

Multi-scale Transactive Control In Interconnected Bulk Power Systems Under High
Renewable Energy Supply and High Demand Response Scenarios

by

David P. Chassin

BSc., Building Science, Rensselaer Polytechnic Institute, 1987

MASc., Mechanical Engineering, University of Victoria, 2015

A Dissertation Submitted in Partial Fulfillment of the
Requirements for the Degree of

DOCTOR OF PHILOSOPHY

in the Department of Mechanical Engineering

© David P. Chassin, 2017
University of Victoria

All rights reserved. This thesis may not be reproduced in whole or in part, by
photocopying or other means, without the permission of the author.

Multi-scale Transactive Control In Interconnected Bulk Power Systems Under High
Renewable Energy Supply and High Demand Response Scenarios

by

David P. Chassin

BSc., Building Science, Rensselaer Polytechnic Institute, 1987

MASc., Mechanical Engineering, University of Victoria, 2015

Supervisory Committee

Dr. Ned Djilali, Supervisor
(Department of Mechanical Engineering)

Dr. Yang Shi, Departmental Member
(Department of Mechanical Engineering)

Dr. Panajotis Agathoklis, Outside Member
(Department of Electric Engineering)

ABSTRACT

This thesis presents the design, analysis, and validation of a hierarchical transactive control system that engages demand response resources to enhance the integration of renewable electricity generation resources. This control system joins energy, capacity and regulation markets together in a unified homeostatic and economically efficient electricity operation that increases total surplus while improving reliability and decreasing carbon emissions from fossil-based generation resources.

The work encompasses: (1) the derivation of a short-term demand response model suitable for transactive control systems and its validation with field demonstration data; (2) an aggregate load model that enables effective control of large populations of thermal loads using a new type of thermostat (discrete time with zero deadband); (3) a methodology for optimally controlling response to frequency deviations while tracking schedule area exports in areas that have high penetration of both intermittent renewable resources and fast-acting demand response; and (4) the development of a system-wide (continental interconnection) scale strategy for optimal power trajectory and resource dispatch based on a shift from primarily energy cost-based approach to a primarily ramping cost-based one.

The results show that multi-layer transactive control systems can be constructed, will enhance renewable resource utilization, and will operate in a coordinated manner with bulk power systems that include both regions with and without organized power markets. Estimates of Western Electric Coordinating Council (WECC) system cost savings under target renewable energy generation levels resulting from the proposed system exceed US\$150B annually by the year 2024, when compared to the existing control system.

Contents

Supervisory Committee	ii
Abstract	iii
Table of Contents	iv
List of Tables	viii
List of Figures	x
Nomenclature	xiv
Acknowledgements	xxiv
Funding Sources	xxix
Dedication	xxx
1 Introduction	1
1.1 Motivation	4
1.2 Main Contributions	5
1.3 Outline of the Thesis	6
2 Problem Statement	8
2.1 Barriers To Integrated Demand Response	9
2.2 Achieving Optimality in Transactive Systems	14
3 Demand Response	17
3.1 Background	18
3.2 Random Utility Model	21
3.2.1 First-principles Model	22

3.2.2	Model Assumptions	24
3.2.3	Equilibrium Demand Response	26
3.3	Model Validation	27
3.4	Summary of Results	33
4	Aggregation	34
4.1	Background	34
4.2	Aggregate Load Curtailment Model	36
4.2.1	Aggregate Load Model	37
4.2.2	Load Control Model	41
4.2.3	Open-Loop Response	42
4.2.4	Model Identification	43
4.3	Aggregate Demand Response Controller Design	45
4.3.1	Proportional Control	47
4.3.2	Proportional-Derivative Control	48
4.3.3	Unity Damped Control	50
4.3.4	Deadbeat Control	52
4.3.5	Pole Placement Control	52
4.3.6	Integral Error Feedback	53
4.4	Agent-based Simulation Results	54
4.5	Summary of Results	57
5	Regulation	58
5.1	Methodology	58
5.1.1	Frequency control mechanism	59
5.1.2	Transactive control platform	60
5.1.3	Demand response integration in the 5-minute market	61
5.1.4	\mathcal{H}_2 -optimal control policy	64
5.2	Implementation	67
5.2.1	State-space realization	70
5.2.2	Model Validation	71
5.3	Control Performance	72
5.4	Summary of Results	75
6	Dispatch	76
6.1	Background	77

6.2	Methodology	78
6.3	Optimal Dispatch Controller	82
6.4	Performance Evaluation	84
6.5	Case Study: WECC 2024	87
6.6	Summary of Results	93
7	Discussion	94
7.1	Demand Response	95
7.1.1	Model Limitations	95
7.1.2	Technical and Regulatory Impacts	99
7.2	Aggregation	101
7.3	Regulation	103
7.3.1	Robustness to FADR Uncertainty	103
7.4	Dispatch	105
7.5	Ramping Market Price	107
7.5.1	Unified Market Design	109
7.6	Recommendations	111
8	Conclusions	113
8.1	Principal Contributions and Findings	115
8.2	Future Research	116
8.2.1	Demand Response	117
8.2.2	Aggregate Thermostatic Load Control	117
8.2.3	Regulation	118
8.2.4	Dispatch	119
8.3	Summary	120
A	Demand Elasticity	121
A.1	Analysis	122
A.2	Discussion	123
A.2.1	Device Control	123
A.2.2	Business Models	124
A.2.3	Regulatory Oversight	124
A.3	Conclusion	125
B	Price Negotiation Convergence	126

B.1	Transactive Price-Discovery	127
B.2	Stable Mechanism Design	130
B.2.1	Demand Response Uncertainty	132
B.3	Conclusions	134
C	Model Specifications	137
D	Simulation Results	143
	Bibliography	144

List of Tables

Table 3.1	Feeder characteristics	29
Table 3.2	Olympic data analysis results	30
Table 3.3	Columbus analysis results for Feeders 120, 140, 160 and 180	31
Table 3.4	Columbus analysis results for only non-experiment days	32
Table 4.1	House thermal parameters.	41
Table 4.2	Controller design configurations.	47
Table 4.3	Maximum attenuating proportional control gains for various conditions.	48
Table 4.4	Proportional-derivative controller design parameters.	49
Table 4.5	Controller design parameters for peak load (-15°C).	55
Table 5.1	Cost, generation, and net export impacts of ACE control versus \mathcal{H}_2 -optimal control	74
Table 6.1	Marginal prices and marginal costs for 105 GWh schedule at 100 GW initial power and 10 GW/h ramp for cases in Figure 6.5	90
Table 6.2	Single hour cost savings under low and high renewable for a ramp from 100 GW to 110 GW at 5 minute discrete dispatch control update rate, with varying energy error redispatch	90
Table 6.3	WECC 2024 cost savings from optimal dispatch under different transmission constraint and renewable scenarios	91
Table 6.4	Summary of energy and price impacts of optimal dispatch control for the WECC 2024 base case	91
Table 7.1	Bid price entropy statistics	99
Table C.1	Aggregate Load Model Parameters	137
Table C.2	WECC 2024 demand forecast and internal area losses [1]	138
Table C.3	WECC 2024 aggregated installed supply capacity [1]	139

Table C.4	WECC 2024 producer cost and surplus difference for 100% in-elastic load (in M\$/year) [1]	140
Table C.5	WECC 2025 production cost per unit in \$/MWh [1]	141
Table C.6	WECC 2024 model inter-area transfer limits [1]	142
Table D.1	Estimated, simulated, and errors of aggregate thermostat load state transition probability ρ using joint PDF (N=normal, Ln=log normal, Log=logistic) using 10^6 homes and $t_d = 1$ minute. . . .	143

List of Figures

Figure 2.1	Top-to-bottom rethink of electricity infrastructure, including providers of transmission and distribution infrastructure, system operators and resource aggregators.	10
Figure 2.2	The California “Duck Curve” [Source: CAISO]	11
Figure 2.3	Inter-temporal data flow diagram.	14
Figure 3.1	Demand curves for steady state thermostatic loads in a trans-active control system.	26
Figure 3.2	Example of demand function model validation with Columbus demonstration data. The bids shown are from 2013-06-22 22:45 EDT. The clearing price P_C and quantity Q_C are indicated by the circle. The expected price \bar{P} and quantity \bar{Q} are indicated by the plus sign. The standard deviation of price \tilde{P} and quantity \tilde{Q} are indicated by the ellipse.	29
Figure 4.1	State-space model of aggregate conventional thermostatic loads in heating regime with refractory states n_{on}^* and n_{off}^* . $\Delta\tau$ is the difference between the indoor and outdoor temperatures and δ is the hysteresis band limit.	36
Figure 4.2	General state-space model of discrete-time zero-deadband aggregate thermostatic loads.	37
Figure 4.3	Discrete-time heating thermostat transition probabilities for <i>on</i> and <i>off</i> states with PDF of 10^6 homes with a setpoint change of -0.1°F	39
Figure 4.4	Zero and pole locations of Equation (4.4) for a random population of 1 million homes with thermal parameters given in Table IV at various outdoor air temperatures.	42

Figure 4.5	Open loop impulse (left), decay (center), and step (right) response of aggregate load model compared to agent-based simulation for 100,000 thermostats per unit input u at -10°C	43
Figure 4.6	General structure of the controller (top): Block (A) is the aggregate load model, (B) is the reduced-order observer, (C) is the load controller, and (D) is the integral error feedback.	46
Figure 4.7	Discrete-time root-locus of aggregate $T\delta_0$ thermostatic loads.	47
Figure 4.8	100 MW proportional load control step response with maximum attenuating proportional control gains based on the load parameters in Table 4.1 (left) and proportional-derivative control step response (right).	49
Figure 4.9	Unity damped system diagram.	51
Figure 4.10	Unity damped (left) and deadbeat (right) responses of aggregate load controllers.	51
Figure 4.11	100 MW impulse response (left) and proportional control (right) at -15°C	56
Figure 4.12	100 MW unity damping load control (left) and deadbeat load control (right) at -15°C	56
Figure 4.13	100 MW tuned load control (left) and integral feedback control (right) at -15°C	57
Figure 5.1	System frequency control diagram.	60
Figure 5.2	Inter-area transfer flows within an interconnected system consisting of N control areas.	61
Figure 5.3	Five-minute resource dispatch with supply (red) and demand (blue) response to a loss of generation (ΔQ_s).	62
Figure 5.4	Real-time response of generation and load to a disturbance.	63
Figure 5.5	Standard system for \mathcal{H}_2 -optimal control synthesis	64
Figure 5.6	System frequency and control area export regulation control diagram.	68
Figure 5.7	Model component frequency (f p.u. nominal frequency), load (l p.u. area load) and generation (g p.u. system load) responses to a local disturbance (ΔQ_s p.u. system load).	68
Figure 5.8	Control area model in standard form.	70

Figure 5.9	ACE control (black) and \mathcal{H}_2 -optimal (blue) control performance for design conditions (5% FADR), showing the raw ACE signal (a p.u. area load), area generation output (p p.u. system load), and system frequency (f p.u. nominal frequency) response to a loss of generation within the control area.	72
Figure 5.10	ACE control (left) and \mathcal{H}_2 -optimal control (right) model validation for varying demand response level with generation response (g p.u. area load), demand response (l p.u. area load), and generation regulation cost (cg in \$/h p.u. area load).	73
Figure 5.11	ACE control (left) and \mathcal{H}_2 -optimal control (right) cost and dispatch for varying demand response levels, where C_a is the total control cost (in \$), C_Q is the power control response cost (in \$), and C_f is the frequency control response cost (in \$).	74
Figure 6.1	Power (left) and ramp (center and right) price functions. . . .	79
Figure 6.2	Optimal dispatch controller diagram with discrete update time t_s	82
Figure 6.3	Optimal discrete time control for various values of ω at $t_s = 5$ minutes.	84
Figure 6.4	Conventional power dispatch for base case: (top) significant negative schedule error requiring over-production, (center) small negative, zero and positive schedule error requiring over (red), normal (black) and under (blue) production, and (bottom) significant positive schedule error requiring under-production. . . .	86
Figure 6.5	Single hour optimal dispatch for low (top) and high (bottom) renewables with a ramp from 100 GW to 110 GW using a 10-minute discrete-time dispatch control rate, with hourly energy schedule correction errors varying from -5% to $+5\%$	89
Figure 6.6	WECC 2024 load duration (top) and optimal dispatch savings duration (bottom) using discrete optimal control at 5-minute dispatch rate for the unconstrained (left) and constrained (right) high renewables scenario. The scatter plots are the corresponding cost (top) and load (bottom) values for the durations curves shown.	92
Figure 7.1	Physical (top) and temporal (bottom) system diagram of trans-active control systems.	95

Figure 7.2	Demand (left) and elasticity (right) curves for a nominal case with $\eta_D = -2.5$ and $\bar{P} = 0.5$	96
Figure 7.3	Demand response state diversity duration curves for the Olympic feeder and Columbus feeder numbers 120, 140, 160, and 180.	99
Figure 7.4	Sensitivity of savings to marginal price of ramping resources.	106
Figure 7.5	Energy market with supply curve (red) and demand curve (blue), with the price P_C at which the export ΔQ is realized.	108
Figure 7.6	Ramping market with supply curve (red) and demand curve (blue), with the price P_C at which the export ΔQ is realized.	109
Figure 7.7	Price-based control over-actuation of aggregate load	109
Figure 7.8	Market-based unified price-based control of aggregate load	110
Figure B.1	Logistic map iteration sequence of price discovery mechanisms in a transactive system for $a < -b$ (left) and $a > -b$ (right).	128
Figure B.2	Advanced negotiation strategy diagram with quantity constraint tracking	131
Figure B.3	Simulation of stable (left), marginal (center) and unstable (right) negotiations without (dotted) and with (solid) corresponding deadbeat negotiation strategies.	132
Figure B.4	Discrete-time system diagram of advanced negotiation strategy diagram with demand curve uncertainty	133
Figure B.5	Integral error feedback negotiation strategy for the same cases as Figure B.3 with a +10% error in the demand response curve estimate \hat{b}	134

Nomenclature

\bar{P}	Expected price average; in \$/MWh (demand response).
\bar{Q}	Mean heating system power; in W (aggregation).
\bar{q}	The population average device load; in MW (aggregation).
\bar{R}	The population average duty cycle; unitless (aggregation).
\bar{r}_{off}	The unweighted mean rate of temperature change of all devices that are <i>off</i> ; °C/s (aggregation).
\bar{r}_{on}	The unweighted mean rate of temperature change of all devices that are <i>on</i> ; °C/s (aggregation).
$\bar{\rho}$	Mean thermostat duty cycle; unitless (demand response).
β	Price response function parameter; unitless (demand response).
$\ddot{\tau}$	The second time derivative of the indoor air temperature; °C/s ² (aggregation).
\ddot{Q}	Ramping rate of change; in MW/h ² (dispatch).
$\Delta\tau$	The actual temperature deviation from the desired temperature τ_D ; in °C (aggregation).
$\Delta f(t)$	System frequency deviation; in Hz (regulation).
$\Delta Q(t)$	Area net power exports deviation; in MW (regulation).
$\Delta Q_s(t)$	Disturbance magnitude; in MW (regulation).
Δt	The elapsed time in a state; in seconds (aggregation).
δ	The hysteresis band half-width; in °F (aggregation).

$\dot{\tau}$	The rate at which a device temperature changes; in °C/s (aggregation).
\dot{Q}	Ramping; in MW/h (dispatch).
\dot{Q}^*	Discrete power at next time step; in MW (dispatch).
\dot{Q}_0	Initial ramping; in MW/h (dispatch).
\dot{Q}_T	Terminal ramping; in MW/h (dispatch).
η_D	Demand elasticity; unitless (demand response).
$\hat{a}_s(s)$	Filtered ACE signal in s -domain (regulation).
$\hat{f}(s)$	Interconnection frequency response in s -domain (regulation).
$\hat{f}(s)$	System frequency; in s -domain (regulation).
$\hat{g}_d(s)$	Droop-controlled generation response; in s -domain (regulation).
$\hat{g}_r(s)$	ACE-controlled generation response response in s -domain (regulation).
$\hat{l}(s)$	Load response in s -domain (regulation).
$\hat{Q}(s)$	Interconnection power response in s -domain (regulation).
λ	Lagrange multiplier (including Q_z); in \$/MWh (dispatch).
\mathbf{c}	The output gains $\begin{bmatrix} c_1 & c_2 \end{bmatrix}$; 2×1 vector (aggregation).
\mathbf{c}	The output vector $\begin{bmatrix} \bar{q} & 0 \end{bmatrix}$; 2×1 vector (aggregation).
\mathbf{h}	The input gains $\begin{bmatrix} h_1 \\ h_2 \end{bmatrix}$; 2×1 vector (aggregation).
\mathbf{K}_c	The observer output control gain model vector; 2×1 vector (aggregation).
\mathbf{x}	The system state $\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$; 2×1 vector (aggregation).
$\tilde{\mathbf{c}}$	The observer output model vector; 2×1 vector (aggregation).
$\tilde{\mathbf{h}}$	The observer input model vector; 2×1 vector (aggregation).
μ	Lagrange multiplier (excluding Q_z); in \$/MWh (dispatch).
ν	Linear demand response scale function; unitless (demand response).

ω	Square root of energy to ramping marginal price ratio; in h^{-1} (dispatch).
ρ	Duty cycle of thermostatic load; unitless (demand response).
ρ'_{off}	The complimentary load-weighted population average rate of temperature change for devices that are <i>off</i> ; $^{\circ}\text{C}/\text{s}$ (aggregation).
ρ'_{on}	The complimentary load-weighted population average rate of temperature change for devices that are <i>on</i> ; $^{\circ}\text{C}/\text{s}$ (aggregation).
ρ_{off}	The load-weighted population average rate of temperature change for devices that are <i>off</i> ; in $^{\circ}\text{C}/\text{s}$ (aggregation).
ρ_{on}	The load-weighted population average rate of temperature change for devices that are <i>on</i> ; in $^{\circ}\text{C}/\text{s}$ (aggregation).
$\sigma(Q)$	Load entropy; unitless (demand response).
σ^2	The variance of the rate of change of indoor air temperature; $^{\circ}\text{C}/\text{s}$ (aggregation).
σ_0	Maximum load entropy; unitless (demand response).
σ_{off}^2	The variance of the rate of change of indoor air temperature when the heating/cooling system is <i>off</i> ; $^{\circ}\text{C}/\text{s}$ (aggregation).
σ_{on}^2	The variance of the rate of change of indoor air temperature when the heating/cooling system is <i>on</i> ; $^{\circ}\text{C}/\text{s}$ (aggregation).
τ	The device temperature; in $^{\circ}\text{C}$ (aggregation).
τ_A	The temperature of the indoor air; $^{\circ}\text{C}$ (aggregation).
τ_D	The desired device temperature; in $^{\circ}\text{C}$ (aggregation).
τ_M	The temperature of the solid mass; $^{\circ}\text{C}$ (aggregation).
τ_O	The outdoor air temperature; $^{\circ}\text{C}$ (aggregation).
τ_{hys}	Thermostatic control hysteresis; in $^{\circ}\text{F}$ (demand response).
τ_{obs}	Observed thermostat state; in $^{\circ}\text{F}$ (demand response).

τ_{off}	Duration of ‘off’ period for thermostatic load; in hours (demand response).
τ_{off}	The temperature of a device that is <i>off</i> ; °C (aggregation).
τ_{off}	Thermostatic steady state when ‘off’; in °F (demand response).
τ_{on}	Duration of ‘on’ period for thermostatic load; in hours (demand response).
τ_{on}	The temperature of a device that is <i>on</i> ; °C (aggregation).
τ_{on}	Thermostatic steady state when ‘on’; in °F (demand response).
τ_{set}	Thermostat set point ; in °F(demand response).
\tilde{P}	Expectedated price standard deviation; in \$/MWh (demand response).
A	A cost parameter; unit varies according to context (dispatch).
a	Load price response function zero-order constant; unitless (demand response).
a	Marginal price of energy; in \$/MW ² ·h (dispatch).
a	The principal pole of the discrete-time system model; in s ⁻¹ (aggregation).
$A(t)$	Raw ACE signal; in MW (regulation).
B	A cost parameter; unit varies according to context (dispatch).
B	Bid price; in \$/MWh (demand response).
B	Frequency control bias; in MW/Hz (regulation).
b	Load price response function first-order constant; unitless (demand response).
b	Marginal price of power; in \$/MW ² (dispatch).
b	The principal zero of the discrete-time system model; in s ⁻¹ (aggregation).
C	A cost parameter; unit varies according to context (dispatch).
C	The system controllability matrix; 2 × 2 matrix (aggregation).
c	Marginal price of ramping; in \$·h/MW ² (dispatch).
$C(t)$	Cost over the time interval 0 to t ; in \$ (dispatch).

- C^* Cost associated with discrete time control; in \$ (dispatch).
- c_1 The average load of a unit of n_{on} devices; in MW (aggregation).
- c_2 The average load of a unit of n_{off} devices; in MW (aggregation).
- C_A The heat capacity of the indoor air; in J/K (aggregation).
- C_M The heat capacity of the solid mass; in J/K (aggregation).
- C_{base} Cost associated with base case control; in \$ (dispatch).
- COP HVAC system efficiency; unitless (aggregation).
- D A cost parameter; unit varies according to context (dispatch).
- D Interconnection damping constant; per unit (regulation).
- d Disturbance magnitude; in MW (regulation).
- $d(kt_d)$ Slope of the load demand curve at the dispatch point; in \$/MW²h (regulation).
- D_q Load diversity at order q ; unitless (demand response).
- $E(t)$ Energy over the time interval 0 to t ; in MWh (dispatch).
- E_{Δ} Energy demand parameter; in MWh (dispatch).
- E_T Energy over T ; in MWh (dispatch).
- $eta(P)$ Demand elasticity at the price P ; unitless (demand response).
- $F(s)$ Low-pass ACE control signal filter transfer function (regulation).
- $f(t)$ System frequency; in Hz.
- F_d Fraction of total load that can be responsive to frequency (regulation).
- F_r Fraction of generating units that are ACE-controlled (regulation).
- f_s Nominal or scheduled system frequency; in Hz.
- G The state transition matrix of the population of devices; 2×2 matrix (aggregation).

- $g(k)$ Load probability function; unitless (demand response).
- $G(t, Q, \dot{Q})$ Cost Lagrangian; in \$ (dispatch).
- $G_d(s)$ Droop-controlled generation response transfer function (regulation).
- $G_r(s)$ ACE-controlled generation resource transfer function (regulation).
- G_{pd} The proportional-derivative control transfer function; unitless (aggregation).
- H Shannon entropy of load; unitless (demand response).
- h The observer scalar reference input gain; unitless (aggregation).
- $H(s)$ Interconnection overall transfer function (regulation).
- h_1 The net number of devices added to the controlled on population of devices by a unity input signal; unitless (aggregation).
- h_2 The net number of devices added to the controlled off population of devices by a unity input signal; unitless (aggregation).
- K Bid comfort control setting; in \$/MWh. $^{\circ}$ F (demand response).
- K The closed-loop proportional control gain; unitless (aggregation).
- k Discrete time step; in p.u. t_s (dispatch).
- k The discrete time index such that $t = kt_s$; unitless (aggregation).
- k Thermostatic device population count; unitless (demand response).
- k_1 The derivative gain of the closed-loop proportional-derivative control; in seconds (aggregation).
- k_2 The proportional gain of the closed-loop proportional-derivative control; unitless (aggregation).
- K_d Fraction of total load that is armed by 5-minute dispatch (regulation).
- K_l Load control recovery time constant; in seconds (regulation).
- K_p Load quasi-steady rebound response time constant; in seconds (regulation).

- K_q The observer integral error feedback control gain model scalar; unitless (aggregation).
- $L(s)$ Load transfer function (regulation).
- M Interconnection inertial constant; in seconds (regulation).
- N Number of thermostats under control; unitless (demand response).
- N_C Number of controllers; unitless (aggregation).
- n_{off} Number of devices that are off but not locked; unitless (aggregation).
- n_{off}^* Number of devices that are locked off; unitless (aggregation).
- n_{on} Number of devices that are on but not locked; unitless (aggregation).
- n_{on}^* Number of devices that are locked on; unitless (aggregation).
- O The system observability matrix; 2×2 matrix (aggregation).
- $P(Q)$ Power price function; in \$/MWh (dispatch).
- $P(t)$ Regulation energy price; in \$/MWh (regulation).
- P_d 5-minute dispatch energy price; in \$/MWh.
- P_s Hourly schedule energy price; in \$/MWh.
- Q Heating system power demand; in W (aggregation).
- Q Load; in MW (demand response).
- q The augmented state for integral error feedback control; in J (aggregation).
- q The total heat added to the device; in W (aggregation).
- $Q(t)$ Actual net exports from a control area; in MW (regulation).
- $Q(t)$ Total power; in MW (dispatch).
- Q^* Discrete power; in MW (dispatch).
- Q_0 Initial load; in MW (dispatch).

- Q_{Δ} Power demand parameter; in MW (dispatch).
- Q_E Scheduled load; in MW (dispatch).
- q_H Heating system output; in W (aggregation).
- q_I The heat added from internal, solar and ventilation gains; in W (aggregation).
- Q_R Most probable load; in MW (demand response).
- Q_s Scheduled net exports from a control area; in MW (regulation).
- q_S The heat added/removed by the heating/cooling system; in W (aggregation).
- Q_T Terminal load; in MW (dispatch).
- Q_U Unresponsive load; in MW (demand response).
- Q_z Must-take generation; in MW (dispatch).
- R Droop control constant; unitless (regulation).
- r Thermostatic state decay rate ; in °F/h (demand response).
- $R(Q, \dot{Q})$ Ramping price function; in \$/MW (dispatch).
- r_{off} The rate at which a device temperature changes when off; in °C/s (aggregation).
- r_{on} The rate at which a device temperature changes when on; in °C/s (aggregation).
- s Frequency domain complex variable; in h^{-1} (dispatch).
- s The Laplace domain complex variable s ; in s^{-1} (aggregation).
- $s(kt_d)$ Slope of the generation supply curve at the dispatch point; in \$/MW²h (regulation).
- T Interval terminating time; in hours (dispatch).
- t Real time variable; in seconds.
- t Time domain real variable; in hours (dispatch).

t_d	Dispatch control discrete-time sampling rate; in minutes (scheduling).
T_f	ACE control signal filter time constant in seconds (regulation).
T_g	Generation resource governor time constant; in seconds (regulation).
T_l	Load control derivative response gain (regulation).
t_r	Discrete control discrete-time sampling rate; in seconds (scheduling).
t_s	Discrete control discrete-time sampling rate; in hours (scheduling).
t_s	The discrete controller sampling time; in seconds (aggregation).
t_s	Time step; in seconds (dispatch).
T_{ch}	Generation resource steam chest time constant; in seconds (regulation).
t_{min}	The minimum control lockout time; in seconds (aggregation).
U	The z -domain transformation of the input u (aggregation).
U_A	The thermal conductivity between indoor and outdoor air; in W/K (aggregation).
U_M	The thermal conductivity between indoor air and solid mass; in W/K (aggregation).
$W(z)$	Lambert W-function; unitless (demand response).
x_1	The first state of the state vector \mathbf{x} , which is the number of devices on, n_{on} ; unitless (aggregation).
x_2	The first state of the state vector \mathbf{x} , which is the number of devices off, n_{off} ; unitless (aggregation).
Y	The z -domain transformation of the output y (aggregation).
y	The net load of the population of controlled devices.
z	The discrete-time z -domain variable $z = e^{-st}$; unitless (aggregation).
z_1, z_2	The desired poles for the integral error feedback control design; in s^{-1} (aggregation).

z_q The desired dominant pole for the integral error feedback ground; in s^{-1} (aggregation).

ACKNOWLEDGEMENTS

Most honest stories about how one gets something done probably should start with a confession, and mine is a long one. I consider it above average hubris when I say I hope to accomplish something that others might find useful enough for me to be remembered well after I am gone. I think this is what normal people of a certain age do. But mine is actually no greater a sin than looking for a change of scenery when in 1992 after 10 years in upstate New York I quit graduate school and left my successful start-up company after I became disillusioned with both architecture as a field of research and the stressful life of an entrepreneur.

I went west to the desert of central Washington State to start over with no expectations and no grand vision. I certainly did not know anything about, least of all expect to be more than an “extra” in the momentous changes that were coming to the electric power industry. In any event, I certainly had no inkling how my decision to start over could lead me through Victoria to Menlo Park writing these words. Now I only hope to set the record as straight as I recall it.

When I arrived at Pacific Northwest National Laboratory I began a twenty year series of more or less random encounters and collaborations with people who together would change the way I saw and understood the world of large-scale engineered systems and set me thinking about how and why one would go about trying to change the way we design and operate them.

In the mid-1990s, sensing diminished research returns in the already-crowded field of building energy efficiency I began looking for something new and exciting on which to work. At the Laboratory Richard Quadrel was beginning ground-breaking research on architecture and engineering design tools that employed the latest artificial intelligence methods. Michael Brambley was just starting a new program in building system diagnostics. Together they turned me to the world of engineering design tools and the question of how one uses automation to improve system operation through advanced controls and diagnostics. Robert Pratt was just coming off the seminal end-use load characterization and assessment program for Bonneville Power Administration, which provided a wealth of data about building energy consumption at the end-use level. Although budget cuts were a constant threat, I was very fortunate to lead the Building Sciences group just at a time when we began thinking about what would be the next big thing in buildings research.

At that time I shared an office suite with Landis Kannberg who managed the

electric power engineering group. Jeff Dagle and Matt Donnelly were among the people who began asking similar questions from the electric grid perspective and it was inevitable that these two groups would join forces. By 1999 the nugget of an idea borne of long hours thinking and talking about what might happen if buildings were more actively part of the power system operations.

This notion became the Energy Systems Transformation Initiative. Initially, Steve Hauser was hired to manage the initiative. He brought a grand vision and a wealth of connections to people who shared it. (Among them was Jesse Berst who Steve credits with coining the term “transactive”, although I didn’t learn that until years later). Soon after his arrival Steve drew on my whiteboard a simple taxonomy for how end-use devices are controlled in energy systems. This taxonomy has become one the key elements of what we (and some now regret) call the “smart grid”.

Passive: Devices that simply react to their environment and cannot take action autonomously to adjust in any non-trivial way. A conventional household thermostat is passive because if the price of electricity goes up or the frequency of the grid drops suddenly, it keeps on going without regard to anything other than what the temperature of your home is.

Active: Devices that engage in more intelligent responses such as reducing consumption when prices rise or frequency drops, but do so in a completely autonomous manner without input from the user of the device.

Interactive: Devices that act on the basis of interactions with consumers. Such a device might not reduce consumption as much when prices are high if you are having a dinner party or your elderly parent is staying with you during a heat wave.

Transactive: Devices that exchange information with other devices to help the system as a whole find a more efficient way to operate.

Steve made one vital contribution to my understanding of this taxonomy: it doesn’t apply a value judgment on devices in the different categories, i.e., active isn’t better than passive or worse than interactive. But it does tell us something about what we should expect from devices in each class. The next 10 years of research became about understanding those expectations and how one would go about designing and operating an entirely new kind of system using them.

Soon after the initiative began, Ross Guttromson joined the Laboratory. He brought a pragmatic approach to engineering from his years in the nuclear navy and working for large industrial engineering firms. One of his first assignments was to work with me to develop a new kind of simulation environment that would allow us to study what this new system might do. The product of that collaboration became the US Department of Energy's premiere tool for simulating the smart grid. Now known as GridLAB-D, I managed its development until 2013 and so it naturally ends up figuring prominently in the present work. Eric Lightner at the US Department of Energy's Office of Electricity realized early on the need for such a tool and provided the funding needed to ensure it was delivered and supported for the many years it needed to mature and become accepted by the smart grid research community.

Meanwhile, Donald Hammerstrom was chosen in 2005 to lead a project to demonstrate some of the Laboratory's initial thinking in this area. Together with Lynne Kiesling and Preston Michie, we planned, designed, deployed, operated and analyzed what is widely regarded now as the first operational retail real-time pricing system using a distribution capacity double auction mechanism. This project was completed in 2007 and became known as the Olympic Peninsula demonstration and serves to this day as a significant milestones in the development of transactive technology. Most of the data from that project is used in the present work and serves as the reference model for many subsequent studies in transactive control, including this one.

From about 2006 on, I had the great fortune to work with and learn from some of the best power engineers in the world, both at PNNL and as a part of my work on the Western Electric Coordinating Council's Load Modeling Task Force. I credit much of the practical knowledge I have of power systems to Jason Fuller, Frank Tuffner, and Kevin Schneider at PNNL, Clint Gerkenmeyer at Benton Rural Electric Cooperative, Dmitry Kosterev at Bonneville Power Administration, Richard Bravo at Southern California Edison, Bernie Lesieutre at the University of Wisconsin in Madison, Joe Eto at Lawrence Berkeley National Laboratory and John Undrill at the University of Arizona.

In 2009 the American Renewable and Reinvestment Act provided funding for two additional demonstrations of transactive control, one of which was managed by Steve Widergren. The AEP gridSMART demonstration project in Columbus Ohio provided the second data set for a retail capacity double auction and formalized many of the lessons learned in the Olympic Peninsula. Much of the credit for the results of that project go the large team of anonymous but outstanding engineers and operators at

American Electric Power and for their fortitude and perseverance in generating an extremely valuable data set for research, for which I must offer my deepest thanks as well.

It became clear after two field trials that transactive control system implementations were too “far ahead of the headlights”. I felt that it was time to step back and understand how and why this all really worked. In the fall of 2012 I applied to the University of Victoria graduate mechanical engineering program to begin pursuing this question as a PhD topic. Then quite unexpectedly, PNNL was looking for a new investment opportunity and in January 2013, PNNL dusted off an old idea I had proposed back in 2001. A new laboratory initiative called the “Control of Complex Systems” was started in April 2013, one month before I was scheduled to begin my coursework. To his credit, Suresh Baskaran placed enormous faith in our ability to sell the idea that control theory needed a new kick to meet the demands of the coming 21st Century grid, that PNNL was uniquely positioned to take on that challenge, and that DOE would soon be making a commitment to a new program in advanced controls. During the summer and fall of 2013 with support and encouragement from Marylin Quadrel and many others at PNNL I led the initiative development team from Victoria while carrying a full graduate course-load. By October 2013 the new Control of Complex Systems Initiative plan was submitted and accepted by the Laboratory. In January 2014 Jakob Stoustrup was recruited from Aalborg University in Denmark to manage the new initiative. One of Jakob’s first decisions was to make transactive control the centerpiece of the new initiative. Thus it came to pass that my research on transactive control at the University of Victoria was aligned so well with the research agenda at PNNL. I am deeply grateful to the support that the Laboratory provided me during the first 3 years of this research.

But in an unexpected twist, Sila Kiliccote and Mark Hartney at SLAC National Accelerator Laboratory offered to support my research and join the new Grid Integration Systems and Mobility Group. It was an offer I could not refuse and I moved to Silicon Valley in the summer of 2016. I had the great fortune then to work with a completely new group of truly brilliant scientists and engineers, including Ram Rajagopal, Abbas El Gamal, Claudio Rivetta, Emre Can Kara, and Mayank Malik, and with the support of Arum Majumdar, the co-director of the Stanford’s Precourt Institute for Energy I was able complete this work in late 2017.

Among all those who supported my efforts and influenced it the most I must acknowledge Sahand Behboodi, whose collaboration and contribution to my thinking

about the challenges we faced in completing our respective research is unmatched by anyone else. He provided a crucial combination of deep insight, broad thinking and highly disciplined approach to modeling and simulation that complement my own predilections well. Our collaboration over the years has led to the series of papers that are the basis of this thesis and which I am optimistic will be the starting point for many years of follow-up work in this area.

I must also thank the faculty at the University of Victoria who helped me “retool” so that I would be better equipped to conduct this research. I particularly want to recognize Panajotis Agathoklis, Wu-Sheng Lu, Ben Nadler, Daniel Rondeau, Andrew Rowe, Yang Shi, and Hong-Chuan Yang, all of whom I hold in the highest regard for their consummate dedication to the craft and teaching me the finer points in the practice of their field.

I would also like to offer a special thanks Susan Walton and Pauline Shepherd of the Institute for Integrated Energy Systems for their endless patience and direct support of me and Norma over the years.

Above all, my supervisor Ned Djilali receives special recognition not only for patiently corralling my wide-ranging interests and channeling it toward a finished product, but also for encouraging me to collaborate so much with other students and faculty at UVic and elsewhere.

Finally, I cannot offer enough thanks to my family, who supported, pushed, cajoled, threatened, and otherwise assisted me in accomplishing my lifelong goal. To my wife Norma, my mother Ann, her husband Jeffrey, my father John and my children Nik, Forrest and Isaac, all of whom contributed in their own way, I now offer you all an official record of the recognition you deserve for the roles you played in helping me get it done.

Whether the world is better as a result of the present work is a judgment I can only leave to others. As for myself, I am content to submit this thesis as a record of my best attempt to heed the teachings of my forefathers...

*You are not expected to complete the task [of creation],
but neither are you free to desist from trying.*

— R. Tarfon, *Pirkei Avot*

FUNDING SOURCES

Funding for this research is from both US and Canadian sources, including

The US Department of Energy through Pacific Northwest National Laboratory, which is operated by Battelle Memorial Institute under Contract No. DE-AC05-76RL01830, and through SLAC National Accelerator Laboratory, which is operated by Stanford University under Contract No. DE-AC02-76-SF00515.

The Pacific Institute for Climate Solutions through the Institute for Integrated Energy Systems at the University of Victoria.

DEDICATION

To my wonderful, wise and serene wife Norma,
For trustingly following me to the edge of the Earth,
For confidently guiding me back when I got too close to it,
For stoically supporting me when I was too weak to continue,
And for lovingly walking with me through all the good times.

Chapter 1

Introduction

The National Academy of Engineering considers bulk electric power generation and distribution to be the single most important engineering achievement of the 20th century [2]. Every part of today's trans-national economy is supported in some way by electrification. A myriad power plants convert primary energy sources such as fossil and nuclear fuels, hydrological cycles, wind and solar radiation to high quality, reliable, and versatile electric energy that is used to drive an economic engine without parallel in all of human history. Arguably, not since the discovery of fire or the invention of writing has a single idea so dramatically affected every aspect of the human condition.

But the story of electricity is still being written. The 20th century model of the top-down utility and the interconnected transmission network is being challenged by 21st century problems. Climate change concerns, environmental impacts, low natural gas prices and the lack of prospective sites for large reservoirs are driving out coal plants, nuclear reactors and large-scale hydro-electric facilities as the primary sources of electric power. In many systems solar and wind resources have become an increasingly significant share of the generation resource mix.

However, with these new resources come new planning and operation problems. In particular, prime mover intermittency and the lack of control over their power output are putting strains on interconnected systems and forcing system operators to either place limits on intermittent generation resources or engage new kinds of resources in the control of the system. Over the last few decades extensive research has been conducted on how controllable loads in particular can be engaged in the various planning and operation processes in bulk power systems. Many solutions to key elements of the problem have been proposed, some of which have been successfully demonstrated

in field trials, and a few of which have found their way to full-scale operation by utilities. Nearly all of these solutions utilize centralized “top-down” control methods, and most operate in an open-loop control regime. These solutions are often simple, and are demonstrably sufficient for modest levels of renewable penetration.

But centralized methods can be inflexible and lack robustness to variability and availability of both renewable supply and controllable demand resources. Decentralized control approaches are already widespread in bulk power systems, as for example in the case for regional scheduling using organized energy markets, or control area generation regulation and bulk system frequency support. Taking decentralized approaches to their logical limit, fully distributed approaches have been proposed, particularly for managing local capacity limits and under-frequency load shedding.

This thesis examines the feasibility of scaling-up to an entire interconnection a particular distributed method of integrating controllable resources called “transactive control”. The problem is considered in the context of deep decarbonization of the bulk electric power systems, with particular attention to the use of loads as resources in scheduling, dispatch and regulation processes. While the problem involves regulatory, economic, and policy considerations, the main focus of this thesis is on the technical problem and solutions that support a flexible approach to regulation, economic and policy questions. This thesis proposes general methods that might be used to implement solutions that are widely applicable to the multi-scale approaches enabled by a transactive control paradigm.

Among the most important drivers in the evolution of modern electricity infrastructure is the effect of carbon emissions on the environment. For decades we have been aware of the effect of power plant emissions on the air quality, rain water acidification, and the concentration of carbon dioxide in the atmosphere. Measures to reduce the effect of soot, nitrogen and sulfur oxides in the atmosphere downwind from fossil fueled power plants have largely been successful. Clean air policies and regulations have enabled some recovery of watersheds in the Northeast that were once suffering from acidification despite some persistent long term effect in certain ecosystems [3, 4].

The 2015 Paris climate accord had raised hopes that a global comprehensive atmospheric carbon policy was at hand, and the adoption of the Clean Power Plan by the United States was a positive sign that together with China, the world’s leading economies would take effective measures to push toward effective mitigation of the impact of carbon at global temperatures rise. The recent policy reversals are widely regarded as a setback in this regard. But many consider the trend toward greater

dependence on renewable resources irreversible, simply because the social, policy and market conditions increasingly favor renewable electricity generation resources [5].

Unfortunately, increasing demand for renewable electricity generation resources does not automatically bring about adoption of technologies that mitigate the climate impact of fossil-based electricity generation and satisfy ever growing electric system load. Each class of renewable generation comes with one or more disadvantages that limit the extent to which it can be integrated in bulk system operation. Hydro-electric generation has long been employed as a significant renewable source of electricity. But climate change may jeopardize the magnitude and certainty with which the existing asset base can meet demand [6, 7], while population displacement, habitat destruction and fish stock degradation limit the growth of new assets. Shifts in both load and hydro generation potentially increase uncertainty in long term planning and further enhance the need for technical approaches that support operational flexibility [8].

Similar issues arise with wind and solar generation resources. Wind power has seen rapid growth in recent years, but system reliability requirements can limit the penetration of wind generation without additional mitigation measures such as firming resources [9]. Solar resources are also becoming increasingly available but have intermittency challenges similar to those of wind. In addition, residential rooftop solar resources are challenging the classical utility revenue model [10] and are known to cause voltage control issues in distribution systems [11].

Finally, the reliable, robust control and optimal operation of an increasingly complex bulk electricity system have become very real concerns [12]. Many see 100% penetration of renewable generation as the principal objective. But this may only be possible in certain regions and only if there is a nearby bulk electricity interconnection on which such a region could rely when renewable intermittency causes shortfalls in energy supply. There are good reasons why it may not be possible for an entire interconnection to be supplied 100% by wind and solar energy, leaving open a role for nuclear, hydro, and natural gas.

The traditional utility approach to renewable intermittency is to allocate additional firm generation resources to replace all potentially non-firm renewable resources. These firm resources are often fast-responding thermal fossil resources and hydro resources when and where available. For new renewable resources the impact of this approach is quantified as an intermittency factor, which discounts the contribution of wind in addition to its capacity factor and limits the degree to which it can contribute to meeting peak demand [13]. However, the intermittency factor does

not account for the ramping requirements created by potentially fast-changing renewable resources [14]. The need for fast-ramping resources discourages the dispatch of slower high-efficiency fossil and nuclear generation assets while promoting faster low-efficiency fossil and hydro, where available, for regulation and reserve services [15].

1.1 Motivation

The motivation for this thesis is two-fold. First, the long-term average cost of new renewables energy resources must be covered by the short-term price volatility in electricity markets. However, as the penetration of renewables increases, the cleared energy price is more frequently zero, or even negative as fossil-fired generators attempt to remain online while waiting for prices to rise. This decline in revenues could place an economic throttle on the growth of renewables [16] that can only be mitigated by enabling new sources for the revenues necessary to pay for utility infrastructure, particularly in the presence of high levels of distributed renewable resources. Furthermore, ramping resource scarcity may induce high price volatility in the ancillary services markets, which may lead to price shocks to the system overall. This work seeks to enable mechanisms that significantly reduce costs, stabilize prices, and enable new revenues from the coordination of the resources that support efficient system scheduling, dispatch, and regulation.

Second, although there is an existing mechanism called “transactive control” that has shown promise in field demonstrations, to date this technology has not been scaled fully to an entire interconnection. There are many reasons why this has not yet occurred, but chief among them are the obstacles to modeling, simulating and evaluating the performance of transactive systems at scale. This work seeks to identify models and methods that can be used to perform such evaluations.

Market-based mechanisms lie at the heart of transactive control systems. So deploying transactive control at the interconnection scale will require the almost ubiquitous use of markets to allocate any and all scarce resources in the systems. The long-term barrier to accomplishing such an ambitious goal is that the existing transactive system design is focused almost exclusively on the energy market-based dispatch of demand-side resources. But the value of demand response alone may be quite small [17]. The same may be said for energy storage [18]. As a result we must find other ways to provide incentives for needed resources to participate in market-based solu-

tions. This includes markets to facilitate shifting costs (or revenues) away from (or to) non-energy markets, such as power/capacity or ramping/regulation markets. If the volume of services traded in these three markets were more balanced, then we should expect a more adaptable, equitable, and stable economic regime. Hence, through the concept of transactive control we should expect a more adaptable, equitable, and stable technical operation as well.

