Coordinated Aggregation of Distributed Demand-Side Resources

Final Project Report

Power Systems Engineering Research Center

Empowering Minds to Engineer
the Future Electric Energy System

Coordinated Aggregation of Distributed Demand-Side Resources

Final Project Report

Project Team

Eilyan Bitar, Project Leader
Cornell University

Kameshwar Poolla and Pravin Varaiya
University of California at Berkeley

PSERC Publication 14-12

December 2014

For information about this project, contact:

Eilyan Bitar
School of Electrical and Computer Engineering
Cornell University
326 Rhodes Hall
Ithaca, NY 14853-5401
Email: eyb5@cornell.edu
Phone: 607-255-7156

Power Systems Engineering Research Center

The Power Systems Engineering Research Center (PSERC) is a multi-university Center conducting research on challenges facing the electric power industry and educating the next generation of power engineers. More information about PSERC can be found at the Center's website: http://www.pserc.org.

For additional information, contact:

Power Systems Engineering Research Center
Arizona State University
527 Engineering Research Center
Tempe, Arizona 85287-5706
Phone: 480-965-1643
Fax: 480-965-0745

Notice Concerning Copyright Material

PSERC members are given permission to copy without fee all or part of this publication for internal use if appropriate attribution is given to this document as the source material. This report is available for downloading from the PSERC website.

© 2014 Cornell University. All rights reserved

Acknowledgements

This is the final report for the Power Systems Engineering Research Center (PSERC) research project titled "Coordinated Aggregation of Distributed Demand-Side Resources" (project S-52). We express our appreciation for the support provided by PSERC's industry members. In particular, we wish to thank the following industry members for their critical feedback and support during the project:

- Bahman Daryanian (GE Energy)
- But-Chung Chiu (ALSTOM Grid Inc.)
- Bill Timmons (WAPA)
- Jim Price (CAISO)
- Heather Sanders (CAISO)
- Peter Klauer (CAISO)
- Benjamin Kroposki (NREL)
- Marissa Hummon (NREL)

Executive Summary

The legacy electricity system will involve into a *grid with intelligent periphery (GrIP)*. There will be IP addressable sensors everywhere, including every appliance, generator, storage resource, and key nodes on electric transmission and distribution systems. These sensors connect to the internet. Data from them are available anywhere, anytime. The distribution infrastructure will partially support reconfigurable bi-directional power flow through intelligent switches. New expansion will take place in the periphery. This expansion will be subject to local community control. It includes renewable micro-generation, storage, combined heat and power, and highly adjustable demand. More power will be generated locally and consumed locally. Largely self-sufficient communities, or balanced clusters, that generate and consume most of their power locally will emerge. In the following research summary, we describe a new architecture to enable this vision of a *grid with an intelligent periphery* (GrIP). We design novel coordination strategies, analyze their benefit, and simulate their performance. Our research demonstrates systemic value in managing resource clusters, and offers a novel market mechanism design to monetize this value.

In **Chapter 1**, we offer a detailed description of the *GrIP architecture* and the various actors involved in its operation.

In **Chapter 2**, we outline a systematic procedure to construct control-oriented models of demand-side resource flexibility for both *deferrable loads* and *Thermostatically Controlled Loads* (TCLs). It is widely accepted that TCLs can be used to provide regulation reserve to the grid. We first argue that the aggregate flexibility offered by a collection of TCLs can be succinctly modeled as a stochastic battery with dissipation. We next characterize the power limits and energy capacity of this battery model in terms of TCL parameters and random exogenous variables such as ambient temperature and user-specified set-points. In the proceeding chapter, we then describe a direct load control architecture for regulation service provision where we use a priority-stack-based control framework to select which TCLs to control at any time. The control objective is for the aggregate power deviation from baseline to track an automatic generation control signal supplied by the system operator.

In **Chapter 3**, we describe a direct load control architecture for regulation service provision where we use a priority-stack-based control framework to select which of the Thermostatically Controlled Loads) TCLs to control at any time. The control objective is for the aggregate power deviation from baseline to track an automatic generation control signal supplied by the system operator. Simulation studies suggest the proposed methodology works well in practical settings.

In **Chapter 4**, we describe a novel forward market for differentiated electric power services, where consumers consent to deferred service of pre-specified loads in exchange for a reduced per-unit price for energy. The longer a consumer is willing to defer, the larger the reduction in price. The proposed forward contract provides a guarantee on the aggregate quantity of energy to be delivered by a consumer-specified deadline. Under the earliest-deadline-first (EDF) scheduling policy, which is shown to be optimal for the supplier, we explicitly characterize a deadline-differentiated pricing scheme that yields an efficient competitive equilibrium between supply and demand. We further show that this efficient pricing scheme, in combination with EDF scheduling, is incentive compatible (IC) in that every consumer would like to reveal her true deadline to the supplier, regardless of the actions taken by other consumers.

Project Publications

Doctoral Theses

1. Anand Subramanian, *Mitigating Renewable Variability through Control and Optimization Techniques.* PhD Dissertation, University of California, Berkeley, May 2013.

Journal Publications

1. E. Bitar, Y. Xu. *Deadline Differentiated Pricing of Delay-Tolerant Demand.* Submitted, 2014.
2. S. Bose, E. Bitar. *The Locational Marginal Value of Energy Storage.* submitted to IEEE Transactions on Automatic Control, 2014.
3. H. Hao, B. Sannadaji, K. Poolla, T. Vincent. *Aggregate Flexibility of Thermostatically Controlled Loads.* IEEE Transactions on Power Systems, to appear, 2014.
4. R. Louca, E. Bitar. *Nondegeneracy and Inexactness of Semidefinite Relaxations of Optimal Power Flow.* submitted to IEEE Transaction on Power Systems, 2014.
5. D. Munoz-Alvarez, E. Bitar. *Financial Storage Rights in Electric Power Networks.* submitted to IEEE Transaction on Power Systems, 2014.
6. Ashutosh Nayyar, Matias Negrete-Pincetic, Kameshwar Poolla and Pravin Varaiya. *Duration-differentiated Energy Services with a Continuum of Loads.* submitted to IEEE Transaction on Networked Systems, 2014.
7. Ashutosh Nayyar, Matias Negrete-Pincetic, Kameshwar Poolla, Pravin P. Varaiya. *Rate-constrained Energy Services: Allocation Policies and Market Decisions.* submitted to IEEE Transactions on Automatic Control, 2014.
8. A. Subramanian, M. Garcia, D. Callaway, K. Poolla, and P. Varaiya. *Real-time Scheduling of Distributed Resources.* IEEE Trans. on Smart Grid, vol.4, no. 4, pp: 2122 - 2130, 2013.

Conference Publications

1. E. Bitar and S. Low. *Deadline-Differentiated Pricing of Deferrable Electric Power Service.* Proc. of the 51st IEEE Conference on Decision and Control, Invited Paper, Hawaii, 2012.

2. E. Bitar and Y. Xu. *On Incentive Compatibility of Deadline Differentiated Pricing for Deferrable Demand.* Proceedings of the 52nd IEEE Conference on Decision and Control, Invited Paper, Florence, Italy, 2013.

3. E. Bitar and K. Poolla. *Selling Wind Power in Electricity Markets: The status today, the opportunities tomorrow.* Proc. of the IEEE American Control Conf., Montreal, 2012.

4. H. Hao, B. Sanandaji, K. Poolla, and T. Vincent. *Frequency Regulation from Flexible Loads: Potential, Economics, and Implementation.* Proceedings of the American Control Conference, Portland, June 2014.

5. H. Hao, B. Sannadaji, K. Poolla, T. Vincent. *A Generalized Battery Model of a Collection of Ther- mostatically Controlled Loads for Providing Ancillary Service.* 51st Annual Allerton Conference on Communication, Control, and Computing, 2013.

6. W. Lin, E. Bitar. *Forward Electricity Markets with Uncertain Supply: Cost Sharing and Efficiency Loss.* IEEE Proc. of the CDC, Los Angeles, CA, 2014.

7. P. de Martini, A. Wierman, S. Meyn and E. Bitar. *Integrating Distributed Energy Resource Pricing and Control.* Proceedings of CIGRE USNC Grid of the Future Symposium, 2012.

8. D. Munoz-Alvarez, E. Bitar. *Financial Storage Rights: Basic Definition and Properties.* IEEE Proceedings of the North American Power Symposium, Pullman, WA, 2014.

9. D. Munoz-Alvarez, J. Wang, E. Bitar, L. Tong. *Piecewise Affine Policies for Multi-Period Economic Dispatch Under Uncertainty.* Proceedings of the Power and Energy Society General Meeting, Washington D.C., 2014.

10. A. Nayyar, J. Taylor, A. Subramanian, K. Poolla, and P. Varaiya. *Aggregate Flexibility of a Collection of Loads.* Proceedings of the Conference on Decision and Control, pp:5601-7, Florence, Italy, December 2013.

11. A. Nayyar, K. Poolla, and P. Varaiya. *A statistically robust payment sharing mechanism for an aggregate of renewable energy producers.* European Control Conference, pp:3025 - 3031, Zurich, July 2013.

12. B. Sanandaji, H. Hao, and K. Poolla. *Fast Regulation Service Provision via Aggregation of Thermostatically Controlled Loads.* Hawaii International Conference on System Sciences, pp: 2388 - 2397, Kona, January 2014. Best paper prize finalist.

13. A. Subramanian, J. Taylor, E. Bitar, D. Callaway, K. Poolla, and P. Varaiya. *Optimal Power and Reserve Capacity Procurement Policies with Deferrable Loads.* Proceedings of the Conference on Decision and Control, pp:450-456, Maui, December 2012.

Contents

Acknowledgements i

Executive Summary ii

Project Publications iv

Contents vii

List of Figures x

List of Tables xi

1 The Role of Flexible Demand-Side Resources 1
- 1.1 An Example of the Power of Coordinated Aggregation . . . 2
- 1.2 The Grid with Intelligent Periphery (GrIP) . . . 2
  - 1.2.1 Clusters . . . 4
  - 1.2.2 Architecture . . . 4
  - 1.2.3 Information Flow in Wide-Area Control Loops . . . 5

2 Modeling Demand-Side Flexibility 6
- 2.1 Thermostatically Controlled Loads . . . 6
  - 2.1.1 Main Contributions . . . 8
  - 2.1.2 Related Work . . . 9
  - 2.1.3 Individual TCL Models . . . 10
    - 2.1.3.1 Dead-Band Model . . . 10
    - 2.1.3.2 Continuous Power Model . . . 11
- 2.2 Deferrable Loads . . . 12
  - 2.2.1 Introduction . . . 12
  - 2.2.2 Problem Formulation . . . 14
  - 2.2.3 Generation Modeling . . . 14
    - 2.2.3.1 Renewable generation (w_k) . . . 15
    - 2.2.3.2 Bulk power (B_k) . . . 15
    - 2.2.3.3 Reserve generation (r_k) . . . 15

2.2.4 Resource Modeling . . . 15
2.2.4.1 Static Loads . . . 16
2.2.4.2 Deferrable Loads . . . 16
2.2.4.3 Storage . . . 17
2.2.5 Scheduling Policies . . . 17
2.2.6 Performance Metric . . . 18
2.2.7 Policy Optimization Problem . . . 19
2.3 Generalized Battery Models of Aggregate Flexibility . . . 20

3 Controlling Demand-Side Flexibility 24
3.1 Thermostatically Controlled Loads . . . 24
3.1.1 Priority Stack Scheduling . . . 25
3.1.2 Practical Considerations . . . 27
3.1.3 Simulation Studies . . . 28
3.1.4 Tracking Performance . . . 28
3.1.5 Battery Model Conservatism . . . 31
3.1.6 The Effect of Ambient Temperature . . . 33
3.1.7 Conclusions and Future Work . . . 34
3.2 Deferrable Loads . . . 35
3.2.1 Earliest-Deadline-First (EDF) Scheduling . . . 35
3.2.2 Least-Laxity-First (LLF) Scheduling . . . 36
3.2.3 Reserve scheduling: Zero - Laxity . . . 37
3.2.4 Receding Horizon Control (RHC) . . . 38
3.2.4.1 Decision Variables . . . 38
3.2.4.2 Objective Function . . . 39
3.2.4.3 RHC Optimization Problem . . . 39
3.2.4.4 Constraints . . . 40
3.2.5 Simulation Study . . . 40
3.2.6 Parameters . . . 41
3.2.6.1 Load . . . 41
3.2.6.2 Generation . . . 41
3.2.6.3 DA Market Clearing . . . 42
3.2.7 Comparison of Scheduling Algorithms . . . 42
3.2.8 Deferrable load penetration . . . 43
3.2.9 Load deferrability . . . 44
3.2.10 Practical considerations . . . 45
3.2.11 Conclusions and Future Work . . . 46

4 A Market for Flexibility: Deadline Differentiated Pricing 48
4.1 A Market for Quality-Differentiated Electric Power Service . . . 49
4.1.1 Summary and Contributions . . . 51
4.1.2 Relation to the Literature . . . 52
4.2 Delay-Tolerant Demand Model . . . 54
4.2.1 Consumer Surplus . . . 57

4.3 Supply Model . . . 60
4.3.1 Optimal Scheduling . . . 61
4.3.1.1 Inter-class scheduling policies . . . 61
4.3.1.2 Intra-class scheduling policies . . . 62
4.3.1.3 Supplier profit . . . 63
4.3.1.4 Optimal scheduling policy . . . 64
4.3.2 Marginal Cost Pricing . . . 64
4.4 Incentive Compatibility and Market Equilibrium . . . 67
4.4.1 Incentive Compatibility . . . 68
4.4.2 Implementation and Market Equilibrium . . . 69
4.4.2.1 A mechanism design approach . . . 69
4.4.2.2 A market equilibrium approach . . . 70
4.5 Conclusion . . . 72

A Proofs for Chapter 2 **74**
A.1 Proof of Theorem 2.12 . . . 74
A.2 Proof of Theorem 2.13 . . . 75

B Proofs for Chapter 4 **77**
B.1 Intra-class Scheduling Policies . . . 77
B.2 Proof of Theorem 4.6 . . . 78
B.3 Proof of Lemma 4.8 . . . 83
B.4 Proof of Theorem 4.14 . . . 84

Bibliography **88**

List of Figures

1.1 A stylized example demonstrating the potential of coordinated aggregation to absorb variability from renewable generation. 2
1.2 Layered Architecture. 3
1.3 Data flow. 3
1.4 Distributed Control Loops. 3

3.1 Control architecture for regulation service provisoin. 25
3.2 The ON and OFF priority stacks. The stacks are sorted by temperature distance. . . 26
3.3 Tracking of a regulation signal succeeds when it is within the power limits and energy capacity of the (sufficient) battery model $\mathbb{B}(\phi_1)$. 29
3.4 Tracking of a regulation signal fails when it exceeds the power limit of the (necessary) battery model $\mathbb{B}(\phi_2)$. 30
3.5 Tracking of a regulation signal fails when it exceeds the energy capacity of the (necessary) battery model $\mathbb{B}(\phi_2)$. 31
3.6 Tracking of a typical 6-hour regulation signal from PJM that is within the power limits and energy capacity of the (sufficient) battery model $\mathbb{B}(\phi_1)$. 32
3.7 Conservatism in the bounds on energy capacity from the necessary and sufficient models. The green dots show the "true" energy capacity as calculated numerically. The red bars extend between the energy capacity bounds given by the sufficient and necessary battery models. 33
3.8 Participation percentage vs. ambient temperature. 34
3.9 Power limits and effective capacities vs. ambient temperature. 35
3.10 Load profiles comparing resource scheduling under EDF, LLF, and RHC to no scheduling base case 43
3.11 Percentage reductions in up reserve cost achieved by RHC-based scheduling at various levels of deferrable load penetration (α) 44
3.12 Percentage reductions in up reserve cost achieved by RHC-based scheduling at various levels of task deferrability (ϕ) 45
3.13 Cumulative distribution function on the maximum number of on-off switches under RHC-based scheduling. 46

4.1 Plot of four utility functions that satisfy Assumption 4.2.1. 56

List of Tables

2.1 Typical parameter values for a residential AC unit. 11
2.2 Interpretation of parameters ϕ. 20
2.3 Comparison of generalized battery models for a collection of heterogeneous Thermostatically Controlled Loads (TCLs) 23

3.1 Power limits and energy capacities. 28
3.2 Prediction performance of the sufficient battery model 33
3.3 Comparison of scheduling algorithms showing percentage reductions in 4 reserve cost metrics (compared to no coordination baseline). 43
3.4 Comparison of scheduling algorithms showing computation time (s) required for different levels of deferrable load penetration (α). 46

Chapter 1

The Role of Flexible Demand-Side Resources

As the electric power industry transitions to a greater reliance on intermittent and distributed energy resources, there is an increasing need for flexible resources that can respond dynamically to weather impacts on wind and solar photovoltaic output. These renewable generation sources have limited controllability and production patterns that are intermittent and uncertain. Such variability represents one the most important obstacles to the deep integration of renewable generation into the electricity grid. The current approach to renewable energy integration is to balance variability with dispatchable generation. This works at today's modest penetration levels, but it cannot scale, because of the projected increase in reserve generation required to balance the variability in renewable supply [12]. If these increases are met with combustion fired generation, they will both be counterproductive to carbon emissions reductions and economically untenable.

As wind and solar energy penetration increases, how must the assimilation of this variable power evolve, so as to minimize these integration costs, while maximizing the net environmental benefit? Clearly, strategies which attenuate the increase in conventional reserve requirements will be an essential means to this end. One option is to harness the flexibility in consumption on the demand side. As such, significant benefits have been identified by the Federal Energy Regulatory Commission [18] in unlocking the value in coordination of demand-side resources to address the growing need for firm, responsive resources to provide supply-demand balancing services (ancillary services) for the bulk power system.

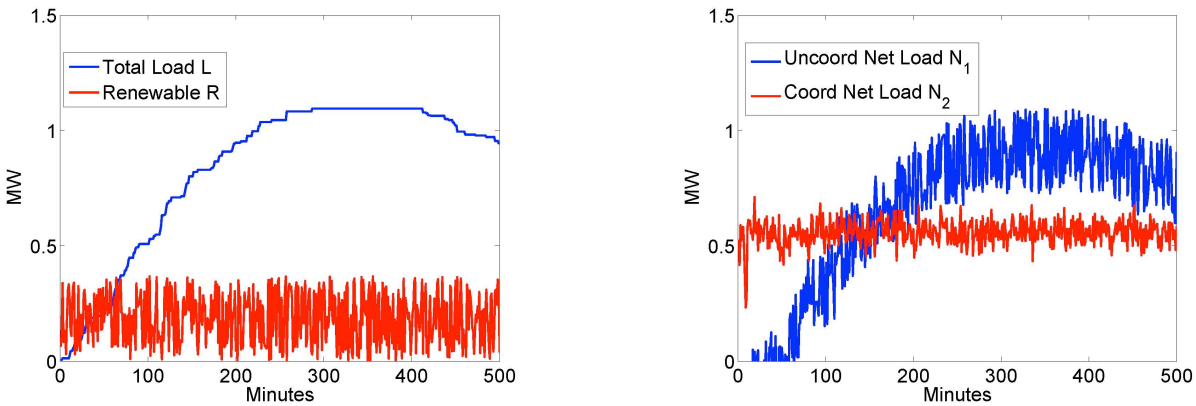

FIGURE 1.1: A stylized example demonstrating the potential of coordinated aggregation to absorb variability from renewable generation.

1.1 An Example of the Power of Coordinated Aggregation

Consider 100 EVs being charged over 8 hours. Maximum charging rates were 1.65KW. Each EV has different energy needs over different service intervals. Service intervals were chosen randomly with Gaussian arrivals/departures centered at 6pm/8am with 1.5 hour standard deviation. The residual energy charge of the EV batteries was uniformly distributed between 5% and 95% of capacity. The uncoordinated total load under uniform charging presented by these vehicles is the curve L. Renewable generation R services 25% of the total energy needs of the EVs. The remaining net load N1 = L - R is supplied by the core grid. If coordinated aggregation is used under the Earliest Deadline First (EDF) scheduling policy, the loads absorb much of the variability. The remaining coordinated net load N2 must be scheduled by the ISO. Notice in the above figure how dramatically coordinated aggregation reduces the net load variability presented to the core grid . A more nuanced simulation would reveal that coordinated aggregation improves net load predictability, reducing the necessary operating reserve capacity.

1.2 The Grid with Intelligent Periphery (GrIP)

The legacy electricity system will involve into a *grid with intelligent periphery*. There will be IP addressable sensors everywhere, including every appliance, generator, storage resource, and key nodes on electric transmission and distribution systems. These sensors connect to the internet. Data from them are available anywhere, anytime. The distribution infrastructure will partially support reconfigurable bi-directional power flow through intelligent switches. New expansion will take place in the periphery. This expansion will be subject to local community control. It includes

renewable micro-generation, storage, combined heat and power, and highly adjustable demand. More power will be generated locally and consumed locally. Largely self-sufficient communities, or balanced clusters, that generate and consume most of their power locally will emerge.

Information and decisions in today's grid are highly centralized [77]. Data such as power demand, generator supply bids, and status of the grid infrastructure flow to a central system operator, where they are processed by optimization algorithms that calculate control commands to be issued in return. While sophisticated in design and occurring on multiple time scales, from hourly generation schedules to second-by-second load following and frequency regulation, this control regime is essentially unidirectional and addresses a limited number of points. How will this change? With the number and significance of peripheral resources [including small renewable generators, storage devices, controllable loads, and network protection devices] growing into the hundreds of thousands of relevant sensing and control nodes, the centralization of information flow and decision-making becomes impractical for reasons of sheer scale, as well as resilience. Decision authority and information flow both must be distributed.

How can this dual architecture of distributed decision and information flow be designed? Decentralized Smart Grid control architectures are not new [78–80]. Our approach suggests the clustering of resources, the functional organization of entities, and the factorization of information flow into cyber information services and physical system application clients. Our research aims to illuminate this architecture and supply detailed algorithms resident in these services and clients that explicitly account for variability.

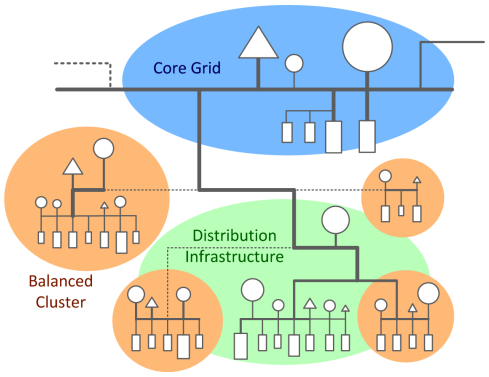

FIGURE 1.2: Layered Architecture.

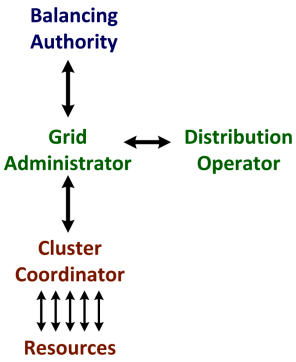

FIGURE 1.3: Data flow.

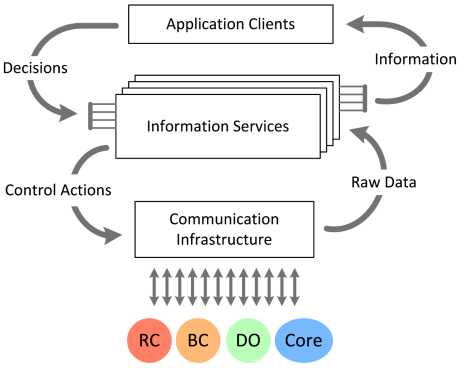

FIGURE 1.4: Distributed Control Loops.

1.2.1 Clusters

Peripheral resources and devices can be grouped functionally. *A resource cluster* [RC] is a [geographically dispersed, but usually homogenous] collection of resources that supply a service [ex: storage, demand response]. Examples include micro-grids, community solar farms [81], and commercial refrigeration resources offering demand response that are managed collectively [82].

A *Balanced Cluster* [BC] is a [local, but diverse] collection of resources that largely achieve internal power balance. Any remaining mismatch between power consumption and production may be met by contracted imports and exports. Balanced cluster aggregate their resources to present reduced and managed variability to the grid. Thus, they mirror on a small scale the function of a larger grid balancing authority. Current examples: zero-net-energy buildings [83] and Community Choice Aggregation jurisdictions [84]. BCs relieve the core grid of the need to acquire reserves to counter internal demand and supply variations. Further, a BC may increase overall system resilience if it can be disconnected [islanded] without affecting grid stability.

1.2.2 Architecture

How can the resource clusters aggregate their capabilities? How can a resource cluster allocate its resources to meet a service request? How can balanced clusters best be managed in order to benefit the overall grid? How do we coordinate power balance with balanced clusters? Optimal solutions to these question will require an *architecture to organize data acquisition and decision making.* Our proposed GrIP architecture is depicted in Figure 1.2 and includes:

- *Cluster Coordinators* [CC] manage resource clusters in order to meet a variety of objectives [ex: conservation, increase profit, or promote a green mandate by maximizing renewable generation]. To maintain internal balance, the CC must have sufficient control over each resource and device within the cluster; in reality, the degree of control authority will vary.
- *Grid Administrators* [GA] represent clusters in a geographical region to a larger balancing area authority or system operator of the core transmission system. The GA may enter into contracts as a supplier of generation or ancillary services [ex: storage, reserves] to the grid, as a load-serving entity, or as a dispatchable load, by combining the offers of its constituent clusters.
- *The Distribution Operator* [DO] services the local infrastructure in the legacy grid. The DO may lease lines, but is not usually an active market participant.

What operational flexibilities does this architecture afford? What are its advantages and shortcomings? A key function of the cluster coordinator is to prevent potential conflicts that arise in peer-to-peer interactions. The intermediary role of the grid administrator is to present a single point of contact - thus relieving the communications burden - to the system operator responsible for wide-area stability across an extensive geographic region. Most importantly, we assert that this layered architecture allows us to realize the core power system objectives of resilience and efficiency with diverse, distributed, uncertain resources. *Our research aims to comprehensively justify this claim through analysis, simulation, and experiment.*

1.2.3 Information Flow in Wide-Area Control Loops

Resources and devices transmit data about their current state and forecast their future state to CCs, who run Information Services to achieve the cluster's control objectives. This represents a dynamic optimization problem under uncertainty. The resulting information is sent to the GA, who aggregates the forecasts of the individual clusters and executes Application Clients. As discussed below, these applications include resource management as well as infrastructural operations such as network reconfiguration, which the GA undertakes in cooperation with the DO. The layered information architecture provides the information appropriate to the decisions made by CC's and GA's. The transformation from data to commands forms a feedback loop as shown in Figure 1.4.

How is this information flow optimally managed? In a naïve implementation there would be hundreds of thousands of loops. We propose to factor these loops into common information services that process the data into a set of sufficient statistics which summarize the state of the peripheral devices. These statistics are used by CC's and GA's to run applications that produce the commands to actuate the resources. *How are information services defined?* A crucial research problem is to define minimal sets of information services and the sub-routines that produce them, and applications and associated algorithms.

Chapter 2

Modeling Demand-Side Flexibility

It is widely accepted that Thermostatically Controlled Loads (TCLs) can be used to provide regulation reserve to the grid. In this chapter, we first argue that the aggregate flexibility offered by a collection of TCLs can be succinctly modeled as a *stochastic battery* with dissipation. We next characterize the power limits and energy capacity of this battery model in terms of TCL parameters and random exogenous variables such as ambient temperature and user-specified set-points. In the next chapter, we then describe a direct load control architecture for regulation service provision where we use a priority-stack-based control framework to select which TCLs to control at any time. The control objective is for the aggregate power deviation from baseline to track an automatic generation control signal supplied by the system operator.

2.1 Thermostatically Controlled Loads

Ancillary services provide for resources to handle supply-demand imbalances at various time-scales, maintain power quality, and assure reliable power delivery under contingencies. There are many different ancillary services. Of these, regulation reserve (or frequency regulation) and load-following are the key services for maintaining the power balance under normal operating conditions. Load following is a feed-forward approach and handles predictable and slower changes in load. Regulation is a feedback strategy that mitigates faster and unpredictable changes in system load and corrects unintended fluctuations in generation [49].

It is widely accepted that the deep penetration of renewable generation will substantially increase the need for ancillary services [50–52]. Recent studies [53] project that the spring time maximum up-regulation reserve needed to accommodate California's 33% renewable penetration mandate will

increase from 277 MW to 1,135 MW. Similar increases in down-regulation reserve are projected. The maximum load-following requirement will need to increase from 2.3 GW to 4.4 GW. If these additional ancillary services are provided by traditional fossil fuel generators, it will diminish the net carbon benefit from renewables, reduce generation efficiency, and be economically untenable.

There is an emerging consensus that demand side resources must play a key role in supplying zero-emission regulation service that is necessary for deep renewable integration. These include TCLs, electric vehicles, and strategic storage. Existing programs, such as the SmartAC^{TM} program of Pacific Gas and Electric (PG&E), aggregate residential Air Conditioners (ACs) for peak load shaving and emergency load management [54]. Because these load control mechanisms are primarily concerned with very low frequency changes in demand (i.e., the changes occur over hours timescale), they are invoked infrequently and offer limited financial value. In contrast, there is an enormous untapped potential for flexible loads to offer more lucrative fast ancillary services such as frequency regulation or load-following.

Residential TCLs such as ACs, heat pumps, water heaters, and refrigerators, represent about 20% of the total electricity consumption in the United States [55, 56], and thus offer significant potential for provision of various ancillary services. TCLs have inherent thermal storage, so their electricity consumption can be modulated while still meeting the desired temperature requirements of the end user.

This chapter aims to provide a foundation for a practical method by which TCLs can be utilized to provide regulation reserve to the grid. In the scheme we consider, an aggregator would contract with owners of individual TCLs to *directly control* the electricity consumption of their units. Units owners are assured that the temperature limits they specify through thermostat set-point selection will be respected. We provide an important tool that would be useful to the aggregator: *a simple, compact battery model* that characterizes the set of power profiles that the collection of TCLs can accept while meeting their local constraints. This model depends on exogenous random processes such as ambient temperatures, user specified set-points, and unit participation. The aggregator could forecast these random processes and predict the admissible set of power profiles for the TCL collection on a future window. This flexibility in power consumption is monetized by offering it in ancillary service markets to provide regulation reserve. If the offer is accepted, the aggregator is obligated to deliver the service at delivery time. The aggregator then uses a *priority stack control strategy* to track an exogenous regulation signal supplied by the system operator.

2.1.1 Main Contributions

TCLs are flexible in the sense that a variety of power trajectories are able to meet user-specified temperature constraints. We consider a heterogeneous collection of TCLs indexed by k. We describe TCLs with a simplified continuous power model. For each TCL, we define a nominal power profile P_o^k and model its flexibility as the set $\mathbb{E}^k$ of permissible deviations from this nominal that result in temperature trajectories which respect the dead-band constraints of the unit. The aggregate flexibility $\mathbb{U}$ of the collection is the sum of the flexibilities of the individual units, i.e. $\mathbb{U} = \sum_k \mathbb{E}^k$. This set has a complex structure.

We first introduce a generalized battery model $\mathbb{B}(\phi)$ which serves as a compact, abstract representation of a set of power signals. The battery model $\mathbb{B}(\phi)$ has parameters $\phi = (C, n_-, n_+, \alpha)$, where C is the battery energy capacity, n_- and n_+ are the charge/discharge rate limits, and α is the dissipation rate. This simple model is attractive as it is defined by a few physically intuitive parameters. The main contribution of our work is to establish that the set $\mathbb{U}$ can be bounded by two generalized battery models:

$$\mathbb{B}(\phi_1) \subseteq \mathbb{U} \subseteq \mathbb{B}(\phi_2).$$

We offer explicit formulae for the battery parameters ϕ_1 and ϕ_2 in terms of individual TCL parameters and exogenous random processes such as ambient temperature and user defined set-points. We further show that the gap between these two models vanishes as the population of TCLs becomes homogeneous.

We next offer a direct-load control architecture for an aggregator to provide frequency regulation service from a collection of TCLs. This requires an *ex ante* contractual agreement with the system operator that specifies two items on the forward delivery window: (a) the "nominal" or baseline power consumption $n(t)$ of the aggregate, and (b) the class of admissible regulation signals $r(t)$. In the event the offered service is accepted, the aggregator receives an Automatic Generation Control (AGC) signal from the system operator and is obligated to ensure that the collection of TCLs consumes power profile $r(t) + n(t)$ in aggregate. The aggregator is also obligated to ensure that power profiles allocated to each TCL must respect their temperature dead-band constraints. If $r(t) \in \mathbb{B}(\phi_1)$, this is possible. If $r(t) \notin \mathbb{B}(\phi_2)$, this is not possible. Thus, the aggregator will forecast the battery parameters ϕ_1 and offer the regulation capability $\mathbb{B}(\phi_1)$ which it is assured that it can reliably deliver.

