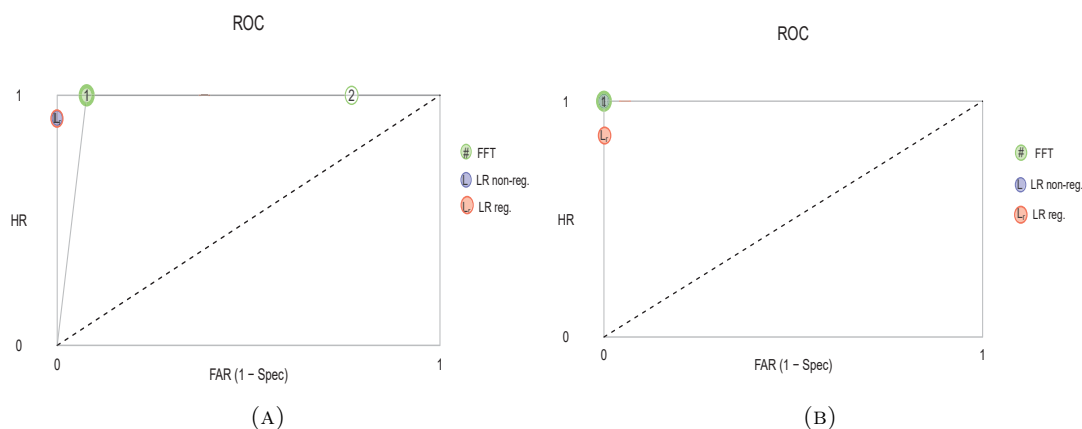


FIGURE C.4: Performance comparison of the FFT and LR models of (a) consumer 1 and (b) consumer 2.

correctly in all 24 cases (100% correct). Additional performance statistics, including Specificity ($=1-\text{FPR}$), Sensitivity ($=\text{TPR}$), d-prime (D') ($=Z(\text{TPR})-Z(\text{FPR})$)², and AUC ($=(\text{TPR}-\text{FPR}+1)/2$) (area under the curve of ROC (described below)) can be estimated. These statistics show the cumulative performance of the models. Table C.6 shows their estimated values for all fitted models. We observe that the LR_{nr} and FFT models of consumer 2 are classification tests with perfect fitting accuracy.

Finally, Fig. C.4a, C.4b show two ROC³ (Receiver Operating Characteristic) curves which compare the performance of the resulting models (FFTs (in green), LR_{nr} (in purple) and LR_r (in orange)). Specifically, Fig. C.4a shows that the performance $(\text{FAR}_{\text{FFT}_1}, \text{HR}_{\text{FFT}_1})=(0.08, 1)$ (i.e., $J_{\text{FFT}_1}=0.92$) of the FFT of consumer 1 is comparable to that of the LR_{nr} and LR_r models of the same consumer with $(\text{FAR}_{\text{LR}_{nr,1}}, \text{HR}_{\text{LR}_{nr,1}})=(\text{FAR}_{\text{LR}_r,1}, \text{HR}_{\text{LR}_r,1})=(0,0.91)$ (i.e., $J_{\text{LR}_{nr,2}}=J_{\text{LR}_r,1}=0.91$). From Fig. C.4b, we observe that the LR_{nr} and FFT models built for consumer 2 have identical performance, $(\text{FAR}_{\text{FFT}_2}, \text{HR}_{\text{FFT}_2})=(\text{FAR}_{\text{LR}_{nr,2}}, \text{HR}_{\text{LR}_{nr,2}})=(0,1)$ (i.e., $J_{\text{FFT}_2} = J_{\text{LR}_{nr,2}}=1$ (perfect fit)), and that they outperform the LR_r model of the same consumer with $(\text{FAR}_{\text{LR}_r,2}, \text{HR}_{\text{LR}_r,2})=(0,0.83)$ (i.e., $J_{\text{LR}_r,2}=0.83$).

²Function $Z(p)$, $p \in [0, 1]$, is the inverse of the cumulative distribution function of the Gaussian distribution. Here, we adjust any extreme values (i.e., 0 or 1) of TPRs and FPRs, to prevent infinite values or indeterminate forms of D' , according to an approach presented in [153]. Namely, rates of 0 are replaced with $0.5/n$ and rates of 1 are replaced with $(n - 0.5)/n$, where n is the number of positive or negative training instances.

³A ROC curve shows the ability of a quantitative diagnostic test to correctly classify subjects as the decision threshold is varied. A diagnostic test able to perfectly identify subjects with and without the positive condition (here, C_1) produces a curve that passes through the upper left corner (0, 1) of the plot. A diagnostic test with no ability to discriminate better than chance, produces a diagonal line from the origin (0, 0) to the top right corner (1, 1) of the plot. Most tests lie somewhere in-between these extremes [154].