Thus we are motivated to understand first how a more balanced transactive system might function, second how much benefit it provides globally, and finally whether all concerned parties are better off participating in it than withdrawing from it. This thesis therefore focuses on how we model and evaluate the performance of key elements of transactive systems when operated at the interconnection scale.

1.2 Main Contributions

The first major contribution of this thesis develops and assesses the performance of a short-term demand response (DR) model for utility load control with applications to resource planning and control design. Long term demand models tend to underestimate short-term demand response when it is induced by prices. This has two important consequences. First, planning studies tend to undervalue DR and often overlook its benefits in utility demand management program development. Second, when DR is not overlooked, the open-loop DR control gain estimate may be too low. This can result in overuse of load resources, control instability and excessive price volatility. Our objective is therefore to develop a more accurate and better performing short-term demand response model. We construct the model from first principles about the nature of thermostatic load control and show that the resulting formulation corresponds exactly to the Random Utility Model employed in economics to study consumer choice.

The second major contribution of this thesis demonstrates a utility-scale direct load control problem, where the controlled loads are discrete-time zero-deadband residential thermostats that allow frequent utility-dispatched micro-adjustments to the consumer's heating/cooling setpoints. These new digital thermostats can serve as the basis for highly accurate and stable closed-loop direct load control systems, as well as price-based indirect load control systems. A new aggregate load model for discrete-time zero-deadband thermostats is developed and its fundamental characteristics are described from first-principles.

The third major contribution of this thesis develops an \mathcal{H}_2 -optimal power regulation scheme for balancing authorities to provide regulation services using both generation and load resources in the presence of a significant amount of intermittent renewable generation. The optimal controller is designed to minimize the loss of total economic surplus due to deviations from the schedule and dispatch resulting from system contingencies.

The fourth major contribution of this thesis considers the optimal resource dispatch problem for distribution-level resources that are sensitive to both energy and ramping prices. Resource aggregators and load-serving entities that use price-based resource control must solve an economic optimization problem to determine the optimal dispatch of distributed generation, storage, and load resources during each scheduling interval. The solution to this problem provides the basis for significant cost savings at the interconnection level.

1.3 Outline of the Thesis

The main elements of this thesis are presented in six parts. Chapter 2 introduces the challenges of including and optimizing the scheduling, dispatch, and regulation of aggregated controllable demand resources in the presence of multiple price signals from the wholesale and retail markets, and the transactive approach to solving this class of problem. Chapter 3 presents an economic model of demand response under transactive control. This model focuses on the bid behavior and price responses of the aggregate load resources that participate in retail energy markets. Chapter 4 develops a statistical model of aggregate load dynamics. The purpose of this model is to enable modeling of aggregate dynamics of loads after they receive price signals from retail energy markets. Chapter 5 examines a new control model of aggregate load and uses it to design an optimal frequency response control strategy for a control area that includes fast-acting demand response resources. Chapter 6 derives an optimal dispatch strategy and evaluates its performance under hour-ahead scheduling from wholesale markets. The strategy is developed to facilitate economically optimal dispatch when energy resources are plentiful but ramping resources are scarce. Finally, in Chapter 7 the results of these approaches are discussed and some concepts for future research are presented.

Supporting material may be found in the appendices. Appendix A and B briefly present background material on ramping price elasticity and price stability in trans-

active systems. The remaining appendices provide supporting information to assist in reproducing and building upon the results of this research. An auxiliary report is available on arXiv for those who seek background information of power system operations, demand response and transactive control [19].

Chapter 2

Problem Statement

In 2003 Economics Nobel Laureate Vernon Smith published an editorial with Lynne Kiesling in the Wall Street Journal [20] summarizing the consensus in the wake of the California Electricity crisis. In their view the crisis was in part precipitated by the lack of customer engagement in electricity pricing mechanisms [21]. Reflecting on the technical and regulatory supply-side response to the crisis, they wrote “What is inadequately discussed, let alone motivated, is the [other] option – demand response”. It is now widely accepted that demand-side resources can mitigate the market power of energy suppliers. More importantly, demand response presents a real opportunity for improvement in electricity planning and operations. Research on short-term demand-side resources in particular has increased as the growth of intermittent wind and solar resources further exacerbates the problem of managing the balance between supply and demand in power systems [22].

For demand-side resources to serve as a reliable option for utilities to mitigate the renewable resources intermittency, system operators prefer to control distributed loads in real-time using strategies similar to those used for generators. This is an emerging challenge in systems where demand-side resources are expected to play a significant role in mitigating the adverse impacts of renewable intermittency on key system control functions like frequency regulation, schedule tracking, and local voltage support [23]. Transactive control was conceived as an efficient approach to integrate demand resources, as well as other distributed resources that could benefit system operations, such as rooftop photovoltaics and electric vehicle battery chargers [24, 25]. The multi-scale and multi-temporal paradigm can efficiently integrate wholesale energy, capacity, and regulation markets at the bulk system level with distribution operations, where demand response resource are aggregated and dispatched [26].

2.1 Barriers To Integrated Demand Response

Demand response has long been considered a low-cost alternative to added generation capacity [27]. Demand is now also being considered as an alternative to fast-response generation reserves to reduce the dispatch of inefficient generation resources [28]. But load control strategies for demand response applications can be challenging to plan and operate, and little has been done to quantify their economic impact at the interconnection level. This is in part because the competing objectives of local and global control [29, 30]. It is also in part because of the complexity of the models and the simplifications required to make them analytically tractable [31], numerically feasible in simulations for large-scale resource planning [32], and realizable in renewable integration studies [33].

Effective and widely used strategies for optimizing the scheduling and operation of bulk-system resources have used markets to solve the cost-minimizing security-constrained resource allocation problem since they were proposed in the early 1980s [34]. Market-based control strategies were later adapted to building control systems [35], generalized for power balancing [36], applied to feeder-scale operations [24], then utility-scale operations [25], and most recently proposed for ancillary services [37, 38]. In addition, there is a rich literature describing models of varying complexity that have been used to study the control of aggregate loads in these cases [39, 40, 41]. The design of utility-based generation-following load control systems, either by direct command and control or by indirect price-based control, remains an active area of research.

The trend toward a more integrated and interconnected complex energy system is inexorable. Progress on the 21st century's infrastructure of complex interlocking energy resource, transformation, information, service, social, and economic networks is challenging our current understanding of these systems and our ability to design and control them. Transactive control was introduced to help address this transformation by enabling a more integrated system where all the costs of delivering energy to customers could be considered in real-time. An illustration of such a top-to-bottom restructuring of electricity delivery based on transactive control signals is shown in Figure 2.1. According to this vision of the future system, resource producers and consumers have equal access to the infrastructure provided, while operators determine the prices at which resources are efficiently allocated without violating physical limits, and aggregators group smaller resources together to balance the market power and

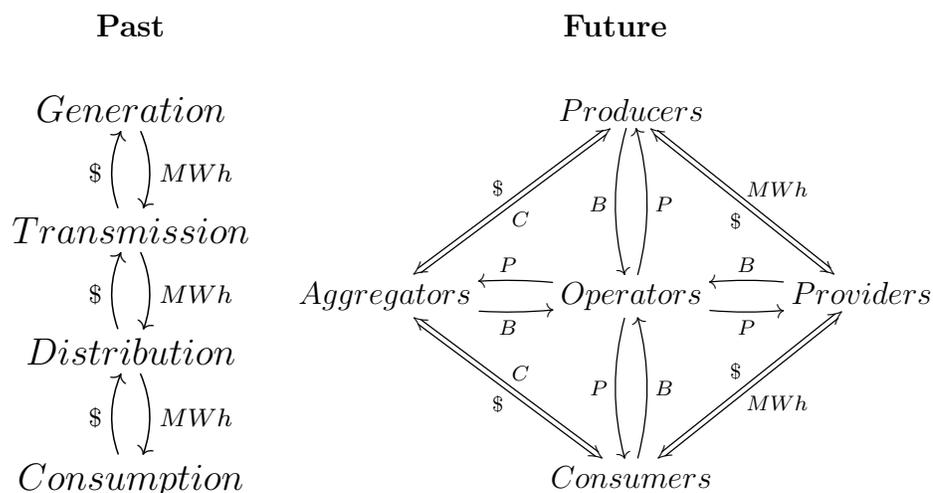


Figure 2.1: Top-to-bottom rethink of electricity infrastructure, including providers of transmission and distribution infrastructure, system operators and resource aggregators.

physical influence of larger resources.

Fuller defines Transactive Control as [42]

Utilizing a central control and distributed agent methodology [...] to act on behalf of consumers, sending information and automatically adjusting settings in response to a centralized signal.

To remain simple and general, this definition deliberately omits considerations of the temporal and physical hierarchies of power system operation. Neither does it specify any particular requirement to satisfy existing or anticipated challenges to the system. For example, while transactive control is widely believed to help address ramping problems, very little work has been done to show how it does so at the system level.

Recently, it seems no new work on renewable integration and demand response can fail to mention the California ISO forecast of the net load shape through the year 2020. The shape of the curve shown in Figure 2.2 has led to its colloquial name “the Duck Curve”. But this genial name does not properly convey the significance of the finding: a load ramp in the late afternoon of 13,000 MW over three hours is an operational challenge that should not be underestimated. In his report on the subject, Lazar proposes ten strategies to address this challenge [43].

Strategy 1: Target energy efficiency to the hours when load ramps up sharply;

Strategy 2: Acquire and deploy peak-oriented renewable resources;

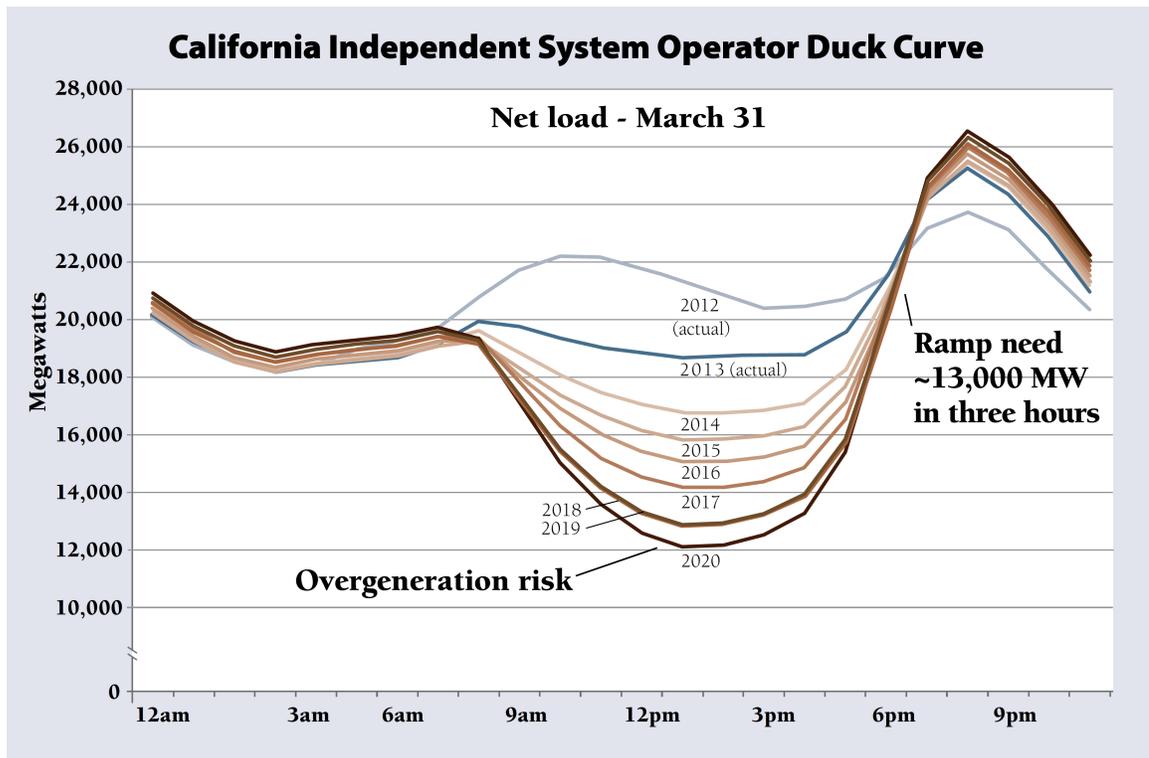


Figure 2.2: The California “Duck Curve” [Source: CAISO]

Strategy 3: Manage water and wastewater pumping loads;

Strategy 4: Control electric water heaters to reduce peak demand and increase load at strategic hours;

Strategy 5: Convert commercial air conditioning to ice storage or chilled-water storage;

Strategy 6: Focus utility prices on the “ramping hours” to enable price-induced changes in load;

Strategy 7: Deploy electrical energy storage in targeted locations;

Strategy 8: Implement aggressive demand-response programs;

Strategy 9: Use inter-regional power exchanges to take advantage of diversity in loads and resources; and

Strategy 10: Retire inflexible generating plants with high off-peak must-run requirements.

Among these, this thesis focuses primarily on the technical mechanisms that support Strategies 4, 6, 7, 8, and 9, all of which call for a more integrated approach to system planning and operation.

Significant challenges and research opportunities remain in load modeling and simulation, understanding of the impact of consumer behavior on demand response, the foundational theory for controlling widely dispersed demand response resources, and the verification, validation, monitoring and metering of demand response systems in utility operations.

Overall, it is clear that we are entering a period of increased electric utility receptiveness and growing innovation in the methods and strategies for turning a largely passive customer base into an active part of electric system operation. Although new customer-owned distributed generation and storage resources are expected to become increasingly significant, recruitment of existing controllable loads is still the most widely available resource base available to engage the customer in system control.

The impact of controllable load on system operation can be deduced from studies on the impact of variable generation. The studies to date suggest that the benefits of variable generation outweigh the costs for reasonable mixes of variable generation relative to conventional resources [16].

Many of the adverse impacts of variable generation are positive impacts for controllable load in the sense that the magnitude of the cost or impact as a function of generator variability is a cap on the magnitude of the benefit of load as a function of load controllability.

Controllable load exhibits the further advantage of high downward substitutability and thus can be significantly favored under liberalized ancillary service markets. This feature of controllable load suggests that well-designed ancillary service markets along with market-based load control strategies could be a very powerful combination.

Significant further research on how to structure such energy and ancillary service markets, design load control strategies, and model the systems in which they operate is required to further elucidate the benefits of this approach. Ultimately our ability to plan and operate bulk power systems that utilize such resources will depend on our ability to understand both the system as a whole as well as the details of the economic, electromechanical, and human components which comprise it.

The transactive system architecture can potentially allow for the aggregation and control of all the necessary resources, both supply and demand, at every level from transmission to end-use devices, as well as all the necessary capability, energy, capac-

ity, and ramping, at the necessary time-horizons from days-ahead to real-time. The comprehensive nature of the structure should alleviate concerns of present day system planners and operators regarding controllability of distributed smart grid assets, allowing them to be fully incorporated into system operations to achieve multiple objectives:

- Higher utilization of generation, transmission, and distribution assets, by changing on-peak load behavior;
- Lower wholesale market costs or power production costs, especially during high price periods;
- Lower ancillary service costs by engaging distributed assets to supply them;
- Lower cost for integrating new solar and wind generation into system operations by mitigating their variability and uncertainty;
- Higher environmental benefits from more efficient asset utilization and the potential to easily internalize environmental costs; and
- Increased reliability at both the bulk grid and distribution levels, from coordinating the engagement of distributed assets by multiple operating entities, by (1) increasing available reserve margins, (2) incorporating them into bulk grid wide-area control schemes, and (3) integrating them with distribution level voltage control and reconfiguration schemes.

The transactive architecture should allow increased penetration of demand response and other distributed assets, resulting from their significantly enhanced economic viability, by allowing them to provide a complete set of services on par with traditional large-scale transmission-level resources. This architecture also helps sustain utility revenue requirements, stabilizes utility customer costs at low rates made possible by lower cost distributed assets that displace the need for additional conventional infrastructure. Thus the vision of enabling overall cost effectiveness and environmentally sound grid infrastructure can be realized. While minimizing the information content of data transferred, it enhances overall cyber-security and customer privacy.

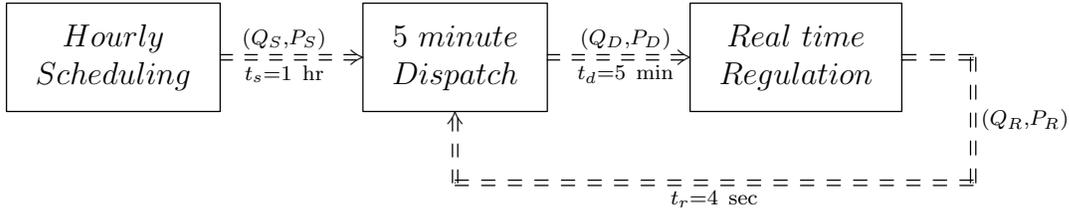


Figure 2.3: Inter-temporal data flow diagram.

2.2 Achieving Optimality in Transactive Systems

Several open questions remain when consider the optimal area control design problem in the presence of significant demand response resources that autonomously respond to frequency deviations caused by intermittent generation. Autonomous frequency control using responsive loads was proposed by Schweppe et al. [44] and demonstrated in the Olympic project, which showed its potential to mitigate generation loss. Autonomous load control can provide much faster response to frequency deviation than generation resources or dispatched load control can. However the aggregate control gain and economic elasticity of responsive loads vary over time because these loads are typically thermostatic (e.g., waterheaters, heat-pump compressors) that have both time-of-day and weather sensitivities. Thus it seems necessary to investigate how the standard ACE control or the previously considered optimal area control designs would operate in the presence of autonomous demand response.

The question of what constitutes optimality under transactive control is complicated by the lack of consensus in the definition of what is “transactive” control [45]. For this thesis we start from the definition proposed by Fuller because of its generality and simplicity. This definition does not specify any particular physical or temporal control architecture, leaving us free to choose what is most suitable for the problem at hand. We use the temporo-physical hierarchy defined in [46] as illustrated in Figure 2.3, which fits well with Fuller’s definition and provides a relatively simple data flow between physical and temporal scales. Using this approach the total generation and load is scheduled hourly such that, for each control area, a uniform price is obtained at which supply is equal to load plus net exports. This schedule is used to set each area’s price schedule P_S and net exports Q_S , which are in turn used by 5-minute dispatch markets [47] to allocate remaining resources in response to deviations from the hourly schedule. The allocation of Q_D additional exports at the price P_D is then used as a basis for the real-time price P_R and quantity Q_R .

Because we wish to consider the behavior of the system when demand response is active, we are motivated to find control strategies that maintain the maximum total economic surplus established by the schedule. Based on the transactive system design demonstrated by Hammerstrom et al. [48], Kiani and Annaswamy [49] proposed a hierarchical transactive control model for renewable integration that incorporates primary, secondary and tertiary frequency control that is consistent with the architecture in Figure 2.3. This model was successful in describing not only the primary regulation response of steam turbines to a loss of wind using a transactive model, but also the disequilibrium process of the secondary and tertiary responses. Because transactive control incorporates economic signals, Kiani’s model can be used to evaluate the impacts of transactive controls on total economic surplus both with and without their proposed tertiary control. Evaluating the surplus impacts provides a useful alternative to the typical optimization objective of minimizing frequency deviations, generator response, or regulation cost, especially in the context of transactive controls where the joint energy, power, and ramp responses have different time-varying cost functions and are considered over different time-horizons.

The transactive control dispatch system is used to solve two concurrent problems.

- (i) **Schedule tracking:** The hourly schedule is set by the unit commitment process [50]. From this process we obtain two important parameters that are used to dispatch retail resources at the 5 minute time-scale¹: (1) the schedule price and net exports, and (2) the participation factor for each generator and responsive load to match real-time demand and supply within a control area.

The schedule price P_s is determined from the hour-ahead supply and demand bids, and corresponds to the control area’s net export schedule Q_s , which becomes the control reference for dispatching units over the coming hour.

- (ii) **Resource dispatch:** Every five minutes generation and load resources are re-dispatched and the regulation control is reset to establish the basis for the control of system frequency and area exports over the next 5-minute interval. Units with non-zero participation factors bid into the dispatch market to allow the schedule to be adjusted so that recent resource state changes can be considered. Contribution and participation factors are computed and used (1) to reset the power output for generation units, (2) to reset the state of demand

¹The 5-minute dispatch interval is chosen because it was used in the Columbus and Olympic systems and allows for easier comparison of simulations with data obtained from operations.

response, and (3) set the frequency regulation gain for both generation and demand response.

With both the scheduling and dispatch strategies available we have all the necessary elements required to consider a regulation response strategy that minimizes schedule deviation and tracks a surplus maximizing schedule with adjustments from the last 5-minute redispatch operation. The implementation of solutions to the scheduling surplus maximization problem in the interconnection scale is addressed in [1]. According to this formulation the goal is to minimize the initial over-production of power which reduces later under-production and allows the system to track its schedule more cost-effectively.

We therefore need to address one specific aspect of the larger transactive control design problem, namely the integration of the 5-minute dispatch control with the automatic control mechanism that regulates system frequency in the presence of demand resources that are frequency sensitive [38]. To this end, this thesis proposes among other things an approach to regulating frequency and area exports, and minimizes the loss of economic surplus resulting from deviations from the hour-ahead schedule.

The approach proposed in this thesis addresses four important elements of this solution: (1) modeling aggregated demand response behavior in transactive systems, (2) controlling aggregated demand response, (3) optimal area control using demand response as a regulation resource, and (4) optimal dispatch of demand resources at the control area level.

Chapter 3

Demand Response

Historically, demand response programs have taken the form of so-called “demand side management” (DSM) activities. DSM seeks to alter electricity demand load shapes to make them better match the available supply and reduce load peaks so as to defer costly capacity expansion investments. Traditional DSM programs include increased building and appliance efficiency standards, as well as equipment replacement/retrofit programs.

In this chapter a derivation of a short-term demand response model suitable for transactive control systems is presented, followed by its validation with field demonstration data. A discussion of what the newly gained understanding of short term demand response might mean in terms of technology development, consumer acceptance, regulatory policy, and research opportunities is presented in Section 7.1 .

The first contribution in this chapter is the development of an analytic function for short-term demand in residential thermostatic loads that are responsive to real-time prices. The development of the demand function reflects first principles regarding the nature of thermostatic load control. We show that this model reduces to the Random Utility Model (RUM) employed in economics to study consumer choices and the valuation of non-market goods [51].

The second contribution in this chapter is a validation of the model against data obtained from the Olympic and Columbus field demonstrations. These demonstration projects implemented residential level thermostatic inputs to a price-based market clearing transactive control system on a five minute time-scale. The results of the field demonstrations show that customers could exhibit positive short-term demand response to short-term price variations. We show that the demand model can be easily calibrated to give an accurate representation of the market data from these

experiments.

Finally, the model is compared to four alternative models of demand response (DR): no DR, half DR, full DR, and demand elasticity from Faruqui’s 2010 survey of DR programs [52]. For all models, we compute the error in predicting the amount of load shed at the 5-minute real-time price produced by the double-auctions of the Olympic and Columbus experiments. The results show that the RUM outperforms the alternative models for common “steady-state” demand conditions. In more extreme situations where load state diversity is low or when large price deviations occur over a very short time frame, the performance of the RUM is comparable to that of the competing models.

3.1 Background

Load shifting has long been recognized as a second approach to modulate demand response. Whereas traditional energy efficiency programs aim to reduce overall consumption, load shifting focuses specifically on changing the time of day when energy is used in order to favor times when costs are lower. Programs that focus on load shifting typically require mechanisms such as time-of-use (TOU) pricing or real-time pricing to induce transient changes in consumer behavior, such as those described by Vardakas [53]. TOU and seasonal rates focus on the customer’s response to simple static price signals [54]. The Electric Power Research Institute (EPRI) carried out a major study of the top five experiments in the United States in the early 1980s and concluded that consumers indeed responded to higher prices by shifting some of their load to off-peak periods [55]. Later experiments produced similar results. The City of Anaheim Public Utilities conducted a residential dynamic pricing experiment and found that for a peak-time rebate of \$0.35/kWh they could reduce electricity use by 12% during critical-peak days [56]. California’s Advanced Demand Response System pilot program used a critical peak pricing (CPP) tariff using the GoodWatts system to obtain peak reductions as high as 51% on event days with a CPP rate and 32% on non-event days with TOU rate. Enabling technology was identified as an important driver for load reductions [52]. This observation was also made in the Olympic Peninsula Project, where both TOU and real-time price (RTP) tariff were tested [57]. Similar results over a large number of studies have been widely reported and are summarized in a survey published by Faruqui et al. [52].

Since the introduction of *homeostatic utility control* by Schweppe et al. [44], it

has been understood that key system state variables such as frequency and voltage in large-scale interconnections could be regulated using price signals. Prices have since been used primarily to schedule and dispatch generation resources using power markets [50]. Both energy efficiency programs and time-of-use rates have consistently been shown to effectively reduce loads on time-scales greater than one hour [58].

To avoid unfair pricing in the presence of demand response, David et al. [59] and later Kirschen et al. [60] examined how the elasticity of demand could be considered in wholesale scheduling systems. Initial work applying market-based mechanisms to building control systems showed that the notion of market-based demand control was feasible and effective for more granular systems [35]. The general concept of transactive control was initially proposed [61, 62] as a method of coordinating very large numbers of small resources using market-like signals at the electricity distribution level. The theory is essentially the same as for wholesale markets. However, realizations can be quite different insofar as more frequently updated price signals are typically used to manage distribution system constraints such as feeder capacity limits. These prices can dispatch both distributed generation, energy storage and demand response resources at much higher temporal and physical granularity than is possible with wholesale markets.

A number of previous studies have considered the operational impact of using retail price signals for controlling load in electric power systems. Glavitsch et al [63] showed that nodal pricing could find a socially optimal operating point for power markets. Following up on this work Alvarado [64] considered the question of whether power systems could be controlled entirely using prices, and found that price signals could indeed work. But the results came with some caveats, the most significant of which is the question of stability of the feedback control over the entire system.

The feasibility of transactive control methods was demonstrated in the Olympic and Columbus projects using distribution resource control systems that dispatched distribution-level resources in quasi real-time using price signals. These experiments yielded a trove of high-resolution data about the behavior of load resources in response to short-term price variations.

Overall, two important lessons have been learned from decades of utility research, development, and field experimentation with demand response [53]:

1. Consumer interest and sustained participation is essential to the success of demand response programs. Too many programs showed too little consumer

interest and participation. This drives up program costs and reduces effectiveness. Tools to keep customers engaged and responsive to utility priorities are needed. Substantive contract diversity and meaningful incentives need to be available for customers to choose and actively engage in programs.

2. Programs should not provide rewards and incentives on the basis of complex baseline or reference models. Mechanisms that provide or enable endogenous sources of counterfactual prices and quantities should be preferred by utilities.

Although transactive systems are similar to wholesale markets, the price signals are applied to different resources, affect consumer needs differently and are applied at much higher temporal and physical granularity than is possible in wholesale markets. Short-term consumer response to price variations is also understood to be quite distinct from long-term demand response. Long-term demand response is typically associated with changes in consumer behavior and the conversion to more energy efficient houses and appliances.

On the other hand, short-term demand response is primarily in the form of time-shifting and often requires automation. Short term demand response can be very different from long-term demand response because controllable load resources can be quickly exhausted, leading to control saturation. As a result, short and long term consumer responses are not generally comparable. In practice, long-term demand response models tend to underestimate the magnitude of the controllable resources and overestimate their endurance [24, 25]. This has two important consequences: (1) Planning studies tend to undervalue the potential contribution of short-term demand response system which is often overlooked in utility program development; and (2) when it is not overlooked, the open-loop control gain is underestimated, resulting in over-control, instability and excessive price volatility.

The lack of solid theoretical basis for short-term performance claims has emerged as a significant challenge [65, 28]. Using static long-term own-price elasticities can be expected to give rise to erroneous short-term demand response control because short-term elasticities are more often substitution elasticities where the substitute is obtained in time rather than by an alternative product, a distinction which was made evident by Fan's study of Australian price elasticities [66], among others. Own-price elasticities represent averages over long periods of time. These averages may fail to capture the magnitude and variability possible at any given time. For example, Reiss and White [67] developed a household electricity demand model for assessing

the effects of rate structure change in California and found that a small fraction of households respond to the price changes with elasticities as large as -2 , which far exceeds the average long-term elasticity of -0.14 found in Faruqui's survey of DR programs [58]. Unfortunately, computing the elasticity of demand for short-term demand response to real-time prices has proved challenging because the counterfactual price and demand are difficult to determine in the absence of a short-term feedback signal that elucidates the loads' willingness to pay [68]. Thus, more accurate models of short-term demand response are necessary for utility load control planning and design if these systems are to be deployed effectively.

3.2 Random Utility Model

Short-term consumer response to price variations is generally regarded as quite distinct from long-term demand response. The primary difference stems from the fact that long-term demand response is typically associated with enduring changes in consumer habits, whereas short-term demand response usually requires automation to support temporary changes in device behavior when prices are high. As a result, the consumer's considerations when, whether, and how to respond are not generally comparable, nor are they necessarily mutually exclusive.

Unfortunately, most studies of demand response in the electricity sector have focused on the static long-term elasticity of consumer demand [58]. Lacking alternative sources for short-term demand elasticity measures, utilities tend to use existing long-term elasticities as the basis for load control program evaluation and control systems design. Two important consequences arise from any discrepancy between the two elasticities.

1. Over- or underestimation of the program value. If the short-term elasticity is greater than the long-term elasticity, then an indirect load control program would tend to be under-valued and would be less likely adopted.
2. Over- or underestimation of control gain. If the short-term elasticity is greater than the long-term elasticity, then any attempt to mitigate instability from feedback signals would likely underestimate the open-loop control gain and would result in incorrect design of the closed-loop control system. This can potentially lead to less stable system operations and higher price volatility.

These discrepancies, and the results of the Olympic and Columbus studies showing that customers could exhibit positive short-term demand response, are the primary motivations for developing a new model of short term electricity demand.¹

3.2.1 First-principles Model

We model the general behavior of thermostats governed by consumer preference based on an engineering model of houses' responses in the short-term given a consumer's static setting for comfort. The average duty cycle is based on the fraction of time the system is *on* relative to the total cycle time. For a thermostat operating within its deadband, the fraction is very closely approximated by

$$\rho = \frac{\tau_{on}}{\tau_{on} + \tau_{off}} \quad (3.1a)$$

where

$$\tau_{on} = \frac{1}{r} \ln \frac{\tau_{set} - \frac{1}{2}\tau_{hys} - \tau_{on}}{\tau_{set} + \frac{1}{2}\tau_{hys} - \tau_{on}} \quad (3.1b)$$

and

$$\tau_{off} = \frac{1}{r} \ln \frac{\tau_{set} + \frac{1}{2}\tau_{hys} - \tau_{off}}{\tau_{set} - \frac{1}{2}\tau_{hys} - \tau_{off}}. \quad (3.1c)$$

In these equations, r is the indoor air temperature decay rate time constant, τ_{set} is the temperature set-point, τ_{hys} is the thermostat's hysteresis, τ_{on} is the steady-state temperature when the heating system is *on* (or cooling is *off*), and τ_{off} is the steady-state temperature when the heating system is *off* (or cooling is *on*). This duty cycle corresponds to the probability that we observe a device to be *on* at any given time.

The transactive control system assumes thermostats submit bids such that the probability of clearing a lower retail electricity price is the duty cycle required to maintain consumer comfort. Thus the system can be expected to run with the duty cycle needed while preferentially running when retail electricity prices are lower. This

¹It is worth noting that the effect of short-term demand response does not necessarily result in lower energy demand. It often results in lower peak load at the expense of increased total energy use. The reason is that consumers tend to respond to variations in price about a mean or expected price, increasing demand when prices are low and decreasing it when prices are high. Comfort is achieved by "storing" thermal resources (i.e., heat or cool) during low price periods and releasing it during high price period. Because the store/release process is not expected to be 100% efficient [69] any price-induced reduction in peak is typically associated with an increase in total energy use.

comfort tracking cost minimizing strategy is embodied in the bid-response function

$$B = \bar{P} + K \frac{\tilde{P}}{\tau_{obs} - \tau_{set}} \quad (3.2)$$

where \bar{P} is the expectation value for the clearing price, \tilde{P} is the standard deviation, τ_{obs} is the actual indoor air temperature, and K is the customer's comfort control setting. The consumer-controlled variable K expresses how sensitive the household is to the trade-off between money (the cost of energy) and comfort (distance from ideal temperature). A high value of K signals a higher sensitivity to price fluctuations. It gives the customer more opportunities to reduce costs at the expense of reduced short-term comfort, embodied by the thermostat's tracking error $\tau_{obs} - \tau_{set}$.

For a population of N thermostats with mean duty cycle $\bar{\rho}$, the probability of finding k devices in the *on* state is proportional to the binomial distribution of the count

$$g(k) = \frac{N!}{k!(N-k)!} \rho^k (1-\rho)^{N-k}. \quad (3.3)$$

The design of thermostatic loads requires significant oversizing of equipment, so we assume that on peak the mean duty cycle $\bar{\rho} = 0.5$. Given the mean thermostatic device load \bar{q} we can define the load deviation $Q = (k - \frac{1}{2}N)\bar{q}$ from the most probable load \bar{Q} and the total thermostatic load $Q_R = N\bar{q}$. The aggregate load entropy is then given by

$$\lim_{\rho \rightarrow 0.5} \sigma(Q) = \sigma_0 - 2 \frac{Q^2}{Q_R} \quad (3.4)$$

where

$$\sigma_0 = \frac{Q_R}{\bar{q}} \ln \frac{Q_R}{2\bar{q}} - \frac{1}{2} \ln 2\pi$$

is the maximum entropy corresponding to the most probable load $\bar{Q} = \bar{\rho}Q_R$. The probability of observing any given load Q is the probability

$$2^{-N} g(k) = \frac{1}{1 + e^{-2\sigma(Q)}}. \quad (3.5)$$

Using the standard definition of demand elasticity for a load Q at price P we obtain

$$\eta = \frac{P}{Q} \frac{\partial Q}{\partial P} = P \frac{\partial \ln Q}{\partial P} = P \frac{\partial \sigma(Q)}{\partial P}. \quad (3.6)$$

We observe that the maximum demand responsiveness $\frac{\partial Q}{\partial P}$ occurs when the entropy

is at its maximum. Therefore at the most probable price and load it must be that

$$\left. \frac{\partial^2 Q}{\partial P^2} \right|_{\substack{P=\bar{P} \\ Q=\bar{Q}}} = \left. \frac{\partial \sigma(Q)}{\partial Q} \right|_{Q=\bar{Q}} = 0.$$

Integrating twice we find

$$\sigma(\bar{Q}) = a + b\bar{P}.$$

where a and b are unknown constants which must be determined from boundary conditions or from a fit to data. We substitute this result into Eq. (3.5) and deduce that the load as a function of price is

$$Q(P) = \frac{Q_R}{1 + e^{a+bP}} + Q_U \quad (3.7)$$

where Q_U is the unresponsive load not subject to price-responsive behavior.

3.2.2 Model Assumptions

This model has the same form as McFadden's random utility model [51]. The RUM has been used extensively in economics to study consumer choice and in the valuation of non-market goods [70]. That it can be derived independently from the first principles of thermostatic controls establishes a strong tie between the engineering approach and state of the art economic modeling of consumer preferences.²

The random utility model (and thus subsequent manipulations below) makes two important assumptions.

1. A consumer's choice is a discrete event in the sense that a device acting on the consumer's behalf must make an all-or-nothing decision. The consumer can either run or not run an air-conditioner. The device cannot be run at part-load for the next interval.
2. The consumer's (or device's) attraction to a particular choice is affected by a random error with a type 1 extreme value (Gumbel) distribution. In this case we use the term *attraction* in the retailing sense but we could just as well use the term *utility* to be consistent with economic theory. The randomness of the utility's observation of the current comfort preference is assumed to arise from

²McFadden received the 2000 Nobel prize in economics for his pioneering work on economic choices.

the devices acting on behalf of consumers. The devices rationally choose the outcomes with the highest utility based on the consumer's indicated preference for comfort.

Although the random utility model has been derived using various methods, McFadden points out that according to Luce [71] it is axiomatic that the relative odds in a binary choice will remain the same for independent alternatives when additional alternatives become available. Therefore, the selection probability can always be written in the form

$$P_i \frac{e^{\nu(z_i)}}{\sum_{n=1}^N e^{\nu(z_n)}}$$

where $\nu(z)$ are scale functions of the stimulus z . When ν is linear in parameters, this result is the multinomial logit formula found in the standard statistical literature.

In the absence of prior knowledge of the quantities demanded by consumers, the derivation of the aggregate demand curve is based on this discrete choice statistic for consumer demand [72]. Thermostats act as agents on behalf of consumers whose utility is assumed to be composed of an observable component based on their comfort setting, which follows an extreme value distribution and an unobservable component that has zero mean from the perspective of the market. Thermostats are expected to bid a price that will give the heating/cooling system a probability of running satisfying the duty cycle required to maintain indoor air temperature. A thermostat's response to a market clearing price is an exclusive choice based solely on the bid it submitted [73]. For a dichotomous choice, the reasoning is as follows: U is the consumer benefit (*utility* in economic theory) that the thermostat obtains from taking a particular action given the consumer's preferences. This net benefit depends on an unobservable characteristic α that has a zero mean distribution and an observable characteristic β that is a known decreasing function of price. The net benefit is defined as $U = \beta x + \alpha$ where x is the binary choice ($x = 1$ or 0) β represent the marginal utility of the anticipated change in comfort and cost of electricity, and α represents all other unobservable variables that can give rise to error in choices. The action corresponding to $x = 1$ is taken if $U > 0$.

From a logistic regression we find that the probability of taking the action is then

$$\rho\{x\} = \frac{1}{1 + e^{-\beta x}}. \quad (3.8)$$

The optimal consumer bid from Eq. (3.2) is the utility maximizing price. The random

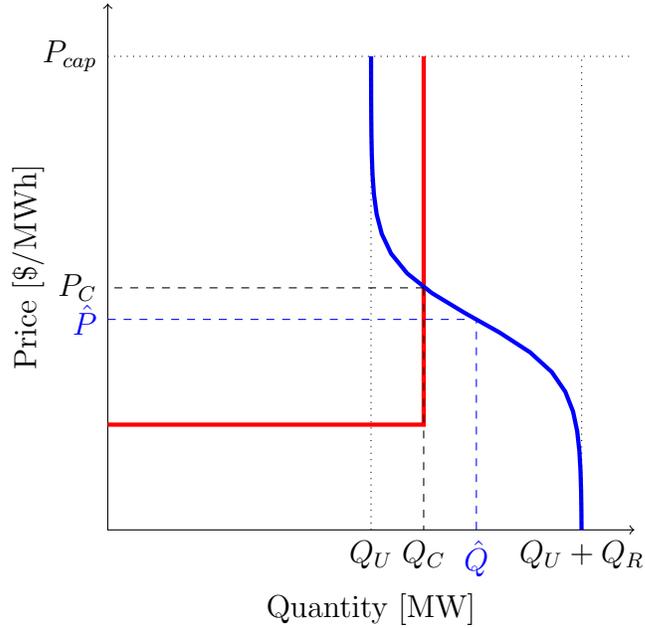


Figure 3.1: Demand curves for steady state thermostatic loads in a transactive control system.

utility model gives the same result as the load probability in Eq. (3.5).

3.2.3 Equilibrium Demand Response

The transactive control system used in the demonstration projects is quiescent when load state diversity is maximized and the total load is steady. This steady-state condition occurs when the distribution of bids is symmetric about the mean price with the same relative variance. We assume that states are uniformly distributed over the thermostat deadbands because without price disturbances, thermostatic devices follow the standard duty cycle regime. Devices will be *on* for the time required to rise from the *on* boundary of the thermostat deadband to the *off* boundary and *off* for the time required to go the other way. Undiversified states will tend to randomize under the influence of diverse physical parameters and state diversity grows until the thermostats settle into the equilibrium demand regime. After re-diversification, prices return to obeying the logistic distribution of Equation (3.8). We rescale the physical quantities for an arbitrary system with Q_U unresponsive load and Q_R responsive load at the prices p . Finally we rewrite Equation (3.8) to obtain Equation (3.7) again, where a and b are the demand curve's shape parameters.

The most probable demand elasticity at steady state occurs at maximum diversity

when $p = -\frac{a}{b}$ and

$$\hat{\eta}_D = \eta_D \left(-\frac{a}{b} \right) = \frac{a}{2} \quad (3.9)$$

as shown in Figure 3.1. Using this we can estimate the demand function parameters for any set of N bids by fitting a linear function to the bids within the central 60% of the demand response range, i.e., from $0.2Q_R$ to $0.8Q_R$. The most probable price \hat{P} is found at the mid-point $\hat{Q} = Q_U + \frac{1}{2}Q_R$. The demand elasticity $\eta_D = 2\hat{P}/Q_R d$ where d is the demand response slope obtained from the linear fit. From this we find the curve parameters

$$a = 2\eta_D \quad \text{and} \quad b = -a/\hat{P} \quad (3.10)$$

which can be obtained by applying the definition of elasticity to Eq. (3.7) such that $\eta_D = \frac{p}{Q(p)} \frac{dQ(p)}{dp} \Big|_{p=\hat{P}}$ and observing that the most probable price occurs when $p = -a/b$.

3.3 Model Validation

To test the validity of the Random Utility Model, we compare its performance to that of models currently employed by utilities. The four comparator models are based on three different levels of static participation from the installed demand response capacity (i.e., 0%, 50% and 100% of $Q_R - Q_U$), as well as the -0.14 long-term demand elasticity proposed by Faruqui [58]. The key difference between the different models used by utilities is the assumption made about the short term elasticity of demand:

1. **Zero elasticity (No DR)** may be preferred when the demand response program is not expected to be a significant fraction of the total load. However, in such a case the open-loop control gain is essentially zero, and both program valuation and feedback control design are not possible. Depending on the shape of the demand curve, zero elasticity may arise at more than one quantity.
2. **Maximum elasticity (Full DR)** may be preferred when the purpose of the study is to design the feedback control so as to avoid instability. The maximum elasticity would correspond to the maximum open-loop gain and the choice of closed-loop gain would therefore be such that instability could be avoided for all physically realizable load conditions.
3. **Most probable elasticity (Half DR)** may be preferred when the purpose of the study is to evaluate a load control program's long-term value. Depending

on the shape of the demand curve, the most probable elasticity may also be the maximum elasticity.

We supplement these standard models by adding a fourth comparator:

4. **Elasticity $\eta = -0.14$.** This last model imposes the long term price elasticity estimated by Faruqui [58].

It can be difficult to augment these static elasticity models with what is known about consumer choices in energy purchasing. In a study by Goett and Hudson published in 2000 [74], customers were found to consider a long-term marginal energy price increase more onerous when the price is low than when it is high. However, energy consumers do not go so far as to consider price changes strictly in proportional terms either. In addition, the authors noted that bonus or coupon inducements can affect consumer choices. But these inducements are only relevant to long-term decisions such as tariff or supplier choice. Other variables often considered by consumers include contract duration (where long terms are viewed negatively), variable rates (also viewed negatively) with shorter-term fluctuations being viewed more negatively.

The demand function model proposed in Section 3.2 was validated using data obtained from the Olympic and Columbus projects. Bids received in both projects were fit using the method described above. An example of this fitting process is shown in Figure 3.2. The Olympic data set includes 103,842 market clearing events from April 1, 2006 to March 31, 2007. A total of 1,174,923 bids were received from 38 customers. Real-time prices were computed from 5-minute double auctions on a single feeder. The Columbus data set includes four separate feeders with over 5-minute market clearing events from June 1, 2013 to September 30, 2013. The data for both studies is summarized in Table 3.1.

The performance of the RUM and four alternative models are evaluated by comparing the quantity predicted by each model at the observed market price to the actual quantity observed. The -0.14 long-term elasticity is unlikely to be a good approximation for fast-acting demand response, but it provides a clear indication of the errors potentially introduced when using long term elasticities in studies of short term demand response.

The results of the Olympic experiment are presented in Table 3.2. We observe that the random utility model outperforms the alternative models for all performance metrics. The difference is particularly significant for the mean error and bias error,

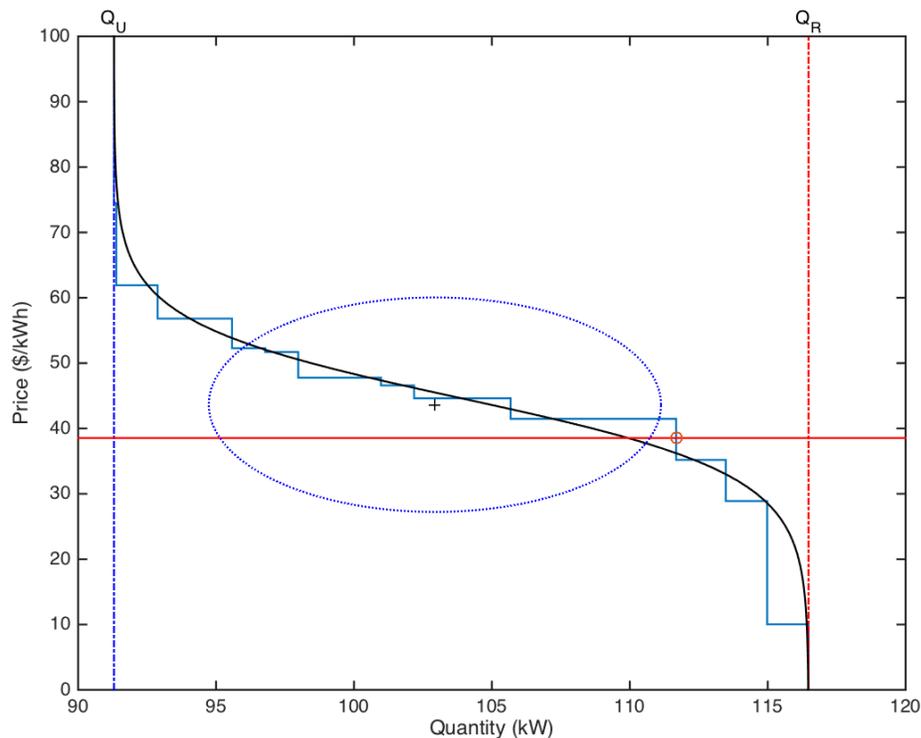


Figure 3.2: Example of demand function model validation with Columbus demonstration data. The bids shown are from 2013-06-22 22:45 EDT. The clearing price P_C and quantity Q_C are indicated by the circle. The expected price \bar{P} and quantity \bar{Q} are indicated by the plus sign. The standard deviation of price \tilde{P} and quantity \tilde{Q} are indicated by the ellipse.