If the offered service is accepted, the aggregator receives an AGC signal $r(t) \in \mathbb{B}(\phi_1)$ from the system operator, and must devise a *run-time* strategy to allocate the power $r(t) + n(t)$ to the various TCLs. This allocation must be *causal* and meet the user-defined temperature needs of

the TCLs. We propose an allocation strategy based on priority stacks. Priorities are computed based on temperature distance to switching boundary or time to reach the switching boundary. The priority-stack-based strategy naturally accommodates operational constraints such as avoiding short-cycling on each unit, and most importantly offers a generic architecture suitable for direct load control of various flexible loads including electric vehicles and deferrable appliances. We explore some details of the practical implementations of our control strategy, and illustrate our principal ideas with simulation studies.

Our *set-theoretic characterization* of the aggregate flexibility of a collection of TCLs using battery models allows us to naturally treat regulation service provision. Ex ante, it is used to conservatively represent the set of feasible regulation signals that the aggregator can support. The resulting run-time control problem to follow the external regulation signal r becomes *extremely simple*. There is no need to account for the dynamics of the TCL population.

2.1.2 Related Work

Early work showing potential of TCLs for mitigating the variability from renewables may be found in [58]. Here, the author shows that the aggregate power consumption of a collection of TCLs can be made to follow a high-frequency signal such as the power output of a wind-farm.

There is a substantial body of work on models for TCLs and TCL aggregations [59–64]. These models fall along a spectrum of complexity and fidelity, and have been used for a variety of simulations, control design, and validation efforts. There are some papers that propose battery models to treat the aggregate flexibility for a collection of TCLs, see for example [65]. These papers do not qualify their battery models as approximations, nor do they attempt to quantify the quality of the approximation. The battery models derived usually do not include a dissipation term and therefore understate the aggregate flexibility of TCL collections. To the best of our knowledge, our set-theoretic characterization of aggregate flexibility together with quantifying the associated error using generalized battery models is new. Temperature-based and/or time-based priority control methods that are closely related or identical to this work were developed in [63, 66, 67].

Another class of models commonly used to study aggregations of TCLs are population-bin transition models, for example [63, 65]. These control-oriented models are often high order (depending on the number of bins), linear Ordinary Differential Equations (ODEs). These researchers then study regulation as a control system design problem with this model. There are three difficulties with population bin models: (a) they are complex, not universal portable, (b) they do not easily handle participation effects (i.e. modifying the model to account for individual TCLs coming on line

or dropping out), (c) they require much more complex control strategies that take into account population dynamics.

While there is considerable literature on *indirect load control* through price proxies [68, 69], it is our view that these schemes are less reliable, as resource availability and tracking accuracy of a dispatch signal are not guaranteed [70]. Moreover, with large scale deployment of price-based demand response programs, power system stability becomes a serious concern [71].

2.1.3 Individual TCL Models

2.1.3.1 Dead-Band Model

We first present a standard hybrid-system model for a TCL. We will use this model for our simulation studies. In this model, the temperature evolution of a TCL is described as

$$\dot{\theta}(t) = \begin{cases} -a(\theta(t) - \theta_a) - bP_m + w, & \text{if } q(t) = 1, \\ -a(\theta(t) - \theta_a) + w, & \text{if } q(t) = 0. \end{cases} \tag{2.1}$$

The state of the unit is captured by the binary signal $q(t)$. We say the unit is in the ON state at time t if $q(t) = 1$, and in the OFF state if $q(t) = 0$. The unit cycles between ON and OFF states when the temperature crosses user-specified temperature thresholds:

$$\lim_{\epsilon \to 0} q(t + \epsilon) = \begin{cases} q(t), & |\theta(t) - \theta_r| < \Delta, \\ 1 - q(t), & |\theta(t) - \theta_r| \geq \Delta. \end{cases}$$

Here, Δ is the dead-band, θ_r is the set-point, and the process noise w accounts for disturbances. The constants

$$a = \frac{1}{R_{th}C_{th}}, \quad b = \frac{\eta}{C_{th}},$$

can be expressed in terms of the thermal resistance R_{th}, thermal capacitance C_{th}, and coefficient of performance η. See [58, 59, 63] for more details on the model. The rated power P_m is positive for cooling devices, and it is negative for heating devices. Table 2.1 describes the parameters and their typical values for a residential AC [72]. This simple first-order hybrid model is an approximation. The cooling dynamics of west facing houses with exposed attic spaces require higher order models. A detailed and explicit treatment of modeling uncertainty in (2.1) is outside the scope of this work.

TABLE 2.1: Typical parameter values for a residential AC unit.

Parameter	Description	Value	Unit
C_{th}	thermal capacitance	2	kWh/°C
R_{th}	thermal resistance	2	°C/kW
P_m	rated electrical power	5.6	kW
η	coefficient of performance	2.5	
θ_r	temperature setpoint	22.5	°C
Δ	temperature deadband	0.3	°C

The average power consumed by a TCL over a cycle is

$$P_a = \frac{P_m T_{\text{ON}}}{T_{\text{ON}} + T_{\text{OFF}}}, \tag{2.2}$$

where T_{ON} is the ON state duration, and T_{OFF} is the OFF state duration per cycle. It is straightforward to show that

$$\begin{aligned} T_{\text{ON}} &= R_{th} C_{th} \ln \frac{\theta_r + \Delta - \theta_a + R_{th} P_m \eta}{\theta_r - \Delta - \theta_a + R_{th} P_m \eta}, \\ T_{\text{OFF}} &= R_{th} C_{th} \ln \frac{\theta_r - \Delta - \theta_a}{\theta_r + \Delta - \theta_a}. \end{aligned}$$

2.1.3.2 Continuous Power Model

As a more convenient abstraction of the dead-band model, we consider a continuous thermal model. We will use this approximate model for all our analytical work. Here, the TCL accepts any continuous power input $p(t) \in [0, P_m]$ and the dynamics are:

$$\dot{\theta}(t) = -a(\theta(t) - \theta_a) - bp(t). \tag{2.3}$$

Note that in this model, as common in the literature, the disturbance w is assumed to be zero mean [58, 59]. If w has non-zero mean, it could be treated by an appropriate change to θ_a. Maintaining the temperature $\theta(t)$ within the user-specified dead-band $\theta_r \pm \Delta$ is treated implicitly as a *constraint* on the power signal $p(t)$. When evaluating the trajectory $\theta(t)$ of this continuous model, it is assumed that $\theta(0) = \theta_r$. The parameters that specify this continuous power model are $\chi = (a, b, \theta_a, \theta_r, \Delta, P_m)$.

The nominal power required to keep a TCL at its set-point is given by

$$P_o = \frac{a(\theta_a - \theta_r)}{b} = \frac{\theta_a - \theta_r}{\eta R_{th}}. \tag{2.4}$$

Simple calculations with typical parameters reveal that the nominal power P_o under the continuous power model closely follows the average power P_a under the dead-band model (as in (2.2) for a wide range of operating conditions.

It can be shown that the continuous model is asymptotically exact for a collection of *homogeneous* TCLs represented by the hybrid model. More precisely, a load trajectory that is feasible for one is feasible for the other within a constant tolerance that decreases with the size of the collection plus a time varying tolerance that decreases to zero at least as fast as the time constant a of the TCLs. The fact that the feasible load trajectory for the collection with the hybrid model is feasible (in sum) for the continuous model follows simply by linearity. In this case, the variable θ for the continuous model is exactly the average temperature of the homogeneous TCL collection. The reverse is true because a feasible trajectory for the continuous model can be divided among a collection of hybrid units in such a way as to drive the temperatures of the hybrid units to asymptotically approach the temperature of the continuous model, and in the worst case the convergence is bounded by a term proportional to e^{-at}. For a collection of *heterogeneous* TCLs, we can cluster them into groups so that TCLs in each group have similar parameters. Each of these groups with the hybrid model can be approximated by that with the continuous model.

2.2 Deferrable Loads

In this section, we develop and analyze real-time scheduling algorithms for coordinated aggregation of deferrable loads and storage. These distributed resources offer flexibility that can enable the integration of renewable generation by reducing reserve costs. In the next chapter, we then present three scheduling policies: earliest deadline first (EDF), least laxity first (LLF), and receding horizon control (RHC).

2.2.1 Introduction

Motivated by a combination of environmental, energy security, and economic concerns, many countries have committed to substantially increasing the use of clean, renewable energy resources. This presents serious challenges to existing grid operations. Renewable sources of energy are intermittent, non-dispatchable, and uncertain. To achieve deep penetrations of renewable energy, grid operations must economically address the *variability* associated with renewable generation.

The current operating paradigm is to absorb this variability in operating reserves. This approach works at modest penetration levels, but fails when 30% or more of total energy generation comes

from renewables. A recent study indicates that California will require increases in maximum spring time load-following down-reserve capacity from 3,275 MW to 5,283 MW in order to meet its 33% renewable energy penetration target [103]. The California Independent System Operator (CAISO) forecasts comparable increases in various other reserve metrics [104]. These increases in reserves are economically untenable and defeat the environmental benefit of renewables.

Much of this renewable generation is being deployed at the distribution level with rooftop solar photovoltaics (PVs) and small-scale wind farms [88, 89]. Distributed generation along with small-scale storage, deferrable loads, and advanced power electronics are examples of distributed energy resources (DERs). Recent years have witnessed increased deployment of these distribution level assets [105, 106]. These assets offer a range of benefits for power system operation such as (a) relieving transmission and distribution network congestion, (b) improving system reliability, and (c) mitigating renewable variability. *Coordinated aggregation* refers to the intelligent control of DERs and offers the promise of effectively mitigating the increased reserve costs of deep renewable penetration.

Coordinated aggregation schemes broadly fall into two classes: (a) indirect load control (ILC) where DERs respond, in real-time, to price proxy signals that induce desired behavior, and (b) direct load control (DLC) where DERs cede physical control of devices to operators who determine appropriate actions. In both cases, the central problem is management of a *resource cluster* – a networked collection of DERs. In DLC, a *Cluster Manager* (CM) coordinates operation of the various assets present in a resource cluster. The CM participates in ex-ante markets to procure generation to meet load requirements, performs resource scheduling in real-time, and presents the capabilities of all managed DER assets to the system operator as a dispatchable resource. This hierarchical architecture is necessary as system-wide aggregation of DERs involves prohibitive communication costs [100], and falls outside the scope of system operator business models.

In this section and also in the next chapter, we focus on the CM's real-time scheduling function and analyze its impact on *operating reserves*. In this context, operating reserves refers to load-following reserves, procured in real-time energy markets, and regulation reserves on time scales of a few minutes. In particular, this does not include any contingency or ramping reserves dispatched in response to rare and severe events. Our contributions are as follows. We develop two resource scheduling algorithms (inspired by processor time allocation) – earliest deadline first (EDF) and least laxity first (LLF). We then formulate a receding horizon control (RHC) scheduling algorithm that explicitly accounts for operating reserve costs. We then present simulation studies to assess the performance of these resource scheduling policies in the metrics of reserve energy and capacity costs. We quantify the marginal benefits of deferrable load penetration and of load scheduling flexibility.

We conclude that resource scheduling under any algorithm offers compelling reserve energy cost reductions while RHC-based scheduling also offers significant reserve capacity cost reductions. We conclude that, while reserve cost reductions offered by coordinated aggregation depend on the statistical nature of the renewable generation process, *the majority of these benefits can be realized with modest levels of deferrable load penetration and load deferrability.*

The development and analysis of coordinated aggregation schemes is an active research area. Recent studies have developed algorithms [112], and quantified the numerous operational benefits [107] afforded by ILC. There have also been studies focused on the economic efficiency [108, 110] and stability of price signals [111], as well as the practical issues associated with implementing ILC programs [109]. Current research in DLC focuses on developing and analyzing algorithms for coordinated aggregation [96, 113, 114, 118]. For example, [116] develops a distributed scheduling protocol for electric vehicle charging; [115] uses approximate dynamic programming to couple wind generation with deferrable loads; and [98–100] suggest the use of RHC approaches for resource scheduling. Our contribution is not in suggesting the usage of RHC control strategies, but rather in the development of specific cost functions which explicitly minimize operating reserve costs.

2.2.2 Problem Formulation

Consider the problem of balancing a cluster's power generation and load over the operating window $[0, T]$. We assume this balancing occurs N times within this window at times indexed by $k \in \{1, 2, \ldots, N\}$. At each time k, the CM performs two functions. First, the CM allocates surplus available generation after meeting static load requirements to deferrable loads and storage. This is called *resource scheduling*. Second, the CM determines the amount of reserves required (r_k) to meet load requirements. As we ignore the impact of rare and severe events, all energy imbalances are managed using operating reserves. This amount of operating reserves is determined at each time k and is constant for a time interval $\Delta t = \frac{T}{N}$ until the next opportunity to balance load and generation. This is called *reserve scheduling*.

2.2.3 Generation Modeling

We assume the CM procures generation to meet load requirements from three sources: (a) inexpensive renewables within its cluster, (b) bulk power purchased in conventional electricity markets, and (c) operating reserves.

2.2.3.1 Renewable generation (w_k)

Renewable generation (w_k) is drawn from sources (ex: rooftop PV) within the CM's resource cluster. We assume this generation is free (zero marginal cost) but exhibits significant variability. The CM accepts all power generated by renewable sources of energy within its resource cluster. Let $\mathbf{w} = \{w_k\}_{k=1}^N$ denote the realized sequence of renewable generation in the operating interval.

2.2.3.2 Bulk power (B_k)

Bulk power (B_k) is acquired from the bulk transmission system contracted ex-ante in a forward market. For each hour-long interval, the CM receives a constant amount of power. There is no uncertainty associated with bulk power delivery. Let $\mathbf{B} = \{B_k\}_{k=1}^N$ denote the sequence of bulk power procurements. We assume $\mathbf{B}$ is *constant* over hour-long intervals and separate amounts of bulk power are purchased for each such interval within the operating window.

2.2.3.3 Reserve generation (r_k)

Reserve generation (r_k) is dispatched to ensure load requirements of the resource cluster are met at each balancing time k. Reserve generation can be positive (up regulation) or negative (down regulation). Let $\mathbf{r} = \{r_k\}_{k=1}^N$ refer to the sequence of reserve dispatch decisions made during the operating interval.

There are two cost components associated with reserve generation: energy and capacity. We assume the CM pays a price p_r (\$/MWhr) for each unit of reserves (up or down) dispatched to perform load balancing (energy cost) and a price p_C (\$/MW) for the maximum reserves dispatched at any balancing time (capacity cost). Our results can be generalized to the case of asymmetric reserve costs.

2.2.4 Resource Modeling

A CM's resource cluster will contain heterogeneous loads. These may include: (a) power on-demand loads which offer no scheduling flexibility, (b) deferrable loads such as plug-in electric vehicles (EVs) and thermostatically controlled loads (TCLs), and (c) distributed electricity storage. Detailed modeling of loads to varying degrees of fidelity is an active research area [85, 86]. There are three kinds of loads:

2.2.4.1 Static Loads

Static Loads have no flexibility associated with their admissible power profile. The CM must ensure that adequate generation is present to meet this load requirement at each time k within the operating interval $[0, T]$. We model the aggregate power requirement of all static loads as a power profile $\mathbf{l}^S = \{l_k^S\}_{k=1}^N$ that must be satisfied at each balancing time k.

2.2.4.2 Deferrable Loads

Deferrable Loads require a certain energy be delivered over a specified time interval. The energy needs of deferrable loads can be modeled as *tasks*.

Definition 2.1. A *task* T_i is characterized by a service interval $\{a_i, \ldots, d_i\}(\subseteq \{1, \ldots, N\})$ over which a quantity of energy E_i must be delivered with a maximum power transfer rate m_i. Let p_{ik} be the power delivered to task T_i at time k. The task requirement can be expressed as:

$$\sum_{k=a_i}^{d_i} p_{ik}\Delta t \;=\; E_i\,, \;\; 0 \leq p_{ik} \leq m_i \;\forall k \in \{a_i, \ldots, d_i\} \tag{2.5}$$

Each task T_i is parametrized by (E_i, m_i, a_i, d_i). This parametrization succinctly captures the degree of flexibility present in deferrable loads. For instance, E_i, in the case of an EV, corresponds to a user-specified battery state-of-charge desired at the end of the service interval $[a_i, d_i]$, and m_i is a power transfer rate limit imposed by the charging equipment or other power delivery constraints. We make the following assumptions:

A1 Tasks are *preemptive*, i.e., they can be serviced with interruptions,

A2 The power consumed by a task T_i (p_{ik}) admits values on the continuous interval $[0, m_i]$.

While the ability to continuously modulate power delivery to each deferrable load (Assumption A2) may not be entirely accurate, it reasonably approximates reality. Crucially, it enables the computation of power allocation schedules using convex optimization methods. For practical implementations of such algorithms, we propose setting p_{ik} to the nearest discretized power level.

We now introduce some key definitions.

Definition 2.2. Consider a task T_i parametrized by (E_i, m_i, a_i, d_i). Let p_{ik} be the power delivered to task T_i at time k. The *energy state* of task T_i at time k is:

$$e_{ik} = E_i - \sum_{n=a_i}^{k} p_{in}\Delta t. \tag{2.6}$$

The energy state of a task at time k is the remaining energy requirement for the task at the end of the kth time-step.

Definition 2.3. A task T_i is *active at time* k if $a_i \leq k \leq d_i$ and $e_{ik} > 0$. For a collection of tasks $\mathbb{T} = \{T_i\}_{i=1}^{M_T}$, with T_i parametrized as (E_i, m_i, a_i, d_i), the set of active tasks at time k is written $\mathbb{A}_k$.

Equivalently, a task is active at time k if it can and needs to be serviced at that time.

2.2.4.3 Storage

We model electricity storage as *devices*.

Definition 2.4. A *device* D_j is characterized by maximum (E_j^+) and minimum (E_j^-) energy capacities, and maximum charging (m_j^+) and discharging $(-m_j^-)$ power transfer rates. Let p_{jk} be the power delivered to device D_j. Device requirements can then be expressed as:

$$E_j^- \leq \sum_{n=1}^{k} p_{jn}\Delta t \leq E_j^+ , \;\; -m_j^- \leq p_{jk} \leq m_j^+ \;\; \forall k \tag{2.7}$$

The maximum and minimum energy capacities can be readily calculated from the total storage capacity and initial energy level of the device. Thus, each device D_j can be parametrized by $(E_j^-, E_j^+, m_j^-, m_j^+)$. We remark that, unlike tasks, devices are always active. The set of all devices present during the operating window is $\mathbb{D} = \{D_j\}_{j=1}^{M_D}$.

Tasks and devices are distinguished by their energy requirements. The energy requirement for a task is expressed as a single equality constraint; the corresponding requirement for a device is a set of $2N$ linear inequality constraints.

2.2.5 Scheduling Policies

The CM must meet all load requirements over the operating window. Accordingly, the CM first assigns procured generation (renewables and bulk power) to meet static load requirements. The

remaining generation, which is used to satisfy deferrable load requirements, is called *available generation.*

Definition 2.5. Available generation to meet deferrable load requirements, $\mathbf{g} = \{g_k\}_{k=1}^N$, is the bulk and renewable power remaining after satisfying static load requirements.

$$g_k = w_k + B_k - l_k^S, \quad \forall k. \tag{2.8}$$

Scheduling policies determine allocations of available generation to deferrable loads and storage.

Definition 2.6. A scheduling policy σ is an algorithm that computes reserves, and allocations of available generation and reserves, to satisfy the requirements of all tasks over the operating window. Specifically, for a collection $\mathbb{T}$ of tasks, and $\mathbb{D}$ of devices, the scheduling policy is:

$$\sigma(\mathbf{g}, \mathbb{T}, \mathbb{D}) = (\mathbf{r}, \{\mathbf{p}_k\}_{k=1}^N), \tag{2.9}$$

where $\mathbf{p}_k$ is the set of power allocations at time k to all active tasks $\{p_{ik}\}_{i \in \mathbb{A}_k}$, and devices $\{p_{jk}\}_{j \in \mathbb{D}}$. Clearly, balancing of generation and load results in:

$$g_k + r_k = \sum_{i \in \mathbb{A}_k} p_{ik} + \sum_{j \in \mathbb{D}} p_{jk}. \tag{2.10}$$

A scheduling policy σ is causal if its allocations at time k depend only on the information state $\mathcal{I}_k$ at time k.

Definition 2.7. The **information state** $\mathcal{I}_k$ at time k consists of:

(a) Task parameters of tasks active at time k: (E_i, m_i, a_i, d_i) $\forall i \in \mathbb{A}_k$,

(b) Device parameters of all devices: $\left(E_j^-, E_j^+, m_j^-, m_j^+\right)$ $\forall j \in \mathbb{D}$,

(c) Energy states of tasks active at time k: e_{ik} $\forall i \in \mathbb{A}_k$,

(d) Realized values of available generation: $\{g_n\}_{n=1}^k$.

2.2.6 Performance Metric

The performance metric used to assess the performance of different scheduling policies is the cost of procuring reserve generation ($\mathbf{r}$) to meet load requirements. As mentioned earlier, there are two components to this cost:

$$J(\mathbf{r}) = p_r \sum_{k=1}^{N} |r_k| \Delta t \; + \; p_C \max_k |r_k| \tag{2.11}$$

The first term penalizes the total amount of reserve generation procured to meet load requirements. The second term captures the non-energy costs associated with making generators that provide reserve generation available. In practice, such reserve *capacity* is procured either through bilateral contracts, where the CM enters into long-term agreements with select generating facilities, or through the clearing of ex-ante capacity markets [87]. The prices p_r and p_C negotiate the relative importance of the cost components. As our primary focus is in quantifying the impact of the CM's real-time scheduling function on reserve costs, we do not consider the procurement costs associated with bulk power in our analysis.

2.2.7 Policy Optimization Problem

Our objective is to develop scheduling policies that reduce reserve costs as captured by the metric (2.11). This can be written as a functional optimization problem (2.12) over the function space of possible scheduling policies. Task and device constraints limit the feasible region of this function space.

$$\min_{\sigma} \quad J(\mathbf{r}) \quad \text{subject to:} \tag{2.12}$$

$$(\{\mathbf{p}_k\}_{k=1}^{N}, \mathbf{r}) = \sigma(\mathbf{g}, \mathbb{T}, \mathbb{D}) \tag{a}$$

$$\forall k, g_k + r_k = \sum_{i \in \mathbb{A}_k} p_{ik} + \sum_{j \in \mathbb{D}} p_{jk}, \tag{b}$$

$$\forall i \in \mathbb{T}, \quad \begin{array}{l} \sum_{k=a_i}^{d_i} p_{ik} \Delta t \; = \; E_i, \\ 0 \leq p_{ik} \leq m_i \; \forall k \in \{a_i, \ldots, d_i\}, \end{array} \tag{c}$$

$$\forall j \in \mathbb{D}, \quad \begin{array}{l} E_j^- \; \leq \sum_{n=1}^{k} p_{jn} \Delta t \; \leq \; E_j^+, \\ -m_j^- \leq p_{jk} \leq m_j^+ \; \forall k \in \{1, \ldots, N\}. \end{array} \tag{d}$$

We seek causal scheduling policies. Specifically, we want to solve a modified version of the optimization problem (2.12) where we constrain the policy to make decisions solely based on the current information state. We achieve this by replacing the constraint (2.12-a) with $(\mathbf{p}_k, r_k) = \sigma(\mathcal{I}_k), \forall k$.

Theorem 2.8. *The optimal scheduling policy that solves (2.12) is* not causal.

TABLE 2.2: Interpretation of parameters ϕ.

parameter	meaning
n_-, n_+	charge/discharge power limits
C	energy capacity
α	dissipation rate

Proof: This is a straightforward extension of the adversarial argument we used to show [90, Theorem 1] and is omitted. ■

This result forces us to be content with causal but necessarily sub-optimal heuristics for resource scheduling.

2.3 Generalized Battery Models of Aggregate Flexibility

Generalized battery models offer a compact representation for a *set of power signals*. These models will prove useful to represent the aggregate flexibility of a collection of loads.

Definition 2.9. A *Generalized Battery Model* $\mathbb{B}$ is a set of signals $u(t)$ that satisfy

$$-n_- \leq u(t) \leq n_+, \quad \forall\, t > 0,$$
$$\dot{x} = -\alpha x - u,\ x(0) = 0\ \Rightarrow |x(t)| \leq C, \quad \forall\, t > 0.$$

The model is specified by the non-negative parameters $\phi = (C, n_-, n_+, \alpha)$, and we write this compactly as $\mathbb{B}(\phi)$.

Remark 2.10. If we regard $u(t)$ as the power drawn from the battery and $x(t)$ as its state-of-charge at time t, the parameters ϕ have natural interpretations as summarized in Table 2.2. For this reason, we view $\mathbb{B}$ as a generalized electricity storage. We refer to the charge/discharge rate limits n_-, n_+ as the *power limits* and C as the *energy capacity* of the battery. It will happen that the parameters ϕ are random processes that depend on ambient temperature and participation rates. As a result, we regard $\mathbb{B}(\phi)$ as a stochastic battery. □

While the state-of-charge is defined to be bounded by C, it may be more severely constrained by the bounds on power signal $u(t)$. This observation motivates the following:

Definition 2.11. The *effective up and down energy capacities* C_+ and C_- of the stochastic battery model $\mathbb{B}(\phi)$ are defined as

$$C_+ = \min\{C,\ n_+/\alpha\}, \quad C_- = \min\{C,\ n_-/\alpha\}.$$

Consider a diverse collection of TCLs indexed by k. Let P_o^k denote the nominal power consumed by the kth TCL. Each TCL can accept perturbations $e^k(t)$ around its nominal power consumption that will meet user-specified comfort bounds. Define

$$\mathbb{E}^k = \left\{ e^k(t) \middle| \begin{array}{c} 0 \leq P_o^k + e^k(t) \leq P_m^k, \\ P_o^k + e^k(t) \text{ keeps } |\theta^k(t) - \theta_r^k| \leq \Delta^k \end{array} \right\}. \tag{2.13}$$

This set of power signals represents the flexibility of the kth TCL with respect to nominal. The *aggregate flexibility* of the collection of TCLs is defined as the Minkowski sum

$$\mathbb{U} = \sum_k \mathbb{E}^k. \tag{2.14}$$

The geometry of this set is, in general, unwieldy. Our objective is to develop succinct characterizations of the aggregate flexibility set. In our central result, we show that $\mathbb{U}$ can be bounded by generalized battery models as

$$\mathbb{B}(\phi_1) \subseteq \mathbb{U} \subseteq \mathbb{B}(\phi_2).$$

We have the following:

Theorem 2.12. *Consider a collection of* heterogeneous *TCLs modeled by the continuous-power model with parameters χ^k. The aggregate flexibility $\mathbb{U}$ of the collection satisfies*

$$\mathbb{U} \subseteq \mathbb{B}(\phi_2),$$

where the parameters $\phi_2 = (C, n_-, n_+, \alpha)$ are given by

$$C = \sum_k \left(1 + \left|1 - \frac{a^k}{\alpha}\right|\right) \frac{\Delta^k}{b^k},$$
$$n_- = \sum_k P_o^k, \quad n_+ = \sum_k P_m^k - P_o^k,$$

where $P_o^k = a^k(\theta_a^k - \theta_r^k)/b^k$ and $\alpha > 0$ is arbitrary.

Proof: See Appendix A. ■

We next examine sufficient characterizations of $\mathbb{U}$. There are many choices of battery parameters ϕ such that $\mathbb{B}(\phi) \subseteq \mathbb{U}$. We have the following:

Theorem 2.13. *Consider a collection of* heterogeneous *TCLs modeled by the continuous-power model with parameters χ^k. Fix $\alpha > 0$, and define*

$$f^k = \frac{\Delta^k}{b^k\left(1 + \left|\frac{\alpha - a^k}{a^k}\right|\right)}.$$

Fix $\beta^k \geq 0, k = 1, \ldots, N$, with $\sum_k \beta^k = 1$. Let (C, n_-, n_+) be any triple that satisfies the constraints

$$\left.\begin{array}{rcl} \beta^k n_- & \leq & P_o^k \\ \beta^k n_+ & \leq & P_m^k - P_o^k \\ \beta^k C & \leq & f^k \end{array}\right\}. \tag{2.15}$$

Then, the aggregate flexibility $\mathbb{U}$ of the collection satisfies

$$\mathbb{B}(\phi_1) \subseteq \mathbb{U},$$

where $\phi_1 = (C, n_-, n_+, \alpha)$.
Further, if $u \in \mathbb{B}(\phi_1)$, the causal power allocation strategy

$$e^k(t) = \beta^k u(t),$$

satisfies the dead-band constraints $|\theta^k(t) - \theta_r^k| \leq \Delta^k$.

Proof: See Appendix A. ■

Theorem 2.13 informs us that there are many battery models $\mathbb{B}(\phi_1)$ that offer sufficient characterizations of the aggregate flexibility set $\mathbb{U}$. In some situations we may seek a sufficient battery model with largest energy capacity C, and in others with largest charge power limit n_-, or with largest discharge power limit n_+. Table 2.3 summarizes the three extreme cases: maximize capacity C, charge rate n_-, or discharge rate n_+. The results follow from Theorem 2.13 by setting β^k as $\frac{f^k}{\sum_k f^k}$, $\frac{P_o^k}{\sum_k P_o^k}$ and $\frac{P_m^k - P_o^k}{\sum_k (P_m^k - P_o^k)}$, respectively.

Since $\mathbb{B}(\phi_1) \subseteq \mathbb{U} \subseteq \mathbb{B}(\phi_2)$, there is, in general, a gap between our necessary and sufficient characterizations of aggregate flexibility. We explore the conservatism in our battery models using simulations in Section 3.1.3. For a population of *homogeneous* TCLs, our battery model characterizations are exact. More precisely, we have:

TABLE 2.3: Comparison of generalized battery models for a collection of heterogeneous TCLs

	(Necessary) Battery $\mathbb{B}(\phi_2)$	(Sufficient) Battery $\mathbb{B}(\phi_1)$		
		Maximize C	Maximize n_-	Maximize n_+
C	$\sum_k (1 + \lvert 1 - \frac{a^k}{\alpha} \rvert) \frac{\Delta^k}{b^k}$	$\sum_k f^k$	$(\sum_k P_o^k) \min_k \frac{f^k}{P_o^k}$	$(\sum_k P_m^k - P_o^k) \min_k \frac{f^k}{P_m^k - P_o^k}$
n_-	$\sum_k P_o^k$	$(\sum_k f^k) \min_k \frac{P_o^k}{f^k}$	$\sum_k P_o^k$	$(\sum_k P_m^k - P_o^k) \min_k \frac{P_o^k}{P_m^k - P_o^k}$
n_+	$\sum_k P_m^k - P_o^k$	$(\sum_k f^k) \min_k \frac{P_m^k - P_o^k}{f^k}$	$(\sum_k P_o^k) \min_k \frac{P_m^k - P_o^k}{P_o^k}$	$\sum_k P_m^k - P_o^k$

Corollary 2.14. *Consider a collection of N* homogeneous *TCLs, modeled by the continuous-power model with parameters $\chi = (a, b, \theta_a, \theta_r, \Delta, P_m)$. Then,*

$$\mathbb{U} = \mathbb{B}(C, n_-, n_+, \alpha),$$

where

$$C = N\Delta/b, \;\; n_- = NP_o, \;\; n_+ = N(P_m - P_o), \;\; \alpha = a,$$

and $P_o = a(\theta_a - \theta_r)/b$.

Proof: Follows immediately from Theorems 2.12 and 2.13 by setting $\alpha = a^k$. ■

Remark 2.15. The gap between the battery models $\mathbb{B}(\phi_2)$ and $\mathbb{B}(\phi_1)$ increases with TCL heterogeneity. To obtain tighter models, we can *cluster* TCLs with similar parameters into g groups, and compute a battery model for each group. The aggregate flexibility is then represented as the union of g generalized battery models. Theorem 2.13 guides the choice of metric for clustering TCLs. □

Chapter 3

Controlling Demand-Side Flexibility

In this chapter, we then describe a direct load control architecture for regulation service provision where we use a priority-stack-based control framework to select which TCLs to control at any time. The control objective is for the aggregate power deviation from baseline to track an automatic generation control signal supplied by the system operator. Simulation studies suggest the practical promise of our methods.

3.1 Thermostatically Controlled Loads

Our objective is to reliably deliver frequency regulation service to the grid by actively controlling an aggregation of TCLs. The regulation signal or AGC command $r(t)$ is typically a command determined by the system operator at 4-second sampling based on the area control error [49]. We adopt a centralized control architecture. This choice is dictated by the reliability requirements necessary to participate in regulation ancillary service markets. At each sample time, the aggregator compares the regulation signal $r(t)$ with the aggregate power deviation $\delta(t) = P_{\text{agg}}(t) - n(t)$. Here $P_{\text{agg}}(t) = \sum_k q^k(t) P_m^k$ is the instantaneous power drawn by the TCLs, and $n(t) = \sum_k P_o^k$ is their baseline power. This requires a contractual *ex ante* agreement on what the "baseline" power consumption $n(t)$ of the aggregation will be on the forward delivery window.