Bibliography

- [1] F. Birol, “Power to the People: The World Outlook for Electricity Investment”, *IAEA Bulletin*, vol. 46, no. 1, pp. 9-12, June 2004. Available at: <https://www.iaea.org/sites/default/files/publications/magazines/bulletin/bull46-1/46104990912.pdf>
- [2] ABB. (Apr. 20 2016). “What is a smart grid?” [Online]. Available: <http://new.abb.com/smartgrids/what-is-a-smart-grid>
- [3] I. Koutsopoulos and L. Tassiulas, “Optimal Control Policies for Power Demand Scheduling in the Smart Grid”, *IEEE Journal on Selected Areas in Communications (JSAC)*, vol.30, no.6, pp. 1049 - 1060, 2012.
- [4] P. Davis, “What Will European Electricity Look Like in 2050? Changes in 2050 Electricity Use”, *Imperial College London, Energy Future Labs*, August 7, 2015 [Online]. Available: <http://wwwf.imperial.ac.uk/blog/energyfutureslab/2015/08/07/what-will-european-electricity-look-like-in-2050-changes-in-2050-electricity-use/>
- [5] M. Pickavet et al., “Worldwide Energy Needs for ICT: The rise of Power-Aware Networking”, *Proc. IEEE International Symposium on Advanced Networks and Telecommunication Systems (ANTS)*, pp. 1-3, Mumbai, India, December 2008.
- [6] International Tecommunication Union, “ICTs and Climate Change” [Online]. Available: http://www.itu.int/themes/climate/docs/report/02_ICTandClimateChange.html
- [7] P. Corcoran and A. Andrae, “Emerging Trends in Electricity Consumption for Consumer ICT”, National University of Ireland, Galway, Connacht, Ireland, Tech. Rep., July 2013. Available at: <http://hdl.handle.net/10379/3563> [Accessed on April 21, 2016].
- [8] Greenpeace Inc. (May 2015). “Clicking Clean: A Guide to Building the Green Internet” [Online]. Available: <http://www.greenpeace.org/usa/wp-content/uploads/legacy/Global/usa/planet3/PDFs/2015ClickingClean.pdf>
- [9] R. Kumar and L. Mieritz, “Conceptualizing “Green IT” and Data Center Power and Cooling Issues”, Gartner Research Paper No. G00150322, September 2007.

- [10] S. Keshav and C. Rosenberg, “How Internet Concepts and Technologies can help Green and Smarten the Electrical Grid”, *Proc. ACM SIGCOMM Green Networking Workshop*, pp. 35-40, New Delhi, India, August 2010.
- [11] C. Thrampoulidis, S. Bose, and B. Hassibi, “Optimal Placement of Distributed Energy Storage in Power Networks”, *IEEE Transactions on Automatic Control*, vol. 61, no. 2, pp. 416-429, February 2016.
- [12] I. Koutsopoulos, V. Hatzi, and L. Tassiulas, “Optimal Energy Storage Control Policies for the Smart Power Grid”, *Proc. IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pp. 475-480, Brussels, Belgium, October 2011.
- [13] V. Hatzi, I. Koutsopoulos, and L. Tassiulas, “Optimal Energy Storage Management Policies for Smart Grid Systems”, *in preparation for submission*.
- [14] I. Koutsopoulos, T. Papaioannou, and V. Hatzi, “Modeling and Optimization of the Smart Grid Ecosystem”, *Foundations and Trends in Networking*, vol. 10, no. 2-3, pp. 115-316, June 2016.
- [15] D. P. Bertsekas and R. G. Gallager, *Data Networks (2nd ed.)*. Englewood Cliffs, NJ: Prentice Hall, 1992.
- [16] A. Giannitrapani, S. Paoletti, A. Vicino, and D. Zarrilli, “Bidding Strategies for Renewable Energy Generation with Non Stationary Statistics”, *Proc. World Congress of the International Federation of Automatic Control (IFAC)*, pp. 10784-10789, Cape Town, South Africa, August 2014.
- [17] S. H. Low, “Convex Relaxations of Optimal Power Flow—Part I: Formulations and Equivalence,” *IEEE Transactions on Control of Network Systems*, vol. 1, no. 1, pp. 15-27, March 2014.
- [18] Y. Baghzouz, “Power Flow Analysis”, Lecture Notes, 2011. Available at: [http://www.egr.unlv.edu/eebag/ECG 740 Power Flow Analysis.pdf](http://www.egr.unlv.edu/eebag/ECG%20740%20Power%20Flow%20Analysis.pdf)
- [19] D. Kirschen, “Optimal Power Flow”, Lecture Notes, 2011. Available at: www.ee.washington.edu/research/real/Library/Teaching/06-OPF.pptx.
- [20] R. Weron, *Modeling and Forecasting Electricity Loads and Prices: A Statistical Approach*. Hoboken, NJ: Wiley, October 2006.
- [21] A. Sfetsos, “A Comparison of Various Forecasting Techniques Applied to Mean Hourly Wind Speed Time Series”, *Renewable Energy*, vol. 21, no. 1, pp. 23-35, September 2000.