Table 3.1: Feeder characteristics

Feeder Id	Customers	Demand Response		
		Bids	kW	(% peak load)
<i>Olympic (heating)</i>				
–	38	1,174,923	1,193	(52.4)
<i>Columbus (cooling)</i>				
120	11	281,045	20	(0.7)
140	30	699,241	45	(1.0)
160	53	1,478,148	83	(1.2)
180	8	210,241	12	(0.2)

but less significant for the standard deviation. Results for the full Columbus data sets are presented in Table 3.3. They are mixed. We note that the Full DR model outperforms the RUM for mean error. On the other hand, the RUM outperforms the Full DR model for bias error and the standard deviation on Feeder 160 (the largest

Table 3.2: Olympic data analysis results

Model	Error		Standard
	Mean	Bias	Deviation
Random Utility	-0.35%	1.74%	7.24%
Half DR	-0.67%	5.63%	10.10%
Full DR	-7.38%	6.83%	14.14%
No DR	6.04%	5.49%	8.89%
$\eta = -0.14$	6.62%	6.83%	14.14%

and most diverse in terms of number of consumers). The picture that emerges out of Columbus is that the RUM model performs best in some instances but is comparable to the static models overall.

There is, however, an important qualifier to the Columbus results. Saturation of demand response (either all *on* or all *off*) is associated with diminishing load state diversity. This violates the assumption of the RUM. For this reason, we expect feeders that frequently saturate the demand response resource control to not be well represented by the RUM. These feeders would produce data more often consistent with the No DR or Full DR models (depending on conditions). Unlike the Olympic study, the experimental protocol for the Columbus study deliberately probed these limits of control every other day. This resulted in frequent loss of load state diversity in violation of the zero-mean assumption.

The experimental protocol in the Columbus demonstration is expected to have introduced additional errors. This has been verified to first order by analyzing the data excluding the experiment days. The results are shown in Table 3.4, where the error on Feeder 160 is reduced from 1.00% to 0.67%. However, a second-order effect is now observed insofar as the Full DR model seems to still perform better than the random utility model with the error reduced from 0.81% to 0.34%. This can be explained by the second-day recovery during which thermostats receive relatively lower prices compared to the previous day and tend to respond more aggressively to them. This hypothesis cannot be directly verified as the experimental protocol did not normally include a third day during which neither an experiment nor a recovery was taking place.

The model as presented is valid only for steady-state conditions. We therefore ought to consider whether transient demand response behavior influences the accuracy of the random utility model. Two factors are known to influence the magnitude of the demand response transient: (1) the fraction of devices that respond to the price

Table 3.3: Columbus analysis results for Feeders 120, 140, 160 and 180

<i>Feeder 120</i> Model	Error		Standard
	Mean	Bias	Deviation
Random Utility	0.78%	0.24%	0.31%
Half DR	0.75%	0.25%	0.33%
Full DR	0.53%	0.24%	0.30%
No DR	0.98%	0.29%	0.37%
$\eta = -0.14$	14.53%	0.24%	0.30%

<i>Feeder 140</i> Model	Error		Standard
	Mean	Bias	Deviation
Random Utility	0.97%	0.29%	0.36%
Half DR	0.99%	0.33%	0.41%
Full DR	0.71%	0.29%	0.37%
No DR	1.27%	0.38%	0.48%
$\eta = -0.14$	14.71%	0.29%	0.37%

<i>Feeder 160</i> Model	Error		Standard
	Mean	Bias	Deviation
Random Utility	1.00%	0.28%	0.36%
Half DR	1.00%	0.33%	0.41%
Full DR	0.81%	0.31%	0.39%
No DR	1.19%	0.36%	0.45%
$\eta = -0.14$	14.81%	0.31%	0.39%

<i>Feeder 180</i> Model	Error		Standard
	Mean	Bias	Deviation
Random Utility	0.61%	0.22%	0.26%
Half DR	0.44%	0.18%	0.24%
Full DR	0.34%	0.15%	0.20%
No DR	0.53%	0.22%	0.28%
$\eta = -0.14$	14.34%	0.15%	0.20%

signal, and (2) the diversity of device states when a price signal is received.

The results suggest that although the random utility model is valid for predicting steady-state demand response behavior, its accuracy is limited in the case of large magnitude price fluctuations that tend to drive a significant majority of responding devices to a single common state. This kind of state degeneracy violates the parameter distribution assumptions of the random utility model and reduces its accuracy for

Table 3.4: Columbus analysis results for only non-experiment days

<i>Feeder 120</i> Model	Error		Standard
	Mean	Bias	Deviation
Random Utility	0.78%	0.25%	0.31%
Half DR	0.74%	0.26%	0.33%
Full DR	0.52%	0.24%	0.30%
No DR	0.97%	0.29%	0.37%
$\eta = -0.14$	14.52%	0.24%	0.30%

<i>Feeder 140</i> Model	Error		Standard
	Mean	Bias	Deviation
Random Utility	0.79%	0.21%	0.26%
Half DR	0.77%	0.21%	0.28%
Full DR	0.52%	0.20%	0.26%
No DR	1.02%	0.24%	0.31%
$\eta = -0.14$	14.52%	0.20%	0.26%

<i>Feeder 160</i> Model	Error		Standard
	Mean	Bias	Deviation
Random Utility	0.67%	0.24%	0.30%
Half DR	0.59%	0.30%	0.38%
Full DR	0.34%	0.28%	0.37%
No DR	0.84%	0.33%	0.42%
$\eta = -0.14$	14.34%	0.28%	0.37%

<i>Feeder 180</i> Model	Error		Standard
	Mean	Bias	Deviation
Random Utility	0.69%	0.19%	0.22%
Half DR	0.58%	0.18%	0.24%
Full DR	0.45%	0.15%	0.21%
No DR	0.70%	0.21%	0.27%
$\eta = -0.14$	14.45%	0.15%	0.21%

predicting the load after an abrupt large-magnitude change in price is observed.

In general, we expect normal utility operations to be more like the Olympic conditions than the Columbus conditions. We thus conclude that while the Columbus results only weakly support the random utility model, they mainly point out the importance of the steady-state assumptions in Section 3.2.

3.4 Summary of Results

We have developed a logistic demand curve for short term electricity consumption derived from the first principles of controllable thermostatic electric loads operating under the transactive control paradigm. We have shown that this model corresponds to the Random Utility Model commonly used in the economics of consumer choice. The model's performance is compared to results from two US Department of Energy demonstration projects in which short-term demand response data was obtained. We find that the random utility model predicts the total demand response to smaller price fluctuations very well, but that model performance degrades as the magnitude and frequency of price excursions increases and as the diversity of load states decreases. We conclude that the random utility model is suitable for demand response studies that utilize steady state conditions for most situations with only infrequent and modest price excursions.

In its present form the random utility model provides a robust framework that is well-founded in the engineering principles of how thermostatic devices behave in price-based control environments. By joining the engineering and economic behavior of such devices, the random utility model is set to become an essential element in the planning, design and eventual deployment of large-scale load control strategies.

Chapter 4

Aggregation

This chapter proposes an aggregate load model of thermal loads controlled by thermostats that have no deadband and synchronize their control update interval to transactive market. This mechanism, which is denoted $T\Delta_0$, is used to address standard control design questions regarding the aggregate control of $T\delta_0$ thermostats in utility-scale demand response systems. Section 4.2 presents a new state-based dynamic model of the aggregate load control problem for these new types of thermostats. Section 4.3 examines the aggregate load controller design problem more generally and considers the performance various conventional load control designs using $T\delta_0$ thermostats. Section 4.4 tests the aggregate load model and controller designs using a large-scale agent-based simulation of $T\delta_0$ demand response. Further discussion of the general findings based on the simulation and suggestions for possible variations in the aggregate load controller designs are deferred to Section 7.2.

4.1 Background

Significant changes in generation mix must occur to meet growing load and mitigate the climate-change impacts of fossil-based electricity generation. Demand response control has the potential to displace some and possibly all generation resources used for regulation and contingency reserves. However, the current standard practice for both direct and indirect control of thermostatic load relies primarily on so-called “one-shot” load shedding strategies for emergency peak load relief only. This approach uses a controllable subset of all thermostatic loads in a particular class, e.g., water heaters or air-conditioners, which are transitioned to a curtailed regime that reduces the

population average power demand. After a time, these responsive loads are released and return to their normal operating regimes.

This strategy is known to exhibit fluctuations in aggregate load during the initial response as well as demand recovery rebounds after the loads are released. To mitigate this behavior, “one-shot” direct load control strategies are sometimes enhanced by either centralized load diversification mechanisms, such as using multiple subgroups of the responsive loads dispatched in a sequence that smooths the overall response of the load control system [75], or distributed mechanisms, such as using stochastic control strategies [76]. Many of these mechanisms require some knowledge of the aggregate thermal response of the buildings in which the loads are operating [77]. To solve the more general tracking problem where load “follows” intermittent generation [78] these mechanisms must address response saturation and loss of diversity [79], high sensitivity to modeling errors and noise [80], and stability considerations due to feedback delays [81, 82]. In particular proportional control [83] and integral control [77] strategies have been proposed to overcome many of the problems identified.

Aggregating building thermal loads are known to provide a potentially significant resource for balancing purposes [34] and have been used as the primary resource for many demand response strategies, including those that seek to use real-time prices to continuously regulate loads based on their bids, as in so-called transactive control systems. When field demonstrations of transactive control using real-time prices were conducted [24], the results sometimes revealed significant tracking errors in the discrete-time response of the aggregate load control. The cause of the error was found to be bidding strategies that didn’t or couldn’t account for the thermostat hysteresis. Compensated bidding strategies developed to address these problems did not fully mitigate these tracking problems [25], in part because of the complexities accounting for hysteresis. The hysteresis of standard thermostats not only requires a switched-mode representation of the individual building thermal response, but also requires so-called “refractory states”, meaning that states are locked in for a certain time after being entered [84]. These locked states are associated with transition delays rather than thermal parameters. Tractable state space models of aggregate loads can be obtained using model-order reduction strategies that linearize the system model and limit the number of state variables required to represent responsive loads [39, 40], as illustrated in Figure 4.1. The rate at which devices turn on and off is determined by (1) the rate $\dot{\tau} = r_{on}$ and $\dot{\tau} = r_{off}$ at which they respectively cross the hysteresis band limits $\Delta\tau + \delta$ and $\Delta\tau - \delta$, as well as (2) the rate at which the lockout times t_{min}

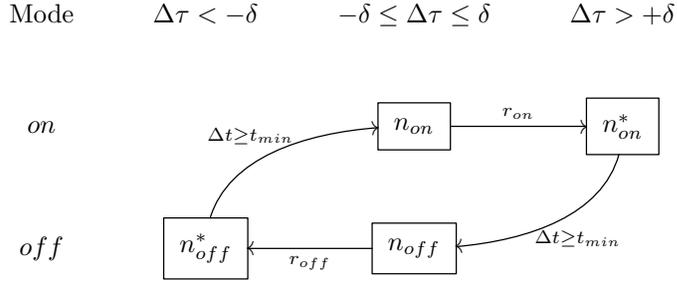


Figure 4.1: State-space model of aggregate conventional thermostatic loads in heating regime with refractory states n_{on}^* and n_{off}^* . $\Delta\tau$ is the difference between the indoor and outdoor temperatures and δ is the hysteresis band limit.

expire. Such state-space models minimally represent any thermostat with non-zero deadband. However they also require model parameter identification to be used in formulating bidding strategies.

An alternative thermostatic controller design strategy was proposed to overcome these modeling issues while not compromising the advantages of hysteresis control of thermal loads [85]. This thermostat design uses a discrete-time zero-deadband ($T\delta_0$) concept that has no refractory states and synchronizes the state transition times with external signals such as those coming from real-time retail double-auctions. The new thermostat provides significant fast-acting DR resources and the same comfort and cost savings as conventional thermostats when operated under real-time price tariffs. By using suitably selected sampling rates to limit fast-cycling of equipment, $T\delta_0$ thermostats were hypothesized to give rise to readily linearized aggregate load models. However, the aggregate control of these loads has yet to be analyzed and simulated in detail to resolve steady state control error issues and achieve utility-scale functionality.

4.2 Aggregate Load Curtailment Model

In this section we develop a model of aggregate load when using $T\delta_0$ thermostats and propose a general controller design approach that will allow various aggregate load control strategies to be explored.

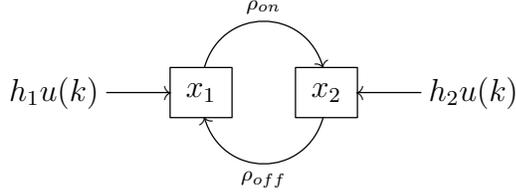


Figure 4.2: General state-space model of discrete-time zero-deadband aggregate thermostatic loads.

4.2.1 Aggregate Load Model

The current standard practice for direct load control can be readily applied to discrete-time thermostats when operating with zero deadband. When the sampling time t_s exceeds the minimum heating/cooling system refractory state time t_{min} we can reduce the state-space model to a second-order model as follows. Because there is no deadband we can ignore the refractory states n_{off}^* and n_{on}^* shown in Figure 4.1. We then derive the aggregate load response using a discrete-time state-transition representation for $T\delta_0$ thermostats

$$\begin{aligned} n_{on}(k+1) &= (1 - \rho_{on})n_{on}(k) + \rho_{off}n_{off}(k) \\ n_{off}(k+1) &= \rho_{on}n_{on}(k) + (1 - \rho_{off})n_{off}(k) \end{aligned} \tag{4.1}$$

where k is given in units of the sampling interval t_s , and ρ_{off} is the rate at which systems move out of the *off* state and ρ_{on} is the rate at which they move out of the *on* state, which we derive from the population average properties of individual homes' thermal responses.

We can now consider a discrete-time model, as shown in Figure 4.2, where the states $x_1 = n_{on}$ and $x_2 = n_{off}$ represent the number of responsive devices in the *on* and *off* states, respectively. The time t_{min} is generally regarded to be in the range of 1 to 2 minutes, so we cannot consider designs where $t_s < 1$ minute without having to reintroduce the refractory states in the model. The rate parameters ρ_{on} and ρ_{off} represent the fraction of those devices whose indoor air temperature τ crossed the indoor temperature setpoint τ_D is any given interval t_s . The rate parameters of the discrete-time model are determined from how the thermostat setpoint threshold τ_D divides the population occupying each state. We represent the rates at which devices are added into (or removed from) the controlled device population from (or to) the general uncontrolled device population by $\mathbf{h}u(k) = \begin{bmatrix} h_1 \\ h_2 \end{bmatrix} u(k)$.

From Equation (4.1) we develop a single-input/single-output demand response system state space representation for the net change in scalar load $y(k) < 0$ based on the scalar load control signal $u(k) > 0$

$$\mathbf{x}(k+1) = \underbrace{\begin{bmatrix} 1 - \rho_{on} & \rho_{off} \\ \rho_{on} & 1 - \rho_{off} \end{bmatrix}}_G \mathbf{x}(k) + \mathbf{h} u(k) \quad (4.2)$$

$$y(k) = \mathbf{c} \mathbf{x}(k)$$

where G represents the state transition matrix for the population of thermostats, \mathbf{h} represents the aggregate load control input vector and \mathbf{c} represents the aggregate load output vector. In general the input vector \mathbf{h} will be determined by the utility's choice of which control signals are sent to thermostats and how thermostats curtail loads. The particulars of the output matrix \mathbf{c} are determined by the nature of the response that is of interest, e.g., total load reduced or increased, or net change in load.

In the case of residential thermostats, we compute the rates ρ_{off} and ρ_{on} from the population statistics of the rates r_{on} and r_{off} at which indoor air temperature deviation $\Delta\tau = \tau - \tau_D$ changes in a single home. The rates of change of temperature deviation are determined from the second-order thermal response [86]

$$q(t) = \left(\frac{C_A C_M}{U_M} \right) \ddot{\tau} + \left[C_A + C_M \left(1 + \frac{U_A}{U_M} \right) \right] \dot{\tau} + U_A \tau \quad (4.3)$$

where U_A the thermal conductance of the indoor air to the outdoor air, C_A is the heat capacity of the indoor air, U_M is the thermal conductance of the indoor air to the building's solid mass, and C_M is the heat capacity of the building's solid mass. The heat function $q(t)$ includes both the internal, solar and ventilation heat gains and losses $q_I(t)$, as well as the heat gain or loss $q_H(t)$ resulting from operation of the heating/cooling system. From this we can derive the rates

$$r_{off}(t) = \dot{\tau}_{off} = -\frac{U_A}{C_A} \tau_A(t) - \frac{U_M}{C_A} \tau_M(t) + \frac{1}{C_A} q_I(t)$$

$$r_{on}(t) = \dot{\tau}_{on} = r_{off} + \frac{1}{C_A} q_H(t)$$

when the heating/cooling system is *off* and *on*, respectively. We assume that the heating/cooling systems are sized appropriately so that $r_{off}(t) < 0 < r_{on}(t)$ when

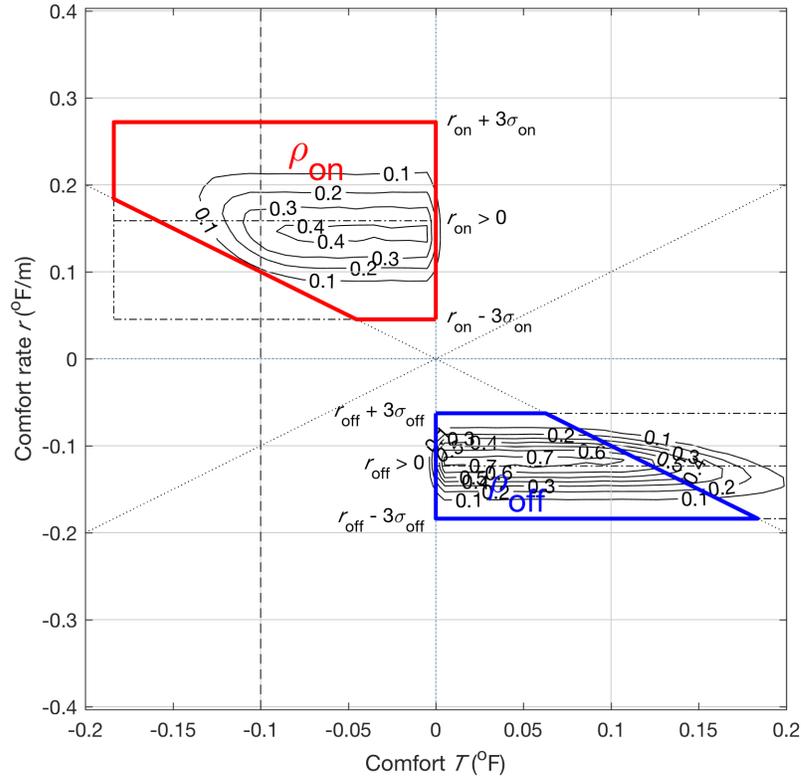


Figure 4.3: Discrete-time heating thermostat transition probabilities for *on* and *off* states with PDF of 10^6 homes with a setpoint change of -0.1°F .

heating and $r_{on}(t) < 0 < r_{off}(t)$ when cooling for all t . In the remainder of this chapter we will consider only the heating case, with the understanding that the cooling case is similar in every respect, with sign changes where appropriate.

The computation of ρ_{on} and ρ_{off} is illustrated in Figure 4.3. We now assume that immediately following a control action all the devices in a particular state can be inscribed in a rectangle, the horizontal dimension of which covers the range of indoor air temperatures τ and the vertical dimension of which covers the range of its derivative $r = \dot{\tau}$. The mean rates of devices during the interval k to $k+1$ are denoted $\bar{r}_{on}(k)$ and $\bar{r}_{off}(k)$. Device operating at the lowest rate are denoted $\bar{r}_{on}(k) - 3\sigma_{on}(k)$ and $\bar{r}_{off}(k) - 3\sigma_{off}(k)$ where $\sigma_{on}(k)$ and $\sigma_{off}(k)$ are the standard deviations of rates $r_{on}(k)$ and $r_{off}(k)$, respectively. These devices have a lower probability of crossing the setpoint threshold τ_D than ones running at the highest rates $\bar{r}_{on}(k) + 3\sigma_{on}(k)$ and $\bar{r}_{off}(k) + 3\sigma_{off}$. We assume that the distribution of devices within the rectangle has

virtually zero skew and we further assume that very nearly all of the device rates in the population fall within the ranges $\bar{r} \pm 3\sigma$ for both the *on* and *off* states. The zero skew assumption may not be reasonable for large changes in setpoint, but such consideration is beyond the scope of this chapter. In addition we assume that $3\sigma < \bar{r}$ for both the *on* and *off* states, a condition which is expected to be satisfied because we assume that the devices are suitably oversized for their applications, as is the common practice.

Two distinct cases must be considered depending on whether all the faster devices cross the τ_D threshold. In the first case (shown for the *on* state in Figure 4.3) only the devices in the blue region ρ_{off} will transition to the off state. We also know that the fastest devices in the complementary mode will overshoot no further than $\tau_{off}(k+1) = \tau_D + \bar{r}_{on}(k) + 3\sigma_{on}(k)$. From this we can define the probabilities of devices transitioning out of the *off* and *on* states as

$$\rho_{off} = \frac{\bar{r}_{off}}{\bar{r}_{on} + 3\sigma_{on}} \vee 1 \quad \text{and} \quad \rho_{on} = \frac{\bar{r}_{on}}{\bar{r}_{off} + 3\sigma_{off}} \vee 1$$

respectively, where $\vee 1$ denotes the unity saturation limit for the fraction of devices that can transition from a particular state during a single time interval t_s .

In the second case (shown for the *off* state in Figure 4.3) the devices in the red region ρ_{on} will transition to the *on* state. In this fast transition case we have

$$\rho'_{off} = 1 - \frac{3\sigma_{off}}{4\bar{r}_{off}} \quad \text{and} \quad \rho'_{on} = 1 - \frac{3\sigma_{on}}{4\bar{r}_{on}}$$

where the different form arises from the truncation of region **B** as compared to region **A**. The choice of which value of ρ to use is based on which state has the faster devices, which can vary dramatically with outdoor air temperature and heating/cooling system performance. When $\bar{r}_{on} > \bar{r}_{off}$, then ρ'_{on} and ρ_{off} are used, and when $\bar{r}_{on} < \bar{r}_{off}$, then ρ'_{off} and ρ_{on} are used.

Note that consequently the values of ρ_{on} and ρ_{off} arise from the aggregate behavior of the populations of devices whose temperatures move at the rates r_{on} and r_{off} with variances σ_{on}^2 and σ_{off}^2 , respectively.

Table 4.1: House thermal parameters.

Parameter	Unit	-3σ	Mean	$+3\sigma$
U_A	[BTU/°F.h]	200	350	500
C_A	[BTU/°F]	1551	2000	2449
U_M	[BTU/°F.h]	503	2000	3497
C_M	[BTU/°F]	7007	10000	12994
T_S	[°F]	69	72	75
Q_H	[BTU/h]	2764	11552	20340

4.2.2 Load Control Model

The basic “one-shot” load curtailment control strategy that is typically implemented by utilities can be described using Equation (4.2) with $\mathbf{h} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$ and $\mathbf{c} = \begin{bmatrix} 0 & 1 \end{bmatrix}$. Such strategies turn off $u(0)$ devices that are on, after which we observe by how many devices the load has reduced. Given knowledge of the average kW load \bar{Q} per device, these quantities can be given in kW if desired. We obtain the pulse transfer function for the “one-shot” load curtailment when discrete-time zero-deadband thermostats are employed:

$$\frac{Y(z)}{U(z)} = \frac{(z - b)}{(z - 1)(z - a)} \quad (4.4)$$

where $a = 1 - \rho_{off} - \rho_{on}$ and $b = 1 - \rho_{on}$. We make the following observations about this system.

1. The system is marginally stable. The dominant non-integrating pole is stable because $0 < \{\rho_{off}, \rho_{on}\} < 1 \implies -1 < a < 1$.
2. The system has a minimum-phase because $0 < \rho_{on} < 1 \implies 0 < b < 1$.
3. The dominant pole is always to the left of the zero because $0 < \{\rho_{off}, \rho_{on}\} < 1 \implies a < b$.

The relationship of the pole to the zero for various outdoor temperature conditions can be obtained using house thermal parameters such as the ones presented in Table 4.1. For the values of Table 4.1 the results are illustrated in Figure 4.4.

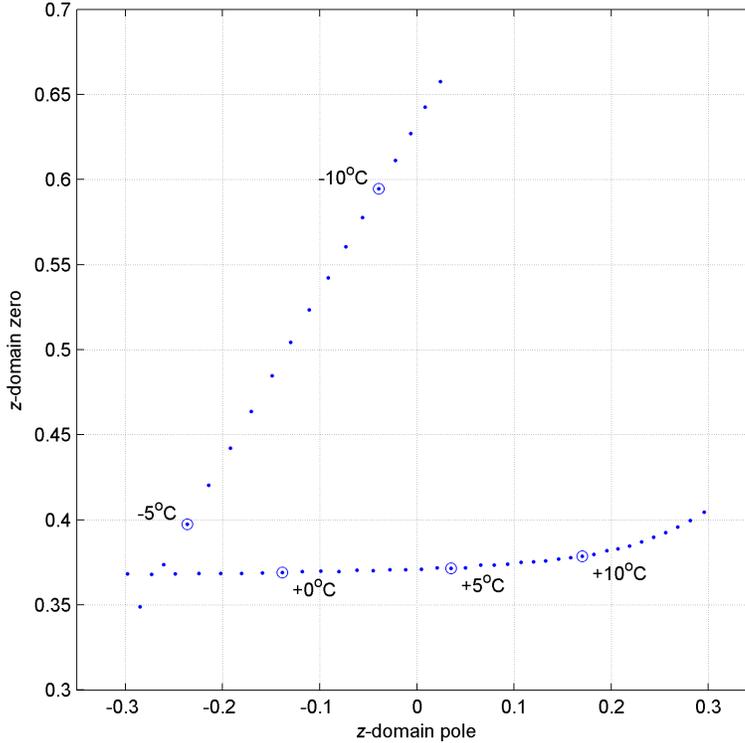


Figure 4.4: Zero and pole locations of Equation (4.4) for a random population of 1 million homes with thermal parameters given in Table IV at various outdoor air temperatures.

4.2.3 Open-Loop Response

The impulse response of the open loop system, Equation (4.4), for an impulse $u(0) = 1$ is :

$$y(k) = \frac{1 - b + (a - b)a^{k-2}}{1 - a} \quad \text{for } k = 1, 2, 3, \dots \quad (4.5)$$

with $y(0) = 0$, which will always be the response of a “one-shot” load curtailment signal when the loads are controlled by discrete-time thermostats with no deadband.

The steady state response is

$$y(\infty) = \frac{\rho_{on}}{\rho_{off} + \rho_{on}}$$

which we observe is the population average duty cycle R and is independent of $u(k)$ for $k > 0$ provided that $\sum_{j=0}^{\infty} u(j) > 0$. We also note that any signal $u(k) > 0$ will add more devices to the controlled population while $u(k) < 0$ will remove devices from the controlled population. For any $k > 0$ we only require $u(k) = -\sum_{j=0}^{k-1} u(j)$

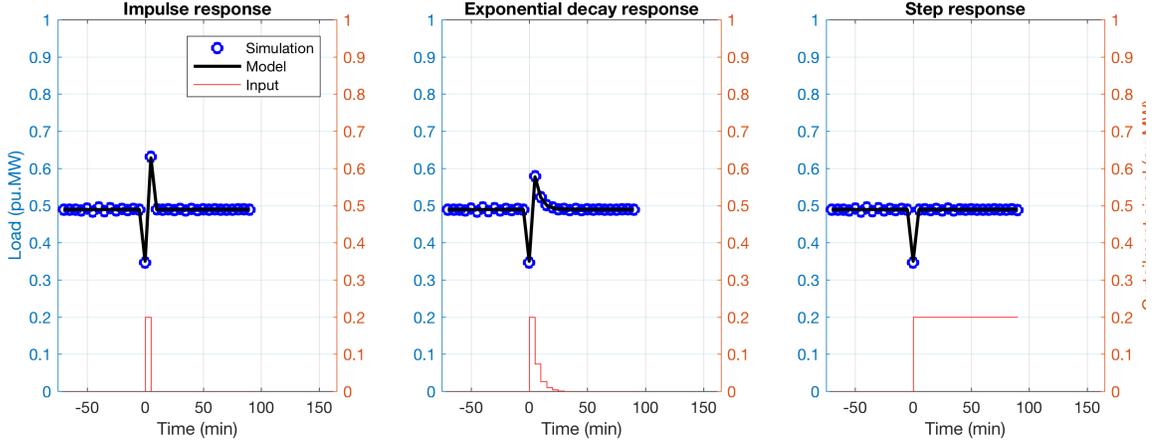


Figure 4.5: Open loop impulse (left), decay (center), and step (right) response of aggregate load model compared to agent-based simulation for 100,000 thermostats per unit input u at -10°C .

to return to the initial condition $\mathbf{x}(0)$ and when $\sum_{j=0}^{\infty} u(j) = 0$, we will always have $y(\infty) = \mathbf{c}\mathbf{x}(0)$. The model responses to various inputs are compared to simulations with 100,000 thermostats using an agent-based simulation. The results illustrated in Figure 4.5.

4.2.4 Model Identification

The performance of a utility-scale implementation of aggregate load controllers to be designed in the next section depends on the estimation accuracy of model parameters ρ_{on} and ρ_{off} . Values for these parameters can be obtained by comparing the response to an impulse input with Equation (4.5). We can show that the responses at $k = 2$ and 3 are sufficient to give an estimate for the observer parameters.

A single impulse response can be used to provide a relatively quick and simple method of model parameter identification. After a single impulse $u(0)$ and initial conditions $\hat{\mathbf{x}}(0) = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$, the system's initial response is observed to be $\hat{\mathbf{x}}(1) = \begin{bmatrix} 0 \\ y(1) \end{bmatrix}$. After a second time-step, the system is observed to be in the state $\hat{\mathbf{x}}(2) = \begin{bmatrix} x_1(2) \\ y(2) \end{bmatrix}$. Given Equation (4.1) we can estimate

$$\hat{\rho}_{off} = 1 - \frac{y(2)}{u(0)}$$

and similarly after a third observation $y(3)$ is obtained we estimate

$$\hat{\rho}_{on} = \frac{y(3) - \frac{[y(2)]^2}{u(0)}}{u(0) - y(2)}.$$

This method of estimating the model parameters makes two assumptions that must be considered:

1. The initial conditions are $\hat{\mathbf{x}}(0) = \mathbf{0}$. This condition is achieved by releasing all the devices currently under control and waiting for the normal settling time of controlled devices to elapse to ensure that the uncontrolled population is roughly in both state and thermal equilibrium.
2. Only a single control impulse $u(0)$ is sent at time $k = 0$ and then no control signals $u(k)$ for $k = 1, 2, 3, \dots$ are sent so the impulse response can be clearly discerned in the outputs $y(2)$ and $y(3)$.

These conditions are relatively easy to create and the impulse $u(0)$ need not be large to obtain useful measurements, particularly if the test is repeated multiple times for each outdoor air temperature. Using this method a database of model parameters can be obtained and used to estimate model statistics as well. Furthermore, the magnitude $y(1)$ will give an estimate of the product $\mathbf{h}\mathbf{c}$, while observation of $y(4)$ permits the estimation of \mathbf{h} and \mathbf{c} separately, if needed.

Finally, it is not necessary to probe the system response at all outdoor conditions because the relationship of a and b is well known, particularly for infrequent peak load conditions that can be more difficult to observe. The linear relationship of a and b over the range of low outdoor air temperatures is seen in Figure 4.4 and allows reliable extrapolation from more frequent conditions to more rarely observed and more critical peak load conditions.

In the likely case that measurement noise is present, a mean square approximation of these parameters may be considered by producing a series of impulses spaced apart by a sufficient interval to guarantee that Assumption (1) above is satisfied. Under high duty-cycle conditions, the linear relationship of the zero and pole assures that varying conditions are not an obstacle to determining the slope of the line that relates them using a least-squares fit.

4.3 Aggregate Demand Response Controller Design

In the previous section we proposed an aggregate load model and discussed its main open-loop properties. It is important to this study that we examine the typical range of control strategies suited to direct dispatch of demand response resources and assess the degree to which these strategies work satisfactorily for aggregations of $T\delta_0$ thermostats. Therefore, in this section we examine various controller designs, all of which are variations implemented on the general controller design shown in Figure 4.6. The controller design parameters for this general controller are as follows:

- \mathbf{h} is the system input vector for the response to the scalar signal $u(k)$. This is generally a curtailment signal and indicates how many devices are turned *off*.
- \mathbf{c} is the system output vector for the scalar load $y(k)$ arising from the internal states $\mathbf{x}(k)$.
- $\tilde{\mathbf{h}}$ is the observer input vector.
- $\tilde{\mathbf{c}}$ is the observer output vector.
- h is the scalar reference input gain.
- \mathbf{K}_c is the observer gain vector.
- K_q is the integral error feedback gain scalar.

The flexible design of controller allows for many of the basic control strategies that are typically employed for discrete-time linear time-invariant system. This is done with the understanding that some of the parameters may change over time intervals much longer than the time horizon over which most demand response control objectives are stipulated. In particular it is expected that the state transition rates ρ_{on} and ρ_{off} will change as a function of outdoor air temperature, but that the relationship will be relatively easy to obtain for the aggregate population and that it will be sufficiently consistent between seasons to allow simple system identification approaches to provide accurate long term model parameters. A simple method of identifying these parameters was discussed in the previous section.

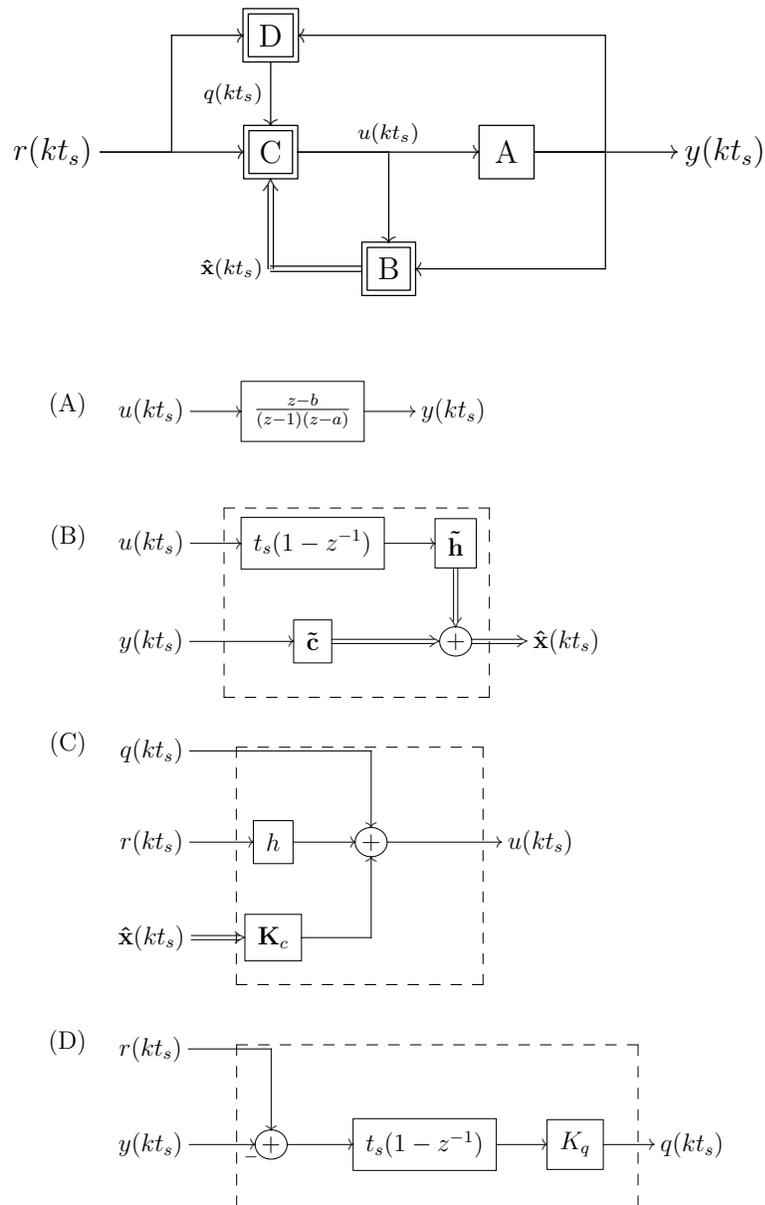
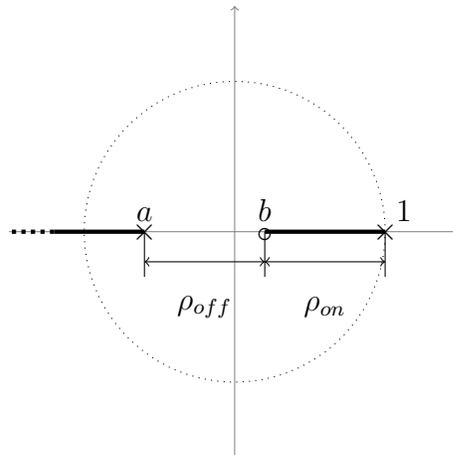


Figure 4.6: General structure of the controller (top): Block (A) is the aggregate load model, (B) is the reduced-order observer, (C) is the load controller, and (D) is the integral error feedback.

Table 4.2: Controller design configurations.

Configuration	Observer		Controller		Error feedback
	$\tilde{\mathbf{h}}$	$\tilde{\mathbf{c}}$	K_c	h	K_q
Proportional	$\begin{bmatrix} 0 \\ 0 \end{bmatrix}$	$[0 \ 1]$			
$[0 \ \rho_{on} + \frac{1}{2}\rho_{off}]$	0	$\rho_{on} + \frac{1}{2}\rho_{off}$			
Unity damping	$\begin{bmatrix} 1 \\ 0 \end{bmatrix}$	$[0 \ 1]$	$[\hat{\rho}_{on} \ 1 - \hat{\rho}_{on} - \hat{\rho}_{off}]$	$1 - \hat{\rho}_{on} - \hat{\rho}_{off}$	0
Deadbeat	$\begin{bmatrix} 1 \\ 0 \end{bmatrix}$	$[-1 \ 1]$	See Eq. (4.7)	See Eq. (4.8)	0
Pole placement	$\begin{bmatrix} 1 \\ 0 \end{bmatrix}$	$[-1 \ 1]$	See Eq. (4.9)	See Eq. (4.10)	0
Integral error feedback	$\begin{bmatrix} 1 \\ 0 \end{bmatrix}$	$[-1 \ 1]$	See Eq. (4.11)	See Eq. (4.12)	See Eq. (4.11)

Figure 4.7: Discrete-time root-locus of aggregate $T\delta_0$ thermostatic loads.

4.3.1 Proportional Control

We can now consider the behavior of proportional control by examining the root locus of the closed-loop system. With $-1 < a < b < 1$, the root-locus in Figure 4.7 indicates that the system has two real poles and a zero that is always between the poles.

Since the open-loop system in Equation (4.4) has a discrete integrator, the steady state error for a step input is zero. Unfortunately, numerical methods do not find

Table 4.3: Maximum attenuating proportional control gains for various conditions.

Temperature	Load	Gain	Zero	Dominant pole	Gain margin
[°C]	[% peak]	K	[/h]	[/h]	[dB]
-15	79	0.45	0.78	0.93	3.1
-10	70	0.40	0.66	0.91	3.9
-5	61	0.20	0.49	0.93	11.1
0	33	0.70	0.30	0.65	0.8
5	25	0.80	0.27	0.56	0.3
10	17	0.85	0.26	0.49	0.6
15	10	0.90	0.24	0.41	0.6

values of K with acceptable phase and gain margins. However, gains can be found for the fastest possible attenuation for various outdoor conditions, as shown in Table 4.3.

The Jury-Marden test gives us the stability constraint on the closed-loop gain

$$\rho_{on} < K < \rho_{off} + \rho_{on}$$

which can be a very narrow range and highly dependent on accurate knowledge of the value of ρ_{on} , particularly when ρ_{off} is small. Small values of K may lead to slow response under certain conditions. PID control can yield a stable aggregate load controller under conditions we expect to encounter in a realistic utility operational setting. However, the response can also be somewhat oscillatory under higher load conditions, when both fast and reliable aggregate load control is most needed, as shown in Figure 4.8 (left). The narrow band of constraints on K limit at high loading conditions limits the possibility of improving performance to such an extent that proportional control seems largely impractical for direct load control. Indeed at peak load only marginally stable control can be achieved when $K = \rho_{on}$.

4.3.2 Proportional-Derivative Control

Faster response than proportional control is often achieved by using a proportional derivative controller such that

$$G_{pd}(z) = \frac{(z - b)(k_1 z + k_2)}{(z - 1)(z - a)}$$

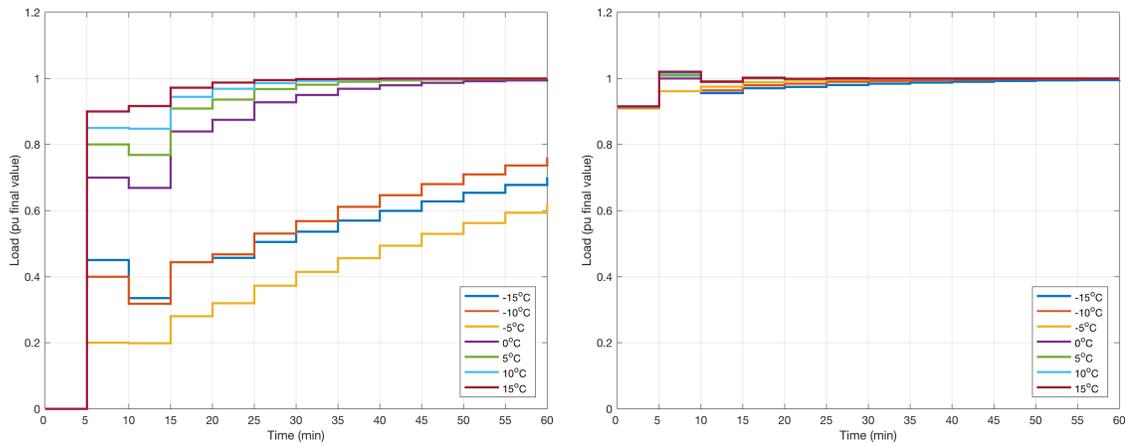


Figure 4.8: 100 MW proportional load control step response with maximum attenuating proportional control gains based on the load parameters in Table 4.1 (left) and proportional-derivative control step response (right).

Table 4.4: Proportional-derivative controller design parameters.

T_O	k_1	k_d	z_1	z_2	p_1	p_2
-15	9.94	4.38	0.78	-0.44	0.80	-0.40
-10	9.95	4.38	0.66	-0.44	0.68	-0.41
-5	9.95	4.38	0.49	-0.44	0.54	-0.42
0	10.52	4.63	-0.44	0.30	-0.38	0.33
5	10.66	4.69	-0.44	0.27	-0.36	0.30
10	10.77	4.74	-0.44	0.26	-0.35	0.27
15	10.87	4.79	-0.44	0.24	-0.33	0.25

Solving for the fastest possible response with zero poles we obtain

$$k_1 = \frac{a}{b^2} - \frac{a}{b} - \frac{1}{b} \quad \text{and} \quad k_2 = \frac{a}{b}$$

However, the stability margin of the system is not suitable for operation in noisy conditions.

Using pole placement for damping $\xi = 0.8$ and settling time $t_s = 1$ hour, as shown in Table 4.4, does not offer improvements in the system's stability characteristics in spite of satisfactory time-domain response to step inputs, as shown in Figure 4.8 (right).

Lead and/or lag compensator design for the wide variety of conditions that the controller is required to operate under would require an adaptive controller and was

not considered here.

