

**Assessing Impact of Large-Scale Distributed Residential
HVAC Control Optimization on Electricity Grid Operation
and Renewable Energy Integration**

by

Charles D. Corbin

B.S., Cornell University, 1999

M.S., University of Colorado, 2009

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy

Department of Civil, Environmental and Architectural Engineering

2014

This thesis entitled:
Assessing Impact of Large-Scale Distributed Residential HVAC Control Optimization on
Electricity Grid Operation and Renewable Energy Integration
written by Charles D. Corbin
has been approved for the Department of Civil, Environmental and Architectural Engineering

Prof. Gregor P. Henze, Ph.D.

Michael J. Brandemuehl, Ph.D.

Tim X. Brown, Ph.D.

Clemens Felsmann, Ph.D.

Moncef Krarti, Ph.D.

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Corbin, Charles D. (Ph.D., Architectural Engineering)

Assessing Impact of Large-Scale Distributed Residential HVAC Control Optimization on Electricity
Grid Operation and Renewable Energy Integration

Thesis directed by Prof. Gregor P. Henze, Ph.D.

Demand management is an important component of the emerging Smart Grid, and a potential solution to the supply-demand imbalance occurring increasingly as intermittent renewable electricity is added to the generation mix. Model predictive control (MPC) has shown great promise for controlling HVAC demand in commercial buildings, making it an ideal solution to this problem. MPC is believed to hold similar promise for residential applications, yet very few examples exist in the literature despite a growing interest in residential demand management. This work explores the potential for residential buildings to shape electric demand at the distribution feeder level in order to reduce peak demand, reduce system ramping, and increase load factor using detailed sub-hourly simulations of thousands of buildings coupled to distribution power flow software. More generally, this work develops a methodology for the optimization of residential HVAC operation using a distributed but directed MPC scheme that can be applied to today's programmable thermostat technologies to address the increasing variability in electric supply and demand. Case studies incorporating varying levels of renewable energy generation demonstrate the approach and highlight important considerations for large-scale residential model predictive control.

Dedication

To my late father, Charles Kay Corbin Jr. (April 2, 1933 – February 20, 2014), who provided me every possible advantage, despite disadvantages of his own.

Acknowledgements

This work is the result of contributions from a great number of individuals who inspired, guided, critiqued, supported and challenged me over the course of many years. Without their brilliance, this work would not be half the work it is. Sincere gratitude goes to: Dr. Gregor Henze for many insightful conversations, for sharing an encyclopedic knowledge of control theory, and for patience as I found my way; committee members Dr. Michael Brandemuehl, Dr. Moncef Krarti and Dr. Tim Brown for their invaluable perspectives on all aspects of simulation and modeling; Vince Cushing for his thought leadership and vision, and for providing an avenue for early commercial building MPC investigations through Clean Urban Energy; Erik Greensfelder for providing early inspiration regarding the large-scale impacts of commercial MPC; Dr. Clemens Felsmann for providing an opportunity to apply MPC techniques to TABS buildings through the LowEx:Monitor project; my BSP colleagues Dr. Peter May-Ostendorp, Rois Langner and Simon Olivieri for advancing the state-of-the-art in their research and for providing an opportunity for me to contribute; the Fleming crew (Neal Kruis, Ryan Tanner, Anthony Florita) for being a faithful sounding board against which I could bounce my many ideas and frustrations; Jim Hauswirth and Dylan Cutler for helping me clear my head on our weekly runs and for their valuable outside perspectives on my research; my mother and father for giving me the foundations that have allowed me to succeed in this and other endeavors; and last but certainly not least, my partner Shan Wo for her unwavering support over the years as I struggled through the not-always-easy-or-fun process we call a doctoral degree.

Contents

Chapter

1	Introduction	1
1.0.1	Significance of Research	3
1.0.2	Research Questions	7
1.1	Research Objectives	8
1.2	Survey of Literature	8
1.2.1	Optimal Control of Building Systems	9
1.2.2	Reduced Order Modeling	10
1.2.3	Residential Demand Side Management and Control Optimization	13
1.3	Organization of Document	17
2	Methodology	19
2.1	Research Methodology	19
2.2	Metrics for Evaluating Optimization Performance	21
2.2.1	Total Electricity Consumption	21
2.2.2	Peak Demand	22
2.2.3	Peak to Valley Ratio	22
2.2.4	Load Factor	22
2.2.5	System Ramping	22
2.2.6	Feeder Demand Spectral Density	23

2.3	Simulation Environment	23
2.3.1	Background	23
2.3.2	Simulation Framework	24
2.3.3	Electric Grid Simulation Software	26
2.4	Simulation Hardware	27
2.5	Experimental Plan	28
2.5.1	Demand Response	28
2.5.2	Demand Limiting Optimization	29
2.5.3	Dynamic Price Optimization	29
2.5.4	Load Shape Optimization	29
2.5.5	Rooftop Solar	29
2.5.6	Utility Scale Wind	30
3	Model Selection and Development	31
3.1	Building Model	31
3.1.1	Envelope Model	31
3.1.2	Solar Gains	35
3.1.3	Glazing Model	36
3.1.4	Shading Model	36
3.1.5	HVAC Models	37
3.1.6	Schedules	39
3.1.7	Internal Gains	41
3.1.8	Thermostat Model	42
3.1.9	Simulation Flow	45
3.1.10	Annual Validation of Building Model	46
3.1.11	Sub-Hourly Validation of Building Model	50
3.2	Electric Grid Model	55

3.2.1	Selection of Climate and Feeder Models	55
3.2.2	Populating Feeder Model with Residences	58
3.2.3	Modeling Home as ZIP Load	62
3.2.4	Selection of ZIP Fractions	64
3.2.5	Feeder Sensitivity to ZIP Fractions	66
3.2.6	Development of Hybrid Model	68
3.2.7	Validation of Hybrid Model	69
4	Model Predictive Control and Optimization	74
4.1	Methodology	74
4.2	Optimization Process	75
4.3	Optimization Algorithm	77
4.4	Refinement of Search Space Mapping	78
5	Demand Response	81
5.1	Methodology	82
5.2	Selection of Demand Day	83
5.2.1	Houston	83
5.2.2	Los Angeles	84
5.2.3	New York	84
5.3	Participant Sampling and Number of Simulations	86
5.4	Interpretation of Metrics Tables	88
5.5	Houston	88
5.6	Los Angeles	92
5.7	New York	95
6	Demand Limiting Optimization	99
6.1	Methodology	99

6.1.1	Objective Function Definition	101
6.2	Houston	102
6.3	Los Angeles	108
6.4	New York	112
6.5	Precooling Investment	116
7	Dynamic Price Optimization	124
7.1	Methodology	125
7.1.1	Day-Ahead Price Modeling	125
7.1.2	Objective Function Definition	130
7.2	Methodology Refinement	131
7.3	Day-Ahead Price Simulations	138
7.3.1	Simulation Results	138
7.3.2	Summary and Conclusions	143
7.4	Synthetic Price Simulations	143
7.4.1	Simulation Results	144
7.4.2	Summary and Conclusions	151
7.5	Demand Oscillations	151
8	Load Shape Optimization	155
8.1	Methodology	155
8.1.1	Generation of Feeder Reference Demand Curve	156
8.1.2	Generation of House Reference Demand Curve	158
8.1.3	Objective Function Definition	159
8.2	Houston	160
8.3	Los Angeles	167
8.4	New York	172
8.5	Summary and Conclusions	176

9	Rooftop Solar Electric Generation	177
9.1	Methodology	179
9.2	Model Development and Validation	179
9.3	Desired PV Penetration	183
9.4	Sizing and Distribution of PV Systems	184
9.5	Houston	185
9.6	Los Angeles	195
9.7	New York	201
9.8	Summary and Conclusions	206
10	Utility Scale Wind Generation	207
10.1	Methodology	207
10.2	Model Selection and Description	208
10.3	Desired Wind Penetration	210
10.4	Sizing Wind Systems	210
10.5	Houston	211
10.6	Los Angeles	219
10.7	New York	225
10.8	Summary and Conclusions	231
11	Conclusions and Future Work	232
11.1	Summary of Results & Conclusions	232
11.2	Contributions & Future Work	235

Bibliography	240
---------------------	-----

Appendix

A Supplement to Demand Response	248
B Supplement to Demand Limiting Optimization	252
B.1 Houston	252
B.1.1 Demand Limiting as a Function of Mass Level	261
B.2 Los Angeles	263
B.3 New York	267
C Supplement to Dynamic Price Optimization	272
C.1 Houston – Day-Ahead Price	272
C.2 Los Angeles – Day-Ahead Price	276
C.3 Los Angeles – Synthetic Price	286
C.4 New York – Day-Ahead Price	292
C.5 New York – Synthetic Price	302
D Supplement to Load Shape Optimization	308
D.1 Houston	308
D.2 Los Angeles	310
D.3 New York	314
E Supplement to Rooftop Solar Electric Generation	319
E.1 Houston	319
E.2 Los Angeles	322
E.3 New York	328

F	Supplement to Utility Scale Wind Generation	336
F.1	Houston	336
F.2	Los Angeles	340
F.3	New York	345
G	Assessment of Building Thermal Storage Efficiency	350
H	Assessment of Zero Demand Optimization	353

Tables

Table

3.1	NRMSE between model and EnergyPlus, two mass levels, cooling step change.	54
3.2	Climate characteristics of cities selected for study.	55
3.3	Characteristics of Feeder 8.	57
3.4	Characteristics of Feeder 13.	57
3.5	Characteristics of Feeder 22.	58
3.6	ZIP fractions for three feeder models.	66
3.7	ZIP fractions and case names for sensitivity test.	66
3.8	Percentage error in total electricity consumption for three feeders.	70
5.1	Performance metrics, demand response, Houston.	92
5.2	Performance metrics, demand response, Los Angeles.	95
5.3	Performance metrics, demand response, New York.	95
6.1	Cooling set point boundaries for demand limiting optimizations.	100
6.2	Performance metrics, demand limiting optimization, Houston, 70%.	104
6.3	Performance metrics, demand limiting optimization, Los Angeles, 70%.	108
6.4	Performance metrics, demand limiting optimization, New York, 70%.	112
7.1	Model parameters for price prediction CART.	126
7.2	Performance metrics, day-ahead price optimization, Houston, 70%.	139
7.3	Performance metrics, day-ahead price optimization, Houston, zero-degree, 70%.	139

7.4	Performance metrics, day-ahead price optimization, Houston, ramp-return, 70%. . . .	140
7.5	Performance metrics, synthetic price optimization, Houston, zero-degree, 70%.	145
7.6	Performance metrics, synthetic price optimization, Houston, ramp-return, 70%.	145
7.7	Performance metrics, synthetic price optimization, Houston, zero-degree, 30%.	145
8.1	Performance metrics, load shaping optimization, Houston, 70%.	161
8.2	Performance metrics, load shaping optimization, Houston, 30%.	161
8.3	Performance metrics, load shaping optimization, Los Angeles, 70%.	167
8.4	Performance metrics, load shaping optimization, New York, 70%.	172
9.1	Default photovoltaic system parameters.	181
9.2	Photovoltaic system parameters for model validation.	181
9.3	PV Penetration for feeders.	184
9.4	Performance metrics, load shaping optimization, Houston, high solar, 70%.	186
9.5	Performance metrics, load shaping optimization, Houston, low solar, 70%.	186
9.6	Performance metrics, load shaping optimization, Houston, high solar, 30%.	191
9.7	Performance metrics, load shaping optimization, Houston, low solar, 30%.	191
9.8	Performance metrics, load shaping optimization, Los Angeles, high solar, 70%.	195
9.9	Performance metrics, load shaping optimization, Los Angeles, low solar, 70%.	198
9.10	Performance metrics, load shaping optimization, New York, high solar, 70%.	204
9.11	Performance metrics, load shaping optimization, New York, low solar, 70%.	204
10.1	Wind turbine scaling factors.	211
10.2	Performance metrics, load shaping optimization, Houston, high wind, 70%.	212
10.3	Performance metrics, load shaping optimization, Houston, low wind, 70%.	212
10.4	Performance metrics, load shaping optimization, Los Angeles, high wind, 70%.	222
10.5	Performance metrics, load shaping optimization, Los Angeles, low wind, 70%.	222
10.6	Performance metrics, load shaping optimization, New York, high wind, 70%.	225

10.7	Performance metrics, load shaping optimization, New York, low wind, 70%.	228
B.1	Performance metrics, demand limiting optimization, Houston, 30%.	252
B.2	Performance metrics, demand limiting optimization, Los Angeles, 30%.	263
B.3	Performance metrics, demand limiting optimization, New York, 30%.	267
C.1	Performance metrics, day-ahead price optimization, Houston, ramp-return, 30%. . . .	272
C.2	Performance metrics, day-ahead price optimization, Los Angeles, 70%.	276
C.3	Performance metrics, day-ahead price optimization, Los Angeles, zero-degree, 70%. . .	276
C.4	Performance metrics, day-ahead price optimization, Los Angeles, ramp-return, 70%. .	277
C.5	Performance metrics, day-ahead price optimization, Los Angeles, ramp-return, 30%. .	277
C.6	Performance metrics, synthetic price optimization, Los Angeles, zero-degree, 70%. . . .	286
C.7	Performance metrics, synthetic price optimization, Los Angeles, ramp-return, 70%. . .	286
C.8	Performance metrics, synthetic price optimization, Los Angeles, zero-degree, 30%. . . .	287
C.9	Performance metrics, day-ahead price optimization, New York, zero-degree, 70%. . . .	292
C.10	Performance metrics, day-ahead price optimization, New York, 70%.	292
C.11	Performance metrics, day-ahead price optimization, New York, ramp-return, 70%. . . .	293
C.12	Performance metrics, day-ahead price optimization, New York, ramp-return, 30%. . . .	293
C.13	Performance metrics, synthetic price optimization, New York, zero-degree, 70%.	302
C.14	Performance metrics, synthetic price optimization, New York, ramp-return, 70%.	302
C.15	Performance metrics, synthetic price optimization, New York, zero-degree, 30%.	303
D.1	Performance metrics, load shaping optimization, Los Angeles, 30%.	310
D.2	Performance metrics, load shaping optimization, New York, 30%.	314
E.1	Performance metrics, load shaping optimization, Los Angeles, high solar, 30%.	322
E.2	Performance metrics, load shaping optimization, Los Angeles, low solar, 30%.	322
E.3	Performance metrics, load shaping optimization, New York, high solar, 30%.	328
E.4	Performance metrics, load shaping optimization, New York, low solar, 30%.	328

F.1	Performance metrics, load shaping optimization, Houston, high wind, 30%.	336
F.2	Performance metrics, load shaping optimization, Houston, low wind, 30%.	336
F.3	Performance metrics, load shaping optimization, Los Angeles, high wind, 30%.	340
F.4	Performance metrics, load shaping optimization, Los Angeles, low wind, 30%.	340
F.5	Performance metrics, load shaping optimization, New York, high wind, 30%.	345
F.6	Performance metrics, load shaping optimization, New York, low wind, 30%.	345
H.1	Performance metrics, zero demand optimization, Houston, 70%.	356
H.2	Performance metrics, zero demand optimization, Houston, high wind, 70%.	356

Figures

Figure

1.1	Growth in residential electricity retail sales.	5
1.2	Projected electricity consumption.	6
1.3	Projected carbon dioxide emissions.	6
3.1	Building envelope model expressed as a thermal network.	32
3.2	Example of HVAC cycling, low cooling set point.	43
3.3	Example of HVAC cycling, high cooling set point.	44
3.4	Comparison of annual electricity use in cooling physics test cases.	48
3.5	Comparison of annual gas use in heating physics test cases.	49
3.6	Comparison of monthly electricity use in cooling and heating physics test cases.	50
3.7	Sub-hourly heating validation of reduced order model.	52
3.8	Sub-hourly cooling validation of reduced order model.	53
3.9	Sub-hourly validation of model, cooling step change, two mass levels.	54
3.10	Distribution of home characteristics for Houston feeder model.	59
3.11	Distribution of home characteristics for Los Angeles feeder model.	60
3.12	Distribution of home characteristics for New York feeder model.	61
3.13	Sensitivity of Feeder Demand to ZIP Fractions.	67
3.14	Monthly validation of hybrid GridMPC model.	70
3.15	Feeder demand validation for Houston feeder model.	71

3.16	Feeder demand validation for Los Angeles feeder model.	72
3.17	Feeder demand validation for New York feeder model.	73
4.1	Illustration of search space mapping.	80
5.1	Houston feeder peak demand day, July 29.	83
5.2	Los Angeles feeder peak demand day, September 24.	84
5.3	New York feeder peak demand day, July 16.	85
5.4	Convergence of feeder demand profile NRMSE.	87
5.5	Houston feeder demand curves for 2hr, 30% DR event.	90
5.6	Houston feeder demand curves for 2hr, 70% DR event.	90
5.7	Houston feeder demand curves for 6hr, 30% DR event.	91
5.8	Houston feeder demand curves for 6hr, 70% DR event.	91
5.9	Los Angeles feeder demand curves for 2hr, 30% DR event.	93
5.10	Los Angeles feeder demand curves for 2hr, 70% DR event.	93
5.11	Los Angeles feeder demand curves for 6hr, 30% DR event.	94
5.12	Los Angeles feeder demand curves for 6hr, 70% DR event.	94
5.13	New York feeder demand curves for 2hr, 30% DR event.	96
5.14	New York feeder demand curves for 2hr, 70% DR event.	96
5.15	New York feeder demand curves for 6hr, 30% DR event.	97
5.16	New York feeder demand curves for 6hr, 70% DR event.	97
6.1	Demand limiting optimization, Houston, 70%.	103
6.2	Total spectral power, demand limiting optimization, Houston, 70%.	105
6.3	Demand limiting optimization, Houston, July 2, 70%.	105
6.4	Demand limiting optimization, Houston, July 21, 70%.	106
6.5	Demand limiting optimization, Los Angeles, 70%.	109
6.6	Total spectral power, demand limiting optimization, Los Angeles, 70%.	110

6.7	Demand limiting optimization, Los Angeles, July 10, 70%.	110
6.8	Demand limiting optimization, Los Angeles, July 29, 70%.	111
6.9	Demand limiting optimization, New York, 70%.	113
6.10	Total spectral power, demand limiting optimization, New York, 70%.	114
6.11	Demand limiting optimization, New York, July 16, 70%.	114
6.12	Demand limiting optimization, New York, July 17, 70%.	115
6.13	Precooling investment, Houston, noon, 70%.	118
6.14	Precooling investment, Los Angeles, noon, 70%.	120
6.15	Precooling investment, Los Angeles, July 10, 70%.	121
6.16	Precooling investment, New York, noon, 70%.	123
7.1	Historical and modeled CAISO prices.	127
7.2	Historical and modeled ERCOT prices.	127
7.3	Historical and modeled NYISO prices.	128
7.4	Predicted CAISO prices.	128
7.5	Predicted ERCOT prices.	129
7.6	Predicted NYISO prices.	129
7.7	Day-ahead price optimization, Houston, first test, 70%.	132
7.8	Day-ahead price optimization, Houston, second test, 70%.	133
7.9	Day-ahead price optimization, Houston, 70%.	135
7.10	Day-ahead price optimization, Houston, zero-degree, 70%.	136
7.11	Day-ahead price optimization, Houston, ramp-return, 70%.	137
7.12	Total spectral power, day-ahead price optimization, Houston, zero-degree, 70%.	140
7.13	Total spectral power, day-ahead price optimization, Houston, 70%.	141
7.14	Total spectral power, day-ahead price optimization, Houston, ramp-return, 70%.	141
7.15	Day-ahead price optimization, Houston, July 1, zero-degree, 70%.	142
7.16	Day-ahead price optimization, Houston, July 26, zero-degree, 70%.	142

7.17	Total spectral power, synthetic price optimization, Houston, zero-degree, 70%.	146
7.18	Total spectral power, synthetic price optimization, Houston, ramp-return, 70%.	146
7.19	Total spectral power, synthetic price optimization, Houston, zero-degree, 30%.	147
7.20	Synthetic price optimization, Houston, zero-degree, 70%.	148
7.21	Synthetic price optimization, Houston, ramp-return, 70%.	149
7.22	Synthetic price optimization, Houston, zero-degree, 30%.	150
7.23	Oscillations in demand profile.	153
8.1	Example of feeder reference demand and signal.	159
8.2	Total spectral power, load shape optimization, Houston, 70%.	162
8.3	Total spectral power, load shape optimization, Houston, 30%.	162
8.4	Load shape optimization, Houston, 70%.	163
8.5	Load shape optimization, Houston, 30%.	164
8.6	Load shape optimization, Houston, July 2, 70%.	165
8.7	Load shape optimization, Houston, July 7, 70%.	166
8.8	Total spectral power, load shape optimization, Los Angeles, 70%.	168
8.9	Load shape optimization, Los Angeles, 70%.	169
8.10	Load shape optimization, Los Angeles, July 9, 70%.	170
8.11	Load shape optimization, Los Angeles, July 11, 70%.	170
8.12	Load shape optimization, Los Angeles, July 29, 70%.	171
8.13	Total spectral power, load shape optimization, New York, 70%.	173
8.14	Load shape optimization, New York, 70%.	174
8.15	Load shape optimization, New York, July 4, 70%.	175
8.16	Load shape optimization, New York, July 24, 70%.	175
9.1	Example of the Duck Curve produced in systems with high PV penetration.	178
9.2	PV model output, July 14–21.	182
9.3	Load shape optimization, Houston, high solar, 70%.	187

9.4	Load shape optimization, Houston, low solar, 70%.	188
9.5	Load shape optimization, Houston, high solar, 30%.	189
9.6	Load shape optimization, Houston, low solar, 30%.	190
9.7	Total spectral power, load shape optimization, Houston, high solar, 70%.	191
9.8	Total spectral power, load shape optimization, Houston, low solar, 70%.	192
9.9	Total spectral power, load shape optimization, Houston, high solar, 30%.	192
9.10	Total spectral power, load shape optimization, Houston, low solar, 30%.	193
9.11	Load shape optimization, Houston, July 1, high solar, 70%.	193
9.12	Load shape optimization, Houston, July 21, high solar, 70%.	194
9.13	Load shape optimization, Los Angeles, high solar, 70%.	196
9.14	Load shape optimization, Los Angeles, low solar, 70%.	197
9.15	Total spectral power, load shape optimization, Los Angeles, high solar, 70%.	199
9.16	Total spectral power, load shape optimization, Los Angeles, low solar, 70%.	199
9.17	Load shape optimization, Los Angeles, July 13, high solar, 70%.	200
9.18	Load shape optimization, Los Angeles, July 15, high solar, 70%.	200
9.19	Load shape optimization, New York, high solar, 70%.	202
9.20	Load shape optimization, New York, low solar, 70%.	203
9.21	Total spectral power, load shape optimization, New York, high solar, 70%.	204
9.22	Total spectral power, load shape optimization, New York, low solar, 70%.	205
10.1	Normalized wind turbine output.	209
10.2	Pre-composited optimized load shape, Houston, high wind, 70%.	213
10.3	Load shape optimization, Houston, high wind, 70%.	214
10.4	Load shape optimization, Houston, low wind, 70%.	215
10.5	Total spectral power, load shape optimization, Houston, high wind, 70%.	216
10.6	Total spectral power, load shape optimization, Houston, low wind, 70%.	216
10.7	Load shape optimization, Houston, July 6, high wind, 70%.	217

10.8	Load shape optimization, Houston, July 14, high wind, 70%.	217
10.9	Load shape optimization, Houston, July 20, high wind, 70%.	218
10.10	Load shape optimization, Los Angeles, high wind, 70%.	220
10.11	Load shape optimization, Los Angeles, low wind, 70%.	221
10.12	Total spectral power, load shape optimization, Los Angeles, high wind, 70%.	222
10.13	Total spectral power, load shape optimization, Los Angeles, low wind, 70%.	223
10.14	Load shape optimization, Los Angeles, July 9, high wind, 70%.	223
10.15	Load shape optimization, Los Angeles, July 13, high wind, 70%.	224
10.16	Load shape optimization, New York, high wind, 70%.	226
10.17	Load shape optimization, New York, low wind, 70%.	227
10.18	Total spectral power, load shape optimization, New York, high wind, 70%.	229
10.19	Total spectral power, load shape optimization, New York, low wind, 70%.	229
10.20	Load shape optimization, New York, July 3, high wind, 70%.	230
10.21	Load shape optimization, New York, July 10, high wind, 70%.	230
A.1	Annual Houston feeder demand.	249
A.2	Annual Los Angeles feeder demand.	250
A.3	Annual New York feeder demand.	251
B.1	Total spectral power, demand limiting optimization, Houston, 30%.	253
B.2	Demand limiting optimization, Houston, 15 min., 70%.	254
B.3	Demand limiting optimization, Houston, 30 min., 70%.	255
B.4	Demand limiting optimization, Houston, 90 min., 70%.	256
B.5	Demand limiting optimization, Houston, 120 min., 70%.	257
B.6	Demand limiting optimization, Houston, 30%.	258
B.7	Demand limiting optimization, Houston, July 29, 70%.	259
B.8	Precooling investment, Houston, July 29, 70%.	260
B.9	Demand limiting optimization by mass level, Houston, July 29.	262

B.10	Total spectral power, demand limiting optimization, Los Angeles, 30%.	264
B.11	Demand limiting optimization, Los Angeles, 30%.	265
B.12	Demand limiting optimization, Los Angeles, July 24, 70%.	266
B.13	Total spectral power, demand limiting optimization, New York, 30%.	268
B.14	Demand limiting optimization, New York, 30%.	269
B.15	Demand limiting optimization, New York, July 23, 70%.	270
B.16	Precooling investment, New York, July 16, 70%.	271
C.1	Total spectral power, day-ahead price optimization, Houston, ramp-return, 30%.	273
C.2	Day-ahead price optimization, Houston, ramp-return, 30%.	274
C.3	Day-ahead price optimization, Houston, July 29, zero-degree, 70%.	275
C.4	Total spectral power, day-ahead price optimization, Los Angeles, 70%.	277
C.5	Total spectral power, day-ahead price optimization, Los Angeles, zero-degree, 70%.	278
C.6	Total spectral power, day-ahead price optimization, Los Angeles, ramp-return, 70%.	278
C.7	Total spectral power, day-ahead price optimization, Los Angeles, ramp-return, 30%.	279
C.8	Day-ahead price optimization, Los Angeles, 70%.	280
C.9	Day-ahead price optimization, Los Angeles, ramp-return, 70%.	281
C.10	Day-ahead price optimization, Los Angeles, zero-degree, 70%.	282
C.11	Day-ahead price optimization, Los Angeles, ramp-return, 30%.	283
C.12	Day-ahead price optimization, Los Angeles, July 7, 70%.	284
C.13	Day-ahead price optimization, Los Angeles, July 24, 70%.	284
C.14	Day-ahead price optimization, Los Angeles, July 28, 70%.	285
C.15	Total spectral power, synthetic price optimization, Los Angeles, zero-degree, 70%.	287
C.16	Total spectral power, synthetic price optimization, Los Angeles, ramp-return, 70%.	288
C.17	Total spectral power, synthetic price optimization, Los Angeles, zero-degree, 30%.	288
C.18	Synthetic price optimization, Los Angeles, zero-degree, 70%.	289
C.19	Synthetic price optimization, Los Angeles, ramp-return, 70%.	290