- [22] M. Adamou and S. Sarkar, "A Framework for Optimal Battery Management for Wireless Nodes", *Proc. IEEE Conference on Computer Communications (INFOCOM)*, vol.3, pp. 1783-1792, New York, USA, June 2002.
- [23] L. Lin, N. B. Shroff, and R. Srikant, "Asymptotically Optimal Power-Aware Routing for Multihop Wireless Networks with Renewable Energy Sources," *Proc. IEEE Conference on Computer Communications (INFOCOM)*, vol. 2, pp. 1262-1272, Miami, FL, USA, March 2005.
- [24] M. Gatzianas, L. Georgiadis, and L. Tassiulas, "Control of Wireless Networks with Rechargeable Batteries", *IEEE Transactions on Wireless Communications*, vol. 9, no. 2, pp. 581-593, February 2010.
- [25] L. Huang and M. J. Neely, "Utility Optimal Scheduling in Energy Harvesting Networks", *Proc. 12th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, Paris, France, May 2011.
- [26] P. M. V. D. Ven, N. Hegde, L. Massoulié and T. Salonidis, "Optimal Control of Residential Energy Storage Under Price Fluctuations", *Proc. IARIA International Conference on Smart Grids, Green Communications and IT Energy-aware Technologies (ENERGY)*, Venice/Mestre, Italy, May 2011.
- [27] P.M.V.D. Ven, N. Hegde, L. Massoulié, and T. Salonidis, "Optimal Control of End-User Energy Storage", *IEEE Transactions on Smart Grid*, vol. 4, no. 2, pp. 789-797, June 2013.
- [28] A. K. Mishra, D. E. Irwin, P. J. Shenoy, J. Kurose, and T. Zhu, "SmartCharge: Cutting the Electricity Bill in Smart Homes with Energy Storage", *Proc. 3rd International Conference on Future Energy Systems (e-Energy)*, pp. 1-10, Madrid, Spain, May 2012.
- [29] L. Huang, J. Walrand and K. Ramchandran, "Optimal Demand Response with Energy Storage Management", *Proc. IEEE 3rd International Conference on Smart Grid Communications (SmartGridComm)*, pp. 61-66, Tainan city, Taiwan, November 2012.
- [30] T. Erseghe, A. Zanella and C. G. Codemo, "Optimal and Compact Control Policies for Energy Storage Units With Single and Multiple Batteries," *IEEE Transactions on Smart Grid*, vol. 5, no. 3, pp. 1308-1317, May 2014.
- [31] H. M. Soliman and A. Leon-Garcia, "Game-Theoretic Demand-Side Management With Storage Devices for the Future Smart Grid", *IEEE Transactions on Smart Grid*, vol. 5, no. 3, pp. 1475-1485, May 2014.

- [32] Y. Wang, W. Saad, N. B. Mandayam, and H. V. Poor, "Integrating Energy Storage Into the Smart Grid: A Prospect Theoretic Approach", *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7779-7783, Florence, Italy, May 2014.
- [33] Y. Zhou, A. Scheller-Wolf, N. Secomandi, and S. Smith, *Managing Wind-Based Electricity Generation with Storage and Transmission Capacity*. Tepper Working Paper, Technical Report 2011-E36. Carnegie Mellon University, Pittsburgh, 2013.
- [34] K. Wang, F. Ciucu, C. Lin, and S. H. Low, "A Stochastic Power Network Calculus for Integrating Renewable Energy Sources into the Power Grid", *IEEE Journal on Selected Areas in Communications*, vol. 30, no. 6, pp. 1037-1048, July 2012.
- [35] N. Gast, D.-C. Tomozei, and J.-Y. Le Boudec, "Optimal Generation and Storage Scheduling in the Presence of Renewable Forecast Uncertainties", *IEEE Transactions on Smart Grid*, vol. 5, no. 3, pp. 1328-1339, May 2014.
- [36] S. Lakshminaryana, T. Q. S. Quek, H. V. Poor, "Cooperation and Storage Trade-offs in Power Grids with Renewable Energy Resources", *IEEE Journal on Selected Areas in Communications*, vol. 32, no. 7, pp. 1386-1397, July 2014.
- [37] R. Uргаonkar, B. Uргаonkar, M. J. Neely and A. Sivasubramaniam, "Optimal Power Cost Management Using Stored Energy in Data Centers", *Proc. ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, pp. 221-232, San Jose, CA, USA, June 2011.
- [38] Y. Guo, Z. Ding, Y. Fang, D. Wu, "Cutting Down Electricity Cost in Internet Data Centers by Using Energy Storage", *Proc. IEEE Global Telecommunications Conference (GLOBECOM)*, pp. 1-5, Houston, TX, USA, December 2011.
- [39] L. Gkatzikis, G. Iosifidis, I. Koutsopoulos and L. Tassiulas, "Collaborative Placement and Sharing of Storage Resources in the Smart Grid", *Proc. IEEE 3rd International Conference on Smart Grid Communications (SmartGridComm)*, Venice, Italy, November 2014.
- [40] W. Saad, Z. Han and H. V. Poor, "Coalitional Game Theory for Cooperative Micro-Grid Distribution Networks", *Proc. IEEE International Conference on Communications Workshops (ICC)*, pp. 1-5, Kyoto, Japan, May 2011.
- [41] Y. Zhang, N. Gatsis, and G. B. Giannakis, "Robust Management of Distributed Energy Resources for Microgrids with Renewables", 2012, *arxiv CoRR*, arXiv:1207.4831.
- [42] B. B. Cambazoglu and R. A. Baeza-Yates, "Scalability Challenges in Web Search Engines", *Synthesis Lectures on Information Concepts, Retrieval, and Services*, Morgan & Claypool Publishers, December 2015.