4.3.3 Unity Damped Control

We can design a direct aggregate load control strategy for $T\delta_0$ thermostats that will maintain a constant desired load curtailment $r(k) > 0$ for $k > 0$, assuming that $y(0) = r(0) = 0$. We desire only that

$$y(k) = r(k) \quad \text{for } k = 1, 2, 3, \dots$$

which gives us

$$u(k) = (\mathbf{c}\mathbf{h})^{-1}[y(k) - r(k) + \mathbf{c}(\mathbf{I} - \mathbf{G})\hat{\mathbf{x}}(k)] \quad \text{for } k = 1, 2, 3, \dots$$

where \mathbf{G} is the state transition matrix for the population of thermostats and $\hat{\mathbf{x}}(k)$ is an estimate of $\mathbf{x}(k)$.

For any curtailment control system the accumulated inputs from $u(0)$ to $u(k-1)$ represents the total number of devices N that have been curtailed up to the time k . So $\hat{x}_2(k) = y(k)$ represents the load that is still *off* at the time k . Therefore we must have

$$\hat{x}_1(k) = \sum_{j=0}^{k-1} u(j) - y(k) \quad \text{for } k = 1, 2, 3, \dots,$$

which represents the load that has returned to the *on* state at the time k . This gives us the estimated state

$$\hat{\mathbf{x}}(k) = \begin{bmatrix} \sum_{j=0}^{k-1} u(j) - y(k) \\ y(k) \end{bmatrix} \quad \text{for } k = 1, 2, 3, \dots$$

This state can be found from the input $u(k)$ and output $y(k)$ using a reduced order observer described by:

$$\hat{\mathbf{x}}(k) = \tilde{\mathbf{b}} \sum_{j=0}^k bu(j) + \tilde{\mathbf{c}} y(k) \quad \text{for } k = 1, 2, 3, \dots$$

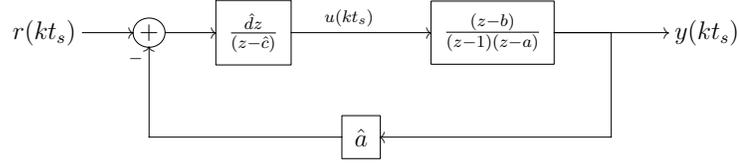


Figure 4.9: Unity damped system diagram.

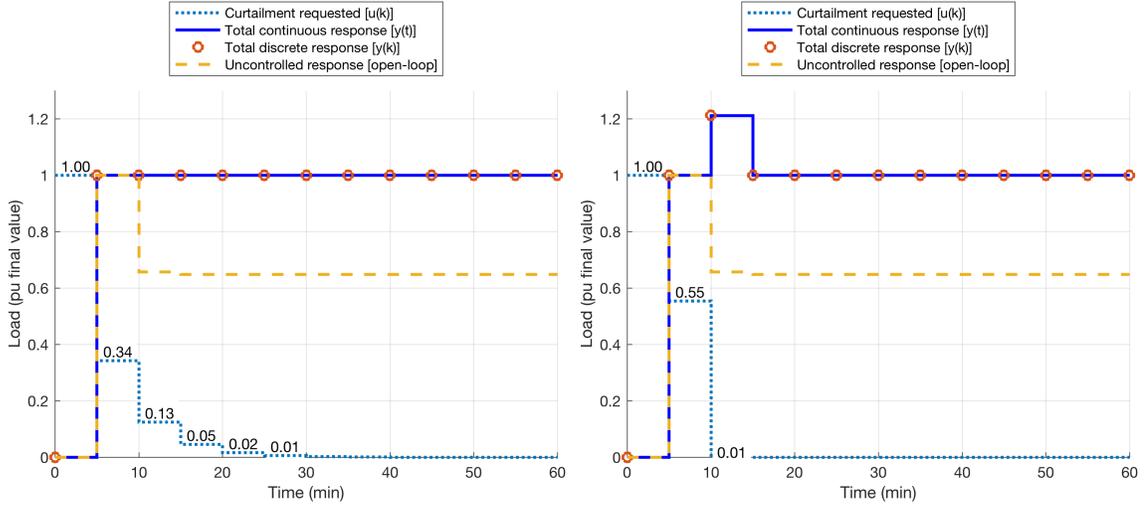


Figure 4.10: Unity damped (left) and deadbeat (right) responses of aggregate load controllers.

where $\tilde{\mathbf{b}} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$ and $\tilde{\mathbf{c}} = \begin{bmatrix} 1 \\ -1 \end{bmatrix}$. From this we can determine the load control signal:

$$u(k) = r(k) - \hat{a}y(k) - (1 - \hat{b}) \sum_{j=0}^{k-1} u(j) \quad \text{for } k = 1, 2, 3, \dots, \quad (4.6)$$

where $\hat{a} = (1 - \hat{\rho}_{off} - \hat{\rho}_{on})$ and $\hat{b} = 1 - \hat{\rho}_{on}$ with $\hat{\rho}_{off}$ and $\hat{\rho}_{on}$ being the estimates of the aggregate load response. This controller can be implemented as shown in Figure 4.9 with $\hat{c} = \frac{1}{1 + \hat{\rho}_{on}}$ and $\hat{d} = \frac{\hat{\rho}_{on}}{1 + \hat{\rho}_{on}}$. We observe that $0 < \hat{d} < 0.5 < \hat{c} < 1$. Thus the added pole is stable and the added zero does not affect the minimum phase property of the system. We note that this controller should reach steady state on the first iteration and thus has a damping ratio of 1.0 and settling time of t_s , as shown in Figure 4.10 (left).

The introduction into the system of model parameters \hat{c} and \hat{d} creates a source of constant disturbances in the system that can result in a steady state error. This

problem limits general applicability of this controller design unless integral error feedback control is used. An alternative approach to mitigate model error is to include information obtained directly from controllable devices. This would be the case if bidding mechanisms are used, such as when retail markets are implemented using double auctions [36, 24, 25].

4.3.4 Deadbeat Control

An alternative type of controller is a deadbeat controller that uses only two load control impulses to achieve steady state [87]. The controller has the advantage that it does not continually draw on the uncontrolled population of devices to achieve the control objective. It has the disadvantage that it may overshoot on the second time-step, as shown in Figure 4.10 (right).

The state $\hat{\mathbf{x}}(k)$ and output $y(k)$ are determined using the matrices

$$\tilde{\mathbf{h}} = \begin{bmatrix} \frac{(1-\hat{b})z}{z-1} \\ 0 \end{bmatrix} \quad \text{and} \quad \tilde{\mathbf{c}} = \begin{bmatrix} 0 \\ \hat{a} \end{bmatrix}$$

We solve for the feedback gain for zero poles using

$$\mathbf{K} = \begin{bmatrix} \hat{a} + 1 & -\hat{a} \end{bmatrix} \tilde{\mathbf{A}}^{-T} \mathbf{C}^{-1} \quad (4.7)$$

where $\tilde{\mathbf{A}} = \begin{bmatrix} 1 & 0 \\ \hat{\rho}_{on} + \hat{\rho}_{off} - 2 & 1 \end{bmatrix}$ is a Toeplitz matrix and $\mathbf{C} = \begin{bmatrix} 0 & \hat{\rho}_{off} \\ 1 & 1 - \hat{\rho}_{off} \end{bmatrix}$ is the controllability matrix. The tracking reference input gain is

$$h = \frac{1}{\mathbf{c}(\mathbf{I} - \mathbf{A} + \mathbf{hK})^{-1}\mathbf{h}} \quad (4.8)$$

4.3.5 Pole Placement Control

In the general case of a pole-placement controller [87], we have the same controllability matrix and Toeplitz matrix as the deadbeat controller above. Given a desired damping coefficient ξ and settling time t we compute the desired pole locations z_1 and z_2 and we obtain the tuned controller gains

$$K_c = \begin{bmatrix} \hat{\rho}_{on}[\hat{\rho}_{off} + \hat{\rho}_{on}(1 + z_1 + z_2)] + z_1 z_2 - z_1 - z_2 - 1 \\ -\hat{\rho}_{off} - \hat{\rho}_{on} - z_1 - z_2 \end{bmatrix}^T \quad (4.9)$$

and

$$h = \frac{z_1 z_2 - z_1 - z_2 - 1}{\hat{\rho}_{on}} \quad (4.10)$$

As in the case of unity-damped and deadbeat controllers, any error in $\hat{\rho}_{on}$ and $\hat{\rho}_{off}$ is expected to result in a steady state error.

The performance of the five different controllers discussed in this section depends on the accuracy of the estimated system parameters ρ_{on} and ρ_{off} . If these are not sufficiently accurate, this may lead to unacceptably large steady-state errors. Assuming that the values of the system parameters change more slowly than the time constants of the system, steady-state error can be mitigated using an integral error feedback controller. The design of such a controller is discussed next.

4.3.6 Integral Error Feedback

To correct for the steady state error in the pole-placement controller we implement integral error feedback using an augmented state

$$q(k+1) = q(k) + t_s[r(k) - y(k)].$$

We include the integral feedback error in the state-space representation using the augmented controllability matrix

$$\mathbf{C} = \begin{bmatrix} 0 & \rho_{off} & 0 \\ 1 & 1 - \rho_{off} & 0 \\ 0 & -t_s & 1 \end{bmatrix}$$

and the characteristic polynomial is $a(z) = (z-1)^2(z-a)$ or

$$a(z) = z^3 - (a+2)z^2 + (2a+1)z - a$$

The augmented Toeplitz matrix is therefore

$$\tilde{\mathbf{A}} = \begin{bmatrix} 1 & 0 & 0 \\ \rho_{off} + \rho_{on} - 3 & 1 & 0 \\ 2\rho_{off} + 2\hat{\rho}_{on} - 3 & \rho_{off} + \rho_{on} - 3 & 1 \end{bmatrix}.$$

The desired characteristic polynomial is simply $\alpha(z) = (z-z_1)(z-z_2)(z-z_q)$

where z_1 , z_2 , and z_q are the desired poles of the closed-loop system. Thus we have

$$\alpha(z) = z^3 - (z_1 + z_2 + z_q)z^2 + (z_1z_2 + z_1z_q + z_2z_1)z - (z_1z_2z_q)$$

from which we obtain controller gains based on the estimated model

$$\begin{bmatrix} \mathbf{K}_c & K_q \end{bmatrix} = \begin{bmatrix} \hat{\rho}_{on} + \hat{\rho}_{off} - 3 - z_1 - z_2 - z_q \\ z_1z_2 + z_1z_q + z_2z_q - 3 + 2\hat{\rho}_{on} + 2\hat{\rho}_{off} \\ 1 - \hat{\rho}_{on} - \hat{\rho}_{off} - z_1z_2z_q \end{bmatrix}^T \tilde{\mathbf{A}}^{-T} \mathbf{C}^{-1} \quad (4.11)$$

with the reference input gain

$$h = \frac{1}{z_1z_2 + z_1z_q + z_2z_q + 2\hat{\rho}_{on} + \hat{\rho}_{off} - 3} \quad (4.12)$$

This control design eliminates the steady-state error induced by model errors in $\hat{\rho}_{on}$ and $\hat{\rho}_{off}$ with a settling time determined by the pole z_q .

4.4 Agent-based Simulation Results

The controller designs were tested on an agent-based simulation [88] of 100,000 residential thermostats using a second-order building thermal model, including internal and solar gains, and ventilation losses. (The thermal properties of these loads are given in Table C.1.) The second order models are linearized for the given outdoor temperature resulting in first-order models for each house such that the individual homes have distinct air temperature change rates as a function of the state of the heating system. Note that the thermal model used in the agent-based simulation presented in this section is not an aggregate model as the one used for controller design in the previous section. This allows the performance of the controllers designed in the previous section to be evaluated using a plant model which better reflects reality. Therefore, the effects of disturbances caused by errors arising from model order reduction of the design model as well as measurement noise are considered. The results of the PDF model estimates are given in Table D.1.

To implement direct load control, thermostat setpoint changes are applied to a subset of uncurtailed heating units. The magnitude of the setpoint change is generally a function of the fastest rate of change, which at peak load is approximately r_{off} . The magnitude of this value was chosen to ensure that the impulse response resulted

Table 4.5: Controller design parameters for peak load (-15°C).

Control	Poles	Gains			Errors	
		K_c	K_q	h	Average	Maximum
Impulse	[0]	[0.00 0.00]	0.09	0.00	69%	71%
Proportional	[0]	[0.00 0.31]	0.09	0.31	44%	87%
Damped	[0.72]	[0.09 0.46]	0.09	0.91	18%	61%
Deadbeat	[0.72]	[0.28 1.46]	0.09	2.93	35%	40%
Pole placement	[0.43]	[0.60 0.59]	0.09	3.67	84%	92%
Integral	[0 0.72]	[-0.30 1.61]	0.09	1.00	13%	41%

in a 100% response at the first time step. The number of homes curtailed is based on the average heating system load when on such that $u(k) = 1$ is equivalent to 1 MW of load, or approximately $N_C = 10^6/\bar{Q}$, where \bar{Q} is the mean value of the heating unit load $Q = q_H/COP$, where COP is the heating unit efficiency.

When a negative value of $u(k)$ is obtained, units are released into the uncurtailed population. The simulation first releases systems that have been curtailed the longest, ensuring that the released population is the most diversified and exhibits the least rebound oscillation after returning to the uncurtailed population.

The controller design parameters discussed in Section 4.3 are generated for peak load conditions using the thermal parameters shown in Table 4.1. A summary of the controller design parameters studied are shown in Table 4.5.

The impulse response for $\tau_O = -15^\circ\text{C}$ is shown in Figure 4.11 (left). The response is different from the one expected from the aggregate model and illustrates the effect of the errors and noise induced by state fluctuations in the system which are not captured by the second-order aggregate load model. Further, the steady-state of the response and the settling time for the devices to reach their normal diversity are also indicated. The proportional control response does not have a steady state error, but this is not clearly visible because of the very slow response, as shown in Figure 4.11 (right).

The response of the unity damping controller is shown in Figure 4.12 (left). The effect of model error can be seen in the initial response, as a result of which it fails to quickly achieve the desired level of curtailment. The response of deadbeat control has the expected significant overshoot, but also exhibits large steady state error, as shown in Figure 4.12 (right).

The response of the tuned controller using pole placement shows a compromise

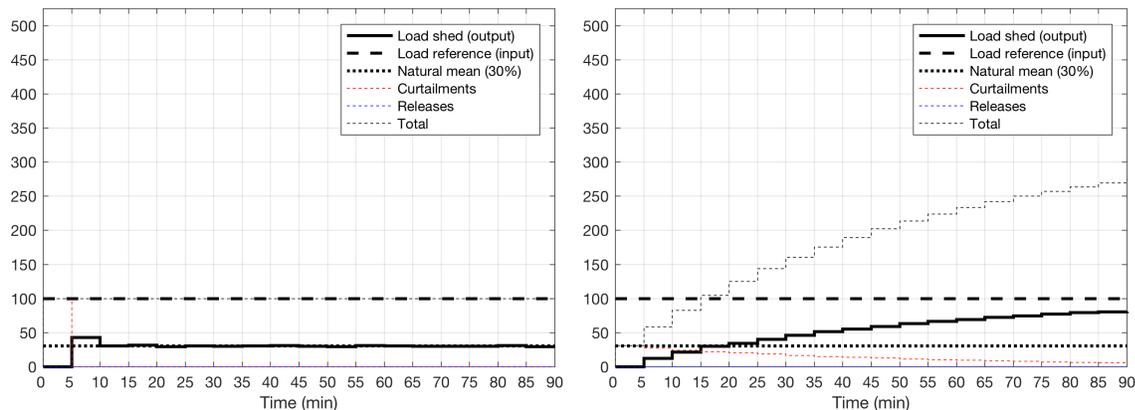


Figure 4.11: 100 MW impulse response (left) and proportional control (right) at -15°C .

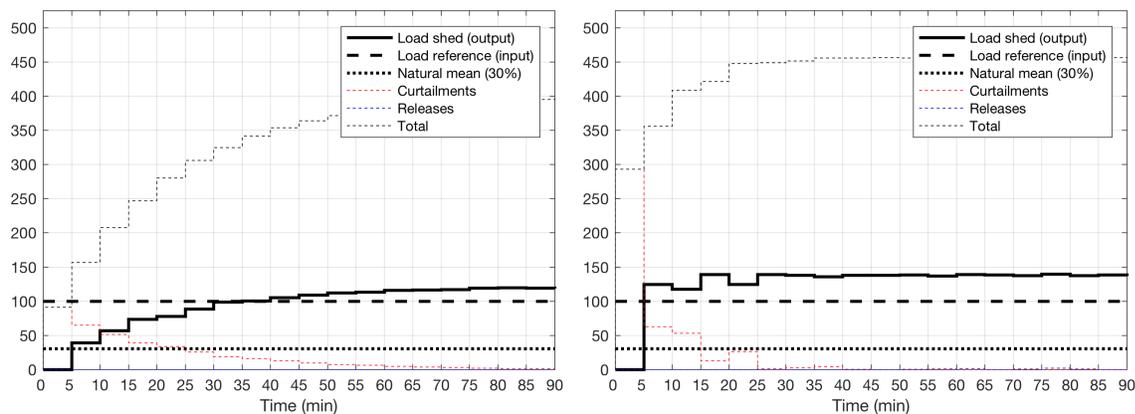


Figure 4.12: 100 MW unity damping load control (left) and deadbeat load control (right) at -15°C .

between the unity damping and deadbeat controller designs, but still exhibits a large steady state error, as shown in Figure 4.13 (left). The integral error feedback control response shown in Figure 4.13 (right) addresses the problems identified in the previous controller designs. The system exhibits an acceptable level of overshoot and maintains the desired curtailment level for more than 90 minutes.

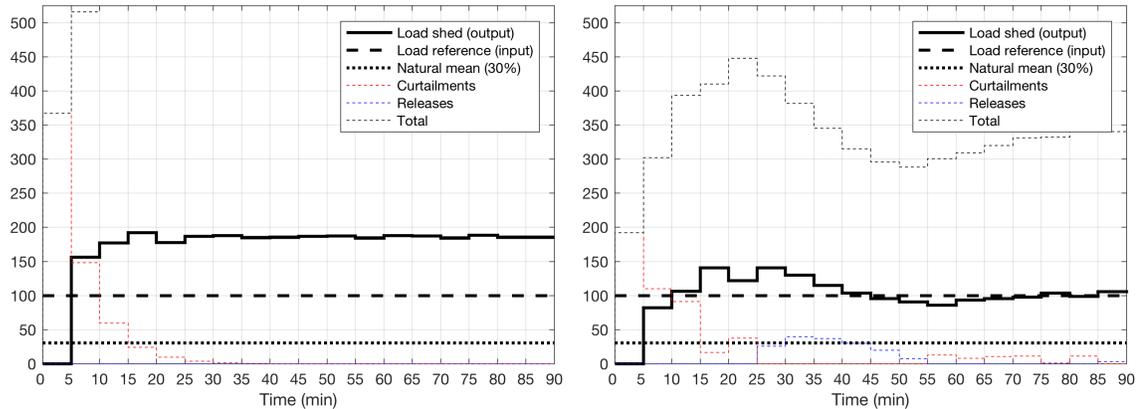


Figure 4.13: 100 MW tuned load control (left) and integral feedback control (right) at -15°C .

4.5 Summary of Results

In this chapter we considered the utility-scale direct load control problem for the situation when the controlled loads employ discrete-time zero-deadband ($T\delta_0$) residential thermostats. We have shown how $T\delta_0$ thermostats allow utility dispatchers to use small adjustments to the consumer’s setpoint to modulate the total load with greater precision and endurance than typically possible using current setback control of thermostats with non-zero deadbands. These new digital thermostats can serve as the basis for highly accurate direct load control systems, as well as price-based indirect load control systems.

We have derived a new linear aggregate load model based on the dynamics of load states and used its fundamental characteristics to consider a number of benchmark aggregate load control designs from first-principles. We used this model to design a simple closed-loop aggregate controller for a discrete-time utility-scale demand response dispatch system that is compatible with the requirements for both direct and indirect load control systems and tested the control design using a large-scale agent-based model of demand response based on thermostatic loads. We showed that the aggregate controlled load is stable, controllable and observable and exhibits both the transient and steady-state response characteristics necessary to serve equally well for utilities that seek to control load using either direct load control or price-based indirect demand response strategies.

Chapter 5

Regulation

Under the transactive control paradigm, retail markets for energy, capacity, and regulation services are deployed to provide an analogous realization of wholesale markets at the distribution level. In spite of the conceptual similarity, the behavior of retail markets differs significantly from that of wholesale markets and remains an active area of research [89]. In particular, load behavior usually dominates the behavior of retail systems, which contrasts with wholesale systems where generation is dominant. In addition, there are a number of important processes in bulk power interconnection operations that have yet to be integrated fully into the transactive paradigm. Two such processes are system frequency regulation and control area import/export schedule tracking.

In Section 5.1 we introduce the interconnection operation control platform. We then present the methodology for optimally controlling an area's response to system frequency deviations while tracking scheduled area exports. In Section 5.2 we propose the structure of the model and the design solution for an optimal area control policy. Finally in Section 5.3 we evaluate the performance of the optimal control policy when compared to the conventional control policy under varying demand response conditions.

5.1 Methodology

System operators that wish to use demand response resources to mitigate renewable intermittency must have the means to control responsive loads in much the same way they control responsive generating units. This can be done by updating the load

control system gains every few minutes given the available demand response resources committed to frequency regulation. Given these load control gain settings, the loads' responses to frequency deviations can be autonomous without requiring the use of an analog to AGC for loads.

It is quite feasible with today's technology to dispatch load control gain settings to load aggregators who then disseminate specific setpoints to the loads they control without having to dispatch AGC signals to each load directly. However, doing so requires adjustments to the existing frequency and inter-area power exchange control system. This section details how this is accomplished within the structure of modern power system control.

5.1.1 Frequency control mechanism

Primary control of bulk electric power systems is driven in part by deviations in frequency at the system level and modeled by the system transfer function $\frac{1}{Ms+D}$, where M represents the system's inertial response and D represent the system's damping response. Each control area implements a combination of speed-droop control on conventional generating units, under-frequency load shedding, and grid-friendly loads to provide primary control. Renewable generating units provide no frequency regulation capability because they cannot control their prime movers (wind and solar). Secondary control of frequency and area exports is provided by units equipped with master controllers and is based on the area control signal a using the conventional ACE formula

$$A(t) = \Delta Q(t) + B\Delta f(t), \quad (5.1)$$

where $A(t)$ is the raw ACE signal, $\Delta Q(t) = Q(t) - Q_s$ is the deviation of the net energy exported from the scheduled net exports over tie lines and $\Delta f(t) = f(t) - f_s$ is the interconnection's frequency deviation from scheduled frequency. If an area's net exports deviates from its schedule (because of either an internal or external disturbance), the area adjusts its generation (and potentially also its load) such that it eventually will zero out $A(t)$, while also providing adjustments necessary to support system-wide corrections to frequency deviation.

In most realizations the ACE signal is updated by the SCADA system about every 4 seconds and further passed through a smoothing filter so that it changes with a time-constant well in excess of the generating units' fastest response, e.g., 30-90 seconds, with the purpose of reducing wear and tear on generating unit governor motors and

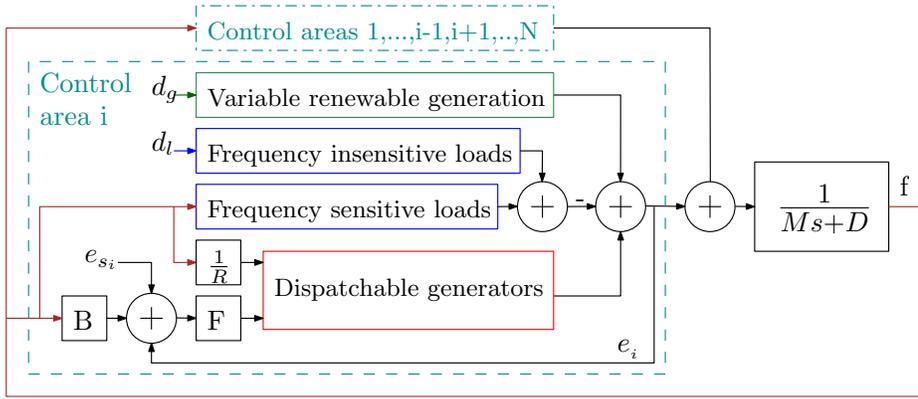


Figure 5.1: System frequency control diagram.

turbine valves [90]. In addition, this control action is subject to control performance standards (i.e., CPS-1 and CPS-2 [91]), although these are not considered in this study.

The fast response of frequency-sensitive loads (grid-friendly smart appliances) to the frequency deviation enables the grid operator to redispatch generating units in a more economically-efficient manner, although demand response may saturate relatively quickly. Figure 5.1 illustrates the system’s frequency regulation diagram with variable renewable generation and frequency sensitive demand response.

5.1.2 Transactive control platform

As noted in Chapter 2, Fuller’s definition of Transactive Control does not specify any particular physical or temporal control architecture. We chose the hierarchy illustrated in Figure 2.3 because it provides a relatively simple data flow between physical and temporal scales. Using this approach, the total generation and load is scheduled hourly such that for each control area a uniform price is obtained at which supply is equal to load plus net exports. Figure 5.2 illustrates an interconnection including N wholesale markets each belongs to a control area that exchange electricity through system tielines to increase the economic surplus. This schedule is used to set each area’s price P_s and net exports Q_s which are in turn used by 5-minute dispatch markets [24] to reallocate resources in response to deviations from the hourly schedule. Depending on the events that have occurred during the preceding 5-minute time period, the state of operation of generators and loads at the end of the period is not necessarily the same as at the beginning of the period. For example, the water

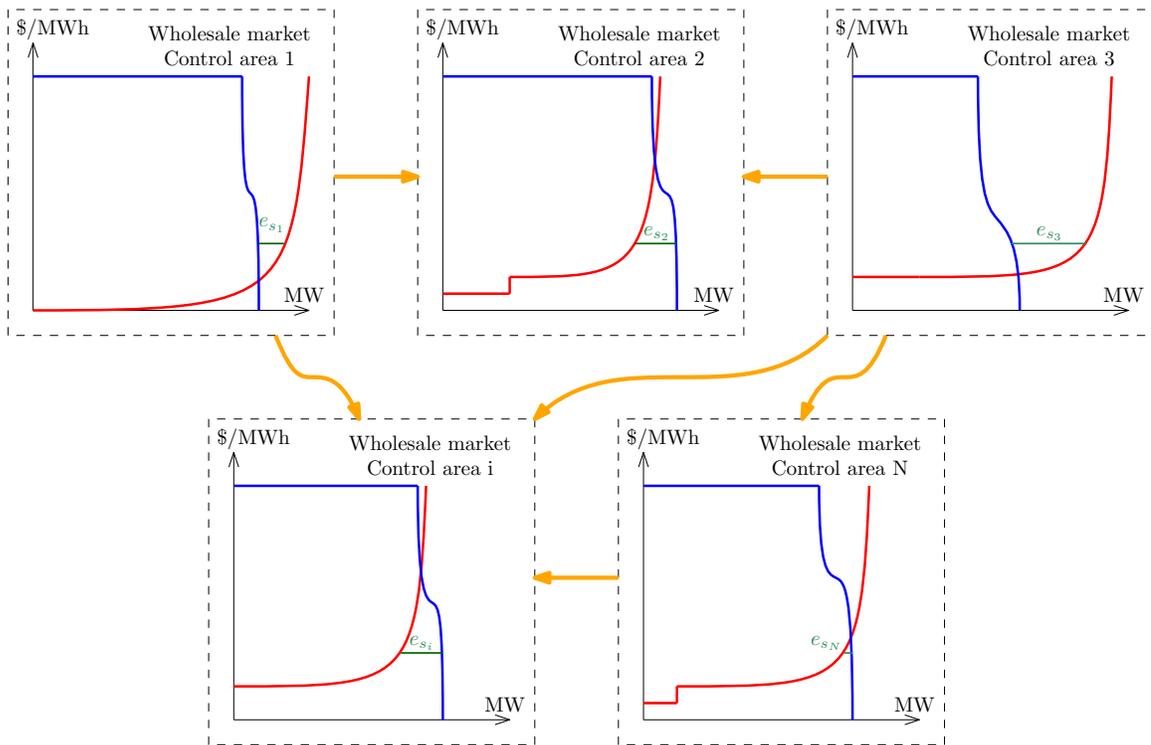


Figure 5.2: Inter-area transfer flows within an interconnected system consisting of N control areas.

temperature of a hot water tank whose heater was switched off at the last time period has lowered, and we expect this load might submit a higher bid to the market to avoid staying in the off mode and satisfy a higher level of demand urgency. Accordingly, generation and load resources may participate in the market with new bid prices, and as a result the clearing price P_D would change. However, the area export e_a should remain as close as possible to the hourly schedule.

5.1.3 Demand response integration in the 5-minute market

Figure 5.3 illustrates the impact of a supply disturbance on the 5-minute market settlement process. The blue and the red dashed curves represent the demand and the supply curves, respectively, in the next market cycle. In this example a portion of renewable generation (in the flat segment of supply curve) is lost and the supply curve is shifted to left by the magnitude of the disturbance ΔQ_s ¹. Load participation in

¹Other kinds of disturbances include non-renewable generation loss or changes in load, and these will have similar impact with only particular details differing. The choice of renewable generation loss is preferred in this study because (1) it is a common concern for which demand response is

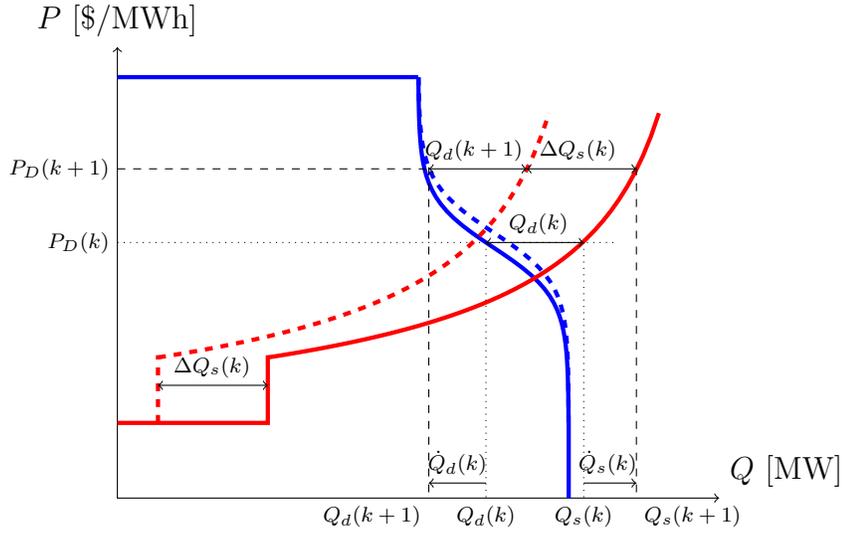


Figure 5.3: Five-minute resource dispatch with supply (red) and demand (blue) response to a loss of generation (ΔQ_s).

frequency response prior to the redispatch operation causes the shape of the demand curve to change slightly and present only the remaining available demand response to the next 5-minute redispatch market. Moving from the market k to the market $k+1$, the clearing price increases from $P_D(k)$ to $P_D(k+1)$ so that the dispatched load changes by $\dot{Q}_d(k) = Q_d(k+1) - Q_d(k) = (1-\alpha)\Delta Q_s(k)$ and the dispatched generation changes by $\dot{Q}_s(k) = Q_s(k+1) - Q_s(k) = -\alpha\Delta Q_s(k)$ where $0 \leq \alpha \leq 1$ to satisfy the physical constraint that $\dot{Q}_s(k) - \dot{Q}_d(k) = -\Delta Q_s(k)$ or $Q_d(k+1) = Q_d(k) = e_s$.

To respond efficiently to a frequency deviation, generation units must change their output by $\dot{Q}_s(k)$ as their contribution to restoring the area's export power to the committed hourly-scheduled level, as shown in Figure 5.4. Concurrently load must change demand by $\dot{Q}_d(k)$ as its contribution to efficiently restoring system frequency. The export power error at the time t is $\Delta Q(t) = Q(t) - Q_d(kt_d) = \Delta Q_s(kt_d) + Q_s(t) - Q_d(t)$. The economically optimal response is that for which the marginal cost of the generation response is equal to the marginal cost of the demand response. We compute the regulation response price, $P(t)$ to quantify the marginal cost of deviations from the hourly schedule in real time:

$$P(t) = P_D(kt_d) + \frac{s(kt_d)}{d(kt_d) - s(kt_d)} \frac{d(kt_d)}{s(kt_d)} \Delta Q(t).$$

often cited as a potential solution, and (2) it provides a clearer illustration of the various effects on transactive system behavior.

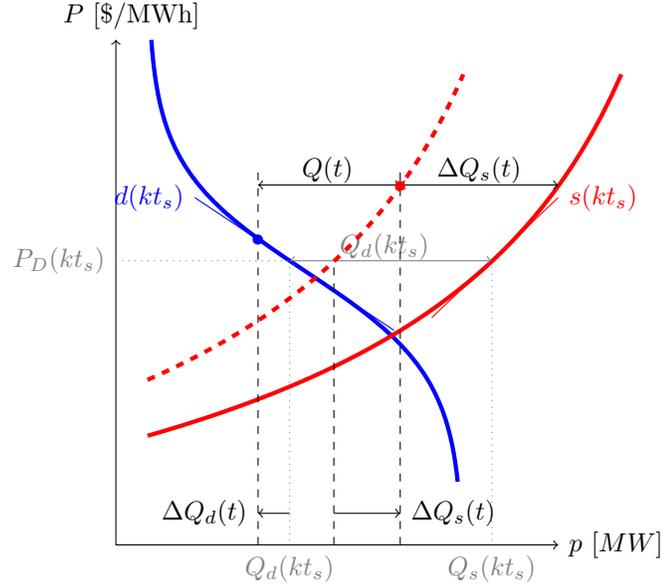


Figure 5.4: Real-time response of generation and load to a disturbance.

where $s(kt_d)$ and $d(kt_d)$ are the slopes of the generation supply and load demand curves at the time of dispatch kt_d , respectively, for the redispatch exports $Q_d(kt_d)$ for the next 5 minutes, and $Q(t)$ is the actual exports at the time t . In the 5-minute dispatch market k , $Q_d(k)$, $s(k) > 0$ and $d(k) < 0$ are updated every 5-minutes.

At the system level a deviation in net power will be associated with a deviation in frequency as well, regardless of whether the power deviation is endogenous to the local control area. For this reason we incorporate two additional cost components, one for the frequency deviation itself and the other for the control response arising from the ACE signal. The net cost taken over the entire system is zero in the sense that the payments made to areas mitigating a deviation are equal to payments by the areas contributing to it.

The total balance of payments is

$$\int_0^{300} P_D[\Delta Q(t) + B\Delta f(t)] - P(t)[A(t) - \Delta Q_s(t)]dt, \quad (5.2)$$

where $P(t)$ is the cost of the over/under-response to the ACE signal $A(t)$ as a result of the disturbance $\Delta Q_s(t)$. The value of $P(t)$ will depend on the mix of generation (e.g., hydro, coal, nuclear, combine cycle gas turbine) that responds to the ACE signal. Any non-zero payments by any party represents a deviation from the surplus-maximizing condition represented by the schedule and therefore represents a loss of

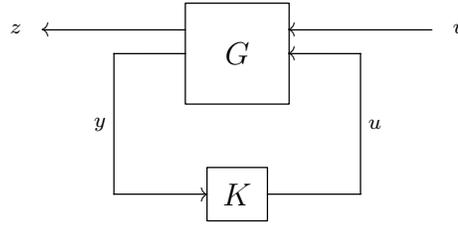


Figure 5.5: Standard system for \mathcal{H}_2 -optimal control synthesis

surplus. Our objective then is to minimize these payments by expressing them as a weighted squared sum of the three cost components, Bf , p , and $P(t)(a - \Delta Q_s)/P_D$. This 2-norm minimization in the wake of a generation contingency of magnitude ΔQ_s is expressed by the objective

$$\min_{A(t)} \int_0^{300} \omega_1 [\Delta f(t)]^2 + \omega_2 \Delta_e^2(t) + \omega_3 (A(t) - \Delta Q_s(t))^2 dt. \quad (5.3)$$

where $\omega_1 = B$, $\omega_2 = 1$ and $\omega_3 = \frac{P(t)}{P_D}$.

When this objective is satisfied, we can be assured that we have also maximized the total surplus given the prevailing conditions: by minimizing the individual payments or receipts on both sides of the balance of payments we have minimized the deviation from the surplus-maximizing schedule and therefore minimized the surplus loss due to regulation.

5.1.4 \mathcal{H}_2 -optimal control policy

We now have the conditions necessary to synthesize the \mathcal{H}_2 -optimal control policy for a control area that minimizes the costs associated with operating the system as it returns to the scheduled set-point, including frequency regulation in the presence of FADR resources. We require the individual component models within the control area used to synthesize the optimal control policy.

We now introduce the state-space solution of the \mathcal{H}_2 -optimal control problem [92]. We consider the standard control system in Figure 5.5 and we partition of the plant G according to

$$\begin{bmatrix} z \\ y \end{bmatrix} = \begin{bmatrix} G_{11} & G_{12} \\ G_{21} & G_{22} \end{bmatrix} \begin{bmatrix} v \\ u \end{bmatrix}.$$

The closed-loop system has the transfer function

$$z = F(G, K)v$$

where $F(G, K) = G_{11} + G_{12}(I - KG_{22})^{-1} + KG_{21}$. The \mathcal{H}_2 -optimal control problem consists of finding the causal controller K that stabilizes the plant G while minimizing the cost function

$$J_2(K) = \|F(G, K)\|_2^2$$

where $\|F(G, K)\|_2$ is the \mathcal{H}_2 norm of the transfer function from v to z .

To obtain this transfer function we begin with non-dispatchable generation, primarily wind and solar resources. These resources do not contribute to either droop or ACE-control responses and have null responses to both frequency and ACE signal. The fraction of non-dispatchable generation in the study area is denoted F_w and for the design case we will use 75% renewable resource penetration to exemplify an extreme situation.

The thermal generating unit response to the filtered area generation control signal a_s is given by the simplified transfer function [90]

$$G_r(s) = \frac{\hat{g}_r(s)}{\hat{a}_s(s)} = \frac{1}{(1 + sT_g)(1 + sT_{ch})}, \quad (5.4a)$$

where \hat{g}_r is the thermal unit's output, T_g is the governor time constant and T_{ch} is the time constant of the steam chest. Typical values for steam turbine units are [90]

$$T_{ch} = 0.3 \text{ sec} \quad \text{and} \quad T_g = 0.2 \text{ sec},$$

which gives the ACE-controlled generation response transfer function

$$G_r(s) = \frac{16.7}{(s + 5.00)(s + 3.33)}. \quad (5.4b)$$

The fraction of units that respond to the area generation control signal is denoted F_r which is set to 25% for the design case. All dispatchable units that do not respond to the area control error are provided with droop response such that [90]

$$G_d(s) = \frac{\hat{g}_d(s)}{\hat{f}(s)} = \frac{1}{R}G_r(s), \quad (5.4c)$$

where \hat{g}_d is the droop-controlled unit's output and $R = -0.05$ is the conventional frequency droop response of generating units. Given that we have selected a design case with the extreme of 75% renewable generation, we expect the number of droop units to be zero and this component is omitted from the design case model.

The filtered area generation control signal a_s is computed by sending the raw ACE signal through a low pass filter F to avoid excessive actuation of the regulating units. For the purposes of this study the values

$$B = 21 \quad \text{and} \quad F(s) = \frac{1}{1 + sT_f}, \quad (5.4d)$$

are used with $T_f = 0.02$ seconds.

Based on data obtained from field tests [48], the conventional grid-friendly control response exhibits two important behaviors. The first behavior is the primary underfrequency event response, which acts like a strong derivative response reaching maximum within a few seconds followed by a very slow recovery using integral error feedback over a period of a minute or more. These are approximated satisfactorily using the demand response transfer function

$$L(s) = \frac{\hat{l}(s)}{\hat{f}(s)} = \frac{F_d K_d s + K_p}{T_l s^2 + s + K_l}, \quad (5.4e)$$

where \hat{l} is the load response, F_d is the fraction of total load that is responsive, and for the design condition $K_d = 1/F_d$ is the fraction of responsive load that is armed by the 5-minute redispatch², and K_p is the quasi-steady state rebound response. The derivative response time constant $T_l = 0.17$ seconds and the recovery time constant $K_l = 0.01$ per second are based on the responses observed in the grid-friendly controllers studied in the Olympic test [48]. This gives us

$$L(s) = \frac{59s}{(s + 5.8)(s + 0.1)}, \quad (5.4f)$$

as the general fast-acting demand response transfer function. The response of the loads is initially very fast and very strong, but it decays within a few minutes, and it is therefore minimally described as a second-order response with derivative control. The rebound response K_p is excluded in this model because it is expected to be

²Note $K_D F_D$ is unity at 5% DR but when F_D is changed K_D is not changed.

addressed by the redispatch operation after a maximum of 5 minutes. The load response is therefore not net-energy neutral over the 5-minute period. This allows us to suppress the non-minimum phase behavior that can emerge from thermostatic loads when their curtailment signal is released and they settle into a higher load condition for a prolonged demand response rebound period. For the controller design condition we use 5% controllable load resources, but the total demand response availability is varied from 0% to 50% for the control robustness analysis below.

The interconnected system's overall response to net power deviations is given by the damped inertial response transfer function

$$H(s) = \frac{\hat{f}(s)}{\hat{Q}(s)} = \frac{1}{Ms + D}, \quad (5.4g)$$

where \hat{Q} is the response of all system generator output power, $D = 1$ is a typical value for the load damping constant, and $M = 6$ is a typical value for the system with somewhat low inertia [90].

Figure 5.6 illustrates the system. Each control area is modeled with: (1) loads L controlled by frequency through the controller K_L/R ; (2) frequency droop generators G_d controlled by the droop gain K_G/R , and (3) ACE-controlled generators using the controller K which we will design. The ACE input to the controller K considers the frequency droop $-1/R$, system damping D , and the export error $E_A - E_S$, while the frequency is obtained from the system inertial response $1/(Ms + D)$. The frequency input to ACE is defined as the bias $B = D - 1/R$.

The interconnected system's open-loop frequency, generation and load responses to a nearly step disturbance are shown in Figure 5.7. For the design condition we have set the demand response control gain to match the generation control gain as expected from the 5-minute market dispatch of the regulation contribution factors.

5.2 Implementation

We can now consider the \mathcal{H}_2 -optimal control design problem [93] for the system shown in Figure 5.6 and arranged in the standard form shown in Figure 5.8. The controller K provides coordinated dispatch of regulation response for generation resources G_r . The controller measures the system frequency f and the control area's net export schedule power deviation p . The current ACE control policy from Eq. (5.1) is the

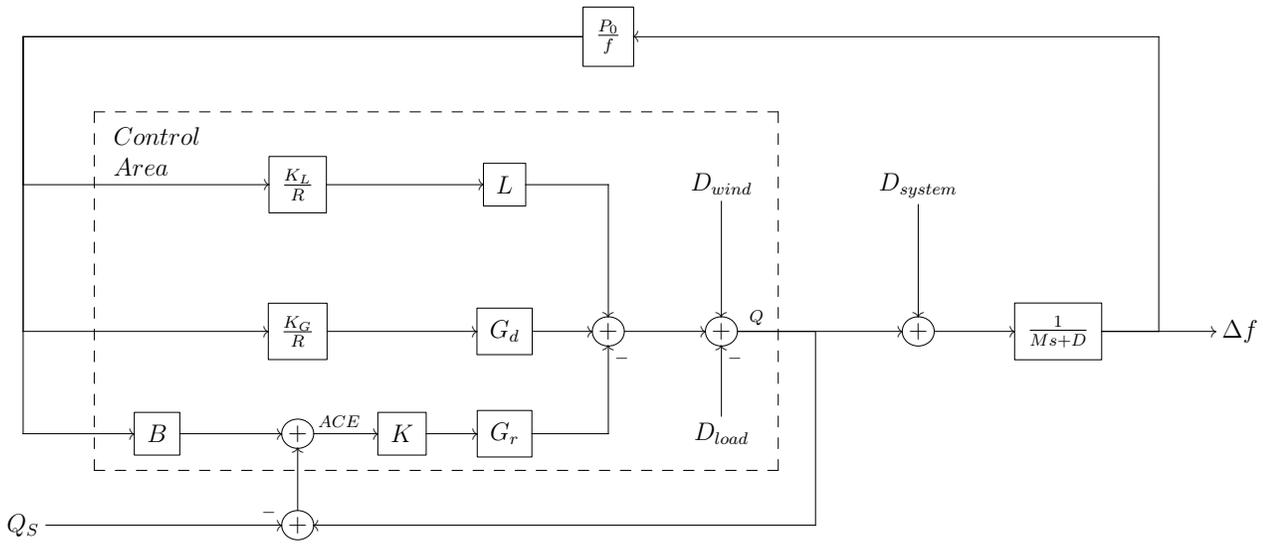


Figure 5.6: System frequency and control area export regulation control diagram.