C.20	Synthetic price optimization, Los Angeles, zero-degree, 30%.	291
C.21	Total spectral power, day-ahead price optimization, New York, zero-degree, 70%. . . .	293
C.22	Total spectral power, day-ahead price optimization, New York, 70%.	294
C.23	Total spectral power, day-ahead price optimization, New York, ramp-return, 70%. . .	294
C.24	Total spectral power, day-ahead price optimization, New York, ramp-return, 30%. . .	295
C.25	Day-ahead price optimization, New York, zero-degree, 70%.	296
C.26	Day-ahead price optimization, New York, ramp-return, 70%.	297
C.27	Day-ahead price optimization, New York, 70%.	298
C.28	Day-ahead price optimization, New York, ramp-return, 30%.	299
C.29	Day-ahead price optimization, New York, July 14, 70%.	300
C.30	Day-ahead price optimization, New York, July 17, 70%.	300
C.31	Day-ahead price optimization, New York, July 19, 70%.	301
C.32	Total spectral power, synthetic price optimization, New York, zero-degree, 70%. . . .	303
C.33	Total spectral power, synthetic price optimization, New York, ramp-return, 70%. . . .	304
C.34	Total spectral power, synthetic price optimization, New York, zero-degree, 30%. . . .	304
C.35	Synthetic price optimization, New York, zero-degree, 70%.	305
C.36	Synthetic price optimization, New York, zero-degree, 30%.	306
C.37	Synthetic price optimization, New York, ramp-return, 70%.	307
D.1	Load shape optimization, Houston, July 21, 70%.	308
D.2	Load shape optimization, Houston, 120 min., 70%.	309
D.3	Total spectral power, load shape optimization, Los Angeles, 30%.	311
D.4	Load shape optimization, Los Angeles, 120 min., 70%.	312
D.5	Load shape optimization, Los Angeles, 30%.	313
D.6	Total spectral power, load shape optimization, New York, 30%.	315
D.7	Load shape optimization, New York, 120 min., 70%.	316
D.8	Load shape optimization, New York, 30%.	317

D.9	Load shape optimization, New York, July 15, 70%.	318
E.1	Houston feeder demand, July 1–7, high solar.	319
E.2	Annual Houston feeder demand, high solar.	320
E.3	Load shape optimization, Houston, July 6, high solar, 70%.	321
E.4	Total spectral power, load shape optimization, Los Angeles, high solar, 30%.	323
E.5	Total spectral power, load shape optimization, Los Angeles, low solar, 30%.	323
E.6	Annual Los Angeles feeder demand, high solar.	324
E.7	Load shape optimization, Los Angeles, high solar, 30%.	325
E.8	Load shape optimization, Los Angeles, low solar, 30%.	326
E.9	Load shape optimization, Los Angeles, July 3, high solar, 70%.	327
E.10	Total spectral power, load shape optimization, New York, high solar, 30%.	329
E.11	Total spectral power, load shape optimization, New York, low solar, 30%.	329
E.12	Annual New York feeder demand, high solar.	330
E.13	New York feeder demand, July 1–7, high solar.	331
E.14	Load shape optimization, New York, low solar, 30%.	332
E.15	Load shape optimization, New York, high solar, 30%.	333
E.16	Load shape optimization, New York, July 2, high solar, 70%.	334
E.17	Load shape optimization, New York, July 13, high solar, 70%.	334
E.18	Load shape optimization, New York, July 20, high solar, 70%.	335
F.1	Total spectral power, load shape optimization, Houston, high wind, 30%.	337
F.2	Total spectral power, load shape optimization, Houston, low wind, 30%.	337
F.3	Load shape optimization, Houston, high wind, 30%.	338
F.4	Load shape optimization, Houston, low wind, 30%.	339
F.5	Total spectral power, load shape optimization, Los Angeles, high wind, 30%.	341
F.6	Total spectral power, load shape optimization, Los Angeles, low wind, 30%.	341
F.7	Load shape optimization, Los Angeles, high wind, 30%.	342

F.8	Load shape optimization, Los Angeles, low wind, 30%.	343
F.9	Load shape optimization, Los Angeles, July 23, high wind, 70%.	344
F.10	Total spectral power, load shape optimization, New York, high wind, 30%.	346
F.11	Total spectral power, load shape optimization, New York, low wind, 30%.	346
F.12	Load shape optimization, New York, low wind, 30%.	347
F.13	Load shape optimization, New York, high wind, 30%.	348
F.14	Load shape optimization, New York, July 24, high wind, 70%.	349
G.1	Histogram of envelope thermal storage efficiency, Los Angeles.	352
H.1	Total spectral power, zero demand optimization, Houston, 70%.	357
H.2	Total spectral power, zero demand optimization, Houston, high wind, 70%.	357
H.3	Zero demand optimization, Houston, 70%.	358
H.4	Zero demand optimization, Houston, high wind, 70%.	359

Chapter 1

Introduction

Despite continuous upgrades and improvements to the systems that form the nation's electricity grids, the fundamental operation of these systems has changed relatively little over almost one hundred years. This is to say that the systems which supply electricity to the residential, commercial and industrial sectors do so according to demand. As the absolute magnitude and diversity of this demand has increased, so too has the size and diversity of generation required to meet our electricity needs. Traditionally, this has meant that generation capacity was built to satisfy peak demand, using generation that was less efficient and expensive to operate, leaving idle or underutilized, much of the generation capacity until needed. In practical terms, this has created a large disparity in electricity price from hour to hour, as systems are daily brought online to satisfy demand, and taken offline when demand has subsided. It is a well known fact that the disparity between off-peak and on-peak spot prices may exceed several orders of magnitude, with prices dropping below zero in some extreme cases.

With the introduction of large-scale renewable energy generation, a new conflict between renewable sources and traditional generation operated according to demand has been created. These new generation sources compete with the constraints of traditional generation which cannot be easily or quickly throttled without impacts to long-term reliability. This leads to curtailment of wind generated electricity, preventing the cleaner, renewable energy source from replacing dirtier traditional sources. Because of this, and the inherent intermittency and variability of renewable sources, the amount of renewable energy present on the electric system is effectively capped unless

strategies can be employed to utilize renewable generation when it is available.

Obviously, one solution to both problems, i.e. the daily variation in demand combined with the intermittency of renewable generation, is grid-scale storage. While there is significant interest and promise in this solution, the cost of deploying large storage facilities is prohibitive and at least decades away from being practical. Recent studies [32, 46, 51] on electric and plug-in hybrid vehicles suggest that a distributed approach to grid-scale storage may provide a solution to the problem. However, this solution too is many years in the making, and poses a different problem of increased electricity demand [92, 88, 53] at peak periods unless managed properly.¹

Assuming that neither grid-scale, nor distributed storage are deployed at scale for a decade or more, and that renewable generation will continue to become an increasing component in the generation mix, the problem becomes one of how best to enable demand side management of electricity consumers given an existing mix of traditional and renewable generation.

Distilling the problem statement above yields a three-part hypothesis which informs the direction of this research:

1. The use of model predictive control (MPC) for residential heating, ventilation and air conditioning (HVAC) control may eliminate traditional, reactive, demand response by proactively and predictably shaping the aggregate demand of residential building “collectives” in anticipation of supply side constraints, allowing optimal planning of resources.
2. In systems with relatively high renewable energy penetration, or during catastrophic events, generation capacity may suddenly disappear. Short-term, i.e. day-ahead, predictive control of residential loads allows demand flexibility and graceful response under grid stress events.
3. Renewable energy is often available when it is not needed, resulting in low utilization. The widespread adoption of predictive controls applied to residential buildings’ thermal mass

¹ One may note that the methods employed and researched in this work could be easily extended to the control of electric vehicle (EV) charging. As of now, application of the techniques to EVs is not within the scope of the work.

offers the promise of deeper penetration of renewable energy without an investment in electric storage systems and without destabilizing electric grid operations.

1.0.1 Significance of Research

Model Predictive Control (MPC) has enjoyed great popularity in the research of demand limiting techniques for commercial buildings. The literature is full of examples in which this approach has been successfully applied and validated [12, 55, 10, 68, 56, 57, 44, 40, 42, 41, 43], giving rise to companies such as BuildingIQ [18] and QCoefficient, Inc. [26] seeking to capitalize on MPC as a service for building owners and managers. Research on the application of MPC to residential buildings is relatively sparse in the literature; a good example of which is reported by Hubert et al. [47], who claim building-level optimization algorithms are critical to the efficient and robust operation of the Smart Grid. Realizing the potential for application in the residential sector, companies such as EcoFactor [34] and Nest Labs [69] who have taken very different approaches to implementing MPC, but have not attempted to tackle the problem of optimizing building collectives in aggregate.

By and large, the residential sector has only realized “predictive control” through residential demand response (DR) programs implemented by utilities, whereby the utilities predict the demand reduction available based on the aggregated response to historical event opt-ins, and issue demand events that effectively disable HVAC systems to achieve a demand reduction. Obviously this is a crude approach to predictive control that does not necessarily consider the impacts on a residence’s occupants. Still, the application of residential DR deployment has only been realized at pilot scales at best as the efficacy of this approach is unproven, perhaps because of the uncertainty of occupant behavior and frequent opt-outs. The author believes that traditional DR is an outdated approach to shaping residential load which is better suited to MPC. Traditional DR is a one-way reactive technology called upon when grid demand exceeds capacity. When MPC of residential HVAC is implemented, the traditional concept of DR will likely become obsolete: rather than issuing demand events, MPC will allow continuous optimization of residential load that can be forecast

with potentially greater accuracy and control, creating a two-way communication between utility and building, i.e. supply and demand.

The importance of improved residential energy controls, resulting in electricity consumption and demand reduction cannot be understated. According to the United States Energy Information Administration (EIA) [1], residential electricity consumption has been on a consistently increasing trajectory for more than sixty years at a rate of 20 TWh per year (Figure 1.1). Not surprisingly, this corresponds to an increase in residential air conditioning, which is present in 87% of homes as of 2009, compared with 57% in 1980. It is estimated that residential air conditioning exceeds 293 TWh annually. While projections suggest that on average, the energy use intensity of residential buildings is on the decrease due to envelope and equipment energy efficiency improvements, the total electricity used by the residential sector will increase over the next 25 years [2] under a “business as usual” reference case scenario (Figure 1.2). Under this scenario, residential electricity use by air conditioning will increase by 363 TWh annually compared to 2009 levels. The EIA projects that with best available demand reduction technology, space cooling savings on the order of 4,400 TWh over the next 25 years are possible. Indeed, these savings include effects from a number of energy reduction technologies including high-efficiency equipment and better insulated building envelopes, but some fraction will inevitably come from improved operation; after all, the most efficient air conditioner is the one that does not run at all.

The EIA projects that over the next 25 years, the addition of roughly 48 GW in generating capacity will be required to satisfy summer peak demand. While the exact number attributed to residential electricity demand is harder to determine, it stands to reason that a non-trivial percentage of this demand will come from residential cooling end-uses, considering the temporal coincidence of residential cooling need with peak grid demand. Considered individually, residential cooling demand reductions during peak demand events may only measure in the kilowatts, but considered in aggregate, a large number of homes represent megawatts and potentially gigawatts of demand that can be shifted or eliminated. Enabling residential consumers to participate in the shaping of electric demand allows greater flexibility of grid operation that may reduce the need for

additional peak generation.

Finally, the EIA projects that while gas displaces coal in the short-term and renewable energy penetration increases under the reference case scenario, the overall increase in residential electricity consumption results in a (albeit small) net increase in electricity related emissions (Figure 1.3). Emissions reductions from this sector must still occur in order to meet existing emissions targets proposed by the President in 2009 [72] and formalized in the President’s Climate Action Plan [36]. The potential of MPC to reduce emissions directly through reduced residential consumption, or through improved operation that allows increased renewable utilization, is indeed a promising approach to achieving these targets in a cost-effective way.




Figure 1.1: Growth in electricity retail sales. Source: EIA.




Figure 1.2: Projected electricity consumption under reference, high economic growth and high demand reduction technology cases. Source: EIA.

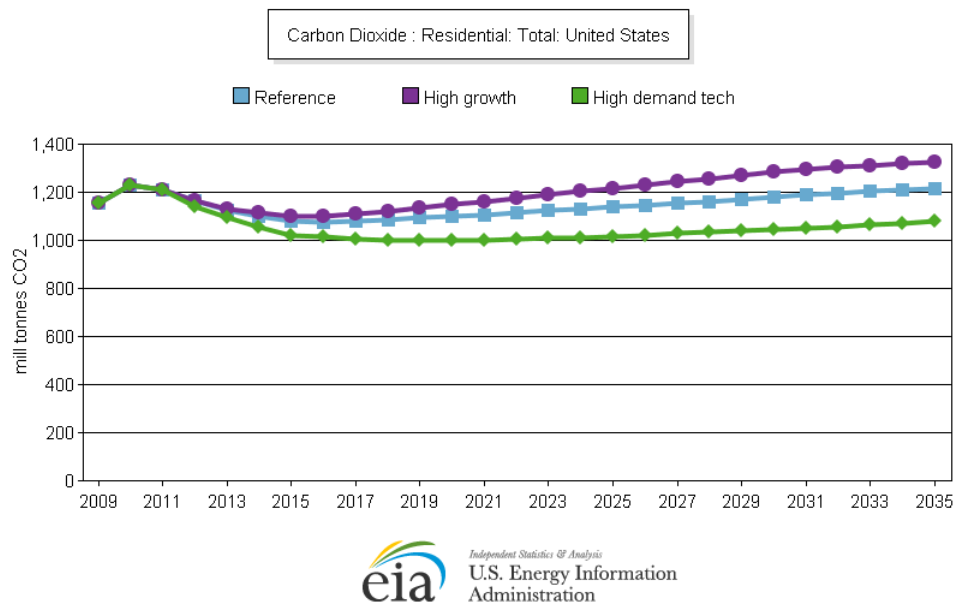


Figure 1.3: Projected carbon dioxide emissions under reference, high economic growth and high demand reduction technology cases. Source: EIA.

1.0.2 Research Questions

The impact of large-scale MPC applied to residential HVAC is not well understood, and the methods by which the maximum benefit to the home owner and utility is realized, are unknown. This results in a number of interesting research questions which will be considered in this research. These are:

1. The size of the MPC search space grows exponentially with each new decision variable. At scales exceeding tens of homes, a “supervisory” optimization approach combining all homes in a single optimization is intractable at a practical time-scale. This suggests that the optimization must be distributed, if not at the premises, then across multiple processors or computers at the least. How would distributed HVAC MPC be performed and can this be a template for real-world implementation?
2. Real-time price, being a proxy for electric demand, is an obvious candidate for use as a driving function in optimizing load. However, from the perspective of the utility, individual optimizations performed by each home based on real-time pricing results in decisions that, when aggregated across a large number of homes, has the potential to result in a less-than-optimal load shape because system level interactions are ignored. What, then, is the signal or driving function that must be fed to individual homes such that the aggregate response meets the objectives of the utility?
3. Assuming that the aggregated load from a collective of homes can be molded into a predictable and desirable shape, can distributed MPC be utilized to drive particular aspects of grid operation? That is, can distributed MPC be used to allow peaking plants to remain off during peak periods? Put another way, can distributed MPC be used for day-ahead resource planning, given a set of installed traditional generation?
4. Following logically from the assumptions above, what impact would distributed MPC have on electric grid operation in the presence of distributed generation, and would distributed

MPC allow a higher penetration of renewable energy to be utilized?

1.1 Research Objectives

In consideration of these questions, the following objectives define the major tasks undertaken in this research:

1. Develop and validate a simplified reduced-order residential building model suitable for large-scale HVAC MPC studies.
2. Select a power flow simulation package, obtain and modify distribution feeder models representative of the US electric power grid.
3. Demonstrate a distributed model predictive control framework for optimizing thousands of buildings that integrates reduced-order building models with power flow simulation.
4. Explore the impact of distributed model predictive control on electricity grid operation under demand-shaping objectives, with cases considering high levels of renewable energy contribution.

1.2 Survey of Literature

The following section presents a summary of literature relevant to this research, divided into three subsections. The first, Optimal Control of Building Systems, surveys various examples of MPC to building control with an emphasis on commercial applications. The second section enumerates a number of approaches to simplifying building simulation models, informs the simulation model presented in this research and provides support for the selected methodology. The final section surveys a number of recent studies involving demand side management in the residential context, highlighting the areas lacking in research that this work attempts to address.

1.2.1 Optimal Control of Building Systems

Model predictive control (MPC) of building energy systems has been studied over the previous two or more decades by combining mathematical programming techniques with a variety of building simulation methodologies. By incorporating simplified thermal and plant models with a direct search algorithm, [9] has shown potential for both energy cost savings and peak power reduction through optimal thermal mass control, depending on utility rate structure and building characteristics. Further examples by [37] and [44] have combined the Optimization Toolbox in [87] with TRNSYS [84] to perform similar studies on passive and active thermal storage systems in buildings, with similarly promising results. Another approach to MPC optimization has been demonstrated by [101] by incorporating the Nelder-Mead simplex algorithm into EnergyPlus [30] to show cooling cost savings associated with peak electric demand reductions.

Other software environments have been developed that provide additional means of investigating control optimization. GenOpt [93] can iteratively execute simulation programs using plain text input/output files until an optimal solution is found. GenOpt is designed to perform a single optimization per execution, rather than a sequence of optimizations continuously, e.g. a month of single day optimizations in sequence. [27] incorporated GenOpt into an MPC framework to examine zone temperature control optimization in office buildings and the subsequent effect on HVAC demand.

The “Energy Management System” was added to EnergyPlus to allow users to implement custom control strategies, yet this functionality is not readily extended to MPC. The Building Control Virtual Test Bed (BCVTB) [94] allows coupling between EnergyPlus and several other simulation environments. The BCVTB is designed to manage communication between EnergyPlus and the stand-alone application, where the latter provides control decisions at each simulation time step to be evaluated by EnergyPlus. This approach is best suited to the evaluation of pre-defined control algorithms.

The MATLAB-EnergyPlus framework developed by this author prior to this work has been

used to generate optimal window opening schedules for a small mixed-mode building. From these optimal results, [65] created a generalized linear model to extract near-optimal heuristics that could be implemented directly in direct digital control systems. This work was subsequently expanded in [66] and [64], using the optimization framework to generate optimized control strategies to which generalized linear models, classification and regression trees, and adaptive boosting techniques were applied. In [28], this author used the framework to optimize pump schedule and supply water temperature for a thermally activated building structure (TABS) simulation to significantly reduce energy consumption.

1.2.2 Reduced Order Modeling

Most of the examples found in the literature rely on lower fidelity, reduced order models, with few applications to detailed sub-hourly simulation models. Reduced order models, due to their simplicity, low computational cost, and relative ease of automated tuning, are a promising tool for researching MPC applications.

Using a five-parameter lumped capacitance-resistance model, [31] demonstrated that a second-order model captures the thermal response characteristics of a building. Parameter identification was performed using an iterative least-squares estimation technique using the Gauss-Newton method and site measured data from unconditioned experimental test cells. Data collected included easily measured indoor and outdoor temperature, but also heat-flux from solar radiation, wall surface temperatures and interior wall temperatures. Although good agreement was found between calibrated model prediction and measured data, the authors note several significant problems including sensitivity to parameters' initial guess value and large variation of parameter values that result in equivalent model fit. The authors caution that the applicability of this approach to heated multi-zone buildings must be questioned, and suggest that continuous re-calculation of parameters by a building management system may overcome such limitations.

Mathews, Richards and Lombard [63] propose a 2R-1C reduced order model which lumps all building mass into a single capacitance and all thermal resistances into interior and exterior

resistances. In the model, ventilation is considered to be an additional resistance. Convective and radiative heat transfer components are combined into constant values. Rather than calculating solar heat-flux directly, the model uses sol-air temperature as a driving function. To account for low-mass elements not included in the lumped capacitance, the authors introduce a phase-shift correction. Using measured indoor and outdoor temperatures, and solar insolation and wind speed measurements from nearby weather stations, the model was validated for 32 buildings ranging from test cells to small commercial. Lumped model parameters were calculated based on physical characteristics and phase-correction calculated from temperatures measured during free-float. The authors report that 80% of all temperature predictions fell within 2K of measured temperature, and 95% within 3K.

Gouda et al. [39] develop first and second order element models to be compared with a 20 order reference model, and results measured using a room in a campus building at the University of Northumbria. Using a nonlinear constrained optimization method, the model parameters of the first and second order model are estimated using the output of the previously calibrated 20th order model of the room. In each reduced order model, separate parameters were calculated for each of the two external constructions. To simulate the capacitance of the room air, an additional first order capacitance model was added. Results agreed well between first and second order and 20th order reference models using a step test applied to each of the three room models. The step response revealed significant differences between first and second order models, but little difference between second and 20th order model. Using inside and external temperature, and insolation data at 15 minute resolution measured in the room over a 30 day free-float period, the two reduced order models were tuned using the nonlinear optimization method. As with the simulation study, excellent agreement between the measured temperatures and the second order model were achieved. The second order model consistently predicted temperatures less than 0.2K lower than measured, but the first order model fared somewhat poorer, with predicted temperatures 1-2K lower than measured. The authors suggest that the second order element model is sufficiently accurate in modeling the short-term thermal response required for simulation of control systems, at only 6%

of the computational cost of the more detailed benchmark model.

Braun and Chaturvedi [11] present a thermal network model of a multi-zone commercial building for the purpose of online model predictive control. In the model each zone is treated separately, having 4 wall elements composed of a 3R-2C model each. Window glazing was treated separately as a simple massless resistance element. The solution of the heat balancing equations describing the combined multi-zone model was expressed using a state-space formulation converted into a transfer function. The model was then trained using a combination of direct search global optimization and local nonlinear regression from cooling loads generated by a TRNSYS building model. The objective function minimized in training was the sum of RMSE of cooling loads for the entire training period. The authors found that 14 days of training data generated by the detailed TRNSYS model was sufficient for calibration. A separate field study was conducted in Chicago using a single zone model, calibrated using two weeks of measured temperature and calculated cooling loads. Using a four week validation data set, the authors found that cooling loads could be predicted within 9%.

Xu and Wang [98] develop a new hybrid model combining a conduction transfer function (CTF) model of the external envelope with a 2R-2C thermal network model of the internal mass. The latter has parameters tuned automatically by a genetic algorithm (GA). The authors suggest that the new hybrid model allow more convenient construction of the CTF model based on detailed physical properties of the building envelope. Because the estimation of the internal mass is much more difficult due to the number and variety of elements in a physical building, the thermal network approach allows easy parameter estimation. Model calibration and validation was performed using two summer weeks and one winter week of temperature and solar data collected on a 50 story building on Hong Kong. Cooling load and internal temperatures from the model prediction compared favorably with measured data cooling load and average internal temperature with approximately ten percent relative root mean square error typical of the former and 0.58K typical of the latter. The authors compare results from their previous studies of the same building [89, 97] using a 3R-2C model which predicted nearly identical cooling loads and internal temperature.

By applying an Extended Kalman Filtering technique, O'Neill et al. [74] develop a state space representation of a reduced order model in order to estimate the internal loads of zones within a class and office room building on the University of California at Merced. The reduced order model adapts the 3R-2C wall model from the work of Braun et al. [13], adding a parallel resistance path for glazing elements and an air-node mass balance. Solar distribution is applied using an area-weighted fraction. The model is validated using a previously constructed and calibrated EnergyPlus model of the building. Internal loads estimated by the EKF technique are within 10% of the load calculated by EnergyPlus 93% of the time. The authors suggest the differences are caused by 3R-2C vs CTF and absence of long-wave radiation exchange. The technique was then applied to the building by taking supply and room air measurements, including temperature and humidity, taken at two minute resolution in seven rooms. Real-time weather data at the same resolution from external sensors was used for the load estimation. Unfortunately, the supply air temperature to each room was not available, making internal load measurement difficult. No clear conclusions regarding the accuracy of the method for estimating load of the building were presented.

It is fairly clear from the literature that significant computational expense can be avoided by using a variety of reduced-order modeling techniques, at little loss in accuracy. A common theme among many of these studies is the use of R-C networks, with 3R-2C envelope elements dominating the literature. The literature demonstrates that these models are suitable for capturing the physical processes of heating and cooling well, suggesting they are suitable for HVAC control studies. Given the simplicity of both specifying and developing these models, this approach is adopted for this work.

1.2.3 Residential Demand Side Management and Control Optimization

A number of studies have been recently published which combine residential demand side management (DSM) and optimal power flow modeling, where DSM is accomplished through direct load control or price-responsive automatic control. Several consider scenarios in which renewable energy contributes to the generation mix. These studies provide both theoretical and practical

foundations for the proposed research, but also highlight areas of opportunity.

In a discussion of electricity dispatch optimization, Wang et al. [90] propose a “Ramp Market” to provide ancillary services for electric grids where significant wind generation and demand response participation are present. The proposed market would allow both fast and slow ramping technologies to bid individual ramping capacity in order to match supply to demand and vice-versa. In order to accomplish this, the authors propose a multi-tiered dispatch optimization in which each tier has different objectives that are coordinated between both parent and child tier. Examples of tiers include: RTO/ISO tier for optimizing regional transmission objectives; a market tier which provides portfolio optimization and aggregation; a plant tier that conducts optimal resource scheduling; an aggregator tier with detailed neighborhood based load optimization; a household tier with appliance, electric vehicle and building load optimization.

Hindi et al. [45] describe a “multi-rate” MPC methodology for reducing demand-supply imbalance. Multi-rate MPC combines two types of load control services: regulation, which occurs on the order of seconds and minutes, and so-called fast demand response, which is defined as demand response that takes effect immediately such as disabling an air-conditioner compressor. The authors suggest that the combination of these two types of services provides a new ancillary service that can be implemented using direct demand-supply imbalance signals as well as indirect price-based signals, e.g. real-time pricing. Simulation results show that trade-offs occur between the two services when they are combined.

In [95], Widergren et al. consider a distributed optimization method for demand side management where residential HVAC systems and water heaters are supplied real-time pricing signals. These systems participate autonomously in a five minute market at the feeder level, by bidding capacity in real-time. Because decisions are made autonomously, market bidding results in distributed optimization without centralized control. Using GridLAB-D, 300 building types based on 140 typical customer demand shapes were simulated both with and without the feeder-level market using a synthetic real-time pricing signal. Simulation results show the potential for significant bill savings by individual households. The authors do not address the issue of system stability.

Li, Chen and Low [60] propose a method whereby households exposed to real-time pricing optimally schedule power consumption in order to maximize individual benefit. The authors describe a distributed optimization scheme in which the utility sets the price of electricity according to the marginal cost, each household then optimizes and the resulting load is fed back to the utility. Together the households and the utility iteratively compute prices, operation schedules, and consumption for the next day until converged. The advantage of this approach is that the utility does not need to know the underlying constraints and operational details of the individual households, just the aggregate result. Simulation studies using this approach consider eight homes each with various appliances, air conditioning, PHEVs and storage batteries. The studies find that households coordinated indirectly by real-time pricing to flatten total demand, reduce peak load, and minimize variation. Simulations in which additional battery storage is present show that the additional storage allows more flexible scheduling and increased load factor. Although the demand leveling benefit increases with an increase in participating households, the benefit to household and utility eventually shows diminishing returns. Perhaps the most interesting result is that the method also results in maximum utility benefit, i.e. social welfare.