- [43] M. de Kunder. *The Size of the World Wide Web (The Internet)* [Online]. Available: <http://www.worldwidewebsite.com/>
- [44] Verbatique. (Mar. 25, 2015). *Average Power Use Per Server* [Online]. Available: <http://www.verbatique.com/average-power-use-server>
- [45] Verbatique. (Oct. 15, 2009). *Carbon Footprints of Servers Can Vary by 10X* [Online]. Available: <http://www.verbatique.com/carbon-footprints-servers-can-vary-10x>
- [46] S. Panwar, D. Towsley and J. Wolf, “Optimal Scheduling Policies for a Class of Queues with Customer Deadlines to the Beginning of Service”, *Journal of ACM*, vol. 35, no. 4, pp. 832-844, October 1988.
- [47] S. Shakkottai and R. Srikant, “Scheduling Real-time Traffic With Deadlines Over a Wireless Channel”, *Proc. ACM International Workshop on Wireless Mobile Multimedia (WoWMoM99)*, pp. 35-42, Seattle, WA, USA, August 1999.
- [48] T. Ren, I. Koutsopoulos, and L. Tassiulas, “QoS Provisioning for Real-time Traffic in Wireless Packet Networks”, *Proc. IEEE Global Telecommunications Conference (GLOBECOM)*, pp. 1673-1677, Taipei, Taiwan, November 2002.
- [49] V. Hatzi, B. B. Cambazoglu and I. Koutsopoulos, “Web Page Download Scheduling Policies for Green Web Crawling”, *Proc. 22nd International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, September 2014.
- [50] V. Hatzi, B. B. Cambazoglu, and I. Koutsopoulos, “Optimal Web Page Download Scheduling Policies for Green Web Crawling”, *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 5, pp. 1378-1388, May 2016.
- [51] L. Page, S. Brin, R. Motwani, and T. Winograd, *The PageRank Citation Ranking: Bringing Order to the Web*. Technical Report 1999-66. Stanford InfoLab, November 1999. Available at: <http://ilpubs.stanford.edu:8090/422/1/1999-66.pdf>
- [52] Woodbank Communications Ltd. (2005). “Solar Power (Technology and Economics)”, *Electropaedia: Battery and Energy Technologies* [Online]. Available: http://www.mpoweruk.com/solar_power.htm
- [53] B. Aksanli, J. Venkatesh, L. Zhang, and T. Rosing, “Utilizing Green Energy Prediction to Schedule Mixed Batch and Service Jobs in Data Centers”, *Proc. 4th Workshop on Power-Aware Computing and Systems (HotPower)*, article no. 5, Cascais, Portugal, October 2011.
- [54] J. Cho and H. Garcia-Molina, “Effective Page Refresh Policies for Web Crawlers”, *ACM Transactions on Database Systems (TODS)*, vol. 28, no. 4, pp. 390-426, December 2003.

- [55] J. Edwards, K. McCurley, and J. Tomlin, "An Adaptive Model for Optimizing Performance of an Incremental Web Crawler", *Proc. 10th International Conference on World Wide Web (WWW)*, pp. 106-113, Hong Kong, May 2001.
- [56] Q. Tan and P. Mitra, "Clustering-Based Incremental Web Crawling", *ACM Transactions on Information Systems (TOIS)*, vol. 28, no. 4, article no. 17, November 2010.
- [57] K. Radinsky and P. N. Bennett, "Predicting Content Change on the Web", *Proc. 6th ACM International Conference on Web Search and Data Mining (WSDM)*, pp. 415-424, Rome, Italy, February 2013.
- [58] J. L. Wolf, M. S. Squillante, P. S. Yu, J. Sethuraman, and L. Ozsen, "Optimal Crawling Strategies for Web Search Engines", *Proc. 11th International Conference on World Wide Web (WWW)*, pp. 136-147, Honolulu, Hawaii, USA, May 2002.
- [59] S. Pandey and C. Olston, "User-Centric Web Crawling", *Proc. 14th International Conference on World Wide Web (WWW)*, pp. 401-411, Chiba, Japan, May 2005.
- [60] C. Olston and S. Pandey, "Recrawl Scheduling Based on Information Longevity", *Proc. 17th International Conference on World Wide Web (WWW)*, pp. 437-446, Beijing, China, April 2008.
- [61] A. Dua and N. Bambos, "Downlink Wireless Packet Scheduling with Deadlines", *IEEE Transactions on Mobile Computing*, vol. 6, no. 12, pp. 1410-1425, December 2007.
- [62] W. Yuan and K. Nahrstedt, "Energy-Efficient Soft Real-Time CPU Scheduling for Mobile Multimedia Systems", *Proc. 19th ACM Symposium on Operating Systems Principles (SOSP)*, pp. 149-163, Bolton Landing, NY, USA, October 2003.
- [63] A. El Gamal, C. Nair, B. Prabhakar, E. Uysal-Biyikoglu, and S. Zahedi, "Energy-efficient Scheduling of Packet Transmissions over Wireless Networks", *Proc. IEEE Conference on Computer Communications (INFOCOM)*, pp. 1773-1782, New York, USA, June 2002.
- [64] L. Wang and Y. Xiao, "A Survey of Energy-Efficient Scheduling Mechanisms in Sensor Networks", *Mobile Networks and Applications*, vol. 11, no. 5, pp. 723-740, October 2006.
- [65] Z. Wang, N. Tolia, and C. Bash, "Opportunities and Challenges to Unify Workload, Power, and Cooling Management in Data Centers", *ACM SIGOPS Operating Systems Review*, vol. 44, no. 3, pp. 41-46, July 2010.