If $r(t) < \delta(t)$, the population of TCLs needs to "discharge" power to the grid, which means some of the ON units will be turned OFF. Conversely, if $r(t) > \delta(t)$, then the population of TCLs must consume more power. This requires turning ON some of the OFF units. The selection of the most appropriate TCLs that must be turned ON or OFF is done through a priority-stack-based control strategy which will be described in the next section. We stress that this is a *feedback* control strategy

which offers robustness against disturbances w due to occupancy patterns, solar radiation, etc. and modeling errors in the dynamics of the TCLs. The overall control architecture is depicted in Fig. 3.1.

3.1.1 Priority Stack Scheduling

Suppose at time t we have $r(t) < \delta(t)$. The population of TCLs must reduce their power consumption. We must therefore turn OFF some units that are ON. It is most appropriate to turn OFF those units that will most imminently turn OFF. Imminence can be measured naturally by *temperature distance* to the switching boundary, i.e. by $\pi^k(t) = (\theta^k(t) - \underline{\theta}^k)/\Delta^k$, where $\underline{\theta}^k = \theta_r^k - \Delta^k$. The temperature distance is normalized to account for heterogeneity. We can therefore construct the ON *priority stack* which consists of units that are ON. The TCLs in this stack are ordered by their priority criterion, $\pi^k(t)$. Analogously, we construct the OFF priority stack of units that are OFF, ordered by $\pi^k(t) = (\overline{\theta}^k - \theta^k(t))/\Delta^k$, where $\overline{\theta}^k = \theta_r^k + \Delta^k$.

Imminence can also be measured by *time* to the switching boundary. For example, in the ON priority stack, units that will turn OFF autonomously the soonest receive the highest priority.

The unit with the highest priority will be turned ON (or OFF) first, and then units with lower priorities will be considered in sequence until the desired regulation is achieved. This priority-stack-based control strategy minimizes the ON/OFF switching action for each unit, which avoids short-cycling and reduces wear and tear of the mechanical equipment. Priority stacks are illustrated in Fig. 3.2. We index the units available for manipulation in the ON stack from *bottom to top* by $\{1, 2, \cdots, N_1\}$, and the units available for manipulation in the OFF stack from *top to bottom* by

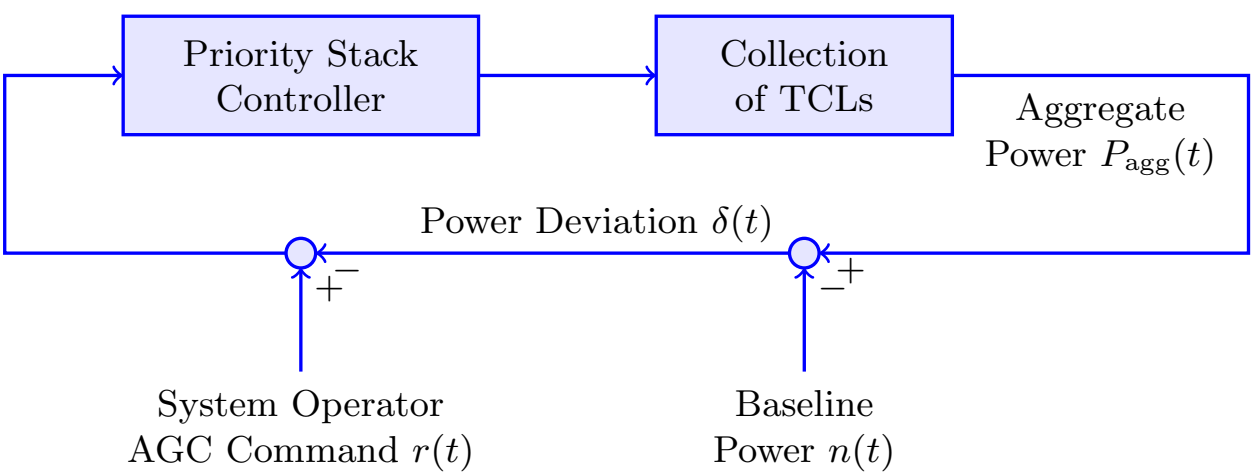

FIGURE 3.1: Control architecture for regulation service provisoin.

$\{1, 2, \cdots, N_0\}$. Units are turned ON or OFF when their *real time* power consumption $P^k(t)$ matches the difference between the regulation signal $r(t)$ and aggregate power deviation $\delta(t)$. The associated control algorithm is summarized in Algorithm 1.

Algorithm 1 Priority-stack-based control algorithm

```
loop
    receive π^k(t) and P^k(t);
    construct priority stacks;
    read r(t);
    compute δ(t) = P_agg(t) − n(t);
    if δ(t) < r(t) then
        find j* = min {j | Σ_{i=1}^{j} P^i(t) ≥ r(t) − δ(t)};
        turn ON units indexed by {1, 2, ··· , j*}.
    else if δ(t) > r(t) then
        find j* = min {j | Σ_{i=1}^{j} P^i(t) ≥ δ(t) − r(t)};
        turn OFF units indexed by {1, 2, ··· , j*}.
    end if
end loop
```

Remark 3.1. The priority-stack-based control offers a generic architecture suitable for direct load control of various classes of flexible loads. For example, common scheduling strategies for electric

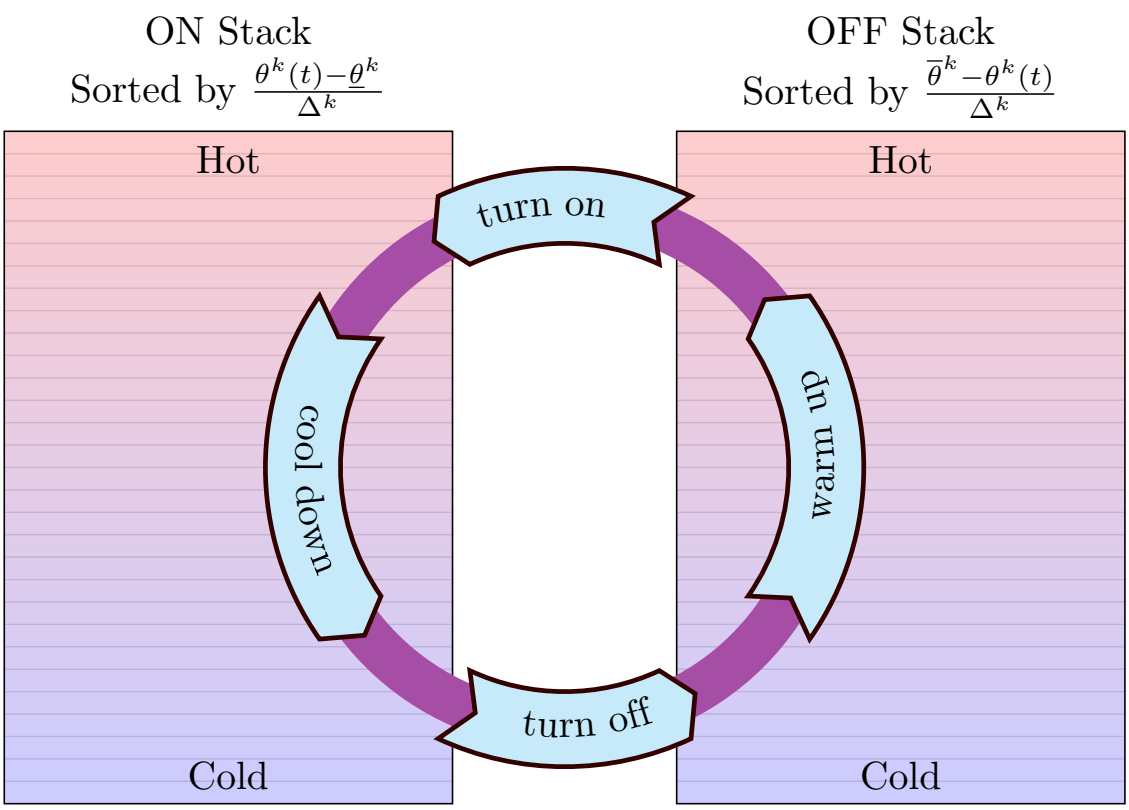

FIGURE 3.2: The ON and OFF priority stacks. The stacks are sorted by temperature distance.

vehicle fleet such as Earliest-Deadline-First define a priority stack [73]. Priorities can also be used to encode pricing and service quality: consumers who wish to exercise greater control over their loads could opt to receive low compensation when called to supply regulation services in exchange for a lower priority position.

3.1.2 Practical Considerations

To implement the proposed direct load control strategy, we require (at a minimum) measurements of power $P^k(t)$ and temperature $\theta^k(t)$ at a sampling rate of 0.25 Hz for each TCL. While $\theta^k(t)$ and the set-point θ_r^k are directly available from the thermostat, measuring the power $P^k(t)$ requires additional hardware infrastructure. For each TCL, run-time system identification algorithms can be used to estimate the operating state $q^k(t)$, the ambient temperature θ_a^k, and model parameters a^k, b^k, Δ^k from the temperature time series $\theta^k(s), s \leq t$. Using this information, a local embedded controller computes $\pi^k(t)$ for each TCL.

The priority criterion $\pi^k(t)$ and power consumption $P^k(t)$ are transmitted to the aggregator. The aggregator forms the priority stack from the collated data and computes the control action for the next sample. This is broadcast to the TCLs where the local controller implements the action. Latency in the control loop will determine the quality of the offered regulation service in terms of power ramp rates. This scheme has modest computation and communication overhead.

Remark 3.2. Measuring the power consumption of each TCL necessitates nontrivial capital cost as power meters are expensive. *In our view, this is unavoidable.* Alternate and realistic schemes, Scenarios 1 and 2 of [63], have been proposed that use population-bin models and requires measuring only the TCL ON/OFF state. This requires moderately simpler sensing infrastructure, but requires a more complex control strategy. Other schemes have been proposed where the aggregate power $P_{\text{agg}}(t)$ is *estimated* using population models [64], or disaggregated from substation measurements [58], and Scenarios 3 and 4 of [63]. The former scheme is open-loop in character. The latter schemes must contend with the fact that participating TCLs represent a small fraction of the connected loads at a substation, making it very difficult to infer their power consumption from an aggregate measurement. These schemes therefore face big challenges in meeting the reliability requirements necessary to participate in the regulation ancillary service market [74].

Remark 3.3. Our scheme (as well as Scenarios 1 and 2 of [63]) requires real-time telemetry to transmit power measurements to the aggregator. This requires very little bandwidth and is becoming increasingly inexpensive. Since telemetry costs are dominated by sensing infrastructure costs, transmission represents a minor expense. Real-time telemetry is not always needed when a conventional generator resource is providing regulation. The key difference is that with TCLs we have *distributed*

TABLE 3.1: Power limits and energy capacities.

	Sufficient Battery $\mathbb{B}(\phi_1)$	Necessary Battery $\mathbb{B}(\phi_2)$
n_-	1.9 MW	1.9 MW
n_+	2.8 MW	3.7 MW
C	0.19 MWh	0.26 MWh

resources. Telemetry is needed to close the loop, i.e. send measurements to the controller which is implemented by the remote aggregator. For generators, the control is local: sensing and actuation are at the same place, so real-time telemetry may not be needed.

3.1.3 Simulation Studies

We consider a population of 1000 diverse ACs. In our simulations, we use the more accurate hybrid model (2.1) for each TCL. The nominal model parameters are listed in Table 2.1. For a heterogeneous collection, we assume that the TCL parameters are drawn from a uniform distribution with 10% heterogeneity around their nominal values. For example, $R_{th}^k \sim U(0.95R_o, 1.05R_o)$ where R_o is the nominal value of the thermal resistance. The ambient temperature is assumed to be $32°C$, and the initial temperatures and operating states of the population of TCLs are randomized.

The power limits and energy capacities of the collection of TCLs using the (necessary) battery model $\mathbb{B}(\phi_2)$ and the (sufficient) battery model $\mathbb{B}(\phi_1)$ are listed in Table 3.1. In both models, the dissipation parameter α is given as the average of the time constants of individual TCLs. Formally, we assume $\alpha := \frac{1}{N}\sum_{k=1}^{N} 1/(R_{th}^k C_{th}^k)$. For the sufficient battery model, we maximize the charge rate n_- (the fourth column in Table 2.3). Note that the battery parameters are derived based on the continuous power model (2.3).

We apply our priority-stack-based control strategy to track a one-hour long regulation signal $r(t)$ from PJM (Pennsylvania-New Jersey-Maryland Interconnection) [75]. The magnitude of the PJM signal is scaled appropriately to match the power limits and energy capacity of 1000 ACs.

3.1.4 Tracking Performance

Fig. 3.3 shows that if the regulation signal $r(t)$ has both power and capacity requirements within the analytic bounds of the (sufficient) battery model $\mathbb{B}(\phi_1)$, the population of TCLs delivers excellent tracking. The maximum tracking error is less than 1% of the maximum magnitude of the regulation signal. Additional simulation results (not reported here) reveal that even with one sample (4 sec) communication delay, good tracking is still achieved with a maximum tracking error less than 5% of

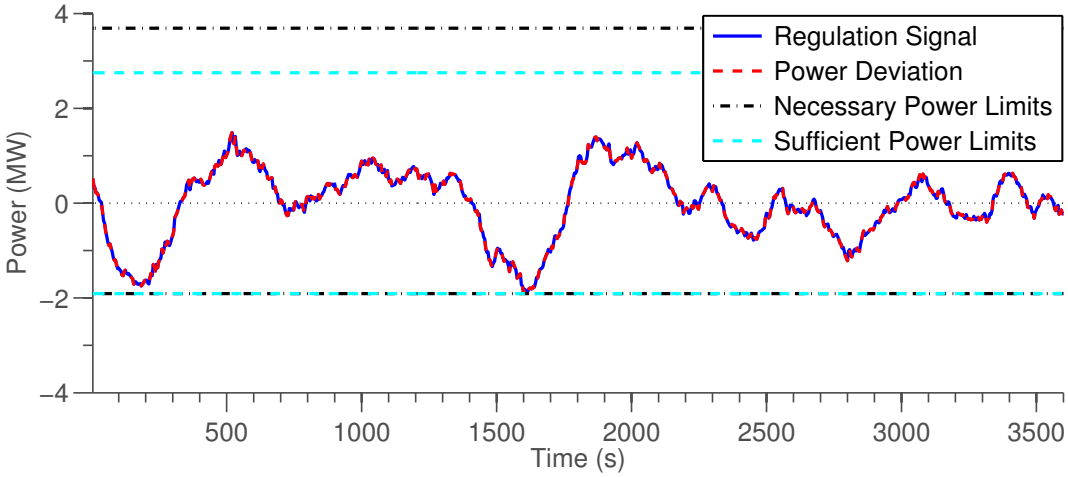

(a) Power trajectory

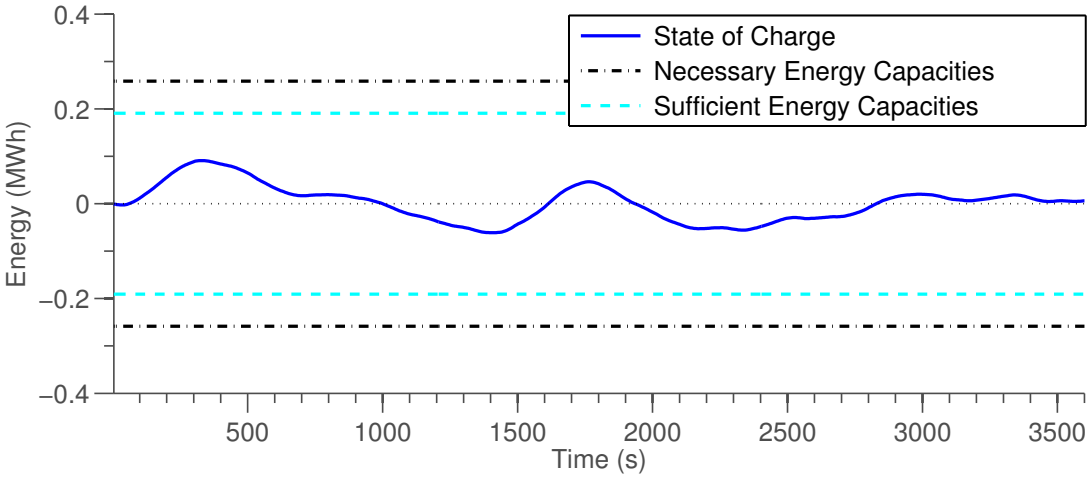

(b) State of charge

FIGURE 3.3: Tracking of a regulation signal succeeds when it is within the power limits and energy capacity of the (sufficient) battery model $\mathbb{B}(\phi_1)$.

the maximum magnitude of the regulation signal. If the regulation signal violates either the power limits or energy capacity of the (necessary) battery model $\mathbb{B}(\phi_2)$, the population of TCLs fails to track the regulation signal. Figs. 3.4 and 3.5 show that when the regulation signal exceeds the power limits or the energy capacity respectively, we cannot track the regulation signal. Extensive simulations (not reported here) using other regulation signals yield similar conclusions.

We use a typical 6-hour long regulation signal from PJM (shown in Fig. 3.6 (a)) that is fairly close to the power limits and energy capacity to test the prediction performance of our sufficient battery model. Specifically, we examine the effect of tracking of a regulation signal that (just) satisfies

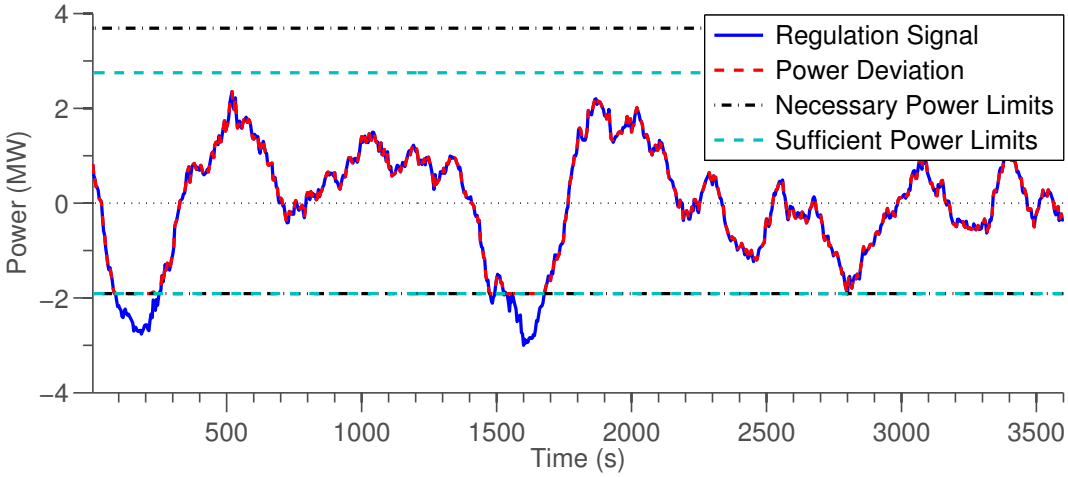

(a) Power trajectory

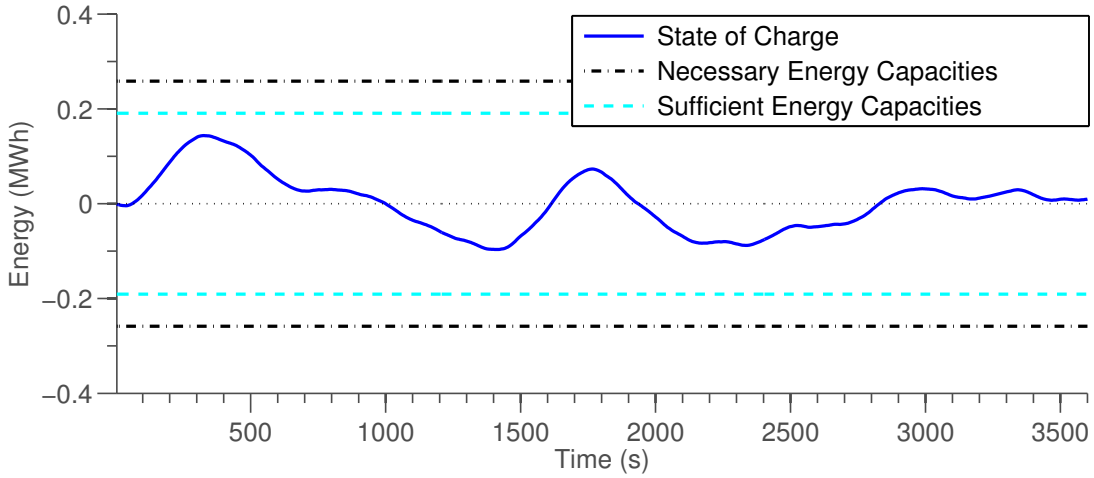

(b) State of charge

FIGURE 3.4: Tracking of a regulation signal fails when it exceeds the power limit of the (necessary) battery model $\mathbb{B}(\phi_2)$.

the sufficient battery model on the number of additional ON/OFF switchings that occur above nominal, and occurrence of short cycling events. Table 3.2 shows the performance statistics. The average number of ON/OFF switching in the 6-hour period without providing regulation service (nominal) is 30 times, while the average switching in the same period with regulation using the priority-stack controller is 65 times, or about twice of the nominal value. Additionally, we assume for each unit, its minimum ON-time and OFF-time are 1 minute. With provision of frequency regulation, the minimum, average, and maximum number of short cycling events among all units in the 6-hour period are respectively 0, 1.1, and 8. Extensive simulations using other regulation signals are also conducted, and similar statistics are obtained. We observe that the number of short

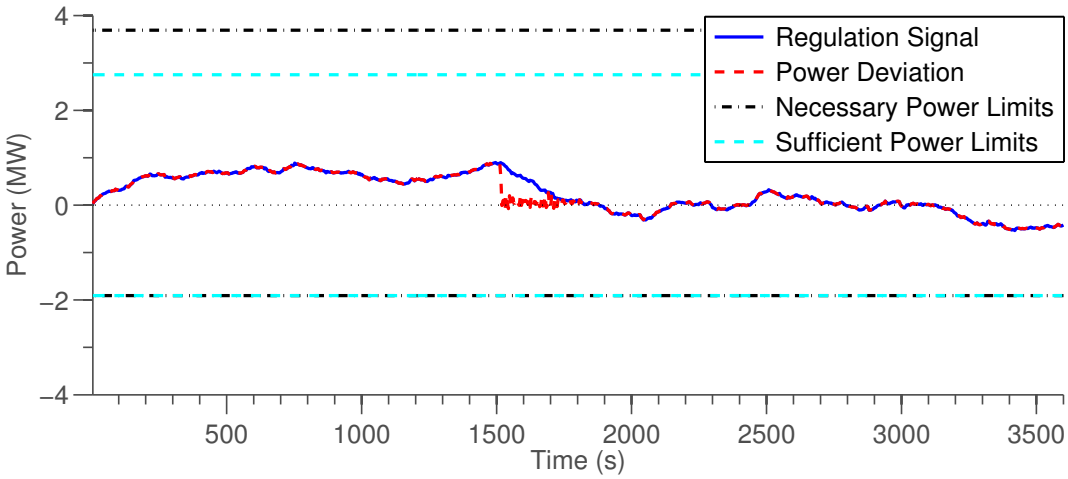

(a) Power trajectory

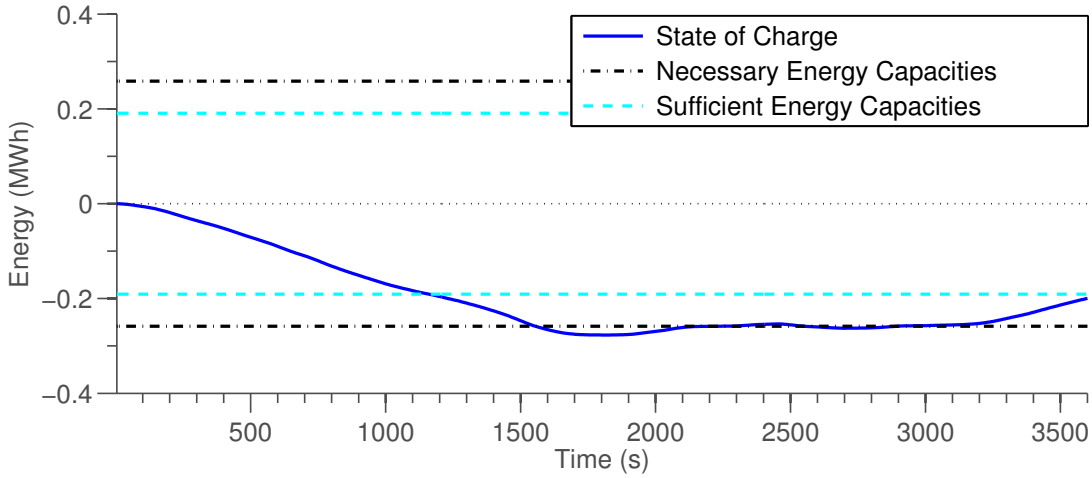

(b) State of charge

FIGURE 3.5: Tracking of a regulation signal fails when it exceeds the energy capacity of the (necessary) battery model $\mathbb{B}(\phi_2)$.

cycling events is relatively small, and would be drastically reduced with a less aggressive regulation signal. Note also that short cycling events are closely connected to large, high frequency oscillations in the regulation signal, and an exact characterization of feasibility with a short cycling constraint is reported in [76].

3.1.5 Battery Model Conservatism

Recall that the gap between the necessary and sufficient battery models is due to heterogeneity in the collection of TCLs (see Corollary 2.14). We synthetically vary the diversity d in the TCL parameters from 1% to 40%. In each case, we then scale the AGC signal $r(t)$ until our control

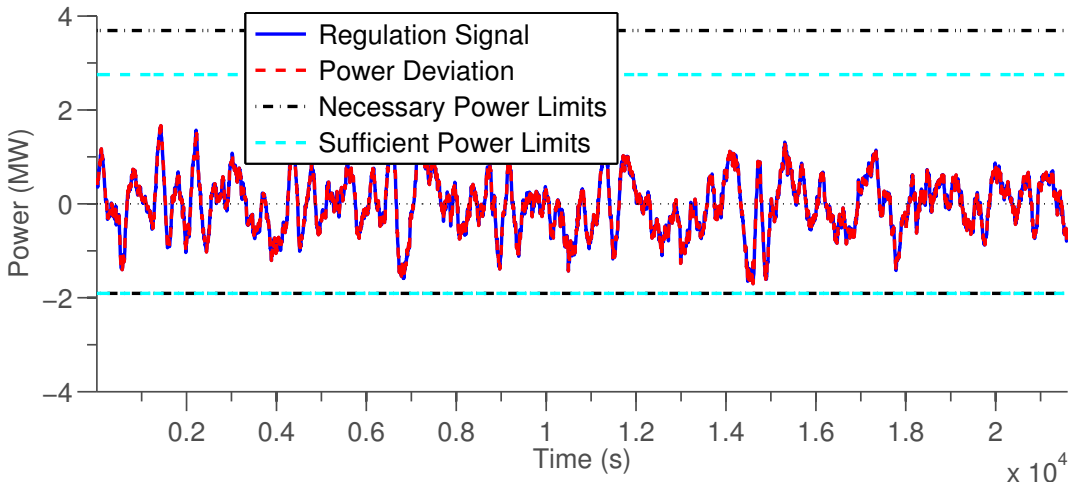

(a) Power trajectory

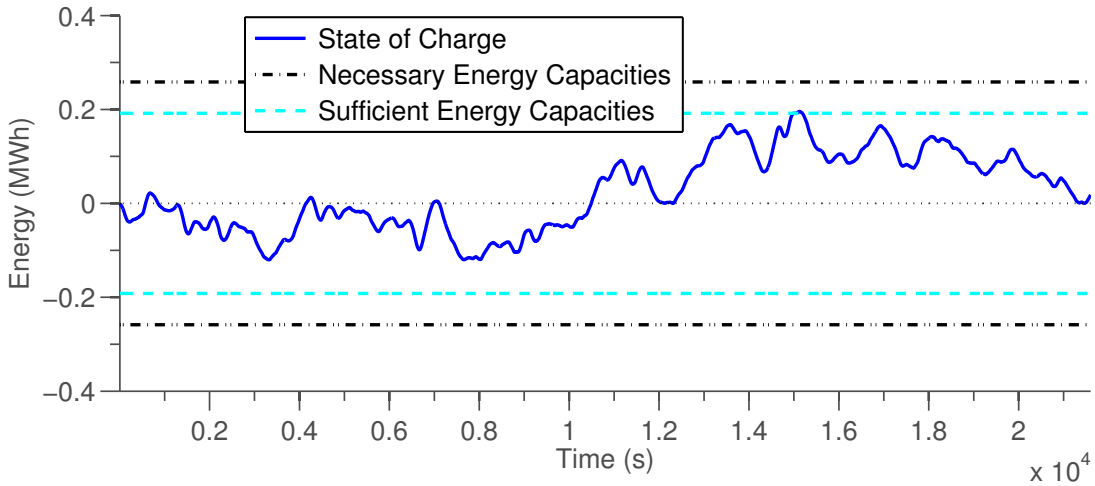

(b) State of charge

FIGURE 3.6: Tracking of a typical 6-hour regulation signal from PJM that is within the power limits and energy capacity of the (sufficient) battery model $\mathbb{B}(\phi_1)$.

system fails to provide tracking on a 1-hour window. More precisely, we numerically compute (for each d),

$$\max \gamma : \quad \text{tracking error } \leq 1\% \text{ for AGC command } \gamma r(t)$$

We declare the "true" battery power limits and energy capacity to be

$$n_-^{\text{true}} = -\gamma \min_t r(t), \quad n_+^{\text{true}} = \gamma \max_t r(t),$$
$$C^{\text{true}} = \max |x(t)|, \quad \dot{x} = -\alpha x - \gamma r(t), \; x(0) = 0.$$

TABLE 3.2: Prediction performance of the sufficient battery model

	Switchings		Short Cyclings		
	nominal	regulation	min	mean	max
Sufficient Battery	30	65	0	1.1	8

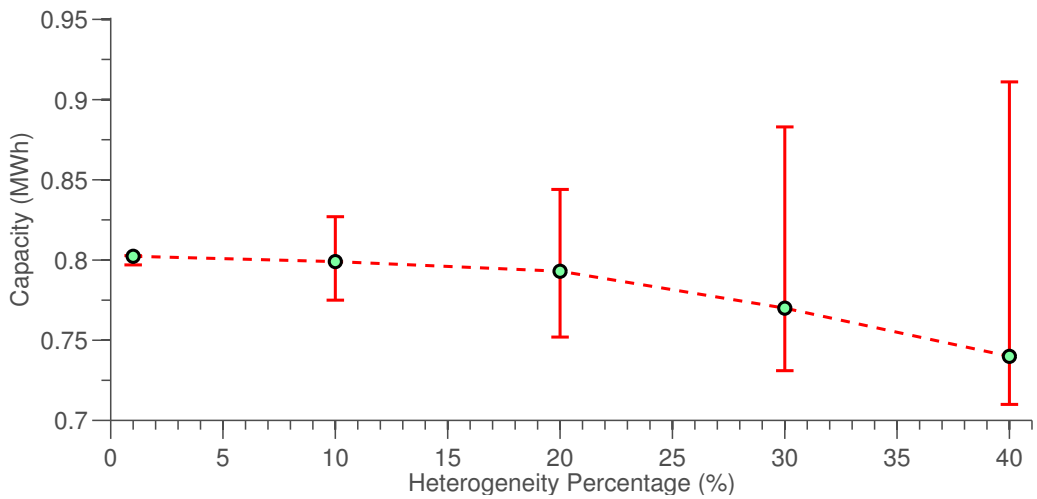

FIGURE 3.7: Conservatism in the bounds on energy capacity from the necessary and sufficient models. The green dots show the "true" energy capacity as calculated numerically. The red bars extend between the energy capacity bounds given by the sufficient and necessary battery models.

Fig. 3.7 compares the numerical energy capacity to the bounds given by the necessary and sufficient battery models as a function of diversity d. We note that for $d < 10\%$, the models capture the aggregate flexibility quite well. Similar plots can be obtained for power limits also.