Fuller, Schneider and Chassin [38] investigate the effects of a transactive residential HVAC controller in a centralized double-auction market. The transactive HVAC controller is allowed to adjust cooling set point according to price determined by a market with five minute clearing interval, and bid the price back into the market. Using 1497 single-family residences and IEEE 123-node test feeder modeled in GridLAB-D, a simulation study is performed for the month of August in Seattle, WA. In the study, transactive controllers are limited to 381 of the residences. Results are compared to simulations in which the controllers are allowed to adjust set points in response to price, but are not allowed to bid; this is known as an “active” market. In the active market, price encourages the controllers to delay cooling until price decreases in the afternoon, creating a peak demand that is higher than the peak seen in flat rate simulations. Transactive control simulations reduce peak demand, but do not result in reduced energy consumption and increase cost to the consumer. Although not intended to be an optimal control scheme, the simulation suggests that

detailed simulation studies such as these must be performed to understand the potential impacts of smart grid operation.

Oveido et al. [77] investigate a decision tree-based approach using heuristic rules for managing the demand in residences with plug-in hybrid electric vehicles (PHEV), local renewable electricity generation (both wind and solar) and time-of-use (TOU) pricing. The rules govern what source of energy is used, with priority given first to renewable generation if available, then to PHEV battery storage, and finally electricity supplied by the grid. The simulation study suggested that these simple rules when paired with a TOU pricing program could reduce grid-supplied electricity consumption by 14.8%, resulting in cost savings of up to 22.6%. A second study investigated the robustness of grid stability under two conditions, and showed that the heuristics reduced the potential for overloading the local feeder, provided that a communication network between PHEVs on the same feeder prevented concurrent charging. When communication did not occur, the probability for destabilizing the feeder increased. Although PHEV controls are not considered in the research presented here, this study provides an interesting future topic for exploration using the techniques developed in this work.

Zhang et al. [100] propose a simplified 2R-2C building model to understand the aggregate behavior of a mixed population of buildings under demand response events. A detailed simulation model was created in GridLAB-D from two thousand detailed building models whose physical parameters were drawn randomly from normal distributions. K-means cluster analysis was then used to generate five clusters representing the states and properties of the aggregate model. Both models were subjected to a simulated thermostat set-up demand response event and the aggregate model correctly captured the behavior of the detailed model both during the event, and in the rebound immediately following. The authors also show that the aggregate model is able to capture the effects of the detailed model under a series of set point changes.

Cecati et al. [23] evaluate an Energy Management System (EMS) that combines optimal power flow simulation, demand side management and active management schemes (via real-time pricing) to incorporate different types of renewable resources, improve elasticity, and reduce con-

sumer energy costs. The EMS operates at two time-scales: day-ahead scheduling subject to anticipated market prices; real-time optimization that modifies the schedule according to current demand. The proposed system evaluated a 19-bus radial distribution system (incorporating two wind turbines) simulated with the MATLAB SimPowerSystems toolbox. Simulation results suggest the system is able to increase utilization of wind generation through active distribution.

Despite the interest in residential demand side management, very few residential MPC applications appear in the literature. As noted above, studies focus either on direct load control, i.e. demand response, or on indirect market-based control mechanisms. While a number of these studies do consider the large-scale impacts of both centralized and distributed control mechanisms, the MPC studies that do appear tend to favor price-based optimization. It is far from clear that price-based optimization is the ideal solution to managing demand, illustrating the need for studies exploring alternative methods. Further, the interaction of large-scale DSM with renewable energy sources is not well represented in the literature, suggesting an additional area of fertile research.

It is the objective of this work to address these needs by investigating the opportunities afforded by large-scale distributed residential HVAC MPC for shaping electric demand, as an alternative to those already presented in the literature. Coupling distributed HVAC MPC with power flow simulation at the distribution feeder level will illustrate the potential for minimizing aggregate peak demand, and highlight important considerations for DSM, regardless of approach. This work will also consider scenarios in which rooftop solar and utility scale wind is present, where the variability in generation can result in feeder demand characteristics undesirable to the utility.

1.3 Organization of Document

The organization of this work is as follows: First, a methodology is described outlining the various aspects of the research, including an overview of the simulation environment, and an experimental plan enumerating the cases considered. Next, a description of the optimization framework and the models developed for the research are presented. A chapter investigating a simple demand response scenario motivates the need for model predictive control. Four chapters

describing the optimization cases and their results are then presented in detail. Two chapters follow investigating the potential for the methodology developed in this work to address high and low levels of renewable energy penetration. The work concludes with a high-level summary of findings, conclusions and suggestions for future work.

Chapter 2

Methodology

This chapter presents an overview of the methodology developed for this work. A research methodology describing the general approach is presented first, followed by a description of key metrics used to evaluate results throughout. The simulation environment developed during this research is then described in detail. The chapter concludes with a description of the experimental cases considered.

2.1 Research Methodology

This research considers the potential for a large number of residential buildings to shape electrical demand according to arbitrary global objectives through model predictive control. Here, “arbitrary” is intended to mean an objective which can vary from day to day depending on operating conditions, but which has a very specific goal of manipulating demand to alleviate the pressures placed on electric generation assets. “Global” in this context is defined as an objective considering a representative portion of the United States electric grid, specifically, a set of prototypical distribution feeders developed for such investigations. These objectives are designed to result in a paradigm shift: from generation that follows load to load that follows supply.

As there are a lack of tools that are capable of performing these investigations *in toto*, a significant portion of this work was devoted to developing the framework with which the research is performed. This research has required the development of new software coupling the simulation and optimization of many thousands of residential buildings with an open-source distribution system

simulation package. The requirement for optimizing many thousands of buildings concurrently has necessitated that a distributed optimization approach be adopted, and that a reduced-order building model be developed that is relatively simple and fast, yet realistic in its behavior. The latter concern is addressed through model cross-validation against several well-validated simulation software packages. To capture the cycling behavior of conditioning equipment, this work requires the simulation of HVAC control at sub-hourly time steps. Minimum compressor run-time, cooling coil and fan staging, and thermostat hysteresis have been modeled to better simulate HVAC electric demand.

Electric distribution feeder models utilized in this work are selected based primarily on climate, and secondarily on their residential load characteristics. This work is concerned principally with the control of residential air-conditioning loads, therefore, the set of feeders selected are dominated by residential load, and are physically located in areas of the country where air-conditioning load is a major contributor to peak demand. Additional consideration is given to areas in which there exist relatively large renewable energy penetration or potential. Out of the twenty-four feeder models made available by PNNL [82], three are studied here.

The selected feeder models are stripped of all residential loads and modified such that residential loads calculated by the reduced-order building model can be injected as loads from an equivalent electrical circuit. Home characteristics, originally selected by random sampling of US Census data, are maintained in the process. This model, which combines the original feeder with the injected residential loads, is referred to as the “hybrid model”. Following validation of the hybrid model against the original, the hybrid model can then be simulated in distribution power flow software to study the aggregated system behavior under different optimal control scenarios.

This work first briefly investigates the impact of a demand response event during a peak demand day on the distribution feeder(s). The demand response study is intended to illustrate a worst case control scenario and motivates the need for improved control and system-level coordination. The distributed optimization framework is then used to minimize peak demand at each home, resulting in a feeder-wide reduction in peak demand. An investigation into the demand shifting

potential assists in explaining the result.

Electricity cost minimization under dynamic pricing structures is then studied to determine the impact on feeder demand. Using historical price and weather data, day-ahead electricity price is modeled with classification and regression trees (CARTs) for use with TMY weather data. Based on the observed behavior, experiments are conducted using a synthetic price signal intended to prevent the undesirable characteristics that emerge with day-ahead price optimization. A further experiment using synthetic prices illustrates problems associated with electricity price optimization.

A novel new approach to optimizing the aggregate response of buildings is developed and studied. This approach relies upon a reference signal, containing demand forecast information at the feeder level, to shape the demand at each individual residence. This method is then used to show how demand can be manipulated at the residences to alleviate problems associated with distributed generation, namely rooftop solar electric and utility scale wind.

2.2 Metrics for Evaluating Optimization Performance

Many of the results presented in this work are time series graphs. A few key performance metrics are proposed to assist in summarizing and interpreting the findings. These metrics, which are all reported at the feeder distribution transformer, are described below.

2.2.1 Total Electricity Consumption

Total electricity consumption is simply the sum of daily (or monthly) demand in megawatts divided by the time step in seconds. Total energy consumption difference between base case and optimized case gives an indication of the “efficiency” of the optimization. None of the optimization objectives studied in this research explicitly consider net energy consumption, so this metric serves to show how optimization changes total energy consumption.

2.2.2 Peak Demand

Peak demand is the maximum electricity demand in megawatts, measured from simulation time series results. The metric is calculated for each 15-minute period of the simulation time frame using a simple moving average. Peak demand determines the maximum generation requirements of the feeder. Values presented in this work are assumed to be 15-minute peak demand unless otherwise specified.

2.2.3 Peak to Valley Ratio

Peak to valley ratio is simply the ratio of the peak demand to the minimum demand calculated over a given simulation time frame. This metric is calculated for each 15-minute period of the simulation time frame. Peak to valley ratio gives a measure for the amount of generation that must be brought online to satisfy demand. A value of one indicates a constant load. This metric gives a sense for the magnitude of the dispatchable generation required to meet demand.

2.2.4 Load Factor

Load factor is a ratio of the average demand over the peak demand calculated over a given simulation time frame. Load factor indicates how well the generation assets in the system are utilized. Load factors approaching unity indicate a nearly constant use of the assets. The metric is calculated using the 15-minute moving average of demand consistent with other metrics.

2.2.5 System Ramping

System ramping is a measure of the total absolute change in electricity demand from one measurement to another over a given simulation time frame. This metric is calculated on a 15-minute basis by smoothing the five-minute simulation result using a simple moving average. This metric, presented in megawatts, indicates how much demand changes in total over the course of a day or month. High values indicate large time step to time step variations, but do not quantify how quickly demand may change.

2.2.6 Feeder Demand Spectral Density

This metric quantifies the temporal effects of optimization on feeder demand using an approach adopted from signal analysis. Whereas the system ramping metric previously described quantifies the absolute demand change over given simulation time frame, this metric helps explain how that demand is distributed over the time frame in the frequency domain. That is, it illustrates how much of the demand change occurs at frequencies ranging from minutes to hours. A large value of signal power at a one hour frequency would indicate significant demand fluctuations each hour.

To calculate this metric, the total optimized feeder demand time series is first converted into the frequency domain using a discrete-time Fourier transform (DTFT). The transformed time series coefficients are used to develop the power spectral density (PSD) of the original feeder demand curve. This spectral density is then normalized so that the percentage of power present in the signal between any two frequencies can be calculated and compared to the base case result.

2.3 Simulation Environment

As the tools required for the simulation studies in this research only partially existed, a simulation environment capable of combining model predictive control, reduced-order building simulation, and distribution feeder simulation was developed. The following section describes the overall structure of this environment, followed by a description of the power flow simulation software and its selection. This section covers the environment at a high-level; specific details related to the building and distribution feeder model development, and optimization scheme are presented in Chapters 3 and 4.

2.3.1 Background

In previous work [28], the author describes an MPC framework which relied upon MATLAB as the platform linking optimization to modeling engine. This approach was used with great success

for QCoefficient, Inc. (formerly Clean Urban Energy, Inc.), for investigating optimal TABS control strategies, and for the research by group members, most notably in [65, 66, 64, 28, 86, 73].

While successfully used for more than three years, this framework was not well suited for the studies undertaken in this work for a number of reasons:

1. The framework was intended for single building optimizations; introducing thousands of buildings into the simulation would result in a major re-write of the code.
2. Although extensions to MATLAB have been added that make it a much more object oriented language, other more mature languages exist that would simplify the code significantly.
3. From a practical standpoint, MATLAB is expensive, and free alternatives to many of the algorithms and toolboxes are available.
4. Related to the previous point, parallelization in MATLAB is limited to a fixed number of cores; for fast computation, many more would be required.

Weighting potential time-savings against the added effort required to write new code, the decision was made to develop a new framework in Java.

2.3.2 Simulation Framework

The simulation framework, termed “GridMPC”, developed for this research is loosely modeled after the environment described in [28], but differs in implementation details. GridMPC includes utilities for parsing simulation input file and weather files (TMY2 or EPW), and for writing files required by the power flow simulation software.

The simulation input file contains the characteristics of each individual residential building, the optimization objective and parameters, and the distribution feeder used by the power flow software. Depending on the size of the distribution feeder studied, this file measures between

45,000 and 65,000 lines of text. Additional tools and scripts written in various languages have been developed to programmatically generate the input file, as editing each by hand is not tenable.

The output of GridMPC consists of: the optimized simulation results for each home, with a value at each time step (row), for each end-use (column); the input “player” file required by the power flow simulation, containing the sum of all electricity demand at each time step; a CSV file containing the daylight-time-corrected weather interpolated to each time step; a power flow simulation input file linking the distribution feeder templates with the building simulation and weather files.

GridMPC is initiated from a shell script at the command line. The simulation follows a process somewhat different from its predecessor:

1. Read simulation and weather files.
2. Generate a unique building model for each home described by the input file.
3. Create a pool of simulation threads to execute the following steps in parallel for each building:
 - 3.1. Auto-size HVAC equipment if size is not specified in input file.
 - 3.2. Warm-up building to initial state unless continuing from previous simulation.
 - 3.3. Generate candidate decision vector, a time series of control decisions.
 - 3.4. Simulate planning horizon using candidate vector
 - 3.5. Evaluate fitness and exit criteria. If exit criteria are satisfied, continue to step 3.6, otherwise return to step 3.3.
 - 3.6. Write simulation result and “player” files.
 - 3.7. Advance one execution horizon. If no more planning horizons are defined by input, continue to step 4, otherwise, return to step 3.3.
4. Write CSV weather and power flow files.

5. Initiate the power flow simulation.

In the predecessor framework, an initialization horizon was required to set the thermal state of the modeling engine (EnergyPlus). With a reduced-order model now in Java, state can be set explicitly. Furthermore, when evaluating a new candidate vector, the software need only rewind the simulation clock; the previous building state is maintained. Since the initialization horizon is no longer required, considerable computation is avoided.

2.3.3 Electric Grid Simulation Software

Software capable of accurately simulating the electric distribution system is required for understanding the aggregate effects of optimal control strategies. Both free and commercial software packages are available depending on the type of studies required. Given limited funding, free and open-source software (FOSS) was a gating requirement. Luckily, a number of relatively mature FOSS options exist. A report authored by CanmetENERGY [67] evaluates several open-source options including GridLAB-D [24], OpenDSS [75], and APREM [76]. While all three would be suitable for the needs of this research, GridLAB-D was selected for a number of reasons, including:

1. Active online help forums and code development
2. Ability to run on common cluster computer platforms (Linux)
3. Relatively fast simulation time
4. Use of common weather file formats
5. Extensive online documentation and training materials
6. Distributed generation models
7. Availability of prototypical distribution feeder models

GridLAB-D is a time series power system simulation tool developed by PNNL for the purpose of evaluating Smart Grid technologies. The project is funded by the Department of Energy Office

of Electricity Delivery and Energy Reliability. The software models all components of the distribution system in detail, from switches, to distribution lines, transformers, and residential end-uses. Validation of GridLAB-D has been performed at the aggregate level using measured data from the Olympic Peninsula Demonstration Project [16]. Version 3.0¹ was used for this research.

Although the software contains a well validated residential building model derived from Sonderegger’s Ph.D. research [85], the decision was made to perform the simulation of building load and optimization outside of GridLAB-D. Utilities in the software for injecting external loads, and the modular agent-based structure make this task relatively easy. Details regarding the method for injecting loads into the power flow simulation are provided in Section 3.2.3.

The selection of GridLAB-D as the grid simulation framework necessitates the simulation of thousands of residential buildings. This level of detail is necessary in the power flow simulation to accurately simulate the distribution system.² Admittedly, this research does not leverage the capabilities of GridLAB-D to the fullest extent. For example, the second order effects of control optimization on transformer and line losses can be easily quantified by GridLAB-D. This could be an interesting area of research made capable by the simulation environment created for this work, and is suggested for future consideration.

2.4 Simulation Hardware

A dedicated computer cluster was constructed to perform the simulation experiments outlined below. The cluster is composed of four nodes, having 8 virtual cores (4 hyper-threading physical cores) running at 3.4GHz with 8GB of RAM each. The nodes are connected via gigabit ethernet to 180GB of shared solid-state storage.

The mini-ITX form-factor of each node allows all four nodes to be housed in a 3U rack-mount

¹ When this research began, version 2.3 was the latest stable release. Version 3.0 was released as a candidate late in 2013 and the current stable version released at the beginning of 2014.

² Whether or not this requires unique building models is a matter of debate. Section 11.2 identifies this as an area of future research.

server chassis. Final cost of the cluster, including chassis, cables and router, totaled slightly under two thousand dollars. The purchase of the hardware was enabled by a grant-in-aid from ASHRAE.

The topology of the cluster is modeled after that described in [17]. The names of the four nodes are, appropriately, Kennedy, Bellman, Glover and Holland. When fully utilized, the cluster is capable of running 32 simultaneous optimizations. Optimizations are coordinated by TORQUE Resource Manager [4]. A one month long optimization of a feeder model containing 1500 homes completes in roughly one hour.

2.5 Experimental Plan

This work considers a large number of experimental simulation cases for three climate and distribution feeder pairs (see Section 3.2.1 for description). Each simulation case considers two levels of participation measuring 30% and 70% of the feeder residential population. Participants optimize HVAC operation according to the specified objective. Non-participants simulate HVAC operation according to the default cooling schedule. With the exception of the first set of demand response experiments, all experiments consider the entire month of July, a sufficiently hot month in each of the three climates.

The following subsections summarize the experiments conducted and describes each experiment at a high level. Each experiment is presented in a separate chapter, with its own section describing the methodology in more detail.

2.5.1 Demand Response

The first experiments consider the annual peak demand day under two and six hour demand response event durations. This experiment is intended to illustrate an extreme case of the current practice, and provide motivation for improved HVAC control and system-level coordination.

2.5.2 Demand Limiting Optimization

These experiments apply a simple peak demand minimizing objective function to each home in order to reduce overall feed peak demand. Because demand at the individual home level is not smooth due to the cycling behavior of the HVAC, experiments considering various time-averaging windows are performed. These experiments are intended to examine a completely decentralized and uncoordinated optimization approach in contrast to the previous demand response experiments.

2.5.3 Dynamic Price Optimization

Dynamic price optimization might also be described as day-ahead market price (DAM) optimization. These experiments utilize hourly electricity price forecasts in order to indirectly coordinate HVAC operation. Each building seeks to minimize daily electricity costs, leading to unintended consequences under certain conditions. Additional experiments are conducted using a smooth synthetic price signal in an attempt to remove the undesirable system-level features. These experiments are intended to show how coordination can be achieved in a decentralized optimization, and illustrate the potential problems with such an approach.

2.5.4 Load Shape Optimization

These experiments examine a novel new method of decentralized optimization which improves upon the former dynamic price optimization results. The experiments employ a “reference signal” to coordinate HVAC operation in order to achieve a desired aggregate feeder demand profile. Multiple iterations of the concept are explored, resulting in various levels of performance. The best performing case is adopted as the optimization method for the remaining experiments.

2.5.5 Rooftop Solar

In the previous experiments, the general objective was to shape feeder demand in a way that minimizes demand, increases load factor, or otherwise improves the metrics described previously. In this set of experiments, generation in the form of rooftop solar electric is introduced at two levels of

penetration to show the effectiveness of the reference signal approach to: 1) shaping feeder demand under these conditions, and 2) alleviating the problems associated with intermittent generation within the distribution feeder.

2.5.6 Utility Scale Wind

These experiments show the effectiveness of the reference signal method to shaping feeder demand in the presence of utility scale wind generation. In these experiments, the reference signal contains information about the system load *outside* of the distribution feeder, providing an example load shaping to satisfy needs at a higher level. Two levels of penetration corresponding to 9.4% and 25% of annual feeder electricity consumption are examined.

Chapter 3

Model Selection and Development

This chapter details the development and selection of the simulation models used throughout the research. Those models created for this research are cross-validated using accepted validation procedures; models adopted from previous research which have been modified for this research are cross-validated against the original. In some cases, models are assumed to be well validated and simply adopted outright. HVAC models adopted from existing building energy simulation packages fall into this category.

3.1 Building Model

The building model developed for this research combines a number of elements to estimate the electric demand of the residence. These elements include: an envelope model for estimating heating and cooling requirements; HVAC equipment models to calculate electric demand associated with heating and cooling; a thermostat model to control the heating and cooling operation of the HVAC equipment; simplified end-use models for the calculation of electric demand (and the heat gains they produce) from miscellaneous electric loads, lights, appliances and electric hot water heaters. Each of the elements are described in the following section.

3.1.1 Envelope Model

Using an electrical circuit analogy, a building envelope may be expressed as a thermal network of resistive and capacitive elements. The model chosen for this work consists of six components

that represent roof, walls, glazing, floor, air and internal mass as illustrated in Figure 3.1.1 below.




Figure 3.1: Building envelope model expressed as a thermal network.

Energy balances may be formulated around individual elements to create a system of ordinary differential equations, which can be discretized in time and solved analytically. For the exterior wall node, the energy balance can be expressed as:

$$\sum Q_{wo} = Q_{solwall} + h_{wo}A_w(T_o - T_{wo}) + \frac{A_w}{R_w}(T_{wi} - T_{wo}) + C_{wo}\frac{dT_{wo}}{dt} \quad (3.1)$$

where:

$\sum Q_{wo}$ is the energy balance at the outside wall node

$Q_{solwall}$ is the energy gain due to insolation

h_{wo} is the exterior film coefficient

A_w is the wall area

T_o is the exterior dry-bulb temperature

T_{wo} is the exterior wall temperature

R_w is the wall thermal resistance

T_{wi} is the interior wall temperature

C_{wo} is the exterior wall thermal capacitance

$\frac{dT_{wo}}{dt}$ is change in exterior wall temperature with respect to time

Similarly, an energy balance around the wall interior node can be formulated. Here, the assumption is that solar gains through the windows are equally distributed to the envelope and interior mass nodes according to the respective areas of the elements.

$$\sum Q_{wi} = \frac{Q_{solwin}A_w}{A_t} + h_{wi}A_w(T_i - T_{wi}) + \frac{A_w}{R_w}(T_{wo} - T_{wi}) + C_{wi}\frac{dT_{wi}}{dt} \quad (3.2)$$

where:

ΣQ_{wi} is the energy balance at the inside wall node

Q_{solwin} is the total insolation passing through all windows

h_{wi} is the interior film coefficient

A_w is the interior wall area

A_t is the total interior wall, roof, floor and mass area

T_i is the interior dry-bulb temperature

C_{wi} is the interior wall thermal capacitance

$\frac{dT_{wi}}{dt}$ is change in interior wall temperature with respect to time

It follows that the same equations can be formulated for the interior and exterior ceiling and floor nodes. The exterior floor node is an exception if in contact with the ground, (i.e. no crawlspace). In this case, the energy balance can be expressed as:

$$\sum Q_{fo} = \frac{A_f}{R_s}(T_s - T_{fo}) + \frac{A_f}{R_f}(T_{fi} - T_{fo}) + C_{fo}\frac{dT_{fo}}{dt} \quad (3.3)$$

where:

ΣQ_{fo} is the energy balance at the outside floor node

A_f is the floor area

T_{fi} is the interior floor temperature

R_f is the floor thermal resistance

T_{fo} is the exterior floor temperature

R_s is the soil thermal resistance

T_s is the deep ground temperature

C_{fo} is the exterior wall thermal capacitance

$\frac{dT_{fo}}{dt}$ is change in exterior floor temperature with respect to time

The energy balance formulated from the perspective of the mass node is:

$$\sum Q_m = \frac{Q_{solwin}A_m}{A_t} + h_{si}A_m(T_i - T_m) + C_m \frac{dT_m}{dt} \quad (3.4)$$

where:

ΣQ_m is the energy balance at the mass node

A_m is the mass area

T_m is the mass temperature

C_m is the mass thermal capacitance

$\frac{dT_m}{dt}$ is change in mass temperature with respect to time

And finally, the energy balance formulated from the perspective of the air node is:

$$\begin{aligned} \sum Q_i = Q_{inf} + Q_{int} + Q_{hvac} + h_{wi}A_m(T_m - T_i) + h_{ri}A_{ri}(T_{ri} - T_i) + \\ h_{fi}A_{fi}(T_{fi} - T_i) + h_{wi}A_{wi}(T_{wi} - T_i) + \frac{A_g}{\frac{1}{h_{gi}} + R_g + \frac{1}{h_{go}}}(T_o - T_i) + C_i \frac{dT_i}{dt} \end{aligned} \quad (3.5)$$

where:

ΣQ_i is the energy balance at the air node

Q_{inf} is the energy gain/loss due to infiltration

Q_{int} is the energy gain due to internal sources

Q_{hvac} is the energy gain/loss due to HVAC

A_g is the area of windows

h_{gi} is the interior window film coefficient

R_g is the thermal resistance of the windows

h_{go} is the exterior window film coefficient

C_i is the air thermal capacitance

$\frac{dT_i}{dt}$ is change in air temperature with respect to time

3.1.2 Solar Gains

Surfaces in the model may be designated as solar exposed. For each of the solar exposed surfaces, total solar insolation is calculated from the beam, diffuse, and horizontal components using the model proposed by Liu and Jordan for an isotropic clear sky[61].

$$I_t = \left[I_b \cos(i) \alpha F_{SL} + I_d \cos^2 \left(\frac{\beta}{2} \right) + I_h \rho \sin^2 \left(\frac{\beta}{2} \right) \right] A_{surface} \quad (3.6)$$

where:

I_t is the total insolation on the surface

I_b is the beam component of the solar radiation

$\cos(i)$ is the angle of incidence of beam radiation

F_{SL} is the sunlit fraction

α is the absorptance of the surface

I_d is the diffuse component of the solar radiation

β is the tilt angle of the surface

I_h is the total horizontal solar radiation

ρ is the ground albedo

$A_{surface}$ is the area of the surface

3.1.3 Glazing Model

The glazing model is a very simple extension of the opaque surface model. For each solar exposed glazing, total solar insolation passing through the glazing is calculated from:

$$I_t = \left[I_b \cos(i) F_{SL} + I_d \cos^2 \left(\frac{\beta}{2} \right) + I_h \rho \sin^2 \left(\frac{\beta}{2} \right) \right] A_{glazing} SHGC \quad (3.7)$$

where:

$SHGC$ is the solar heat gain coefficient

$A_{glazing}$ is the area of the glazing

Strictly speaking, $SHGC$ varies with incidence angle, and is different for direct and diffuse components of insolation. The simplification made here in order to reduce simulation time, will tend to overestimate solar gains at high incidence angles.

3.1.4 Shading Model

Shading by overhangs and fins is calculated for all solar exposed surfaces following the procedure outlined in [6]. The sunlit area of a surface can be calculated by first finding the length and height of the shade-line projection from adjacent fins and overhangs on the surface in question:

$$S_W = P_V |\tan \Delta| \quad (3.8)$$

$$S_H = P_H |\tan \Omega| \quad (3.9)$$

where:

S_W is the length of the shadow from the vertical fin

P_V depth of the fin

Δ is the fin projection profile angle

S_H is the length of the shadow from the horizontal overhang

P_H depth of the overhang

Ω is the overhang projection profile angle

and then applying the following equation that relates surface, fin and overhang geometry:

$$A_{SL} = [W - (S_W - R_W)] [H - (S_H - R_H)] \quad (3.10)$$

where:

A_{SL} is the sunlit area of the surface

W is the surface width

H is the surface height

R_W is the distance between surface and fin measured horizontally

R_H is the distance between surface and overhang measured vertically

The sunlit fraction, F_{SL} , is calculated as simply the ratio of sunlit area to total surface area, and then used to multiply the beam solar radiation component of Equations 3.6 and 3.7. Shading fractions are pre-calculated in the model for each surface and hour in the simulation time frame, avoiding re-calculation during optimization.