- [66] D. Kliazovich, P. Bouvry, and S. U. Khan, "DENS: Data Center Energy-Efficient Network-Aware Scheduling", *Cluster Computing*, vol. 16, no. 1, pp. 65-75, March 2013.
- [67] D. Xu and X. Liu, "Geographic Trough Filling for Internet Datacenters", *Proc. IEEE Conference on Computer Communications (INFOCOM)*, pp. 2881-2885, Orlando, FL, USA, March 2012.
- [68] J. Shuja, K. Bilal, S. A. Madani, M. Othman, R. Ranjan, P. Balaji, and S. U. Khan, "Survey of Techniques and Architectures for Designing Energy-Efficient Data Centers", *IEEE Systems Journal*, vol. 10, no. 2, pp. 507-519, July 2014.
- [69] C. Ren, D. Wang, B. Urgaonkar, and A. Sivasubramaniam, "Carbon-Aware Energy Capacity Planning for Datacenters", *Proc. IEEE 20th International Symposium on Modeling, Analysis & Simulation of Computer and Telecommunication Systems (MASCOTS)*, pp. 391-400, Arlington, Virginia, USA, August 2012.
- [70] K. Le, O. Bilgir, R. Bianchini, M. Martonosi, and T. D. Nguyen, "Managing the Cost, Energy Consumption, and Carbon Footprint of Internet Services", *Proc. ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, pp. 357-358, New York, USA, June 2010.
- [71] Z. Liu, M. Lin, A. Wierman, S. H. Low and L. L. H. Andrew, "Geographical Load Balancing with Renewables", *ACM SIGMETRICS Performance Evaluation Review*, vol. 39, no. 3, pp. 62-66, December 2011.
- [72] X. Li, Z. Qian, S. Lu, and J. Wu, "Energy Efficient Virtual Machine Placement Algorithm With Balanced and Improved Resource Utilization in a Data Center", *Mathematical and Computer Modelling*, vol. 58, no. 5-6, pp. 1222-1235, September 2013.
- [73] W. Yue and Q. Chen, "Dynamic Placement of Virtual Machines with Both Deterministic and Stochastic Demands for Green Cloud Computing", *Mathematical Problems in Engineering*, vol. 2014, Article ID 613719, 11 pages, July 2014.
- [74] D. Hatzopoulos, I. Koutsopoulos, G. Koutitas, and W. Van Heddeghem, "Dynamic Virtual Machine Allocation in Cloud Server Facility Systems with Renewable Energy Sources", *Proc. IEEE International Conference on Communications (ICC)*, pp. 4217-4221, Budapest, Hungary, June 2013.
- [75] J. Doyle, R. Shorten, and D. O' Mahony, "Stratus: Load Balancing the Cloud for Carbon Emissions Control", *IEEE Transactions on Cloud Computing*, vol. 1, no. 1, pp. 116-128, August 2013.

- [76] T. Marsh, "Serious Games Continuum: Between Games for Purpose and Experiential Environments for Purpose", *Entertainment Computing*, vol.2, no. 2, pp. 61-68, 2011.
- [77] E. Knol and P. W. de Vries, "EnerCities: A Serious Game to Stimulate Sustainability and Energy Conservation:Preliminary Results", *eLearning Papers*, no. 25, July 2011. Available at: <https://ssrn.com/abstract=1866206>
- [78] B. Reeves, J. J. Cummings, and D. Anderson, "Leveraging the Engagement of Games to Change Energy Behavior", *Proc. CHI Workshop on Gamification*, Vancouver, BC, Canada, May 2011.
- [79] A. Bourazeri and J. Pitt, "Serious Game Design for Inclusivity and Empowerment in Smartgrids", *Proc. 1st International Workshop on Intelligent Digital Games for Empowerment and Inclusion (IDGEI)*, Chania, Crete, Greece, May 2013.
- [80] R. S. Brewer, Y. Xu, G. E. Lee, M. Katchuck, C. A. Moore, and P. M. Johnson, "Energy Feedback for Smart Grid Consumers: Lessons Learned From the Kukui Cup", *Proc. IARIA International Conference on Smart Grids, Green Communications and IT Energy-aware Technologies (ENERGY)*, pp. 120-126, Lisbon, Portugal, March 2013.
- [81] Y. Lurie. (Oct. 20, 2014). *Gamification for Utilities: Captivating Customers with Energy Information* [Online]. Available: <http://simpleenergy.com/gamification-for-utilities-captivating-customers-with-energy-information/>
- [82] T. G. Papaioannou, V. Hatzi and I. Koutsopoulos, "Optimal Design of Serious Games for Demand Side Management", *Proc. IEEE International Conference on Smart Grid Communications (SmartGridComm)*, Venice, Italy, November 2014.
- [83] T. Papaioannou, V. Hatzi, and I. Koutsopoulos, "Optimal Design of Serious Games for Consumer Engagement in the Smart Grid", *IEEE Transactions on Smart Grid*, 2016.
- [84] D. Kozikowski, A. Breidenbaugh, and M. Potter, "The Demand Response Baseline", *EnerNOC OPS Publication*, v.1.75, 2006.
- [85] S. Rottiers, "The Sociology of Social Recognition: Competition in Social Recognition Games", Working paper no. 10/04, *Heman Deleeck Center of Social Policy, University of Antwerp*, September 2010.
- [86] L. Gkatzikis, I. Koutsopoulos, T. Salonidis, "The Role of Aggregators in Smart Grid Demand Response Markets", *IEEE Journal on Selected Areas in Communications*, vol.31, no.7, pp.1247-1257, July 2013.