3.1.6 The Effect of Ambient Temperature

We have assumed that all of the available TCLs participate in offering regulation services. In practice, for certain types of TCLs such as ACs and heat pumps, participation rates depend strongly on ambient temperature θ_a. For ACs, we expect little participation when θ_a is low, and significant participation when θ_a is high. Modeling participation is extremely complex, requires large amounts of data, and any resulting models will likely have limited predictive power. For purposes of illustration, suppose we synthetically model participation using an inverse tangent function as shown in Fig. 3.8. This captures the intuitive observation that more people are likely to use their AC units at higher ambient temperatures. This does not account for occupancy which exhibits daily and

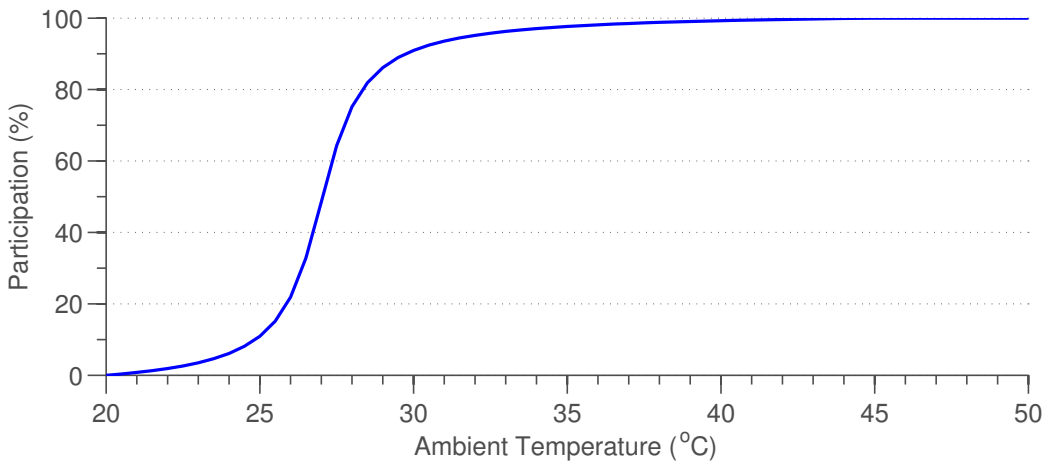

FIGURE 3.8: Participation percentage vs. ambient temperature.

hourly patterns. With this participation model, we can compute the sufficient battery parameters ϕ_1 as a function of ambient temperature. These are shown in Fig. 3.9 for air-conditioning loads.

3.1.7 Conclusions and Future Work

In this work, we illustrated that (a) the generalized battery model provided a succinct and powerful framework to characterize the aggregate flexibility of a population of TCLs, (b) the power and capacity bounds derived from the continuous model accurately captured the aggregate flexibility of TCLs with the hybrid model, and (c) the priority-stack-based control strategy yielded excellent tracking performance and good robustness.

The enormous potential of TCLs presents a tremendous opportunity for providing regulation service to the grid. There are several important research issues that must be addressed to realize this vision. These include: (a) deriving battery models that account for TCL model uncertainty, (b) an exploration of suitable and low-cost firmware and communication infrastructure to implement direct-load control, (c) understanding the quality of regulation service provided in terms of latency and ramp rates, (d) estimating the overall hourly availability of TCLs using historic measurement data, and (e) developing fair schemes to compensate loads for participating in regulation services.

For a signal x with support $[0, \infty)$, we define the standard norms $\|x\|_\infty = \max_t |x(t)|, t \geq 0$ and $\|x\|_1 = \int_0^\infty |x(t)| dt$.

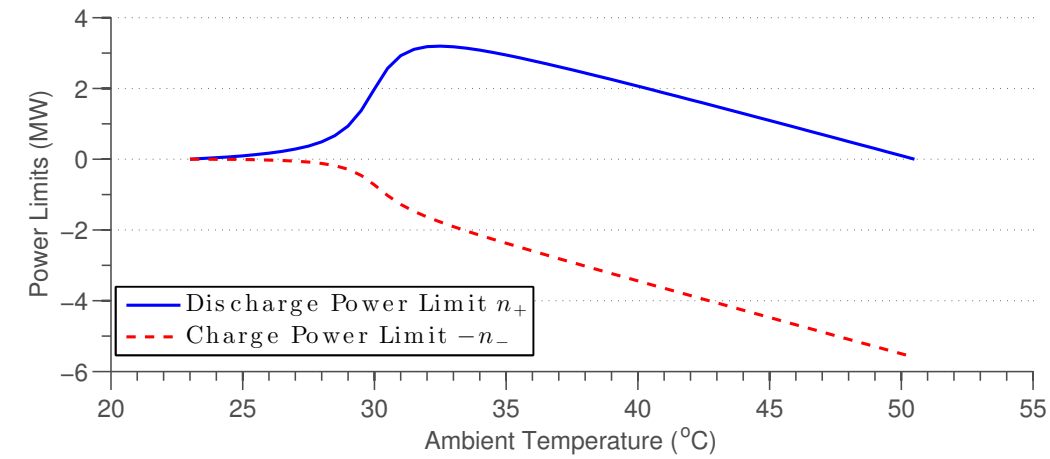

(a) Power limits

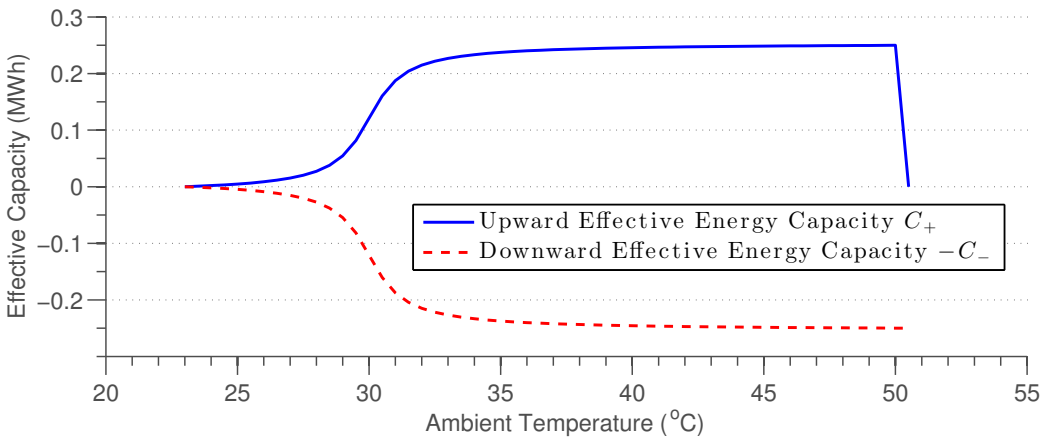

(b) Effective energy capacities

FIGURE 3.9: Power limits and effective capacities vs. ambient temperature.

3.2 Deferrable Loads

In this section, we first describe two causal heuristics for allocating available generation. We then develop a policy for procuring and allocating reserve generation called zero-laxity. Finally, we formulate a receding horizon control (RHC) algorithm inspired by the functional optimization problem presented in Section 2.2.7.

3.2.1 Earliest-Deadline-First (EDF) Scheduling

Earliest deadline first is a well-known real-time scheduling policy for processor time allocation (PTA). The objective in PTA is to schedule a collection of computational tasks, on single or multiple

processors, such that all task deadlines are met. In this context, EDF is known to be an optimal scheduling policy for a single processor in the following sense: if some scheduling policy can meet all deadlines of a collection of pre-emptive tasks, then the schedule for those tasks under EDF also satisfies all deadlines [91]. In addition, Liu and Layland showed that EDF scheduling meets the deadlines of all periodic tasks as long as total processing requirements do not exceed processor capacity [92].

The resource scheduling problem considered in this chapter is similar to PTA with available generation as the analog to processing time. However, there are three primary differences between the two scheduling problems. First, while total available processing time is constant over an operating window, its counterpart in resource scheduling, available generation, is variable. Second, resource scheduling involves task rate limits which constrain power service to particular tasks. Comparable scheduling constraints do not exist for PTA. Third, available capacity on a processor, at a given time, is devoted fully to a single task. In comparison, multiple tasks can be scheduled and serviced using available generation in resource scheduling. As a result, many of the performance guarantees for EDF present in the PTA literature, such as optimality, do not apply to an EDF resource scheduling policy.

Consider the scheduling problem at a time k where available generation g_k is allocated to active tasks $\mathbb{A}_k$ and devices $\mathbb{D}$. The EDF scheduling policy allocates available generation g_k to the task T_I with the *m*ost imminent deadline, i.e. $I = \text{argmin}_{i \in \mathbb{A}_k} d_i$. Available generation g_k in excess of the rate limit for task T_I is allocated to the active task with the next imminent deadline. Ties are broken arbitrarily. This process continues either until all generation g_k is expended or all active tasks are serviced at their rate limits.

The device scheduling component of this algorithm is as follows. If all generation is expended, active tasks not serviced at their rate limits receive power from energy stored in devices. This process continues either until all active tasks are serviced at their rate limits or until device constraints (discharge rate m_j^- or minimum energy level E_j^-) limit any further power service. If all active tasks are being serviced at their rate limits, excess available generation is allocated to devices $\mathbb{D}$. This process continues either until all available generation is expended or until device constraints (charging rate m_j^+ or maximum energy capacity E_j^+) limit any further service.

3.2.2 Least-Laxity-First (LLF) Scheduling

EDF allocates resources only on the basis of task deadlines. It does not account for task energy states. Least laxity first (LLF) is another causal scheduling policy that specifically incorporates

remaining task requirements in allocation decisions.

LLF has also been analyzed as a dynamic scheduling algorithm for PTA [94, 95]. Mok showed that, similar to EDF, LLF is also an optimal scheduling policy for PTA in the single processor case [93]. Certain studies have also explored the use of an LLF scheduling policy for resource scheduling. For instance, Chen et al. study the performance of satisfying EV charging needs in a parking garage through LLF-based scheduling [96].

Definition 3.4. The *laxity* for each task T_i at time k, ϕ_{ik}, is defined as:

$$\phi_{ik} = (d_i - k) - \frac{e_{ik}}{m_i} \tag{3.1}$$

Laxity is the difference between the amount of time remaining and minimum time required to satisfy task requirements. This is a measure of the degree of deferrability in scheduling a task. Tasks with larger laxities offer greater scheduling flexibility.

Consider the scheduling problem at a time k where available generation g_k is allocated to active tasks $\mathbb{A}_k$ and devices $\mathbb{D}$. The LLF scheduling policy allocates available generation g_k to the task T_I with the least laxity, i.e. $I = \text{argmin}_{i \in \mathbb{A}_k} \phi_{ik}$. Available generation g_k in excess of the rate limit for task T_I is allocated to the active task with the next smallest laxity. Ties are broken arbitrarily. This process continues either until all generation g_k is expended or all active tasks are serviced at their rate limits. The device scheduling component of LLF is identical to that of EDF described in Section 3.2.1.

3.2.3 Reserve scheduling: Zero - Laxity

The EDF and LLF policies described above determine allocations of available generation to tasks and devices. A resource scheduling policy must also make reserve generation procurement decisions and allocate operating reserves to various loads. We use task laxities to develop the reserve scheduling component for these policies.

A task cannot be completed by its deadline if the minimum time required to satisfy its energy requirement exceeds the time until deadline. Specifically, a task T_i is infeasible at time k if and only if $\phi_{ik} < 0$. The zero-laxity (ZL) policy, a causal heuristic for reserve scheduling, ensures that no task laxities become negative.

Consider the scheduling problem at time k. The CM first ensures that the static load requirement at time k is satisfied. If static load exceeds the sum of bulk and renewable generation ($g_k < 0$),

the CM must procure sufficient reserves to meet static load requirements: $r_k = -g_k$. In addition, for each active task T_i with laxity equal to 0, the CM procures and allocates sufficient reserve generation to ensure that T_i is serviced - from both available and reserve generation - at its rate limit m_i. No reserves are assigned to devices.

We employ the ZL reserve scheduling policy in conjunction with EDF or LLF for scheduling available generation. The computation in these algorithms is modest: sorting task deadlines (EDF) or laxities (LLF).

3.2.4 Receding Horizon Control (RHC)

Receding horizon control (RHC), or model predictive control (MPC), is a widely used and effective strategy for state and input constrained control problems [97]. RHC involves solving finite horizon optimization problems successively at each time step to determine appropriate control actions. In this context, an optimization problem is solved at each time-step k to yield cost-minimizing generation scheduling allocations over some time horizon $\{k, \ldots, k+H\}$. These decisions are based on generation forecasts and information from all active tasks and devices. Tasks are then scheduled based only on computed allocation decisions for the first time-step k. This process is repeated at the next time-step $k+1$ with an updated set of active tasks, devices, and generation forecasts on a future horizon $\{k+1, k+1+H\}$. Determining power allocations in this iterative fashion enables the incorporation of updated task information and more precise generation forecasts.

Consider the optimization problem solved for RHC scheduling at each time-step k. Let the *horizon length* H be the number of Δt time-steps between k and the largest deadline in the set of active tasks ($H = \max_{i \in \mathbb{A}_k} d_i$).

Let M_T and M_D be the number of active tasks and devices respectively. Without loss of generality, we set $k = 1$. In addition, let $\hat{\mathbf{g}}$ refer to forecasted values of available generation through the horizon $\{\hat{g}_1, \hat{g}_2, \ldots, \hat{g}_H\}$. Each optimization problem attempts to meet all tasks requirements by allocating only forecasted available generation.

3.2.4.1 Decision Variables

(a) $G \in \mathbb{R}_+^{M_T \times H}$ where G_{it} is the amount of available generation assigned to task T_i at time-step t,

(b) $R \in \mathbb{R}_+^{M_T \times H}$ where R_{it} is the amount of reserve generation assigned to task T_i at time-step t,

(c) $S \in \mathbb{R}^{M_D \times H}$ where S_{jt} is the amount of power assigned to device D_j at time-step t.

3.2.4.2 Objective Function

$$
\begin{aligned}
J(G,R,S) = \alpha_E \sum_{t=1}^{H} \Big(\sum_{i=1}^{M_T} R_{it} \Big) + \Big(\hat{g}_t - \sum_{i=1}^{M_T} G_{it} - \sum_{j=1}^{M_D} S_{jt} \Big) \\
+ \alpha_C \Big(\max_t \sum_{i=1}^{M_T} R_{it} + \max_t \Big(\hat{g}_t - \sum_{i=1}^{M_T} G_{it} - \sum_{j=1}^{M_D} S_{jt} \Big) \Big) \\
+ \sum_{t=1}^{H} \sum_{i \in \mathbb{A}_t} (N - \phi_{it})^2
\end{aligned} \tag{3.2}
$$

The first and second terms capture up and down reserve *energy* costs respectively while the third and fourth terms capture up and down reserve *capacity* costs respectively. The fifth term in (3.2) maximizes task laxities at subsequent time-steps within the horizon H. Effectively, this incentivizes earlier allocations of available generation, and discourages deferral of task energy requirements. The parameters α_E, and α_C negotiate the relative importance of the objective function components.

3.2.4.3 RHC Optimization Problem

$$
\min_{G \geq 0, R \geq 0, S} J(G,R,S) \quad \text{subject to:}
$$

$$
\forall\, t,\ \sum_{i=1}^{M_T} G_{it} + \sum_{j=1}^{M_D} S_{jt} \ \leq\ \hat{g}_t \tag{3.3}
$$

$$
\forall\, i,\ \sum_{t=1}^{d_i} G_{it} + R_{it} \ =\ E_i \tag{3.4}
$$

$$
\forall\, t,\ \begin{cases} 0 \leq G_{it} + R_{it} \leq m_i \Delta t, & \forall\, i : t \leq d_i \\ G_{it} = 0, R_{it} = 0, & \forall\, i : t > d_i \end{cases} \tag{3.5}
$$

$$
\forall\, j, t,\ E_j^- \ \leq \sum_{n=1}^{t} S_{jn} \ \leq\ E_j^+ \tag{3.6}
$$

$$
\forall\, j, t,\ -m_j^- \ \leq\ S_{jt} \ \leq\ m_j^+ \tag{3.7}
$$

$$
\phi_{it} = d_i - t - \frac{e_{it}}{m_i}, e_{it} = E_i - \sum_{n=1}^{t} G_{in} + R_{in} \tag{3.8}
$$

3.2.4.4 Constraints

(a) *Generation*: (3.3) ensures that the sum of all power allocated to tasks and devices cannot exceed forecasted available generation ($\hat{g}_t$) at any time t. Furthermore, when satisfying tasks using energy stored in devices, (3.3) guarantees total power delivered to tasks does not exceed total energy discharged from devices.

(b) *Task: Total Energy Requirement*: (3.4) ensures each task's energy requirement (E_i) is met through allocation of available and reserve generation.

(c) *Task: Rate Limits*: (3.5) ensures power is only allocated to active tasks. Moreover, the allocation for a task is non-negative and bounded by the task rate limit (m_i).

(d) *Device: Capacity*: (3.6) enforces maximum (E_j^+) and minimum (E_j^-) energy levels on the device energy state at all times within the horizon.

(e) *Device: Rate Limits*: (3.7) enforces discharge ($-m_i^-$) and charging (m_i^+) rate limits on allocations to and from devices.

(f) *Laxity*: (3.8) is used to compute task laxities ϕ_{ik} which are present in the objective function.

This optimization problem is a quadratic program (QP) which must be solved at each time-step k. The computations associated with RHC scheduling are substantial when compared to resource scheduling under the EDF or LLF policies. However, our simulation studies reveal that RHC scheduling offers significant reductions in reserve capacity costs.

3.2.5 Simulation Study

We now demonstrate the value of resource scheduling in reserve cost reduction using simulation studies. All reported results are averaged over 50 test cases. In each of these cases, a CM manages a resource cluster consisting of significant levels of solar PV generation, static loads, and EVs as deferrable loads. We implement the scheduling policies described in Section 3.2 to meet load requirements at 10 minute intervals over a 24 hour window. We quantify algorithm performance in terms of reserve energy and capacity costs.

3.2.6 Parameters

3.2.6.1 Load

The load is comprised of two components: static and deferrable. Static load profiles are created from time-series data of total California ISO system demand sampled at 10 minute intervals. This data is normalized to a peak load of 2 MW. We generate static load profiles by randomly selecting one day from 8 days of June 2012 data.

A specified fraction (α) of the total load energy requirement is assumed to be deferrable. Deferrable loads are modelled as tasks. Task parameters (arrivals, deadlines, and energy needs) are randomly generated based on typical EV charging characteristics [102]. These parameters are chosen to ensure initial task feasibility ($E_i \leq m_i(d_i - a_i + 1)\Delta t$). We use a constant maximum rate of charge m_i for all tasks consistent with the SAE J1772 AC Level 1 EV charging standard [101]. We create a nominal profile for deferrable load $\{l_k^D\}_{k=1}^N$ assuming each task receives constant power over its entire service interval.

$$l_k^D = \sum_{i \in \mathbb{A}_k} \frac{E_i}{d_i - a_i + 1} \tag{3.9}$$

This profile is used in bulk procurement decisions and serves as the baseline for computing reserve reductions achieved through resource scheduling.

3.2.6.2 Generation

We use solar PV generation data, obtained from the PV integrator SolarCity, aggregated from 30 different but proximate installations in California to create renewable generation profiles. This time series data of total PV generation is scaled to a peak output of 750 kW sampled at 1 minute intervals. Day-long renewable generation profiles are randomly selected from this data set.

As the CM determines load allocations every 10 minutes, scheduling decisions are made at each balancing time assuming the current amount of renewable generation is constant over the following 10 minute interval. However, reserve requirements for load balancing are computed using the actual renewable generation profile.

We create synthetic renewable generation forecasts for RHC scheduling by adding Gaussian noise to these solar generation profiles: the forecast $\hat{w}_t$ made at time k is

$$\hat{w}_t = w_t + \sum_{n=k+1}^{t} \epsilon_n \ , \ \ t = \{k+1, \ldots, k+H\} \tag{3.10}$$
$$\epsilon_n \sim \mathcal{N}\left(0, \sigma_n^2\right) \ , \ \hat{w}_k = w_k$$

where H is the horizon length of the RHC scheduling algorithm. We assume σ_n^2 increases *linearly* with the prediction window length n. This applies to all k corresponding to times between 06:00 and 20:00 in the day. Outside this interval, we assume there is no renewable generation.

3.2.6.3 DA Market Clearing

We simulate bulk power purchase decisions based on hourly forecasts of load and renewable data. For static loads, we use CAISO data containing hourly DA load forecasts ($\bar{l^S}$) for each hour of the operating window. For deferrable loads, we create DA forecasts ($\bar{l^D}$) by adding zero-mean Gaussian noise to hourly averages of nominal deferrable load (l^D) with a variance of 3% of deferrable load.

For hourly renewable generation forecasts, the DA forecast $\bar{w}$ employed is the average over the past 5 days of the mean generation for a particular hour-long interval. The bulk generation procured at each time-step k, constant over hour-long intervals, is then:

$$B_k = \bar{l^S}_k + \bar{l^D}_k - \bar{w}_k \tag{3.11}$$

3.2.7 Comparison of Scheduling Algorithms

Figure 3.10 illustrates the performance of EDF, LLF, and RHC scheduling policies in reducing reserves for a typical test case. Each subfigure compares the load profiles achieved by a particular resource scheduling algorithm to the baseline load and procured generation profiles. The generation profile shown in each sub-figure is the sum of bulk and renewable power. These plots reveal that resource scheduling can modify load profiles to closely match generation and thus reduce the need for reserves.

Qualitatively, the load profiles under the various scheduling policies vary significantly towards the end of the operating window (16:00-24:00 hours). Specifically, scheduling under EDF and LLF results in greater up reserve requirements during this period than scheduling under RHC. We note

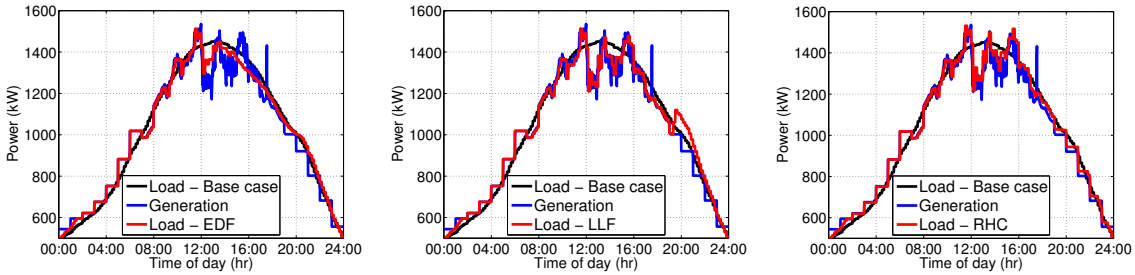

(a) EDF (b) LLF (c) RHC

FIGURE 3.10: Load profiles comparing resource scheduling under EDF, LLF, and RHC to no scheduling base case

	EDF	LLF	RHC
Up Reserve: Energy	35.6	38.6	40.4
Up Reserve: Capacity	18.8	1.9	27.2
Down Reserve: Energy	32.5	28.8	37.5
Down Reserve: Capacity	-2.6	4.6	11.0

TABLE 3.3: Comparison of scheduling algorithms showing percentage reductions in 4 reserve cost metrics (compared to no coordination baseline).

that Figure 3.10(b) indicates an initial spike in reserve procurement. This is an artifact of using a zero-laxity reserve policy in conjunction with LLF. Here, the CM procures up reserves only when task laxities approach 0. As the laxities of all tasks are *identical* under LLF, serving a particular task with reserve generation implies all tasks must be served with reserves. In contrast, RHC-based scheduling is characterized by balanced reserve procurements over the entire operating window (Figure 3.10(c)).

Table 3.3 shows average percentage decreases in four cost metrics for each of these algorithms: up reserve energy, up reserve capacity, down reserve energy, and down reserve capacity. Resource scheduling clearly achieves significant reductions in reserve *energy* requirements. Scheduling under RHC outperforms LLF and EDF modestly in both energy metrics. In contrast, scheduling algorithms have significantly different performance in reserve *capacity* metrics where RHC is the decisive winner. We attribute the small up reserve capacity reductions achieved by LLF (1.9%) to large instantaneous reserve procurements resulting from the zero-laxity heuristic.

3.2.8 Deferrable load penetration

We now investigate the marginal benefit of scheduling additional deferrable loads. Specifically, we vary the proportion of total load that is deferrable (α) and quantify the reserve cost reductions

generated through RHC-based scheduling. Figure 3.11 shows percentage reductions in up reserve energy and capacity costs. Similar results for down reserves are omitted due to space constraints.

Figure 3.11 shows that additional amounts of deferrable load enable greater reductions in the up reserve metric. Moreover, the marginal benefit of having additional deferrable loads clearly decreases with increasing deferrable load penetration. This suggests that the primary impact of resource scheduling on reserve generation is evident even at low levels of deferrable load penetration. At high deferrable load penetrations, energy reserve costs stem from surpluses or shortfalls in total energy procured to meet load over the entire operating window. These imbalances are caused by errors in mean values of load and generation forecasts employed in ex-ante markets. While resource scheduling can address the reserve costs stemming from generation intermittency, it has limited impact at mitigating these reserve costs.

Figure 3.11 exhibits no clear pattern with respect to capacity cost reductions. We remark that reserve capacity computations are very sensitive to high frequency fluctuations of renewable generation. Resource scheduling decisions, made every 10 minutes, do not account for fluctuations in renewable generation at shorter time-scales. However, reserve capacity is computed based on actual generation data sampled at these shorter time scales. Accordingly, achieving large reserve capacity cost reductions requires increases in the frequency at which resource scheduling decisions are made.

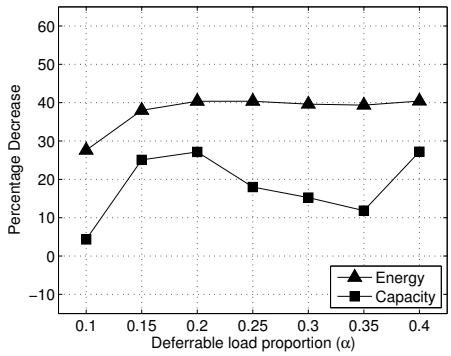

FIGURE 3.11: Percentage reductions in up reserve cost achieved by RHC-based scheduling at various levels of deferrable load penetration (α)

3.2.9 Load deferrability

Next, we investigate the impact of load deferrability, as captured by task laxity, on reserve requirements. In these simulations, all tasks have the same laxity (ϕ) upon arrival. For a fixed deferrable load proportion ($\alpha = 0.2$), we vary task laxity (ϕ) and compute the reserve energy and capacity cost reductions achieved through RHC-based scheduling.

Figure 3.12 shows percentage decreases in up reserve energy and capacity costs at various degrees of scheduling flexibility. Clearly, greater load deferrability has a compelling impact on reserve energy requirements. Specifically, up reserve energy costs can be reduced by a further 20% through additional scheduling flexibility. In contrast, increased flexibility only enables modest improvements in up reserve capacity. Similar results can be obtained for down energy and capacity reductions. These simulations demonstrate the value of additional scheduling flexibility in reducing reserve energy, and to a lesser extent, capacity cost reductions.

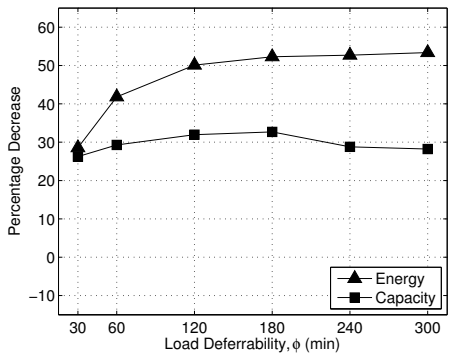

FIGURE 3.12: Percentage reductions in up reserve cost achieved by RHC-based scheduling at various levels of task deferrability (ϕ)

3.2.10 Practical considerations

Finally, we examine two practical considerations of implementing such coordinated scheduling algorithms: computational requirements and impact of switching on EVs.

Table 3.4 shows the total time taken, by each algorithm, to compute charging schedules over the entire operating interval. All simulations are performed using Matlab on a Dell XPS Studio 8100 machine with a processing speed of 2.93 GHz and 12.0 GB of RAM. We use the Gurobi optimization solver to compute the RHC schedules. These values are computed for different levels of deferrable load penetration (α). Clearly, the computational cost increases with the number of deferrable loads (α). Moreover, coordinated scheduling under RHC is far more computationally intensive than under EDF and LLF. This is expected as the RHC algorithm involves solving a sequence of QPs. We remark that while the total computation time over the entire operating interval (24 hours) for RHC is high (716s for $\alpha = 0.4$), the amount of computation done at each balancing time is much shorter. Indeed, RHC-based scheduling is computationally tractable for problems of the size described in this chapter.

α	**0.1**	**0.2**	**0.3**	**0.4**
EDF	0.06	0.12	0.17	0.23
LLF	0.06	0.11	0.15	0.20
RHC	41	124	399	716

TABLE 3.4: Comparison of scheduling algorithms showing computation time (s) required for different levels of deferrable load penetration (α).

Coordinated scheduling algorithms may cause increased on-off switching of deferrable loads. Frequent switching of EV batteries is undesirable as it could adversely impact battery life and long-term operation [117]. Figure 3.13 shows a cumulative distribution function on the maximum number of on-off switches under RHC-based scheduling empirically computed over 100 test cases. One can see that most tasks (50%) are operated continuously once scheduled, and a small number (10%) are switched more than four times. Results for EDF- and LLF-based scheduling, omitted due to space constraints, are qualitatively similar to Figure 3.13. As compared to RHC, EDF exhibits less switching while LLF results in more.

Sporadic battery charging may have negative effects on battery state of health, though less so than frequent charging and discharging. If the state of health impacts can be quantified, it would be straightforward to extend the RHC framework to penalize frequent cycling in the objective function. Future work involves addressing these issues.

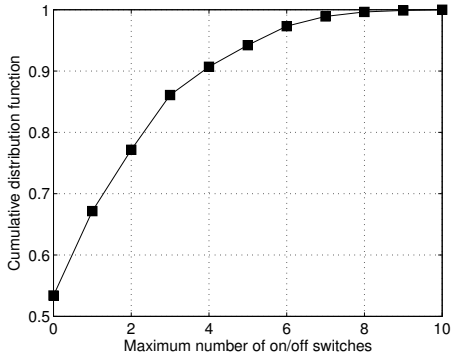

FIGURE 3.13: Cumulative distribution function on the maximum number of on-off switches under RHC-based scheduling.

3.2.11 Conclusions and Future Work

We have examined scheduling algorithms for coordinated aggregation of deferrable loads and storage. We modeled deferrable loads and storage as tasks and devices respectively. We then developed EDF and LLF scheduling algorithms for resource scheduling, and a laxity-based heuristic for reserve

scheduling. We offered a novel cost metric for RHC-based scheduling that explicitly handles operating reserve costs. These are the costs associated with managing non-contingency energy imbalances. Through simulation studies, we find that scheduling under *any* of these algorithms reduces reserve energy costs, while RHC-based scheduling also reduces reserve capacity costs. Most importantly, we conclude that the benefits of *coordinated aggregation can be realized from modest levels of both deferrable load participation and flexibility.*

Realizing the benefits of resource aggregation requires the convergence of technology infrastructure to support coordination, and market mechanisms to elicit and reward resource participation. How should loads be compensated for their flexibility? Who pays for storage? What are fair pricing mechanisms for flexibility? Are these sufficient to incentivize participation? These economic aspects deserve detailed study. Implementation of CM algorithms requires an underlying communication/computation architecture. How much computation, bandwidth, latency, and reliability is required? Should the control architecture be centralized or (partially) distributed? These essential technological aspects also require exploration.

Chapter 4

A Market for Flexibility: Deadline Differentiated Pricing

Clearly, there is an opportunity to transform the current operational paradigm, in which supply is tailored to follow demand, to one in which *demand is capable of reacting to variability in supply* – an approach which is generally referred to as demand response (DR) [2], with the primary challenge being the *reliable extraction of the desired response* from participating demand resources on time scales aligned with traditional bulk power balancing services.

Today, most demand response programs are largely limited to peak shaving and contingency-based applications, with the two most common economic paradigms for customer recruitment and control being: (1) *direct load control* under which the system operator or utility companies procure the capability of load adjustment through a forward transaction (e.g. call options for interruptible load) and (2) *indirect load control* under which consumers themselves adjust energy consumption according to time-varying prices (e.g. through time of use or dynamic pricing). Dynamic (real-time) pricing is often identified as a priority for the implementation of indirect load control [21]. Despite its potential to improve the economic efficiency of electricity markets [9, 22, 25, 41], it subjects consumers with the risk of paying high peak prices, and can result in unreliable aggregate (consumer) response [38]. Closer to the focus of this paper on delay-tolerant demand, a recent empirical study [40] indicates that for plug-in hybrid electric vehicle (PHEV) charging, dynamic pricing may perform worse than a flat-rate tariff (in term of fuel costs and emissions impacts). A recent market monitoring report expresses the concern of California Independent System Operator (CAISO) on dynamic real-time pricing:

"While there are many economists that are enthusiastic about DR for all consumers, we are not aware of a reported success of real-time pricing for a big, heterogeneous population area that could serve as a benchmark. Mobilizing retail level demand side flexibility to reduce operating and investment cost in the electricity sector by employing smart grid technologies and market mechanism is still regarded as work in progress."