3.1.5 HVAC Models

Definitions for the HVAC component models summarized below are described in detail in the EnergyPlus Engineering Reference [3] and will not be reproduced here. These models are largely derived from the ASHRAE HVAC 2 Toolkit [8]. HVAC components implemented in the reduced-order model include:

1. Simple Hot Water Boiler. A constant efficiency boiler with configurable nominal capacity and efficiency that can use electricity, natural gas, propane or oil as heating sources.
2. Single-Speed Electric DX Air Cooling Coil. A single speed direct expansion air cooling coil with configurable nominal capacity and coefficient of performance (COP). Energy input ratio (EIR), part-load ratio (PLR) and capacity correction coefficients are configurable.

3. Electric Air Heating Coil. A constant efficiency electric heating coil with configurable nominal capacity and efficiency.
4. Gas Air Heating Coil. A constant efficiency gas heating coil with with configurable nominal capacity and efficiency. The model has been extended to simulate propane as well.
5. Single-Speed Electric Heat Pump DX Air Heating Coil. A single speed direct expansion air heating coil with configurable nominal capacity and COP. EIR, PLR and capacity correction coefficients are configurable.
6. Water Source Electric DX Air Cooling Coil. A single speed direct expansion air cooling coil with nominal capacity and COP. EIR, PLR and capacity correction coefficients are configurable.
7. Water Source Electric Heat Pump DX Air Heating Coil. A single speed direct expansion air heating coil with nominal capacity and COP. EIR, PLR and capacity correction coefficients are configurable.
8. Hot Water Baseboard Heater with Only Convection. A hot water baseboard with configurable overall heat transfer coefficient (UA) and flow rate.
9. Electric Baseboard Heater with Only Convection. A constant efficiency electric baseboard with nominal capacity and efficiency.
10. Direct Evaporative Cooler. A simple evaporative cooler model with configurable nominal saturation efficiency.
11. Constant Volume Fan. A simple fan model with configurable nominal flow rate, pressure rise, fan efficiency and motor efficiency.
12. Constant Speed Pump. A simple pump model with configurable nominal pressure rise, flow rate, motor efficiency and total efficiency.

13. Dual-Mode Single Set Point Thermostat. A thermostat model allowing separate heating and cooling schedules.

From the individual components, HVAC system models are constructed. These systems include:

1. Forced air furnace. Combines constant volume fan with electric, natural gas, propane or oil air heating coil, and dual set point thermostat.
2. Baseboard heating system. Uses electric baseboard heater (convection only) which has been extended to natural gas and propane sources as well. The latter are assumed to vent exhaust gas to the outside.
3. Boiler with hot water radiator. Combines boiler and constant speed pump with hot water baseboard heater, and dual set point thermostat.
4. Air-source heat pump with electric backup. Combines single-speed air source heat pump DX heating and cooling coils with constant volume fan, electric resistance heating coil, and dual set point thermostat.
5. Water-source heat pump. Combines single-speed water-source heat pump DX heating and cooling coils with constant volume fan, constant speed pump, and dual set point thermostat.
6. Central air conditioner. Combines single-speed DX air cooling coil with constant volume fan and dual set point thermostat.
7. Window/wall air conditioner. Combines single-speed DX air cooling coil with constant volume fan and dual set point thermostat.

3.1.6 Schedules

Individual component models rely on schedules to determine operation (thermostat set point), usage (washer, dryer, etc) or occupancy. Because schedules can vary from hour to hour, each hour of

the year, a flexible scheduling model was created that relies on a compact serialized representation as input. The format borrows from the EnergyPlus input format. For example:

```
06/30:Mondays Tuesdays Wednesdays Thursdays Fridays:6@0|18@1|24@0;
06/30:Saturdays AllOtherDays:10@0|16@1|24@0;
12/31:AllDays:24@1
```

defines the following schedule:

1. **From** January 1 **Through** June 30:

For Mondays, Tuesdays, Wednesdays, Thursdays, and Fridays:

Scheduled value is 0 from midnight until 6am, 1 from 6am until 6pm, and 0 from 6pm until midnight.

2. **From** January 1 **Through** June 30:

For Saturdays and all other days that have not yet been specified:

Scheduled value is 0 from midnight until 10am, 1 from 10am until 4pm, 0 from 4pm until midnight.

3. **From** July 1 **Through** December 31:

For all days:

Scheduled value is 1 from midnight until midnight.

When a value is specified “at” an hour, it is assumed that the value applies for all hours since the previously specified hour, “until” the specified hour. The same applies for dates except that instead of “until” the specified date, the schedule is read as “through” the specified date. Hours are indexed starting with 1; literally, 4 means the “fourth hour of the day”, not 4am as implied by the interpretation.

Supported day types include: Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, Sunday, Weekdays, Weekends, AllDays and AllOtherDays. AllOtherDays is a catch-all for any day

type that has not been specified. In the example above, AllOtherDays ensures that a schedule is defined for Sunday.

A schedule is agnostic to its ultimate usage. A schedule may contain binary values (0,1), temperature values (32-100), or fractional values (0.0-1.0). The interpretation of the schedule is left to the model using the schedule. For example, many equipment schedules are expressed as probability density functions (PDFs) represented as discrete hourly probabilities whose values when summed equal 1.0¹ for a 24-hour period. In the case of lights, consumption is determined by schedule, wattage, and the number of hours used during a day. Expressing the schedule as a PDF means that the number of hours and wattage to be simply multiplied and distributed across the day according to the density function, allowing the same schedule to be used regardless of the number of hours of usage per day.

3.1.7 Internal Gains

Energy consumption and internal gains from equipment such as appliances and lights can be modeled using nominal energy demand (in watts), schedule, fuel type, and sensible heat fraction. As noted above, the schedule value may also be a fraction representing the probability that the equipment is on during a given hour. The energy consumption of equipment is calculated as:

$$E_i = E_{nom} S_i \tag{3.11}$$

where:

E_i is the energy consumed at time step i

E_{nom} is the nominal energy demand

S_i is the schedule value (fractional) at time step i

¹ There are of course special cases where one may require that the sum is more or less than 1.0. Weekends are such a case where usage is typically a factor more than weekday usage.

and the heat gain from the equipment is simply:

$$G_i = E_i F_{sensible} \quad (3.12)$$

where:

G_i is the heat gain at time step i

$F_{sensible}$ is the fraction of energy converted to heat

Heat gains from occupants are also modeled using Equations 3.11 and 3.12, assuming E_{nom} is the heat produced by one relaxed seated person (MET). Fuel type has little meaning in this context and is therefor not defined for occupants. Note that latent gains from equipment and occupants are not modeled.

3.1.8 Thermostat Model

The thermostat modeled in this work represents a simple dual-mode programmable single set point thermostat. The term “single” indicates that only one set point, either heating or cooling, may be active at any given time. As implemented, the thermostat may be “programmed” to have a different set point value at each hour of the year using a set point schedule.

The thermostat has a single hysteresis (set point deadband) value, defaulted to 0.5K, which is configurable. This results in the on/off cycling common in residential systems. Additionally, it implements minimum run time and coil/fan staging logic, both of which are obviously important to the sub-hourly electricity demand of the system. Minimum compressor run time forces the system to run for a minimum fixed time even if the target temperature has been achieved. In practice, this prevents short-cycling of the equipment in order to reduce wear and tear. The value is defaulted to 10 minutes.

Staging logic allows the heating and cooling equipment to be simulated as having multiple stages. The logic implemented is fairly simple. For a two stage cooling system, the thermostat

compares the current indoor temperature² to the target temperature (the set point temperature plus hysteresis). If the absolute difference is greater than a pre-defined threshold, both stages of cooling are engaged, otherwise, only one stage is engaged. This results in single-stage operation when the system is maintaining a fixed set point, and two-stage operation when the set point is changed. This threshold value is defaulted to 2K.

Two examples of cooling operation are shown in Figures 3.2 and 3.3, illustrating the three features of the thermostat. The thermostat deadband can be clearly seen in the upper panel as it oscillates around the cooling set point. Minimum run time can be observed in the early morning hours of the lower panels. Both continuous one-stage and intermittent two-stage operation can be observed in early afternoon in Figure 3.2; Figure 3.3 shows two-stage operation only when the set point is stepped down in the late afternoon.




Figure 3.2: Example output from reduced order model showing HVAC cycling, thermostat hysteresis, staging and minimum run time. Low cooling set point results in cycling early in the morning and frequent second cooling stage operation in the afternoon.

The thermostat parameter values cited above are derived from a limited survey of HVAC

² This is actually the anticipated indoor temperature for the current time step after the free float calculations are performed. Residential thermostats capable of controlling multistage systems use an anticipated internal temperature concept as well. This results in two-stage operation when loads are high, i.e. during the peak afternoon hours.




Figure 3.3: Example output from reduced order model showing HVAC cycling, thermostat hysteresis, staging and minimum run time. High cooling set point delays cooling operation until mid-morning and avoids second cooling stage operation in the afternoon.

equipment specifications [22, 48], and from the author’s professional experience reviewing embedded thermostat code.

3.1.9 Simulation Flow

The simulation process is similar in many ways to that of whole-building simulation packages used throughout industry. A small number of code optimizations (not to be confused with mathematical optimization) have been made which improve performance within the model predictive control environment described previously. The general simulation process and optimizations are as follows:

1. Building geometry is constructed from input parameters that define the aspect ratio of the building, total floor area, number of levels, size and distribution of glazing, and overhangs.
2. Solar angles, incident solar radiation and sunlit fractions are pre-calculated from hourly weather data. Once initialized, the model does not re-calculate these values during simulation. For a single annual run, this step does not introduce any computational efficiency gains, however, when the model is re-simulated during optimization, significant computational expense is avoided.
3. Internal gains from equipment, lights and occupants are pre-calculated. As with solar angles and sunlit fractions, these values are not re-calculated during simulation unless they are changed via the simulation process.
4. The building is warmed-up by repeatedly simulating the first day until the mass and air temperatures converge. The model assumes infinite equipment capacity during this step.
5. The HVAC equipment is auto-sized. The model is simulated for an entire year assuming infinite equipment capacity, the largest heating and cooling loads are determined and scaled by a configurable fraction, and the sizing of individual components performed. For heating and cooling coils, boilers and baseboards, these scaled values become the nominal capacity.

For fans and pumps, air and water flow are calculated assuming nominal temperature differentials across down-stream components. Auto-sizing is optional.

6. For each time step of the simulation, building state is determined. First, the zone free-float temperature is found by simulating without HVAC equipment. If the zone temperature exceeds the cooling set point or drops below the heating set point, the energy required to bring the zone back to temperature is calculated. Because the capacity of the HVAC equipment may vary depending on ambient conditions, the total available heating and cooling capacity is determined for the current time step. If this capacity is less than required to achieve or maintain the set point temperature, new mass and zone temperatures are calculated.
7. Fuel consumption by the HVAC equipment is calculated given the delivered heating and cooling energy from each time step.
8. At the end of the simulation, the building model can be written to file in a serialized format, i.e. a format that allows the model to be reconstructed with exactly the same state. This allows the model to be loaded and re-run without any of the overhead of initialization, auto-sizing or warm-up. Because the state of the building is saved for each time step, a subsequent simulation may start at any arbitrary time in the correct state.

Note that the simulation of zone and HVAC are tightly coupled, reflecting the available heating and cooling capacity at each time step. This is particularly important when control strategies are employed that exercise the limits of the heating and cooling system as is common during optimization. Also note that the model is capable of sub-hourly simulation time steps down to one second.

3.1.10 Annual Validation of Building Model

BESTEST-EX [50] is a simulation model testing standard derived from ASHRAE/ANSI 140 (BESTEST) [49], specifically created for testing the accuracy of residential building modeling

software for existing buildings. The testing procedure consists of two parts. The first includes a suite of building physics test cases meant to validate the candidate software's ability to correctly capture building heating and cooling load calculations on an annual basis for a variety of different building configurations. The second includes a suite of tests meant to validate automatic building calibration methods. In these tests, a number of building envelope characteristics are allowed to vary within a given range. Calibration software should be able to match the results from a set of annual simulations in which the actual envelope characteristics are not disclosed. It can be argued that the second part of the procedure indirectly tests the ability of the candidate model to simulate different building characteristics. Only the first part of the testing procedure, i.e. the set of physics test cases, is performed for the reduced-order model. The procedure requires that both heating and cooling dominated climates are simulated, represented by Boulder, CO and Las Vegas, NV.

The base case reduced order model is constructed using the geometry and material characteristics documented in the testing procedure. The building is assumed to be 1,539 sq. ft with a single conditioned zone, a raised floor exposed to air and an unconditioned attic. The building is 57 ft. wide by 27 ft deep, with the front facing due south. Walls are assumed to be uninsulated, the raised floor is insulated with R-19 batt, and the ceiling floor is insulated with R-11 batt. Window to wall ratio is approximately 26% and windows are assumed to be single paned with a wood frame. Heating and cooling capacity is assumed to be infinite, with efficiency and COP fixed at values of 70% and 3.0, respectively. Further details regarding the modeling assumptions can be found in the documentation.

Nine additional cases are constructed according to the specifics defined in the procedure. These test cases include: improved air sealing; improved attic insulation; improved wall insulation; thermostat set-back; improved windows; higher envelope solar absorptivity; lower envelope solar absorptivity; shading from overhangs; a case combining all measures.

For validation, rough calibration of the reduced-order model is performed by hand until good agreement between the base case annual simulation and benchmark simulation is achieved. It has been found that only the internal mass to floor area ratio and internal to external envelope capac-

itance ratio need be adjusted to arrive at the calibrated results presented below. Once calibrated, the entire suite of test cases is simulated. The difference between the calibrated and original internal mass to floor area ratio (calculated from the detailed building envelope specifications in the testing procedure) is on the order of 10%. The internal to external envelope capacitance ratio was found to be 1:4.

Figure 3.4 shows an annual comparison of cooling loads from three benchmark software simulations (EnergyPlus, SUNREL, DOE2.1E) and the reduced order model (3R-2C) for all test cases. As seen in the upper panel, agreement between the reduced-order model and the benchmark is quite good. The lower panel shows the difference in load from each test case to the base case. A separate analysis not included here confirms that the results are within the acceptable error proscribed by the testing procedure.

Figure 3.5 shows an annual comparison between heating loads for the same cases in the heating climate. Again, agreement between the benchmarks and the reduced-order model is acceptable.




Figure 3.4: Comparison of annual electricity use in cooling physics test cases for EnergyPlus, SUNREL, DOE2.1E, and the reduced-order model.

Although not specifically required by the testing procedure, the results of the reduced-order model can be compared with the benchmark software on a monthly basis to validate seasonal




Figure 3.5: Comparison of annual gas use in heating physics test cases for EnergyPlus, SUNREL, DOE2.1E, and the reduced-order model.

variations in energy consumption. Figure 3.6 demonstrates that the reduced-order model is capable of capturing these variations as well as the benchmarks.




Figure 3.6: Comparison of monthly electricity use in cooling and heating physics test cases for EnergyPlus, SUNREL, DOE2.1E, and the reduced-order model.

3.1.11 Sub-Hourly Validation of Building Model

Using the hand-calibrated reduced-order model from the annual validation above, and the BESTEST-EX supplied EnergyPlus model, a heating and cooling load comparison is performed to validate that the sub-hourly dynamics of the two models are in agreement. The comparison is performed for the same two locations, representing both heating and cooling dominated climates. Time series results are shown in the figures below.

Figure 3.7 shows total heating energy for both the reduced order model (RC) and the EnergyPlus model at a 5-minute time step. Only those days in which heating is required are plotted. Consistent with the annual validation, the reduced order model tends to under-predict heating loads compared to EnergyPlus. This is most noticeable when loads are highest. Nevertheless, the shape closely matches, and trends are consistent. Normalized root mean square error (NRMSE) between the two simulations for the entire time series (excluding days where heating is disabled) is

7.4%.

Figure 3.8 shows a similar comparison for cooling load. Here, the bias is towards over-predicting cooling energy, but NRMSE (for all days where cooling is enabled) is smaller at 5.5%. Where models differ, loads tend to be lower.

Despite the differences, visual comparison and NRMSE values suggest that the reduced order model sufficiently captures the dynamics of both heating and cooling loads with a bias towards being too ‘warm’. Further calibration would likely improve the result, but is not necessary for the purpose of this work as the NRMSE value is already very good.

The response of the model to changes in thermostat set point is important to the investigations studied in this work. To validate the model under these conditions, both the reduced order model and the EnergyPlus model are subjected to a +3K cooling set point step change between 13:00 and 17:00 during each day of the cooling season. Two variations of the model are tested, representing high and low internal mass to floor area ratios. These mass levels represent the range of values observed in the population sampling process outlined in 3.2.2, and correspond to 2.5 and 4.0 times building floor area.

Results from a single day, July 1, are shown in Figure 3.9. This chart shows cooling load for both models at two mass levels at each five-minute simulation time step. In general, the reduced order model follows the EnergyPlus model very well. The reduced order model tends to lag the EnergyPlus model following the temperature set-up at 13:00, and shows a slightly higher peak when the set point is returned to the previous value at 17:00.

Normalized root mean square error in cooling load for the entire cooling season is shown in Table 3.1. Values of 5.3% and 5.1% indicate very good agreement between the models for a given mass level. From Figure 3.9, it appears that the reduced order model shows less sensitivity to the mass level than the EnergyPlus model. However, a comparison between high and low mass values for a given model shows that over the cooling season, both the reduced order model and EnergyPlus show the same sensitivity. At 3.1% and 3.5%, neither model shows much sensitivity to internal mass level.




Figure 3.7: Heating load profile calculated by reduced order model compared to EnergyPlus, October through May.




Figure 3.8: Cooling load profile calculated by reduced order model compared to EnergyPlus, May through October.




Figure 3.9: Cooling load profile calculated by reduced order model compared to EnergyPlus at two internal mass levels with step change in cooling set point.

Table 3.1: Normalized root mean squared error between reduced order model and EnergyPlus at two mass levels with step change in cooling set point, May through October.

	EnergyPlus Low Mass	EnergyPlus High Mass	RC High Mass
RC Low Mass	5.3%	NA	3.1%
RC High Mass	NA	5.1%	NA
EnergyPlus Low Mass	NA	3.5%	NA

3.2 Electric Grid Model

This work considers the electricity generation, transmission and distribution system from the perspective of a distribution feeder serving commercial, agricultural, industrial and residential loads. Feeder models are adapted from the PNNL Modern Grid Initiative Distribution Taxonomy Report [82]. This report defines 24 prototypical feeders common throughout different climate regions in the United States. The following section describes the selection and characteristics of the three feeders considered in this research, and the modifications made to the feeders in order to integrate with the simulation framework. The section concludes with a validation of the modified feeder models.

3.2.1 Selection of Climate and Feeder Models

The Distribution Taxonomy Report (DTR) classifies feeders according to both ASHRAE climate region and physical characteristics, using a cluster analysis for each of the climates zones defined by [15] and [14]. This classification corresponds with that used in the location selection exercise documented in ASHRAE Research Project 1313-RP [41]. The research presented in the present work relates in many ways to that undertaken in ASHRAE 1313-RP, therefore, the list of initial candidate locations was adopted from that source. From the list of 15 cities in ASHRAE 1313-RP, three cities were identified that 1) represented three different regions according to the feeder taxonomy report and 2) provided a wide range of cooling degree days. These cities are: Houston, TX; Los Angeles, CA; New York, NY. A summary of relevant climate design data from the 2009 ASHRAE Fundamentals Handbook [6] is shown in Table 3.2.

Table 3.2: Climate characteristics of cities selected for study.

City	ASHRAE Climate Zone	CDD50	CDD65
Houston	5	4043	1667
Los Angeles	3	2674	343
New York	2	1911	543

The 24 feeder models were then evaluated based on their physical characteristics. As this

work is concerned primarily with residential cooling demand, feeders were evaluated first on the number of residential buildings served by the feeder and second by the percentage of annual load represented by the residential buildings. Consideration was also given to the topology of the feeder: those typical of urban and suburban distribution feeders were preferred over those typical of rural or lightly populated areas. Three feeders were selected corresponding to the three climate zone and location pairs. The following descriptions are taken directly from the distribution taxonomy report. The number of residences was determined by counting the number of homes defined in the original feeder model. The annual percentage of residential electricity consumption was calculated from a full annual feeder simulation by measuring the total number of megawatt hours reported by the residential meters and comparing to the total number of megawatt hours reported by the distribution transformer.

It should be noted that the description in the report does not always align exactly with the data in the table. For example, Feeder 8 is described as having light industrial loads, however there are no industrial transformers in the feeder model. Additionally, the nominal load for this feeder is well below the load measured in the original feeder simulation. Communication with the development team at PNNL confirms that the nominal feeder load should be closer to 7.4MW and that the report is in error.

3.2.1.1 Feeder 8, New York

Feeder 8 is described as “a representation of a lightly populated suburban area . . . composed of single family homes, light commercial loads, light industrial loads, and some agricultural loads”. The description also notes that due to connections with adjacent feeders, loading of this feeder should be limited to 60% or 4.4MW.

Table 3.3: Characteristics of Feeder 8.

Nominal voltage (kV)	12.47
Nominal load (MW)	7.4
Commercial transformers	6
Industrial transformers	0
Agricultural transformers	0
Residential transformers	396
Number of residences	1506
Percent of residential consumption	86%

3.2.1.2 Feeder 13, Los Angeles

Feeder 13 is described as “a representation of a heavily populated suburban area . . . composed mainly of single family homes with some light agricultural loads”. Feeder loading should be limited to 75% or 5.9MW for similar reasons.

Table 3.4: Characteristics of Feeder 13.

Nominal voltage (kV)	12.47
Nominal load (MW)	7.8
Commercial transformers	0
Industrial transformers	0
Agricultural transformers	107
Residential transformers	1491
Number of residences	1326
Percent of residential consumption	78%

3.2.1.3 Feeder 22, Houston

Feeder 22 is described as “a representation of a heavily populated suburban area with a moderately populated urban area . . . composed mainly of single family residences with some moderate commercial loads”. Feeder loading in this case should be limited to 66% or 7.9MW.

Table 3.5: Characteristics of Feeder 22.

Nominal voltage (kV)	22.9
Nominal load (MW)	12.0
Commercial transformers	14
Industrial transformers	0
Agricultural transformers	0
Residential transformers	284
Number of residences	2146
Percent of residential consumption	80%

3.2.2 Populating Feeder Model with Residences

The GridLAB-D development team provides MATLAB scripts to automatically generate a population of residential buildings for each distribution feeder model based on feeder nominal load characteristics and climate region. The method uses Residential Energy Consumption Survey [5] data collected and made available by the US Energy Information Administration. The scripts randomly sample RECS data across a number of characteristics which include floor area, dwelling type, heating type, cooling type, heating set point, cooling set point, hot water usage, envelope integrity, and other characteristics important to residential energy use. Distributions of several major characteristics are shown in Figures 3.10, 3.11, and 3.12. Note that homes are grouped according to the electrical phase to which they are attached.

After sampling, the scripts generate a file containing the residential building definitions in GLD input format for each building. This file is then transformed into the input that GridMPC requires using a simple search and replace script. This allows GridMPC to simulate the residential buildings and insert the results into the power flow simulation. Section 3.2.6 describes briefly how this is accomplished.




Figure 3.10: Distribution of home characteristics for Houston feeder model.




Figure 3.11: Distribution of home characteristics for Los Angeles feeder model.



Figure 3.12: Distribution of home characteristics for New York feeder model.

3.2.3 Modeling Home as ZIP Load

Building energy modeling from an energy engineering perspective is typically not concerned with the behavior of individual electric devices on the electric system as a whole. Often, an electrical load is simply modeled as a time varying demand, without regard to the nature of the load. That is, all loads regardless of their inductive, capacitive or resistive characteristics are treated the same. However, in electric system modeling, these are the characteristics which define the behavior of the system. Therefore, power systems engineers will model individual devices in a way that accurately characterizes their behavior from an electrical engineering perspective. As with energy models, these models can range from very simple to very complex. The ZIP load model implemented in GridLAB-D (Equations 3.13 and 3.14) is relatively simple, but sufficiently accurate to represent the electric load from this perspective.

$$P = \left| \frac{V_a}{V_n} \right|^2 |S_n| Z_{\%} \cos(Z_{\theta}) + \left| \frac{V_a}{V_n} \right| |S_n| I_{\%} \cos(I_{\theta}) + |S_n| P_{\%} \cos(P_{\theta}) \quad (3.13)$$

$$Q = \left| \frac{V_a}{V_n} \right|^2 |S_n| Z_{\%} \sin(Z_{\theta}) + \left| \frac{V_a}{V_n} \right| |S_n| I_{\%} \sin(I_{\theta}) + |S_n| P_{\%} \sin(P_{\theta}) \quad (3.14)$$

where:

P is the real power

Q is the reactive power

V_a is actual voltage

V_n is nominal voltage

S_n is apparent power at the nominal voltage

$Z_{\%}$ is the percent of the load that is constant impedance

Z_{θ} is the phase angle of the constant impedance fraction

$I_{\%}$ is the percent of the load that is constant current

I_{θ} is the phase angle of the constant current fraction

$P_{\%}$ is the percent of the load that is constant power

P_{θ} is the phase angle of the constant power fraction

A ZIP load model is defined by constant impedance ($Z_{\%}$), current ($I_{\%}$) and power ($P_{\%}$) fractions. The values of these fractions are specific to the type of load being simulated. In a ZIP model, one would specify electric baseboard heat (high constant impedance fraction) differently from compact fluorescent lights (high constant power fraction). Thus, for every electric device in the home, a different ZIP model would be specified.

While it is possible to simulate all devices separately in this research, and indeed possible within the GridLAB-D framework, it is unwieldy in the context of the optimization framework developed for this work. Instead, one seeks to reduce the complexity of the electric model while still maintaining the overall characteristics of the electric load, aggregated at the whole-home level. In this work, each whole home is modeled as a single ZIP load with representative (average) time-invariant ZIP fractions. This approach requires several simplifying assumptions to be adopted:

1. ZIP fractions are fixed for the home. At the whole home level, ZIP fractions vary from one instant to another as various electric devices turn on and off. Modeling the home as a single ZIP load with fixed fractions neglects the effects on the system caused by the mix of loads present at any instant.
2. ZIP fractions are representative of the average home. Across the population homes, the presence (or absence) of hot water heaters, air conditioners, heat pumps, and the specific mix of electric devices result in different whole-home ZIP characteristics. Deriving a set of average ZIP fractions for all homes neglects these effects.
3. Voltage dependence can be ignored. The ZIP load model does not exhibit any voltage dependence. That is, as the voltage on the distribution system changes over time (as a result of the load on the distribution feeder), the load behavior does not. For example, in a voltage-dependent model a hot water heater cycles more quickly at higher voltage, and less quickly at lower voltage. In a model without voltage dependence, the cycling behavior is unaffected.

Regarding assumptions 1 and 2 above, while these simplifications affect the absolute value of real power demand on the distribution feeder from one time step to the next, the observed variation measured at the feeder transformer is within a relatively small range of absolute values. Section 3.2.4 addresses this aspect in more detail.

Assumption 3 is somewhat more difficult to justify considering that the voltage supplied to the homes can affect the cycling behavior of the HVAC equipment. However, the sensitivity of HVAC cycling to this variation would be relatively small compared with the sensitivity to building envelope characteristics and operational constraints enforced by the thermostat, i.e. minimum cycle time and set point hysteresis. Further, because the effect of voltage variation on feeder demand is much smaller than the expected magnitude of change induced by load manipulation, e.g. under optimal control, voltage dependence can be safely neglected.