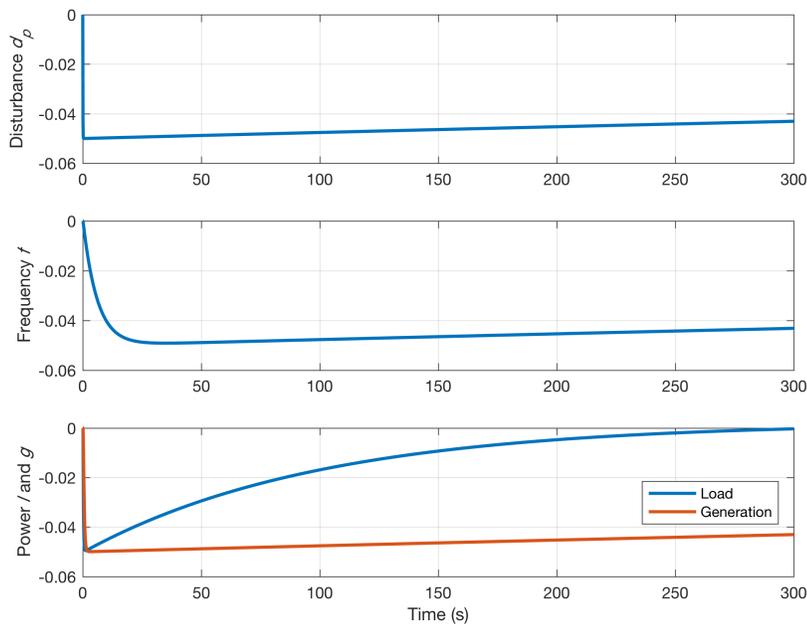


Figure 5.7: Model component frequency (f p.u. nominal frequency), load (l p.u. area load) and generation (g p.u. system load) responses to a local disturbance (ΔQ_s p.u. system load).

baseline control policy for this study. The frequency bias B is computed based on the generation and load characteristics of the control area [90].

A 5% generation local loss input disturbance is modeled as very nearly a step-loss

of generation in the local control area. The input filter for the power disturbance is thus specified as

$$\hat{\Delta}Q_s(s) = \frac{20\Delta Q_s}{s^2 + 20s + 0.01}, \quad (5.5a)$$

where ΔQ_s is the magnitude of the disturbance, which for this study is set at 5% of the total system load. This magnitude disturbance corresponds to a deviation of $\Delta f = \lim_{x \rightarrow \infty} \frac{-d}{D}(1 - e^{-\frac{D}{M}t}) \times f = 3.0$ Hz, which is very significant for a 60 Hz system. This may seem like a large disturbance for a North American system. But it is not atypical for systems in other parts of the world or for microgrids. Demonstrating the effectiveness of transactive control in such systems is useful and therefore a large disturbance is considered. The remaining disturbances are taken as frequency and power measurement noise of magnitude of 1%.

The optimization seeks to minimize regulation deviations from the economically optimal schedule given by the most recent 5-minute dispatch of frequency responsive generation and load resources. Because the maximum surplus is achieved when the dispatch schedule is followed, any deviation from the schedule will reduce the total surplus. We therefore construct the \mathcal{H}_2 control output vector components

$$z(t) = \begin{bmatrix} C_A(t) \\ C_Q(t) \\ C_f(t) \end{bmatrix} = \begin{bmatrix} \frac{P(t)}{P_D} [A(t) - \Delta Q_s(t)] \\ C Q(t) \\ B f(t) \end{bmatrix}, \quad (5.5b)$$

the 2-norm of which we will seek to minimize. The transfer function for the energy cost impact is given as nearly an integrator in the sense that it costs slightly more to provide an early generation response than a late one. The energy cost transfer function is approximated as

$$\hat{c} = \frac{1}{s^2 + 20s + 0.01}. \quad (5.5c)$$

The value $P(t)$ is given in units of \$/MWh and B is given in units of MW/Hz.

The measurement outputs for power and frequency $y(t) = \begin{bmatrix} p(t) \\ f(t) \end{bmatrix}$ are taken directly from the system and the generation+load+disturbance outputs, respectively to which the input disturbance noises are added.

The control input for the raw area control signal is $u(t) = \begin{bmatrix} A(t) \end{bmatrix}$ and will either be the ACE control signal

$$\hat{a} = \frac{1}{s}\hat{Q} + B\hat{f} \quad (5.5d)$$

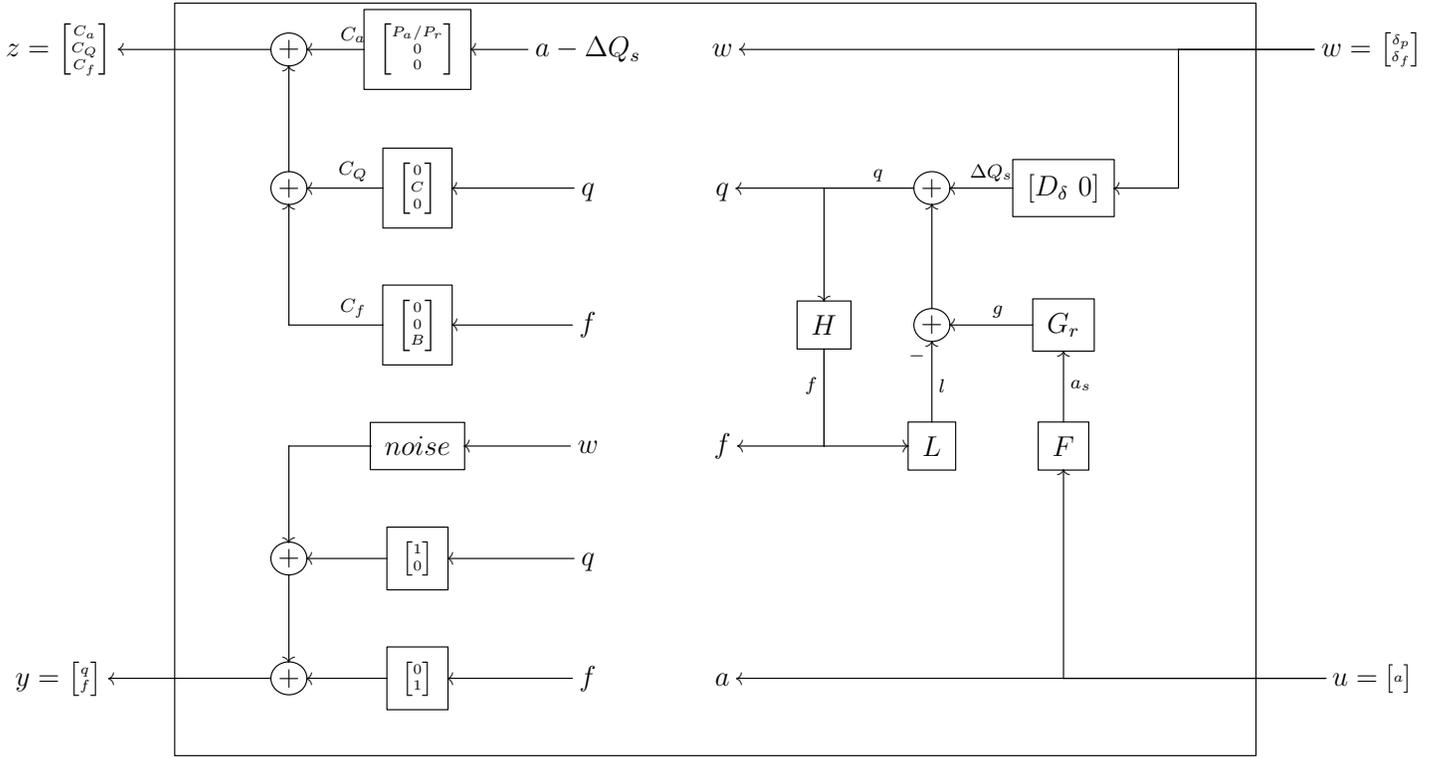


Figure 5.8: Control area model in standard form.

for the baseline model, or the \mathcal{H}_2 -optimal controller output as described below for the study model.

5.2.1 State-space realization

Using the packed matrix notation $G = \left[\begin{array}{c|c} \mathbf{A} & \mathbf{B} \\ \hline \mathbf{C} & \mathbf{D} \end{array} \right] = \mathbf{D} + \mathbf{C}(s\mathbf{I} - \mathbf{A})^{-1}\mathbf{B}$ we obtain the state-space realization for the study model of the control area given by Eq. (5.6a).

$$G = \left[\begin{array}{cccccccccccc|ccc} -0.1667 & 0 & 1.0417 & -1.4706 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0.005 & 0 & 0 \\ 0 & -8.3333 & -4.1667 & 0 & 0 & 0 & 0 & 0 & 0 & 0.1333 & 0 & 0 & 0 & 0 \\ 0 & 4 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.6667 & 0 & 0 & -5.8824 & -0.2353 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.25 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -20 & -0.08 & 0 & 0 & 0 & 0 & 4 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.125 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 8.3333 & -11.7647 & 0 & 0 & -8 & -20 & -0.08 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.125 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -0.0167 & 0 & 0 & 0.0625 \\ \hline 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2 & 0 & 0 & 0 & 0 & 0 \\ 7 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2.0833 & -2.9412 & 0 & 0 & -2 & 0 & 0 & 0 & 0 & 0.01 & 0 & 0 \\ 0.3333 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.01 & 0 \end{array} \right], \quad (5.6a)$$

which yields the synthesized \mathcal{H}_2 -optimal regulation policy for the control area

$$K = \left[\begin{array}{cccccccccccc|cc} -0.1667 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.5 & 0 \\ 0 & -8.3333 & -4.1667 & 0 & 0 & 0 & 0 & 0 & 0 & 0.1333 & 0 & 0 & 0 \\ 0 & 4 & 0 & -0 & 0 & 0 & -0 & 0 & 0 & 0 & 0 & -0 & 0 \\ 0.6667 & 0 & -0 & -5.8824 & -0.2353 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -0 & 0.25 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -547.427 & 772.839 & 0 & -20 & 525.45 & 0 & 0 & 0 & 0 & 262.765 & 0 \\ 0 & 0 & 8.6287 & -12.1817 & 0 & 0.125 & -8.2836 & 0 & 0 & 0 & 0 & -4.1418 & 0 \\ 0 & 0 & 5.4743 & -7.7284 & 0 & 0 & -5.2553 & -20 & -0.08 & 0 & 0 & 1.3723 & 0 \\ 0 & 0 & -0.0863 & 0.1218 & 0 & 0 & 0.0828 & 0.125 & 0 & 0 & 0 & 0.0414 & 0 \\ \hline -0.024 & -0.0082 & -0.0166 & 0.0047 & -0.0667 & 0.0091 & 1.4543 & -0.0006 & -0.1 & -0.0496 & 0 & 0 & 0 \\ -0.3843 & -0.1308 & -0.2652 & 0.0759 & -1.0671 & 0.1453 & 23.2695 & -0.01 & -1.6002 & -0.5265 & 0 & 0 & 0 \end{array} \right], \quad (5.6b)$$

and the corresponding transfer function from power p to area control a

$$\frac{-58.472(s+20)^2(s+5.6947)(s+5)(s+3.3333)(s+0.3625)(s+0.016667)(s+0.0086675)(s+0.0023866)}{(s+24.142)(s+19.999)(s+5.8723)(s+5.2747)(s^2+7.2484s+13.55)(s+0.16667)(s+0.010017)(s+0.0015464)}. \quad (5.6c)$$

The smallest eigenvalue of the closed-loop system GK is -0.0005 and there are conjugate poles at $-3.6242 \pm 0.6443i$. Unlike the ACE control policy, the \mathcal{H}_2 -optimal area control policy does not rely on a frequency input and only requires measurement of net exports from the control area. By way of comparison, the conventional ACE control state-space model is given by

$$K_A = \left[\begin{array}{c|cc} 0 & 1 & 0 \\ \hline 1 & 0 & B \end{array} \right], \quad (5.7)$$

where $B = D - 1/R = 21$, which is used for the baseline model. The smallest eigenvalue of the closed-loop system GK_A is -0.0005 and it has a single pair of complex conjugate poles at $-0.0367 \pm 0.0343j$.

5.2.2 Model Validation

The model of ACE control is verified for varying amounts of demand response under the design conditions, as illustrated in Figure 5.9. The ACE response is adequate for the conditions given insofar as it restores both frequency and power within about 120 seconds of the initial event. The \mathcal{H}_2 -optimal control response fully restores both frequency and power deviations to zero but with significantly less overshoot. Although it does not occur in this particular study, we anticipate that any residual transient error that persists at the end of a 5-minute dispatch interval will be corrected after the next dispatch or scheduling operation.

The area control signal (a) for \mathcal{H}_2 -optimal control is initially faster in its response to the initial event but of lesser magnitude. The power and frequency responses,

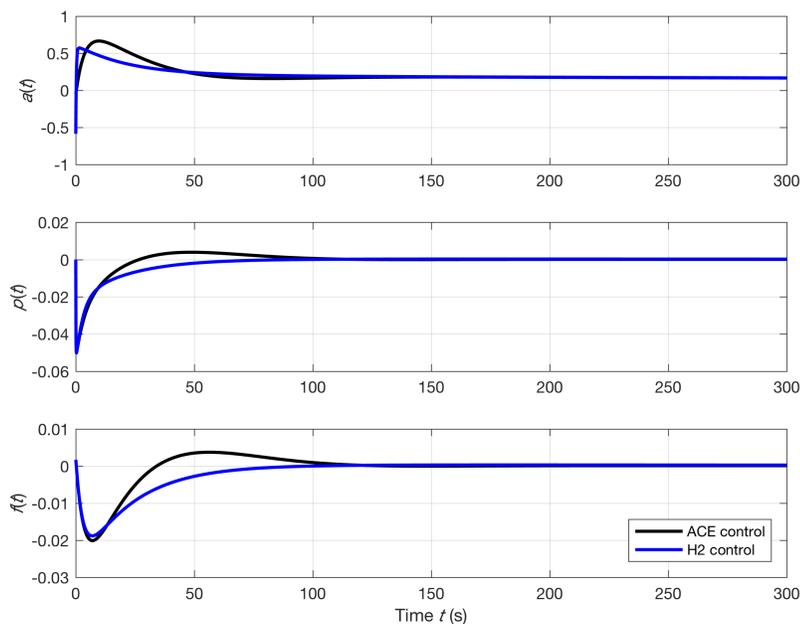


Figure 5.9: ACE control (black) and \mathcal{H}_2 -optimal (blue) control performance for design conditions (5% FADR), showing the raw ACE signal (a p.u. area load), area generation output (p p.u. system load), and system frequency (f p.u. nominal frequency) response to a loss of generation within the control area.

(p and f , respectively) are very similar for the first 10 seconds following the event. However afterwards the power and frequency response to \mathcal{H}_2 -optimal control signal is reduced to minimize costly overshoot.

The steady-state power and frequency deviation for both control policies is zero and both achieve steady-state in approximately the same time. As a result, the area control signal reaches the same steady-state value for both control policies.

5.3 Control Performance

We recognize that demand response resource availability can vary widely from one area to another, from hour to hour, and from season to season. Thus we evaluate the performance of the \mathcal{H}_2 -optimal control policy relative to the conventional ACE control policy by comparing the response of each to widely varying demand response resource availability. The area generation control signal, net power exports and system frequency are compared for a 100 MW load base control area with a nominal energy

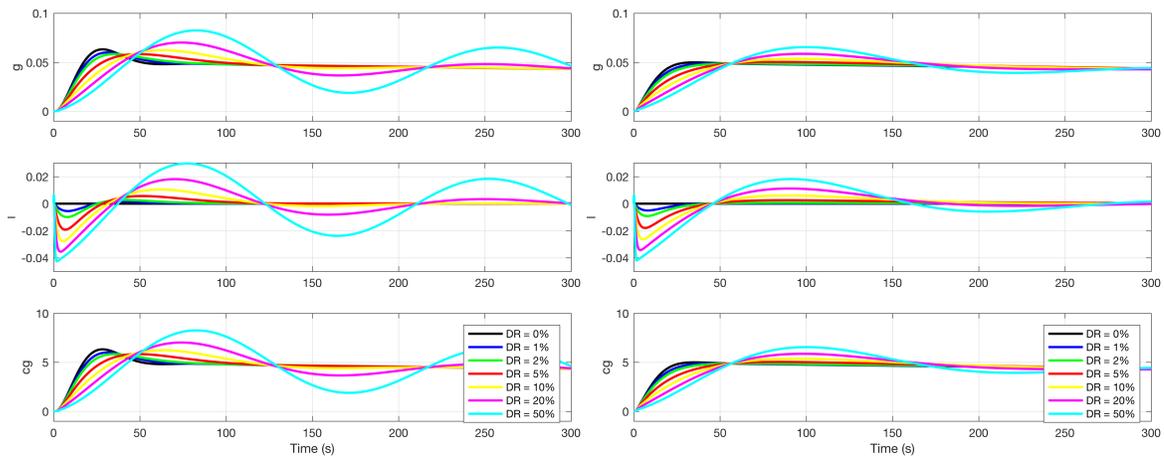


Figure 5.10: ACE control (left) and \mathcal{H}_2 -optimal control (right) model validation for varying demand response level with generation response (g p.u. area load), demand response (l p.u. area load), and generation regulation cost (c_g in $\$/h$ p.u. area load).

price of $\$100/\text{MWh}$. In addition, the cost of regulation and energy used for regulation are compared.

The disturbance response of generation (p per unit area load) and load (l per unit area load) are shown in Figure 5.10 for the conventional ACE signal (left) and \mathcal{H}_2 -optimal control (right). In addition, the generation control cost c_g is shown in units of $\$/h$ per unit area load. We observe decreasing stability of the ACE control policy under higher penetration of fast-acting demand response. This phenomenon is consistent with previously observed results for load controller delays that exceed $1/4$ second [94]. In contrast, the \mathcal{H}_2 -optimal control policy exhibits less oscillation and shorter settling-time performance indicating that it is much less susceptible to overall performance degradation under high demand response scenarios. In every other important respect, and particularly with respect to the steady-state, the \mathcal{H}_2 -optimal control policy is comparable to the ACE control policy.

The comparative costs contribution to the objective function are presented in Figure 5.11. The ACE control policy exhibits significantly more deviations from the schedule, particularly under high demand response conditions and is unable to establish a steady regulation regime under high demand response. In all conditions ranging from no demand response to 50% demand response, the \mathcal{H}_2 -optimal control policy establishing a steady regulation regime that zeros out the deviation of operation from the surplus maximized schedule within about 150-200 seconds.

The cost savings and energy impacts from utilizing \mathcal{H}_2 -optimal control are shown

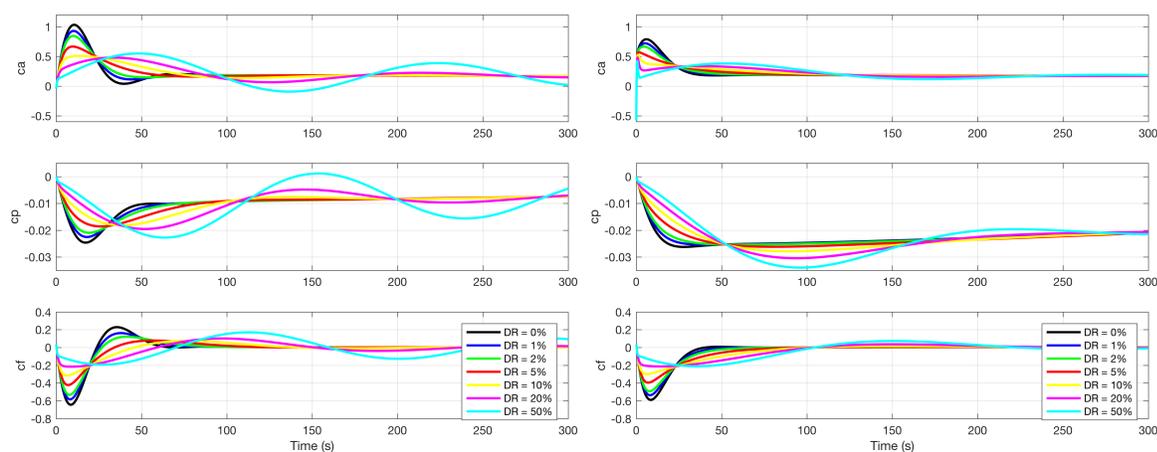


Figure 5.11: ACE control (left) and \mathcal{H}_2 -optimal control (right) cost and dispatch for varying demand response levels, where C_a is the total control cost (in \$), C_Q is the power control response cost (in \$), and C_f is the frequency control response cost (in \$).

Table 5.1: Cost, generation, and net export impacts of ACE control versus \mathcal{H}_2 -optimal control

FADR (%)	Cost			Generation			Exports		
	ACE (\$)	\mathcal{H}_2 (\$)	Saving (%)	ACE (MWh)	\mathcal{H}_2 (MWh)	Reduction (%)	ACE (MWh)	\mathcal{H}_2 (MWh)	Change (%)
0	13750	13443	2.2	13.75	13.44	2.2	13.75	13.75	0.0
1	13734	13389	2.5	13.73	13.39	2.5	13.75	13.79	-0.3
2	13736	13382	2.6	13.74	13.38	2.6	13.75	13.79	-0.3
5	13754	13427	2.4	13.75	13.43	2.4	13.75	13.77	-0.1
10	13761	13469	2.1	13.76	13.47	2.1	13.75	13.74	0.1
20	13822	13451	2.7	13.82	13.45	2.7	13.76	13.82	-0.4
50	14264	13369	6.3	14.26	13.37	6.3	13.81	14.33	-3.7

in Table 5.1. It is noteworthy that in all cases generation and energy costs are reduced, while exports are largely unchanged. We note that at very high FADR levels cost and generation are reduced more significantly while exports are impacted more clearly. This suggests that further study of the system behavior at very high levels of demand response may be desired in the future

5.4 Summary of Results

In this chapter we presented a \mathcal{H}_2 -optimal approach to synthesizing the control policy for control areas in bulk electricity interconnections. The approach is suited to controlling both generation and demand response in areas that have a high penetration of both intermittent renewable resources and fast-acting demand response. The implementation of the \mathcal{H}_2 -optimal control policy is compatible with, and indeed depends on the transactive control 5-minute dispatch strategy such as demonstrated by the Olympic and Columbus studies.

The transactive \mathcal{H}_2 -optimal area control policy is shown to be superior to the conventional ACE control policy in that it is (1) significantly more robust to uncertainty in the amount of fast-acting demand response that is available, (2) always less costly and less energy intensive, and (3) minimizes deviation from any surplus maximizing schedule.

Chapter 6

Dispatch

This chapter proposes a system-wide optimal resource dispatch strategy that enables a shift from primarily energy cost-based approach to primarily ramping cost-based one. This optimal dispatch answers the question of what power schedule to follow during each hour as a function of the marginal prices of energy, power and ramping over the hour¹. The main contributions of this chapter are (1) the derivation of the formal method to compute the optimal sub-hourly power trajectory for a system when the cost of energy and ramping are both of the same order, (2) the development of an optimal resource allocation strategy based on this optimal trajectory, and (3) a simulation method to evaluate the cost savings of choosing the optimal trajectory over the conventional sub-hourly dispatch used in today’s system operation.

In Section 6.2 we develop the optimal control function in both time and frequency domains. In the case of the frequency domain optimal control function the solution is presented as a continuous function. A discrete-time solution suitable for periodic feedback control systems is presented in Section 6.3. In Section 6.4 we examine the performance of this optimal dispatch solution in terms of varying prices for a given “typical” hour, and in Section 6.5 we analyze the cost savings in an interconnection that models the Western Electric Coordinating Council (WECC) system for the year 2024 under both low (13%) and high (50%) renewable generation scenarios. A detailed discussion and synthesis of the consequences that appear to arise from this new paradigm and our perspective on possible future research on this topic are deferred Chapter 7.

¹The marginal price of a product or service is the change in its price when the quantity produced or delivered is increased by one unit.

6.1 Background

The conventional utility approach to addressing renewable generation variability is to allocate additional firm generation resources to replace all potentially non-firm renewables resources. This approach can have significant financial impacts in reducing the revenue from energy supply while increasing the costs of providing ramping response to variable generation. This chapter examines an approach to mitigating these financial impacts.

Firm resources are typically fast-responding thermal fossil resources or hydro resources when and where available. For new renewable resources the impact of this approach is quantified as an intermittency factor, which discounts the contribution of wind in addition to its capacity factor and limits the degree to which they can contribute to meeting peak demand [13]. However, the intermittency factor does not account for the ramping requirements created by potentially fast-changing renewable resources [14]. The need for fast-ramping resources discourages the dispatch of high-efficiency fossil and nuclear generation assets and can encourage reliance on low-efficiency fossil-fuel resources for regulation services and reserves [15].

One solution to overcoming the renewable generation variability at the bulk electric level is to tie together a number of electric control areas into a super-grid so that they can share generation and reserve units through optimal scheduling of system inertias [1]. In an interconnected system, the combined power fluctuations are smaller than the sum of the variations in individual control areas. Furthermore, fast-acting energy storage systems and demand response programs can provide required ancillary services such as real-time power balancing [95] and frequency regulation [96] if they are equipped with suitable control mechanisms. A competitive market framework in which energy resources participate to sell and buy ancillary service products can accelerate the transition to a high-renewable scenario by supporting the long-term economic sustainability of flexible resources.

Concerns about the financial sustainability of utilities under high level of renewables are also beginning to arise. The question is particularly challenging when one seeks solutions that explicitly maximize social welfare rather than simply minimizing production cost [97]. The growth of low-marginal cost renewable resources can lead one to expect utility revenues to decline to the point where they can no longer recover their long term costs. But this conclusion may be erroneous if one fails to consider both the impact of demand own-price elasticity, as well as the impact of load control

automation on substitution elasticity. The latter type of demand response can significantly increase the total ramping resource on peak and decrease ramping resource scarcity. One option for replacing energy resource scarcity rent is increasing fixed payments. But this may lead to economic inefficiencies as well as an unraveling of the market-based mechanisms built so far. Another option is to enable payments based on ramping resource scarcity rent through existing markets for ancillary services. At the present time, the majority of resources continue to be dispatched based on the energy marginal cost merit order. But it is not unreasonable to consider how one might operate a system in which the marginal cost of energy is near zero and resources are dispatched instead according to the ramping cost merit order.

In the presence of high levels of variable generation, the scheduling problem is a co-optimization for allocating energy and ramping resources [98]. Under existing energy deregulation policies, there is usually a market in which energy producers compete to sell energy, and a separate market in which they compete to sell power ramping resources for flexibility. Producers get paid for their energy deliveries in the energy market and for power ramping flexibility in the flexibility market. But today's dual-pricing mechanism is dominated by the energy markets, which drives generation resources to secure revenue primarily in the energy market, and only deliver residual ramping resources in the flexibility market. Meanwhile poor access to energy markets leads loads and storage to seek participation primarily in the flexibility market while only revealing their elasticities to the energy market. This relegates loads and storage to only a marginal role in the overall operation of the system, which is the motivation for seeking policy solutions to improving their access to wholesale energy markets, such as FERC Orders 745 and 755.

6.2 Methodology

Consider a utility's cost minimization problem over a time interval T . The utility's customers purchase their net energy use $E(T)$ at a pre-determined retail price. So in today's systems, profit maximization and cost minimization are essentially the same problem. For each hour the utility pays for energy delivered at a real-time location-dependent wholesale price that is also dependent on demand under typical deregulated nodal pricing markets. The utility's scheduled energy use is forecast for each hour based on its customers' expected net energy use, which is then used to compute the utility's net load over that hour. We assume that over any interval T

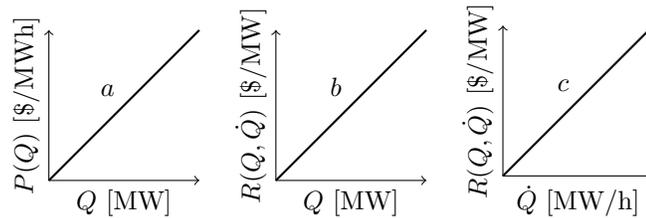


Figure 6.1: Power (left) and ramp (center and right) price functions.

the utility may incur additional costs for any deviation in actual net load from the scheduled load.

The price function at the operating point is split up into the marginal price of energy $a = \frac{\partial P}{\partial Q}$ (measured in $\$/\text{MW}^2 \cdot \text{h}$), the marginal price of power $b = \frac{\partial R}{\partial Q}$ (measured in $\$/\text{MW}^2$), and the marginal price of ramping $c = \frac{\partial R}{\partial \dot{Q}}$ (measured in $\$ \cdot \text{h}/\text{MW}^2$). In order to reflect resource scarcity all cost functions are assumed to be quadratic so that the price function for each is linear as shown in Figure 6.1. The marginal prices a and b determine prices as a function of the power demand Q , and the marginal price c determines prices based on the ramp rates \dot{Q} . The cost parameters arise from the schedule and may vary from hour to hour, but do not change within any given hour. Any of the marginal prices may be zero or positive depending on the market design and prevailing conditions in the system. For the purposes of this chapter, we will assume that they cannot be negative.

Over the time interval T the total cost of both the power trajectory $Q(t)$ and the ramping trajectory $\dot{Q}(t)$ given the power price $P(t) = aQ(t)$ and ramp price $R(t) = bQ(t) + c\dot{Q}(t)$, respectively, is given by

$$C(T) = \int_0^T P[Q(t)]Q(t) + R[Q(t), \dot{Q}(t)]\dot{Q}(t)dt. \quad (6.1)$$

Given the dispatch from $Q(0)$ to $Q(T)$ and the scheduled energy use $E(T) = \int_0^T Q(t)dt$ we augment the cost function with the Lagrange multiplier λ so that we have

$$\begin{aligned} \int_0^T a(Q - Q_z)Q + b(Q - Q_z)|\dot{Q}| + c\dot{Q}^2 + \lambda Q dt \\ = \int_0^T G(t, Q, \dot{Q})dt, \end{aligned}$$

where the $|\dot{Q}|$ represents the magnitude of the ramp rate \dot{Q} , and Q_z is the amount

of must-take generation having zero or effectively zero marginal energy cost. Then the optimal dispatch trajectory $Q(t)$ is the critical function obtained by solving the Euler-Lagrange equation

$$\frac{\partial G}{\partial Q} - \frac{d}{dt} \frac{\partial G}{\partial \dot{Q}} = 0.$$

From this we form a second-order ordinary differential equation describing the critical load trajectory

$$\ddot{Q} - \frac{a}{c}Q = \frac{\mu}{2c}.$$

where $\mu = \lambda - aQ_Z$. Using the Laplace transform we find the critical system response in s -domain is

$$\hat{Q}(s) = \frac{Q_0 s^2 + \dot{Q}_0 s + \frac{\mu}{2c}}{s(s^2 - \omega^2)}, \quad (6.2)$$

where $\omega^2 = \frac{a}{c}$. The general time-domain solution for the critical function over the interval $0 \leq t < T$ is

$$Q(t) = \left(Q_0 + \frac{\mu}{2a}\right) \cosh \omega t + \frac{\dot{Q}_0}{\omega} \sinh \omega t - \frac{\mu}{2a}, \quad (6.3)$$

where Q_0 and \dot{Q}_0 are initial power and ramp values.

We can determine whether this solution is an extremum by computing the second variation

$$\begin{aligned} \frac{\partial^2 C}{\partial Q^2}(T) &= \int_0^T [\alpha(v)^2 + 2\beta(vv') + \gamma(v')^2] dt \\ &= \int_0^T H(t) dt, \end{aligned}$$

with $H(t) > 0$ for all $v \neq 0$ subject to $v(0) = 0 = v(T)$. We then have

$$\alpha = \frac{\partial^2 G}{\partial Q^2} = 2a, \quad \beta = \frac{\partial^2 G}{\partial Q \partial \dot{Q}} = b, \quad \gamma = \frac{\partial^2 G}{\partial \dot{Q}^2} = 2c.$$

Thus for all $a, b, c > 0$, $H(t) > 0$ and $Q(t)$ is a minimizer. Since the only physical meaningful non-zero values of a and c are positive, this is satisfactory. We will examine

cases when a and c are zero separately. Note that when $\dot{Q} < 0$, we have $b < 0$, so that the sign of b does not affect the general solution.

Given the constraints $\int_0^T Q(t)dt = E_T$ and $Q(T) = Q_T$, which come from the hour-ahead schedule, we obtain the solution for μ and \dot{Q}_0 for the case where $a, c > 0$:

$$\begin{bmatrix} \mu \\ \dot{Q}_0 \end{bmatrix} = \begin{bmatrix} A & B \\ C & D \end{bmatrix}^{-1} \begin{bmatrix} E_\Delta \\ Q_\Delta \end{bmatrix}, \quad (6.4)$$

where

$$\begin{aligned} A &= \frac{\sinh \omega T - \omega T}{2a\omega} & B &= \frac{\cosh \omega T - 1}{\omega^2} \\ C &= \frac{\cosh \omega T - 1}{2a} & D &= \frac{\sinh \omega T}{\omega} \\ E_\Delta &= E_T - \frac{\sinh \omega T}{\omega} Q_0 & Q_\Delta &= Q_T - Q_0 \cosh \omega T. \end{aligned}$$

When $a = 0$, the cost of energy is zero and only the ramping cost is considered. Then the time-domain solution is

$$Q(t) = \frac{\mu}{4c}t^2 + \dot{Q}_0 t + Q_0, \quad (6.5)$$

with

$$\begin{aligned} A &= \frac{T^3}{12c} & B &= \frac{T^2}{2} \\ C &= \frac{T^2}{4c} & D &= T \\ E_\Delta &= E_T - Q_0 T & Q_\Delta &= Q_T - Q_0, \end{aligned}$$

which gives the critical response in s -domain

$$\hat{Q}(s) = \frac{\mu}{4cs^3} + \frac{\dot{Q}_0}{s^2} + \frac{Q_0}{s}.$$

When $c = 0$, there is no scarcity for ramping so that the ramping price is based only on the marginal energy cost of additional units that are dispatched. Then we have the time-domain solution

$$Q(t) = -\frac{\mu}{2a}, \quad (6.6)$$

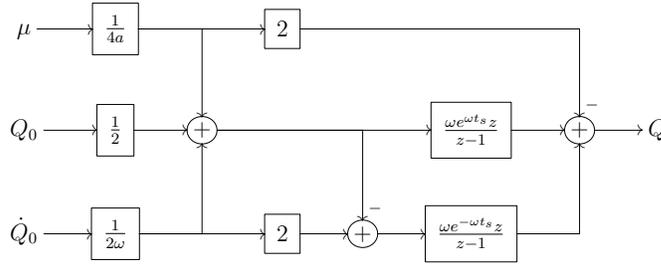


Figure 6.2: Optimal dispatch controller diagram with discrete update time t_s .

with

$$\mu = -\frac{2aE_T}{T}.$$

This gives the critical response in s -domain

$$\hat{Q}(s) = -\frac{\mu}{2as},$$

and the initial and final ramps from $Q(0)$ to $-\frac{\mu}{2a}$ and from $-\frac{\mu}{2a}$ to $Q(T)$ are limited by the ramping limits of the responding units.

6.3 Optimal Dispatch Controller

The partial fraction expansion of Eq. 6.2 is

$$\frac{K_1}{s + \omega} + \frac{K_2}{s} + \frac{K_3}{s - \omega}, \quad (6.7)$$

where $K_1 = \frac{Q_0}{2} - \frac{\dot{Q}_0}{2\omega} + \frac{\mu}{4a}$, $K_2 = -\frac{\mu}{2a}$, and $K_3 = \frac{Q_0}{2} + \frac{\dot{Q}_0}{2\omega} + \frac{\mu}{4a}$, with the values of the parameters are computed from Eq. 6.4.

The initial response of the optimal controller is dominated by the forward-time solution

$$K_1 e^{-\omega t} = \mathcal{L}^{-1} \left[\frac{\frac{Q_0}{2} - \frac{\dot{Q}_0}{2\omega} + \frac{\mu}{4a}}{s + \omega} \right] (s),$$

which handles the transition from the initial system load Q_0 to the scheduled load

$Q_E = -\frac{\mu}{2as}$. The central response is dominated by the scheduled load solution

$$K_2 = \mathcal{L}^{-1} \left[-\frac{\mu}{2as} \right] (s).$$

Finally, the terminal response is dominated by the reverse-time solution

$$K_3 e^{\omega t} = \mathcal{L}^{-1} \left[\frac{\frac{Q_0}{2} + \frac{\dot{Q}_0}{2\omega} + \frac{\mu}{4a}}{s - \omega} \right] (s),$$

which handles the transition from the scheduled load to the terminal load Q_T . A discrete-time controller that implements the solution of Eq. 6.7 is shown in Figure 6.2. The controller implements the three main components to the optimal solution with step inputs μ , Q_0 , and \dot{Q}_0 . Note that the marginal prices a , b and c for the entire hour are constants in the controller blocks, which makes the controller design linear time-invariant within each hour, but time-variant over multiple hours. The discrete-time solution is then

$$Q^*(k) = \begin{cases} K_1 \tau^k + K_2 + K_3 \tau^{-k} & : a > 0, c > 0 \\ \frac{\mu}{4c} t_s^2 k^2 + \dot{Q}_0 t_s k + Q_0 & : a = 0, c > 0 \\ -\frac{\mu}{2a} & : a > 0, c = 0 \end{cases}$$

where $\tau = e^{\omega t_s}$.

The discrete-time dispatch control is illustrated in Figure 6.3 for various values of $\omega = \sqrt{a/c}$. When the value of ω is large, the optimal dispatch is dominated by the energy cost and the cost of high ramp rates is negligible compared to the energy cost. The result is a dispatch that moves as quickly as possible to scheduled load Q_E . In the limit of zero ramping cost, the optimal response is a step function². As the cost of ramping increases relative to the energy cost, the optimal dispatch begins to reduce the ramp rate while still following a trajectory that satisfies the hourly energy delivery requirement. In the limit of zero energy cost, the optimal dispatch trajectory is parabolic.

²Step responses are only possible by generation or load tripping, which is not considered as part of the conventional control strategy.

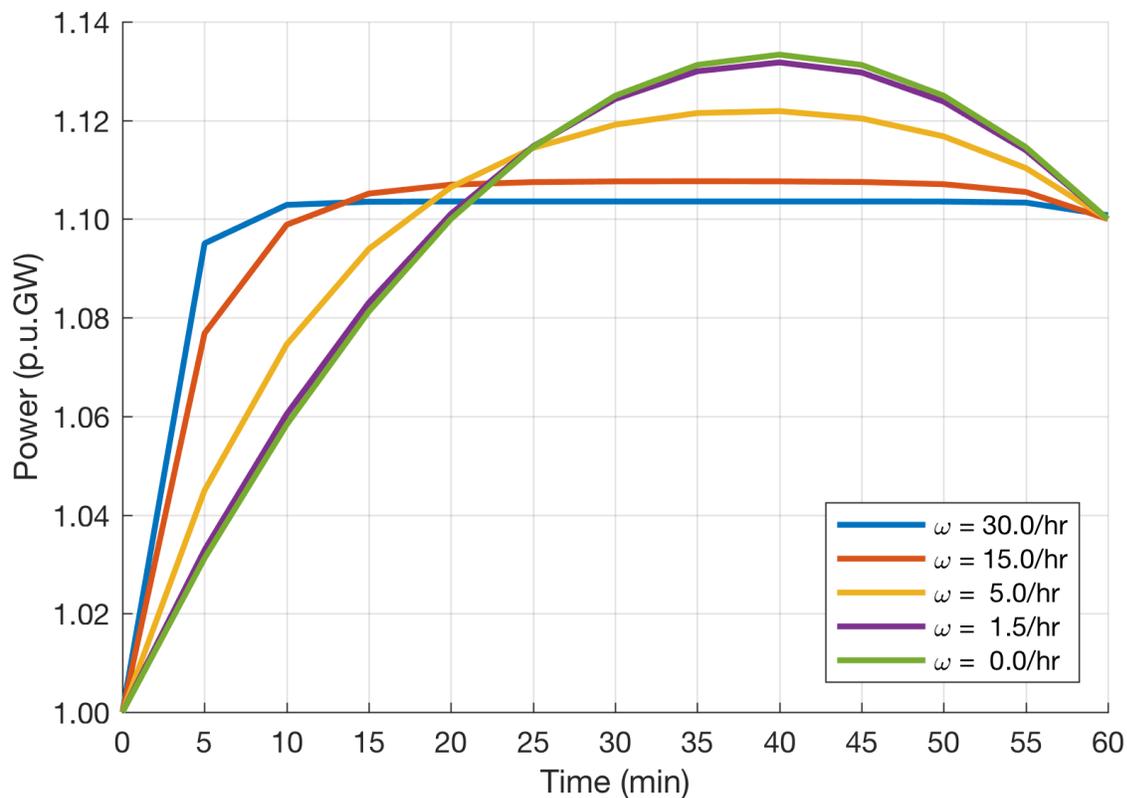


Figure 6.3: Optimal discrete time control for various values of ω at $t_s = 5$ minutes.

6.4 Performance Evaluation

In this section we develop the cost performance metric of the optimal dispatch control design. The optimal dispatch cost function is found by evaluating Equation 6.1 using

Equations 6.3, 6.5 and 6.6. Thus when $a, b, c > 0$ we have³

$$\begin{aligned}
C(T) &= \frac{\sinh 2\omega T}{2\omega} [a(A^2 + B^2) + bAB\omega] \\
&+ \frac{\sinh^2 \omega T}{2\omega} [b(A^2 + B^2)\omega + 4aAB] \\
&+ \frac{\cosh \omega T - 1}{\omega} [(bA\omega + 2aB)C - (aB + bA\omega)Q_z] \\
&+ \frac{\sinh \omega T}{\omega} [(bB\omega + 2aA)C - (aA + bB\omega)Q_z] \\
&+ [aC^2 - aCQ_z] T.
\end{aligned}$$

where $A = Q_0 + \mu/2a$, $B = \dot{Q}_0/\omega$ and $C = -\mu/2a$. For the case when $a = 0$ we have

$$C(T) = \frac{b\mu^2}{32c^2} T^4 + \frac{3b\mu\dot{Q}_0 + \mu^2}{12c} T^3 + \left(\frac{b\dot{Q}_0^2}{2} + \frac{b\mu(Q_0 - Q_z)}{4c} + \frac{\mu\dot{Q}_0}{8c^2} \right) T^2 + \left(b\dot{Q}_0(Q_0 - Q_z) + c\dot{Q}_0^2 \right) T.$$

When $c = 0$ we have

$$C(T) = aE_T \left(\frac{E_T}{T} - Q_z \right).$$

We use as the base case a conventional unit dispatch strategy that requires generators ramp to their new operating point during the 20 minutes spanning the top of the hour. Accordingly the generators begin ramping 10 minutes before the hour and end ramping 10 minutes after the hour. In the aggregate for a given hour this strategy is illustrated in Figure 6.4 where

$$Q_E = \frac{6}{5} \left(E_T - \frac{Q_0 + Q_T}{12} \right),$$

with the initial and terminal ramp rates

$$\dot{Q}_0 = 6(Q_E - Q_0) \quad \text{and} \quad \dot{Q}_T = 6(Q_T - Q_E).$$

Three cases are shown: overproduction to compensate for a lack of generation in previous hours (top), scheduled production (center), and underproduction to compensate for extra generation in previous hours (bottom).

³Note that if the ramp rate \dot{Q} changes sign at the time $t_c = \frac{1}{\omega} \tanh^{-1}(-\frac{B}{A})$ and $0 < t_c < T$, then we must divide the cost integral into two parts to account for the absolute value of \dot{Q} on b terms.

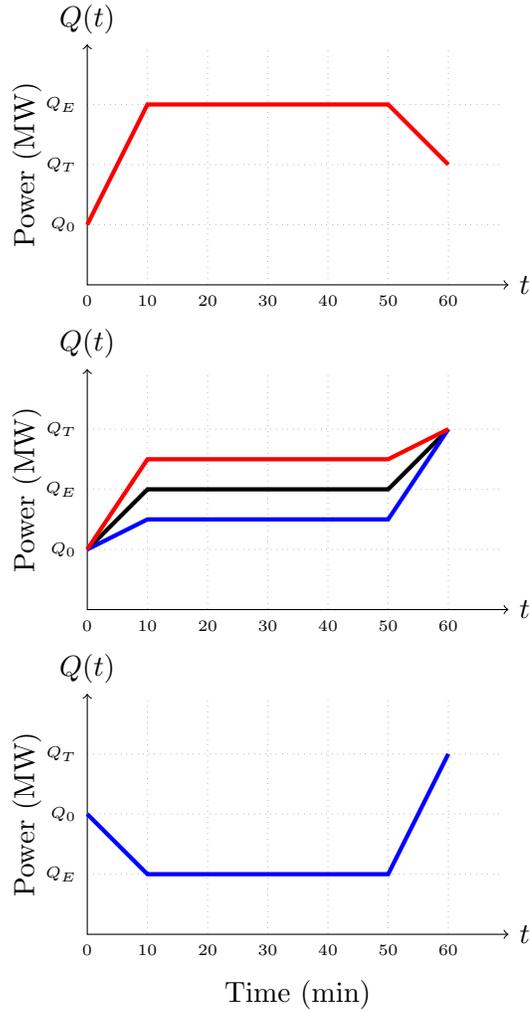


Figure 6.4: Conventional power dispatch for base case: (top) significant negative schedule error requiring over-production, (center) small negative, zero and positive schedule error requiring over (red), normal (black) and under (blue) production, and (bottom) significant positive schedule error requiring under-production.