In short, performance based on indirect load control (e.g., dynamic real-time pricing) may not provide the level of assurance required to avoid the use of conventional generation to manage the electric power system. In the following sections, we propose a market framework that incentivizes end-use consumers to truthfully reveal their preferences (e.g., their maximum capability to delay energy consumption), which are important information needed for efficient implementation of direct load control. The proposed market framework centers on the provisioning of deadline differentiated energy services to end-use customers, whose quality differentiation maps to flexibility in the family of feasible power profiles capable of satisfying said service. By offering a family of differentiated services with forward contracts the coordinating entity implicitly purchases the right to manage in real-time the delivery of power to participating customers. In this way, it can align its operational requirements with the vast heterogeneity in end-use customer needs.

4.1 A Market for Quality-Differentiated Electric Power Service

Flexibility in consumption can be interpreted as a *continuum of feasible power profiles* capable of preserving the end-use function of a demand resource. A basic question, is how to design a market that enables a coordinating entity the ability to "extract this flexibility" for execution of real-time control applications – e.g. balancing variability in renewable supply?

One possibility resides in the construction of a market for *quality-differentiated electric power services*, where the price to a consumer for receiving a particular service is a monotonic function of the desired quality-of-service (QoS). Naturally, a reduction in QoS is accompanied by a reduction in price. For example, in the concrete setting of deferrable loads, deadline would be a natural specification of QoS. The longer a customer is willing to delay the receipt of a specified quantity of energy, the less that customer pays (per-unit) for said energy. As a QoS specification maps directly to a set of feasible power profiles capable of servicing a load, the aggregator can imbed its extraction of flexibility from the demand-side in its delivery of differentiated services with a guaranteed QoS to participating consumers.

The general concept of service differentiation is not new [29, 30]. Many have studied the problem of centrally coordinating the response of a collection of loads for load-following or regulation services

– all while ensuring the satisfaction of a pre-specified QoS to individual resources [11, 14, 28, 31, 37, 43]. However, there has been little work in the way of designing market mechanisms that endogenously price the flexibility being offered by the demand side, while incentivizing consumers to truthfully reveal their preferences (e.g. the capability to delay energy consumption) to the operator. Several classic [13, 44] and more recent [6] papers have explored the concept of *reliability-differentiated pricing* of interruptible electric power service, where the consumer takes on the risk of interruption in exchange for a reduction in the price for energy. Beyond the apparent issues of moral hazard and difficulty in auditing the delivered reliability of such services, the primary drawback of such an approach stems from the explicit transferal of quantity risk to the demand side, as it requires participating consumers to plan their consumption in the face of uncertain supply. From a consumer's perspective, this amounts to solving a nontrivial problem of sequential decision making under uncertainty.

With the aim of alleviating the aforementioned challenges, we propose a novel forward market for *deadline differentiated energy services*, where consumers consent to deferred service of pre-specified loads in exchange for a reduced per-unit energy price. The forward market for deadline differentiated energy service is described as a three-step process. Time is assumed to be discrete with periods indexed by $k = 0, 1, 2, \ldots, N$.

Step 1 (Scheduling and Pricing). Prior to period $k = 0$, the supplier announces a *mechanism* $(\pi, \boldsymbol{\kappa})$ consisting of both a *scheduling policy* π (cf. its formal definition in Section 4.3), and *pricing scheme* $\boldsymbol{\kappa} = (\kappa_1, \ldots, \kappa_N)$ that maps the aggregate demand bundle $\mathbf{x}$ (cf. its definition in Step 2) into a menu of *deadline-differentiated prices*,

$$p_k = \kappa_k(\mathbf{x}), \qquad k = 1, \ldots, N. \tag{4.1}$$

The price menu stipulates a per-unit price p_k (\$/kWh) for energy guaranteed delivery by period k. At the heart of the mechanism design is the restriction that *prices are nonincreasing in the deadline*. Namely, the longer a customer is willing to defer her consumption, the less she is required to pay. We will use $\mathcal{P} = \{\mathbf{p} \in \mathbb{R}_+^N \mid p_1 \geq p_2 \geq \ldots \geq p_N\}$ to denote the set of feasible deadline-differentiated price bundles.

Step 2 (Purchasing). Each consumer then purchases a bundle, $\mathbf{a} = (a_1, \ldots, a_N)^\top \in \mathbb{R}_+^N$ (kWh), of deadline-differentiated energy quantities, where a_k denotes the quantity of energy guaranteed delivery to the consumer by the deadline k. Before a stage $k = 1, \ldots, N$, the total quantity (of energy) delivered to a consumer who purchases a bundle $\mathbf{a}$ is at least $\sum_{t=1}^{k} a_t$, and is at most $\sum_{t=1}^{N} a_t$ (which is the consumer's aggregate demand over the entire horizon). And naturally, a consumer requesting a bundle $\mathbf{a}$ is required to pay $\mathbf{p}^\top \mathbf{a}$.

We denote the *aggregate demand bundle* (summed over all individual consumer bundles) by $\mathbf{x} = (x_1, \ldots, x_N)^\top \in \mathbb{R}_+^N$. Here, x_k denotes the aggregate quantity requiring delivery by deadline k. The mechanism design is such that price bundle $\mathbf{p}$ is determined by the aggregate demand bundle $\mathbf{x}$, according to the pricing scheme $\boldsymbol{\kappa}$ announced prior to period $k = 0$. A pricing scheme $\boldsymbol{\kappa}$, which has been defined in Eq. (4.1), is therefore a mapping from $\mathbb{R}_+^N$ to $\mathcal{P}$.

Step 3 (Delivery). Finally, the supplier must deliver the requested aggregate demand bundle $\mathbf{x}$ according to the announced scheduling policy π. Essentially, a scheduling policy is said to be feasible if it delivers each individual consumer's requested energy bundle by its corresponding deadlines. The supplier is assumed to have two sources of generation from which he can service demand: *intermittent* and *firm supply*.

(a) *Intermittent generation.* An intermittent supply modeled as a discrete time random process $\mathbf{s} = (s_0, s_1, \ldots, s_{N-1})$ with known joint probability distribution. Here, $s_k \in \mathcal{S} \subset \mathbb{R}_+$ (kWh) denotes the energy produced during period k. The intermittent supply is assumed to have *zero marginal cost.*

(b) *Firm generation.* A firm supply with *constant marginal cost* $c_0 > 0$, assumed to be time-invariant.

A basic challenge addressed in this paper is the design of such a market mechanism, which implements truth-telling at a dominant strategy equilibrium across the population of consumers, while maximizing social welfare between the supplier and consumers.

4.1.1 Summary and Contributions

Before continuing, we provide here a roadmap of the paper together with a summary of our main contributions.

(a) We provide a stylized (yet quite reasonable) model for both delay-tolerant electricity demand and a supplier with both firm and intermittent electricity generation in Sections 4.2 and 4.3, respectively.

(b) We formulate the supplier's scheduling problem as a constrained stochastic optimal control problem and prove optimality of the earliest-deadline-first (EDF) scheduling policy, in terms of minimizing the expected cost of firm supply over all feasible scheduling policies. As a corollary, this result enables the explicit characterization of the supplier's marginal cost curve

in Theorem 4.9, which specifies the bundle of deadline-differentiated energy quantities $\mathbf{x}$ that said supplier would like to provide at a given price bundle $\mathbf{p}$.

(c) It is reasonable to expect that the supplier cannot observe the true deadline of each individual consumer. The presence of asymmetric information may lead to significant welfare loss, if consumers have incentive to lie about their true deadlines. In Section 4.4.1, we show that a mechanism consisting of an EDF scheduling policy, together with a pricing scheme reflecting the supplier's marginal cost (cf. its characterization in Theorem 4.9), implements consumers' *truth-telling* behavior in dominant strategies.

(d) Under an additional mild assumption on each consumer's marginal valuation of energy (cf. Assumption 4.4.2), we show in Section 4.4.2 that a marginal cost pricing scheme, in combination with EDF scheduling, induces an efficient competitive equilibrium between a population of truth-telling consumers and a price-taking supplier. In other words, the pricing scheme characterized in Theorem 4.9, which has been shown to be incentive compatible in Section 4.4.1, results in an efficient market equilibrium, at which the social welfare (the sum of aggregate consumer surplus and supplier profit) is maximized.

(e) In Section 4.4.2, we discuss briefly the possible implementation of the proposed deadline differentiated pricing mechanism, through both a mechanism design and a (competitive) market equilibrium approach.

4.1.2 Relation to the Literature

There have emerged several papers concerned with the design of incentives for delay-tolerant electric load, e.g. the charging of plug-in hybrid electric vehicles (PHEVs). [23] propose a similar idea of offering consumers a discounted per-unit electricity price in exchange for delaying their energy consumption. However, the focus of [23] is not on pricing, but rather on the problem of optimal scheduling faced by the operator, who is assumed to have full information about consumers – e.g. knowledge of their deadlines. Closer to the present paper, [35] propose a greedy online mechanism for electric vehicle charging, which is shown to be incentive compatible and achieve a bounded (worst-case) competitive ratio. However, these theoretical results require a VCG-type (discriminatory-price) payment scheme, as opposed to the uniform-price mechanism proposed in this paper. More strongly, they impose an additional assumption that consumers cannot report a false deadlines exceeding their true underlying deadlines. This substantially simplifies the problem of designing an incentive compatible pricing scheme when consumers are permitted to report arbitrary deadlines – the setting considered in the present paper.

This work is also related to two streams of literature in operations research – the first of which, applies the principal-agent paradigm to study the design of compensation schemes for both inventory and production systems [15, 32], and the second on priority pricing for queuing systems with multiple user classes [10, 24, 26, 34, 45]. The former body of literature deals primarily with problems resulting from systems in which agents' possess a hidden action (that cannot be observed by the principle; also known as moral hazard). Closer to the present paper, the latter stream of literature addresses the scenario in which hidden information is privately held by users in a queuing system. More specifically, this body of work studies the (joint) pricing and scheduling for a queuing system with a heterogeneous population of customers, who are in general characterized by differing delay costs (per unit of time) and service time distributions. While the main result of this paper (a full characterization of a joint pricing and scheduling scheme that is both socially optimal and incentive compatible) is similar in spirit to the results derived in the aforementioned literature on priority pricing for queuing systems, it differs in several important ways.

In the literature on priority pricing for queuing systems, it is usually assumed that there is an infinite quantity of potential customers, who enter the (queuing) system until they are indifferent between balking and entering – i.e. at an equilibrium the marginal surplus (of the last consumer) is zero [10, 24, 26, 45]. This is not the case for retail electricity markets, where there are a large but fixed number of consumers, whose marginal valuation on electricity consumption is commonly higher than the electricity price.

As a result, in this paper we consider a fixed population of consumers, whose types are distributed according to a certain distribution that is exogenous and unknown to the service provider. This setting enables our model to account for the surplus of all consumers and not only those who are admitted into the system. It does, however, generate additional technical challenges, as the aggregate demand depends not only on the pricing scheme (determined by the service provider), but also on the distribution of consumer types.[1]

There are also important technical differences between our model and the queuing system models. On the demand side, while customers in a queueing system are nominally distinguished by their (uniform) delay cost per unit time, our model considers the setting in which each customer is capable of delaying her energy consumption until a specific deadline without utility loss. This is motivated by the fact that an electricity consumer's delay cost is in general not uniform. On the supply side, the operating cost of a service provider in queuing systems is deterministic and known

[1] Note that in the literature on priority pricing for queuing systems, it is usually assumed that the service provider knows the demand curve, and therefore the arrival rates of (different types of) customers are determined by the service-differentiated prices set by the service provider.

[24, 26, 34, 45]. While in our model, we treat the operating cost of a service provider as random to account for the intermittency of supply derived from renewable generation.

Finally, we point out the relation of the present work to recent literature exploring the design of revenue-maximizing pricing mechanisms capable of inducing truth revelation of consumers' privately known deadlines for purchasing a single object. [1] studies the problem of revenue maximization in queueing systems, where consumers hold private information about their preferences. Closer to the present paper, [27] studies the revenue-maximizing sale of a single object to buyers with differing deadline preferences for buying. Considering a special case with two periods and two buyers, the author characterizes a pricing mechanism that is revenue-maximizing among all incentive compatible mechanisms.

4.2 Delay-Tolerant Demand Model

We consider a model involving a continuum of consumers, indexed by $i \in [0, 1]$. Since each individual consumer's action has no influence on the price, she will act as a price taker. This setting is reasonable in the context of retail electricity markets, as each consumer herself is too small to influence the price. We note that the price-taking assumption is somewhat standard in the literature on mechanism design for queueing systems, where each consumer is assumed to be either infinitesimally small [1, 10, 26, 45] or price-taking [24].

Consider a utility model yielding a consumer preference ordering on deadlines. It is natural to assume that the *utility* derived from electricity consumption is non-increasing in the delivery deadline k. Namely, the longer the delay in consumption, the larger the potential loss of utility. In particular, we assume that each consumer has a *single deadline preference*. More specifically, a consumer with deadline preference k derives no disutility from deferring consumption until deadline k and derives zero utility for any consumption thereafter. This assumption is reasonable for electric loads such as plug-in electric vehicles (PEVs), dish washers, and laundry machines, as customers commonly require only that such loads are fully executed before a specific time.

Naturally then, we define the **type of consumer** i as a triple, $\theta_i = (k_i, R_i, q_i)$, consisting of her *deadline* $k_i \in \mathbb{N}$, *marginal utility* $R_i \in \mathbb{R}_+$, and *maximum demand* $q_i \in \mathbb{R}_+$. Clearly, consumer i's utility function depends only on her type θ_i, and is assumed to satisfy the following conditions.

Assumption 4.2.1. The utility received by a type-θ_i consumer $U_{\theta_i}(y)$ depends only on the consumer's cumulative consumption by her deadline k_i, y. The utility function U_{θ_i} is non-negative and

non-decreasing over $[0, q_i]$ with

$$U_{\theta_i}(y) \leq yR_i, \qquad \forall\, y \in [0, q_i],$$

where $R_i = U_{\theta_i}(q_i)/q_i$. Further, $U_{\theta_i}(y) = U_{\theta_i}(q_i)$ for all $y \geq q_i$. □

Note that the marginal utility R_i associated with a consumer type θ_i is defined as the ratio of the maximum utility $U_{\theta_i}(q_i)$ to the maximum demand q_i. The above assumption simply requires that over $[0, q_i]$, the utility function is below the straight line that connects the two points $(0, 0)$ and $(q_i, U_{\theta_i}(q_i))$, and then becomes flat over $[q_i, \infty)$. We note that Assumption 4.2.1 holds for any utility function that is non-decreasing and convex over $[0, q_i]$. On the other hand, incentive compatibility may fail to hold for utility functions that are concave over $[0, q]$ (cf. the discussion in Remark 4.15). In the following example we discuss some forms of utility functions that satisfy Assumption 4.2.1.

Example 4.1. *The following piecewise linear utility function satisfies Assumption 4.2.1*

$$U_{\theta_i}(y) = R_i \cdot \min\{y, q_i\}.$$

Indeed, every utility function that lies below the above piecewise linear function satisfies Assumption 4.2.1, for example, utility functions of the following form

$$U_{\theta_i}(y) = \begin{cases} 0, & \text{if } \ 0 \leq y \leq T_i, \\ \bar{R}_i(y - T_i), & \text{if } \ T_i < y \leq q_i, \\ U_{\theta_i}(q_i), & \text{if } \ q_i < y, \end{cases}$$

where $T_i > 0$ is the minimum amount of energy required by the consumer (e.g., in order to fulfill her next trip). The consumer obtains zero utility if the battery level is below this threshold, and her "marginal utility" R_i is given by $\bar{R}_i(q_i - T_i)/q_i$. In Fig. 1 we plot four utility functions that satisfy Assumption 4.2.1. □

We use Θ to denote the set of all possible types, which is assumed to be finite. Let $\rho : \Theta \to [0, 1]$ denote the **distribution of consumer types** over the space Θ. In other words, for every $\theta \in \Theta$, there is a $\rho(\theta)$ fraction of consumers of type θ. It follows that $\sum_{\theta \in \Theta} \rho(\theta) = 1$.

Definition 4.1 (Consumer action)**.** The **action** of a consumer is a vector $\mathbf{a} = (a_1, \ldots, a_N)^\top \in \mathbb{R}_+^N$, where a_k denotes the amount of energy guaranteed delivery by deadline k. The maximum amount

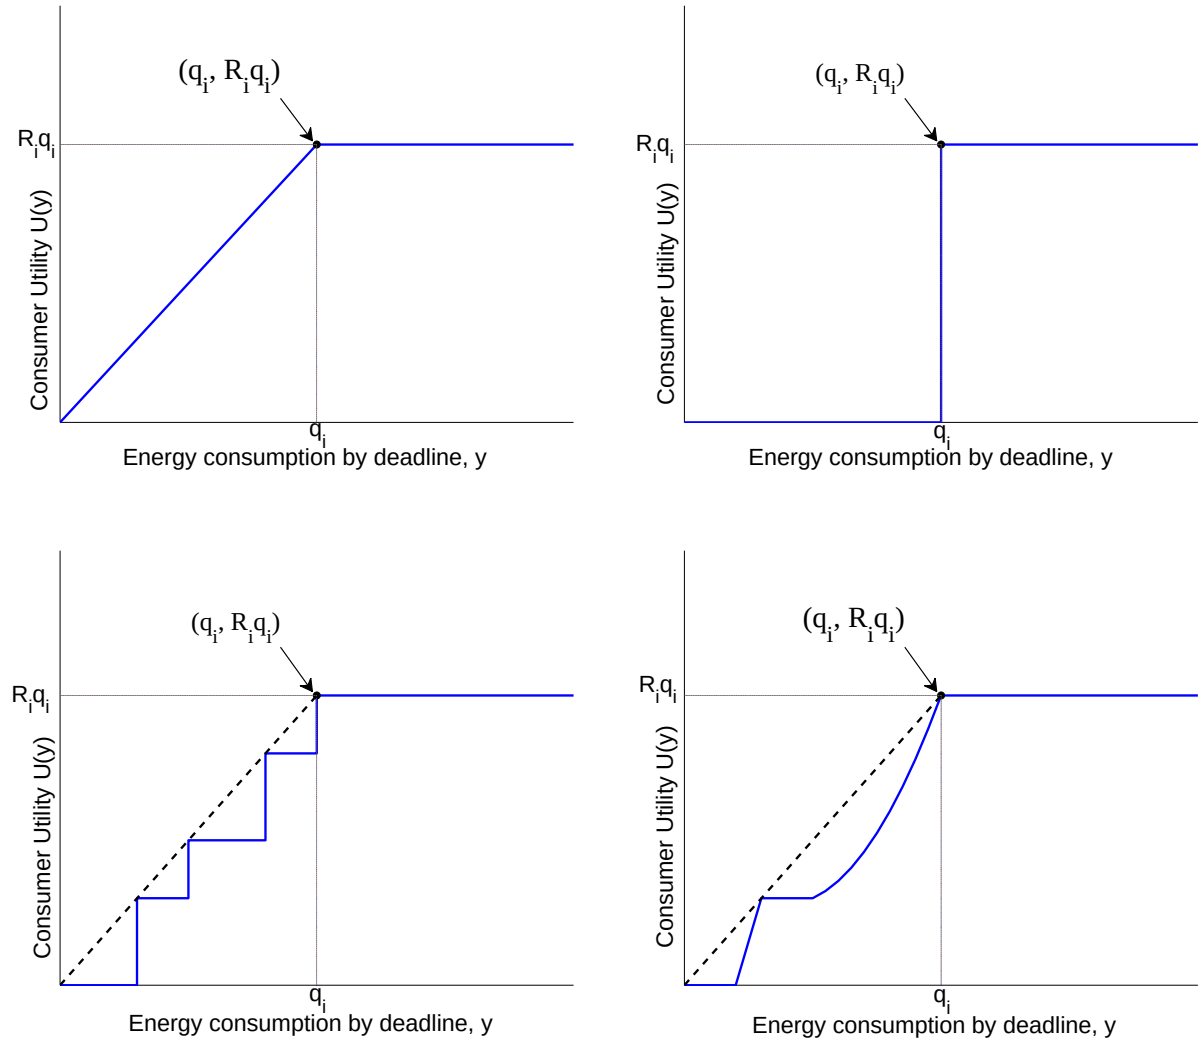

FIGURE 4.1: Plot of four utility functions that satisfy Assumption 4.2.1.

of energy any consumer can request is $Q = \max_{(k,R,q)\in\Theta}\{q\}$. Hence, each consumer's *action space* is restricted to $\mathcal{A} = \left\{\mathbf{a} \in \mathbb{R}_+^N \mid \sum_k a_k \leq Q\right\}$. ■

Note that in the above definition we have that $q \leq Q$ for every type $\theta = (k, R, q)$. In other words, it is feasible for every consumer to request her maximum demand q. Given a fixed scheduling policy and pricing scheme, a consumer's **strategy** $\varphi : \Theta \to \mathcal{A}$ maps her type into an action. In the proceeding analysis, we will be concerned with identifying conditions on both the scheduling policy and pricing scheme that lead to efficient allocations, while inducing consumers to truthfully reveal their underlying deadline preferences. A truth-telling strategy is defined as follows.

Definition 4.2 (Truth-telling)**.** A consumer of type $\theta = (k, R, q)$ is defined to be **truth-telling** if she requests q units of energy at her true deadline k, and nothing else. Formally, a strategy $\mathbf{a}^* = \varphi^*(\theta)$ is said to be truth-telling if it satisfies Eq. (4.2).

$$a_j^* = \begin{cases} q, & j = k, \\ 0, & j \neq k. \end{cases} \tag{4.2}$$

We note that under an arbitrary scheduling policy and pricing scheme, it is possible that a consumer's best response may differ from the truth-telling action defined in (4.2). Our aim is to provide an explicit characterization of a scheduling policy and pricing scheme that implement truth-telling as a dominant strategy for every consumer i. Given the collection of consumer types $\boldsymbol{\theta} = \{\theta_i\}_{i\in[0,1]}$ and a strategy profile $\boldsymbol{\varphi} = \{\varphi_i\}_{i\in[0,1]}$, the aggregate demand bundle is given by the mapping

$$\mathbf{x} = \mathbf{d}(\boldsymbol{\theta}, \boldsymbol{\varphi}) = \int_{i\in[0,1]} \varphi_i(\theta_i)\eta(di), \tag{4.3}$$

where η is the Lebesgue measure defined over $[0, 1]$, and $\mathbf{d} = (d_1, \ldots, d_N)$ maps $(\boldsymbol{\theta}, \boldsymbol{\varphi})$ into an N-dimensional nonnegative vector. Note that we have implicitly assumed that for every $\boldsymbol{\theta}$, the function $\boldsymbol{\varphi} = \{\varphi_i(\theta_i)\}_{i\in[0,1]}$ is Lebesgue integrable in i. This assumption holds, for example, under a symmetric strategy profile according to which all consumers of the same type take the same action[2].

In particular, the aggregate demand under the *truth-telling strategy profile* $\boldsymbol{\varphi}^* = \{\varphi_i^*\}_{i\in[0,1]}$ (cf. Definition 4.2) simplifies to

$$x_j^* = d_j(\boldsymbol{\theta}, \boldsymbol{\varphi}^*) = \sum_{\theta\in\Theta} q \cdot \mathbf{1}_{\{j=k\}} \cdot \rho(\theta), \qquad j = 1, \ldots, N, \tag{4.4}$$

where $\theta = (k, R, q)$ and $\mathbf{1}_{\{\cdot\}}$ denotes the indicator function.

4.2.1 Consumer Surplus

We are now in a position to characterize the expected surplus derived by a consumer, which depends on (i) her own type and strategy, (ii) the remaining population's type and strategy profile, (iii) the pricing scheme, and of course, (iv) the scheduling policy employed by the supplier to deliver the requested demand bundles. Before formally characterizing an individual consumer's surplus in our

[2]We note that it is usually assumed in the literature (on priority pricing for queuing systems) that consumers use a symmetric strategy profile [10, 24, 26, 45].

model, we first introduce pertinent notation and assumptions. Given an aggregate demand bundle $\mathbf{x}$, we define the random variable $\omega^{\pi}_{k,i}(\mathbf{x},\mathbf{a})$ as denoting the amount of energy delivered to consumer i by stage k given her requested bundle $\mathbf{a}$. This random variable, which naturally depends on the scheduling policy π employed by the supplier, is formally defined in Appendix B.1 (cf. Eq. (B.1)). Moreover, we have allowed the supply $\omega^{\pi}_{k,i}(\mathbf{x},\mathbf{a})$ to depend explicitly on the consumer index i, as the supplier may employ a scheduling policy that is consumer index-dependent. This dependency on the scheduling policy is made precise in Section 4.3.

We assume that the requested quantities are always supplied by their corresponding deadlines and the total quantity delivered to a consumer never exceeds said consumer's total demand. This is made precise in the following assumption.

Assumption 4.2.2. Given the supplier's delivery commitments, we require that for each consumer $i \in [0,1]$ taking action $\mathbf{a} \in \mathcal{A}$ that

$$\sum\nolimits_{t=1}^{k} a_t \;\leq\; \omega^{\pi}_{k,i}(\mathbf{x},\mathbf{a}) \;\leq\; \sum\nolimits_{t=1}^{N} a_t, \quad \text{almost surely,}$$

for all $1 \leq k \leq N$ and aggregate demand bundles $\mathbf{x} \in \mathbb{R}^N_+$. □

Given such a market mechanism, it is important to understand when a consumer has incentive to misreport her underlying deadline preference. Consider a consumer i of type $\theta_i = (k, R, q)$ facing a particular scheduling policy and pricing scheme satisfying all prior assumptions. Because of monotonicity of prices and the service guarantee provided by Assumption 4.2.2, said consumer has no incentive to request any quantity of energy before her true deadline k. However, if the potential price differences among different deadlines are large enough, a consumer may have an incentive to report a false later deadline if early delivery is likely (i.e. with high probability) under the specified scheduling policy π. Intuitively, a consumer might have incentive to report a false bundle if the reduction in total expenditure derived by requesting quantities with later deadline exceeds the expected loss of utility incurred by a shortfall in the delivered amount by her true deadline k. As such, in designing a pricing scheme and scheduling policy, it is important that the price differential across deadlines balance the likelihood of early supply under the specified scheduling policy.

To formally define and study the incentive compatibility of a particular mechanism (i.e. pricing scheme and scheduling policy), we now define the expected surplus derived by a consumer under a particular strategy. Given a scheduling policy π and a pricing scheme $\boldsymbol{\kappa}$ employed by the supplier, and the types of all consumers $\boldsymbol{\theta} = \{\theta_i\}_{i\in[0,1]}$, consumer i's **expected surplus** under a strategy profile $\boldsymbol{\varphi} = \{\varphi_i\}_{i\in[0,1]}$ is given by

$$v^{\pi}_i(\boldsymbol{\theta},\boldsymbol{\varphi},\boldsymbol{\kappa}) \;=\; \mathbb{E}\left\{U_{\theta_i}\left(\omega^{\pi}_{k,i}(\mathbf{x},\mathbf{a})\right)\right\} \;-\; \boldsymbol{\kappa}(\mathbf{x})^{\top}\mathbf{a}, \tag{4.5}$$

where k is consumer i's true deadline, $\mathbf{x} = \mathbf{d}(\boldsymbol{\theta}, \boldsymbol{\varphi})$ is the aggregate demand bundle, and $\mathbf{a} = \varphi_i(\theta_i)$ is the action taken by consumer i. Here, expectation is taken over the random variable $\omega^{\pi}_{k,i}(\mathbf{x}, \mathbf{a})$.

Since each individual consumer is assumed to be infinitesimal in size, and thus have no influence on the price, a scheduling policy π together with a pricing scheme $\boldsymbol{\kappa}$ indeed defines a game for the (price-taking) consumer population, with each individual consumer's payoff expressed in (4.5). We note that this is an **aggregative game**, where the payoff function of each player depends on the strategy profile (used by all players) only through the sum of their actions – namely, the aggregate demand bundle $\mathbf{x} = \mathbf{d}(\boldsymbol{\theta}, \boldsymbol{\varphi})$. We can therefore denote the payoff to consumer i under the aggregate demand bundle $\mathbf{x}$ as

$$V^{\pi}_i(\theta_i, \varphi_i, \mathbf{x}, \boldsymbol{\kappa}) = \mathbb{E}\left\{U_{\theta_i}\left(\omega^{\pi}_{k,i}(\mathbf{x}, \mathbf{a})\right)\right\} - \boldsymbol{\kappa}(\mathbf{x})^{\top}\mathbf{a}, \tag{4.6}$$

where $\mathbf{a} = \varphi_i(\theta_i)$. Moreover, it follows from Assumption 4.2.2 that under the truth-telling strategy φ^*_i, the payoff derived by consumer i simplifies to the deterministic quantity

$$V^{\pi}_i(\theta_i, \varphi^*_i, \mathbf{x}, \boldsymbol{\kappa}) = U_{\theta_i}(q) - \kappa_k(\mathbf{x})q, \tag{4.7}$$

where $\theta_i = (k, R, q)$. It is important to note that the expression in (4.7) does not depend on the types and strategies of the other consumers, or the scheduling policy used by the supplier (as long as Assumption 4.2.2 is respected).

Nash equilibrium may not be a plausible solution concept to explore for this game, as it requires each individual consumer to have information regarding the other consumers' types, as well as knowledge of the distribution of $\omega^{\pi}_{k,i}(\mathbf{x}, \mathbf{a})$, which in turn depends on the distribution of the intermittent supply process $\mathbf{s}$. We therefore focus on a much stronger solution concept – namely, the **dominant strategy equilibrium** of the game. A strategy φ_i is a dominant strategy for consumer i of type θ_i, if it maximizes her expected payoff regardless of the actions taken by the other consumers. We have the following definition.

Definition 4.3 (Dominant strategy)**.** A strategy φ_i is said to be a **dominant strategy** for consumer i of type θ_i, if it maximizes her expected payoff regardless of the actions taken by the other consumers, i.e.

$$V^{\pi}_i(\theta_i, \varphi_i, \mathbf{x}, \boldsymbol{\kappa}) \geq V^{\pi}_i(\theta_i, \varphi'_i, \mathbf{x}, \boldsymbol{\kappa}), \qquad \forall\, \varphi'_i, \qquad \forall\, \mathbf{x} \in \mathbb{R}^N_+.$$

Moreover, a mechanism that implements truth-telling in dominant strategies is said to be **incentive compatible**, which we formally define in Definition 4.12. Surprisingly, we show in Section 4.4 that a mechanism consisting of earliest-deadline-first (EDF) scheduling in combination with a pricing

scheme reflecting the supplier's marginal cost is both incentive compatible and induces an efficient competitive equilibrium between the supplier and consumers.

Remark 4.4 (An alternative formulation). One can define an alternative game among the consumer population by leveraging on an additional assumption that the supplier has a priori knowledge of the aggregate demand curve, but not necessarily the type of each individual consumer. We note that this is the typical setting considered in the literature on priority pricing for queuing systems [24, 26, 34, 45]. In this setting, the supplier can simply announce a fixed price bundle $\mathbf{p} \in \mathcal{P}$ (as opposed to a pricing scheme $\boldsymbol{\kappa}(\cdot)$), which also defines a game between the consumers. Our results on incentive compatibility outlined in Section 4.4 imply that with the knowledge of the aggregate demand curve, the supplier can simply set the price bundle to coincide with the intersection of the marginal cost supply and demand curves, thus inducing consumers' truth-telling behavior at an efficient competitive equilibrium between the supplier and consumers. We discuss this alternative approach in more detail in Section 4.4.2.2.

4.3 Supply Model

As one of the primary thrusts of this paper being the *characterization of competitive equilibria* between supply and demand and their efficiency properties, we now consider the role of a *price-taking supplier*, whose aim is to maximize his expected profit under a given price bundle. The expected profit derived by a supplier amounts to the revenue derived from the sale of a bundle of deadline differentiated energy quantities less the cost of firm supply required to service said bundle. In characterizing the supply curve under price taking behavior, the supplier's objectives are two-fold:

Scheduling. Determine a *causal scheduling policy* to allocate the realized intermittent supply across the deadline differentiated consumer classes, in order to *minimize the expected cost of firm supply* required to ensure satisfaction of the aggregate demand bundle.

Pricing. Fixing the optimal scheduling policy, determine a *supply curve* that specifies the bundle of energy he is willing to supply at every price bundle. In Theorem 4.9, we explicitly characterize the supplier's *marginal cost supply curve*, which is used to construct an efficient competitive equilibrium between supply and demand (cf. Definition 4.16).