This research investigates the manipulation of HVAC electric demand which is principally dominated by AC electric motors. It should be noted that a more appropriate motor model for use in power system simulation, such as the WECC composite model [54], is available in GridLAB-D version 3.0. Future investigations in which short term dynamics of the system are of interest may benefit from improved motor modeling.

3.2.4 Selection of ZIP Fractions

To arrive at a representative set of average ZIP fractions, each of the three feeder models is simulated for an entire year in GridLAB-D, using the TMY2 weather selected for the feeder. Simulations are configured to output the sum of real and reactive power for the constant impedance, current and power portions of the load at each time step. Note that these values are not the ZIP fractions, rather, these correspond to each of the six terms added together to equal the real and reactive components of the ZIP model. These six terms are shown in Equations 3.15 through 3.20 below.

$$P_P = |S_n| P_{\%} \cos(P_{\theta}) \quad (3.15)$$

$$P_I = \left| \frac{V_a}{V_n} \right| |S_n| I_{\%} \cos(I_{\theta}) \quad (3.16)$$

$$P_Z = \left| \frac{V_a}{V_n} \right|^2 |S_n| Z_{\%} \cos(Z_{\theta}) \quad (3.17)$$

$$Q_P = |S_n| P_{\%} \sin(P_{\theta}) \quad (3.18)$$

$$Q_I = \left| \frac{V_a}{V_n} \right| |S_n| I_{\%} \sin(I_{\theta}) \quad (3.19)$$

$$Q_Z = \left| \frac{V_a}{V_n} \right|^2 |S_n| Z_{\%} \sin(Z_{\theta}) \quad (3.20)$$

The values obtained from the simulations are then averaged over the year, yielding a single number for P_P , P_I , P_Z , Q_P , Q_I and Q_Z . From the averaged values, the impedance, current and power phase angles are calculated simply with trigonometry. In the ZIP model, the three ZIP fractions must sum to one (Equation 3.21).

$$Z_{\%} + I_{\%} + P_{\%} = 1 \quad (3.21)$$

Introducing this relationship allows the equations to be rearranged to eliminate the unknown V_a , V_n and S_n terms. The resulting Equations, 3.22 and 3.23, can then be solved iteratively to arrive at values for $Z_{\%}$, $I_{\%}$ and $P_{\%}$, representing the time and population averaged ZIP fractions specific to the feeder.

$$F = \frac{P_P Q_P \sin(I_{\theta}) \cos(I_{\theta})}{P_I Q_I \sin(P_{\theta}) \cos(P_{\theta})} \quad (3.22)$$

$$P_{\%}^2 - P_{\%} + P_{\%} I_{\%} + F I_{\%}^2 = 0 \quad (3.23)$$

ZIP fractions calculated for the three feeders are shown in Table 3.6 below. From the perspective of the feeder, homes appear on average as having 2/3 constant power and 1/3 constant impedance behavior. The absence of constant current behavior is interesting given the presence of a constant current component in the fractions specified in the underlying GridLAB-D load models.

Table 3.6: ZIP fractions for three feeder models.

Location	$Z_{\%}$	$I_{\%}$	$P_{\%}$
Houston	0.35	0.04	0.61
Los Angeles	0.25	0.05	0.70
New York	0.29	0.10	0.61

3.2.5 Feeder Sensitivity to ZIP Fractions

As alluded to previously, the simplification introduced by time and population averaging ZIP fractions is expected to affect the overall real power demand on the distribution feeder, but the extent of the error introduced by this simplification is unknown. A set of simple experiments is devised to test the sensitivity of real power demand to the ZIP fractions, and estimate the size of the error. To accomplish this, each home within the feeder is replaced with an equivalent ZIP load whose electricity demand at each time step matches the electricity demand of the home it replaces. Three simulations are performed in which the entire population of homes is given the same set of ZIP fractions. The upper panel of Figure 3.13 shows the real demand on the Houston feeder for a single week in July for each of the three cases. The three cases are given by the ZIP combinations in Table 3.7. A fourth case labeled ‘Average’ is the result of a simulation using the average ZIP fractions found previously for the Houston feeder; this series completely obscures the $I_{\%} = 0$ series. The lower panel of the figure shows the demand percentage error relative to the average demand of the $Z = 1$, $I = 1$, and $P = 1$ cases.

Table 3.7: ZIP fractions and case names for sensitivity test.

Case Name	$Z_{\%}$	$I_{\%}$	$P_{\%}$
Z=1	1	0	0
I=1	0	1	0
P=1	0	0	1

While error can be as large as $\pm 4\%$, error at high demand is roughly half this value. The constant current case demand is centered around the average of the three cases and shows behavior very similar to the averaged ZIP fraction simulation. This suggests that this case could instead




Figure 3.13: Sensitivity of Feeder Demand to ZIP Fractions.

be simulated as a combination of only Z and P fractions. The bias of the averaged ZIP fraction simulation towards the constant power simulation is consistent with the larger $P_{\%}$ value found. From these observations, the following conclusions can be drawn:

1. The $Z_{\%} = 1$ and $P_{\%} = 1$ simulations represent the upper and lower boundaries of real power demand for any arbitrary combinations of ZIP fractions. Were the fractions varied each time step for each home, the result would fall between these boundaries.
2. The feeder exhibits relatively little sensitivity to the ZIP fractions in terms of real power demand. The error introduced by population averaged time-invariant ZIP fractions is relatively small and is acceptable for the purposes of this work.

3.2.6 Development of Hybrid Model

Having simplified each home in GridLAB-D into an equivalent ZIP load model, the homes within the feeder can be simulated simply as an electric demand calculated outside of GridLAB-D by the reduced order building model. This is accomplished by first writing the electric demand of each home at each time step of the simulation into a separate file, then linking these files to the power flow simulation using the `ZIPload` and `player` objects. The new model that results, which combines the loads calculated by GridMPC with the GridLAB-D feeder model, is termed the ‘hybrid model’.

This process is performed programmatically when GridMPC executes as it involves thousands of files. It should be noted that the use of `ZIPload` incurs some unnecessary overhead because it is a child of the `house` object which is not needed in the hybrid model. Alternatively, the `triplex_load` object could be used for this purpose, but would necessitate a significant change to the structure of the original GridLAB-D input file to remove all `house` and `triplex_meter` objects to which they attach.

3.2.7 Validation of Hybrid Model

The hybrid feeder models are validated against the reference feeder models to ensure that the aggregate load characteristics are maintained. Validation is performed at two time-scales: monthly validation of cumulative energy consumption ensures seasonal effects on feeder demand are correctly captured; daily validation of instantaneous feeder demand for a peak week ensures that the nominal design load is not exceeded. The latter is particularly important as exceedance of the system limits can result in an overloading of the feeder and subsequent voltage collapse. The hybrid model must also exhibit a demand shape similar to the reference to ensure that the aggregate behavior correctly reflects daily temperature and solar related effects. Too high or low a demand would suggest significant differences in the two building models, possibly requiring homes to be removed from the feeder, or additional reduced-order models to be added.

Figure 3.14 shows cumulative monthly electricity consumption for reference and hybrid versions of the three feeder models. All three hybrid models show reasonable agreement³ with the reference versions. Seasonal trends in the hybrid model are consistent with the reference in all cases.

Table 3.8 summarizes the percentage error between reference and hybrid models. The Houston models appear to be much less consistent than the others, but annual error is only -5.0%. Errors for Los Angeles and New York are 0.2% and 0.1%, respectively. On a monthly basis, the largest error observed is -9.5% for the Houston feeder, and lower for the others.

Cooling loads in Houston are relatively large compared to the other locations so discrepancies between models will be magnified. Recall that the error observed in the reduced order model validation suggests that cooling error increases as absolute cooling load decreases, so it does not appear that the discrepancy between the original and hybrid model is a manifestation of this.

³ Discrepancies between the residential loads calculated by GridLAB-D and GridMPC were identified early in the testing stage that uncovered two rather significant bugs in GridLAB-D. The bugs relate to handling of daylight saving time and interpolation of weather data. This resulted in a two hour shift in the expected results. A work-around for both issues exists by providing the weather data to GridLAB-D as a CSV file that has been correctly shifted for DST and interpolated.




Figure 3.14: Monthly validation of hybrid GridMPC model.

Table 3.8: Percentage error in total electricity consumption for three feeders.

Feeder	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Total
Houston	-9.0	-8.3	-2.0	9.1	-2.4	-6.7	-9.1	-8.1	-3.6	-4.2	-1.8	-9.5	-5.0
Los Angeles	-3.2	-2.6	-1.0	4.7	3.8	3.7	2.5	2.1	2.7	-0.8	-0.8	-2.1	0.2
New York	-6.3	-6.0	-4.2	6.9	7.5	5.6	1.7	4.3	6.5	0.0	-4.8	-5.6	0.1

Rather, it appears that the reduced order building model tends to underpredict cooling electricity compared to the GridLAB-D building model. As the two are quite different in their assumptions regarding the distribution of solar gains and complexity of envelope elements, this is not terribly surprising.

Figure 3.15 illustrates the difference for a single week in July, one of the hottest weeks of the year. Again, the hybrid model underpredicts electricity consumption relative to the original. Despite what appears to be a significant difference the error is acceptable. This validation is meant to show that the hybrid model does not significantly exceed the nominal feeder load. In fact, it tends to be much more conservative than the original in this regard. Recall that the nominal design load of the Houston feeder is 12.0MW and that it is designed to operate at roughly 66% of this value. While both hybrid and original models exceed 7.9MW during this week, the original exceeds 12.5MW on several occasions, while the hybrid never exceeds 10.1MW.




Figure 3.15: Feeder demand validation for Houston feeder model.

Figure 3.16 shows a similar comparison for the Los Angeles feeder. Here, agreement is much better than in the Houston case. This is not unexpected since the ratio of residential cooling electric demand to total feeder electric demand is lower for this location. With respect to feeder loading, neither the original or hybrid model exceed the nominal feeder load of 7.8MW, nor do they exceed the 75% limit of 5.9MW. Although none of these days show the annual peak feeder demand, it will be shown in Chapter 5 that during normal operation these values are not exceeded.




Figure 3.16: Feeder demand validation for Los Angeles feeder model.

Similar to the Los Angeles observations, New York hybrid feeder demand shown in Figure 3.17 closely matches that of the original. Both models far exceed the nominal value documented in the Distribution Taxonomy Report, even in the middle of the night when cooling loads are lowest. While it is fairly clear that the hybrid model does a very good job approximating the original, it is hard to make any claims about feeder loading.

Agreement between the original and hybrid feeder models could be improved through the

addition of residences. Certainly, the hybrid model should not produce electric demands exceeding the original. Similarly, the hybrid model should not underpredict load by a large percentage. Although it is difficult to put an exact number on a lower threshold, this analysis confirms that the former criterion is satisfied. The determination of the latter is somewhat more subjective, but deemed satisfactory for this research.




Figure 3.17: Feeder demand validation for New York feeder model.

Chapter 4

Model Predictive Control and Optimization

The following chapter describes the model predictive control methodology and optimization algorithm used throughout this research. This process is informed in part by the author's previous work which has been published in [28]. As noted earlier, the software implementing this approach has been written anew, but shares some common features with the predecessor. Those features common to both receive little additional treatment here, and the reader is referred to the aforementioned publication for full details.

4.1 Methodology

This work implements the classic receding horizon model predictive control scheme to minimize an objective which is a function of building electricity demand. This is accomplished by manipulating the cooling set point schedules of thousands of simulated residential buildings in parallel. The size of the decision space makes a supervisory control optimization intractable, requiring that a decentralized, distributed approach be adopted. Therefore, each residential building is subject to a separate receding horizon control optimization independent of the others. The formulation of the objective function becomes critical to the aggregate response as will be illustrated in the experiments.

For all cases studied here, the prediction horizon chosen is twenty-four hours following the findings of [42] for processes with diurnal cycles. The execution horizon has no meaning in these studies as each day is considered in isolation. However, the software is configurable to arbitrary

planning and execution horizon lengths.

In contrast to earlier studies using EnergyPlus as the function evaluator, no initialization or termination horizons are required. The former is not necessary as state is maintained by the software automatically and does not need to be reinitialized. The exclusion of a termination horizon is justified by the relatively small thermal storage capacity of a residential building compared to cases studied previously.

As with previous work, the function evaluator appears as a black-box from the perspective of the optimization algorithm. This requires a metaheuristic approach. The algorithm selected, described in further detail in Section 4.3, is appropriate for implementations of this type.

4.2 Optimization Process

The general process of optimization is very similar to the offline process documented in [28]. In summary, for each building and each planning horizon, a separate optimization is performed according to the following steps:

1. Generate candidate control vector.
2. Simulate building with candidate control vector.
3. Evaluate fitness and exit criteria; return to step 1 or exit.

The first iteration of this process is performed with a simulation “seed” which is unique to each building. The seed corresponds to the upper boundary of the cooling set point schedule. Recall that the cooling set point for a given home is drawn from a distribution according to US Census data. The cooling set point schedule is assumed to be constant for the base case, as previous studies suggest that many homes do not have a programmable thermostat or do not program their thermostat with set-backs [79].

In the optimization case, the upper boundary varies by time of day. Specifically, the set point is raised relative to the normal cooling set point when the home is unoccupied. Homes are assumed

to be unoccupied between the hours of 8am and 6pm.¹ The upper boundary may be specified as an absolute value, e.g. 28°C, or as an offset, e.g. 3K. Throughout this research, the upper boundary is specified as an offset. The lower boundary is specified in a similar manner. The offset values are specified in subsequent chapters.

In previous research, decision variables were divided into “modes” to limit the decision space to a reasonable size. This was a necessary side-effect of the choice to use a relatively slow, but very detailed, building simulation engine. In this research, the simulation is greatly simplified due to model reduction, allowing an increased decision space granularity and number of modes. This allows the optimizer to consider forty-eight decision variables corresponding to each thirty minute period of the day.

The decision to use thirty minute modes is not arbitrary, rather it is a balance between competing needs. The ability of the system (the distribution feeder in this case) to quickly change demand is a function of the mode length. Shorter mode lengths allow for faster response. Ideally, modes would be as short as possible. However, changing thermostat set points more frequently is not practical. First, the minimum run time requirement enforced by the thermostat makes shorter modes pointless. Second, because of thermal storage effects, changes in set point faster than the system time constant will have little effect when cooling equipment capacity is discrete, i.e. in a two-stage air conditioning system. Although the time constants have not been explicitly calculated, Figures 3.2 and 3.3 suggest a value less than an hour, but greater than the ten minute minimum compressor run time. Thirty minutes is a logical compromise.

Mode boundaries for all buildings are aligned on the half-hour beginning at midnight. As will be discussed in Chapter 7, this can have interesting and unexpected implications.

¹ These are nominal values. The unoccupied hours vary up to ± 1 hour for each building

4.3 Optimization Algorithm

The Particle Swarm Optimization (PSO) algorithm has been selected based on promising results from previous work [28, 64]. PSO is a metaheuristic algorithm that behaves in a way that is analogous to the flocking behavior of social organisms. In PSO, a “swarm” of particles searches the decision space according to a combination of randomized and simple rule-based decisions, sharing information about the best solution found along the way.

Because PSO follows fairly simple rules, independent of the underlying objective function or its derivative, it is able to operate on black-box simulation models with decision spaces that contain discontinuities, making it particularly well suited for use with whole-building energy simulation.

The mathematical formulation of the optimization algorithm was first proposed and studied by [33] and subsequently refined for discrete and binary variables by [52]. The formulation used in this research can be expressed as:

$$v_{i,t} = \omega v_{i,t-1} + \gamma_1 \phi_1(p_{i,t-1} - x_{i,t-1}) + \gamma_2 \phi_2(g_{i,t-1} - x_{i,t-1}) \quad (4.1)$$

$$x_{i,t} = x_{i,t-1} + v_{i,t} \quad (4.2)$$

where:

i is the particle index

v is the velocity of the i th particle

x is the position of the i th particle

t is the current iteration

$t - 1$ is the previous iteration

ω is the inertial weighting term associated with velocity

ϕ_1 and ϕ_2 are the ‘learning’ terms that weight personal and global best positions

γ_1 and γ_2 are uniformly distributed random numbers in the range of 0 to 1

p is the personal best position

g is the global best position

Like the predecessor used in previous work, the variant used here contains a number of enhancements to the canonical formulation. Following the previous work, this variant allows the mixed use of discretized, continuous and binary decision variables via inputs defining the increment allowed for each dimension. It should be noted that the implementation has been subsequently revised following the findings of [64]; a description of the revision follows in Section 4.4.

Although this variant does not implement a neighborhood concept, it does allow for multiple seeds to be specified in contrast to the single seed per neighborhood of the predecessor. A taboo list feature is also implemented, as well as the familiar box constraints common to this type of MPC. The use of equation constraints has been eliminated simply because they are not utilized in this research.

4.4 Refinement of Search Space Mapping

The original implementation of the discretized PSO algorithm contained two very significant flaws. The first impacted the optimizer’s ability to handle binary decision variables. In that version, the position of the particle was calculated, rounded to the closest increment, and the particle was moved to the new location which must be an integer multiple of the discretization increment. For example, a particle given an upper boundary of ten, and a lower boundary of one with a increment of one, would only be allowed to move to the positions one, two...ten. For the binary case, the particle could only occupy the locations zero or one.

This is problematic because the velocities of particles may often not be large enough to allow the particle to escape its current position, as it will always be rounded back to the closest increment. Thus, particles become easily stuck and may not explore additional values. This is most obvious in the binary case where velocities would rarely exceed 0.5, but applies more generally to all discretized decision variables as the optimizer converges on a solution and velocities decrease. May-Ostendorp [64] solved this problem by employing a technique specifically for the binary case which uses a probabilistic approach according to [52] to determine the location of the particle position. While this approach solves the binary case, it does not address the more general case.

The second flaw involved the behavior of particles close to the boundaries of their box constraints. As particles approached the boundaries, their position would be fixed to the boundary value, and the velocity reversed. However, the change in direction occurred after the particle was fixed to the boundary and had no effect on position, meaning that particles that would have moved to or beyond the next discretized value would not do so in the current iteration. In combination with the rounding flaw described above, particles would tend to stick to boundaries for many iterations.

These two flaws have been subsequently addressed by two new features. First, the optimizer has been modified such that the underlying decision space remains continuous, i.e. all particle positions and velocities continue to have full double floating-point precision. Discretization, i.e. rounding to the nearest increment, now occurs during the objective function evaluation. This means that for any discretized decision variable, small velocities may continue to move the particle across increment boundaries. For example, a particle in a binary decision space has a velocity of 0.4 and a position of 0.2. In the previous version of the algorithm, 0.2 can not exist, so the particle would start in a location of 0.0, be given a new location of 0.4, but would be placed back at 0.0. In the new algorithm, the particle would start at 0.2 and move to a new position of 0.6, which would be evaluated by the objective function as 1.0. By treating the discretized space as continuous, particles no longer become stuck and unable to escape when velocities are small.

Second, the now continuous decision space has been mapped to an arbitrary range, say zero to 360, which loops back upon itself for values greater than 180 or less than zero. Effectively, the particles now travel along the edge of a circle, where the position represents an angle, with zero equal to the lower boundary of the box constraint, and 180 equal to the upper boundary of the box constraint. Values greater than 180, are evaluated as 360 minus 180, and values less than zero evaluated as their absolute value. By wrapping the decision space, large discontinuities in the objective function can be avoided, and velocities may remain undisturbed. An illustration of this concept is shown in the Figure 4.1.




Figure 4.1: Illustration of search space mapping between angle (inner) and decision variable (outer) values.

Chapter 5

Demand Response

Demand response (DR) in the residential context is very much still in its infancy, with deployments limited primarily to smart grid demonstration projects [78] and voluntary utility programs [7, 71]. While some examples [96] utilize remote on/off switches connected directly to air conditioner compressors at the home, current ‘state of the art’ residential peak demand limiting technologies use programmable communicating thermostats (PCTs). With PCTs, the utility may broadcast a message to a selected group of homes, resulting in the automatic adjustment of thermostat set points. In industry parlance, these are termed “demand response events”. This method of demand limiting results in a much more user-friendly experience compared to direct load control, as the thermostat still enforces that some degree of comfort be maintained.

DR events may be either absolute temperature events, i.e. all thermostats have their cooling set point reset to 28°C, or temperature delta events in which all thermostats have their cooling set point reset to the current schedule value +2K. Operation of residential DR programs is typically regulated by the public utilities commission (PUC) in the regulated markets or similar governing body in the deregulated, municipal and cooperative markets. Fewer than fifteen events per year are typically allowed, all corresponding roughly with days when the ratio of demand to generation capacity is at its highest. These may or may not be the days of the year that have the highest demand, as other factors such as scheduled maintenance of generation assets may affect this ratio.

The purpose of simulating a demand response event in this study is to establish a baseline for the maximum achievable reduction available using current practice. This baseline will then be

compared to subsequent cases in order to highlight the advantages and disadvantages afforded by adopting an MPC approach.

5.1 Methodology

Demand reduction potential is a function of both event duration and participation level. In practice, the effective duration and participation can vary from event to event due to participant opt-outs both before and during the event. This implies that a stochastic approach be adopted that requires a distribution of both parameters. The problem is simplified into four cases for each distribution feeder where both parameters are simulated at both high and low levels. This approach gives a sense of the range of possible responses without the overhead of a full stochastic simulation. Rates of 70% and 30% are selected as the high and low levels of participation; two and six hours provide a reasonable range for event durations.

Selection of the participants for a given rate still requires a stochastic process. Recall that the population characteristics vary according to the distributions shown in 3.2.2, therefore the aggregate response is expected to change depending on the participants. This is accomplished by simply sampling the population to create the group of participants for the specified participation rate. Two questions that arise with this approach are: 1) What effect does the population sample have on the expected savings? 2) How many times should the population be sampled to arrive at the true mean response? Both will be addressed in Section 5.3.

Obviously, the choice of day has a large impact on demand reduction potential. Here too, a large number of simulations is implied to test a range of conditions. As demand response events are concerned principally with peak demand days, the investigation is limited to a single day for each feeder. Event start and end times for the two durations are determined simply by identifying the range of hours with the largest average demand. The selection of the demand days and their characteristics follows.

5.2 Selection of Demand Day

Demand response events will typically be called on a day with high peak demand. Accordingly, a day for study is selected that contains the annual 15-minute peak demand, determined from an annual simulation of the modified feeder model using TMY weather. This is easily accomplished by taking a 15-minute simple moving average of simulated feeder demand and selecting the day with the largest value. Normalized annual feeder demand for each feeder can be seen in Figures A.1, A.2 and A.3.

5.2.1 Houston

For the Houston feeder, peak demand occurs on July 29 and has a 15-minute peak demand of 11.4MW, occurring between 15:35 and 15:50 CDT. The day is characterized by an average temperature of 29.5°C, and peak temperature of 35.6°C at 15:00 CDT. Feeder demand and coincident environmental conditions are shown in Figure 5.1 below.




Figure 5.1: Houston feeder peak demand day, July 29.

5.2.2 Los Angeles

For the Los Angeles feeder, peak demand occurs much later in the year on September 24. This day has a 15-minute peak demand of 4.1MW, occurring between 13:50 and 14:05 PDT. Average and peak temperatures for September 24 are 24.8°C, and 35.0°C, respectively. Similar to Houston, peak temperature leads peak demand at 14:00 PDT. Figure 5.2 shows the peak demand day characteristics.




Figure 5.2: Los Angeles feeder peak demand day, September 24.

5.2.3 New York

For the New York feeder, peak demand occurs on July 16. The feeder peak demand of 6.9MW occurs between 14:40 and 14:55 EDT, following the daily peak temperature of 34.5°C at 14:00 EDT. Average temperature for the day measures 29.6°C. July 16 is shown in Figure 5.3 below.




Figure 5.3: New York feeder peak demand day, July 16.

5.3 Participant Sampling and Number of Simulations

To arrive at the 30% and 70% participation levels, random sampling of the total residential building population is performed for each of the three feeders. The process itself is quite simple: 1) assign a random number between zero and one to each building in the feeder and 2) compare to the participation percentage. Those buildings whose number is less than or equal to the percentage, are participants. Participants are then assigned a new temperature schedule containing the temperature adjustment requested in the demand response event.

Because physical characteristics vary from home to home, the magnitude of demand reduction will vary based on the population sample. To arrive at the expected mean demand reduction (and a distribution thereof), a Monte Carlo method must be employed. The number of samples required is not known *a priori*, therefore an experiment to establish the minimum number of samples required for convergence of the expected mean response is performed.

One hundred samples are drawn randomly from the feeder residential building population. Each population is then simulated for each of the four combinations of event durations and participation rates. To check convergence, normalized root mean square error (NRMSE) between the average demand profile after simulation N and the previous average demand profile after simulation $N - 1$ is then calculated for all $1 < N \leq 100$. This results in a NRMSE value at each iteration of the Monte Carlo experiment. Figure 5.4 shows the results of this process for the Houston feeder.

For all event durations and participation rates, convergence is reached rather quickly, implying that the population response is not as diverse as expected. There are likely several reasons for this observation. This may be a result of the method used to populate the feeder originally. It may also be that, regardless of characteristics, the response of individual homes does not vary much from one to another. Or it may simply mean that the size of the sample is so large that most samples characterize the population very well.

Regardless of the reason, this is good news, as it implies that once a minimum number of samples is determined, the number can be used for all experiments in subsequent MPC studies.

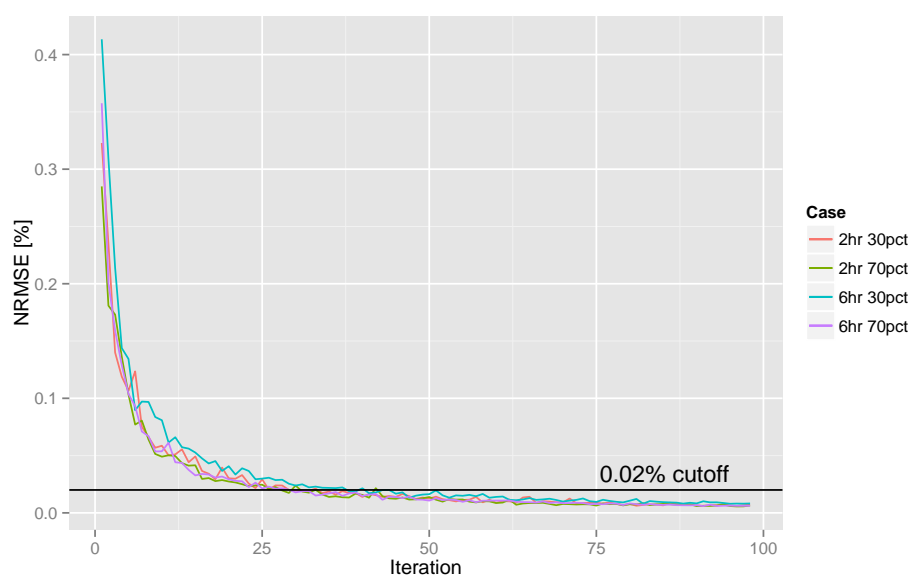


Figure 5.4: Convergence of feeder demand profile NRMSE.

Taking the value at approximately 95% convergence, or roughly 0.02^1 on the chart, gives a reasonable approximation of the number of simulations required. Therefor, throughout the remaining work, 50 samples will be used.