- [87] N. Forouzandehmehr, S. M. Perlaza, Z. Han, and H. V. Poor, “A Satisfaction Game for Heating, Ventilation and Air Conditioning Control of Smart Buildings”, *Proc. IEEE Global Communications Conference Globecom*, Atlanta, GA, USA, December 2013.
- [88] I. L. Glicksberg, “A Further Generalization of the Kakutani Fixed Point Theorem, with Application to Nash Equilibrium Points”, *Proc. American Mathematical Society*, vol. 3, no. 1, pp. 170-174, February 1952.
- [89] N. Li, L. Chen, and S. H. Low, “Optimal Demand Response Based on Utility Maximization in Power Networks”, *Proc. IEEE Power and Energy Society General Meeting*, pp. 1-8, Detroit, Michigan, USA, July 2011.
- [90] D. Geelen, D. Keyson, S. Boess, and H. Brezet, “Exploring the Use of a Game to Stimulate Energy Saving in Households”, *Journal of Design Research*, vol. 10, no. 1-2, pp. 102-120, 2012.
- [91] B. Orland, N. Ram, D. Lang, K. Houser, N. Kling, and M. Coccia, “Saving Energy in an Office Environment: A Serious Game Intervention”, *Energy and Buildings*, vol. 72, pp. 43-52, May 2014.
- [92] G. E. Lee, Y. Xu, R. S. Brewer, and P. M. Johnson, “Makahiki: An Open Source Game Engine for Energy Education and Conservation,” Department of Information and Computer Sciences, University of Hawaii, Honolulu, Hawaii, Tech. Rep. CSDL-11-07, 2012. Available at: <http://csdl.ics.hawaii.edu/techreports/2011/11-07/11-07.pdf>
- [93] P. M. Johnson, Y. Xu, R. S. Brewer, G. E. Lee, M. Katchuck, and C. A. Moore, “Beyond kWh: Myths and Fixes for Energy Competition Game Design”, *Proc. Meaningful Play*, Lansing, Michigan, USA, October 2012.
- [94] J. Costa, R. R. Wehbe, J. Robb and L. E. Nacke, “Time’s Up: Studying Leaderboards For Engaging Punctual Behaviour”, *Proc. 1st International Conference on Gameful Design, Research, and Applications (Gamification)*, pp. 26-33, Toronto, Ontario, Canada, October 2013.
- [95] E. D. Mekler, F. Brühlmann, K. Opwis and A. N. Tuch, “Do Points, Levels and Leaderboards Harm Intrinsic Motivation?: An Empirical Analysis of Common Gamification Elements”, *Proc. 1st International Conference on Gameful Design, Research, and Applications (Gamification)*, pp. 66-73, Toronto, Ontario, Canada, October 2013.

- [96] S. Maharjan, Q. Zhu, Y. Zhang, S. Gjessing, and T. Başar, “Dependable Demand Response Management in the Smart Grid: A Stackelberg Game Approach”, *IEEE Transactions on Smart Grid*, vol. 4, no. 1, pp. 120-132, March 2013.
- [97] T. K. Wijaya, T. G. Papaioannou, X. Liu, and K. Aberer, “Effective Consumption Scheduling for Demand-Side Management in the Smart Grid Using Non-Uniform Participation Rate”, *Proc. 3rd IFIP Conference on Sustainable Internet and ICT for Sustainability (SustainIT)*, Palermo, Italy, October 2013.
- [98] A. Anastopoulou, I. Koutsopoulos, and G. D. Stamoulis, “Efficient Incentive-Driven Consumption Curtailment Mechanisms in Nega-Watt Markets”, *Proc. IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pp. 734-739, Venice, Italy, November 2014.
- [99] A.-H. Mohsenian-Rad , V. W. S. Wong , J. Jatskevich and R. Schober, “Optimal and Autonomous Incentive-Based Energy Consumption Scheduling Algorithm for Smart Grid”, *Proc. IEEE PES Conference on Innovative Smart Grid Technologies (ISGT)*, pp. 1-6, Gothenburg, Sweden, January 2010.
- [100] I. Ullah *et al.*, “An Incentive-based Optimal Energy Consumption Scheduling Algorithm for Residential Users”, *Procedia Computer Science*, vol. 52, no. 1, pp. 851-857, December 2015.
- [101] P. Samadi, A. -H. Mohsenian-Rad, R. Schober, and V. W. S. Wong, “Advanced Demand Side Management for the Future Smart Grid Using Mechanism Design”, *IEEE Transactions on Smart Grid*, vol. 3, no. 3, pp. 1170-1180, September 2012.
- [102] H. A. Aalamia, M. Parsa, and G. R. Yousefi, “Demand Response Modeling Considering Interruptible/Curtailable Loads and Capacity Market Programs”, *Applied Energy*, vol. 87, no. 1, pp. 243-250, January 2010.
- [103] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. NY, USA: Springer-Verlag New York Inc., 2006.
- [104] K. P. Murphy, *Machine Learning: A Probabilistic Perspective*. Cambridge, MA, USA: MIT Press, 2012.
- [105] K. V. Katsikopoulos, “Psychological Heuristics for Making Inferences: Definition, Performance, and the Emerging Theory and Practice”, *Decision Analysis*, vol. 8, no. 1, pp. 10-29, March 2011.
- [106] L. Green and D. R. Mehr, “What Alters Physicians’ Decisions to Admit to the Coronary Care Unit?”, *Journal of Family Practice*, vol. 45, no. 3, pp. 219-226, 1997.