Remark 4.5. Observe that the problem of characterizing a marginal cost supply curve amounts to the explicit solution of a *two-stage stochastic program*, whose expected recourse cost is the optimal value of a *constrained scheduling problem* parameterized by the aggregate demand bundle.

4.3.1 Optimal Scheduling

When considering the problem of scheduling, it's important to distinguish between *intra-class* and *inter-class* scheduling. Loosely speaking, an inter-class scheduling policy (denoted by σ) refers to manner in which available supply is allocated across the deadline differentiated demand classes, while an intra-class scheduling policy (denoted by ϕ) refers to the manner in which available supply is allocated across customers within a given demand class. One can readily see that, given a fixed aggregate demand bundle $\mathbf{x}$, the supplier's expected profit depends only on the inter-class scheduling policy and is invariant under the family of feasible intra-class scheduling policies. However, as one might easily notice, the intra-class scheduling policy has a direct effect on consumer purchase decisions inasmuch as it affects the distribution on each individual's random supply $\omega^{\pi}_{k,i}(\mathbf{x}, \mathbf{a})$, where $\pi = (\sigma, \phi)$ denotes the joint inter and intra-class scheduling policy employed by the supplier.

An important assumption we will make on the supply side is that there is no upper bound on the amount of energy the supplier can deliver within any given time period. This assumption, although appearing strong, is not far from practice, as batteries with high power to energy ratios are rapidly becoming available for electric vehicles. For example, the lithium-ion titanate batteries are capable of recharging to 95% of full capacity within approximately ten minutes [8, 17]. We also note that fast (DC) chargers that can fully charge an electric vehicle within half an hour are being installed in public locations, e.g., parking lots, shopping centers, hotels, theaters, and restaurants [47].

4.3.1.1 Inter-class scheduling policies

We now characterize the optimal inter-class scheduling policy as a solution to a constrained stochastic optimal control problem. First, we define the system **state** at period k as the pair $(\mathbf{z}_k, s_k) \in \mathbb{R}^N_+ \times \mathbb{R}_+$, where the vector $\mathbf{z}_k$ denotes the residual demand requirement of the original aggregate demand bundle $\mathbf{x}$ after having been serviced in previous periods $0, 1, \cdots, k-1$. Define as the **control input** the vectors $\mathbf{u}_k, \mathbf{v}_k \in \mathbb{R}^N_+$, which denote (element-wise in j) the amount of intermittent and firm supply allocated to demand class j at period k, respectively. Naturally then, the state of residual demand evolves according to the discrete time **state equation**:

$$\mathbf{z}_{k+1} \;=\; \mathbf{z}_k - \mathbf{u}_k - \mathbf{v}_k, \qquad k = 0, \ldots, N-1, \tag{4.8}$$

where the process is initialized with $\mathbf{z}_0 = \mathbf{x}$. The delivery deadline constraints manifest in a sequence of nested constraint sets $\mathbb{R}^N_+ \supset \mathcal{Z}_1 \supseteq \mathcal{Z}_2 \supseteq \cdots \supseteq \mathcal{Z}_N = 0$ converging to the the origin,

where the set $\mathcal{Z}_k$ characterizes the **feasible state space** at stage k. More precisely,

$$\mathcal{Z}_k \;=\; \{\mathbf{z} \in \mathbb{R}_+^N \mid z^j = 0, \;\; \forall\, j \leq k\}.$$

In other words, the feasible state space is such that each demand class is fully serviced by its corresponding deadline. We define as the **feasible input space** at stage k the set of all inputs belonging to the set

$$\mathcal{U}_k(\mathbf{z}, s) \;=\; \{(\mathbf{u}, \mathbf{v}) \in \mathbb{R}_+^N \times \mathbb{R}_+^N \mid \mathbf{1}^\top \mathbf{u} \leq s \;\text{ and }\; \mathbf{z} - \mathbf{u} - \mathbf{v} \in \mathcal{Z}_{k+1}\},$$

which ensures one-step state feasibility and that the total allocation of renewable supply does not exceed availability at the current stage. In characterizing the feasible set of causal scheduling policies, we restrict our attention to those policies with *Markovian information structure*, as opposed to allowing the control to depend on the entire history. This is without loss of optimality, as we will later show in Lemma B.1, which reveals, more strongly, that Markovian policies are capable of performing as well as the optimal oracle policy. We describe the scheduling decision at each stage k by the functions

$$\mathbf{u}_k = \boldsymbol{\mu}_k(\mathbf{z}, s) \quad \text{and} \quad \mathbf{v}_k = \boldsymbol{\nu}_k(\mathbf{z}, s),$$

where $\boldsymbol{\mu}_k : \mathcal{Z}_k \times \mathcal{S} \to \mathbb{R}_+^N$ and $\boldsymbol{\nu}_k : \mathcal{Z}_k \times \mathcal{S} \to \mathbb{R}_+^N$. A **feasible inter-class scheduling policy** is any finite sequence of scheduling decision functions $\sigma = (\boldsymbol{\mu}_0, \ldots, \boldsymbol{\mu}_{N-1}, \boldsymbol{\nu}_0, \ldots, \boldsymbol{\nu}_{N-1})$ such that $(\boldsymbol{\mu}_k, \boldsymbol{\nu}_k)(\mathbf{z}, s) \in \mathcal{U}_k(\mathbf{z}, s)$ for all $(\mathbf{z}, s) \in \mathcal{Z}_k \times \mathcal{S}$ and $k = 0, \ldots, N-1$. We denote by $\Sigma(\mathbf{x})$ the **feasible inter-class policy space**. Throughout the paper, we will suppress the explicit dependency of the feasible policy set on the aggregate demand bundle $\mathbf{x}$, when it's clear from the context.

4.3.1.2 Intra-class scheduling policies

Recall that an intra-class scheduling policy ϕ determines the allocation of available supply within each deadline-differentiated demand class, where the supply available to each demand class is determined by the inter-class policy $\sigma \in \Sigma$. We denote the **feasible intra-class policy space** by $\Phi(\sigma)$, which is parameterized by a given inter-class policy $\sigma \in \Sigma$. We denote by $\pi = (\sigma, \phi)$ the joint inter and intra-class scheduling policy employed by the supplier.

It is important to note that the supplier's profit depends *only* on the inter-class scheduling policy. Namely, for any feasible inter-class policy σ, all feasible intra-class scheduling policies $\phi \in \Phi(\sigma)$ yield the supplier the same expected profit. This follows from the supplier's indifference to supply allocation between consumers within a given demand class. Therefore, in characterizing the optimal

scheduling policy for the supplier, we restrict our attention to the characterization of optimal inter-class policies $\sigma \in \Sigma$ for the remainder of the paper. On the other hand, the intra-class scheduling policy employed by the supplier may have significant influence on consumer utility, for example, when the intermittent supply can fulfill only a fraction of the total demand within a demand class. As the formal definition of feasible intra-class policies is not directly relevant to the development of the main results to follow, their precise specification is deferred to Appendix B.1 to maintain continuity in exposition.

4.3.1.3 Supplier profit

We define the *expected profit* $J(\mathbf{x}, \mathbf{p}, \sigma)$ derived by a supplier as the revenue derived from an aggregate demand bundle $\mathbf{x}$ less the expected cost of servicing said demand bundle under a feasible inter-class scheduling policy $\sigma \in \Sigma(\mathbf{x})$. More precisely, let

$$J(\mathbf{x}, \mathbf{p}, \sigma) \;=\; \mathbf{p}^\top \mathbf{x} \;-\; Q(\mathbf{x}, \sigma), \tag{4.9}$$

where Q denotes the expected cost of firm generation incurred servicing $\mathbf{x}$ under a feasible policy $\sigma \in \Sigma$. It follows that

$$Q(\mathbf{x}, \sigma) \;=\; \mathbb{E} \sum_{k=0}^{N-1} \mathbf{c}^\top \mathbf{v}_k^\sigma, \tag{4.10}$$

where $\mathbf{c} = (c_0, \ldots, c_0)$. We write the state and control process as $\{\mathbf{z}_k^\sigma\}$, $\{\mathbf{u}_k^\sigma\}$, and $\{\mathbf{v}_k^\sigma\}$ to emphasize their dependence on the policy σ. We wish to characterize scheduling policies that lead to a minimal expected cost of firm supply. The inter-class scheduling policy $\sigma^* \in \Sigma(\mathbf{x})$ is **optimal** if

$$Q(\mathbf{x}, \sigma^*) \;\leq\; Q(\mathbf{x}, \sigma), \qquad \text{for all} \;\; \sigma \in \Sigma(\mathbf{x}). \tag{4.11}$$

We denote by $Q^*(\mathbf{x}) = Q(\mathbf{x}, \sigma^*)$ the minimum expected cost of firm supply. We proceed with a characterization of the optimal inter-class scheduling policy, σ^*, in the following Section.

4.3.1.4 Optimal scheduling policy

Theorem 4.6 (Earliest-Deadline-First)**.** *Given an aggregate demand bundle* $\mathbf{x} \in \mathbb{R}_+^N$*, the optimal scheduling policy* $\sigma^* \in \Sigma(\mathbf{x})$ *is given by:*

$$\mu_k^{j,*}(\mathbf{z}, s) = \min\left\{z^j, s - \sum\nolimits_{i=1}^{j-1} \mu_k^{i,*}(\mathbf{z}, s)\right\}, \qquad \nu_k^{j,*}(\mathbf{z}, s) = (z^j - \mu_k^{j,*}(\mathbf{z}, s)) \cdot \mathbf{1}_{\{k=j-1\}},$$

for $j = 1, \ldots, N$*,* $k = 0, \ldots, N-1$*, and* $(\mathbf{z}, s) \in \mathcal{Z}_k \times \mathcal{S}$*.*

Qualitatively, the optimal inter-class scheduling policy σ^* is such that the intermittent supply s_k available at each period k is allocated to those unsatisfied demand classes with *earliest-deadline-first* (EDF), while the firm supply is dispatched only as a *last resort* to ensure demand satisfaction. In addition, the optimal policy is *distribution-free* – which is attractive from an implementation perspective, as the optimal input sequence can be computed causally without requiring explicit knowledge of the underlying distribution on $\mathbf{s}$. For the remainder of the paper, we casually refer to σ^* as the EDF scheduling policy.

While a formal proof of Theorem 4.6 is deferred to Appendix B.2, the crux of our argument centers on showing that the earliest-deadline-first (EDF) policy performs (almost surely) as well as any *non-causal* policy with perfect foresight – the so called oracle optimal policy. Technically, the proof relies on showing that any oracle optimal schedule can be inductively mapped to the EDF schedule without causing an increase in the corresponding cost of firm supply, almost surely.

Remark 4.7 (The value of oracle information). While we have made no a priori assumptions regarding the conditional independence structure of the stochastic process $\mathbf{s}$, one can show that the optimal Markovian policy σ^* cannot be improved upon by any other causal policy with complete information history. More strongly, we establish in Lemma B.1 that the cost of firm supply *realized* under the causal policy σ^* *almost surely dominates* the cost realized under any non-casual policy. This ensures that the policy σ^* cannot be improved upon for any sample path realization of intermittent supply process $\mathbf{s}$. Essentially, this implies that there is *no additional value derived from having prescient information* regarding the realization of $\mathbf{s}$.

4.3.2 Marginal Cost Pricing

Given the EDF characterization of the optimal inter-class scheduling policy in Theorem 4.6, we are now in a position to characterize the supplier's optimal supply curve under price taking behavior. To aid in our analysis, we define a *residual random process* induced by the optimal scheduling policy

σ^* as

$$\xi_{k+1}(\mathbf{x}, \mathbf{s}) \;=\; \xi_k(\mathbf{x}, \mathbf{s})^+ \;+\; s_k - x_{k+1}, \tag{4.12}$$

for $k = 0, \ldots, N-1$, where $\xi_0 = 0$. A positive residual, $\xi_k(\mathbf{x}, \mathbf{s}) > 0$, denotes the amount of intermittent supply leftover after having serviced demand class k, according to the EDF inter-class scheduling policy σ^*, by its deadline k. A negative residual, $\xi_k(\mathbf{x}, \mathbf{s}) \leq 0$, denotes the amount by which the intermittent supply fell short or, equivalently, the quantity of firm supply required to ensure satisfaction of the demand class k. While $\xi_k(\mathbf{x}, \mathbf{s})$ depends on both the aggregate demand bundle $\mathbf{x}$ and intermittent supply process $\mathbf{s}$, we omit this dependency when it is clear from the context and compactly write $\boldsymbol{\xi} = (\xi_0, \ldots, \xi_N) \in \mathbb{R}^{N+1}$. Using this newly defined process, we have the following characterization of the minimum expected cost of firm generation under EDF scheduling. First, we require an assumption.

Assumption 4.3.1. The joint probability distribution of the intermittent supply process $\mathbf{s}$ is assumed to be absolutely continuous and have compact support. □

Lemma 4.8. *The expected recourse cost* $Q^*(\mathbf{x})$ (4.11) *derived under and aggregate demand bundle* $\mathbf{x} \in \mathbb{R}^N_+$ *and EDF scheduling policy* $\sigma^* \in \Sigma(\mathbf{x})$ *satisfies*

$$Q^*(\mathbf{x}) \;=\; -c_0 \cdot \sum_{k=1}^{N} \mathbb{E}\left\{\xi_k(\mathbf{x}, \mathbf{s})^-\right\}, \tag{4.13}$$

and is convex over $\mathbb{R}^N_+$ *and differentiable over* $(0, \infty)^N$ *in* $\mathbf{x}$.

Lemma 4.8, the proof of which can be found in Appendix B.3, admits a natural interpretation. Namely, the expected minimum cost of firm supply is equivalent to the amount by which the intermittent supply is expected to fall short for each demand class under EDF scheduling. Moreover, it follows readily that the expected profit $J(\mathbf{x}, \mathbf{p}, \sigma^*)$ derived under EDF scheduling is *differentiable and concave* in $\mathbf{x} \in \mathbb{R}^N_+$. As such, any allocation $\mathbf{x}$ satisfying the first order condition,

$$\nabla_{\mathbf{x}}\, J(\mathbf{x}, \mathbf{p}, \sigma^*)^\top (\mathbf{x} - \mathbf{y}) \;\geq\; 0 \qquad \text{for all} \quad \mathbf{y} \in \mathbb{R}^N_+, \tag{4.14}$$

is profit maximizing for the supplier given a particular price bundle $\mathbf{p} \in \mathcal{P}$. Accordingly, we provide an explicit characterization of the supplier's *marginal cost supply curve* in the following Theorem.

Theorem 4.9 (Marginal cost supply curve)**.** *An allocation* $\mathbf{x}$ *is profit maximizing for a given price bundle* $\mathbf{p}$*, if and only if* $\mathbf{p} = \boldsymbol{\zeta}(\mathbf{x})$*, where the mapping* $\boldsymbol{\zeta} : \mathbb{R}_+^N \to \mathcal{P}$ *satisfies*

$$\zeta_k(\mathbf{x}) \;=\; c_0 \cdot \mathbb{P}(\xi_k \leq 0) \;+\; c_0 \cdot \sum_{t=k+1}^{N} \mathbb{P}(\xi_k > 0, \ldots, \xi_{t-1} > 0, \xi_t \leq 0), \quad k = 1, \ldots, N. \tag{4.15}$$

Proof. Proof. Fix a price bundle $\mathbf{p} \in \mathbb{R}_+^N$. We have previously shown in Theorem 4.6 that the EDF inter-class scheduling policy $\sigma^* \in \Sigma(\mathbf{x})$ is optimal for any demand bundle $\mathbf{x} \in \mathbb{R}_+^N$. Hence, it suffices to show that $\mathbf{p} = \boldsymbol{\zeta}(\mathbf{x})$ satisfies the first order condition for optimality (4.14). Taking the gradient of the supplier's expected profit with respect to $\mathbf{x}$ yields $\nabla_{\mathbf{x}}\, J(\mathbf{x}, \mathbf{p}, \sigma^*) \;=\; \mathbf{p} \;-\; \nabla_{\mathbf{x}}\, Q^*(\mathbf{x})$. It remains to show that

$$\frac{\partial Q^*(\mathbf{x})}{\partial x_k} \;=\; \zeta_k(\mathbf{x})$$

for $k = 1, \ldots, N$. Working the with the simplified form of Q^* established in Lemma 4.8, it follows readily that

$$\frac{\partial Q^*(\mathbf{x})}{\partial x_k} \;= c_0 \cdot \sum_{\ell=k}^{N} \frac{\partial}{\partial x_k}\, \mathbb{E}\left\{\xi_\ell(\mathbf{x}, \mathbf{s})^-\right\}, \tag{4.16}$$

where we've truncated the summation from below at $\ell = k$, as $\xi_\ell(\mathbf{x}, \mathbf{s})$ is wholly independent of x_k for all $\ell < k$. We therefore restrict our attention to $\ell \geq k$ for the remainder of the proof. The next step of the proof relies on the ability to interchange the order of differentiation and expectation in (4.16). It is obvious from construction that $\xi_\ell(\mathbf{x}, \mathbf{s})^-$ is both a continuous function of $(\mathbf{x}, \mathbf{s})$ and piecewise affine in x_k (with a finite number of linear segments) for each $\mathbf{s}$. It follows that $\xi_\ell(\mathbf{x}, \mathbf{s})^-$ is differentiable almost everywhere in $x_k \in \mathbb{R}_+$ and satisfies

$$\left| \frac{\partial}{\partial x_k} \xi_\ell(\mathbf{x}, \mathbf{s})^- \right| \;\leq\; 1$$

almost everywhere. Then, for each $x_k \in \mathbb{R}_+$, we have that $\partial \xi_\ell(\mathbf{x}, \mathbf{s})^- / \partial x_k$ is integrable in x_k and by the dominated convergence theorem

$$\frac{\partial}{\partial x_k}\, \mathbb{E}\left\{\xi_\ell(\mathbf{x}, \mathbf{s})^-\right\} \;=\; \mathbb{E}\left\{ \frac{\partial}{\partial x_k} \xi_\ell(\mathbf{x}, \mathbf{s})^- \right\}.$$

Finally, it is not difficult to see that

$$\frac{\partial}{\partial x_k}\xi_\ell(\mathbf{x},\mathbf{s})^- \;=\; \begin{cases} \mathbf{1}_{\{\xi_k\le 0\}}, & \ell = k, \\ \mathbf{1}_{\{\xi_\ell\le 0\}}\cdot\prod_{t=k}^{\ell-1}\mathbf{1}_{\{\xi_t>0\}}, & \ell > k. \end{cases}$$

And taking expectation, we have the desired result. □

Remark 4.10 (Price monotonicity). In Eq. (4.15), we indeed characterize a **marginal cost pricing scheme $\boldsymbol{\zeta}$**, which maps every aggregate demand bundle in $\mathbb{R}_+^N$ into a deadline differentiated price bundle in $\mathcal{P}$. Moreover, one can readily interpret the pricing scheme as equating the price p_k for energy with deadline k (up to a proportionality constant c_0) with the probability that firm supply will be required to service the bundle $\mathbf{x}$ at any subsequent time period $t \ge k-1$ under the optimal inter-class scheduling policy σ^*. Naturally, the larger the probability of shortfall, the larger the price. Moreover, it is readily verified that the supply curve $\mathbf{p} = \boldsymbol{\zeta}(\mathbf{x})$ yields nonincreasing prices for any demand bundle. More precisely,

$$c_0 \;\ge\; \zeta_1(\mathbf{x}) \;\ge\; \zeta_2(\mathbf{x}) \;\ge\cdots\ge\; \zeta_N(\mathbf{x})$$

for all $\mathbf{x}\in\mathbb{R}_+^N$. This property of price monotonicity is consistent with our initial criterion for constructing such a market system. Namely, the longer a customer is willing to defer her consumption in time, the less she is required to pay per unit of energy.

4.4 Incentive Compatibility and Market Equilibrium

We now explore the game theoretic and welfare implications of the mechanism defined in terms of the scheduling policy and pricing scheme developed in Section 4.3. In particular, we show in Section 4.4.1 that a mechanism consisting of earliest-deadline-first (EDF) scheduling in combination with a pricing scheme reflecting the supplier's marginal cost (4.15) is indeed incentive compatible. In addition, we discuss in Section 4.4.2 the implementation of the proposed deadline differentiated pricing scheme. And, under an additional mild assumption on each consumer's marginal valuation on energy, we show that the marginal cost pricing scheme defined in (4.15) achieves social optimality at a competitive market equilibrium.

4.4.1 Incentive Compatibility

Assumption 4.4.1. We assume that both the inter and intra-class scheduling policies $\pi = (\sigma, \phi)$ and the pricing scheme $\boldsymbol{\kappa}$ employed by the supplier are commonly known by the consumers. □

Remark 4.11. Note that consumers need not to know the joint probability distribution of the intermittent supply process $\mathbf{s}$ and the aggregate demand bundle $\mathbf{x}$. Instead, each consumer is assumed to know *how* the distribution of the intermittent supply process $\mathbf{s}$ and the aggregate demand bundle $\mathbf{x}$ determine both the price bundle and the (random) energy delivery.

Definition 4.12 (Incentive compatibility (IC))**.** A scheduling policy $\pi = (\sigma, \phi)$ and pricing scheme $\boldsymbol{\kappa}$ are deemed **incentive compatible** (IC) for consumers of type θ, if it is a dominant-strategy for consumers of this type to be truth-telling. Formally, a mechanism $(\pi, \boldsymbol{\kappa})$ is incentive compatible for every consumer $i \in [0, 1]$ of type-θ, if

$$V_i^{\pi}(\theta, \varphi_i^*, \mathbf{x}, \boldsymbol{\kappa}) \;\geq\; V_i^{\pi}(\theta, \varphi_i, \mathbf{x}, \boldsymbol{\kappa}), \qquad \forall\, \varphi_i, \quad \forall\, \mathbf{x} \in \mathbb{R}_+^N, \tag{4.17}$$

where φ_i^* is the truth-telling strategy defined in Definition 4.2. ■

Remark 4.13. Unlike the Nash equilibrium characterization of aggregate consumer behavior common to the related literature [10, 24, 26, 45], we require a dominant-strategy implementation in Definition 4.12.

Theorem 4.14 (Incentive compatibility)**.** *Suppose that Assumptions 4.2.2 and 4.4.1 hold. Consider the mechanism $(\pi^*, \boldsymbol{\zeta})$ defined by the marginal cost pricing scheme $\boldsymbol{\zeta}$ characterized in Eq. (4.15), in combination with the scheduling policy $\pi^* = (\sigma^*, \phi)$ consisting of the EDF inter-class scheduling policy σ^* and an arbitrary feasible intra-class policy $\phi \in \Phi(\sigma^*)$. It follows that the mechanism $(\pi^*, \boldsymbol{\zeta})$ is incentive compatible for all consumers of type $\theta = (k, R, q)$ satisfying $R \geq c_0$.*

We defer the proof of Theorem 4.14 to Appendix B.4 for the ease of exposition. Since the optimal price schedule is nonincreasing in deadline, and demand is guaranteed to be met before the requested deadline, every consumer i has no incentive to request a positive amount of electricity before her true deadline k_i. In the proof of Theorem 4.14 we show that a consumer cannot strictly benefit from requesting positive amount of energy after her true deadline k_i, regardless of the actions taken by other consumers. Under the EDF scheduling policy, there is positive probability that the energy requested by some period $k' > k_i$ cannot be delivered by the consumer's true deadline k_i. Under marginal cost pricing, the price difference between the two deadlines, k_i and k', is shown to be not large enough to compensate the expected utility loss resulting from the possible shortfall in energy consumption before deadline k_i. The incentive compatibility result then follows.

Note that the incentive compatibility result established in Theorem 4.14 is strong. The pricing scheme in Eq. (4.15) in combination with EDF scheduling implements truth-telling in dominant strategies, regardless of the type distribution over consumer types and probability distribution on the intermittent supply process $\mathbf{s}$. Moreover, as the pricing scheme $\boldsymbol{\zeta}$ represents the supplier's marginal cost, it can be shown to achieve social optimality, if the condition $R \geq c_0$ holds for every type $\theta \in \Theta$ (cf. Corollary 4.17).

Remark 4.15 (Caveats). We note that incentive compatibility may fail to hold if certain assumptions regarding a consumer's utility function are violated. First, one can readily show that incentive compatibility may fail to hold for a consumer of type $\theta = (k, R, q)$, if the marginal cost of firm supply exceeds her marginal valuation of energy, i.e. $R < c_0$. Second, we note that the result in Theorem 4.14 fails to hold for arbitrary concave utility functions. This is intuitive, as a consumer with a highly concave utility function may prefer to report a false deadline that is later than her true deadline k, if she can obtain a fraction of her demand before stage k (but a utility close to the maximum, because of the large underlying concavity of her utility function) with high probability, at a much lower price.

4.4.2 Implementation and Market Equilibrium

We briefly discuss the implementation of the proposed market for deadline differentiated energy services through both a mechanism design and market equilibrium based approach. First, we make an assumption under which we will operate for the remainder of the paper.

Assumption 4.4.2. We assume that for every $(k, R, q) \in \Theta$, we have $R \geq c_0$. □

We note that Assumption 4.4.2 is quite reasonable for electricity consumers, as their marginal valuation on electricity consumption is commonly higher than the nominal price of electricity, which is represented in our model by c_0. As a result, electricity demand is usually considered to be inelastic [42, 46], especially in the short term [48].

4.4.2.1 A mechanism design approach

One possibility entails its implementation through a classical mechanism-design approach:

1. The supplier announces the mechanism $(\pi^*, \boldsymbol{\zeta})$ (specified in Theorem 4.14) to the population of customers.

2. With common knowledge of the mechanism $(\pi^*, \boldsymbol{\zeta})$, every consumer i reports her type θ_i to the supplier, as it is a dominant strategy to do so under the proposed mechanism.

3. The price bundle is set to $\mathbf{p} = \boldsymbol{\zeta}(\mathbf{x}^*)$, where $\mathbf{x}^* = \mathbf{d}(\boldsymbol{\theta}, \boldsymbol{\varphi}^*)$ is the truth-telling aggregate demand bundle.

4. Each consumer i receives an expected payoff of

$$\mathbb{E}\left\{U_{\theta_i}\left(\omega_{k,i}^{\pi^*}(\mathbf{x}^*, \mathbf{a}^*)\right)\right\} - \mathbf{p}^\top \mathbf{a}^*,$$

   where $\mathbf{a}^* = \varphi_i^*(\theta_i)$ is the truth-telling action of consumer i defined in Eq. (4.2).

It is straightforward to see that the resulting quantity-price pair $(\mathbf{x}^*, \mathbf{p})$ is social welfare maximizing (cf. the proof of Corollary 4.17).

4.4.2.2 A market equilibrium approach

Another approach entails implementation of the deadline differentiated pricing scheme at a *competitive market equilibrium*; namely, the intersection of the "supply" and "demand" curves. Theorem 4.9 indeed characterizes a supply curve. In order to define the demand curve (that specifies the aggregate demand under a given price bundle $\mathbf{p}$), we take the approach considered in the literature on priority pricing for queuing systems [24, 26, 34, 45], where consumers react to a fixed price bundle $\mathbf{p}$ (as opposed to a pricing scheme $\boldsymbol{\kappa}(\cdot)$) announced by the supplier. Given a particular scheduling policy π (announced by the supplier), let a consumer's strategy be a mapping from her type and the observed price bundle to her action; then, a price bundle $\mathbf{p}$ and a scheduling policy π naturally define a game among the consumer population (under the fixed scheduling policy π).

We now define the **truth-telling strategy** of a consumer. Given a price bundle $\mathbf{p} \in \mathcal{P}$ and the consumer's type $\theta = (k, R, q)$, a truth-telling consumer would take an action of the following form:

$$a_j^* = \begin{cases} q \cdot \mathbf{1}_{\{R \geq p_k\}}, & j = k \\ 0, & j \neq k. \end{cases} \tag{4.18}$$

In other words, a truth-telling consumer requests her surplus maximizing quantity $q \cdot \mathbf{1}_{\{R \geq p_k\}}$ to be delivered by her true deadline k. The truth-telling strategy profile is a Nash equilibrium, if the strategy defined in (4.18) maximizes each individual consumer's payoff, provided that the other consumers stick to the same strategy, and that the announced scheduling policy π is employed by the supplier.

Consider now the marginal cost pricing scheme $\boldsymbol{\zeta}$. It follows from Assumption 4.4.2 that, under any price bundle defined by $\boldsymbol{\zeta}$, the aggregate demand bundle of a truth-telling consumer population is inelastic and is given by

$$x_j^* = \sum_{\theta=(k,R,q)\in\Theta} q \cdot \mathbf{1}_{\{j=k\}}\rho(\theta), \qquad j = 1, \ldots, N. \tag{4.19}$$

This follows from the property that $\zeta_k(\mathbf{x}) \leq c_0 \leq R$ for all $(k, R, q) \in \Theta$ and $\mathbf{x} \in \mathbb{R}_+^N$. Namely, the price associated with every deadline k is always no greater than each consumer's marginal utility.

With the information on the aggregate demand bundle $\mathbf{x}^*$ (defined in (4.19)), a system operator (or some other regulatory authority) can set the price bundle to be $\mathbf{p} = \boldsymbol{\zeta}(\mathbf{x}^*)$. It is then optimal for a (price-taking) supplier to announce a scheduling policy π^* (cf. its definition in Theorem 4.14) and serve the aggregate demand $\mathbf{x}^*$. On the demand side, Corollary 4.17 shows that the incentive compatibility result established in Theorem 4.14 implies that the truth-telling strategy profile is a Nash equilibrium of the game defined by π^* and $\boldsymbol{\zeta}(\mathbf{x}^*)$. With the information on the aggregate demand bundle $\mathbf{x}^*$, a system operator can achieve social optimality at a market equilibrium defined as follows (cf. Corollary 4.17).

Definition 4.16 (Market equilibrium)**.** Given the types of all consumers $\boldsymbol{\theta}$, a quantity-price pair $(\mathbf{x}, \mathbf{p})$ constitutes a market equilibrium, if there exists a feasible scheduling policy $\pi = (\sigma, \phi)$ under which the following two conditions holds.

1. Under the price bundle $\mathbf{p}$, there exists a Nash equilibrium at which every consumer is truth-telling, and the resulting aggregate demand bundle is $\mathbf{x}$.
2. Given the price bundle $\mathbf{p}$, the aggregate demand bundle $\mathbf{x}$ and scheduling policy $\pi = (\sigma, \phi)$ maximize the supplier's expected profit (cf. its expression in Eq. (4.9)). ■

In Definition 4.16, the first condition ensures that under the price bundle $\mathbf{p}$, the aggregate demand (resulting from a truth-telling population) is $\mathbf{x}$, and the second condition requires that given the price bundle $\mathbf{p}$, a price-taking supplier would like to employ the scheduling policy $\pi = (\sigma, \phi)$ and supply the bundle $\mathbf{x}$ (i.e., supply equals demand). We will show in the following corollary that the EDF scheduling policy $\pi^* = (\sigma^*, \phi)$ together with the pricing scheme $\boldsymbol{\zeta}$ maximize the social welfare (the sum of aggregate consumer surplus and supplier profit) at a market equilibrium.

Corollary 4.17. *Suppose that Assumptions 4.2.2-4.4.2 hold. For a given type distribution ρ, a mechanism $(\sigma^*, \phi, \boldsymbol{\zeta})$ maximizes the social welfare at the unique market equilibrium $(\mathbf{x}^*, \boldsymbol{\zeta}(\mathbf{x}^*))$.*

Here, $\phi \in \Phi(\sigma^*)$ is an arbitrary feasible intra-class scheduling policy, and the aggregate demand bundle $\mathbf{x}^*$ is defined in Eq. (4.19).

Proof. Proof. Under Assumption 4.4.2, the condition in Eq. (4.17) holds for every $\theta \in \Theta$, and therefore it is optimal for every consumer i to be truth-telling, provided that the price bundle is set to be $\boldsymbol{\zeta}(\mathbf{x}^*)$, and that the other consumers are truth-telling and therefore request the aggregate demand bundle $\mathbf{x}^*$. It follows that for the game defined by the scheduling policy (σ^*, ϕ) and the price bundle $\boldsymbol{\zeta}(\mathbf{x}^*)$, there exists a Nash equilibrium at which every consumer is truth-telling.

Given the price bundle $\boldsymbol{\zeta}(\mathbf{x}^*)$, the aggregate demand bundle $\mathbf{x}^*$ together with the EDF scheduling policy maximizes the supplier's expected profit (cf. Theorem 4.9). Hence, the pair, $(\mathbf{x}^*, \boldsymbol{\zeta}(\mathbf{x}^*))$, constitutes a market equilibrium. Since the supplier marginal cost never exceeds c_0, the aggregate demand (at any price bundle that represents supplier marginal cost) is inelastic, and is expressed in Eq. (4.19). We therefore conclude the uniqueness of a market equilibrium.