5.4 Interpretation of Metrics Tables

The performance metrics described in Chapter 2 are presented in tables like that shown in Table 5.1. With one exception, all values reported in these tables are simply the arithmetic difference between the case being evaluated and the base case. Positive values indicate an increase over the base case, negative values the opposite. Values shown are the average of daily values computed for all days of the month.² Rows labeled “Peak to Valley” are the exception. These show the ratio (expressed as a percentage) of the optimized to the base case peak to valley ratios. Thus, a value of 90% indicates that the optimized peak to valley ratio is 90% of the base case peak to valley ratio. Absolute differences are not shown.

The values for the Load Factor row require special explanation. Load factor is expressed as a percentage of average demand to peak demand. The values reported in the tables are the difference in percentages, not the percentage differences. For example, if the optimized case has a load factor of 55% and the base case a load factor of 50%, 5% would be reported in the table. The interpretation is that the load factor increased by a value of 5% of the peak demand.

5.5 Houston

Figures 5.5, 5.6, 5.7 and 5.8 show the results from a +2K offset demand response event for both participation levels and event durations. The value reported on the ordinate is the total real feeder demand across the three phases, measured at the distribution transformer. The DR case is plotted as a ribbon, representing the upper and lower extents of the demand across the 50

¹ The 95% represents a confidence interval of two standard deviations.

² In the DR cases, only one day is simulated, therefor the values are simply computed for one day.

simulations, and a line representing the average.

The most striking feature of all four charts is the spike in demand that occurs immediately following the event when thermostats return to their pre-event set points. In industry this is known as ‘rebound’ or ‘snap-back’ and is a phenomenon known to occur frequently unless mitigated through a randomized delay in set point return, duty cycling or a ramped return to set point. The maximum five-minute demand seen in the rebound exceeds the nominal feeder design load of 12MW by almost 3MW, and exceeds the base case by approximately 4MW. Maximum five-minute demand reduction is approximately 4MW.

The second obvious feature is the rapid deterioration in demand reduction. This occurs as the internal temperatures in the homes rise to, and then exceed the new cooling set points. The time it takes any given home to float to the new temperature varies, which is reflected in the gradual return to the pre-event demand. In the six hour event, there is virtually no demand relief after a little more than two hours. In the two hour event, some relief remains, but the savings are marginal.

It is worth noting that the amount of electricity saved during the event is almost entirely offset by the electricity required to bring the homes back to set point after the event ends. For example, in the two hour 70% case, savings measure 0.4% of total electricity consumption. In the six hour 70% case, however, the additional savings are incremental at 0.9% total. For long events, some electricity can potentially be saved if the homes return to set point when cooling loads are lower. That said, long events that push into the occupied hours would likely result in increased opt-outs from unhappy residents.

Finally, performance metrics for the demand response event generally show performance worse than the base case. The six hour 30% participation case does show a slight peak demand reduction when averaged over a fifteen-minute window.³ This is consistent with the observations

³ Recall that the values presented in the table are fifteen-minute average values, not the five-minute values shown in the charts.




Figure 5.5: Houston feeder demand curves for 2hr, 30% DR event.




Figure 5.6: Houston feeder demand curves for 2hr, 70% DR event.




Figure 5.7: Houston feeder demand curves for 6hr, 30% DR event.




Figure 5.8: Houston feeder demand curves for 6hr, 70% DR event.

made regarding the charts. Peak demand, peak to valley, load factor and ramping all show values opposite of the desired behavior. Some very small reductions in electricity consumption can be seen, with the greatest reduction in the longer six hour event.

Table 5.1: Performance metrics for Houston feeder demand response events.

	6 Hour		2 Hour	
	70%	30%	70%	30%
Electric Consumption [MWh]	-1.38	-0.64	-0.65	-0.29
Peak Demand [MW]	1.80	-0.05	2.54	0.75
Peak to Valley [%]	116.11	99.55	122.68	106.72
Load Factor [%]	-8.60	0.03	-11.07	-3.80
Ramp [MW]	12.14	4.67	12.09	4.55

5.6 Los Angeles

The Los Angeles feeder simulations show similar characteristics, but smaller demand reduction potential in both magnitude and percentage of peak. In absolute terms, five-minute demand reduction potential measures only about 0.8MW. The obvious reason for this is the relatively small penetration of cooling equipment in this region. Referring back to Figures 3.10 and 3.11, one can see that only slightly more than half of the homes in Los Angeles are equipped with cooling; in Houston, nearly all homes are. This fact has important implications in later studies since the amount of electric demand associated with cooling — the “flexible cooling demand” — limits the potential of all demand reduction measures.

Similar to the Houston feeder, metrics for the Los Angeles feeder generally show an increase in fifteen minute peak demand, and worse performance in the remaining categories. Again, a very slight reduction in the six hour 30% participation case can be seen. Very slight reductions in electricity consumption can also be observed.




Figure 5.9: Los Angeles feeder demand curves for 2hr, 30% DR event.




Figure 5.10: Los Angeles feeder demand curves for 2hr, 70% DR event.




Figure 5.11: Los Angeles feeder demand curves for 6hr, 30% DR event.




Figure 5.12: Los Angeles feeder demand curves for 6hr, 70% DR event.

Table 5.2: Performance metrics for Los Angeles feeder demand response events.

	6 Hour		2 Hour	
	70%	30%	70%	30%
Electric Consumption [MWh]	-0.40	-0.19	-0.19	-0.08
Peak Demand [MW]	0.34	-0.03	0.61	0.20
Peak to Valley [%]	108.37	99.34	115.13	105.03
Load Factor [%]	-5.28	0.22	-8.50	-3.11
Ramp [MW]	1.81	0.62	2.47	0.85

5.7 New York

New York feeder simulations show the same behavior as both Houston and Los Angeles, and fall somewhere between the demand reduction potential they exhibit. Maximum five-minute demand reduction potential is slightly more than 2MW in the two hour 70% case. Again, this is due to the penetration of residential cooling equipment. Similar to the Houston case, demand rebound sets a new peak demand. Here, the effect is somewhat less dramatic with a five-minute peak just slightly more than 8.5MW in the two hour 70% case, compared to the base case value of 6.9MW. Again, metrics for the New York feeder show trends similar to those seen in the previous cases.

Table 5.3: Performance metrics for New York feeder demand response events.

	6 Hour		2 Hour	
	70%	30%	70%	30%
Electric Consumption [MWh]	-0.84	-0.34	-0.41	-0.17
Peak Demand [MW]	0.71	-0.04	1.44	0.50
Peak to Valley [%]	110.54	99.41	121.41	107.46
Load Factor [%]	-6.46	0.16	-11.28	-4.46
Ramp [MW]	5.96	2.14	6.44	2.47

The cases here are obviously extreme examples in which there is no mitigation of the rebound effect. That said, mitigation can only temper the magnitude of the rebound, not eliminate it entirely. DR events can set new demand peaks if participation is large relative to the total demand on the system (the portion of the distribution grid outside of the feeder). To prevent this from




Figure 5.13: New York feeder demand curves for 2hr, 30% DR event.




Figure 5.14: New York feeder demand curves for 2hr, 70% DR event.




Figure 5.15: New York feeder demand curves for 6hr, 30% DR event.




Figure 5.16: New York feeder demand curves for 6hr, 70% DR event.

occurring, events must be carefully crafted and participation limited. Metrics calculated for each feeder suggest that traditional DR can result in increased demand, peak to valley ratio and ramping, and degraded load factor at the feeder level.

Suggestions for improving traditional DR are not within the scope of this work. Rather, these investigations (extreme as they are) provide context for a discussion regarding the alternative approach proposed in this work. The remainder of the work will evaluate several methods for shaping electric demand using model predictive control which could potentially replace or supplement traditional DR.

Chapter 6

Demand Limiting Optimization

The previous chapter highlighted the potential problems associated with traditional demand response implemented across large percentages of the building population. In this chapter, a simple alternative is proposed which seeks to minimize feeder demand without introducing a new peak unintentionally. This method, referred to as ‘demand limiting optimization’, will use a decentralized MPC approach to achieving demand reductions at the distribution feeder level. The approach is first described, followed by results from each of the three distribution feeders. A follow-on study is presented which assists in explaining the results.

6.1 Methodology

Achieving reductions in feeder demand with a decentralized MPC approach may be accomplished in a number of ways. The method utilized here relies on the observation that peak feeder demand is coincident with peak demand at each home. Assuming no interaction exists between homes leading to nonlinear aggregate load behavior, i.e. adopting the principle of superposition, results in a very simple control strategy implemented in the controller: minimize peak demand at each home. In principle, minimization of peak demand at the home level should result in a reduction in peak demand at the feeder level.

To reduce demand at each home, cooling set points can be adjusted up to provide relief to the cooling system. This is not unlike the set point adjustments utilized in the demand response study. However, a return to the desired cooling set point results in a demand spike as cooling

equipment is asked to lower the air temperature within the home. The controller can anticipate these demand spikes, and can therefor make incremental adjustments up and down to the cooling set point to prevent the spikes from occurring. This can be accomplished through precooling, via cooling set point reductions in periods of low demand, to offset the need for cooling in periods of high demand.

In order to allow both cooling relief from set point adjustments up, and demand reduction planning through set point adjustments down, the controller is allowed to make adjustments to the cooling set point within pre-defined boundaries relative to the nominal cooling set point. Recall from Section 4.2, that these boundaries are subject to assumptions regarding occupancy. In this study and those that follow in subsequent chapters, each home is assumed to be unoccupied for ten hours during the day, with occupants departing at $8:00 \pm 1$ hour. The departure time is randomized (but fixed for all simulations thereafter) for each home, leading to a shift in the 48 optimization modes. Randomization is necessary to capture the diversity of schedules, but also to prevent unintended synchronization as shown later in Chapter 7.

The upper and lower boundaries used in this study, specified as deltas from the base case cooling set point, are shown in Table 6.1. The controller may make adjustments in increments of 0.25K, corresponding with the precision of a typical residential thermostat. The -2K lower boundary assumption recognizes that large temperature swings during the occupied period may not be tolerated by occupants.

Table 6.1: Cooling set point boundaries for demand limiting optimizations.

	Occupied	Unoccupied
Upper Boundary	+0K	+3K
Lower Boundary	-2K	-5K

6.1.1 Objective Function Definition

For each home, the objective function to be minimized by the controller is the maximum value of the series found by taking a simple moving average of house electric demand over the planning horizon. This can be expressed as:

$$\min \left[\max \left(\{p_j\}_{j=1}^{N-n+1} \right) \right] \quad (6.1)$$

$$p_j = \frac{1}{n} \sum_{i=j}^{j+n-1} p_i \quad (6.2)$$

where:

p_j is the simple moving average of p at index j

j is the moving average index

N is the length of the time series p

n is the period of the moving average

p_i is the electricity demand of the house at time step i

i is the time step index

House demand contains peaks due to HVAC cycling as seen in Figures 3.2 and Figures 3.3. The length of the moving average period will determine the degree of smoothing, and will therefore have an impact on the optimization. A set of experiments were performed to find the value of the moving average period that resulted in the greatest feeder demand reduction while still yielding a smooth aggregate demand curve. Periods of 15, 30, 60, 90 and 120 minutes were tested, with a period of 60 minutes providing the best balance of demand reduction without introducing the oscillations seen in the 90 and 120 minute cases. Figures B.2, B.3, B.4 and B.5 in Appendix B show the results of the experiments for the Houston Feeder.

6.2 Houston

Figure 6.1 shows the results of the demand limiting optimization for each day¹ in the month of July. Each panel in the grid of results represents a single day of optimization under the demand limiting objective. Visual inspection of the results shows three key characteristics of the optimized feeder demand that is consistent from day to day:

1. Precooling activity can be seen in many of the early morning hours as the controller depresses the cooling set point. This can be confirmed by inspecting the results from individual homes which are not presented here. Depressing the set point allows the controller a greater range of freedom when adjusting the set point up to avoid demand later in the day.
2. Peak demand is reduced in all days of the month. On some days, July 21 for example, the amount of reduction is relatively large compared to others. This is suspected to be related to the amount of flexible cooling load available, and the efficiency of precooling.
3. Troughs in demand are filled-in by the controller actions. Although this was not explicitly coded into the objective function, it is a convenient side-effect. This behavior makes sense intuitively: as the controller seeks to minimize demand, it will tend to spread that demand to adjacent hours in an effort to make demand consistent across the moving average window.

Table 6.2 lists the values of the performance metrics. The optimization shows improvements over the base case across the board, with the exception of total electric consumption. Peak demand is reduced by some 590kW on average, peak to valley ratio is reduced by over 80%, load factor is increased 5% and total system ramping is reduced by 610kW. The increase in total electric consumption can be explained by imperfect thermal storage efficiency of the envelope. That is, the energy invested in precooling the building does not result in an equivalent reduction in cooling

¹ One might notice that July has been conveniently truncated to 30 days for purely aesthetic reasons. The metrics tables are computed using all 31 days in the month.

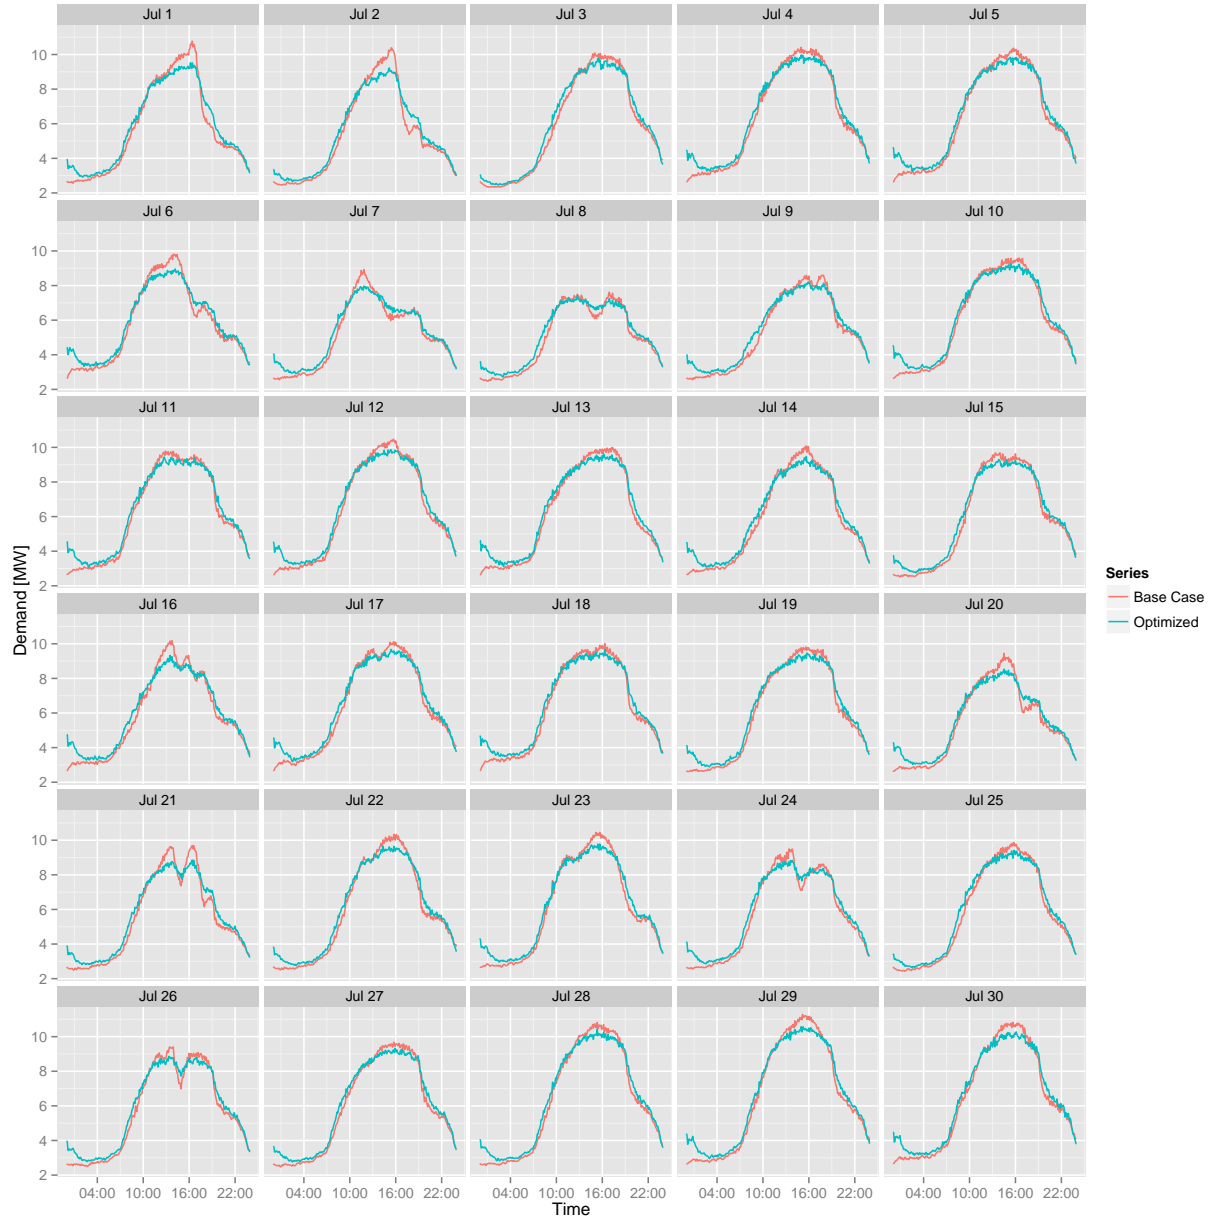


Figure 6.1: Feeder demand profiles for Houston demand limiting optimization, 70% participation.

energy later in the day. This is an expected result as the objective is peak demand minimization, not total energy minimization, and the controller is allowed to use as much or as little energy to reduce demand as possible. A quantification of thermal storage efficiency is presented in Appendix G.

Table 6.2: Performance metrics for Houston feeder demand limiting optimization, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	3.76	3.02	4.53
Peak Demand [MW]	-0.59	-1.26	-0.22
Peak to Valley [%]	81.59	72.35	92.26
Load Factor [%]	5.47	3.18	8.85
Ramp [MW]	-0.61	-3.77	1.60

A plot showing the spectral power distribution of the feeder demand time series shows a different view of system ramping. In Figure 6.2, box plots show the percentage of total spectral power associated with different frequency ranges. Boxes represent 31 daily values, corresponding to the 31 days in July. This analysis shows an increase in the DC portion of the power spectrum for the optimized case. The DC portion is defined here as the range of periods greater than 24 hours. This implies that feeder demand has been slightly flattened with less power associated with periodic trends. In general, a plot showing increased DC power and reduced power in every other frequency bin (compared to the base case) is desirable; a plot like this would indicate an increase in load factor and reduced ramping. This is almost the case here, with only marginal increases in spectral power in the range between twelve and four hours.

Plots of individual days such as those shown in Figures 6.3 and 6.4 give an indication of the variability in the results due to differences in the population of participants. The small amount of variation reflects observations made in the demand response cases. That is, with a relatively large sample each time, the results do not change significantly.² These two days are exemplary in the degree of demand reduction and leveling achieved.

² This is in part related to the sensitivity of the reduced order model to internal mass level. Subsection B.1.1 discusses this in more detail.




Figure 6.2: Total spectral power as a function of frequency bin for Houston feeder demand limiting optimization, 70% participation.




Figure 6.3: Feeder demand profiles for Houston, July 2 demand limiting optimization, 70% participation.




Figure 6.4: Feeder demand profiles for Houston, July 21 demand limiting optimization, 70% participation.

The results are quite amazing considering that there is no coordination between buildings. Instead, each is acting according to its own simple objective. The reason this uncoordinated optimization results in reduce feeder demand is two-fold.

First, the objective of the homes is well aligned with the global feeder objective because home and feeder demand tend to be coincident. Were this not the case, the homes' objective may subvert the feeder's. This has implications for large deployments where the feeder peak demand may not necessary be coincident with the peak of its parent system.

Second, the diversity of buildings and therefor the decisions made by the controller tend to average out any individual home's actions. Recall that cooling demand appears as a square wave with at most two non-zero demand values. Given the formulation of the objective function, specifically that it is a moving average of demand, the controller will seek to flatten the moving average. When the moving average is perfectly flat, the underlying house demand will tend to be a square wave with regular periodicity. Small differences in the set points, building characteristics, etc. will result in offsets between these square waves across the building population. In a large enough population, the result is a peak in demand in one house coincident with a trough in another. On average this results in an overall demand reduction.

It should also be noted that although the controller may only affect the cooling related portion of a home's demand, it sees the whole home demand. This is particularly relevant for homes with electric hot water heaters. In these homes, hot water demand may occur during the middle of the feeder peak demand period. The controller may shift cooling operation away from these times to avoid any coincident cooling and hot water demand, thus helping to reduce feeder peak.

Results for the 30% participation simulations show similar, but smaller reductions in demand. Observations made regarding the 70% simulations apply equally in this case. Results for the 30% participation case can be seen in Appendix B, Figures B.6 and B.1, and Table B.1.

6.3 Los Angeles

Results for the Los Angeles feeder are plotted in Figure 6.5 and summarized in Table 6.3. Demand reduction is very modest for all days in the month of July, with a daily average of only 70kW. Despite the rather small reductions, all performance metrics show some improvement. Very little additional electric consumption results from the optimization, likely due to a lack of precooling by the controller. The small reductions in demand appear to be related to the size of the flexible cooling demand seen in Chapter 5. This issue is further addressed in Section 6.5. It is somewhat surprising that, despite the small reductions in demand, the load factor increases as much as it does. This can be attributed to both the small reduction and peak demand, and the small increase in demand in the hours preceding the peak, thus increasing the average.

Table 6.3: Performance metrics for Los Angeles feeder demand limiting optimization, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	0.62	0.46	0.74
Peak Demand [MW]	-0.07	-0.14	-0.04
Peak to Valley [%]	97.30	93.60	98.70
Load Factor [%]	2.27	1.67	3.68
Ramp [MW]	-0.04	-0.27	0.18

The spectral power analysis shows a much more positive result. With the exception of a minuscule increase in spectral power in the four to six hour frequency range, the plot shows almost ideal behavior. This suggests that the demand limiting optimization, although not suited for large decreases in peak demand in this case, provides benefits in terms of reduced ramping requirements. It is an interesting result considering that the optimization is not incentivized to reduce ramping.

Figures 6.7 and 6.8 serve to illustrate the possible range of results from any given population. Despite the variation in results, both days exhibit modest demand leveling as small amounts of demand are shifted from the peak hours to hours both before and after. On July 29, the precooling activity extends into the early hours of the day, which is uncommon among the days in the month.

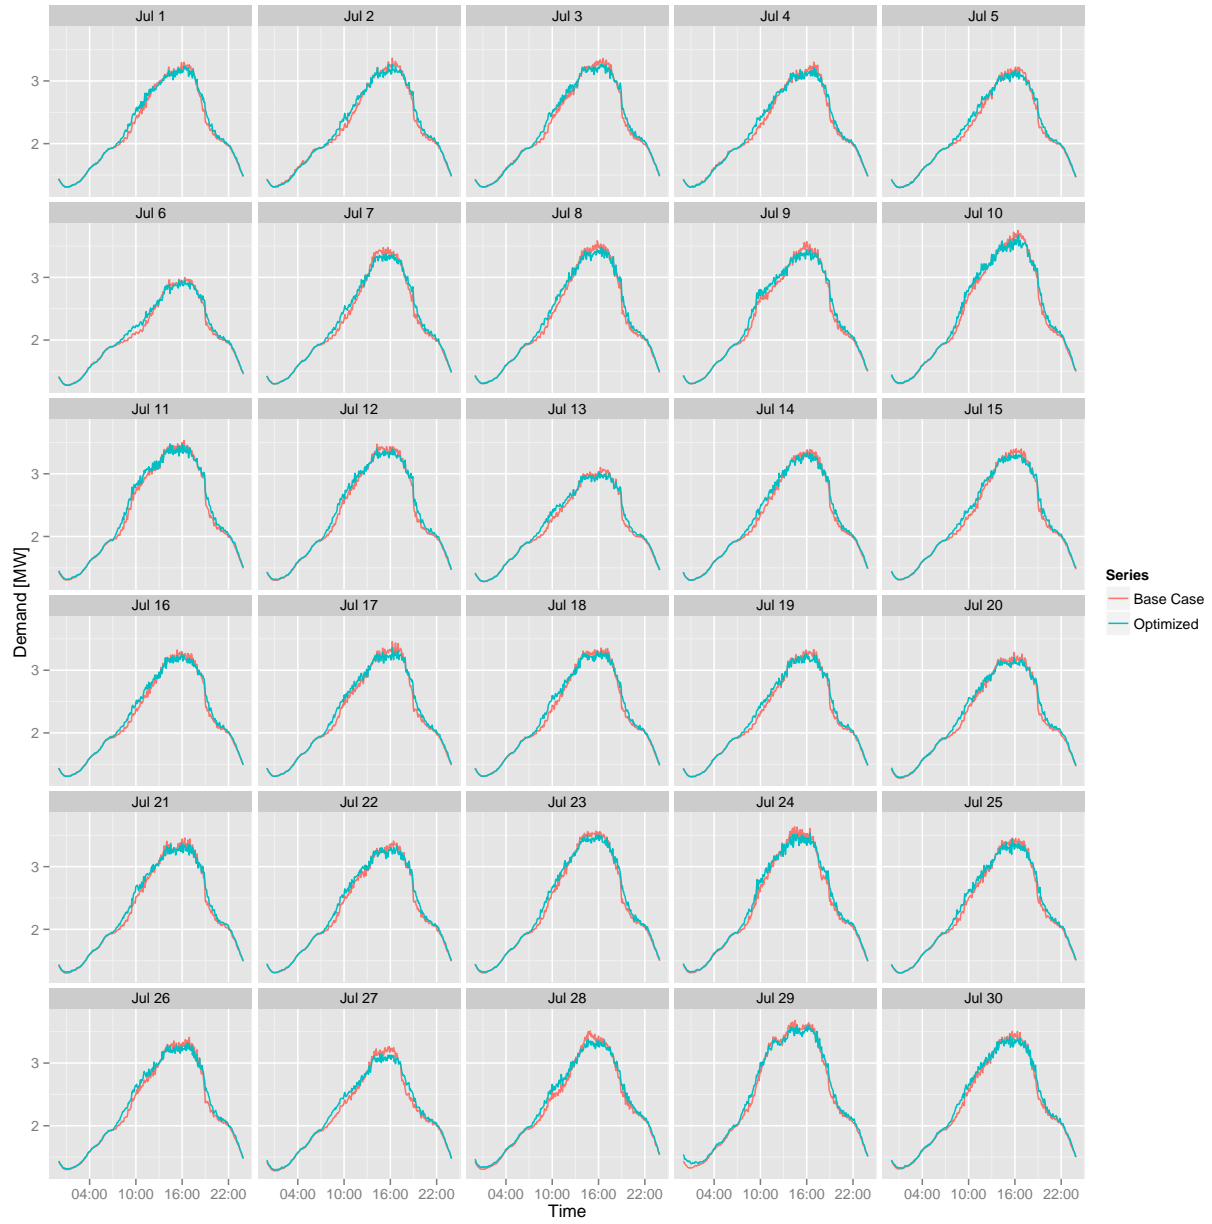


Figure 6.5: Feeder demand profiles for Los Angeles demand limiting optimization, 70% participation.




Figure 6.6: Total spectral power as a function of frequency bin for Los Angeles feeder demand limiting optimization, 70% participation.




Figure 6.7: Feeder demand profiles for Los Angeles, July 10 demand limiting optimization, 70% participation.