- [107] M. K. Dhimi and C. Harries, “Fast and Frugal Versus Regression Models of Human Judgment”, *Thinking & Reasoning*, vol. 7, no. 1, pp. 5-27, 2001.
- [108] H. A. Simon, “Rational Choice and the Structure of the Environment”, *Psychological Review*, vol. 63, no. 2, pp. 129-138, March 1956.
- [109] G. Gigerenzer, and P. M. Todd, “Fast and Frugal Heuristics: The Adaptive Toolbox”, In G. Gigerenzer, P.M. Todd, and the ABC Research Group, *Simple Heuristics That Make Us Smart* (pp. 3-34). New York: Oxford University Press, 1999.
- [110] L. Martignon, O. Vitouch, M. Takezawa, and M. Forster, “Naïve and Yet Enlightened: From Natural Frequencies to Fast and Frugal Decision Trees”, In D. Hardman, & L. Macchi (Eds.), *Thinking: Psychological Perspectives on Reasoning, Judgment, and Decision Making* (pp. 189-211). Chichester, West Sussex, England: John Wiley and Sons, 2003.
- [111] F. Provost and P. Domingos, “Tree Induction for Probability-Based Ranking”, *Machine Learning*, vol.52, no.3, pp. 199-215, September 2003.
- [112] B. Zadrozny, and C. Elkan, “Obtaining Calibrated Probability Estimates From Decision Trees and Naive Bayesian Classifiers”, *Proc. 18th International Conference on Machine Learning (ICML)*, pp. 609-616, Williamstown, MA, USA, June 2001.
- [113] N. V. Chawla, “C4.5 and Imbalanced Data sets: Investigating the Effect of Sampling Method, Probabilistic Estimate, and Decision Tree Structure”, *Proc. 20th International Conference on Machine Learning (ICML): Learning from Imbalanced Data Sets II workshop*, Washington, DC, USA, August 2003.
- [114] V. Hatzi, and I. Koutsopoulos, “Optimal Design of Energy Consumption Reduction Campaigns Through Behavioral Data-Driven Consumer Profiling”, *in preparation for submission*.
- [115] D. P. Bertsekas, *Non Linear Programming*. Belmont, MA, USA: Athena Scientific, 1999.
- [116] L. Martignon, K. V. Katsikopoulos, and J. Woike, “Categorization With Limited Resources: A Family of Simple Heuristics”, *Journal of Mathematical Psychology*, vol. 52, no. 6, pp. 352-361, 2008.
- [117] Wikipedia. (Dec. 4, 2016). *List of knapsack problems* [Online]. Available: https://en.wikipedia.org/wiki/List_of_knapsack_problems
- [118] D. Bertsimas, and J. N. Tsitsiklis, *Introduction to Linear Optimization*. Belmont, MA, USA: Athena Scientific, 1997.

- [119] G. Sierksma, *Linear and Integer Programming: Theory and Practice (2nd ed.)*. New York, NY, USA: Marcel Dekker Inc., 2002.
- [120] Wikipedia. (Dec. 30, 2016). *Linear programming* [Online]. Available: https://en.wikipedia.org/wiki/Linear_programming
- [121] W. Tushar, J. A. Zhang, D. B. Smith, H. V. Poor, and S. Thiebaux, "Prioritizing Consumers in Smart Grid: A Game Theoretic Approach", *IEEE Transactions on Smart Grid*, vol. 5, no. 3, pp. 1429-1438, May 2014.
- [122] L. Chen, N. Li, S. H. Low, and J. C. Doyle, "Two Market Models for Demand Response in Power Networks", *Proc. IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pp. 397-402, Gaithersburg, Maryland, USA, October 2010.
- [123] W. Saad, Z. Han, H. V. Poor, and T. Basar, "Game Theoretic Methods for the Smart Grid", *IEEE Signal Processing Magazine, Special issue on Signal Processing for the Smart Grid*, vol. 29, no. 5, pp. 86-105, September 2012.
- [124] S. Jain, B. Narayanaswamy, and Y. Narahari, "A Multiarmed Bandit Incentive Mechanism for Crowdsourcing Demand Response in Smart Grids", *Proc. 28th National Conference on Artificial Intelligence (AAAI)*, pp. 721-727, Québec City, Canada, July 2014.
- [125] D. Kahneman, P. Slovic, and A. Tversky, *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge, UK: Cambridge University Press, April 1982.
- [126] G. Gigerenzer and R. Selten, *Bounded Rationality: The Adaptive Toolbox*. Cambridge, MA, USA: The MIT Press, 2001.
- [127] G. Gigerenzer, "Fast and Frugal Heuristics: The Tools of Bounded Rationality", In D. J. Koehler and N. Harvey (Eds.), *Blackwell Handbook of Judgment and Decision Making* (pp. 62-88). Oxford, UK: Blackwell, 2004.
- [128] B. Marsh, P. M. Todd, and G. Gigerenzer, "Cognitive Heuristics: Reasoning the Fast and Frugal Way", In J. P. Leighton and R. J. Sternberg (Eds.), *The Nature of Reasoning* (pp. 273-287). Cambridge, UK: Cambridge University Press, 2004.
- [129] J. E. Fischer *et al.*, "Using Simple Heuristics to Target Macrolide Prescription in Children with Community-Acquired Pneumonia," *Archives of Pediatric and Adolescent Medicine*, vol. 156, pp. 1005-1008, October 2002.
- [130] M. A. Jenny, T. Pachur, S. L. Williams, E. Becker, and J. Margraf, "Simple Rules for Detecting Depression", *Journal of Applied Research in Memory and Cognition*, vo. 2, no. 3, pp. 149-157, 2013.