It is straightforward to see that social welfare is maximized at this market equilibrium, because under Assumption 4.4.2, it is optimal to fully serve the aggregate demand expressed in (4.19), and further, the EDF scheduling policy σ^* minimizes the expected cost of servicing $\mathbf{x}^*$. ■ □

With the information on consumers' aggregate demand, the supplier can simply announce the price bundle $\boldsymbol{\zeta}(\mathbf{x})$, and maximize the social welfare at the competitive equilibrium $(\mathbf{x}, \boldsymbol{\zeta}(\mathbf{x}))$. Rigorous treatment of the (dynamic) convergence to a competitive equilibrium dates back to the early 20th century in the economics literature [3, 4, 20]. For example, the instantaneous adjustment process studied in [3, 4], where the derivative of price adjustment is set equal to the aggregate excess (net) demand under the current price, is guaranteed to converge to the efficient competitive equilibrium characterized in Corollary 4.17.

4.5 Conclusion

To explore the flexibility of delay-tolerant electricity loads, we propose a novel deadline-differentiated market for energy that offers discounted (per-unit) electricity prices to consumers in exchange for their consent to defer their electric power consumption, and provides a guarantee on the aggregate quantity of energy to be delivered by a consumer-specified deadline. We provide a full characterization of the joint scheduling and pricing scheme that yields an efficient (competitive) market equilibrium between a (price-taking) supplier and a consumer population. Somewhat surprisingly,

we show that this efficient mechanism is incentive compatible in that every consumer would like to reveal her true deadline to the supplier, regardless of the actions taken by other consumers.

The market we have considered in our analysis is single shot. As a natural extension, it would be of interest to explore the dynamic analog of our formulation in which the market is cleared on a recurrent basis. In addition, such a dynamic setting could provide the foundation on which to explore efficient price discovery schemes.

Acknowledgment

This work was supported in part by NSF CNS-1239178, PSERC under sub-award S-52, and US DoE under the CERTS initiative. The authors would like to thank Costas Courcoubetis, Paul De Martini, Steven Low, Kameshwar Poolla, Pravin Varaiya, and Adam Wierman for their helpful discussions.

Appendix A

Proofs for Chapter 2

A.1 Proof of Theorem 2.12

Let u be any element of $\mathbb{U}$. Then $u = \sum_k e^k$, where $e^k \in \mathbb{E}^k$. The proof requires showing that $u \in \mathbb{B}(\phi_2)$. Define $\vartheta^k(t) = (\theta^k(t) - \theta_r^k)/b^k$. Since $\theta^k(0) = \theta_r^k$, we have from (2.3),

$$\dot{\vartheta}^k(t) = -a^k \vartheta^k(t) - e^k(t), \quad \vartheta^k(0) = 0.$$

Taking a Laplace transform of the above equation, we obtain

$$\vartheta^k(s) = -\frac{1}{s + a^k} e^k(s).$$

Let $\dot{x}(t) = -\alpha x(t) - u(t)$ and $x(0) = 0$. Taking the Laplace transform,

$$\begin{aligned} x(s) &= -\frac{1}{s+\alpha} u(s) = -\frac{1}{s+\alpha} \sum_k e^k(s) \\ &= -\sum_k \frac{1}{s+\alpha} e^k(s) = -\sum_k \frac{s+a^k}{s+\alpha} \frac{1}{s+a^k} e^k(s) \\ &= \sum_k \frac{s+a^k}{s+\alpha} \vartheta^k(s) = \sum_k (1 + \frac{a^k - \alpha}{s+\alpha}) \vartheta^k(s). \end{aligned}$$

Given $Y(s) = H(s)U(s)$, we have $\|y(t)\|_\infty \leq \|h(t)\|_1 \|u(t)\|_\infty$, where $y(t), h(t)$ and $u(t)$ are the inverse Laplace transforms of $Y(s), H(s)$ and $U(s)$ respectively. Therefore,

$$\|x(t)\|_\infty \leq \sum_k (1 + |1 - \frac{a^k}{\alpha}|) \|\vartheta^k(t)\|_\infty.$$

Because $e^k(t) \in \mathbb{E}^k$, $\|\vartheta^k(t)\|_\infty \leq \Delta^k/b^k$, so that

$$\|x(t)\|_\infty \leq \sum_k (1 + |1 - \frac{a^k}{\alpha}|)\Delta^k/b^k. \tag{A.1}$$

Additionally, $-P_o^k \leq e^k(t) \leq P_m^k - P_o^k$, implies that

$$-\sum_k P_o^k \leq u(t) \leq \sum_k P_m^k - P_o^k. \tag{A.2}$$

Inequalities (A.1) and (A.2) verify that $u(t) \in \mathbb{B}(\phi_2)$.

A.2 Proof of Theorem 2.13

In this case, it is sufficient to show that $u(t) \in \mathbb{B}(\phi_1)$ implies $e^k(t) = \beta^k u(t) \in \mathbb{E}^k$. Let $\dot{x} = -\alpha x - u, x(0) = 0$. Now, if $u(t) \in \mathbb{B}(\phi_1)$, the following must hold

$$\begin{aligned} -n_- \leq u(t) \leq n_+, \qquad & \forall\ t > 0, \\ |x(t)| \leq C, \qquad & \forall\ t > 0. \end{aligned}$$

Apply power deviation $e^k(t) = \beta^k u(t)$ to each TCL unit. Define $\vartheta^k(t) = (\theta^k(t) - \theta_r^k)/b^k$ with $\theta^k(0) = \theta_r^k$. From (2.3),

$$\dot{\vartheta}^k(t) = -a^k \vartheta^k(t) - e^k(t), \quad \vartheta^k(0) = 0.$$

Taking the Laplace transform of the above equation,

$$\begin{aligned} \vartheta^k(s) &= -\frac{1}{s + a^k} e^k(s) = -\frac{1}{s + a^k} \beta^k u(s) \\ &= -\beta^k \frac{s + \alpha}{s + a^k} \frac{u(s)}{s + \alpha} = \beta^k (1 + \frac{\alpha - a^k}{s + a^k}) x(s), \end{aligned}$$

where we have used the fact that $x(s) = -\frac{1}{s+\alpha} u(s)$. Using the same bounds as in the proof of Theorem 2.12,

$$\begin{aligned} \|\vartheta^k(t)\|_\infty &\leq \beta^k (1 + |\frac{\alpha - a^k}{a^k}|) \|x(t)\|_\infty \\ &\leq \beta^k (1 + |\frac{\alpha - a^k}{a^k}|) C, \quad \forall\ t > 0. \end{aligned}$$

Since C is chosen so that $\beta^k C \leq \frac{\Delta^k}{b^k\left(1+\left|\frac{\alpha-a^k}{a^k}\right|\right)}$, we have $|\vartheta^k(t)| \leq \Delta^k/b^k$ or equivalently $|\theta^k(t)-\theta_r^k| \leq \Delta^k$. Moreover $\beta^k n_- \leq P_o^k$ and $\beta^k n_+ \leq P_m^k - P_o^k, \forall\, k$, yields $-P_o^k \leq \beta^k u(t) \leq P_m^k - P_o^k$ which shows that $e^k(t) = \beta^k u(t) \in \mathbb{E}^k$.

Appendix B

Proofs for Chapter 4

B.1 Intra-class Scheduling Policies

Formally, for a consumer i who purchases a bundle $\mathbf{a} = (a_1, \ldots, a_N)$, we let $\boldsymbol{\lambda}_{k,i}(\mathbf{x}, \mathbf{a}) \in \mathbb{R}_+^N$ denote (element-wise in j) the amount of energy delivered (at period k) to consumer i so as to satisfy her demand a_j. We denote the intra-class scheduling policy by

$$\phi = \{(\boldsymbol{\lambda}_{0,i}, \ldots, \boldsymbol{\lambda}_{N-1,i}) \mid i \in [0,1]\},$$

where $\boldsymbol{\lambda}_{k,i} : \mathbb{R}_+^N \times \mathcal{A} \to \mathbb{R}_+^N$ for all $i \in [0,1]$ and $k = 0, \ldots, N-1$. Given an inter-class scheduling policy $\sigma \in \Sigma$, an intra-class scheduling policy ϕ is feasible if and only if it satisfies the following constraints.

1. The intra-class scheduling policy should not deliver any supply that is allocated to class j to consumers outside this class. That is, for each demand class $j = 1, \ldots, N$, we have that $\lambda_{k,i}^j(\mathbf{x}, \mathbf{a}) = 0$ for every consumer i such that $a_j = 0$.

2. At every period k, no energy is delivered to demand class j with $j \leq k$ by the feasibility of the inter-class policy $\sigma \in \Sigma$. That is, $\lambda_{k,i}^j(\mathbf{x}, \mathbf{a}) = 0$ for every $i \in [0,1]$, and every $1 \leq j \leq k \leq N-1$.

3. The total supply allocated to demand class j at time period k must be fully utilized, i.e.

$$\int_{[0,1]} \lambda_{k,i}^j(\mathbf{x}, \mathbf{a}) \; \eta(di) \;=\; \mu_k^j + \nu_k^j, \qquad \forall \;\; 0 \leq k < j \leq N-1,$$

where we use the Lebesgue integral[1] (with respect to Lebesgue measure η defined on $[0,1]$), and μ_k^j and ν_k^j denote the amount of intermittent and firm supply allocated to demand class j at period k, according to the inter-class policy σ.

4. Each consumer's individual delivery commitments must be met:

$$\sum\nolimits_{t=1}^{k} a_t \;\leq\; \omega_{k,i}^{\pi}(\mathbf{x},\mathbf{a}) \;\leq\; \sum\nolimits_{t=1}^{N} a_t, \qquad k=1,\ldots,N,$$

where the total energy delivered to consumer i by deadline k is given by

$$\omega_{k,i}^{\pi}(\mathbf{x},\mathbf{a}) = \sum_{t=0}^{k-1}\sum_{j=1}^{N} \lambda_{t,i}^{j}(\mathbf{x},\mathbf{a}). \tag{B.1}$$

Notice that for any feasible inter-class scheduling policy $\sigma \in \Sigma$, it is always possible to ensure the satisfaction of the above constraint.

We denote the **feasible intra-class policy space** by $\Phi(\sigma)$, which is parameterized by a given inter-class policy $\sigma \in \Sigma$.

B.2 Proof of Theorem 4.6

Consider the causal scheduling policy σ^* defined in Theorem 4.6. It is immediate to see that said policy is *feasible* for all $\mathbf{x} \in \mathbb{R}_+^N$ by construction. We prove *optimality* of the policy $\sigma^* \in \Sigma(\mathbf{x})$ by establishing a stronger result. Namely, the causal scheduling policy π^* *almost surely dominates* any non-casual policy in the metric of realized cost of firm supply. In particular, we show in Lemma B.1 that σ^* almost surely achieves the same cost of firm supply incurred under the *optimal oracle policy*, which we now define.

Consider the **oracle problem** (B.2), in which supplier's scheduling actions are chosen with the benefit of perfect foresight. Namely, the control action at each stage is allowed to depend (non-causally) on the entire (realized) sequence of intermittent supply $\mathbf{s} = (s_0, s_1, \ldots, s_{N-1})$. With a slight abuse of notation, we denote by $Q(\mathbf{x},\mathbf{s},\mathbf{u},\mathbf{v})$ the *cost incurred by any sequence of feasible scheduling actions* $\mathbf{u} = (\mathbf{u}_0,\ldots,\mathbf{u}_{N-1})$ and $\mathbf{v} = (\mathbf{v}_0,\ldots,\mathbf{v}_{N-1})$ under any pair $(\mathbf{x},\mathbf{s})$. Computing a sequence of non-causal scheduling actions that minimize the cost of firm supply amounts to solving

[1] Note that we have implicitly required here that $\lambda_{k,i}^{j}(\mathbf{x},\mathbf{a})$ is Lebesgue integrable with respect to i, over the interval $[0,1]$.

a deterministic linear program whose optimal value satisfies

$$Q(\mathbf{x}, \mathbf{s}, \mathbf{u}^{\mathrm{o}}, \mathbf{v}^{\mathrm{o}}) \;=\; \min_{\mathbf{u},\mathbf{v}} \sum_{k=0}^{N-1} \mathbf{c}^{\top}\mathbf{v}_k \quad \text{subject to} \quad \begin{aligned} &\mathbf{z}_{k+1} = \mathbf{z}_k - \mathbf{u}_k - \mathbf{v}_k \\ &(\mathbf{u}_k, \mathbf{v}_k) \in \mathcal{U}_k(\mathbf{z}_k, s_k) \end{aligned} \tag{B.2}$$

for $k = 0, \ldots, N-1$, where $\mathbf{z}_0 = \mathbf{x}$. The *optimal oracle actions* (possibly non-unique) are denoted by $\mathbf{u}^{\mathrm{o}} = (\mathbf{u}_0^{\mathrm{o}}, \ldots, \mathbf{u}_{N-1}^{\mathrm{o}})$ and $\mathbf{v}^{\mathrm{o}} = (\mathbf{v}_0^{\mathrm{o}}, \ldots, \mathbf{v}_{N-1}^{\mathrm{o}})$, which are implicitly parameterized by the pair $(\mathbf{x}, \mathbf{s})$. It's clear to see that the oracle problem (B.2) is feasible for any pair $(\mathbf{x}, \mathbf{s}) \in \mathbb{R}_+^N \times \mathcal{S}^N$. We can now state the following Lemma.

Lemma B.1. *Given any aggregate demand bundle $\mathbf{x} \in \mathbb{R}_+^N$, the causal scheduling policy $\sigma^* \in \Pi(\mathbf{x})$ defined in Theorem 4.6 almost surely achieves the cost of firm supply incurred under the optimal oracle policy. Namely,*

$$Q(\mathbf{x}, \mathbf{s}, \mathbf{u}^*, \mathbf{v}^*) \;=\; Q(\mathbf{x}, \mathbf{s}, \mathbf{u}^{\mathrm{o}}, \mathbf{v}^{\mathrm{o}}) \quad \textit{almost surely,}$$

where $\mathbf{u}_k^* = \boldsymbol{\mu}_k^*(\mathbf{z}_k^{\sigma^*}, s_k)$ *and* $\mathbf{v}_k^* = \boldsymbol{\nu}_k^*(\mathbf{z}_k^{\sigma^*}, s_k)$ *for all* $k = 0, \ldots, N-1$.

Proof. Proof. Let $(\mathbf{x}, \mathbf{s}) \in \mathbb{R}_+^N \times \mathcal{S}^N$. We establish the desired result by showing that any oracle optimal schedule $(\mathbf{u}^{\mathrm{o}}, \mathbf{v}^{\mathrm{o}})$ can be inductively mapped to the EDF schedule $(\mathbf{u}^*, \mathbf{v}^*)$ without causing an increase in the corresponding cost of firm of supply. More precisely, given any oracle optimal schedule whose *first* $k-1$ *inputs* coincide with the corresponding EDF schedule, one can define a mapping that transforms the oracle schedule into another feasible schedule $(\widetilde{\mathbf{u}}, \widetilde{\mathbf{v}})$ whose *first* k *inputs* coincide with EDF scheduling – without causing an increase in the realized cost. Namely,

$$Q(\mathbf{x}, \mathbf{s}, \widetilde{\mathbf{u}}, \widetilde{\mathbf{v}}) \;=\; Q(\mathbf{x}, \mathbf{s}, \mathbf{u}^{\mathrm{o}}, \mathbf{v}^{\mathrm{o}}).$$

(Base Step). We begin by proving the base step. This will, in addition, reveal the crux of the argument required to establish the inductive step. Consider an arbitrary oracle optimal schedule $(\mathbf{u}^{\mathrm{o}}, \mathbf{v}^{\mathrm{o}})$ satisfying (B.2). For clarity in exposition, we restrict our analysis to those oracle policies satisfying $v_k^{j,\mathrm{o}} = 0$ for all $k < j-1$. In other words, we consider only those policies in which the firm supply is deployed only as a last resort. This is clearly without loss of optimality, because of linearity and time invariance of the cost criterion.

Let $(\mathbf{u}^*, \mathbf{v}^*)$ denote the schedule realized under the causal EDF scheduling policy σ^* defined in Theorem 4.6. And let $\boldsymbol{\delta} = \mathbf{u}_0^* - \mathbf{u}_0^{\mathrm{o}}$ denote the discrepancy between the EDF and oracle schedule

at period $k = 0$. Notice that $\mathbf{1}^\top \boldsymbol{\delta} \geq 0$, as EDF scheduling either fully utilizes the available intermittent supply s_0 or completes all tasks at stage $k = 0$.

The Mapping. We now proceed in constructing a mapping that transforms the oracle optimal schedule $(\mathbf{u}^{\mathrm{o}}, \mathbf{v}^{\mathrm{o}})$ into a modified feasible schedule $(\widetilde{\mathbf{u}}, \widetilde{\mathbf{v}})$, whose inputs at stage $k = 0$ coincide with that of EDF scheduling (i.e. $\widetilde{\mathbf{u}}_0 = \mathbf{u}_0^*$ and $\widetilde{\mathbf{v}}_0 = \mathbf{v}_0^*$). And all without causing an increase in the realized cost of firm supply. We construct such a mapping by characterizing a collection of feasible perturbations on the oracle optimal schedule at later stages ($k \geq 1$) to compensate its discrepancy with EDF at the initial stage $k = 0$. More precisely, we consider perturbations of the form

$$\widetilde{\mathbf{u}}_k \;=\; \mathbf{u}_k^{\mathrm{o}} \;+\; \boldsymbol{\beta}_k \quad \text{and} \quad \widetilde{\mathbf{v}}_k \;=\; \mathbf{v}_k^{\mathrm{o}} \;+\; \boldsymbol{\alpha}_k,$$

where $\boldsymbol{\alpha}_k,\ \boldsymbol{\beta}_k \in \mathbb{R}^N$ for $k = 0, \ldots, N-1$. And again, because of objective linearity and time invariance, we assume that perturbations on the firm supply allocation satisfy $\alpha_k^j = 0$ for all $k < j - 1$. In order to enforce consistency with EDF at stage $k = 0$, we necessarily require that the initial perturbations satisfy $\boldsymbol{\beta}_0 = \boldsymbol{\delta}$ and $\alpha_0^1 = -\delta^1$.

Feasible Mappings. Naturally, the perturbations $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ on the oracle optimal schedule must respect certain constraints in order that the modified schedule $(\widetilde{\mathbf{u}}, \widetilde{\mathbf{v}})$ remain feasible. In particular, *deadline constraints on demand satisfaction* require that the perturbations satisfy

$$\alpha_{j-1}^j \;+\; \sum_{k=1}^{j-1} \beta_k^j \;=\; -\delta^j \quad \text{and} \quad \alpha_k^j = \beta_k^j = 0 \quad (k \geq j) \tag{B.3}$$

for all tasks $j = 1, \ldots, N$. In other words, the initial perturbation on the oracle schedule must be compensated by equal and opposite perturbations over subsequent stages. In addition, *limits on the availability of intermittent supply* at each stage are enforced by requiring that $\mathbf{1}^\top \boldsymbol{\beta}_k \;\leq\; s_k - \mathbf{1}^\top \mathbf{u}_k^{\mathrm{o}}$ for all $k = 0, \ldots, N-1$, while *non-negativity* requires that $\boldsymbol{\alpha}_k \geq -\mathbf{v}_k^{\mathrm{o}}$ and $\boldsymbol{\beta}_k \geq -\mathbf{u}_k^{\mathrm{o}}$ for all $k = 0, \ldots, N-1$. It is not difficult to see that the set of feasible perturbations is nonempty for any pair $(\mathbf{x}, \mathbf{s})$.

Mapping the Oracle to EDF. We now establish the desired result by showing the existence of feasible perturbations $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ on the oracle schedule that yield $Q(\mathbf{x}, \mathbf{s}, \widetilde{\mathbf{u}}, \widetilde{\mathbf{v}}) \;\leq\; Q(\mathbf{x}, \mathbf{s}, \mathbf{u}^{\mathrm{o}}, \mathbf{v}^{\mathrm{o}})$. Consider the realized cost of firm supply incurred under any feasible perturbed schedule $(\widetilde{\mathbf{u}}, \widetilde{\mathbf{v}})$. It is straightforward to see that

$$Q(\mathbf{x}, \mathbf{s}, \widetilde{\mathbf{u}}, \widetilde{\mathbf{v}}) \;=\; Q(\mathbf{x}, \mathbf{s}, \mathbf{u}^{\mathrm{o}}, \mathbf{v}^{\mathrm{o}}) \;+\; c_0 \cdot \sum_{j=1}^{N} \alpha_{j-1}^j.$$

Hence, to prove the desired result, it remains to show that $\sum_{j=1}^{N} \alpha_{j-1}^{j} \leq 0$, which implies a nonincrease in the cost of firm supply incurred by the perturbed schedule $(\widetilde{\mathbf{u}}, \widetilde{\mathbf{v}})$. We first define two disjoint sets, Δ^+ and Δ^-, containing those indices where the EDF schedule at stage $k = 0$ either strictly exceeds or falls short of the oracle schedule in terms of its allocation of intermittent supply to each task, respectively. More precisely,

$$\Delta^+ = \{1 \leq j \leq N \mid \delta^j > 0\} \quad \text{and} \quad \Delta^- = \{1 \leq j \leq N \mid \delta^j < 0\}.$$

These sets possess the important property that $\ell < j$ for all $\ell \in \Delta^+$ and $j \in \Delta^-$. Namely, the set of all tasks belonging to Δ^+ have deadlines that strictly precede the deadlines of those tasks belonging to Δ^-. By definition, those tasks $\ell \in \Delta^+$ receive a strict increase in their allocation of intermittent supply under the perturbed schedule (relative to the oracle schedule) at stage $k = 0$. Thus, in order that the perturbed schedule remains feasible according to constraint (B.3), this initial increase in supply must be countered by an equivalent decrease in the allocation of supply to those tasks $\ell \in \Delta^+$ at later stages $k \geq 1$. The opposite is true for those tasks belonging to Δ^-. Accordingly, for those tasks $\ell \in \Delta^+$, we consider feasible reductions on their allocation of firm supply of the form

$$\alpha_{\ell-1}^{\ell} = -\min\left\{\delta^{\ell},\ v_{\ell-1}^{\ell,\circ}\right\} \qquad \forall\ \ell \in \Delta^+. \tag{B.4}$$

In other words, the initial increase in the allocation of intermittent supply to task $\ell \in \Delta^+$ is compensated by the maximum feasible decrease in said task's allocation of firm supply. It follows readily that $\alpha_{\ell-1}^{\ell} \leq 0$ for all $\ell \in \Delta^+$ as both terms, over which the minimization is taken in (B.4), are nonnegative; a fact that will later prove useful. And for those tasks $m \notin \Delta^+ \cup \Delta^-$ (where the oracle schedule agrees with EDF), it suffices to leave the oracle schedule unchanged. Namely, $\beta_0^m = \beta_1^m = \cdots = \beta_{m-1}^m = \alpha_{m-1}^m = 0$ for all $m \notin \Delta^+ \cup \Delta^-$.

We now proceed to establish the desired bound. Summing the demand satisfaction constraint (B.3) over $j \in \Delta^-$, we have that

$$\sum_{j \in \Delta^-} \alpha_{j-1}^{j} = -\sum_{j \in \Delta^-} \delta^j - \sum_{j \in \Delta^-} \sum_{k=1}^{j-1} \beta_k^j. \tag{B.5}$$

One can show that there exists a collection of feasible perturbations $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ on the oracle schedule (consistent with the previous restrictions), which satisfy

$$\sum_{j\in\Delta^-}\sum_{k=1}^{j-1}\beta_k^j \;=\; \min\left\{-\sum_{j\in\Delta^-}\delta^j,\;\; -\sum_{\ell\in\Delta^+}\sum_{k=1}^{\ell-1}\beta_k^\ell\right\}. \tag{B.6}$$

Equation (B.6) can be interpreted as implying the existence of feasible perturbations that enable the *full* transfer of the excess in intermittent supply derived from those tasks $\ell \in \Delta^+$ to cover the aggregate shortfall in supply incurred by those tasks $j \in \Delta^-$. Such perturbations exist, as $\ell < j$ for all $\ell \in \Delta^+$ and $j \in \Delta^-$. In other words, any excess in supply derived from those tasks $\ell \in \Delta^+$ can always be transferred to those tasks $j \in \Delta^-$, as all of the tasks belonging to Δ^+ have deadlines that strictly precede the deadlines of those tasks belonging to Δ^-.

Combining Eqs. (B.5) and (B.6), we have

$$\begin{aligned}
\sum\nolimits_{j\in\Delta^-}\alpha_{j-1}^j \;&=\; \left(-\sum\nolimits_{j\in\Delta^-}\delta^j \;+\; \sum\nolimits_{\ell\in\Delta^+}\sum\nolimits_{k=1}^{\ell-1}\beta_k^\ell\right)^+\\
&=\; \left(-\mathbf{1}^\top\boldsymbol{\delta} \;-\; \sum\nolimits_{\ell\in\Delta^+}\alpha_{\ell-1}^\ell\right)^+\\
&\le\; \left(-\sum\nolimits_{\ell\in\Delta^+}\alpha_{\ell-1}^\ell\right)^+\\
&=\; -\sum\nolimits_{\ell\in\Delta^+}\alpha_{\ell-1}^\ell,
\end{aligned}$$

where the first equality is an application of the identity $(x-y)^+ = x - \min\{x, y\}$ for all $x, y \in \mathbb{R}$, while the second equality follows from a direct substitution of the second term in the parentheses according to Eq. (B.3). And the inequality follows from the fact that $\mathbf{1}^\top\boldsymbol{\delta} \ge 0$. The final equality is a consequence of the previously established property of the chosen perturbations that $\alpha_{\ell-1}^\ell \le 0$ for all $\ell \in \Delta^+$. It follows immediately that

$$\sum_{j=1}^{N}\alpha_{j-1}^j \;\le\; 0,$$

as we've taken $\alpha_{m-1}^m = 0$ for all $m \notin \Delta^+ \cup \Delta^-$. It follows that, under the proposed modified schedule $(\widetilde{\mathbf{u}}, \widetilde{\mathbf{v}})$, we have $Q(\mathbf{x}, \mathbf{s}, \widetilde{\mathbf{u}}, \widetilde{\mathbf{v}}) \;\le\; Q(\mathbf{x}, \mathbf{s}, \mathbf{u}^{\mathrm{o}}, \mathbf{v}^{\mathrm{o}})$, where we necessarily have equality by optimality of the oracle schedule. This establishes the base step.

(Induction Step). For the induction step, we consider an oracle optimal schedule whose *first* $k-1$ *inputs* coincide with the corresponding EDF schedule. Namely, $\mathbf{u}^{\mathrm{o}} = (\mathbf{u}_0^*, \ldots, \mathbf{u}_{k-1}^*, \mathbf{u}_k^{\mathrm{o}}, \ldots, \mathbf{u}_{N-1}^{\mathrm{o}})$

and $\mathbf{v}^{\mathrm{o}} = (\mathbf{v}_0^*, \ldots, \mathbf{v}_{k-1}^*, \mathbf{v}_k^{\mathrm{o}}, \ldots, \mathbf{v}_{N-1}^{\mathrm{o}})$. Proceeding along a path of logic identical to that of the base step, one can readily construct a mapping that transforms the oracle schedule into another feasible schedule $(\widetilde{\mathbf{u}}, \widetilde{\mathbf{v}})$ whose *first* k *inputs* coincide with the EDF schedule – say, $\widetilde{\mathbf{u}} = (\mathbf{u}_0^*, \ldots, \mathbf{u}_k^*, \widetilde{\mathbf{u}}_{k+1}, \ldots, \widetilde{\mathbf{u}}_{N-1})$ and $\widetilde{\mathbf{v}} = (\mathbf{v}_0^*, \ldots, \mathbf{v}_k^*, \widetilde{\mathbf{v}}_{k+1}, \ldots, \widetilde{\mathbf{v}}_{N-1})$ – without causing an increase in the cost of firm supply. Upon realizing that one necessarily has $\mathbf{1}^\top(\mathbf{u}_k^* - \mathbf{u}_k^{\mathrm{o}}) \geq 0$, the construction of such a mapping follows from a direct application of the base step argument to the truncated oracle schedule $(\mathbf{u}_k^{\mathrm{o}}, \ldots, \mathbf{u}_{N-1}^{\mathrm{o}})$ and $(\mathbf{v}_k^{\mathrm{o}}, \ldots, \mathbf{v}_{N-1}^{\mathrm{o}})$. □

B.3 Proof of Lemma 4.8

We first establish the simplified form of Q^* in Eq. (4.13). For notational simplicity, we denote the sequence of optimal inputs by $\mathbf{u}_k^* = \boldsymbol{\mu}_k^*(\mathbf{z}_k, s_k)$ and $\mathbf{v}_k^* = \boldsymbol{\nu}_k^*(\mathbf{z}_k, s_k)$ for $k = 0, \ldots, N-1$. Under the optimal scheduling policy, firm supply is deployed only as a last resort to ensure task satisfaction. It follows that

$$Q^*(\mathbf{x}) \;=\; c_0 \cdot \sum_{k=1}^{N} \mathbb{E}\left\{ v_{k-1}^{k,*} \right\}.$$

To establish the desired result, it suffices to show that $v_{k-1}^{k,*} = -\xi_k^-$. First, define the quantity

$$\delta_k \;=\; \sum_{j=0}^{k-1} s_j \;-\; \sum_{\ell=1}^{k} \sum_{j=0}^{\ell-1} u_j^{\ell,*} \qquad \forall \;\; k = 1, \ldots, N,$$

which denotes the maximum amount of intermittent supply available to demand class $k+1$ across the first k time periods under a sequence of EDF allocations $\mathbf{u}_0^*, \ldots, \mathbf{u}_{k-1}^*$. Clearly, we have that $\delta_k \geq 0$ for all k, given feasibility of the allocations $\mathbf{u}_0^*, \ldots, \mathbf{u}_{k-1}^*$ under the intermittent supply availability constraints. One can readily show via an inductive argument that

$$\xi_k \;=\; \delta_k \;-\; v_{k-1}^{k,*} \qquad \forall \;\; k = 1, \ldots, N.$$

Using this characterization of the residual process $\boldsymbol{\xi}$, we have that

$$\begin{aligned} \delta_k > 0 \;&\Longrightarrow\; v_{k-1}^{k,*} = 0 \;\Longrightarrow\; \xi_k^- = 0, \\ \delta_k = 0 \;&\Longrightarrow\; v_{k-1}^{k,*} \geq 0 \;\Longrightarrow\; \xi_k^- = -v_{k-1}^{k,*}, \end{aligned}$$

which yields the desired result that $\xi_k^- \;=\; -\, v_{k-1}^{k,*}$ and establishes the form of Q^* in Eq. (4.13).

We now establish *convexity* and *differentiability* of the expected recourse cost $Q^*(\mathbf{x})$. First, recall that Lemma B.1 enables the expression of the expected recourse cost as $Q^*(\mathbf{x}) = \mathbb{E}\{Q(\mathbf{x}, \mathbf{s}, \mathbf{u}^o, \mathbf{v}^o)\}$, where $Q(\mathbf{x}, \mathbf{s}, \mathbf{u}^o, \mathbf{v}^o)$ denotes the optimal value of a linear program parameterized linearly by $(\mathbf{x}, \mathbf{s})$. It follows that $Q(\mathbf{x}, \mathbf{s}, \mathbf{u}^o, \mathbf{v}^o)$ is a piecewise affine convex function in $(\mathbf{x}, \mathbf{s})$ [19]. Convexity of $Q^*(\mathbf{x})$ follows, as expectation (take with respect to $\mathbf{s}$) preserves the convexity of $Q(\mathbf{x}, \mathbf{s}, \mathbf{u}^o, \mathbf{v}^o)$ in $\mathbf{x}$. And, given our prior Assumption 4.3.1 that the random process $\mathbf{s}$ have an absolutely continuous joint distribution with compact support, differentiability of $Q^*(\mathbf{x})$ over $(0, \infty)^N$ follows immediately from Proposition 2.9 in [39].

B.4 Proof of Theorem 4.14

Let $\mathbf{x}$ be the aggregate demand of the other consumers (excluding i). Suppose that the supplier uses the EDF (inter-class scheduling) policy σ^*, and an arbitrary intra-class scheduling policy $\phi \in \Phi(\sigma^*)$. We let $\pi^* = (\sigma^*, \phi)$. We consider a consumer i of some type $\theta_i = (k, R, q)$ such that $c_0 \leq R$. We will show that the consumer, who faces an arbitrary aggregate demand bundle requested by other consumers $\mathbf{x}$, would like to take the truth-telling action defined in Eq. (4.2). For the rest of this proof, we will use $\mathbf{p} = \{p_k\}_{k=1}^N$ to denote the price bundle induced by the pricing scheme $\boldsymbol{\zeta}$ (cf. its definition in Eq. (4.15)), at the aggregate demand $\mathbf{x}$.