Figure 6.8: Feeder demand profiles for Los Angeles, July 29 demand limiting optimization, 70% participation.

Results for the 30% participation case can be seen in Appendix B, Figures B.11 and B.10, and Table B.2. Not surprisingly, improvements over the base case are very small compared to previous cases.

6.4 New York

Performance metrics in Table 6.4 and daily plots of demand in Figure 6.9 indicate demand reduction potential between that of the Houston and Los Angeles cases. This observation is consistent with the findings in Chapter 5. Similar to both Houston and Los Angeles, all performance metrics show improvement, with some penalty in electric consumption. Power spectrum characteristics lie between results from the previous feeders, with only very small increases in power in the twelve to four hour frequency range; much smaller than either the increase in the DC component or the decrease in the twenty-four to twelve hour frequency range.

Table 6.4: Performance metrics for New York feeder demand limiting optimization, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	1.84	1.51	2.56
Peak Demand [MW]	-0.22	-0.39	-0.09
Peak to Valley [%]	87.39	76.18	96.90
Load Factor [%]	3.93	2.46	5.98
Ramp [MW]	-0.14	-0.70	0.47

Plots showing individual days in more detail highlight the smoothing effect that the controller has on small irregularities in feeder demand. Disregarding the first hour in both days, this can be seen in Figure 6.11 around 4:00 in the morning, and in Figure 6.12 around 17:00 at night. Small adjustments in feeder demand such as these likely contribute to the improved ramping metrics.

Results for the 30% participation case can be seen in Appendix B, Figures B.14 and B.13, and Table B.3.

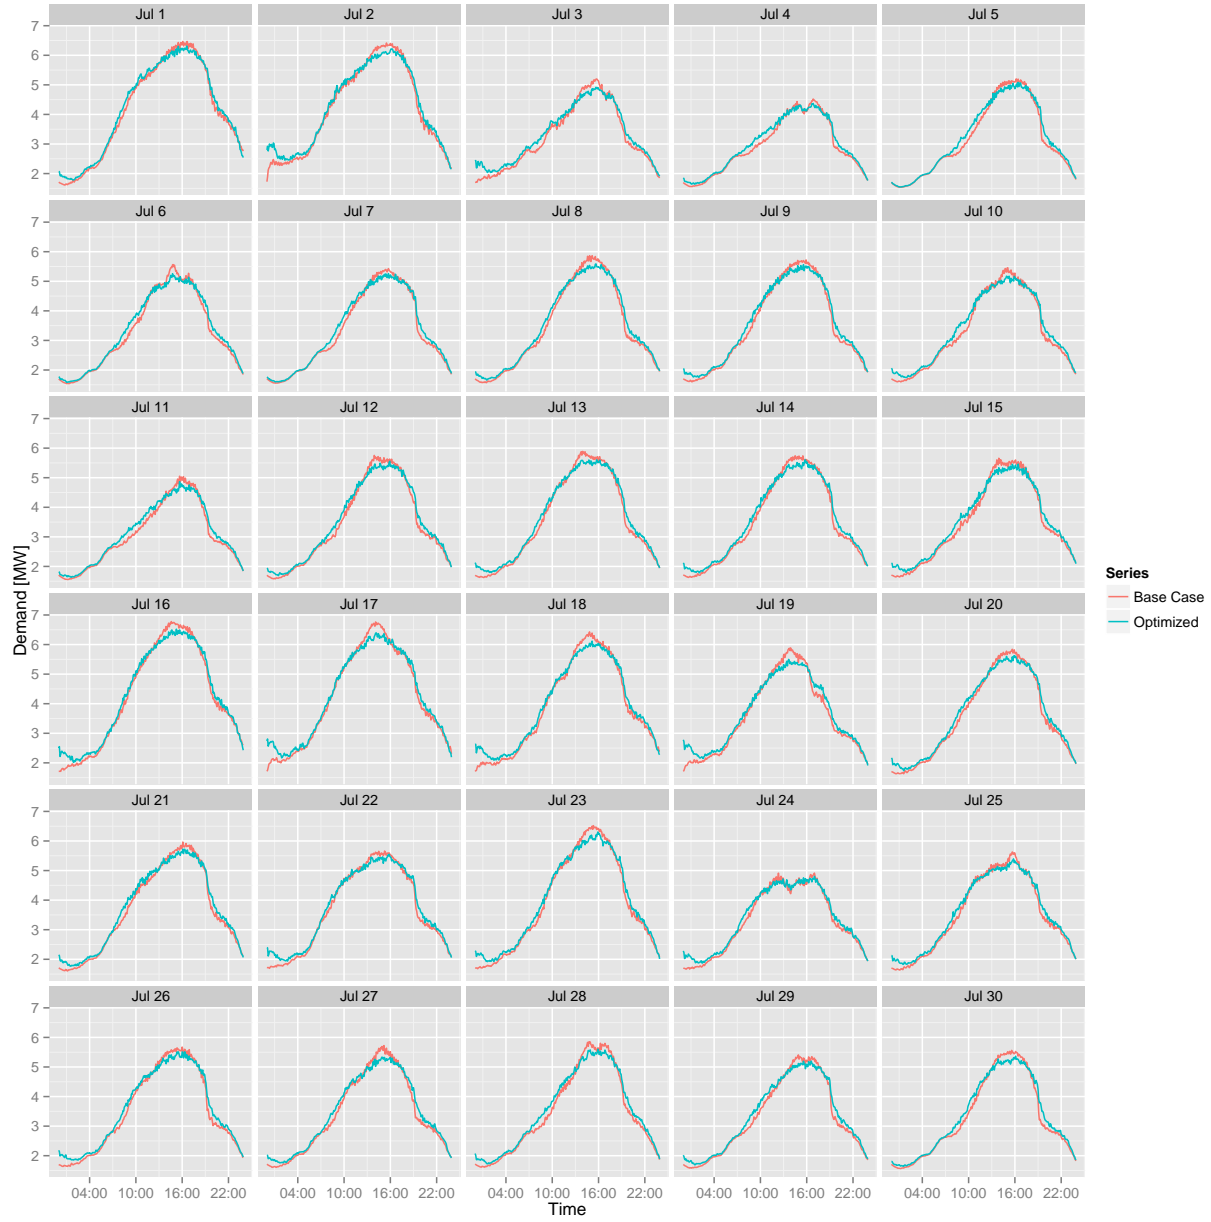


Figure 6.9: Feeder demand profiles for New York demand limiting optimization, 70% participation.




Figure 6.10: Total spectral power as a function of frequency bin for New York feeder demand limiting optimization, 70% participation.




Figure 6.11: Feeder demand profiles for New York, July 16 demand limiting optimization, 70% participation.




Figure 6.12: Feeder demand profiles for New York, July 17 demand limiting optimization, 70% participation.

6.5 Precooling Investment

To better understand the difference in results across the three feeders, an experiment is constructed to measure the amount of precooling the controller may employ to reduce peak demand. Indirectly, the experiment estimates the amount of flexible cooling load available. The experiment works by forcing HVAC systems off during a specified window of time. This methodology is similar to the demand response case, except that the controller has advance knowledge of the event by way of an artificially high demand multiplier during the two-hour window. The experiment is expected to show that both the total amount of precooling invested, and the magnitude cooling demand shifted, is dependent on the feeder, explaining why one feeder performs differently than another.

This is easily accomplished using the MPC environment by defining an objective function that minimizes the sum of demand cost over the day. This can be expressed as:

$$\min \left(\sum_{i=j}^k p_i \cdot c_i \right) \quad (6.3)$$

$$c_i = \begin{cases} 10e^6, & l < i \leq m \\ 1, & \text{otherwise} \end{cases} \quad (6.4)$$

where:

i is the time step index

j is the cost horizon time step start index

k is the cost horizon time step end index

p_i is the electricity demand of the house at time step i

c_i is the demand multiplier at time step i

l is the demand period start index

m is the demand period end index

As with the previous optimization cases in this chapter, the controller is allowed to manipulate the cooling set point to achieve the desired objective. However, in this case, the controller is not

allowed to make adjustments to the set point above the base case cooling set point. Doing so would put the controller at an advantage, as demand can be simply reduced without requiring the controller to prepare for the window. The main goal of the experiment is to measure the amount by which the controller will increase demand prior to the window and prepare the homes for the increased cost. This gives a measure of the amount of precooling investment the controller will make to avoid running the HVAC during the window.

Results of the test for the three feeders are shown in Figures 6.13, 6.14 and 6.16. In this particular test, a two-hour window beginning at 12:00 is simulated for each day of the month.

All feeders show three similar trends: 1) demand is increased by the controller in preparation for the window, 2) demand is significantly reduced during the window and 3) demand spikes at the end of the window when the HVAC systems are reengaged. Where the feeders differ is in the amount of preparation, i.e. precooling, undertaken prior to the window and the size of the demand reduction during.

In the Houston case, nearly every day begins with increased demand over the base case. This is the result of a set point depression below the base case set point. This depression is maintained until occupants leave between 7:00 and 9:00. At this point, the controller tends to depress the set point even further, maintaining an increased demand over the base case. Immediately prior to the beginning of the window, the controller again depresses the set point, creating a large spike in demand. During the window, all HVAC related demand is eliminated. This is followed by a spike in demand, i.e. recovery, at the end of the window as cooling reengages to bring the homes back to set point. It is important to note that demand has been shifted to early portions of the day, when demand is relatively low compared to the late afternoon hours. This helps to explain how the demand limiting optimization above was able to reduce demand during peak hours.

In the Los Angeles case, very little preparation occurs for the two-hour window. Only one day (July 29) exhibits set point depression prior to the end of occupancy. Generally, set point depressions prior to the event result in very modest increases in demand, further suggesting limited precooling activity. Similarly, the demand spike immediately before and after the window, and the

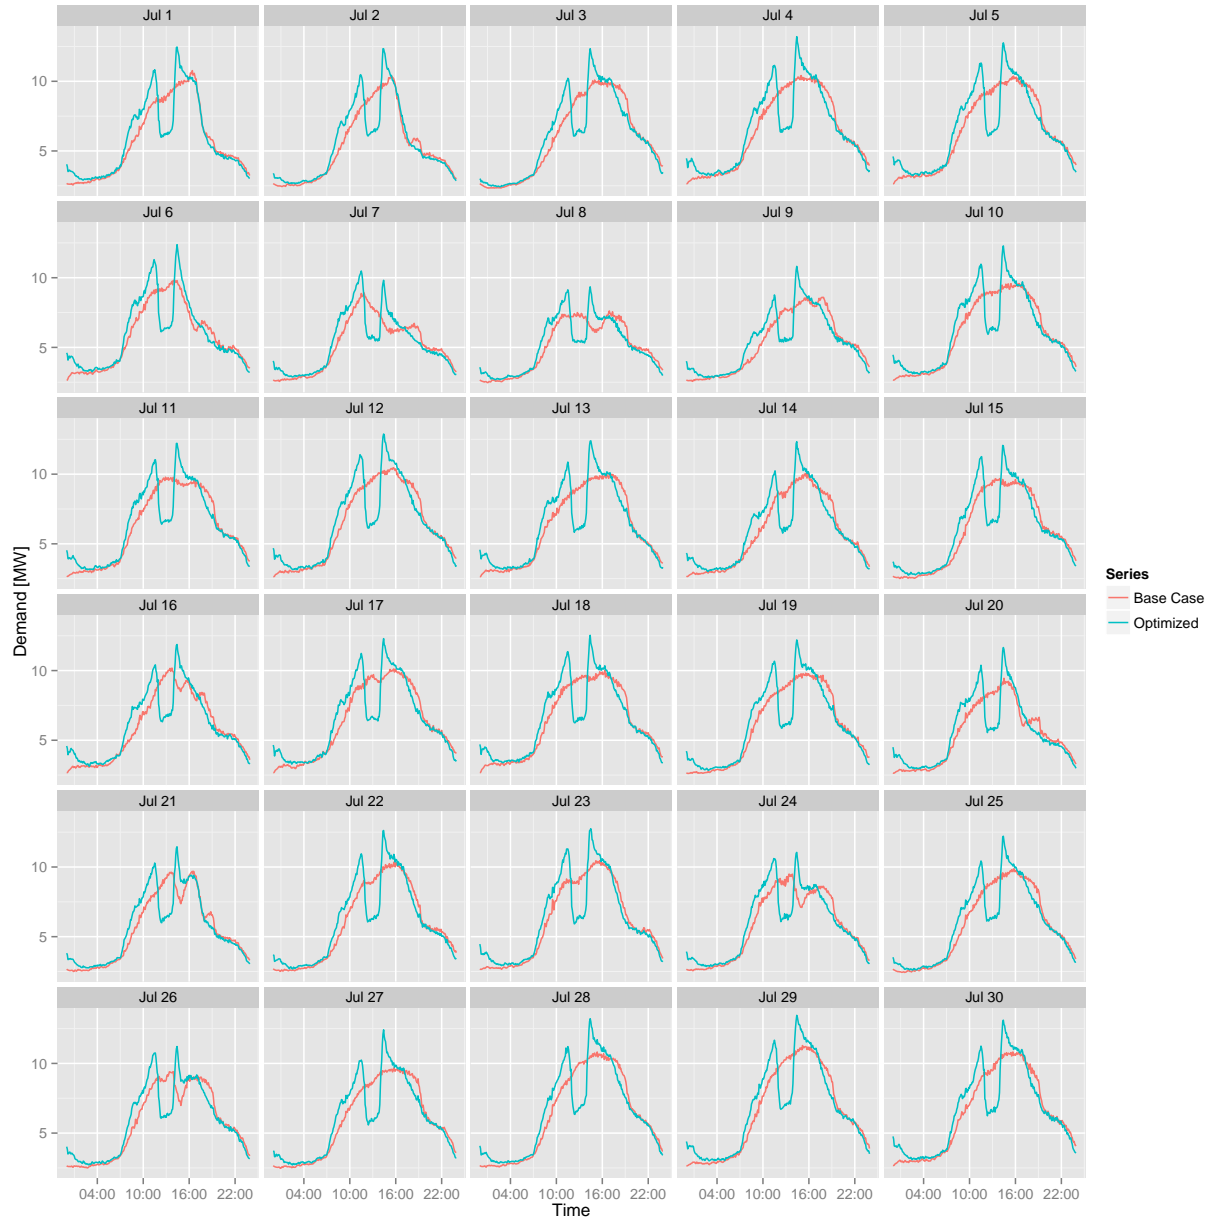


Figure 6.13: Experiment to measure precooling investment and flexible demand in Houston feeder, noon, 70% participation.

demand reduction during, is relatively small compared to Houston. This is consistent with the findings of the demand response cases: there is relatively little flexible cooling demand available. These results suggest precooling is not particularly effective in this house population, perhaps due to losses related to thermal storage efficiency. This is explored with no definite conclusions in Appendix G. While this is indeed possible, it is more likely that precooling is not required to avoid HVAC demand during the window due to the milder climate. This much could be surmised by a lower penetration of air conditioning in this feeder’s house population.

Taking a more detailed look at one day in Los Angeles illustrates the precooling activity in more detail. On this day, July 10, the two-hour window is progressed through the day to better observe how the controller prepares for the event as a function of time of day. The result of this experiment can be seen in Figure 6.15. Each panel in the plot is labeled with the hour of day in which the window begins. The top row shows that flexible cooling load is close to zero until 10:00. This is seen in the lack of precooling prior, the demand reduction during, and the demand spike following the window. Similarly, for windows beginning at 19:00 and 20:00, almost zero precooling and recovery action can be observed. In the remaining panels, precooling is modest. Preparation does begin as early as 7:00, but the majority of precooling action comes immediately before the window. Similarly, the post-window spike is limited to a very short time frame compared to other feeders.³ The effects of the precooling and post-window recovery has all but disappeared by the early evening hours. This result shows that very little of the demand during the peak portions of the day can be shifted to times when demand is significantly lower. This means that demand can not be reduced without creating new spikes in the hours immediately before or after when demand is already high, and explains the inability of the controller to reduce peak demand in the demand limiting optimization.

As with the demand response cases, the New York results shown in Figure 6.9 lie somewhere between those of Houston and Los Angeles. Some days show significant precooling beginning at

³ See Figures B.8 and B.16 in Appendix B for a look at individual days.

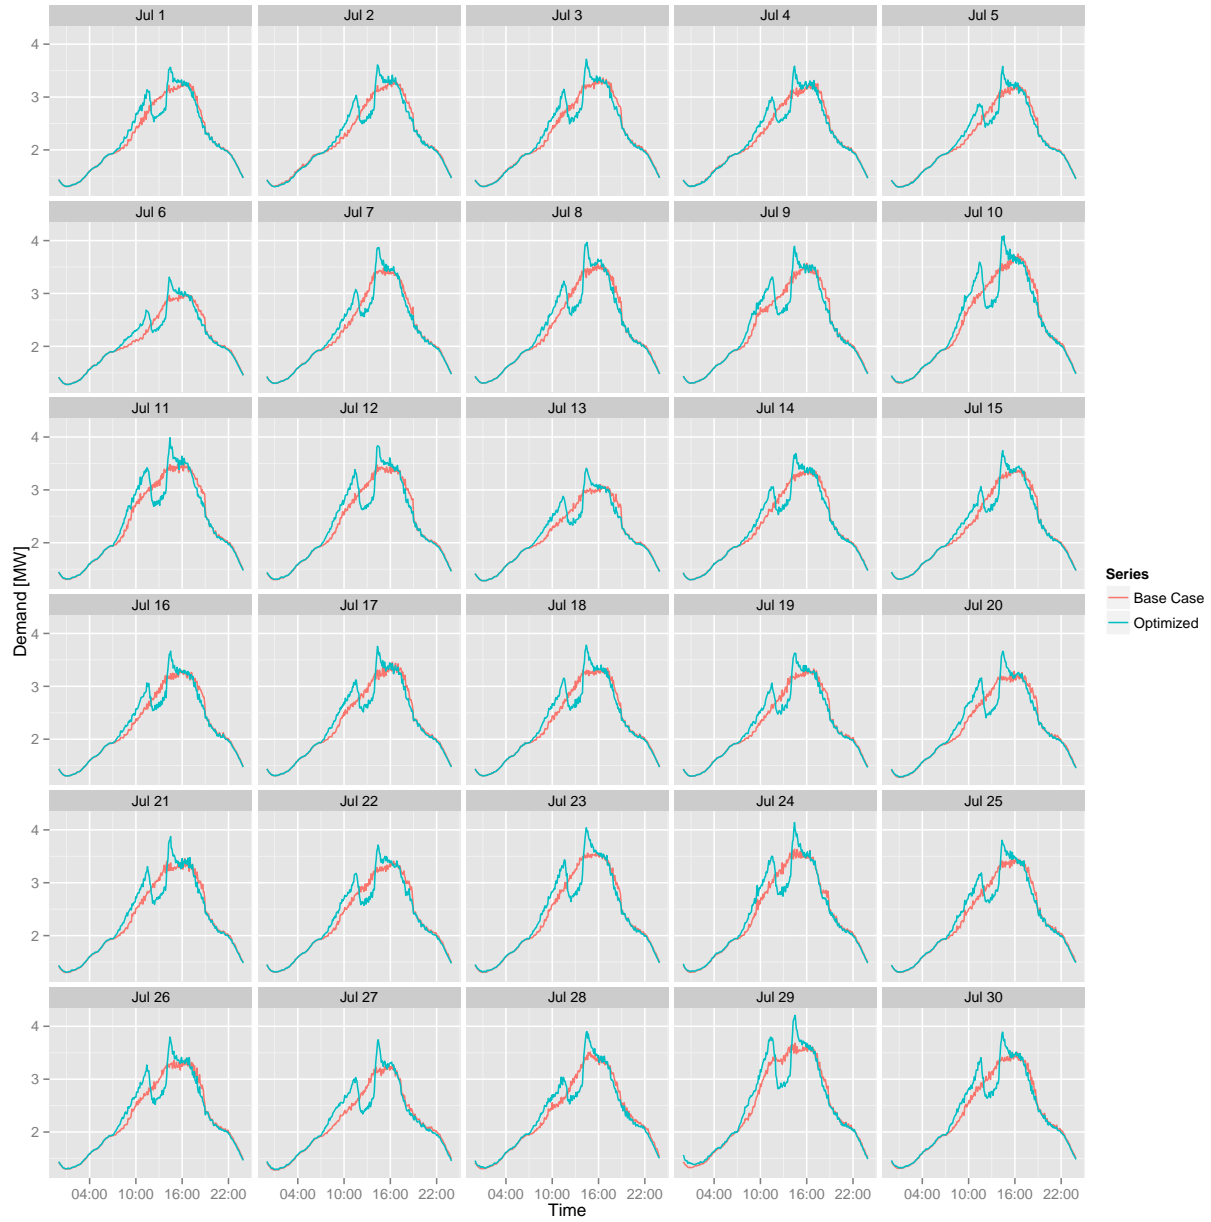


Figure 6.14: Experiment to measure precooling investment and flexible demand in Los Angeles feeder, noon, 70% participation.

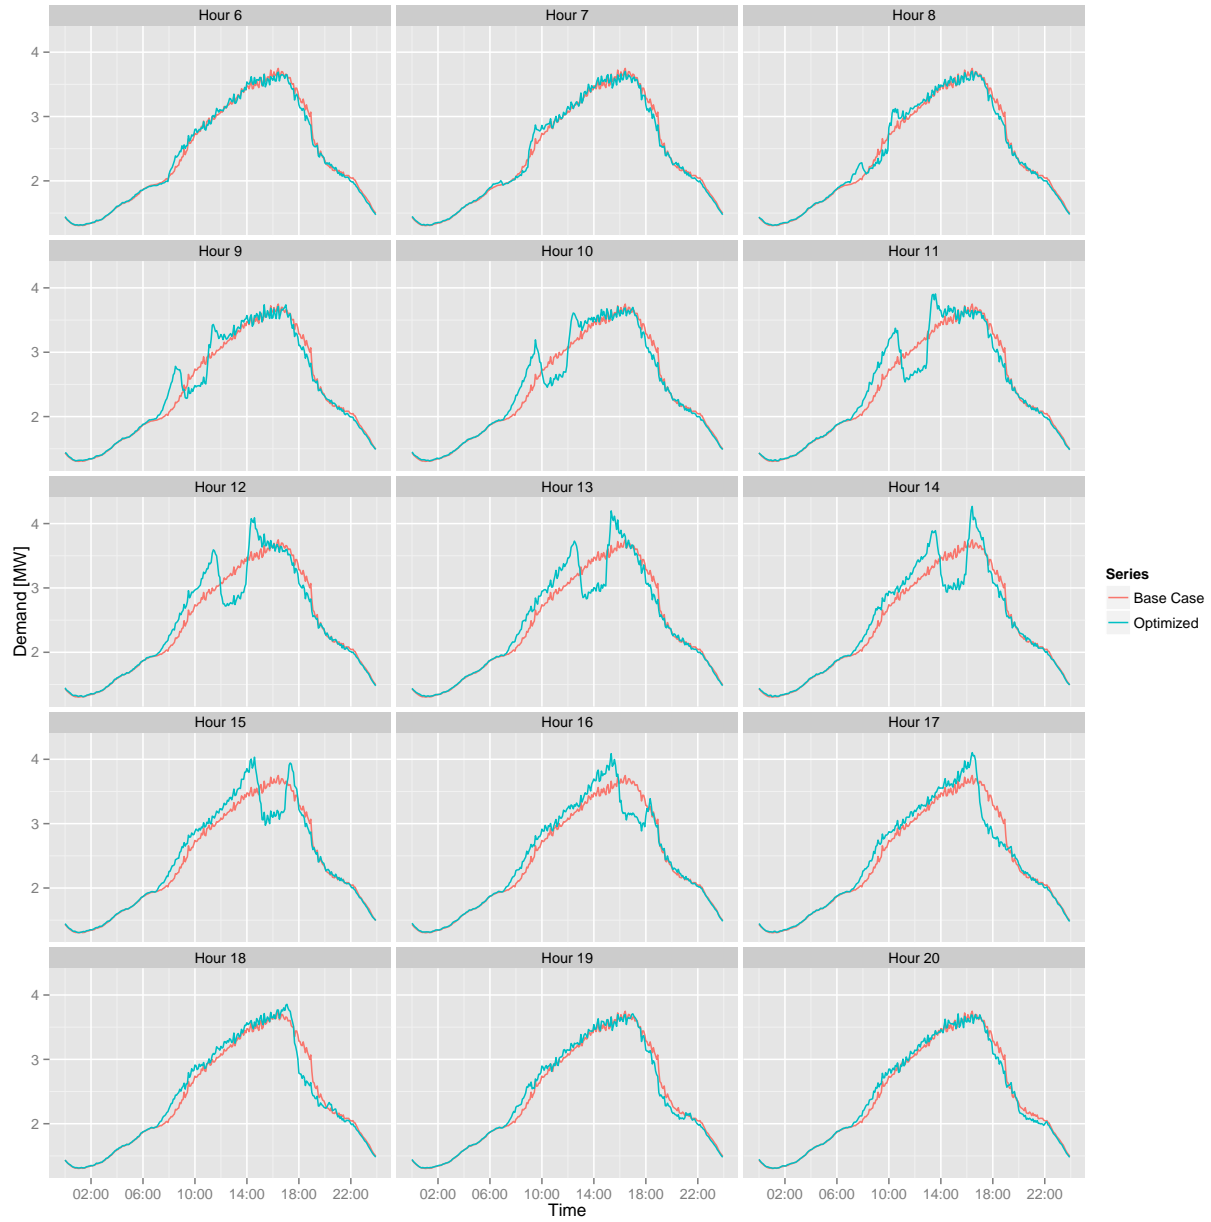


Figure 6.15: Experiment to measure precooling investment and flexible demand in Los Angeles, July 10 feeder, 70% participation.

midnight, while other days show little or no precooling before 7:00. Not surprisingly, the size of the demand spikes and reduction are similar to those reported in the demand response case. Similar to the Houston case, the early precooling activity on some days is present when the same day is subjected to the demand limiting optimization. Taking July 16, for example, the early precooling activity in the demand limiting optimization case is also present in the hourly progression of the window. This is shown in Figure B.16. This precooling activity leads to demand reductions in the afternoon in the demand limiting case. On July 7, however, little precooling activity occurs in both the experiment and in the demand limiting optimization, and little demand reduction is seen in the demand limiting optimization.

The results of this experiment show that precooling activity is specific to the feeder in question. The result is consistent with the observations made regarding the demand limiting optimization. That is, when significant precooling exists early in the day, demand can be shifted, resulting in a reduced peak. By extension, this implies that the ability of the controller to shift demand, regardless of objective function, will vary from one feeder to another.