- [131] H. Brighton, “Robust Inference With Simple Cognitive Models”, *Proc. AAAI Spring Symposium: Between a Rock and a Hard Place: Cognitive Science Principles Meet AI-Hard Problems*, pp. 17-22, Palo Alto, CA, USA, March 2006.
- [132] K. B. Laskey and L. Martignon, “Comparing Fast and Frugal Trees and Bayesian Networks for Risk Assessment”, *Proc. 9th International Conference on Teaching Statistics (ICOTS)*, Flagstaff, Arizona, USA, July 2014.
- [133] V. DeMiguel, L. Garlappi, and R. Uppal, “Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy?”, *The Review of Financial Studies*, vol. 22, no. 5, pp. 1915-1953, 2009.
- [134] R. M. Hogarth and N. Karelaia, “Heuristic and Linear Models of Judgment: Matching Rules and Environments”, *Psychological Review*, vol. 114, no. 3, pp. 733-758, 2007.
- [135] G. Gigerenzer and D. Goldstein, “Reasoning the Fast and Frugal Way: Models of Bounded Rationality”, *Psychological Review*, vol. 103, no. 4, pp. 650-669, 1996.
- [136] G. Gigerenzer, P. M. Todd and the ABC Research Group, *Simple Heuristics That Make Us Smart* (pp. 3-34). New York: Oxford University Press, 1999.
- [137] R. M. Hogarth and N. Karelaia, “Ignoring Information in Binary Choice With Continuous Variables: When is Less “More”?”, *Journal of Mathematical Psychology*, vol. 49, no. 2, pp. 115-124, April 2005.
- [138] N. D. Sintov and P. W. Schultz, “Unlocking the Potential of Smart Grid Technologies with Behavioral Science”, *Frontiers in Psychology*, vol. 6, article 410, April 2015.
- [139] B. Brohmann, S. Heinzle, J. Nentwich, U. Offenberger, K. Rennings, J. Schleich, and R. Wüstenhagen, “Sustainable Energy Consumption and Individual Decisions of Consumers - Review of the Literature and Research Needs”, Working Paper, January 2009.
- [140] E. R. Frederiks, K. Stenner, and E. V. Hobman, “Household Energy Use: Applying Behavioural Economics to Understand Consumer Decision-Making and Behaviour”, *Renewable and Sustainable Energy Reviews*, vol. 41, pp. 1385-1394, January 2015.
- [141] D. Litvine and R. Wüstenhagen, “Helping “Light Green” Consumers Walk the Talk: Results of a Behavioural Intervention Survey in the Swiss Electricity Market”, *Ecological Economics*, vol. 70, no. 3, pp. 462-474, January 2011.

- [142] M. Karaliopoulos, I. Koutsopoulos, and M. Titsias, “First Learn then Earn: Optimizing Mobile Crowdsensing Campaigns Through Data-driven User Profiling”, (*Mobihoc*), Paderborn, Germany, July 2016.
- [143] M. Karaliopoulos, L. Spiliopoulos, and I. Koutsopoulos, “Factoring the Human Decision-Making Limitations in Mobile Crowdsensing”, Proc. Wireless On-demand Network Systems and Services Conference (WONS), Cortina D’ Ampezzo, Italy, 2016.
- [144] J. Yuan Yu, “Crowdsourced Forecasting of Electricity Use - A Solution Proposal for Smart Cities Energy Working Group”, October 2012. Available at: <https://eu-smartcities.eu/sites/all/files/docs/best-practice/smart-cities-sp.pdf>
- [145] Wikipedia. (Dec. 30, 2016). *Gradient descent* [Online]. Available: https://en.wikipedia.org/wiki/Gradient_descent
- [146] Wikipedia. (Dec. 21, 2016). *Sensitivity and specificity* [Online]. Available: https://en.wikipedia.org/wiki/Sensitivity_and_specificity
- [147] D. M. W. Powers, “Evaluation: From Precision, Recall and F-Measure to ROC, Informedness, Markedness & Correlation”, *Journal of Machine Learning Technologies*, vol. 2, no. 1, pp. 37-63, 2011.
- [148] W.J. Youden, “Index for Rating Diagnostic Tests”, *Cancer*, vol. 3, no. 1, pp.32-35, 1950.
- [149] Wikipedia. (Nov. 22, 2016). *Youden’s J statistic* [Online]. Available: https://en.wikipedia.org/wiki/Youden’s_J_statistic
- [150] Wikipedia. (Dec. 9, 2016). *Contingency table* [Online]. Available: https://en.wikipedia.org/wiki/Contingency_table
- [151] Wikipedia. (Oct. 11, 2016). *Confusion matrix* [Online]. Available: https://en.wikipedia.org/wiki/Confusion_matrix
- [152] R Foundation. (Oct. 8, 2016). *FFTrees: Generate, Visualise, and Compare Fast and Frugal Decision Trees* [Online]. Available: <https://cran.r-project.org/web/packages/FFTrees/index.html>
- [153] H. Stanislaw and N. Todorov, “Calculation of Signal Detection Theory Measures”, *Behavior Research Methods, Instruments, & Computers*, vol. 31, n0. 1, pp. 137-149, 1999.
- [154] Analyse-it Software Ltd. (Jan. 8, 2017). *ROC plot* [Online]. Available: <https://analyse-it.com/docs/user-guide/diagnosticperformance/rocplot>