The cost of the base case is then

$$\begin{aligned}
C_{base}(T) &= \frac{aT}{18}(Q_T^2 + Q_T Q_E + 14Q_E^2 + Q_E Q_0 + Q_0^2) \\
&\quad - \frac{aT}{12}(Q_T + 10Q_E + Q_0)Q_z \\
&\quad + \left| \frac{b}{2}(Q_E - Q_0) \right| (Q_E + Q_0 - 2Q_z) \\
&\quad + \left| \frac{b}{2}(Q_T - Q_E) \right| (Q_T + Q_E - 2Q_z) \\
&\quad + \frac{6c}{T}(Q_T^2 - 2Q_T Q_E + 2Q_E^2 - 2Q_0 Q_E + Q_0^2).
\end{aligned}$$

The zero-order hold ramp discrete form of Equation 6.1 gives us the cost of operating with a discrete control time-step t_s , i.e.,

$$\begin{aligned}
C^*(T) &= \sum_{k=0}^{T/t_s} \left(P^*[Q^*(k)]Q^*(k) + R^*[Q^*(k), \dot{Q}^*(k)]\dot{Q}^*(k) \right) t_s \\
&= \sum_{k=0}^{T/t_s} \frac{at_s}{4} \left[Q^*(k)^2 + 2Q^*(k)\dot{Q}^*(k) \right. \\
&\quad \left. + \dot{Q}^*(k)^2 - 2Q_z[Q^*(k) + \dot{Q}^*(k)] \right] \\
&\quad + \frac{1}{2} \left[\left| b(\dot{Q}^*(k) - Q^*(k)) \right| (\dot{Q}^*(k) + Q^*(k) - 2Q_z) \right] \\
&\quad + \frac{c}{t_s} \left[Q^*(k)^2 - 2Q^*(k)\dot{Q}^*(k) + \dot{Q}^*(k)^2 \right]
\end{aligned}$$

where $Q^*(k) = Q(kt_s)$ and $\dot{Q}^*(k) = Q^*(k+1)$. We evaluate the performance of the control strategy for different control update rates t_s using two future scenarios, one for low renewables where $\omega > 1$, and one for high renewables where $\omega < 1$ for both unconstrained and constrained transmission operating conditions.

6.5 Case Study: WECC 2024

In this section we examine the cost savings associated with using the optimal control solution on the WECC 2024 base case model introduced in [1]. The WECC 2024 model is a 20-area base case used by WECC for planning studies. The 20-area model combines a number of smaller control areas based on the anticipated intertie transfer limits reported in the WECC 2024 common case [99]. In this model constraints within control areas are ignored, while internal losses are approximated. The peak load, annual energy production and demand consumption are forecast, including intermittent wind, solar, and run-of-river hydro for the entire year. The demand characteristics, generation capacity, production costs and tieline limits are given in Tables C.2 through C.6.

The model also includes a hypothetical market for each consolidated control area, with a flat zero-cost supply curve for all renewable and must-take generation resources and a constant positive supply curve slope for all dispatchable units. The hourly generation of intermittent resources is provided by the base case model and incorporated into the supply curve so that there is effectively no marginal cost of

production for renewable energy and must take generation. All generating units are paid the hourly clearing price, and when the marginal energy price in a control area is zero then renewable generation may be curtailed. As a result, under the high renewable scenario, zero energy prices are commonplace and renewable generation is curtailed more frequently. Demand response is similarly considered for each control area and the output of this scheduling model provides the hourly nodal prices required to satisfy the transmission constraints, if any.

The low renewables case is the WECC forecast for the year 2024, which correspond to 29.5 GW (16.1%) of renewable capacity and 140.8 TWh (13.4%) of annual renewable generation. The high renewables case is given as 400% of capacity of the WECC forecast for the year 2024, which corresponds to 117.8 GW (63.5%) and 523.9 TWh (49.6%) respectively. The blended energy price of operations is \$130.6/MWh and \$50.2/MWh for the low and high renewables cases, respectively.

The ramping price was not considered in the WECC 2024 base case model. For this study we have assumed that the ramping energy cost is based on the marginal energy cost for the dispatchable generation and the demand response, as well as the cost of changing the dispatchable generation output, as shown in Table 6.1. In both cases, the marginal price of power b is the average marginal price of energy a over the hour. In the low renewables case the marginal price of ramping c is the marginal price of power b multiplied by 12 seconds. In the high renewables case, c is the marginal price of power b multiplied by 49 hours. The value of ω is approximately 121 times greater in the low renewable case than it is in the high renewable case. Note that a is zero when renewables are curtailed while b is assumed to also be zero because curtailed renewables and demand response are presumed to be dispatchable.

The values of the ramping response constant c were also selected such that the overall cost of operating the system remains more or less constant when going from the low to high renewables scenarios under the base case. This allows us to evaluate the impact of the optimal control strategy without involving the question of revenue adequacy under the high renewables scenario. Given that there are few markets from which to determine these values, we must be satisfied with this assumption for now.

The statistical nature of the intermittency and load forecast errors and their connection to load following and regulation was studied at length in [100]. The authors showed that consolidated control of WECC could yield both cost savings and performance improvements. In particular, the study showed that with high accuracy control 1% standard deviation in load forecast was expected, with 0% real-time mean

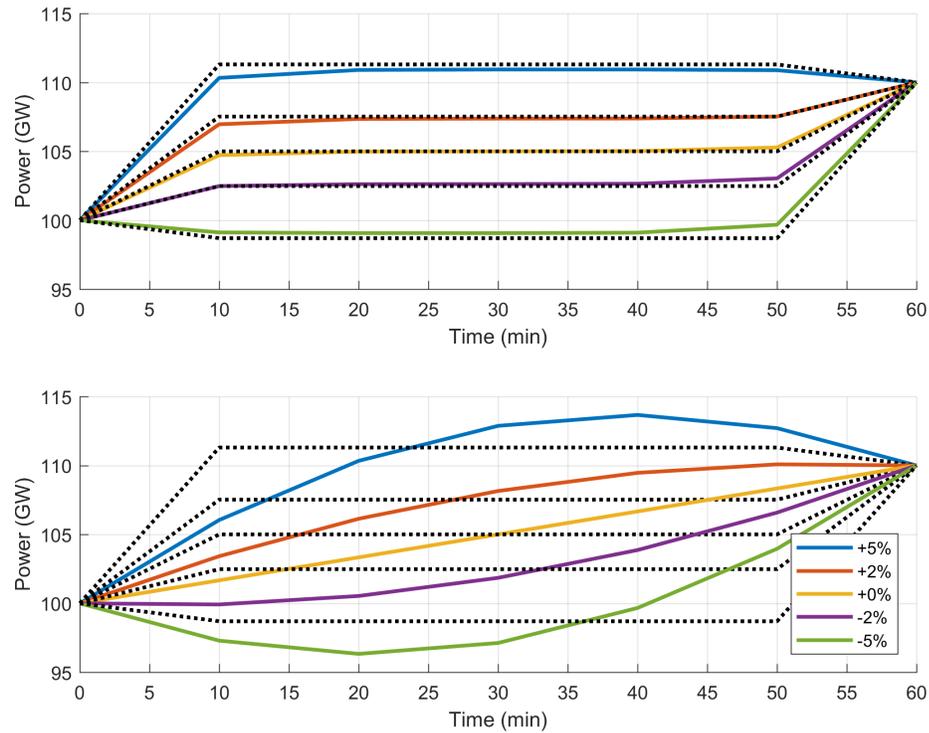


Figure 6.5: Single hour optimal dispatch for low (top) and high (bottom) renewables with a ramp from 100 GW to 110 GW using a 10-minute discrete-time dispatch control rate, with hourly energy schedule correction errors varying from -5% to $+5\%$.

error at 0.15% standard deviation at peak load. However, for the purposes of a preliminary study like the one presented in this chapter, we will consider the scheduling error to be Gaussian with a mean error of 0 MW and a standard deviation of 100 MW. We believe that energy and flexibility markets should be efficient enough to remove all systematic error from the price signals leaving only the random noise that is satisfactorily modeled by Gaussian noise.

The comparison of the conventional and optimal dispatch for a typical case is shown in Figure 6.5. The conventional control strategy is shown in dotted lines, with the 10 minute optimal-dispatch trajectory shown as solid lines. Note that the ramp rate is constant between discrete control updates. The evaluation is completed with the marginal prices and marginal costs at 100 GW, as shown in Table 6.1. The energy schedule changes according to a varying energy error remaining at the end of the previous dispatch interval. A -5% error represents an energy deficit of 5 GWh for a 105 GWh schedule, while a $+5\%$ error represents an energy surplus of 5 GWh.

The marginal prices in Table 6.1 are chosen to satisfy the following conditions:

Table 6.1: Marginal prices and marginal costs for 105 GWh schedule at 100 GW initial power and 10 GW/h ramp for cases in Figure 6.5

Variable	Base case	Study case	Units
<i>Marginal prices:</i>			
a	1.27×10^{-3}	6.34×10^{-4}	\$/MW ² ·h
b	1.27×10^{-3}	6.34×10^{-4}	\$/MW ²
c	4.23×10^{-6}	3.09×10^{-2}	\$.h/MW ²
<i>Marginal costs:</i>			
P	133.09	66.55	\$/MW·h
R	133.13	375.19	\$/MW
ω	17.3	0.1433	h ⁻¹

Table 6.2: Single hour cost savings under low and high renewable for a ramp from 100 GW to 110 GW at 5 minute discrete dispatch control update rate, with varying energy error redispatch

Dispatch Energy (GWh)	Reference Cost				Optimal Cost				Cost Savings		Dispatch Error
	Energy (\$M)	Ramp (\$M)	Total (\$M)	Price (\$/MWh)	Energy (\$M)	Ramp (\$M)	Total (\$M)	Price (\$/MWh)	(\$M/h)	(%)	(%)
<i>Low renewable scenario</i>											
110.3	10.8	1.2	12.0	108.62	10.8	1.1	11.9	107.93	0.1	0.6	-0.8
107.1	10.1	0.9	11.0	102.50	10.1	0.9	11.0	102.50	0.0	0.0	-0.4
105.0	9.6	0.9	10.5	100.00	9.6	0.9	10.5	99.98	0.0	0.0	0.0
102.9	9.1	0.9	10.0	97.53	9.1	0.9	10.0	97.51	0.0	0.0	0.4
99.8	8.4	1.1	9.6	96.06	8.4	1.1	9.5	95.42	0.1	0.7	0.9
<i>High renewable scenario</i>											
110.3	1.6	24.1	25.7	232.77	1.6	13.5	15.1	136.75	10.6	41.3	-2.0
107.1	1.3	11.7	13.0	121.55	1.3	4.8	6.1	57.27	6.9	52.9	-0.8
105.0	1.1	9.4	10.5	100.00	1.1	3.2	4.3	41.25	6.2	58.8	-0.0
102.9	1.0	11.7	12.7	123.35	1.0	4.8	5.8	56.43	6.9	54.2	-0.7
99.8	0.7	24.1	24.8	248.95	0.8	13.4	14.2	142.37	10.6	42.8	-1.9

1. The system operating cost is roughly \$100/MWh at a system load of 100 GW.
2. For the low renewables case, the energy cost is roughly 10 times the ramping cost, while for the high renewables case the ramping cost is roughly 10 times the energy cost for the nominal schedule. This was necessary to ensure that costs were the same for both cases.
3. The marginal power price b for both cases is equal to the marginal energy price a of the respective case.

We considered the performance degradation resulting from longer dispatch in-

Table 6.3: WECC 2024 cost savings from optimal dispatch under different transmission constraint and renewable scenarios

Scenario Model	Cost		
	Base case (\$B/y)	Optimal (\$B/y)	Savings (\$B/y)
<i>Unconstrained:</i>			
Low	126.0	125.9	0.16 (0.1%)
High	108.6	77.8	30.85 (28.4%)
<i>Constrained:</i>			
Low	184.4	184.1	0.26 (0.1%)
High	388.3	231.2	157.12 (40.5%)

Table 6.4: Summary of energy and price impacts of optimal dispatch control for the WECC 2024 base case

Scenario Model	Total Energy (TWh)	Price (\$/MWh)	
		Base case	Optimal
<i>Unconstrained:</i>			
Low	1054.6	119.5	119.35 (-0.1%)
High	1067.2	101.8	67.29 (-51.2%)
<i>Constrained:</i>			
Low	1054.5	174.8	174.55 (-0.2%)
High	1055.7	367.8	87.96 (-318.2%)

tervals by evaluating the performance using 5-minute updates, 1 minute updates, and 4-second discrete control time steps but found no appreciable difference in the economic performance. The results shown in Table 6.2 are shown for the 5-minute dispatch interval. The output of the presented discrete control method is a load profile that does not necessarily lead to the scheduled hourly energy, because the load trajectory over each time interval (which is linear) is slightly different from the optimal load trajectory (that often has a curvature). One approach to deal with this energy deficiency is to use a higher time resolution, so that the trajectories lie on each other more precisely. Another approach is to adjust the targeted load such that it delivers the scheduled energy over each time interval. In this case, the discrete control load is not necessarily equal to the optimal load.

Generally, at low levels of renewables, savings are not possible using the optimal control strategy. The cost savings observed in the extreme low renewables dispatch

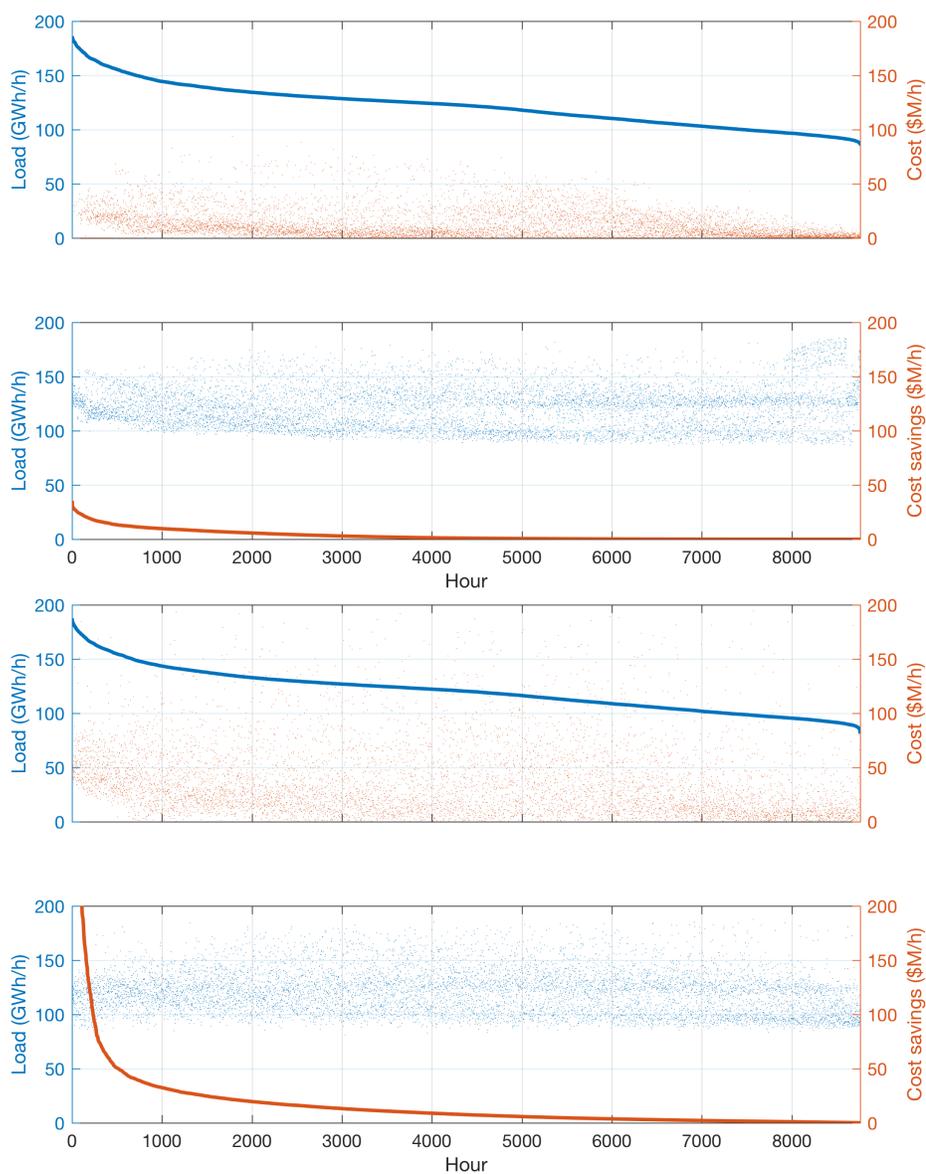


Figure 6.6: WECC 2024 load duration (top) and optimal dispatch savings duration (bottom) using discrete optimal control at 5-minute dispatch rate for the unconstrained (left) and constrained (right) high renewables scenario. The scatter plots are the corresponding cost (top) and load (bottom) values for the durations curves shown.

cases in Table 6.2 are due to the fact that discrete dispatch control follows the optimal trajectory sampling every t_s seconds. This dispatch error can result in small over or underproduction depending on the degree of asymmetry in the optimal trajectory.

At higher levels of renewables the savings are potentially more significant. In addition, the savings are maximum when dispatch tracks the original schedule, which

suggests that there may be a strong economic incentive to avoid carrying over energy tracking error from one schedule interval to the next.

The interconnection-wide scheduling solution in [1] includes a 20-area constrained solution. The hourly energy prices for each area are computed considering both supply and demand energy price elasticities. The energy prices are computed for the interconnection-wide surplus-maximizing schedule over the entire year. The marginal power price is the price of energy for the schedule hour. The marginal price of ramping is 1/300 marginal price of power in the low renewable case, and 49 times the marginal price of power in the high renewable case. The costs, savings and price impact of using this scheduling solution compared to the base case are presented in Tables 6.3 and 6.4. The unconstrained solution is evidently less costly because the combined system-wide fluctuations are smaller than the sum of the individual variations in each balancing authority.

The WECC 2024 system-wide load and savings duration curves⁴ are shown in Figure 6.6. The potential savings are very significant for all scenarios, with the highest savings being found when high levels of renewable resources are available. The savings when more transmission constraints are active are augmented considerably with respect to unconstrained system conditions.

6.6 Summary of Results

The principal result of this chapter is that the use of an optimal dispatch strategy that considers both the cost of energy and the cost of ramping simultaneously leads to significant cost savings in systems with high levels of renewable generation. For the WECC 2024 common case the savings can exceed 25% of total operating costs in the high renewables scenario.

As the bulk power interconnection resource mix shifts from primarily dispatchable non-zero marginal fuel cost resources (e.g., natural gas) to primarily non-dispatchable renewable resources (e.g., hydro, wind, solar) we expect a steady shift in bulk system revenue from energy scarcity rent to ramping scarcity rent. While the total revenue must remain largely the same for financial sustainability, the scarcity pricing mechanism must change.

⁴A duration curve shows the number of hours per year that a time-series quantity is above a particular value. It is obtained by sorting the time-series data in descending order of magnitude and plotting the resulting monotonically descending curve.

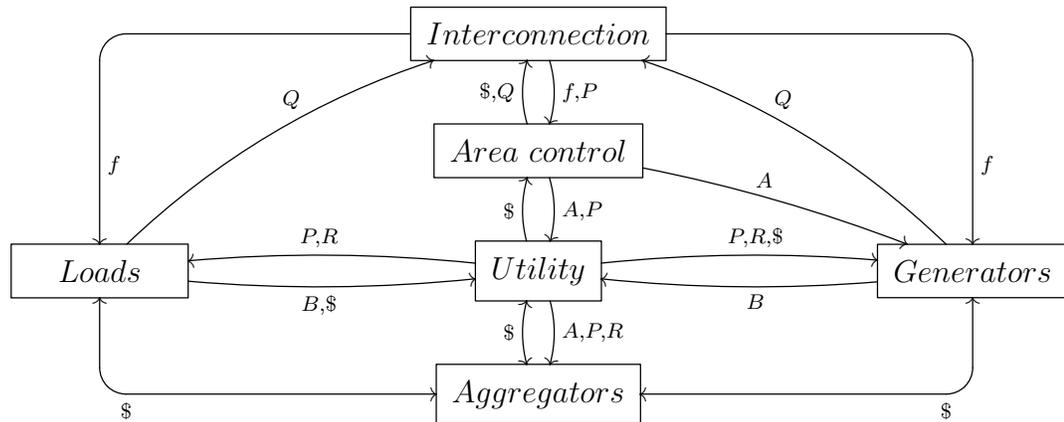
Chapter 7

Discussion

In this chapter we synthesize and discuss more generally the significant results, observations, and outstanding issues arising from the development of a hierarchical inter-temporal solution to transactive control at the interconnection scale.

Transactive control, and more broadly transactive energy require coordination of economic, technical and consumer processes at every scale, both physical and temporal as shown in Figure 7.1. In a conventional power system, scheduling, dispatch, and regulation are established using dissimilar mechanisms according to the prevailing physical or temporal scale. In contrast to this, transactive control systems seek to employ a common scale-free price-based coordination mechanism. This requires a certain degree of uniformity in the coordination approaches at each scale, which in turn requires us to demonstrate that we can model and control control systems that can work with prices over a wide range physical and temporal system scales. This thesis has demonstrated both observation (bidding) and control (response) strategies that support a transactive approaches for several important processes at different scales of grid operation, i.e., aggregate observation and control of periodic demand, control area of system frequency and inter-area power exchange, and optimal control area dispatch.

In this thesis we have examined four important areas of transactive control system design. These are aggregate demand response observation and control, frequency regulation and optimal dispatch. Each section in this chapter discusses these in more detail with Chapters 3, 4, 5, and 6 discussed in Sections 7.1, 7.2, 7.3, 7.4 respectively.



Interarea Transactions

Control Area Transactions

Resource Transactions

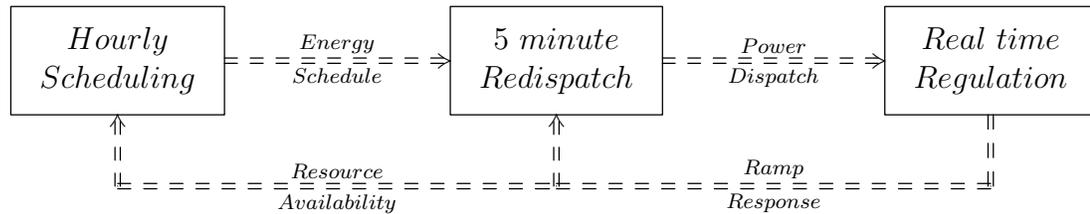


Figure 7.1: Physical (top) and temporal (bottom) system diagram of transactive control systems.

7.1 Demand Response

The availability of a more accurate model of aggregated demand response can be expected to support a wide range of new work on controllable load using real-time pricing. Long-term demand response behavior models do not support the design and analysis of fast-acting demand response as well as the proposed short-term model. In this section we discuss the advantages of using short-term demand response models. As an equilibrium model, the random utility model is expected to be valid for both small signal control stability analysis as well as certain tariff design problems.

7.1.1 Model Limitations

The random utility model is not necessarily valid in its current form for large price disturbance. Roozbehani et al. [82] examined the feedback stability question in the context of wholesale markets and found that real-time wholesale prices could create an unstable closed-loop feedback system for both ex-ante and ex-post settlement

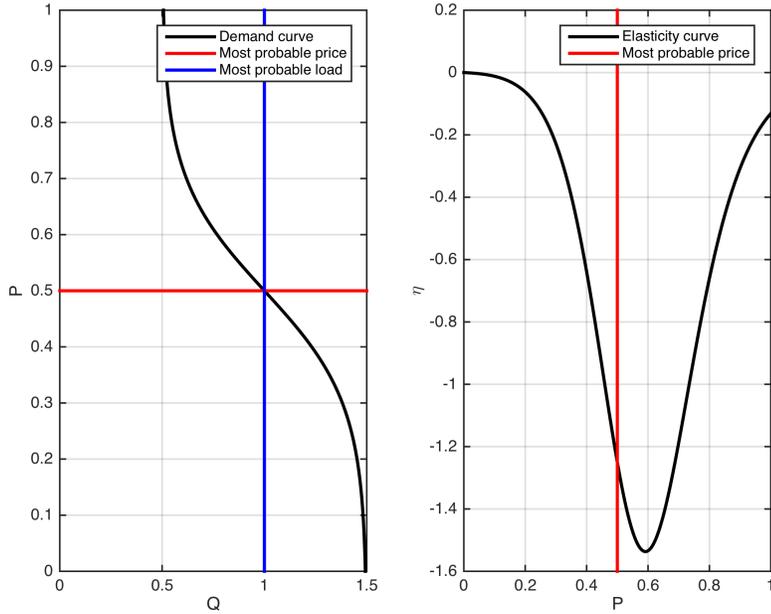


Figure 7.2: Demand (left) and elasticity (right) curves for a nominal case with $\eta_D = -2.5$ and $\bar{P} = 0.5$.

systems. It was established that the absence of an inelastic component in demand contributed to instability, supporting the intuition that increasing feedback gain from price-responsive demand is a concern. Static demand elasticity was also found to lead to loss of efficiency. In follow-up work on price volatility, the authors found that although demand bidding mechanisms eliminate the exogenous feedback delays inherent in settlement-based systems, there remain endogenous load dynamics that can cause bidding mechanisms to exhibit instability [101] and consequently more sophisticated models of demand and consumer response to real-time price dynamics may be required.

The random utility model does not account for the aggregate equilibrium duty cycle of thermostatic loads. In its simplest form, the model assumes a 50% effective duty cycle. The duty cycle tends to skew the demand curve away from the more probable load $\hat{Q} = Q_U + \frac{1}{2}Q_R$. Incorporating this effect would most likely improve the model, particularly with respect to bias error and standard deviation.

When a significant price deviation occurs relative to the natural diversity state, the loads enter a transient response regime. If we compare maximum entropy from Eq. (3.5) to the minimum elasticity from Eq. (3.6) we find that they occur at different

prices. Specifically, load state diversity is maximized when $P = \bar{P}$, but demand elasticity

$$\eta(P) = \frac{-bPe^{a+bP}}{1 + e^{a+bP}}.$$

is minimized when

$$P = \bar{P} \frac{W(e^{2\eta_D-1}) + 1}{2\eta_D}$$

where W is the Lambert W-function

$$W(z) = \sum_{n=1}^{\infty} \frac{(-n)^{n-1}}{n!} z^n.$$

For values of z approaching zero this function is well approximated by z , giving:

$$P \approx \bar{P} \frac{e^{2\eta_D-1} + 1}{2\eta_D}$$

The price at which elasticity is minimized is always greater than \bar{P} for all $0 < Q < Q_R$ when $Q_U > 0$. (Note that Q_U is typically positive, except when a high surplus of uncontrollable distributed generation is present, such as might occur with significant deployment of rooftop solar photovoltaics.) This implies that in equilibrium demand elasticity will tend to increase with prices, as illustrated in Figure 7.2. Under such conditions thermostatic devices no longer follow the equilibrium duty cycle regime and their states diverge from the equilibrium distribution. Consequently their bids depart from the logistic probabilities and no longer follow the bid price distribution of Equation (3.8).

In addition, if the price deviates too quickly, then a diabatic¹ response governs the change in state diversity. As diversity decreases the equilibrium price moves further from the most probable price and the elasticity of demand changes significantly. Decreasing elasticity is observed when loss of diversity favors loads that are *on* and the bid distribution skews left. The distribution skews to the right with increased elasticity when diversity favors loads that are *off*. Note that the periodic behavior of thermostatic loads means that diversity is expected to fluctuate in such a way that elasticity oscillates with damping of about $a/2$ arising from the diversity in the ther-

¹The term ‘diabatic’ here is used in analogy to non-adiabatic processes, i.e., a diabatic process is one in which a significant fraction of the macroscopic state arises from changes in the distribution or arrangement of microscopic states rather than only from the sum of individual states. In other words the total response include a significant contribution from changes system entropy.

mal properties of the home and frequency related to the population average cycling time of the heating/cooling systems.

Device state diversity appears to be a key characteristic that governs demand response. True state diversity can be measured by taking the weighted generalized mean M_{q-1} of the proportional occupancy of states in the population of responsive devices, and then taking the reciprocal of this quantity to obtain the density of devices in states. The diversity of order- q is then defined as

$$D_q \equiv \frac{1}{M_{q-1}} = \left(\sum_{n=1}^N p_n^q n \right)^{1-q^{-1}}. \quad (7.1a)$$

In the limit of $q = 1$, the first-order mean occupancy is well-defined and its natural logarithm converges to

$$H = - \sum_{n=1}^N p_n \ln p_n \quad (7.1b)$$

This is simply the Shannon entropy calculated using natural logarithms instead of base-2 logarithms [102].

Given the linear relationship of bid price to the device state, we use the bid price entropy norm as a measure of state diversity, as shown in Table 7.1. A higher bid price entropy is associated with a higher state diversity. This further explains why the Olympic results fit the equilibrium state assumptions of the random utility model so much better than the Columbus results.

If the demand response resource is very limited it can be expected to saturate more often and lead to reduced diversity and reduced entropy. This effect is illustrated by the reduced diversity duration curves of the Columbus experiments shown in Figure 7.3. In such cases, we expect the full DR and no DR models to be as accurate as the random utility model when DR is called and released, respectively and the half DR model to be more accurate when DR is not called. This condition is more clearly evident in the Columbus data where the total resource was relatively small compared to the total feeder capacity. The observed fluctuations in unresponsive load result in large changes in demand response and lead to DR control saturation, thus making the alternative models satisfactory when compared to the random utility model.

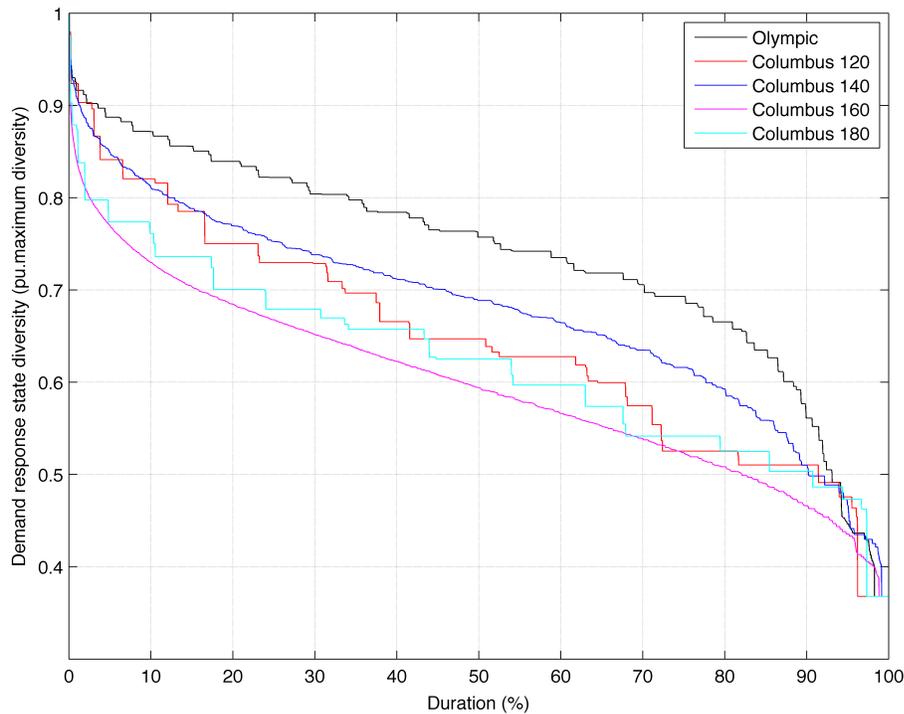


Figure 7.3: Demand response state diversity duration curves for the Olympic feeder and Columbus feeder numbers 120, 140, 160, and 180.

Table 7.1: Bid price entropy statistics
Load State Entropy

System	Mean	Standard Deviation
Olympic	0.68	0.19
Columbus 120	0.54	0.20
Columbus 140	0.60	0.18
Columbus 160	0.47	0.17
Columbus 180	0.51	0.17

7.1.2 Technical and Regulatory Impacts

Interest in transactive control has also resulted in renewed discussion of tariff design and rate-making processes for utilities that wish to adopt the real-time pricing strategy. The Columbus demonstration included a rate case for the real-time price double auction called “RTPda” that was approved by the Public Utility Commission of Ohio [103, 104]. However, real-time price tariffs and rate designs can be difficult to study when loads respond to hourly or sub-hourly prices. These tariffs present new challenges for utilities and regulators alike. Concerns have also been expressed

regarding customer acceptance of real-time prices [58] but utilities could offer a portfolio of tariffs, including a real-time price, from which customers may choose. The rate at which tariffs are adopted then becomes a conventional portfolio optimization problem where opt-in/opt-out incentives and penalties are offered to achieve tariff adoption mixtures that meet utility and regulator objectives [105].

The implications of adopting the random utility demand response model for both utilities and regulations have yet to be explored fully. However, some initial thoughts may foster discussion and suggest further research on methods to effectively incentivize enduring and sustainable demand response from residential customers.

First, the existence of upper and lower price asymptotes is often overlooked in demand response analysis for electricity loads. We recognize that linearization of the demand function is often necessary. But the asymptotes require us to acknowledge the existence of hard upper and lower constraints on short-term responsive load. When prices deviate from the most probable price, the elasticity of demand approaches zero. Thus an important change occurs in fast-acting demand response when prices induce load to move into either curtailment or pre-heat/pre-cool/recovery regimes: load diversity is decreased and important endogenous load oscillations can be induced, which are independent of the oscillations induced by control feedback. Unless technical steps are taken to dampen endogenous load oscillations before they occur, price oscillations can emerge that can only be mitigated by reducing the feedback from load and waiting for load state diversity to be restored. It is therefore incumbent on utilities to identify all the conditions under which instability can emerge and implement either market and/or load control mechanisms to mitigate them.

Second, existing residential electricity tariffs do not adequately support the development of fast-acting demand response or the mechanisms needed to mitigate the potential instabilities associated with demand response. This is particularly a concern in the presence of supply resources at the distribution level, such as rooftop photovoltaic panels and significant amounts of battery storage, such as so-called vehicle-to-grid discharging. Clearly the availability of such resources can be easily incorporated into the market mechanism demonstrated in both the Olympic and Columbus trials. However, the marginal cost of most renewable retail resources is effectively zero, which not only can give rise to revenue adequacy problems for utilities, but can also effectively shut off the very price signaling mechanism needed to control load. Strategic bidding may be necessary for these resources to elicit non-zero price signals, as was demonstrated in the Olympic study to manage minimum and maximum generation

run time limits. However, strategic bidding at the retail level entails an entirely new class of regulatory problems that may require mitigation strategies that are not part of current tariff design and approval procedures.

Demand response in residential settings presents an additional challenge to utilities. Historically, the cost-per-point and cost-per-megawatt for the supporting infrastructure and management of these systems has been significantly higher than for commercial and industrial customers. Recent industry estimates suggest that customer-premises portals and in-home energy management systems will cost between \$US150 and \$US300 by 2030 [106]. But life-cycle cost analysis for transactive technology is not generally available yet, in part due to the lack of simulation tools that can properly evaluate the benefits of the technology [107]. Fast-acting demand response requires higher communication rates than hourly or day-ahead critical peak pricing mechanisms used for commercial and industrial demand management. Automated metering infrastructure has offered the promise of fast and accurate communications with residential loads. But much of that promise has yet to be analyzed in detail or realized in practice, either with incentive-based or even direct load control mechanisms.

The cost per unit of power reduced of behind-the-meter infrastructure for residential loads has been typically difficult to compare to the cost per energy saved using long-term demand response programs. The additional costs per customer tends to delay adoption of residential demand response programs until after the more cost-effective available commercial and industrial demand resources have been exhausted. The advent of smart thermostats like the NESTTM and potentially implicitly smart loads like electric-vehicle chargers can be expected to increase the available low-cost responsive load in the residential settings [108].

7.2 Aggregation

We have shown that closed loop control of aggregate $T\Delta_0$ thermostatic loads can be accomplished provided a suitable control system which curtails $u(k)$ device for $k = 0, 1, 2 \dots$. Each control impulse transfers devices between the unresponsive population and the responsive population, altering the responsive population's state $\mathbf{x}(k+1)$ by simply adding the new population's $x_k(k+1)$ response to the input $u(k)$. By combining load curtailment and load release impulses the aggregate response can be shaped to track an arbitrary reference signal $r(k)$, provided sufficient devices are

available in both the responsive and background unresponsive population to supply the net change required by each impulse $u(k)$.

Overall the results indicate that an aggregate load model of discrete time thermostats can be used to design such an aggregate load control strategy. But not all controllers designed using conventional methods perform equally well. Proportional controllers and unity-damped controllers, while simplest to implement, tend to be slow to respond and can have stability problems at high load. Pole placement designs tend to result in controllers with steady state errors and/or excessive overshoot. Integral error feedback exhibits the smallest average error as well as the smallest overshoot and are therefore deemed the best controllers studied. The main disadvantage of the integral-error-feedback controller design remains the need to implement a parameter estimation method to reduce the effect of model error on tracking performance.

A potentially useful extension to this approach is “symmetric control”, which uses all devices in the unresponsive population instead of only devices that are *on*. This control approach allows the utility to increase loads by turning devices *on* as easily as it can decrease load by turning devices *off*. In some cases, this approach may be more practical for utilities to deploy, and offers the added benefit of addressing possible privacy concerns resulting from any strategy that requires the utility to know whether one particular device is actually *on* before choosing which devices to signal. In addition, it offers the utility an opportunity to aggregate load control for ancillary services, which may require both load *up* and load *down* regulation.

In the cases when we seek full control over the load we must make two important assumptions regarding load curtailment strategies.

1. Devices are selected regardless of their current state.
2. The load is observed based on the number of devices that remain on rather than the number of devices that are turned off.

The demand response strategy is then described by using $\mathbf{h} = [\bar{R}, 1 - \bar{R}]^T$ and $\mathbf{c} = [\bar{q}, 0]$, where $\bar{R} = \frac{\bar{r}_{off}}{\bar{r}_{on} + \bar{r}_{off}}$ is the population average duty cycle, and \bar{q} is the population average load of a single device. The system is controllable when

$$\begin{aligned} |C| &= \begin{vmatrix} \bar{R} & \bar{R}(1 - \rho_{on}) + (1 - \bar{R})\rho_{off} \\ 1 - \bar{R} & \bar{R}\rho_{on} + (1 - \bar{R})(1 - \rho_{off}) \end{vmatrix} \\ &= \bar{R}\rho_{on} - (1 - \bar{R})\rho_{off} \neq 0 \end{aligned}$$

or when

$$\bar{R} \neq \frac{\rho_{off}}{\rho_{off} + \rho_{on}},$$

a condition that is satisfied when $\sigma_{on}, \sigma_{off} > 0$; i.e., the thermal properties of the population are diversified. The system is observable when

$$|O| = \begin{vmatrix} \bar{q} & 0 \\ \bar{q}(1 - \rho_{on}) & \bar{q}\rho_{off} \end{vmatrix} = \bar{q}^2 \rho_{off} \neq 0$$

which is always true when thermostatically controlled load is active.

In future work we anticipate examining parameter identification strategies to facilitate the implementation of practical load control at the utility scale. In addition, adaptive control, model-predictive control and optimal control strategies should be considered. The latter seems particularly interesting in the context of optimal area control and balance frequency regulation objectives with periodic redispatch of load control resources to keep the maximum number of loads available for regulation services over time.

7.3 Regulation

Using the closed-loop system model GK we can compare the proposed control policy's contribution to improving intra-area control robustness to uncertainty in the availability of fast-acting demand response at the time of a generation contingency. Uncertainty in controllable load can be very large and results from significant diurnal and seasonal variations in the load composition [109]. Fast-acting demand response is typically associated with heating, cooling, and more recently, vehicle charging loads because they are flexible in the short term and are usually a significant fraction of the total load at peak times when generation contingencies pose a greater threat to overall system reliability.

7.3.1 Robustness to FADR Uncertainty

The robustness of the \mathcal{H}_2 -optimal control policy relative to the conventional ACE control in the presence of highly varying levels of FADR is apparent from Figure 5.11. This result is highly significant, particularly when used in the context of FADR to mitigate high penetration of renewables. Increasing renewable resources are associated

with declining system inertial response [110] and can lead to more rapid degradation in system stability. As previously discussed, increased FADR can also contribute to deteriorating system stability margins [94]. So while FADR can mitigate renewable intermittency in terms of temporarily relieving thermal generating units from having to quickly ramp, it cannot be concluded that FADR necessarily improves short term system stability without additional measures being applied to the area control policy. The initial results for high-levels of FADR suggest that this is indeed the case and that at the very least the conventional ACE control policy must be reexamined as increasing levels of uncontrollable renewable generation are used and FADR is employed to mitigate the impact.

A further consideration is the selection of the FADR design condition. In this study a 5% FADR level was used. The performance of the new control policy using this design condition is quite robust for a wide range of FADR levels. However, it should be recognized that the new control policy is optimal only when the FADR level is close to 5%. At other levels of FADR availability the performance would be suboptimal, although it still remains much better than the conventional ACE control policy, as the cost and energy savings in Table 5.1 demonstrate. This suggests that careful consideration should be given to the choice of FADR design conditions, especially with respect to (1) the probability distribution of FADR levels over the course of time, (2) the probability of a generation contingency occurring over the course of time, and (3) the relative cost impacts of those contingencies.

It is significant that the new control policy relies only on measurement of import/export power from the control area. For the control policy to be effective, these measurements must be made at very high rate compared to the SCADA measurement rate of 0.25 Hz for ACE. Based on the very fast response rate of the loads, a measurement rate similar to that of phasor measurement units (PMUs) may be necessary for the proposed area control policy. Most PMUs can sample phase angles at 60 Hz, and are capable of point-on-wave measurements in excess of 1000 Hz. However, the control design would have to consider the communications latency from the remote PMUs to the control area's data concentrator [111]. PMU technology and availability is evolving rapidly and the North America Synchrophasor Initiative (NASPI) has considered the possibility of such a requirement in the design and implementation of the current synchrophasor network [112].

7.4 Dispatch

The significance of the results shown in Figure 6.3 cannot be understated. First we observe that when the marginal price of energy a is much larger than the marginal price of ramping c , the optimal response is very similar to the conventional dispatch strategy, giving us some assurance that today's operations are very nearly optimal. However, when $a \ll c$, today's hourly dispatch strategy is not optimal. As the fraction of cost attributed to energy decreases relative to the cost attributed to ramping, we see that ratio $\omega = \sqrt{\frac{a}{c}}$ decreases and the value of changing the dispatch strategy increases dramatically. In the limit of a very high renewable scenario the savings achievable using the optimal dispatch strategy can be extremely significant. Failure to adopt an optimal dispatch such as the one proposed could result in major and likely unnecessary costs. Utilities will inevitably find it necessary to mitigate these costs, either by reducing the amount of renewables, by increasing the revenues from their customers, or by developing some kind of optimal resource allocation strategy such as the one proposed.

A sensitivity analysis of the savings as a function of the marginal price of ramping c shows that the savings are not overly sensitive to changes in our assumption of the cost of ramping scarcity. Figure 7.4 shows that for a 50% decrease in c , we observe a 10.3% decrease in savings, while a 50% increase in c results in a 3.9% increase in savings. This suggests that the savings from employing the optimal dispatch strategy is quite robust to our uncertainty about the marginal price of ramping resources.

In any financially sustainable future scenario, we must consider how the long-term average costs and fixed costs are recovered under the pricing mechanism. We have implicitly assumed in this study that renewable generation and utilities cannot sustainably continue employing complex power purchasing agreements and subsidies to hedge against energy price volatility. Instead all parties should come to rely on separate real-time pricing mechanisms for energy, power and ramping response of the resources they control.

Shifting revenue from resource allocation mechanisms based primarily on energy resource scarcity to ones based primarily on flexibility resource scarcity can be expected to have a significant impact on the cost of subhourly resource dispatch. The optimal strategy for low renewable conditions very closely matches the strategy employed today when moving hour-to-hour from one scheduled operating point to another. Indeed, the optimal dispatch strategy does not offer any significant cost savings

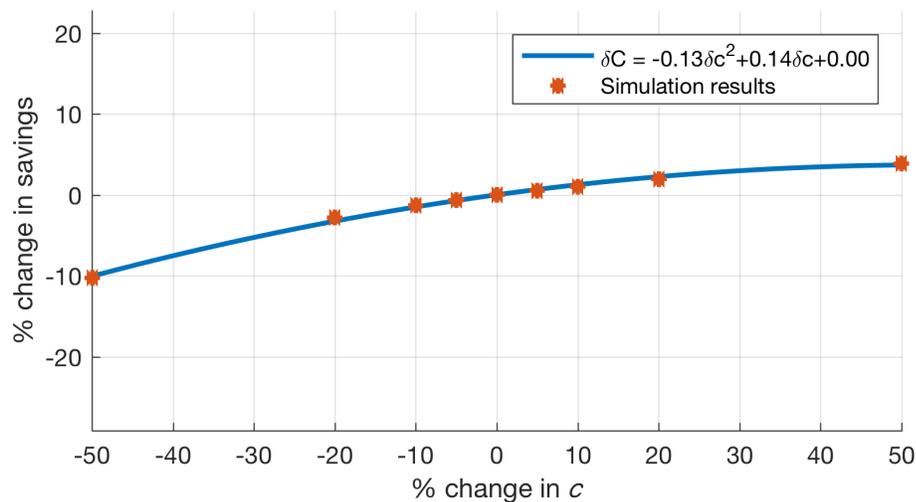


Figure 7.4: Sensitivity of savings to marginal price of ramping resources.

when overall pricing is dominated by energy resource scarcity.