If consumer i is truth-telling, she will request a quantity q before deadline k, and receive an expected payoff of

$$\begin{aligned} V_i^{\pi^*}(\theta_i, \varphi_i^*, \mathbf{x}, \boldsymbol{\zeta}) &= U(q) - qp_k \\ &= qR - qc_0 \left[\mathbb{P}(\xi_k \leq 0) + \sum\nolimits_{t=k+1}^{N} \mathbb{P}(\xi_k > 0, \ldots, \xi_{t-1} > 0, \xi_t \leq 0) \right]. \end{aligned} \tag{B.7}$$

Since the optimal price schedule is nonincreasing in deadline, and demand is guaranteed to be met before the requested deadline, the consumer has no incentive to request a positive amount of electricity at some period t that is earlier than k. We can therefore assume that consumer i takes an action $\mathbf{a}' = \varphi_i'(\theta_i)$ such that

$$a_t' = 0, \qquad t = 0, \ldots, k-1.$$

Since the consumer cannot increase its expected payoff (compared to being truth-telling) by reporting $a_k' \geq q$, we focus on the case where $a_k' < q$. We first write the consumer's expected payoff

(achieved by the action $\mathbf{a}'$) as

$$V_i^{\pi^*}(\theta_i, \varphi_i', \mathbf{x}, \boldsymbol{\zeta}) = \mathbb{E}\left\{U_\theta\left(\omega_{k,i}^{\pi^*}(\mathbf{x}, \mathbf{a}')\right)\right\} - p_k a_k' - \sum_{t=k+1}^{N} p_t a_t'.$$

Showing that $V_i^{\pi^*}(\theta_i, \varphi_i', \mathbf{x}, \boldsymbol{\zeta})$ is no more than the expected payoff in (B.7) is equivalent to

$$\begin{aligned} &Rq - (q - a_k')c_0\left[\mathbb{P}(\xi_k \le 0) + \sum\nolimits_{t=k+1}^{N} \mathbb{P}(\xi_k > 0, \ldots, \xi_{t-1} > 0, \xi_t \le 0)\right] \\ &\ge \mathbb{E}\left\{U_\theta\left(\omega_{k,i}^{\pi^*}(\mathbf{x}, \mathbf{a}')\right)\right\} - \sum\nolimits_{t=k+1}^{N} p_t a_t'. \end{aligned} \tag{B.8}$$

We will derive an upper bound on the right hand side of (B.8), and show that this upper bound cannot exceed the left hand side of (B.8). For notational convenience, we define

$$\eta_t = \{\xi_k > 0, \ldots, \xi_{t-1} > 0, \xi_t \le 0\}$$

for all $t = k+1, \ldots, N-1$ and

$$\eta_k = \{\xi_k \le 0\}, \qquad \eta_N = \{\xi_k > 0, \ldots, \xi_{N-1} > 0\}.$$

Note that these $N - k + 1$ events are mutually disjoint. Further, it is straightforward to see that $\xi_k \le 0$ implies that $\omega_{k,i}^{\pi}(\mathbf{x}, \mathbf{a}') = a_k'$. While, on the other hand, $\xi_k > 0$ implies that one of the (mutually disjoint) events $\{\eta_t\}_{t=k+1}^{N}$ must occur, i.e.,

$$1 = \mathbb{P}\left(\bigcup\nolimits_{t=k}^{N} \eta_t\right) = \sum_{t=k}^{N} \mathbb{P}(\eta_t). \tag{B.9}$$

Under the event η_t, the amount of energy delivered by the EDF scheduling policy before her true deadline k, $\omega_{k,i}^{\pi^*}(\mathbf{x}, \mathbf{a}')$, cannot exceed $\sum_{\tau=k}^{t} a_\tau'$, for $t = k, \ldots, N$. It follows from Assumption 4.2.1 that

$$Rq - U_\theta(x) \ge R(q - x)^+, \qquad \forall\, x \ge 0,$$

where $(\cdot)^+ = \max\{\cdot, 0\}$. We then have

$$Rq - \mathbb{E}\left\{U_\theta\left(\omega_{k,i}^{\pi^*}(\mathbf{x}, \mathbf{a}')\right)\right\} \ge R\sum\nolimits_{t=k}^{N} \mathbb{P}(\eta_t) \cdot \left(q - \sum\nolimits_{\tau=k}^{t} a_\tau'\right)^+. \tag{B.10}$$

We now argue that the right hand side of (B.10) is minimized at some vector $\tilde{\mathbf{a}}$ such that $\sum_{t=k}^{N} \tilde{a}_t \le q$. To see this, suppose that $\sum_{t=k}^{N} a_t' > q$. Let T be the smallest t such that $\sum_{m=k}^{t} a_m' > q$. We

define an alternative vector $\tilde{\mathbf{a}}$

$$\tilde{a}_t = a'_t, \qquad t \leq T-1, \qquad \tilde{a}_T = q - \sum_{m=k}^{T-1} a'_m, \tag{B.11}$$

and $\tilde{a}_t = 0$, for every $t > T$. We have $\sum_{t=k}^{N} \tilde{a}_t = q$, and

$$\sum\nolimits_{t=k}^{N} \mathbb{P}(\eta_t) \cdot \left(q - \sum\nolimits_{\tau=k}^{t} a'_\tau\right)^+ = \sum\nolimits_{t=k}^{N} \mathbb{P}(\eta_t) \cdot \left(q - \sum\nolimits_{\tau=k}^{t} \tilde{a}_\tau\right)^+.$$

Note that if $\sum_{t=k}^{N} a'_t \leq q$, then $\tilde{\mathbf{a}} = \mathbf{a'}$. To validate (B.8), it suffices to show that for any vector $\tilde{\mathbf{a}}$ defined in (B.11),

$$\begin{aligned} & R \sum\nolimits_{t=k}^{N} \mathbb{P}(\eta_t) \cdot \left(q - \sum\nolimits_{\tau=k}^{t} \tilde{a}_\tau\right) \\ \geq\ & (q - \tilde{a}_k) c_0 \left[\mathbb{P}(\xi_k \leq 0) + \sum\nolimits_{t=k+1}^{N} \mathbb{P}(\xi_k > 0, \ldots, \xi_{t-1} > 0, \xi_t \leq 0)\right] - \sum\nolimits_{t=k+1}^{N} p_t \tilde{a}_t, \end{aligned} \tag{B.12}$$

where the left hand side is a lower bound on the loss of expected utility due to the non-truthful action $\tilde{\mathbf{a}}$ (cf. the inequality in (B.10)), and the right hand side is the difference between the payment under the truth-telling action and the action $\tilde{\mathbf{a}}$. We will prove the following inequality that is equivalent to (B.12)

$$\begin{aligned} & R \sum\nolimits_{t=k}^{N} \mathbb{P}(\eta_t) \cdot \left(q - \sum\nolimits_{\tau=k}^{t} \tilde{a}_\tau\right) \\ =\ & R(q - \tilde{a}_k) - R \sum\nolimits_{t=k+1}^{N} \left(\mathbb{P}(\eta_t) \cdot \sum\nolimits_{\tau=k+1}^{t} \tilde{a}_\tau\right) \\ =\ & R(q - \tilde{a}_k) - R \sum\nolimits_{t=k+1}^{N} \left(\tilde{a}_t \cdot \sum\nolimits_{\tau=m}^{N} \mathbb{P}(\eta_m)\right) \\ \geq\ & (q - \tilde{a}_k) c_0 \left[\mathbb{P}(\xi_k \leq 0) + \sum\nolimits_{t=k+1}^{N} \mathbb{P}(\xi_k > 0, \ldots, \xi_{t-1} > 0, \xi_t \leq 0)\right] - \sum\nolimits_{t=k+1}^{N} p_t \tilde{a}_t, \end{aligned} \tag{B.13}$$

where the first equality follows from (B.9), and the second equality is obtained by rearranging terms.

It follows from (4.15) that for $t = k+1, \ldots, N$,

$$\tilde{a}_t p_t \ = \ \tilde{a}_t c_0 \sum_{m=t}^{N} \mathbb{P}(\xi_t > 0, \ldots, \xi_{m-1} > 0, \xi_m \leq 0) \geq \ \tilde{a}_t c_0 \sum_{m=t}^{N} \mathbb{P}(\xi_k > 0, \ldots, \xi_{m-1} > 0, \xi_m \leq 0). \tag{B.14}$$

We then have, for $t = k+1, \ldots, N$,

$$\begin{aligned}
& R\tilde{a}_t \sum\nolimits_{m=t}^{N} \mathbb{P}(\eta_m) - \tilde{a}_t p_t \\
&= R\tilde{a}_t \left(1 - \mathbb{P}(\xi_k \le 0) - \sum\nolimits_{m=k+1}^{t-1} \mathbb{P}(\eta_m)\right) - \tilde{a}_t p_t \\
&\le R\tilde{a}_t \left(1 - \mathbb{P}(\xi_k \le 0) - \sum\nolimits_{m=k+1}^{t-1} \mathbb{P}(\eta_m)\right) - \tilde{a}_t c_0 \sum\nolimits_{m=t}^{N} \mathbb{P}(\xi_k > 0, \ldots, \xi_{m-1} > 0, \xi_m \le 0) \\
&\le R\tilde{a}_t \left(1 - \mathbb{P}(\xi_k \le 0)\right) - \tilde{a}_t c_0 \sum\nolimits_{m=k+1}^{N} \mathbb{P}(\xi_k > 0, \ldots, \xi_{m-1} > 0, \xi_m \le 0) \\
&\le R\tilde{a}_t - \tilde{a}_t c_0 \left(\mathbb{P}(\xi_k \le 0) + \sum\nolimits_{m=k+1}^{N} \mathbb{P}(\xi_k > 0, \ldots, \xi_{m-1} > 0, \xi_m \le 0)\right).
\end{aligned} \tag{B.15}$$

Here, the first inequality follows from (B.14); the second inequality is true, because $R \ge c_0$ and $\eta_m = \{\xi_k > 0, \ldots, \xi_{m-1} > 0, \xi_m \le 0\}$, for $m = k+1, \ldots, t-1$; the last inequality follows from $R \ge c_0$. For any vector $\tilde{\mathbf{a}}$ with $\sum_{t=k}^{N} \tilde{a}_t \le q$, from (B.15) we have

$$\begin{aligned}
& \sum\nolimits_{t=k+1}^{N} \left(-p_t \tilde{a}_t + R\tilde{a}_t \sum\nolimits_{m=t}^{N} \mathbb{P}(\eta_m)\right) \\
&\le \sum\nolimits_{t=k+1}^{N} \tilde{a}_t \left[R - c_0 \left(\mathbb{P}(\xi_k \le 0) + \sum\nolimits_{m=k+1}^{N} \mathbb{P}(\xi_k > 0, \ldots, \xi_{m-1} > 0, \xi_m \le 0)\right)\right].
\end{aligned} \tag{B.16}$$

Since $R \ge c_0$, the right hand side of (B.16) is nondecreasing in $\sum_{t=k+1}^{N} \tilde{a}_t$. The desired result in (B.13) follows from the fact $\sum_{t=k+1}^{N} \tilde{a}_t \le q - \tilde{a}_k$. Since the preceding analysis holds for any action $\mathbf{a}'$ and any aggregate demand bundle $\mathbf{x}$, we conclude that it is a dominant strategy for consumer i to be truth-telling, i.e., the pricing scheme (4.15) is incentive compatible, in the sense of Definition 4.12.

Bibliography

[1] Afèche, P. 2013. Incentive-compatible revenue management in queueing systems: optimal strategic delay, *Manufacturing & Service Operations Management*, **15**(3), 423–443.

[2] Albadi, M. H., E. F. El-Saadany. 2007. Demand response in electricity markets: An overview. *Proc. of Power Engineering Society General Meeting.*

[3] Arrow, K. J., L. Hurwicz. 1958. On the stability of the competitive equilibrium, I. *Econometrica*, **26**, 522–552.

[4] Arrow, K. J., H. D. Block, L. Hurwicz. 1959. On the stability of the competitive equilibrium, II. *Econometrica*, **27**, 82–109.

[5] Bisin, A., P. Gottardi. 1999. Competitive equilibria with asymmetric information. *Journal of Economic Theory*, **87**(1), 1–48.

[6] Bitar, E., P. Khargonekar, K. Poolla, R. Rajagopal, P. Varaiya, F. Wu. 2012. Selling random wind, in *Proc. of 45th Hawaii International Conference on System Science (HICSS).*

[7] Bitar, E., S. Low. 2012. Deadline differentiated pricing of deferrable electric power service, in *Proc. of the 51th IEEE Conference on Decision and Control.*

[8] Botsford, C., A. Szczepanek. 2009. Fast charging vs. slow charging: Pros and cons for the new age of electric vehicles. *EVS24, Stavanger, Norway.*

[9] Borenstein, S. 2005. The long-run effciency of real-time electricity pricing, *The Energy Journal*, **26**(3), 93–116.

[10] Bradford, R. M. 1996. Pricing, routing, and incentive compatibility in multiserver queues, *European Journal of Operational Research*, **89**(2), 226–236.

[11] Callaway, D. S., I. A. Hiskens. 2011. Achieving controllability of electric loads, *Proc. of the IEEE*, **99**, 184–199.

[12] CAISO. 2010. Integration of renewable resources: operational requirements and generation fleet capability at 20% RPS. California Independent System Operator Report.

[13] Chao, H. P., R. Wilson. 1987. Priority service: pricing, investment, and market organization, *The American Economic Review*, **77**(5), 899–916.

[14] Chen, S., T. He, L. Tong. 2011. Optimal deadline scheduling with commitment. *Proc. of Allerton Conference on Communication, Control, and Computing.*

[15] Chen, F. 2000. Sales-force incentives and inventory management. *Manufacturing and Service Operations Management*, **2**(2), 186–202.

[16] Cheng, J. Q., M. P. Wellman. 1998. The WALRAS algorithm: A convergent distributed implementation of general equilibrium outcomes. *Computational Economics*, **12**(1), 1–24.

[17] Etezadi-Amoli, M., K. Choma, J. Stefani. 2010. Rapid-charge electric-vehicle stations. *IEEE Trans. Power Delivery*, **25**(3), 1883–1887.

[18] FERC. 2009. A national assessment of demand response potential. FERC Staff Report.

[19] Gal, T., J. Nedoma. 1972. Multiparametric linear programming. *Management Science*, **18**(7), 406–422.

[20] Hicks, J. R. 1939. *Value and Capital* (Vol. 1). Oxford: Oxford University Press.

[21] Hogan, W. 2010. Demand response pricing in organized wholesale markets. http://www.hks.harvard.edu/fs/whogan/Hogan_IRC_DR_051310.pdf

[22] Ilic, M. D., L. Xie, J. Y. Joo. 2011. Efficient coordination of wind power and price-responsive demand-Part I: Theoretical foundations. *IEEE Tran. on Power Systems*, **26**(4), 1875–1884.

[23] Kefayati, M., R. Baldick. 2011. Energy delivery transaction pricing for flexible electrical loads. *Proc. of IEEE International Conference on Smart Grid Communications (SmartGridComm).*

[24] Lederer, P. J., L. Li. 1997. Pricing, production, scheduling, and delivery-time competition. *Operations Research*, **45**(3), 407–420.

[25] Madaeni, S. H., R. Sioshansi. 2013. Using demand response to improve the emission benefits of wind. *IEEE Trans. on Power Systems*, **28**(2), 1385–1394.

[26] Mendelson H., S. Whang. 1990. Optimal incentive-compatible priority pricing for the M/M/1 queue, *Operations Research*, **38**(5), 870–883.

[27] Mierendorff, K. 2011. Optimal dynamic mechanism design with deadlines. Columbia University, *Working paper*. http://dirkbergemann.commons.yale.edu/files/Mierendorf.pdf

[28] Negrete-Pincetic, M., S. P. Meyn. 2012. Markets for differentiated electric power products in a Smart Grid environment. *Proc. of IEEE Power and Energy Society General Meeting*.

[29] Oren, S. S., S. A. Smith. 1993. *Service Opportunities for Electric Utilities: Creating Differentiated Products*, Kluwer Academic Publishers: Boston.

[30] Oren, S. S., S. A. Smith. 1992. Design and management of curtailable electricity service to reduce annual peaks. *Operations Research*, **40**(2), 213–228.

[31] Papavasiliou, A., and S. Oren. 2010. Supplying renewable energy to deferrable loads: Algorithms and economic analysis. *Proc. of IEEE Power & Energy Society General Meeting*.

[32] Plambeck, E. L., S. A. Zenios. 2003. Incentive efficient control of a make-to-stock production system. *Operations Research*, **51**(3), 371–386.

[33] Prescott, E. C., R. M. Townsend. 1984. Pareto optima and competitive equilibria with adverse selection and moral hazard. *Econometrica*, **52**(1), 21–46.

[34] Radhakrishnan, S., K. R. Balachandran. 1995. Delay cost and incentive schemes for multiple users, *Management Science*, **41**(4), 646–652.

[35] Robu, V., S. Stein, E. H. Gerding, D. C. Parkes, A. Rogers, N. R. Jennings. 2012. An online mechanism for multi-speed electric vehicle charging. *Auctions, Market Mechanisms, and Their Applications, Lecture Notes of the Institute for Computer Sciences* **80**, 100–112.

[36] Rockafellar, R. T., S. Uryasev, M. Zarabankin. 2002. Deviation Measures in Risk Analysis and Optimization, *Research Report 2002-7*, Department of Industrial and Systems Engineering, University of Florida, Gainesville, FL.

[37] Roozbehani, M., A, Faghih, M. Ohannessian, M. A. Dahleh. 2011. The intertemporal utility of demand and price-elasticity of consumption in power grids with shiftable loads, in *Proc. of the 50th IEEE Conference on Decision and Control and European Control Conference*.

[38] Roozbehani, M., M. A. Dahleh, S. K. Mitter. 2013. Volatility of power grids under real-time pricing. To appear in *IEEE Tran. Power Systems*.

[39] Shapiro, A., D. Dentcheva, A. P. Ruszczynski. 2009. *Lectures on Stochastic Programming: Modeling and Theory*, **9**, SIAM.

[40] Sioshansi, R. 2012. Modeling the impacts of electricity tariffs on plug-in hybrid electric vehicle charging, costs, and emissions. *Operations Research*, **60**(3), 506–516.

[41] Spees, K., L. B. Lave. 2007. Demand response and electricity market efficiency. *The Electricity Journal*, **20**(3), 69–85.

[42] Stoft, S. 2002. *Power System Economics: Designing Markets for Electricity*, IEEE Press, Piscataway, NJ.

[43] Subramanian, A., M. Garcia, A. Domínguez-García, D. Callaway, K. Poolla, P. Varaiya. 2012. Real-time scheduling of deferrable electric loads, in *Proc. of American Control Conference (ACC).*

[44] Tan C. W., P. P. Varaiya. 1993. Interruptible electric power service contracts, *Journal of Economic Dynamics and Control*, **17**, 495–517.

[45] Van Mieghem, J. A. 2000. Price and service discrimination in queuing systems: incentive compatibility of Gcμ scheduling, *Management Science*, **46**(9), 1249–1267.

[46] Wilson, R. 2002. Architecture of power markets, *Econometrica*, **70**(4), 1299–1340.

[47] Yilmaz, M., P. T. Krein. 2013. Review of battery charger topologies, charging power levels, and infrastructure for plug-in electric and hybrid vehicles, *IEEE Tran. Power Electronics*, **28**(5), 2151–2169.

[48] Zachariadis, T., N. Pashourtidou. 2007. An empirical analysis of electricity consumption in Cyprus, *Energy Economics*, **29**(2), 183–198.

[49] B. Kirby, "Ancillary services: Technical and commercial insights," Tech. Rep., July 2007.

[50] J. Smith, M. Milligan, E. DeMeo, and B. Parsons, "Utility wind integration and operating impact state of the art," *IEEE Transactions on Power Systems*, vol. 22, no. 3, pp. 900 –908, aug. 2007.

[51] Y. Makarov, C. Loutan, J. Ma, and P. de Mello, "Operational impacts of wind generation on california power systems," *IEEE Transactions on Power Systems*, vol. 24, no. 2, pp. 1039 –1050, may 2009.

[52] S. Meyn, M. Negrete-Pincetic, G. Wang, A. Kowli, and E. Shafieepoorfard, "The value of volatile resources in electricity markets," in *IEEE conference on Decision and Control (CDC)*, 2010.

[53] U. Helman, "Resource and transmission planning to achieve a 33% RPS in California–ISO modeling tools and planning framework," in *FERC Technical Conference on Planning Models and Software*, 2010.

[54] Pacific Gas and Electric Company, "Smart AC program." [Online]. Available: http://www.pge.com/en/myhome/saveenergymoney/energysavingprograms/smartac/index.page

[55] "Buildings energy data book." [Online]. Available: http://buildingsdatabook.eren.doe.gov/default.aspx

[56] "U.S. Energy Information Administration, annual energy review," 2010. [Online]. Available: http://www.eia.gov/totalenergy/data/annual/#consumption

[57] H. Hao, B. Sanandaji, K. Poolla, and T. Vincent, "A generalized battery model of a collection of thermostatically controlled loads for providing ancillary service," in *the 51th Annual Allerton Conference on Communication, Control and Computing*, 2013.

[58] D. S. Callaway, "Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy," *Energy Conversion and Management*, 2009.

[59] R. Malhame and C.-Y. Chong, "Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system," *IEEE Transactions on Automatic Control*, 1985.

[60] S. Kundu, N. Sinitsyn, S. Backhaus, and I. Hiskens, "Modeling and control of thermostatically controlled loads," in *the 17-th Power Systems Computation Conference*, 2011.

[61] S. Bashash and H. K. Fathy, "Modeling and control of aggregate air conditioning loads for robust renewable power management," *IEEE Transactions on Control Systems Technology*, 2013.

[62] S. Shao, M. Pipattanasomporn, and S. Rahman, "Development of physical-based demand response-enabled residential load models," *IEEE Transactions on Power Systems*, vol. 28, no. 2, pp. 607–614, 2013.

[63] J. L. Mathieu, S. Koch, and D. S. Callaway, "State estimation and control of electric loads to manage real-time energy imbalance," *IEEE Transactions on Power Systems*, 2013.

[64] W. Zhang, J. Lian, C.-Y. Chang, and K. Kalsi, "Aggregated modeling and control of air conditioning loads for demand response," *IEEE Transactions on Power Systems*, 2013.

[65] J. L. Mathieu, M. Kamgarpour, J. Lygeros, and D. S. Callaway, "Energy arbitrage with thermostatically controlled loads," in *European Control Conference (ECC)*, 2013.

[66] N. Lu, "An evaluation of the HVAC load potential for providing load balancing service," *IEEE Transactions on Smart Grid*, vol. 3, no. 3, pp. 1263–1270, 2012.

[67] E. Vrettos, S. Koch, and G. Andersson, "Load frequency control by aggregations of thermally stratified electric water heaters," in *3rd IEEE PES International Conference and Exhibition on Innovative Smart Grid Technologies*. IEEE, October 2012, pp. 1–8.

[68] H. Chao, "Price-responsive demand management for a smart grid world," *The Electricity Journal*, vol. 23, no. 1, pp. 7–20, 2010.

[69] H. Shu, R. Yu, and S. Rahardja, "Dynamic incentive strategy for voluntary demand response based on TDP scheme," in *Signal & Information Processing Association Annual Summit and Conference (APSIPA ASC), 2012 Asia-Pacific*. IEEE, 2012, pp. 1–6.

[70] NERC, "Data collection for demand side management for quantifying its influence on reliability: Results and recommendations," Tech. Rep., December 2007. [Online]. Available: http://www.nerc.com/docs/pc/drdtf/NERC_DSMTF_Report_040308.pdf

[71] M. Roozbehani, M. A. Dahleh, and S. K. Mitter, "Volatility of power grids under real-time pricing," *IEEE Transactions on Power Systems*, 2011.

[72] J. Mathieu, M. Dyson, and D. Callaway, "Using residential electric loads for fast demand response: The potential resource and revenues, the costs, and policy recommendations," in *2012 ACEEE Summer Study on Energy Efficiency in Buildings*, 2012.

[73] A. Subramanian, M. Garcia, A. Dominguez-Garcia, D. Callaway, K. Poolla, and P. Varaiya, "Real-time scheduling of deferrable electric loads," in *American Control Conference (ACC)*, 2012, pp. 3643–3650.

[74] J. MacDonald, P. Cappers, D. Callaway, and S. Kiliccote, "Demand response providing ancillary services," in *Grid-Interop 2012*, Irving, Texas, December 2012.

[75] "PJM regulation data." [Online]. Available: http://www.pjm.com/markets-and-operations/ancillary-services/mkt-based-regulation/fast-response-regulation-signal.aspx

[76] B. M. Sanandaji, H. Hao, K. Poolla, and T. L. Vincent, "Improved battery models of an aggregation of thermostatically controlled loads for frequency regulation," *to appear in American Control Conference*, June 2014. [Online]. Available: http://www.eecs.berkeley.edu/~sanandaji/my_papers.html

[77] A.R. Bergen and V. Vittal, "Power Systems Analysis," Prentice-Hall, Englewood Cliffs, NJ, 2000.

[78] H. Farhangi, "The path of the Smart Grid," IEEE Power and Energy, vol.8, no. 1, pp: 18-28, 2010.

[79] I. Foster, C. Kesselman, S. Teucke, "The Anatomy of the Grid: Enabling Scalable Virtual Organizations," International Journal of High Performance Computing Applications, vol. 15, no. 3, pp: 200-22, 2001.

[80] V. Vyatkin *et al.*, "Towards intelligent Smart Grid devices with IEC 61850 interoperability," Transmission and Distribution Conference and Exposition, IEEE PES, 2010.

[81] Wikipedia: Community Choice Aggregation, Community Solar Farms.

[82] Enernoc Corporation. http://www.enernoc.com/

[83] K. Voss [et al], "Load Matching and Grid Interaction of Net Zero Energy Buildings," Proceedings of EUROSUN 2010 International Conference on Solar Heating, Cooling and Buildings, Graz, Austria.

[84] Marin Clean Energy (MCE) in California, www.marincleanenergy.info

[85] J.L. Mathieu, S Koch, and D.S. Callaway, "State estimation and control of electric loads to manage real-time energy imbalance," *IEEE Trans. Power Syst.*, 28(1): 430-440, Feb. 2013.

[86] K. Qian, C. Zhou, M. Allan, and Y. Yuan,"Modeling of load demand due to EV battery charging in distribution systems," *IEEE Trans. Power Syst.*, 26(2): 802-810, May 2011.

[87] D. Brooks, J.Gannon, L. Tran and L. Marshall, "2010 Resource Adequacy Report", California Public Utilities Commission, April 22, 2011.

[88] U.S. Department of Energy, "The potential benefit of Distributed Generation and rate-related issues that may impede their expansion: A study pursuant to Section 1817 of the Energy Policy act of 2005," Feb. 2007.

[89] M. Lehtonen, and S. Nye, "History of electricity network control and distributed generation in the UK and Western Denmark," *Energy Policy*, 37(6): 2338-2345, June 2009.

[90] A. Subramanian, M. Garcia, A. Domínguez-García, D. Callaway, K.Poolla, and P. Varaiya,"Real-time scheduling of deferrable electric loads," *Proc. of the American Controls Conf.*, Montreal, 2012.

[91] S. Baruah, and J. Goossens, "Scheduling real-time tasks: Algorithms and Complexity," in *Handbook of Scheduling: Algorithms, Models and Performance Analysis,* J.Y-T. Leung, Ed. Boca Raton, FL: Chapman Hall / CRC Press, 2004, Chpt. 28.

[92] C.L. Liu, and J.W. Layland, "Scheduling algorithms for multiprogamming in a hard-real-time environment," *Journal of the ACM*, 20(1): 46-61, Jan. 1973.

[93] A.K. Mok, "The design of real-time programming systems based on process models," *Proc. 1984 Real-Time System Symp.*, Austin, Dec. 1984.

[94] M.L. Dertouzos, and A.K. Mok, "Multiprocessor online scheduling of hard-real-time tasks," *IEEE Trans. Software Engineering*, 15(12): 1497-1506, Dec. 1989.

[95] J. Hong, X. Tan, and D. Towsley, " A performance analysis of minimum laxity and earliest deadline scheduling in a real-time system," *IEEE Trans. Computers*, 38(12): 1736-1744, Dec. 1989.

[96] S. Chen, T. He, and L. Tong, "Optimal deadline scheduling with commitment", *49th Allerton Conf. Comm., Control, and Computing*, 2011.

[97] C.E. García, D.M. Prett, and M. Morari, "Model predictive control: Theory and practice - A survey," *Automatica*, 25(3): 335-348, May 1989.

[98] Z. Ma, I. Hiskens, and D. Callaway, "A decentralized MPC strategy for charging large populations of plug-in electric vehicles," *Proc. 18th IFAC World Congr.*, Milan, 2011.

[99] G. Hug-Glanzmann, "Coordination of intermittent generation with storage, demand control and conventional energy sources", *Proc. IREP 2010 - Bulk Power Syst. Dynamics and Control - VIII*, 2010.

[100] M. Galus, R. la Fauci, and G. Andersson, "Investigating PHEV wind balancing capabilities using heuristics and model predictive control", *IEEE Power & Energy Society General Meeting*, 2010.

[101] Society of Automotive Engineers (SAE), "SAE Electric Vehicle Conductive Charge Coupler, SAE J1772 REV. MONTH01," Prepared by the SAE EV Charging Systems Committee, Aug. 2001.

[102] K. Morrow, D. Karner, and J. Francfort, "Plug-in hybrid electric vehicle charging instrastructure review," Report INL/EXT-08-15058, U.S. Department of Energy Vehicle Technologies Program - Advanced Vehicle Testing Activity, Nov. 2008.

[103] U. Helman, "Resource and transmission planning to achieve a 33% RPS in California - ISO Modeling tools and planning framework," *FERC Tech. Conf. on Planning Models and Software*, 2010.

[104] California Independent System Operator (CAISO), "Integration of renewable resources: Operational requirements and Generation fleet capability at 20% RPS," Tech. Rep., Aug. 2010.

[105] J. Driesen, and F. Katiraei, "Design for distributed energy resources," *IEEE Power and Energy Magazine*, 6(3): 30-40, May-June 2008.

[106] N. Hatziargyriou, H. Asano, R. Iravani, and C. Marnay, "Microgrids," *IEEE Power and Energy Magazine*, 5(4): 78-94, July-Aug. 2007.

[107] M. Lijesen, "The real-time price elasticity of electricity," *Energy Economics*, 29(2): 249-258, March 2007.

[108] S. Borenstein, "The long-run efficiency of real-time electricity pricing," *The Energy Journal*, 26(3): 93-116, 2005.

[109] G. Barbose, C. Goldman, and B. Neenan, "A survey of utility experience with real time pricing," *Lawrence Berkeley National Laboratory*, Technical Report LBNL-54238, Dec. 2004.

[110] K. Spees, and L.B. Lave, "Demand response and electricity market efficiency," *The Electricity Journal*, 20(3): 69-85, April 2007.

[111] M. Roozbehani, M. Dahleh, and S.K. Mitter, "On the stability of wholesale electricity markets under real-time pricing," *Proc. 49th IEEE Conf. on Decision and Control*, 2010.

[112] M. Ilic, L. Xie, and J.Y. Joo, "Efficient coordination of wind power and price-responsive demand - Part I: Theoretical Foundations," *IEEE Trans. Power Syst.*, 26(4): 1875-1884, Nov. 2011.

[113] T.F. Lee, M.Y. Cho, Y.C. Hsiao, P.J. Chao, and F.M. Fang, "Optimization and implementation of a load control scheduler using relaxed dynamic programming for large air conditioner loads," *IEEE Trans. Power Syst.*, 23(2): 691-702, May 2008.

[114] Y. Y. Hsu, C. C. Su, "Dispatch of direct load control using dynamic programming," *IEEE Trans. Power Syst.*, 6(3): 1056-1061, Aug. 1991.

[115] A. Papavasiliou, and S. Oren, "Supplying renewable energy to deferrable loads: Algorithms and economic analysis," *IEEE Power & Energy Society General Meeting*, 2010.

[116] L. Gan, U. Topcu, and S. Low, "Stochastic distributed protocol for electric vehicle charging with discrete charging rate," *IEEE Power & Energy Society General Meeting*, 2012.

[117] C. Zhou, K. Qian, M. Allan, and W. Zhou, "Modeling of the cost of EV battery wear due to V2G application in power systems," *IEEE Trans. Energy Convers.*, 26(4): 1041-1050, Dec. 2011.

[118] K. Mets, T. Verschueren, W. Haerick, C. Develder, and F. De Turck, "Optimizing smart energy control strategies for plug-in hybrid electric vehicle charging," *IEEE/IFIP Network Operations and Management Symp. Workshops*: 293-299, 2010