However, as increasing amounts of renewables are introduced, the scarcity rents may shift from energy to flexibility resources. The optimal subhourly dispatch strategy may be expected to change with increasing emphasis on avoiding high ramp rates over sustained periods at the expense of maintaining a constant power level over the hour.

The relationship between existing price signals for various grid services and the three principal price components needed to implement this optimal strategy requires further investigation. It is evident that the marginal price a represents a linearization of the energy price itself at the current operating point. But it is not clear yet whether and to what degree the marginal prices b and c can be connected to any existing price signals, such as the capacity price or the price of ancillary services such as frequency regulation resources, generation reserves, and demand response. The links do suggest themselves based on both the resource behaviors and physical dimensions of the parameters. However, it is not yet certain whether this will be simply a matter of obtaining a linearization of the services' cost functions at the appropriate operating point.

Additionally, it is instructive to note that the marginal price of redispatched power

b is not important to the optimal dispatch strategy, insofar as the parameter does not appear in Eq. 6.2. This leads one to conclude that to the extent capacity limits do not affect either energy or ramping scarcity rents (or are already captured in them), the marginal cost of additional resource capacity is never considered for optimal subhourly dispatch control. This is consistent with the expectation that sunk costs should not be a factor in the selection of which units to dispatch at what level, at least to the extent that these costs are not entering into the energy or ramping costs.

In the presence of significant renewables, the energy marginal cost does not entirely reflect the grid condition without considering the cost of ramping up and down services. Therefore, the energy price cannot be solely used as a control signal to the generation and load units to achieve the optimal utilization of resources. In order to quantify the value of ramping product we suggest using a market framework in which flexible generation and load resources compete to sell their ancillary service products at the bulk electric system level. As renewable level rises the marginal price of energy decreases (smaller a value) because renewables are zero-generation cost resources, but the marginal price of ramping increases (larger c value) because the system requires more flexibility to handle the generation variation. In long run, inflexible units get retired and more flexible units are built to support the renewable integration since flexibility will be a revenue source rather than energy.

The availability of high renewables can lead to situations where low cost energy is being supplied to areas with high cost flexibility resources through constrained interties. The optimal strategy avoids dispatching these high cost flexibility resources to the extent possible by reducing the ramping schedule. The more transmission capacity available, the lower the overall cost, but we note that even when the system is constrained, the cost of optimally dispatching flexibility resources can be significantly lower under the high renewables case than under a low renewables scenario.

7.5 Ramping Market Price

Wholesale electric energy markets have been used widely since they were first introduced in the 1980s. The market design is conceptually relatively simple and is illustrated in Figure 7.5, where renewables have lower marginal cost than fossil units. The supply and demand curves are constructed by summing the marginal cost functions of all supply and demand units when ordered by price. As a result, when the clearing price is obtained, only the supply units with marginal costs below that price

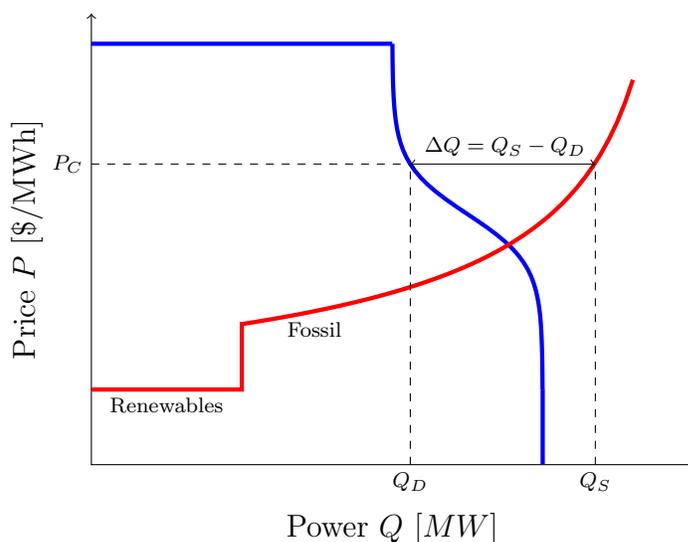


Figure 7.5: Energy market with supply curve (red) and demand curve (blue), with the price P_C at which the export ΔQ is realized.

and demand units with reservation prices above that price are operated.

The key feature of this mechanism for the purposes of transactive control between control areas that if a specific difference ΔQ between supply and demand is desired (e.g., a specific import or export quantity is to be achieved) then we can directly compute the price P_C at which that import or export is realized.

Throughout this thesis ramping supply markets are used assuming (1) that they exist and (2) that they operate in a similar manner. In this section we explore whether these two assumptions are true and to the extent that they are not, what must be done to see them realized.

We illustrate a wholesale ramping market in Figure 7.6. In this market, demand response from loads are the least costly to deploy per unit of ramp rate, followed by storage units, and fossil units are the most costly to deploy. The required rate \dot{Q}_C is obtained from the real-time system operation, and units that have contributed to the ramp response are compensated at the price R_C . All the units with ramping prices below R_C are required to provide the ramp response \dot{Q}_C .

At this point it remains unclear how to design a market clearing mechanism that generates the separate marginal price signals identified in the optimal dispatch strategy in Chapter 6. This remains an important area of ongoing research.

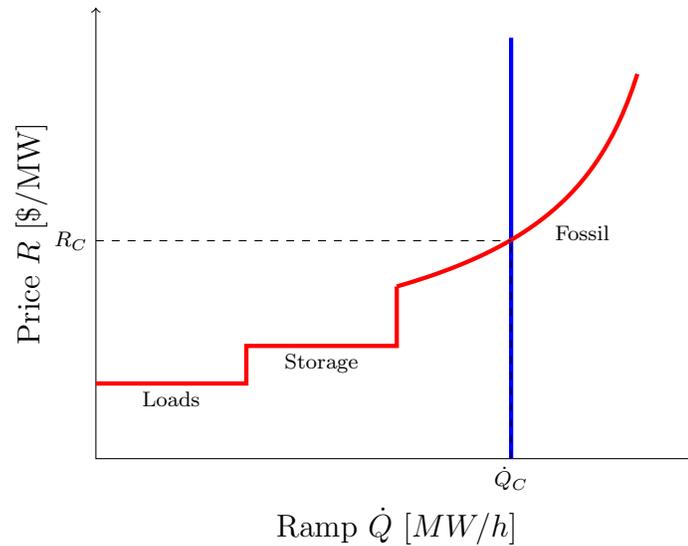


Figure 7.6: Ramping market with supply curve (red) and demand curve (blue), with the price P_C at which the export ΔQ is realized.

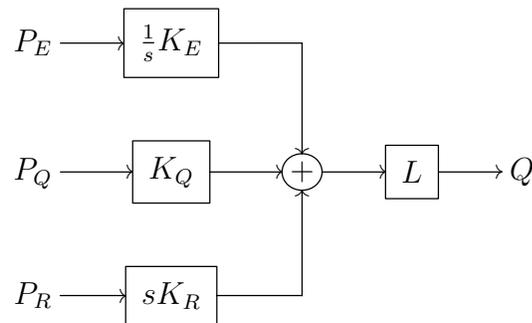


Figure 7.7: Price-based control over-actuation of aggregate load

7.5.1 Unified Market Design

Throughout this thesis we have considered problems where energy, power and ramping are controlled jointly. We have made the assumption that this is necessary based on the observation that $E(t) = \int_0^t L(x)dx$ and $R(t) = \dot{L}(t)$, where E is the total energy consumed by the load L at the time t , and R is the rate of change of the load L at the time t . We then recognized that there can be only one controllable resources in the aggregate at any given timescale, as shown in Figure 7.7, using price control gains K_E , K_Q and K_R . Using independent prices for energy, power, and ramp as control inputs at any given timescale to control the system leads to over-actuation of the system—there are three control inputs (P_E , P_Q , and P_R), but there is only one degree of freedom (Q). Consequently we must consider a unified price-based control

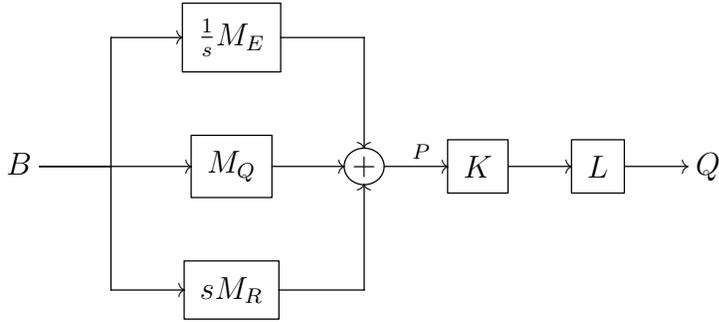


Figure 7.8: Market-based unified price-based control of aggregate load

where there is a single price, which effectively controls the joint energy, power and ramp behavior of the load, as shown in Figure 7.8.

The unified price is given as

$$P(s) = B(s) \left[\frac{1}{s} M_E(s) + M_Q(s) + s M_R(s) \right],$$

where M_E , and M_Q and M_E are the energy, power, and ramp price-discovery mechanisms. The assumption in the literature is that these mechanisms are independent, but the results of this thesis suggest that it cannot be so.

The energy and ramp market mechanisms are well separated in frequency domain, and the results of the optimal dispatch function support this conclusion. However, the bids into the energy market and the bids into the ramp market are not independent, as is often recognized when one attempts co-optimization of the two. In the aggregate, the ramp resources are determined by the cleared energy resources. Thus there is only one set of coordinated bids, B , which come from the aggregated loads and must be cleared through all the markets simultaneously, with each market influencing the response in a separate response band in frequency domain.

What remains to be determined is the exact design of these bid clearing mechanisms so that they jointly determine a unified price. Specifically, we do not have transfer functions that adequately describe the general transfer functions of these markets such that we can bring control theory to bear on the system design problem. In Appendix B we examine the stabilization of a “tatonnement” transfer function for bilateral negotiated price discovery. The study of transfer functions for multilateral and auction-based price discovery processes is a significant open area of research for transactive control systems.

7.6 Recommendations

In summary, we have identified the following ten recommendations for future work in transactive control research.

Recommendation 1: Utilities need to identify the conditions under which price instability can emerge under high renewables and implement either market and/or distributed resource control mechanisms to mitigate them.

Recommendation 2: Strategic bidding may be necessary for resources to elicit nonzero price signals. However, strategic bidding at the retail level entails an entirely new class of regulatory oversight that is not currently envisioned in most jurisdictions.

Recommendation 3: The benefits of automated metering infrastructure for residential customers have yet to be studied in detail or realized in practice, either with indirect price-based or direct load control mechanisms.

Recommendation 4: Symmetric control, which uses all devices in the uncontrolled population instead of only devices that are on, needs to be considered as a possible control approach that allows the utility mitigate renewable intermittency.

Recommendation 5: Adaptive control, model-predictive control and optimal control strategies should be considered for utility-scale resource dispatch, particularly in the context of optimal area control and balance frequency regulation with periodic redispach of distributed resources to maintain provision of sufficient distributed resources for regulation services.

Recommendation 6: Further examination of the choices of fast-acting demand response (FADR) designs is needed, especially with respect to (1) the availability of FADR over time, (2) the probability of a generation contingency occurring over time, and (3) the relative cost impacts of those contingencies.

Recommendation 7: Distributed resource control designs that depend on phasor measurement units (PMUs) need to carefully consider the communications latency from the remote PMUs to the control area's data concentrator before being considered as an element of or alternative to transactive control systems.

Recommendation 8: The relationship between existing price signals for various grid services and the energy, power, and ramp price components needed to implement optimal dispatch strategies need to be investigated. Quadratic cost functions allow a linearization of these prices at the current operating point. But the marginal price of ramping in particular has yet to be connected to any ancillary services market designs or existing cost functions.

Recommendation 9: The problem of how to design a general transactive market clearing mechanism that discovers distinct marginal prices for energy and ramping remains an important unresolved question.

Recommendation 10: The exact design of market mechanisms that produces a unified price signal is unresolved. Specifically, we do not have transfer functions that adequately describes the power response as function of price in the short term such that we can analyze system performance in general. The study of transfer functions for multilateral and auction-based price discovery mechanisms is an important open area of research for transactive control systems.

Chapter 8

Conclusions

The results of this thesis generally show that it is possible to establish both a common transactive basis for control at over a wide range of scales, as well as achieve locally advantageous implementations of transactive control that in many cases outperform the conventional controls used today. This is achieved for example, by leveraging the flexibility of millions of small load devices to relieve the burden on flexible generators to respond to ramping demands stemming from the inflexibility of other generators. Similarly, the results show how dispatch flexibility can be used to both guide more robust and cost-effective frequency regulation response, as well as maintain optimal scheduling from one hour to the next.

In this thesis we have examined four important areas of transactive control system design that will be necessary to the eventual implementation of a more comprehensive transaction energy system. These are (1) aggregate demand response observation and (2) aggregate discrete-time cyclic load control, (3) optimal frequency regulation with area import/export tracking, and (4) optimal area dispatch.

An *ab initio* aggregate load model of fast-acting controllable thermostatic electric loads operating under the transactive control paradigm that corresponds to the Random Utility Model commonly used in the economics of consumer choice. The proposed model is verified with empirical data collected from field demonstration projects and shown to perform better than alternate models commonly used to forecast demand in normal operating conditions. The results suggest more broadly that (1) existing utility tariffs appear may be unable to incentivize sufficient fast-acting demand response to mitigate the impact of high renewables, and (2) existing load control systems may be easily saturated and

become unstable if utilities simply close the control loop using only real-time energy prices.

An aggregate load controller for discrete-time zero-deadband thermostats, a general controller design that allows various aggregate load control strategies to be explored, and specific controller that allow large-scale control of thermostatic loads that have high endurance. The aggregate load controller is studied for a variety of conventional control designs and is used to design a stable discrete-time closed-loop aggregate load controller for a utility-scale demand response dispatch system. The new controller is shown to be stable, controllable, and observable. It is tested using a large-scale agent-based model of residential heating loads, and works equally well with both direct and indirect load control systems.

A frequency regulation strategy using an \mathcal{H}_2 -optimal controller for control areas in bulk electricity interconnections that is suited to jointly controlling generation and demand response in areas that have a high penetration of both intermittent renewable resources and fast-acting demand response. The mechanism offers significant cost savings in systems with very high levels of renewable resources. The results show that the optimal controller outperforms the conventional ACE control policy by 1) providing faster return to the schedule under varying demand response levels, 2) reducing the cost of calling up reserve units for regulation services, and 3) minimizing deviations from the global surplus-maximizing schedule.

An optimal dispatch that considers both the cost of energy and the cost of ramping simultaneously and exploits significant cost savings opportunities in systems with high levels of renewable generation. The optimal price-based control solution responds to both a short-term price of energy for demand dispatch, as well as a short-term price of power for ramp response. The result is (1) a formal method to compute the optimal sub-hourly power trajectory for a system when the cost of energy and ramping are both of the same order, (2) an optimal resource allocation strategy based on this optimal trajectory, and (3) a simulation method to evaluate the cost savings potential of using the optimal trajectory when compared to the conventional sub-hourly dispatch used in today's system operation.

8.1 Principal Contributions and Findings

- (1) **Demand Response** can be modeled as a logistic demand curve for short term electricity consumption derived from the first principles of controllable thermostatic electric loads operating under the transactive control paradigm. We have shown that this model corresponds to the Random Utility Model commonly used in the economics of consumer choice. The model's performance is compared to results from two US Department of Energy demonstration projects in which short-term demand response data were obtained. We found that the random utility model predicts the total demand response to smaller price fluctuations very well, but that model performance degrades as the magnitude and frequency of price excursions increases and as the diversity of load states decreases. We conclude that the random utility model is suitable for demand response studies that utilize steady state conditions for most situations with only infrequent and modest price excursions.

In its present form the random utility model provides a robust framework that is well-founded in the engineering principles of how thermostatic devices behave in price-based control environments. By joining the engineering and economic behavior of such devices, the random utility model is set to become an essential element in the planning, design and eventual deployment of large-scale load control strategies.

- (2) **Load Aggregation** at the utility scale was developed in the direct load control problem for the situation in which the controlled loads employ discrete-time zero-deadband ($T\Delta_0$) residential thermostats. We showed how $T\Delta_0$ thermostats allow utility dispatchers to use small adjustments to the consumer's setpoint to modulate the total load with greater precision than is possible using current setback control of thermostats with non-zero deadbands. These new digital thermostats can serve as the basis for highly accurate direct load control systems, as well as price-based indirect load control systems.

A new linear aggregate load model was found based on the dynamics of load states and its fundamental characteristics were used to consider a number of benchmark aggregate load control designs from first principles. We used this model to design a simple closed-loop aggregate controller for a discrete-time utility-scale demand response dispatch system that is compatible with the re-

quirements for both direct and indirect load control systems and tested the control design using a large-scale agent-based model of demand response based on thermostatic loads. We showed that the aggregate controlled load is stable, controllable, and observable, and exhibits both the transient and steady-state response characteristics necessary to serve equally well for utilities that seek to control load using either direct load control or price-based indirect demand response strategies.

- (3) **Optimal regulation** using an \mathcal{H}_2 -optimal approach was synthesized to obtain the control policy for control areas in bulk electricity interconnections. The approach is suited to controlling both generation and demand response in areas that have a high penetration of both intermittent renewable resources and fast-acting demand response. The implementation of the \mathcal{H}_2 -optimal control policy is compatible with, and indeed depends on the existing methods of transactive control 5-minute dispatch strategy. The transactive \mathcal{H}_2 -optimal area control policy was shown to be superior to the conventional ACE control policy in that it is (1) significantly more robust to uncertainty in the amount of fast-acting demand response that is available, (2) always less costly and less energy intensive, and (3) minimizes deviation from any surplus maximizing schedule.
- (4) **Optimal dispatch** was solved using a strategy that considers both the cost of energy and the cost of ramping simultaneously, which leads to significant cost savings in systems with high levels of renewable generation. For the WECC 2024 common case the savings can exceed 40% of total operating costs in the high renewables scenario. As the bulk power interconnection marginal resources shift from primarily dispatchable non-zero marginal fuel cost resources (e.g., natural gas) to primarily non-dispatchable renewable resources (e.g., wind, solar) we expect a steady shift in bulk system revenue from energy scarcity rent to ramping scarcity rent. While the total revenue must remain largely the same for financial sustainability, the scarcity pricing mechanism must change.

8.2 Future Research

As discussed in Chapter 7 the results achieved do not provide all the elements necessary to realize the vision of a fully transactive system. Significant work remains

at all physical and temporal scales. Notably absent is the computational infrastructure required for transactive implementations. Also significant is that very few pilot projects are planned, particularly addressing coordination by higher level system functions such as control area dispatch or inter-area scheduling. With regard to the specific systems discussed in this these, the following future research opportunities have been identified.

8.2.1 Demand Response

Future work on short-term aggregate demand response must investigate models that account for varying effective duty cycles for large populations of responsive loads and account for the effect of large fast changes in prices. Methods based on the entropy of the population states seem to offer significant promise.

The random utility model supports previous claims that additional research will be required to mitigate the potential instabilities that may emerge when employing real-time pricing signals for closed-loop feedback control of fast-acting residential demand response. This work will necessarily require a contribution from the field of control theory, while maintaining strong support from economists with an interest in mechanism design.

The model also highlights emerging challenges for tariff design and rate approval processes, particularly in cases where significant distributed generation and storage resources participate in price-discovery alongside fast-acting demand response. The question of incentive-compatible retail market design can be expected to become more important as new tariffs and rate structures are developed by utilities and regulators.

8.2.2 Aggregate Thermostatic Load Control

Future work on thermostatic control needs to develop parameter system identification strategies that facilitate the implementation of practical load control at the utility scale. In addition, adaptive control, model-predictive control and optimal control strategies should be considered. The last seems particularly interesting in the context of optimal area control and balance frequency regulation objectives with periodic redispatch of load control resources to keep the maximum number of loads available for regulation services over time.

8.2.3 Regulation

The \mathcal{H}_2 -optimal control design has considered only local disturbances, viz., loss of generation originating within the same generation control area. It will be necessary to consider a wide range of additional disturbances including

1. An internal loss of load (note that this is not equivalent to a negative loss of generation because the responsive resource mix changes),
2. An external loss of generation or load,
3. A loss of a tieline between generation control areas, and
4. Rapid ramp up and down, both internally and externally, due to unexpected changes in renewable resources.

In addition, the new control policy must address the impact of eliminating direct observation of system frequency from the measurement. The role of the secondary generation is to cancel the steady-state frequency deviation and bring each area's net exports back to the scheduled value, either with ACE or \mathcal{H}_2 -optimal control. A component of the \mathcal{H}_2 -optimal area generation control minimizes the frequency deviations so it will cope with some, but not necessarily all of the steady-state frequency deviation. It seems unlikely that such a change to the area control policy could go unnoticed, particularly in the event that system inertia and damping change significantly in a short time. Although generation droop remains in effect for all controlled generating units that do not have a master controller, the slow diminution of droop-only units could lead to situations where there is slow or inadequate control of system frequency even though there is very authoritative control of inter-area power exchanges.

This study examined an extreme case in which 75% of the generation was uncontrolled renewable and only 25% of the generation was regulated using the area control policy. For one thing, such a great percentage of renewables means the system does not have enough inertia to deal with both renewable and load uncertainty, so the existing AGC system is unable to adequately control the system. Although we considered a relatively low inertia systems, we expect that much lower total system inertia should be considered (perhaps as low as $M = 2$ seconds). With such low system inertia it seems even more likely that control areas will require an augmented

or entirely new control policy, and that the advantage of \mathcal{H}_2 -optimal ACE control may be more evident.

In North America, balancing authorities are scored based on how well they contribute to interconnection frequency regulation needs while tracking their export schedule. The optimal response as presented here does not consider how often “zero-crossings” of nominal frequency occur or how noisy frequency control is. Future studies of optimal-ACE policies considering fast-acting demand response control design will need to also consider the control performance standards CPS-1 and CPS-2 [91].

8.2.4 Dispatch

The principal consequence of using the proposed dispatch strategy for control areas considering both energy and ramping costs is an annual cost savings in WECC projected to exceed US\$150B by the year 2024. Extrapolated roughly to all of North America, this may result in savings approach US\$500B annually. It is also clear that the use of energy-only market designs runs counter to the results of this study. Flexibility resource markets may become increasingly important as ramping resources become more scarce. This may be true even in regions that are not dominated by local renewable generation but adequate transmission capacity is available for renewables in remote regions to displace local dispatchable generation. This may give rise to a new set of challenges for utility and system operators as they seek a revenue model that not only provides for operating costs, but also maintains the coupling between retail demand response and wholesale supply and retail delivery constraints. If the cost of the wholesale system becomes increasingly dominated by ramping resource constraints, while retail continues to use energy prices to encourage consumer efficiency, then retail behavior will be not affected as much by short-term wholesale price fluctuations. This trend runs against the desire for more engaged consumers who can respond to system conditions in real time. Clearly a new utility revenue model is needed if the transformation to a high renewable *modus operandi* is to occur successfully in the coming decades.

8.3 Summary

Taken as a whole the results of this thesis show that one can incorporate transactive energy technology into today's system operations without simultaneously transforming all aspects of system operation in all parts of the system. Each solution can be independently deployed profitably and successfully in a heterogenous operating environment while maintaining system reliability, security, and economic stability.

But the advantages of integrating these solutions together grow with the increased use of transactive technologies, with greater availability of renewable energy resources, and with greater participation by consumer-based assets in system operation. In the final analysis, transactive control offers the greatest potential for consumer benefits, the most flexibility approach to evolving the utility business model, and most importantly the most effective means of addressing the global impact of meeting the energy demands of the 21st century economy.

Appendix A

Demand Elasticity

Over the years, research has consistently shown that demand response (DR) can participate in all electricity markets, including energy, capacity, and regulation [28]. Transactive control and Powermatcher were proposed and demonstrated in the field as approaches that could unify the scheduling, dispatch and control of all distributed energy resources, including DR, on the retail side [38]. Since then there is growing interest in the broader concept of transactive energy in both Europe and North America [113]. In addition a number of important technical contributions have been made to support transactive energy systems, including the design of device controllers [73], real-time pricing [114], and retail market designs, in various jurisdictions in North America and elsewhere.

However, the fundamental behaviors and properties of transactive systems remain a largely open research area. From the viewpoint of power system and control engineering, this is a critical gap and to this day it stands as a significant barrier to the widespread adoption of transactive control systems—utilities are hesitant to adopt a new control strategy in the absence of a clear and validated mathematical framework to study and prove the stability, robustness, and reliability against all hazards.

One important aspect of this gap is the nature and role of DR elasticity of energy, power, and ramp resources in electricity markets. This appendix briefly presents a result that may be of significance to researchers in transactive energy and transactive control insofar as it establishes a strong and potentially useful connection between the elasticity of energy, power and ramp response resources in electricity markets, both wholesale and retail. This relationship is examined and its implications are discussed briefly with respect to the design of DR controllers, utility and demand aggregator business models, and regulatory oversight of electricity markets, especially in the

context of retail competition.

A.1 Analysis

Consider the energy consumption $q(p, t)$ at the energy price p and time t , the price varying also in time. The power demand is

$$s(p, t) = \frac{\partial q(p, t)}{\partial t}$$

and the ramp rate is

$$r(p, t) = \frac{\partial s(p, t)}{\partial t} = \frac{\partial^2 q(p, t)}{\partial t^2}.$$

We know the energy demand elasticity is

$$\eta_q(p, t) = \frac{p}{q(p, t)} \frac{\partial q(p, t)}{\partial p}$$

and if we assume that short term energy elasticity (viz., 1 hour or less) is the constant η_q , then we must have

$$\begin{aligned} q(p, t) &= \frac{p}{\eta_q} \frac{\partial q(p, t)}{\partial p} \\ \frac{\partial q(p, t)}{\partial t} &= \frac{\partial}{\partial t} \left(\frac{p}{\eta_q} \frac{\partial q(p, t)}{\partial p} \right) \\ \eta_q s(p, t) &= \frac{\partial p}{\partial t} \frac{\partial q(p, t)}{\partial p} + p \frac{\partial^2 q(p, t)}{\partial p \partial t} \\ &= \frac{\partial q(p, t)}{\partial t} + p \frac{\partial}{\partial p} \frac{\partial q(p, t)}{\partial t} \\ &= s(p, t) + p \frac{\partial s(p, t)}{\partial p} \\ (\eta_q - 1) s(p, t) &= p \frac{\partial s(p, t)}{\partial p} \end{aligned}$$

and thus the power demand elasticity

$$\eta_s = \frac{p}{s(p, t)} \frac{\partial s(p, t)}{\partial p} = \eta_q - 1$$

By similar reasoning we find that the ramp demand elasticity is

$$\eta_r = \frac{p}{r(p,t)} \frac{\partial r(p,t)}{\partial p} = \eta_s - 1$$

and we identify a previously unrecognized but potentially important relation between the energy, power and ramp response elasticities with respect to energy price:

$$\boxed{\eta_q = \eta_s + 1 = \eta_r + 2} \quad (\text{A.1})$$

We make the following observations based on Equation (A.1).

1. **When any one of the ramp, power, or energy demand elasticities is constant in time then they all are constant in time.** Thus it is only necessary to observe one constant elasticity to know them all.
2. **In the limit of absolutely inelastic energy demand, power demand elasticity is unitary.** This result implies that even though energy demand may be nearly inelastic, power and ramp demand can remain highly elastic.
3. **Ramp demand elasticity is always highly elastic.** This suggests that for ramping DR resources, the supplier of ramp resources never has market power.

A.2 Discussion

The absence of concurrent energy, power, and ramp pricing and DR data to validate Equation (A.1) is noted. We believe those with access to field data on DR elasticity should examine their data to confirm whether and under what conditions Equation (A.1) holds. In addition, we suggest the following research questions be considered in light of Equation (A.1) and the observations made above.

A.2.1 Device Control

If a device that demands energy is designed to elicit information from the consumer for the purposes of developing a demand curve, then Equation (A.1) implies that it is sufficient to obtain this information for only one of energy, power, or ramp responses and the other demand curves may be computed directly. It is not necessary to design

a separate transactive control strategy for the energy, power, and ramp behavior of the device.

In addition, this result suggests that policies to prevent resources from bidding concurrently in energy, capacity and regulation markets may not conform to the underlying dynamics of transactive systems. Devices cannot decouple these behaviors from each other even if they wanted to. This fits well with markets that concurrently clear energy, capacity, and regulation prices and ensures that no device is provided conflicting or inconsistent price signals from the markets.

A.2.2 Business Models

Utilities and load aggregators must consider the market power implications of Equation (A.1), insofar as they will often have monopoly market power as suppliers of energy DR resource while having monopsony market power as consumers of ramp response resources. We note in particular that because load is likely to supply ramping resources its market power is expected to be as low as it is for energy demand, while power response resources will likely be close to unitary elasticity. This is particularly important as utilities begin to shift revenue away from energy-based tariffs toward tariffs based on products and services that have greater downward substitutability.

However, if customers are purchasing ramping DR to mitigate the intermittency of their own on-site renewables such as solar photovoltaic (PV) panels, then their market power increases. The presence of large numbers of such PV- and DR-enabled consumers can be expected to create a rich and flexible retail market in which resources can be coordinated using only price signals. If the utility uses a business model based on revenue from trading activity rather than revenue from sales of net energy, capacity scarcity or ramping services, then it is likely to see greater stability in net revenue by becoming a market maker rather than a provider of last resort for these resources.

A.2.3 Regulatory Oversight

From the regulatory perspective, the utility as a market maker presents a new challenge. Historically, regulatory bodies have focused on authorizing retail tariffs because the utility is a natural monopoly. If the utility is a market maker reimbursed only for the cost of operating the system that enables trading and delivery among market

participants, then the regulator now ensures that the utility's market and operation costs are fair and equitable.

However, the regulator is now also concerned with whether the market is being manipulated by any of the participants. As a result, regulators must work with utilities to determine whether any customer or load aggregator is exerting excessive market power. The methods for this kind of monitoring are well-known from other markets. But we expect that certain particulars will be unique for retail electricity markets, particularly in light of the conditions that give rise to Equation (A.1). This problem is complicated by the tight coupling of energy, power, and ramping, as a result of which mitigating the utilities' energy monopoly power may not be sufficient to mitigate their monopsony market power for ramping services, or vice-versa.

A.3 Conclusion

Transactive energy facilitates integrated economic and technical scheduling, dispatch and control of demand response resources as intermittent renewable energy grows and challenges the economic viability, security and reliability of bulk electricity systems. We have shown that there exists an important and simple relationship between the energy price elasticities of energy demand, power capacity demand and ramping resources. Data collection from existing demand response systems is needed to validate this result. But the strong coupling of energy, power, and ramping response elasticities may have important consequences on how we design, deploy, and operate demand response controls, on which utility business models are preferred, and on how we adapt our regulatory oversight mechanisms to better monitor transactive energy systems.

Appendix B

Price Negotiation Convergence

Indirect dispatch systems use incentive signals such as real-time prices to “call” demand response. Most of these systems use day-ahead price signals [115]. Faster-acting 5-minute real-time pricing was also demonstrated successfully in the Olympic and Columbus studies. However, in all such systems computing the incentive signal to be dispatched can be a challenge. In particular, price-feedback systems have been shown to be potentially unstable when they are based on previous responses [82]. In the case of the 5-minute real-time pricing system, a retail double auction was used in which consumer bid prices above which they would forgo consumption for the next five minutes. The advantage of using auctions is that by eliminating the time delay in the feedback, a significant source of the system instability is mitigated.

Unfortunately, auction-based price discovery mechanisms are not always feasible or desirable. In the Pacific study an iterative price-discovery approach was proposed as an alternative to auction-based mechanisms [116]. In this paper we discuss the technical considerations regarding negotiated price-discovery mechanisms when applied to demand response dispatch problems. We address one of the principal problems in fast-acting indirect demand dispatch; i.e., computing the incentive signal necessary to achieve a particular level of demand response. We specifically examine the theoretical basis for a real-time negotiation-based price-discovery mechanisms such as the one tested in the Pacific demonstration and propose an approach for ensuring that such mechanisms robustly and reliably find the retail price at which supply will equal demand.

B.1 Transactive Price-Discovery

We first examine the technical considerations for utilities that wish to dispatch demand response for loads that present responses that are functionally unknown. When this situation arises, utilities must employ one of several possible mechanisms to discover the dispatch signal that will satisfy the physical constraints on the system. In this section we will consider one such mechanism with the understanding that the principles and methods apply more broadly to any iterative price-discovery mechanism a utility may wish to employ. The main contribution of this section is the derivation of a technique to evaluate limitations on iterative mechanisms used to determine the price at which supply equals demand. We examine this limitation to illustrate how one can use it to design stable real-time price discovery mechanisms for retail electricity markets.

Using arbitrary functional models of supply and demand we can analytically examine the behavior of iterative price discovery mechanisms used in systems that do not employ auction clearing for price-discovery. The simplest iterative price-discovery method is a negotiated price, in which the utility offers an initial hypothetical price $p(0)$ to which potential consumers respond either individually or in the aggregate with a hypothetical quantity $q(0)$. The utility follows up with a second proposed price $p(1)$ to which the consumers respond with a proposed quantity $q(1)$, followed by $p(2)$, $q(2)$ and $p(3)$, $q(3)$ and so on until the utility determines that the process has converged on a price that cannot be changed significantly without increasing the mismatch between supply and demand, or that the process must be stopped due to excessive iteration.

In such a process it is presumed that the functions used by suppliers and consumers to convert quantities to prices and prices to quantities, are respectively the supply and demand curves. These curves are not shared in their entirety with the other party, either because they are considered too business-sensitive to reveal (as is often the case with suppliers) or because they are not explicitly known to the party (as is often the case for consumers). The exchange is also presumed to be so limited that neither party can deduce the other's complete curve, while still sufficient to reliably deduce the dispatch price and quantity at which the two curves intersect.

The simplest possible iterative price-discovery process can be described using the

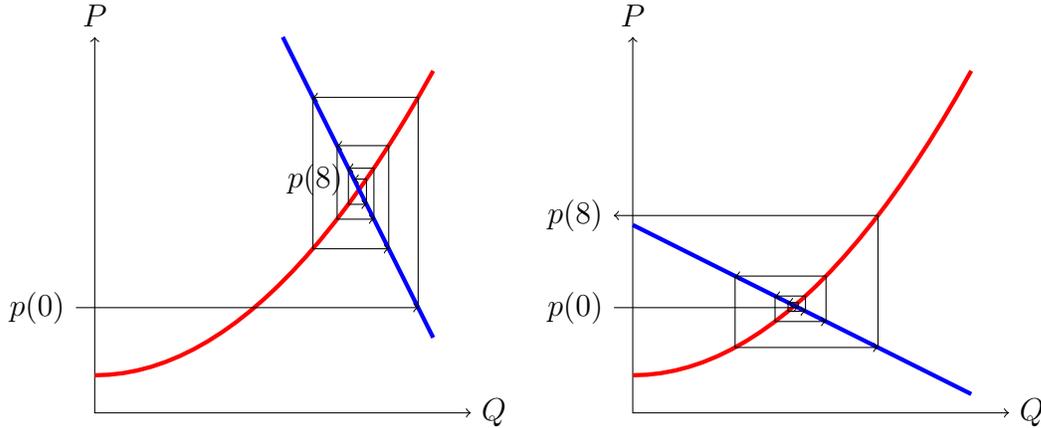


Figure B.1: Logistic map iteration sequence of price discovery mechanisms in a transactive system for $a < -b$ (left) and $a > -b$ (right).

iterative state equations for quantity and price

$$q(k) = \frac{1}{b}[p(k) - P] + Q \quad \text{and} \quad p(k+1) = a[q(k) - Q] + P,$$

where a and b are the slopes of the supply and demand curves, and P and Q are the clearing price and quantity, to which the negotiation process should converge. When the supply and demand curves are linear functions, we can define the iterative price-discovery process using a linear discrete-time state-space representation

$$\begin{bmatrix} p(k+1) \\ q(k+1) \end{bmatrix} = \begin{bmatrix} 0 & a \\ \frac{a}{b^2} & 0 \end{bmatrix} \begin{bmatrix} p(k) \\ q(k) \end{bmatrix} \quad \text{for } k = 0, 1, 2, \dots$$

The stability of this iterative price-discovery mechanism is determined by the magnitude of the roots of the system's characteristic equation $z^2 - \frac{a^2}{b^2} = 0$. Noting that $b < 0 < a$, we conclude that the negotiation can converge on the clearing price and quantity only when $a < -b$.

The impact of this stability condition is illustrated in Figure B.1 for two different combinations of a and b , one stable (top) and one unstable (bottom). When the slopes of the supply curve (a) and demand curves (b) in the neighborhood of the clearing price are such that their ratio (a/b) is less than -1 , the iterative price discovery mechanism fails to converge on the clearing price and quantity¹.

¹The scenario presented assumes that the negotiation always begins with an opening bid from the supplier. If the opening bid is from the consumer, the iteration direction shown in Figure B.1

This stability condition cannot be generally satisfied for all linear supply and demand curves, specifically when supply is very inelastic and/or demand is very elastic. Neither do we expect it to be satisfied for any general non-linear curves. We can evaluate whether a non-linear system will be stable by considering the following approximations for the supply and demand responses

$$\tilde{p}(k+1) = \alpha q(k)^2 + \beta q(k) + P_{min}$$

and

$$\tilde{q}(k) = Q_U + Q_R \frac{p(k)}{P_{max}}$$

respectively, in the neighborhood of the clearing price and quantity. This approximation is the canonical quadratic map problem based on the recurrence equation

$$x_n = \frac{1}{2} (1 - f(n) [r^n f(x)^{-1} (1 - 2x_0)])$$

where $f(n)$ is the time-domain solution function and $f(x)^{-1}$ is its inverse [117]. This system is known to be meta-stable (i.e. it has no single fixed solution point x_n) for values of $r > 3.5$, and chaotic for values of $r > 4$, with closed-form solutions to f known only for $r \in \{-2, 2, 4\}$.

To determine the stability of a linear negotiated market clearing mechanism when the supply and demand curves are not linear functions, we must evaluate the joint spectral radius for the linearized state space representation

$$\begin{bmatrix} p(k+1) \\ q(k+1) \end{bmatrix} = \underbrace{\begin{bmatrix} 0 & A(q(k)) \\ \frac{A(q(k))}{B^2(p(k))} & 0 \end{bmatrix}}_G \begin{bmatrix} p(k) \\ q(k) \end{bmatrix} \quad \text{for } k = 0, 1, 2, \dots$$

where $(A, B) \in G$ are the first order terms of the Taylor expansions

$$\begin{cases} P(x - q) = P(q) + \underbrace{P'(q)}_{A(q)}(x - q) + \frac{1}{2}P''(q)(x - q)^2 + \dots \\ Q(y - p) = Q(p) + \underbrace{Q'(p)}_{B(p)}(y - p) + \frac{1}{2}Q''(p)(y - p)^2 + \dots \end{cases}$$

is reversed and thus the stability condition is reversed as well. This suggests that one approach to addressing stability is to simply reverse the negotiation process when convergence is not achieved. However, this approach does not improve the convergence rate.

of the demand and supply curves about the price and quantity at the iteration k , respectively. Such a system is only stable when the mean spectral radius $\rho(G) < 1$. Thus we can only say that for systems with small amounts of demand response and typical supply curves, convergence can be expected when the supply and demand elasticities at the clearing price and quantity are such that $-B(P_c) < A(Q_c)$. If this is not true, the negotiation will converge only to a boundary region outside of which this condition is satisfied because within that region the process diverges. If the clearing price and quantity are inside the divergence region, then they cannot be discovered by the simple linear negotiation strategy described above.

In the case of logistic demand curves observed in transactive system field demonstrations to date, convergence is possible only when

$$\eta_D > -\frac{1}{\eta_S} \left(\frac{P_c}{Q_c} \right)^2 \quad (\text{B.1})$$

where η_S and η_D are the supply and demand elasticities at the clearing price, respectively. This limits the conditions for which convergence is possible when using linear negotiation strategies to only relatively inelastic demand curves in the neighborhood of the clearing price and quantity. The more elastic supply is, the less elastic demand must be for the negotiation to successfully converge on a price at which supply will equal demand. As has been observed in the field demonstration data, high demand elasticities do occur during period of constrained supply conditions. This is why we conclude that simple linear iterative negotiated price-discovery mechanisms are generally unsuitable for price-based demand response dispatch systems such as those used for transactive control.

B.2 Stable Mechanism Design

We can use the insights gained from the analysis above to devise a price discovery mechanism using an enhanced negotiating strategy for the utility that will satisfy its objective of quickly finding a price at which supply equals demand. The utility's second proposed price in response to the consumer's initial proposed quantity is augmented with a term that includes the last proposed price, such that

$$p(k+1) = [c - k_p]p(k) + [a - k_q]q(k)$$

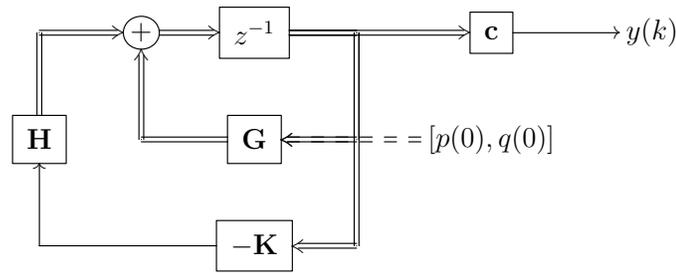


Figure B.2: Advanced negotiation strategy diagram with quantity constraint tracking

where c is a proportional coefficient for the previous price, and $\mathbf{K} = [k_p, k_q]$ are feedback coefficients that we will use to tune the relative input of the previous price and quantity. The design of this system is illustrated in Figure B.2 and the state-space representation of this advanced negotiation strategy for any reference quantity input $r(k)$ is

$$\begin{bmatrix} p(k+1) \\ q(k+1) \end{bmatrix} = \underbrace{\begin{bmatrix} c & a \\ \frac{a}{b^2} + c & 0 \end{bmatrix}}_{\mathbf{G}} \begin{bmatrix} p(k) \\ q(k) \end{bmatrix} - \underbrace{\begin{bmatrix} 1 \\ \frac{1}{b} \end{bmatrix}}_{\mathbf{H}} \underbrace{\begin{bmatrix} k_p & k_q \end{bmatrix}}_{\mathbf{K}} \begin{bmatrix} p(k) \\ q(k) \end{bmatrix} \quad (\text{B.2})$$

with the output quantity

$$y(k) = \underbrace{\begin{bmatrix} 0 & 1 \end{bmatrix}}_{\mathbf{c}} \begin{bmatrix} p(k) \\ q(k) \end{bmatrix}.$$

The negotiation strategy design problem for the utility is then reduced to determining the value \mathbf{K} such that negotiation converges as quickly as possible on the price at which supply equals demand. This “deadbeat” system response requires that the characteristic equation be reduced to $z^2 = 0$, a condition which can be obtained when

$$\mathbf{K} = \left[c - \frac{a^2}{b^2 c (b-1)} - \frac{a}{b}, a + \frac{a^2}{bc(b-1)} \right]. \quad (\text{B.3})$$

An example for three different levels of demand response is shown in Figure B.3, where the demand curve has been linearized in the neighborhood of the clearing price and quantity, as described above. In all three cases, the deadbeat negotiating strategy converges in two iterations to the clearing price and quantity, outperforming the simple linear negotiation strategy for even the stable case. More significantly, the deadbeat strategy converges when the simple strategy fails to converge.

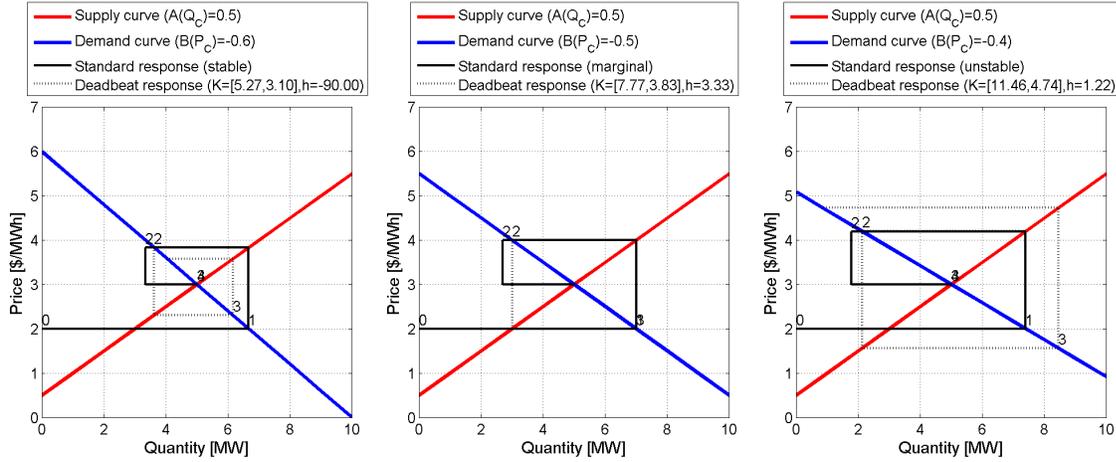


Figure B.3: Simulation of stable (left), marginal (center) and unstable (right) negotiations without (dotted) and with (solid) corresponding deadbeat negotiation strategies.

B.2.1 Demand Response Uncertainty

We note that in Equation (B.3) the slope of the demand curve b is included in the computation of the negotiation strategy parameters of \mathbf{K} . Uncertainty in b can result in an error in the clearing price and quantity. The demand curve may not be known exactly or it may change from time to time. The utility may wish to employ an enhanced negotiation strategy to reduce the clearing price error that may result from inaccurate estimation of this demand curve parameter b .

A reliable design method based on control theory is to compensate for the unknown demand function by implementing integral error feedback [87] in the negotiation process. This approach adds a third state to the state-space representation in Equation (B.2). This new state represents the accumulated error between the most recent quantity from the consumer(s) and the clearing quantity. This error is then multiplied by a gain k_e and the result is added to the price response $p(k+1)$. This approach raises the overall order of the system by one and can be expected to decrease the convergence rate compared to the deadbeat negotiation. But integral error feedback has the significant advantage that it compensates for all constant errors in the system, not just demand elasticity estimation errors. The same method used to find \mathbf{K} above can be used in this case to find the joint feedback gains $[\mathbf{K}, k_e]$ required to obtain the fastest possible convergence using a linear negotiation. The behavior of this price-discovery mechanism is illustrated in Figure B.4.

The solution for $[\mathbf{K}, k_e]$ using this approach can be found given any reasonable

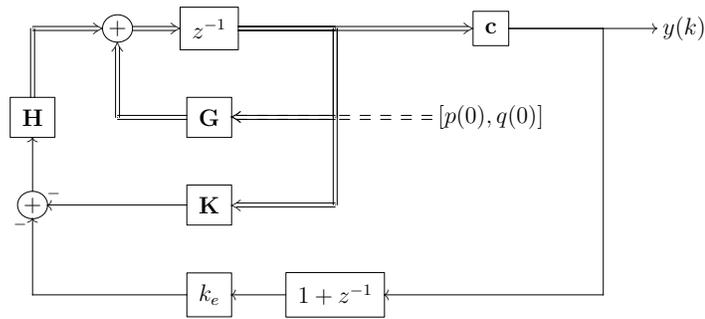


Figure B.4: Discrete-time system diagram of advanced negotiation strategy diagram with demand curve uncertainty

assumption for the value of b . Choosing a value of $\hat{b} = -1$ gives the solution for the feedback gains

$$k_p = 1, \quad k_q = a + \frac{1}{a}, \quad \text{and} \quad k_e = 1 - \frac{1}{a}, \quad . \quad (\text{B.4})$$

The general solution for the augmented state approach is

$$k_p = 1, \quad k_q = a + \frac{\hat{b}^2}{a}, \quad \text{and} \quad k_e = 1 - \frac{\hat{b}^2}{a}$$

where \hat{b} is the utility's estimate for the slope b of the demand response curve, as illustrated in Figure B.5. The general solution can be expressed in terms of a utility's price negotiation strategy as

$$p(k+1) = p(k) + (2a+1)q(k) + \left(1 - \frac{\hat{b}^2}{a}\right) \sum_{j=1}^{k-1} q(j)$$

which will converge in a finite time when $b < 0 < a$ but with no particular restriction on the relative magnitudes of a and b .

The closed-loop stability of this solution depends on the magnitude of the error in the estimate of \hat{b} . Specifically, the characteristic equation of this price-discovery solution is

$$z(z^2 - 1 + \frac{\hat{b}^2}{b^2})$$

which suggests that convergence is guaranteed when

$$\sqrt{2}b < \hat{b} < 0. \quad (\text{B.5})$$

The constraint that the magnitude \hat{b} cannot be more than about 40% greater than b

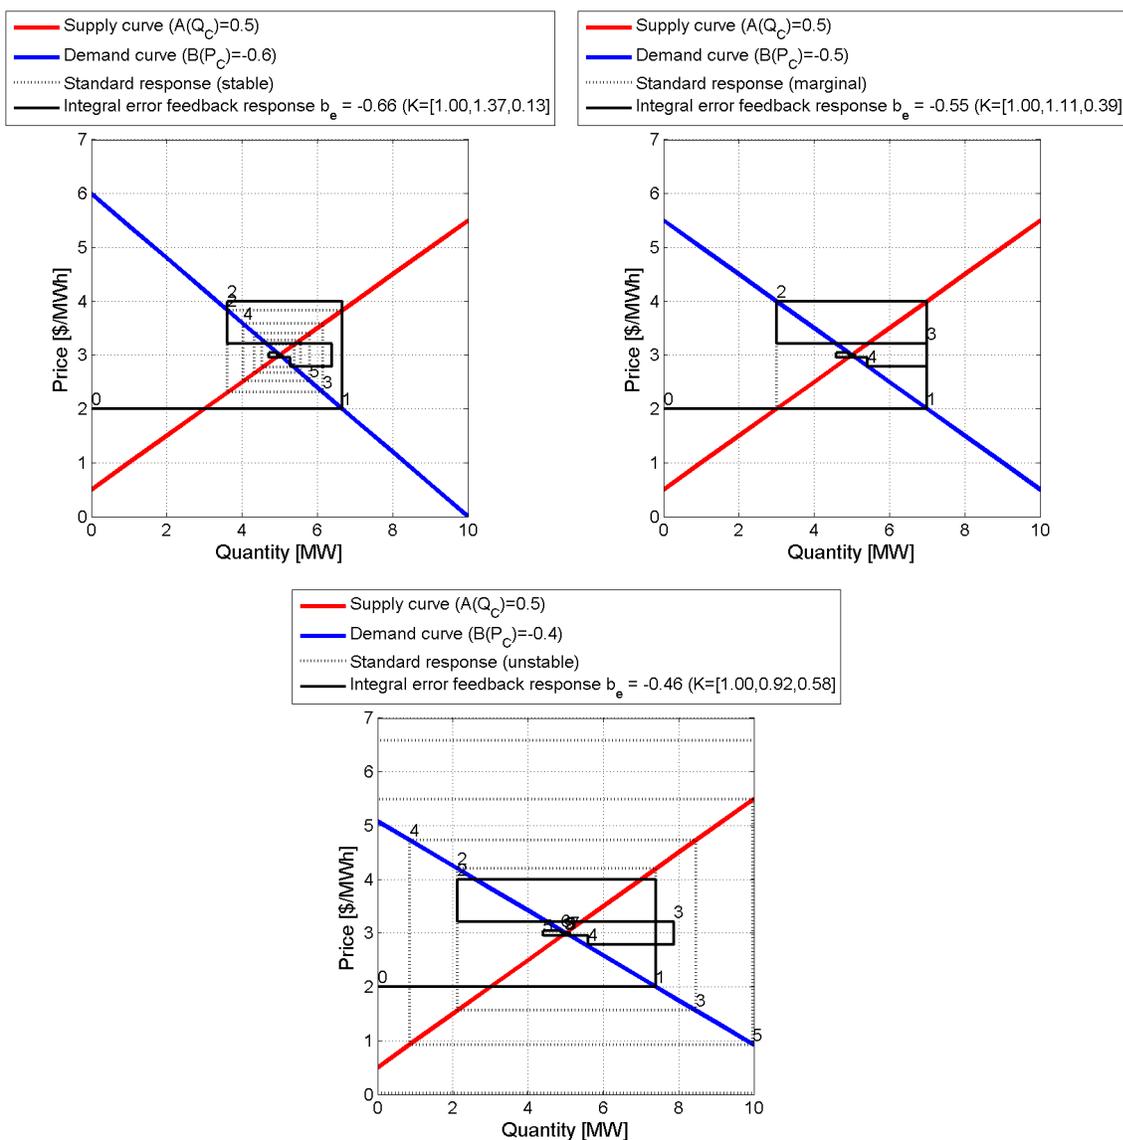


Figure B.5: Integral error feedback negotiation strategy for the same cases as Figure B.3 with a +10% error in the demand response curve estimate \hat{b} .

is a generally reasonable one for the slope of the demand curve in the neighborhood of the solution price and quantity.

B.3 Conclusions

The dependence of deadbeat negotiation strategy on the utility's knowledge of the demand curve slope b suggests that it is possible to compute the clearing price using a

parametric fit of the demand curve from a very large set of bids instead of negotiating with all the loads individually. Instead of performing an auction clearing as in the Olympic and Columbus demonstration projects the utility can indirectly determine the price at which supply equals demand by computing what the negotiation would produce given the demand curve imputed by the bids *without actually performing the negotiation*.

This method does assume that the demand curve fits a mathematical function of some type that can be expected to represent the loads' collective behavior most of the time. For example, the demand curve may generally take the form of a logistic function as observed during quiescent periods in both the Olympic and Columbus demonstrations. This function was derived in Chapter 3 and has the form

$$Q = (1 + e^{\alpha + \beta P})^{-1} Q_R + Q_U$$

where α is the unobservable component of the demand response behavior and β is the observable component. The parameters can be estimated using the logistic regression

$$\alpha = \bar{P} - \beta \bar{\xi}$$

and

$$\beta = \frac{\overline{P\xi} - \bar{P}\bar{\xi}}{\overline{P^2} - \bar{P}^2}.$$

where $\xi = \ln(Q_U + Q_R - Q) - \ln(Q - Q_U)$. Given such a fit we can estimate the slope of the demand curve at the clearing price

$$\hat{b} = -\frac{\beta Q_R e^{\alpha + \beta P}}{(e^{\alpha + \beta P} + 1)^2},$$

a value that can be readily used to compute the clearing price without performing the full market clearing procedure.

In control theory a non-zero reference quantity $r(k)$ is included in Equation (B.2). This term merits further consideration because it can be used by the utility to change the quantity that the negotiation will converge to. However, this quantity does not necessarily converge to a price at which supply will equal demand. Non-equilibrium values of $r(k)$ may be theoretically interesting, but they likely do not have physical meaning that is useful unless the utility intentionally wishes to increase or decrease the load for some operational reason, such as maintaining a net level of imports or exports.

Generally such deliberate manipulation of the negotiation strategy should be regarded with skepticism, particularly if the result would be an economically inefficient price or a power imbalance that results in operational reliability concerns.

Control theory also suggests that when using the integral error feedback negotiation strategy, a better estimate \hat{b} of the demand response slope b results in a faster convergence for the negotiation. But the process will always converge regardless of the error. Here again, collecting real-time demand data from customers can contribute to significantly improving the performance of the price-discovery mechanism by providing information needed to obtain a sufficiently accurate estimate of the demand curve slope.

Appendix C

Model Specifications

Table C.1: Aggregate Load Model Parameters

Parameter	Variable	Distribution	μ	σ	Remarks
Wall conductivity	U_A	Normal	350	50	Truncated at $\pm 3\sigma$
Air heat capacity	C_A	Normal	2000	150	Truncated at $\pm 3\sigma$
Mass conductivity	U_M	Normal	2000	500	Truncated at $\pm 3\sigma$
Mass heat capacity	C_M	Normal	10000	1000	Truncated at $\pm 3\sigma$
Setpoint temperature	T_S	Normal	72	1	Truncated at $\pm 3\sigma$
Design temperature	T_H	None	10	–	
Equipment oversizing	$-\frac{Q_H}{(T_H - T_S)U_A}$	Normal	1.55	0.15	Truncated at $\pm 3\sigma$
Equipment efficiency	COP	Normal	3.0	0.5	Truncated at $\pm 3\sigma$

Table C.2: WECC 2024 demand forecast and internal area losses [1]

Area number	Consolidated area	Peak [MW]	Energy [GWh/year]	Loss factor
1	AESO	16095	115061	1.019
2	BCH	12542	69326	1.020
3	PNW = AVA + BPA + CHPD + DOPD + GCPD + PACW + PGE + PSEI + SCL + TPWR	33384	184103	1.025
4	NWMT	1898	12163	1.023
5	PAWY	1681	11028	1.013
6	NCA = BANC + CIPB + CIPV + TIDC	30626	144848	1.043
7	SPPC	2447	15784	1.026
8	ID = IPFE + IPMV + IPTV	4157	19290	1.036
9	UT = PAID + PAUT	8443	39362	1.028
10	CO = WACM + WAUW	5867	34863	1.023
11	LDWP	7789	34129	1.027
12	NEVP	7034	30083	1.045
13	PSCO	8130	41027	1.028
14	SCA = CISC + VEA	26847	119573	1.040
15	AZ = AZPS + SRP + TEPC + WALC	23596	109534	1.026
16	CISD	5573	26702	1.037
17	IID	1342	4836	1.043
18	PNM	3136	16449	1.026
19	CEF	3255	15452	1.033
20	EPE	2391	11106	1.032

Table C.3: WECC 2024 aggregated installed supply capacity [1]

Area	Non-dispatchable (must-take) generation		Dispatchable generation [MW]
	Intermittent [MW]	Base (invariable) [MW]	
AESO	2275.2	3710	17733
BCH	815.3	3531	17000
PNW	14432.0	5845	36675
NWMT	664.2	1227	1993
PAWY	1315.3	1648	2105
NCA	5092.7	5351	38684
SPPC	800.0	1118	3287
ID	660.3	374	2845
UT	256.5	2183	8136
CO	656.4	1807	4310
LDWP	687.0	89	8727
NEVP	75.7	426	13085
PSCO	2441.1	1616	11645
SCA	6028.1	977	28645
AZ	3220.6	8601	39353
CISD	432.8	34	8558
IID	34.4	1170	1514
PNM	840.3	910	4504
CFE	384.2	697	6600
EPE	1.3	0	3512

Table C.4: WECC 2024 producer cost and surplus difference for 100% inelastic load (in M\$/year) [1]

	Producer cost reduction		Producer surplus reduction	
	Unconstrained	Constrained	Unconstrained	Constrained
AESO	6642	1797	11421	2820
BCH	55	149	79	284
PNW	35	479	-552	588
NWMT	-260	-236	-1622	-1516
PAWY	-297	-246	-2513	-2132
NCA	-393	121	-817	145
SPPC	-297	-254	-1058	-933
ID	1873	1907	2630	2696
UT	453	557	763	997
CO	1244	1296	2978	3147
LDWP	2304	2422	2498	2632
NEVP	-183	-4	-214	-8
PSCO	-454	-306	-1044	-753
SCA	5697	6083	7215	7771
AZ	-4624	-4082	-12111	-10968
CISD	992	1108	1061	1187
IID	-213	-192	-1467	-1375
PNM	-324	-262	-873	-735
CFE	-521	-287	-873	-506
EPE	573	622	573	622
Total	12301	10671	6075	3965

Table C.5: WECC 2025 production cost per unit in \$/MWh [1]

	Standalone		Constrained	
	100% inelastic	95% inelastic	100% inelastic	95% inelastic
AESO	78.48	78.47	67.51	67.50
BCH	32.72	32.71	31.31	31.30
PNW	26.33	26.31	24.64	24.63
NWMT	1.45	1.44	14.78	14.77
PAWY	0.00	0.00	11.11	11.10
NCA	32.14	32.13	31.68	31.67
SPPC	9.76	9.75	21.29	21.28
ID	109.06	109.07	32.83	32.79
UT	35.66	35.62	26.92	26.91
CO	50.93	50.92	19.19	19.18
LDWP	97.06	97.04	51.01	51.01
NEVP	49.45	49.44	49.91	49.90
PSCO	26.24	26.23	31.62	31.61
SCA	76.38	76.37	41.27	41.26
AZ	6.11	6.11	28.10	28.09
CISD	79.17	79.17	52.90	52.90
IID	0.01	0.01	13.82	13.81
PNM	16.58	16.57	27.71	27.71
CFE	21.43	21.43	33.10	33.11
EPE	90.20	90.19	57.33	57.33
Average	41.96	41.95	33.40	33.39

Table C.6: WECC 2024 model inter-area transfer limits [1]

Path number	From area	To area	Minimum [MW]	Maximum [MW]
P01	AESO	BCH	-1200	1000
P03	PNW	BCH	-3150	3000
P08	NWMT	PNW	-2150	3000
P09	NWMT	CO	-2573	2573
P14	ID	PNW	-2250	3400
P16	ID	SPPC	-360	500
P18	NWMT	ID	-256	337
P20	ID	UT	-2250	2250
P22	PNM	AZ	-2325	2325
P24	NCA	SPPC	-150	160
P26	NCA	SCA	-3000	4000
P27	UT	LDWP	-1400	2400
P30	UT	PSCO	-650	650
P31	PSCO	PNM	-690	690
P32	SPPC	UT	-235	440
P35	UT	NEVP	-580	600
P36	CO	PSCO	-1680	1680
P42	IID	SCA	-1500	1500
P43	LDWP	SCA	-4000	4000
P44	CISD	SCA	-2500	2500
P45	CISD	CFE	-800	400
P46N	NEVP	LDWP	-6000	6000
P46S	NEVP	SCA	-5000	5000
P47	EPE	AZ	-1048	1048
P48	EPE	PNM	-1970	1970
P49	AZ	NEVP	-10200	10200
P59	AZ	SCA	-218	218
P65	PNW	LDWP	-3100	3220
P66	PNW	NCA	-3675	4800
P76	PNW	SPPC	-300	300
P78	UT	PNM	-600	600
P79	AZ	UT	-300	265
P80	NWMT	PAWY	-600	600
PP1	PAWY	UT	-1700	1700
PP2	IID	CISD	-150	150
PP3	AZ	CISD	-1160	1650
PP4	AZ	IID	-160	260

Appendix D

Simulation Results

Table D.1: Estimated, simulated, and errors of aggregate thermostat load state transition probability ρ using joint PDF (N=normal, Ln=log normal, Log=logistic) using 10^6 homes and $t_d = 1$ minute.

T_O		N(X)	N(Y)	N(X)	Ln(Y)	Log(X)	N(Y)	Log(X)	Ln(Y)
(°F)		ρ_{off}	ρ_{on}	ρ_{off}	ρ_{on}	ρ_{off}	ρ_{on}	ρ_{off}	ρ_{on}
5	Est	1.00	0.44	1.00	0.44	1.00	0.43	1.00	0.43
	Sim	1.00	0.42	1.00	0.42	1.00	0.42	1.00	0.42
	Err	0.1%	-3.5%	0.1%	-3.5%	0.1%	-2.4%	0.1%	-2.3%
10	Est	0.99	0.55	1.00	0.55	0.99	0.54	1.00	0.54
	Sim	1.00	0.54	1.00	0.54	1.00	0.54	1.00	0.54
	Err	0.5%	-2.3%	0.4%	-2.4%	0.5%	-1.3%	0.5%	-1.4%
15	Est	0.98	0.67	0.98	0.67	0.98	0.67	0.98	0.67
	Sim	1.00	0.67	1.00	0.67	1.00	0.67	1.00	0.67
	Err	1.8%	-0.5%	1.5%	-0.8%	1.9%	0.3%	1.6%	0.1%
20	Est	0.94	0.80	0.95	0.80	0.94	0.79	0.95	0.79
	Sim	0.98	0.81	0.98	0.81	0.98	0.81	0.98	0.81
	Err	4.2%	2.1%	3.4%	1.8%	4.3%	2.8%	3.5%	2.4%
25	Est	0.86	0.90	0.87	0.90	0.86	0.89	0.87	0.90
	Sim	0.91	0.93	0.91	0.93	0.91	0.93	0.91	0.93
	Err	5.5%	3.7%	4.4%	3.4%	5.7%	4.2%	4.7%	3.8%
30	Est	0.75	0.96	0.75	0.96	0.74	0.96	0.75	0.96
	Sim	0.78	0.99	0.78	0.99	0.78	0.99	0.78	0.99
	Err	4.4%	2.7%	3.6%	2.5%	4.6%	3.0%	3.9%	2.8%
35	Est	0.62	0.99	0.62	0.99	0.62	0.99	0.62	0.99
	Sim	0.64	1.00	0.64	1.00	0.64	1.00	0.64	1.00
	Err	2.6%	1.1%	2.4%	1.0%	2.9%	1.2%	2.7%	1.1%
40	Est	0.50	1.00	0.50	1.00	0.50	1.00	0.50	1.00
	Sim	0.51	1.00	0.51	1.00	0.51	1.00	0.51	1.00
	Err	1.7%	0.3%	1.7%	0.3%	1.9%	0.3%	1.9%	0.3%

Bibliography

- [1] S. Behboodi, D. Chassin, N. Djilali, and C. Crawford, “Interconnection-wide hour-ahead scheduling in the presence of intermittent renewables and demand response: a surplus maximizing approach,” Applied Energy, vol. 189, no. 1, pp. 336–351, 2017.
- [2] N. A. of Engineer, “Greatest Engineering Achievement of the 20th Century,” 2017.
- [3] D. C. Josephson, J. M. Robinson, J. Chiotti, K. J. Jirka, and C. E. Kraft, “Chemical and biological recovery from acid deposition within the honnedaga lake watershed, new york, usa,” Environmental Monitoring and Assessment, vol. 186, no. 7, pp. 4391–4409, 2014.
- [4] C. T. Driscoll, K. M. Driscoll, H. Fakhraei, and K. Civerolo, “Long-term temporal trends and spatial patterns in the acid-base chemistry of lakes in the adirondack region of new york in response to decreases in acidic deposition,” Atmospheric Environment, vol. 146, pp. 5 – 14, 2016.
- [5] J. Blazejczak, F. G. Braun, D. Edler, and W.-P. Schill, “Economic effects of renewable energy expansion: A model-based analysis for germany,” Renewable and Sustainable Energy Reviews, vol. 40, pp. 1070 – 1080, 2014.
- [6] M. S. Markoff and A. C. Cullen, “Impact of climate change on Pacific Northwest hydropower,” Climatic Change, vol. 87, no. 3-4, pp. 451–469, 2008.
- [7] M. K. Chandel, L. F. Pratson, and R. B. Jackson, “The potential impacts of climate-change policy on freshwater use in thermoelectric power generation,” Energy Policy, vol. 39, no. 10, pp. 6234 – 6242, 2011.

- [8] S. C. Parkinson and N. Djilali, “Robust response to hydro-climatic change in electricity generation planning,” Climatic Change, vol. 130, no. 4, pp. 475–489, 2015.
- [9] M. A. Ortega-Vazquez and D. S. Kirschen, “Estimating the spinning reserve requirements in systems with significant wind power generation penetration,” Power Systems, IEEE Transactions on, vol. 24, pp. 114–124, Feb 2009.
- [10] G. Blackburn, C. Magee, and V. Rai, “Solar valuation and the modern utility’s expansion into distributed generation,” The Electricity Journal, vol. 27, no. 1, pp. 18–32, 2014.
- [11] M. J. E. Alam, K. M. Muttaqi, and D. Sutanto, “Mitigation of rooftop solar PV impacts and evening peak support by managing available capacity of distributed energy storage systems,” Power Systems, IEEE Transactions on, vol. 28, pp. 3874–3884, Nov 2013.
- [12] I. Dobson, B. A. Carreras, V. E. Lynch, and D. E. Newman, “Complex systems analysis of series of blackouts: Cascading failure, critical points, and self-organization,” Chaos: An Interdisciplinary Journal of Nonlinear Science, vol. 17, no. 2, p. 026103, 2007.
- [13] G. Boyle, Renewable electricity and the grid: the challenge of variability. Earthscan, 2012.
- [14] Y. Makarov, C. Loutan, J. Ma, and P. de Mello, “Operational impacts of wind generation on California power systems,” Power Systems, IEEE Transactions on, vol. 24, no. 2, pp. 1039–1050, 2009.
- [15] M. Nyberg, “Thermal efficiency of gas-fired generation in california: 2012 update,” California Energy Commission, 2013.
- [16] S. Behboodi, D. Chassin, C. Crawford, and N. Djilali, “Renewable resources portfolio optimization in the presence of demand response,” Applied Energy, vol. 162, pp. 139–1418, 2016.
- [17] S. Nolan and M. OMalley, “Challenges and barriers to demand response deployment and evaluation,” Applied Energy, vol. 152, pp. 1–10, 2015.

- [18] S. Nolan, O. Neu, and M. OMalley, “Capacity value estimation of a load-shifting resource using a coupled building and power system model,” Applied Energy, vol. 192, pp. 71–82, 2017.
- [19] D. P. Chassin and N. Djilali, “Multi-scale transactive control in interconnected bulk power systems under high renewable energy supply and high demand response scenarios: Auxiliary report,” arXiv:1711.09704 [nlin.AO], November 2017.
- [20] V. Smith and L. Kiesling, “Demand, Not Supply,” Wall Street Journal Online, vol. 8, no. 20, pp. 1–3, 2003.
- [21] S. Borenstein, “The trouble with electricity markets: understanding California’s restructuring disaster,” Journal of economic perspectives, pp. 191–211, 2002.
- [22] W. W. Hogan, “Electricity market design and efficient pricing: Applications for new england and beyond,” The Electricity Journal, vol. 27, no. 7, pp. 23–49, 2014.
- [23] California ISO, “California ISO demand response and energy efficiency roadmap: Maximizing preferred resources,” technical report, CAISO, December 2013.
- [24] D. J. Hammerstrom, R. Ambrosio, T. a. Carlon, J. G. Desteese, R. Kajfasz, R. G. Pratt, and D. P. Chassin, “Pacific Northwest GridWise(TM) Testbed Demonstration Projects Part I . Olympic Peninsula Project,” Tech. Rep. PNNL-17167, Pacific Northwest National Laboratory, October 2007.
- [25] S. E. Widergren, K. Subbarao, J. C. Fuller, D. P. Chassin, A. Somani, M. C. Marinovici, and J. L. Hammerstrom, “AEP Ohio gridsmart demonstration project real-time pricing demonstration analysis,” PNNL Report, vol. 23192, 2014.
- [26] A. K. Bejestani, A. Annaswamy, and T. Samad, “A hierarchical transactive control architecture for renewables integration in smart grids: Analytical modeling and stability,” Smart Grid, IEEE Transactions on, vol. 5, no. 4, pp. 2054–2065, 2014.

- [27] C. W. Gellings, "Power/energy: Demand-side load management: The rising cost of peak-demand power means that utilities must encourage customers to manage power usage," IEEE Spectrum, vol. 18, pp. 49–52, Dec 1981.
- [28] 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, vol. 50, no. 5, pp. 1389–1400, 2009.
- [29] S. Parkinson, D. Wang, C. Crawford, and N. Djilali, "Comfort-constrained distributed heat pump management," Energy Procedia, vol. 12, no. 0, pp. 849 – 855, 2011.
- [30] D. Wang, S. Parkinson, W. Miao, H. Jia, C. Crawford, and N. Djilali, "Online voltage security assessment considering comfort-constrained demand response control of distributed heat pump systems," Applied Energy, vol. 96, pp. 104–114, 2012.
- [31] J. Mathieu, M. Dyson, D. Callaway, and A. Rosenfeld, "Using residential electric loads for fast demand response: The potential resource and revenues, the costs, and policy recommendations," Proceedings of the ACEEE Summer Study on Buildings, Pacific Grove, CA, vol. 1000, no. 2000, p. 3000, 2012.
- [32] S. Koch, J. L. Mathieu, and D. S. Callaway, "Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services," in Proc. PSCC, pp. 1–7, 2011.
- [33] T. Broeer, J. C. Fuller, F. Tuffner, D. P. Chassin, and N. Djilali, "Modeling framework and validation of a smart grid and demand response system for wind power integration," Applied Energy, vol. 113, pp. 199–207, 2014.
- [34] F. C. Schweppe, R. D. Tabors, J. L. Kirtley, H. R. Outhred, F. H. Pickel, and A. J. Cox, "Homeostatic utility control," Power Apparatus and Systems, IEEE Transactions on, vol. PAS-99, pp. 1151–1163, May 1980.
- [35] B. A. Huberman and S. H. Clearwater, "A multi-agent system for controlling building environments," in ICMAS, pp. 171–176, 1995.
- [36] J. K. Kok, C. J. Warmer, and I. Kamphuis, "PowerMatcher: multiagent control in the electricity infrastructure," in Proceedings of The Fourth International

- Joint Conference on Autonomous Agents and Multiagent Systems, pp. 75–82, ACM, 2005.
- [37] J. E. Dagle, D. W. Winiarski, and M. K. Donnelly, “End-use load control for power system dynamic stability enhancement,” Tech. Rep. PNNL-11488, Pacific Northwest National Laboratory, Richland, WA (United States), February 1997.
- [38] K. Subbarao, J. C. Fuller, K. Kalsi, R. G. Pratt, S. E. Widergren, and D. P. Chassin, “Transactive control and coordination of distributed assets for ancillary services,” Tech. Rep. PNNL-22942, Pacific Northwest National Laboratory, Richland, WA (United States), September 2013.
- [39] N. Lu and D. P. Chassin, “A state queueing model of thermostatically controlled appliances,” in Power Systems Conference and Exposition, 2004. IEEE PES, pp. 59 vol.1–, Oct 2004.
- [40] W. Zhang, J. Lian, C. Y. Chang, and K. Kalsi, “Aggregated modeling and control of air conditioning loads for demand response,” Power Systems, IEEE Transactions on, vol. 28, pp. 4655–4664, Nov 2013.
- [41] J. Mathieu, M. Dyson, D. Callaway, and A. Rosenfeld, “Using residential electric loads for fast demand response: The potential resource and revenues, the costs, and policy recommendations,” Proceedings of the ACEEE Summer Study on Buildings, Pacific Grove, CA, vol. 1000, no. 2000, p. 3000, 2012.
- [42] J. Fuller, K. Schneider, and D. Chassin, “Analysis of residential demand response and double-auction markets,” in 2011 IEEE Power and Energy Society General Meeting, pp. 1–7, IEEE, 2011.
- [43] J. Lazar et al., Teaching the “duck” to Fly. Regulatory Assistance Project, 2016.
- [44] F. C. Schweppe, R. D. Tabors, J. L. Kirtley, H. R. Outhred, F. H. Pickel, and A. J. Cox, “Homeostatic utility control,” IEEE Transactions on Power Apparatus and Systems, vol. PAS-99, pp. 1151–1163, May 1980.
- [45] A. Annaswamy and T. Nudell, “Transactive control – what’s in a name?,” September 2015.

- [46] D. Chassin, S. Behboodi, C. Crawford, and N. Djilali, “Agent-based simulation for interconnection-scale renewable integration and demand response studies,” Engineering, vol. 1, no. 4, pp. 422–435, 2015.
- [47] D. J. Hammerstrom and other, “Pacific northwest smart grid demonstration project technology performance report volume 1: Technology performance,” Tech. Rep. PNWD-4438, Pacific Northwest National Laboratory, Richland WA (USA), June 2015.
- [48] D. J. Hammerstrom, J. Brous, D. P. Chassin, G. R. Horst, R. Kajfasz, P. Michie, T. V. Oliver, T. A. Carlon, C. Eustis, O. M. Jarvegren, et al., “Pacific northwest gridwise™testbed demonstration projects; part ii. grid friendly™appliance project,” PNNL Report 17079, Pacific Northwest National Laboratory (PNNL), Richland, WA (US), 2007.
- [49] A. Kiani and A. Annaswamy, “A hierarchical transactive control architecture for renewables integration in smart grids,” in 2012 IEEE 51st Annual Conference on Decision and Control (CDC), pp. 4985–4990, Dec 2012.
- [50] S. Stoft, “Power system economics,” Journal of Energy Literature, vol. 8, pp. 94–99, 2002.
- [51] D. L. McFadden, “Quantal choice analysis: A survey,” in Annals of Economic and Social Measurement, Volume 5, number 4, pp. 363–390, NBER, 1976.
- [52] A. Faruqui and S. Sergici, “Household response to dynamic pricing of electricity: a survey of 15 experiments,” Journal of regulatory Economics, vol. 38, no. 2, pp. 193–225, 2010.
- [53] J. S. Vardakas, N. Zorba, and C. V. Verikoukis, “A survey on demand response programs in smart grids: pricing methods and optimization algorithms,” IEEE Communications Surveys & Tutorials, vol. 17, no. 1, pp. 152–178, 2015.
- [54] A. Faruqui and J. R. Malko, “The residential demand for electricity by time-of-use: a survey of twelve experiments with peak load pricing,” Energy, vol. 8, no. 10, pp. 781–795, 1983.
- [55] D. W. Caves, L. R. Christensen, and J. A. Herriges, “Consistency of residential customer response in time-of-use electricity pricing experiments,” Journal of Econometrics, vol. 26, no. 1, pp. 179–203, 1984.

- [56] F. Wolak, “Residential customer response to real-time pricing: the Anaheim Critical-Peak Pricing Experiment. Work. Pap. 151, Energy Inst. Cent. Study Energy Mark,” Univ. Calif. Berkeley, 2006.
- [57] L. Kiesling, “Digital technology, demand response, and customer choice: Efficiency benefits,” in NARUC Winter Meetings, 2008.
- [58] A. Faruqui, R. Hledik, and J. Tsoukalis, “The Power of Dynamic Pricing,” The Electricity Journal, vol. 22, pp. 42–56, Apr. 2009.
- [59] A. David and Y. Li, “Effect of inter-temporal factors on the real time pricing of electricity,” IEEE Transactions on Power Systems (Institute of Electrical and Electronics Engineers);(United States), vol. 8, no. 1, 1993.
- [60] D. S. Kirschen, G. Strbac, P. Cumperayot, and D. de Paiva Mendes, “Factoring the elasticity of demand in electricity prices,” IEEE Transactions on Power Systems, vol. 15, no. 2, pp. 612–617, 2000.
- [61] D. P. Chassin, “The abstract machine model for transaction-based system control,” Technical Report PNNL-12082, Pacific Northwest National Laboratory, November 2002.
- [62] D. P. Chassin, J. M. Malard, C. Posse, A. Gangopadhyaya, N. Lu, S. Katipamula, and J. V. Mallow, “Modeling power systems as complex adaptive systems,” Technical Report PNNL-14987, Pacific Northwest National Laboratory, December 2004.
- [63] H. Glavitsch and F. Alvarado, “Management of multiple congested conditions in unbundled operation of a power system,” in Power Industry Computer Applications., 1997. 20th International Conference on, pp. 374–380, IEEE, 1997.
- [64] F. L. Alvarado, “Is system control entirely by price feasible?,” in System Sciences, 2003. Proceedings of the 36th Annual Hawaii International Conference on, p. 53b, IEEE, 2003.
- [65] US DOE, “U.S. Department of Energy LDRD Project List – Annual Report,” 2014.
- [66] S. Fan and R. J. Hyndman, “The price elasticity of electricity demand in south australia,” Energy Policy, vol. 39, no. 6, pp. 3709–3719, 2011.

- [67] P. C. Reiss and M. W. White, “Household electricity demand, revisited,” The Review of Economic Studies, vol. 72, no. 3, pp. 853–883, 2005.
- [68] D. P. Chassin and L. Kiesling, “Decentralized coordination through digital technology, dynamic pricing, and customer-driven control: The gridwise testbed demonstration project,” The Electricity Journal, vol. 21, no. 8, pp. 51–59, 2008.
- [69] I. Beil, I. Hiskens, and S. Backhaus, “Frequency regulation from commercial building HVAC demand response,” Proceedings of the IEEE, vol. 104, no. 4, pp. 745–757, 2016.
- [70] K. G. Pillai and C. Hofacker, “Calibration of consumer knowledge of the web,” International Journal of Research in Marketing, vol. 24, no. 3, pp. 254–267, 2007.
- [71] R. D. Luce, Individual choice behavior: A theoretical analysis. Wiley, New York, 1959.
- [72] K. Train, Qualitative Choice Analysis: Theory, Econometrics, and an Application to Automobile Demand. The MIT Press, 1993.
- [73] D. P. Chassin, J. Stoustrup, P. Agathoklis, and N. Djilali, “A new thermostat for real-time price demand response: Cost, comfort and energy impacts of discrete-time control without deadband,” Applied Energy, vol. 155, pp. 816 – 825, 2015.
- [74] K. E. T. Andrew A. Goett, Kathleen Hudson, “Customers’ choice among retail energy suppliers: The willingness-to-pay for service attributes,” The Energy Journal, vol. 21, no. 4, pp. 1–28, 2000.
- [75] N. Ruiz, I. Cobelo, and J. Oyarzabal, “A direct load control model for virtual power plant management,” Power Systems, IEEE Transactions on, vol. 24, pp. 959–966, May 2009.
- [76] E. Vrettos, C. Ziras, and G. Andersson, “Fast and reliable primary frequency reserves from refrigerators with decentralized stochastic control,” arXiv preprint arXiv:1610.00953, 2016.
- [77] C. Perfumo, E. Kofman, J. H. Braslavsky, and J. K. Ward, “Load management: Model-based control of aggregate power for populations of thermostatically con-

- trolled loads,” Energy Conversion and Management, vol. 55, no. 0, pp. 36 – 48, 2012.
- [78] W. C. Chu, B. K. Chen, and C. K. Fu, “Scheduling of direct load control to minimize load reduction for a utility suffering from generation shortage,” Power Systems, IEEE Transactions on, vol. 8, pp. 1525–1530, Nov 1993.
- [79] D. P. Chassin and J. C. Fuller, “On the equilibrium dynamics of demand response in thermostatic loads,” in System Sciences (HICSS), 2011 44th Hawaii International Conference on, pp. 1–6, Jan 2011.
- [80] S. El-Ferik, S. A. Hussain, and F. M. Al-Sunni, “Identification and weather sensitivity of physically based model of residential air-conditioners for direct load control: A case study,” Energy and Buildings, vol. 38, no. 8, pp. 997 – 1005, 2006.
- [81] S. Bashash and H. K. Fathy, “Modeling and control insights into demand-side energy management through setpoint control of thermostatic loads,” in American Control Conference (ACC), 2011, pp. 4546–4553, June 2011.
- [82] M. Roozbehani, M. Dahleh, and S. Mitter, “On the stability of wholesale electricity markets under real-time pricing,” in Decision and Control (CDC), 2010 49th IEEE Conference on, pp. 1911–1918, IEEE, 2010.
- [83] J. Mathieu, S. Koch, and D. Callaway, “State estimation and control of electric loads to manage real-time energy imbalance,” Power Systems, IEEE Transactions on, vol. 28, pp. 430–440, Feb 2013.
- [84] C. Y. Chang, W. Zhang, J. Lian, and K. Kalsi, “Modeling and control of aggregated air conditioning loads under realistic conditions,” in Innovative Smart Grid Technologies (ISGT), 2013 IEEE PES, pp. 1–6, Feb 2013.
- [85] D. P. Chassin, J. Stoustrup, P. Agathoklis, and N. Djilali, “A New Thermostat for Real-Time Price Demand Response: Cost, Comfort and Energy Impacts of Discrete-Time Control Without Deadband,” Applied Energy, to be published.
- [86] D. P. Chassin, “New residential thermostat for transactive systems,” Master’s thesis, University of Victoria, 2015.

- [87] K. Ogata, Discrete-time control systems, vol. 8. Prentice-Hall Englewood Cliffs, NJ, 1995.
- [88] D. P. Chassin, J. C. Fuller, and N. Djilali, “GridLAB-D: An agent-based simulation framework for smart grids,” Journal of Applied Mathematics, vol. 2014, no. 492320, p. 12, 2014.
- [89] K. P. Schneider, J. C. Fuller, and D. P. Chassin, “Analysis of distribution level residential demand response,” in Power Systems Conference and Exposition (PSCE), 2011 IEEE/PES, pp. 1–6, IEEE, 2011.
- [90] P. Kundur, N. J. Balu, and M. G. Lauby, Power system stability and control, vol. 7. McGraw-hill New York, 1994.
- [91] N. R. Subcommittee, “Balancing and frequency control,” technical document, North American Electric Reliability Corporation, 2011.
- [92] J. C. Doyle, K. Glover, P. P. Khargonekar, and B. A. Francis, “State-space solutions to standard \mathcal{H}_2 and \mathcal{H}_∞ control problems,” IEEE Transactions on Automatic control, vol. 34, no. 8, pp. 831–847, 1989.
- [93] T. Chen and B. A. Francis, Optimal sampled-data control systems. Springer Science & Business Media, 2012.
- [94] S. R. Mattix, M. K. Donnelly, D. J. Trudnowski, and J. E. Dagle, “Autonomous demand response for frequency regulation on a large-scale model of an interconnected grid,” in Power and Energy Society General Meeting, 2012 IEEE, pp. 1–8, July 2012.
- [95] S. Behboodi, D. P. Chassin, N. Djilali, and C. Crawford, “Transactive control of fast-acting demand response based on thermostatic loads in real-time retail electricity markets,” Applied Energy, 2017.
- [96] D. P. Chassin, S. Behboodi, Y. Shi, and N. Djilali, “ \mathcal{H}_2 -optimal transactive control of electric power regulation from fast-acting demand response in the presence of high renewables,” Applied Energy, vol. 205, no. Supplement C, pp. 304 – 315, 2017.

- [97] G. A. Stern, J. H. Yan, P. B. Luh, and W. E. Blankson, “What objective function should be used for optimal auctions in the ISO/RTO electricity market?,” in Power Engineering Society General Meeting, pp. 10–pp, IEEE, 2006.
- [98] Y. T. Tan and D. S. Kirschen, “Co-optimization of energy and reserve in electricity markets with demand-side participation in reserve services,” in Power Systems Conference and Exposition, pp. 1182–1189, IEEE, 2006.
- [99] WECC’s System Adequacy Planning Department, “WECC 2024 common case.”
- [100] R. Diao, N. Samaan, Y. Makarov, R. Hafen, and J. Ma, “Planning for variable generation integration through balancing authorities consolidation,” IEEE Power and Energy Society General Meeting, pp. 1–8, 2012.
- [101] M. Roozbehani, M. A. Dahleh, and S. K. Mitter, “Volatility of power grids under real-time pricing,” Power Systems, IEEE Transactions on, vol. 27, no. 4, pp. 1926–1940, 2012.
- [102] M. O. Hill, “Diversity and evenness: a unifying notation and its consequences,” Ecology, vol. 54, no. 2, pp. 427–432, 1973.
- [103] A. Ohio, “gridSMART(SM) Demonstration Project,” 2016.
- [104] S. E. Widergren, M. C. Marinovici, J. C. Fuller, K. Subbarao, D. P. Chassin, and A. Somani, “Customer engagement in AEP gridSMART residential transactive system,” tech. rep., Pacific Northwest National Laboratory (PNNL), Richland, WA (US), 2014.
- [105] D. P. Chassin, Y. Sun, and A. Somani, “Optimization of customer subscription rates to electric utility tariffs,” in System Sciences (HICSS), 2015 48th Hawaii International Conference on, pp. 2604–2609, IEEE, 2015.
- [106] C. Gellings, “Estimating the costs and benefits of the smart grid,” Technical Report 1022519, Electric Power Research Institute, March 2011.
- [107] D. J. Hammerstrom, C. D. Corbin, N. Fernandez, J. S. Homer, A. Makhmalbaf, R. G. Pratt, and A. Somani, “Valuation of transactive systems,” Technical Report PNNL-25323, Pacific Northwest National Laboratory, May 2016.

- [108] S. Behboodi, D. P. Chassin, C. Crawford, and N. Djilali, “Electric vehicle participation in transactive power systems using real-time retail prices,” in 2016 49th Hawaii International Conference on System Sciences (HICSS), pp. 2400–2407, Jan 2016.
- [109] D. Kosterev, A. Meklin, J. Undrill, B. Lesieutre, W. Price, D. Chassin, R. Bravo, and S. Yang, “Load modeling in power system studies: WECC progress update,” in Power and Energy Society General Meeting-Conversion and Delivery of Electrical Energy in the 21st Century, 2008 IEEE, pp. 1–8, IEEE, 2008.
- [110] S. Sharma, S. H. Huang, and N. Sarma, “System inertial frequency response estimation and impact of renewable resources in ERCOT interconnection,” in Power and Energy Society General Meeting, 2011 IEEE, pp. 1–6, July 2011.
- [111] R. Lira, C. Mycock, D. Wilson, and H. Kang, “PMU performance requirements and validation for closed loop applications,” in Innovative Smart Grid Technologies (ISGT Europe), 2011 2nd IEEE PES International Conference and Exhibition on, pp. 1–7, Dec 2011.
- [112] P. Kansal and A. Bose, “Bandwidth and latency requirements for smart transmission grid applications,” IEEE Transactions on Smart Grid, vol. 3, pp. 1344–1352, Sept 2012.
- [113] Y. Chen and M. Hu, “Balancing collective and individual interests in transactive energy management of interconnected micro-grid clusters,” Energy, vol. 109, pp. 1075–1085, 2016.
- [114] W. Hogan, “Time of use rates and real-time prices,” August 23 2014.
- [115] H. Allcott, “Rethinking real-time electricity pricing,” Resource and energy economics, vol. 33, no. 4, pp. 820–842, 2011.
- [116] P. Huang, J. Kalagnanam, R. Natarajan, D. Hammerstrom, R. Melton, M. Sharma, and R. Ambrosio, “Analytics and transactive control design for the pacific northwest smart grid demonstration project,” in 2010 First IEEE International Conference on Smart Grid Communications, pp. 449–454, Oct 2010.
- [117] S. Wolfram, A new kind of science, vol. 5. Wolfram media Champaign, 2002.