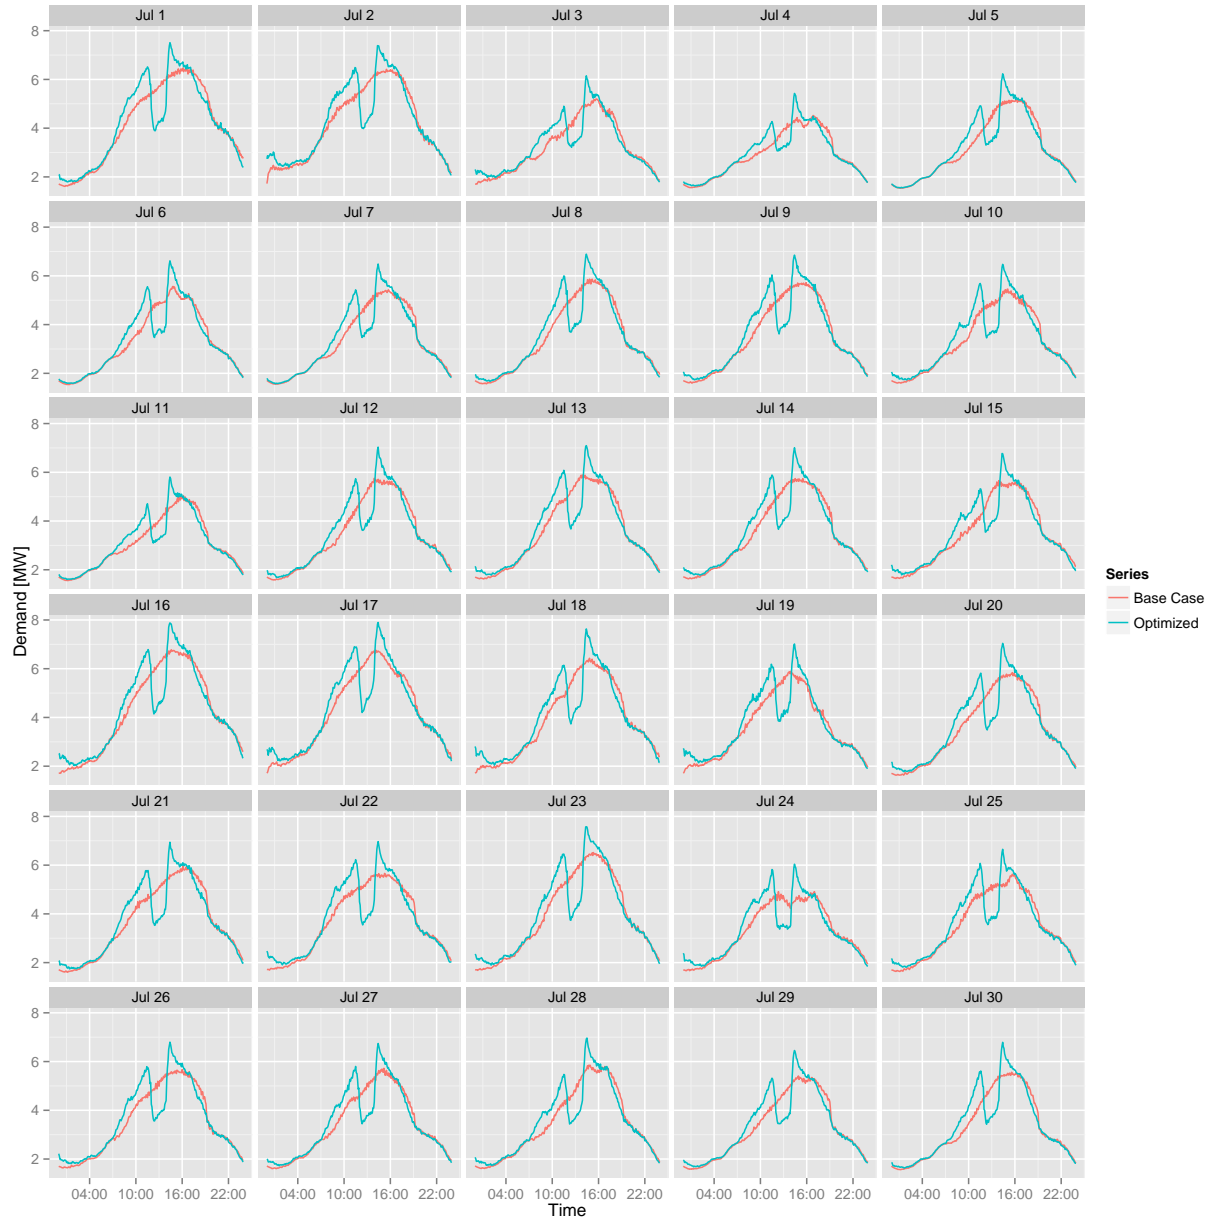


Figure 6.16: Experiment to measure precooling investment and flexible demand in New York feeder, noon, 70% participation.

Chapter 7

Dynamic Price Optimization

Dynamic pricing, including time of use, real-time and day-ahead, has been proposed as a means by which to control demand either directly through price-responsive grid-connected device controllers, or indirectly through behavioral modification. In the former, controllers may be active or passive participants in the market settling price, and demand reduction results from controllers adjusting their consumption in accordance with market price. In the latter, demand reduction results from consumer price awareness via day-ahead notification of high peak price events.

This chapter explores the use of dynamic pricing to shape the electric demand of buildings on the three distribution feeders. Day-ahead pricing is used to inform the model predictive controller so that overall cost may be minimized through price-dependent load planning. The model predictive controller, having the ability to anticipate future demand given weather forecasts and simulated operation, can, in theory, better take advantage of forecast prices in order to minimize operation during the highest price — and therefore highest demand — periods of the day. This is in contrast to auction-based controllers explored in the literature [38, 25, 95] in which the controller is only aware of the current price and demand required by the controlled device, and is unable to consider the impact of decisions on future operation. To give the controller the pricing information it requires, day-ahead price predictions may be used. Historical prices and weather are used to generate a synthetic pricing signal for evaluation.

Results indicate that dynamic-pricing may have unintended consequences that are difficult to predict or control at the distribution feeder level. The chapter concludes with an evaluation

of a synthetic price signal intended to resolve the undesirable traits encountered in the day-ahead pricing cases. An experiment shows the system to be unstable when an objective function of the form studied in this chapter is used, suggesting an alternative approach.

7.1 Methodology

The methodology adopted in this chapter utilizes a dynamic price signal to manipulate feeder demand through daily electricity cost minimization at each home. First, historical day-ahead price and weather are used to train a nonlinear predictive model. This model is then used to generate day-ahead prices using typical weather year data. Next, the predicted prices are provided to the MPC framework as a price signal. The controller in each home then attempts to minimize energy cost using this signal and the predicted energy consumption over the planning horizon. Because high day-ahead electricity prices are historically correlated with periods of high electricity demand, minimization of cost is expected to achieve peak demand reductions. The subsections below describe the price modeling and objective function in additional detail.

7.1.1 Day-Ahead Price Modeling

Hourly electricity prices are modeled using a classification and regression tree (CART) for all three feeder locations. Modeling price is necessary because complete weather data (including insolation) for the locations under study are not available. This requires both historical electric price data, and historical weather on an hourly basis. Historical prices from the Independent System Operators (ISOs) managing the electric grids in the three locations can be obtained from the Electric Reliability Council of Texas (ERCOT) [35], New York Independent System Operator (NYISO) [70] and California Independent System Operator (CAISO) [20] websites, while dry bulb temperatures for these locations are easily obtained from [91].

As illustrated in Equation 7.1, the relationship chosen to model price, P is a function of dry bulb temperature, T , deviation of dry bulb from the monthly mean, ΔT , hour of day, H , day of week, D , and binary variable, W , which indicates if the day is a working day. This relationship

is intended to capture both physical and temporal drivers of electricity price so that the price modeled using TMY weather used in this study is consistent with, and representative of, the historical relationship.

$$\log(P) \sim f(T, \Delta T, H, D, W) \quad (7.1)$$

The statistical package R [80] is used to generate CART rules using the full 2013 year of hourly price and dry bulb temperature history for each location. To obtain a sufficiently well fitting model, **cp** (a measure of model complexity) and **minsplit** (the minimum number of observations in a node before splitting is attempted) are tuned by hand. The resulting parameters are summarized in 7.1 below.

Table 7.1: Model parameters for price prediction CART.

cp	minsplit
3	0.001

Plots showing historical prices and those predicted when provided the original training data are shown in Figures 7.1, 7.2 and 7.3. While the extreme values are not captured, the trends between the original price data and the modeled price data are consistent. That is, when ambient temperature is high, prices tend to be high.

Next, the prices for typical weather year data are predicted using the generated CARTs. Figures 7.4, 7.5 and 7.6 show examples of modeled prices. Consistent with the model training results, modeled price is predicted to be low when ambient temperatures are low, and high when ambient temperatures are high.

Price could have been easily modeled as a function of historical demand, and predicted using simulated feeder demand. However, as the prices obtained are representative of distribution hubs with a much larger diversity of loads, it would not be valid to predict price using the residential feeders simulated here.

The method presented here is very simplistic in its assumption that price is determined




Figure 7.1: Historical and modeled CAISO prices using historical weather.




Figure 7.2: Historical and modeled ERCOT prices using historical weather.




Figure 7.3: Historical and modeled NYISO prices using historical weather.




Figure 7.4: Predicted CAISO prices using typical weather.




Figure 7.5: Predicted ERCOT prices using typical weather.



Figure 7.6: Predicted NYISO prices using typical weather.

solely by this small set of parameters. Indeed, hourly electricity price is a function of a great number of parameters not considered here, including fuel price, generator availability, scheduled and unscheduled outages, and system congestion. To be clear, the purpose of this exercise is not to present a perfect price prediction model. This methodology is presented as a way to simply estimate price given typical weather data, recognizing that the use of historical price without correction ignores a primary predictor.

7.1.2 Objective Function Definition

The objective function to be minimized is simply the sum of energy cost over the planning horizon as shown in Equation 7.2. As implemented, this objective function is minimized by the optimal controller in every home of the participating group individually. Optimization parameters, e.g. optimization boundaries, number of decision variables, etc., are identical to those used in Chapter 6 unless otherwise specified.

$$\min \left(\sum_{i=j}^k e_i \cdot c_i \right) \quad (7.2)$$

where:

i is the time start index

j is the cost horizon time step start index

k is the cost horizon time step end index

e_i is the electricity used at time step i

c_i is the cost of electricity at time step i

When cost is uniform over the planning horizon, the controller seeks to minimize total energy use. When presented with artificially large prices, e.g. a price of \$1000.00/kWh during the afternoon hours, the controller will avoid all operation during these times. In cases between, the controller will trade-off additional energy used for precooling against higher costs hours at a subsequent time.

7.2 Methodology Refinement

Initial optimization results using the day-ahead prices lead to three significant refinements in the methodology. The first two described below were used in all optimizations presented in this work. The remaining refinement is adopted for only the cases presented in this chapter. Note that the results presented earlier in Chapter 6 incorporate the first two refinements as well; those optimizations were re-run following the findings shown below.

The first refinement regards the alignment of occupancy and mode boundaries. Initial optimizations assumed occupancy times for all homes to be identical. Aside from being totally unrealistic, this had the effect of aligning the mode boundaries for all homes. This resulted in some very interesting aggregate feeder demand profiles that contained step changes in feeder demand. An example of this can be seen in Figure 7.7. In addition to the step chart appearance, there appeared to be synchronization behavior involving large number of homes, as evidenced by the large swings in demand at each half hour boundary.

By applying a more realistic assumption regarding occupied hours and randomizing the departure times within a ± 1 hour range, the step-change nature of the aggregate demand is eliminated. By shifting the departure time of the occupants, the mode boundaries shift as well, reducing the likelihood that the mode boundaries of one home align with the mode boundaries of another. This change results in a more realistic simulation of aggregate demand. Figure 7.8 shows the results from the same optimization using the randomly shifted departure times. Note that while the step changes have been removed, the synchronization still exists.

An explanation for the synchronization is not obvious at first glance. One clue is in the period of the oscillations created by synchronization. It turns out that the oscillations are on a perfect one-hour frequency, suggesting a driver at this particular frequency. In fact there is: day-ahead price is defined on an hourly basis. Weather data, schedules and all other aspects of the simulation are interpolated to prevent discontinuities at the hour boundary, but price is not.

The explanation for why this causes the oscillations is actually quite simple. The controller,

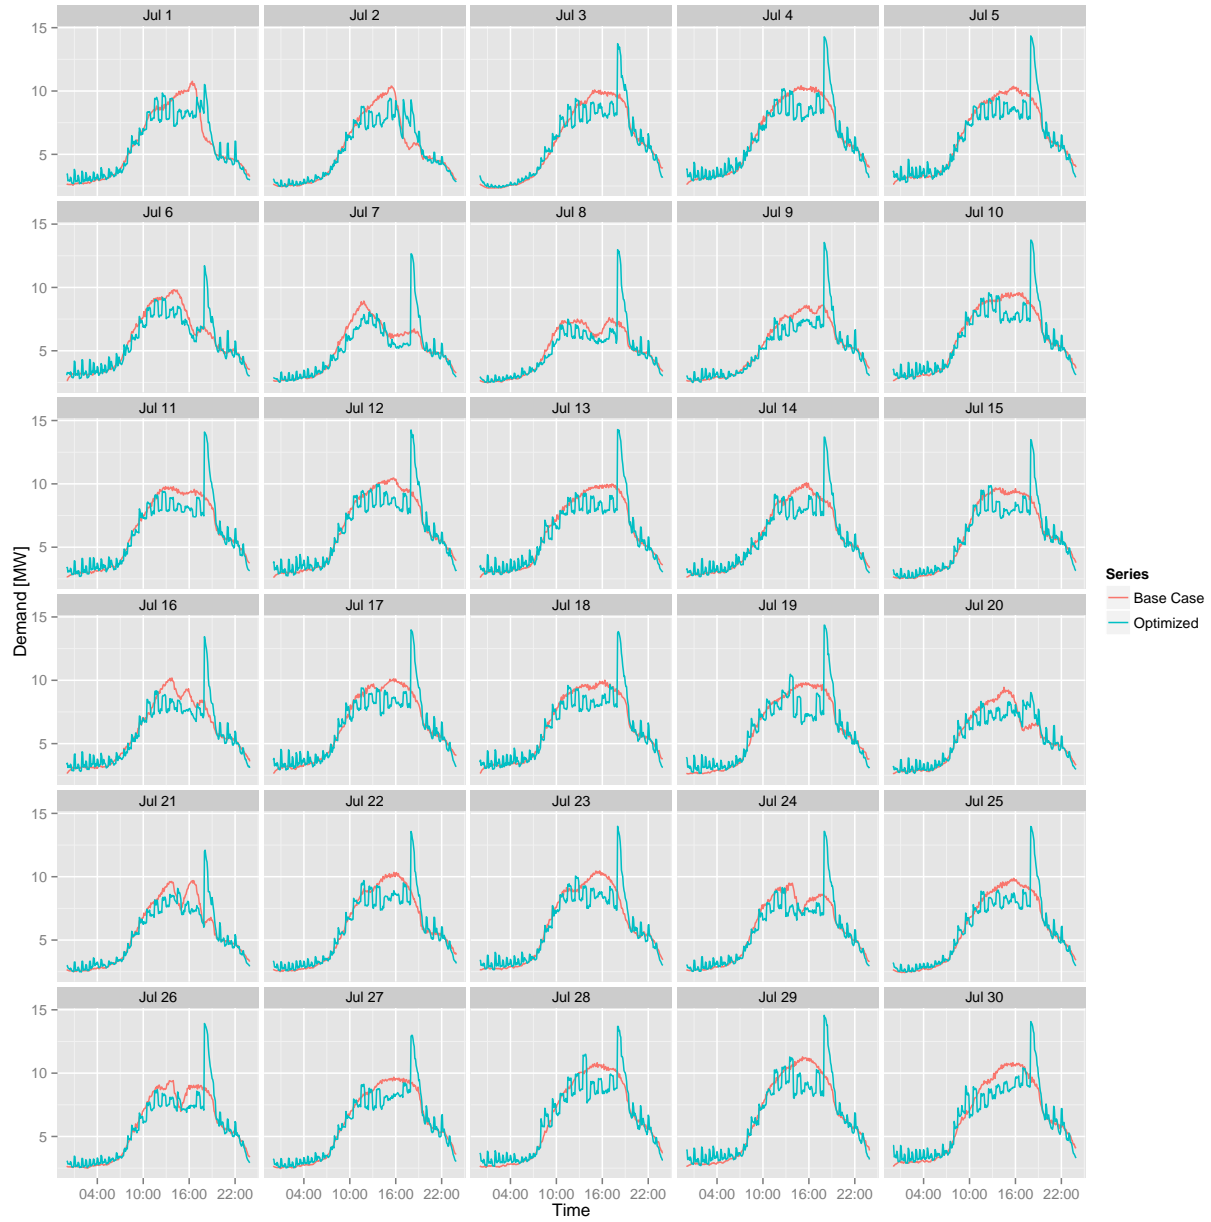


Figure 7.7: Feeder demand profiles for Houston day-ahead price optimization, first test, 70% participation.

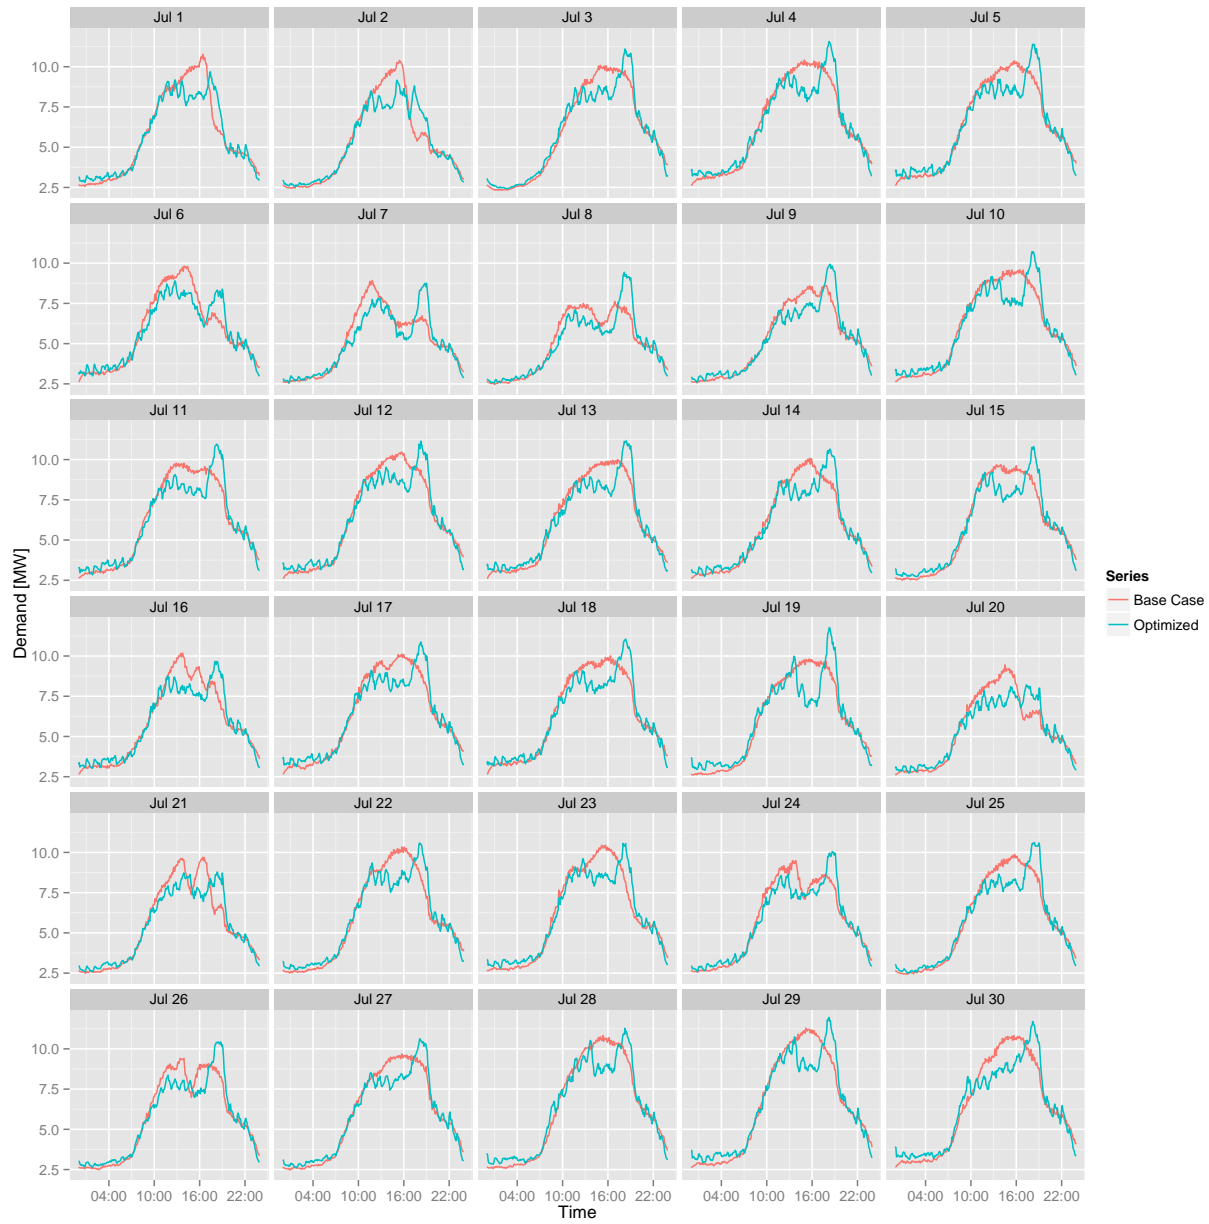


Figure 7.8: Feeder demand profiles for Houston day-ahead price optimization, second test, 70% participation.

seeing a price increase or decrease on the hour boundary adjusts set points in a way that allows it to precool the building in the hour or half hour before a price change, and float through the higher price period. This results in a duty-cycling behavior that, in aggregate, appears to be synchronization between homes. As an experiment, price data is interpolated to a five-minute interval and all 288 values are provided to the objective function to test if this can be eliminated. In fact, it is eliminated, as shown in Figure 7.9.

Interpolation of price data, and more generally the interpolation of all time series signals provided to the controller, is the second refinement to the methodology. While this might strike some as a false solution, it provides some very important guidance for the implementation of price-responsive controllers. That is, large discontinuities in price can result in unintended behavior at the aggregate level, and caution should be exercised in large deployments of price-based controllers without first understanding the larger system-wide ramifications.

The final refinement to the methodology regards the boundaries of the optimization. Specifically, the upper boundary in the hours preceding the arrival of occupants and the return to the occupied cooling set point. Figures 7.7, 7.8 and 7.9 consistently show a large demand spike centered around 18:00 when the upper boundary of the optimization is depressed. This spike in demand is obviously not desirable from a peak demand limiting perspective.

In an attempt to remove the spike, the upper boundary is manipulated to give two variations. Recall from Table 6.1 that the upper boundary is defined as a delta from the normal occupied cooling set point. In the first variation, the upper boundary of the optimization is defined as the occupied cooling set point. This is referred to as the “zero-degree” case. The zero-degree case forces the controller to precool in order to provide relief during high price periods. In the second variation, the upper boundary steps down from the +3K offset by 0.5K per half hour in the three hours preceding the occupants’ return. This is referred to as the “ramp-return” case. Example results from these variations are shown in Figures 7.10 and 7.11.

Simulations using the two variations indicate that the zero-degree case generally results in lower peak demand with a small number of exceptions. The ramp-return case, however, fails to

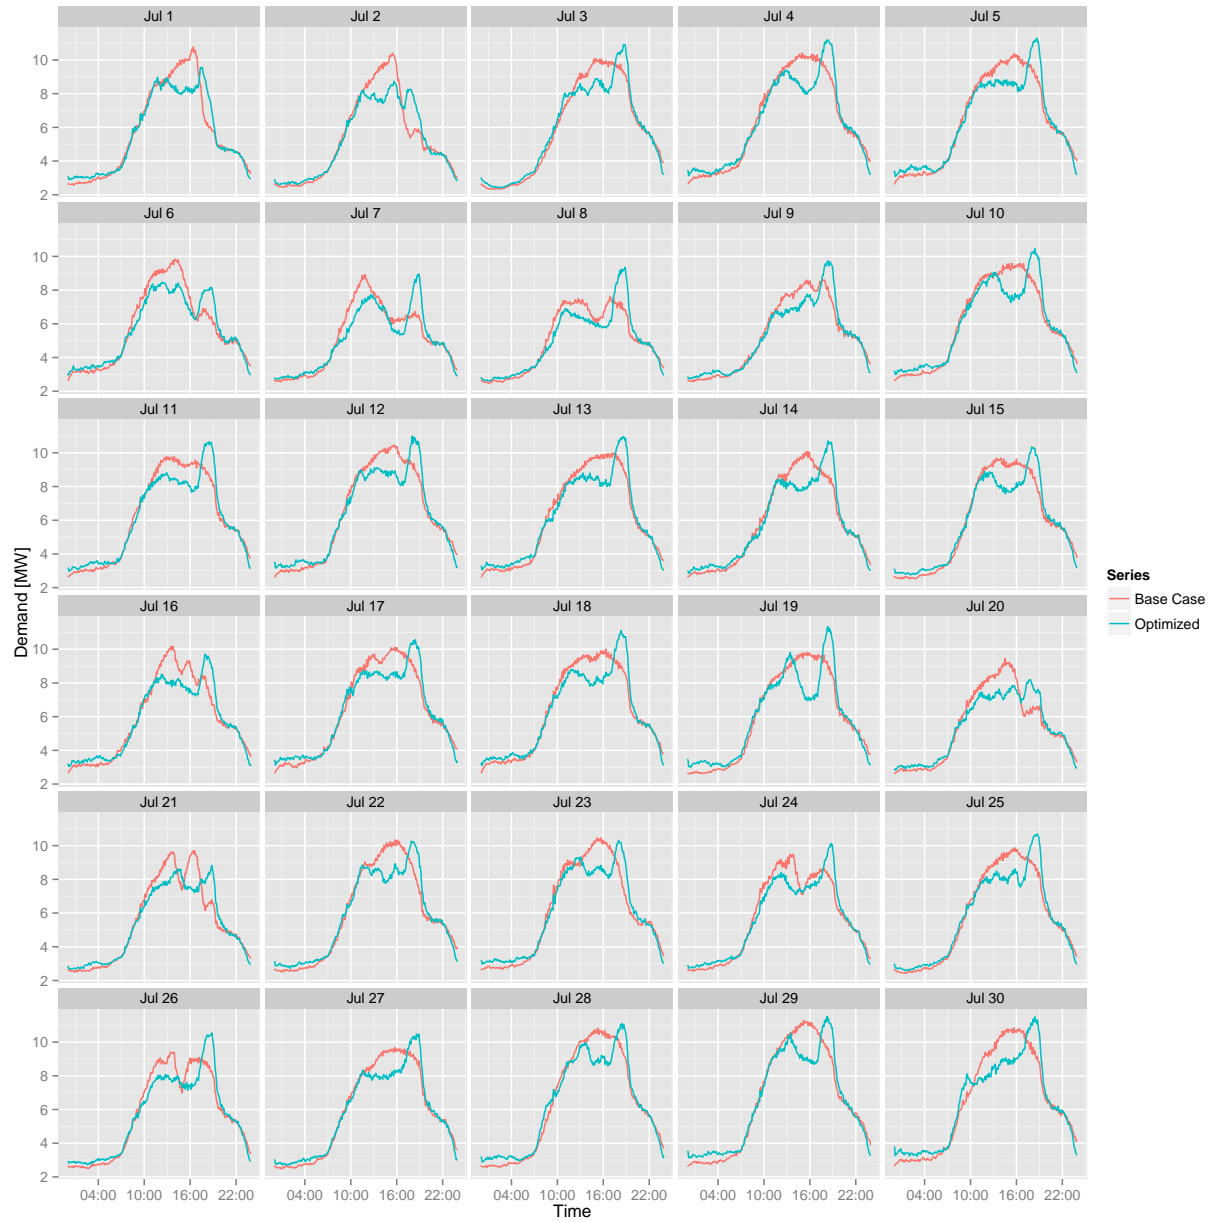


Figure 7.9: Feeder demand profiles for Houston day-ahead price optimization, 70% participation.

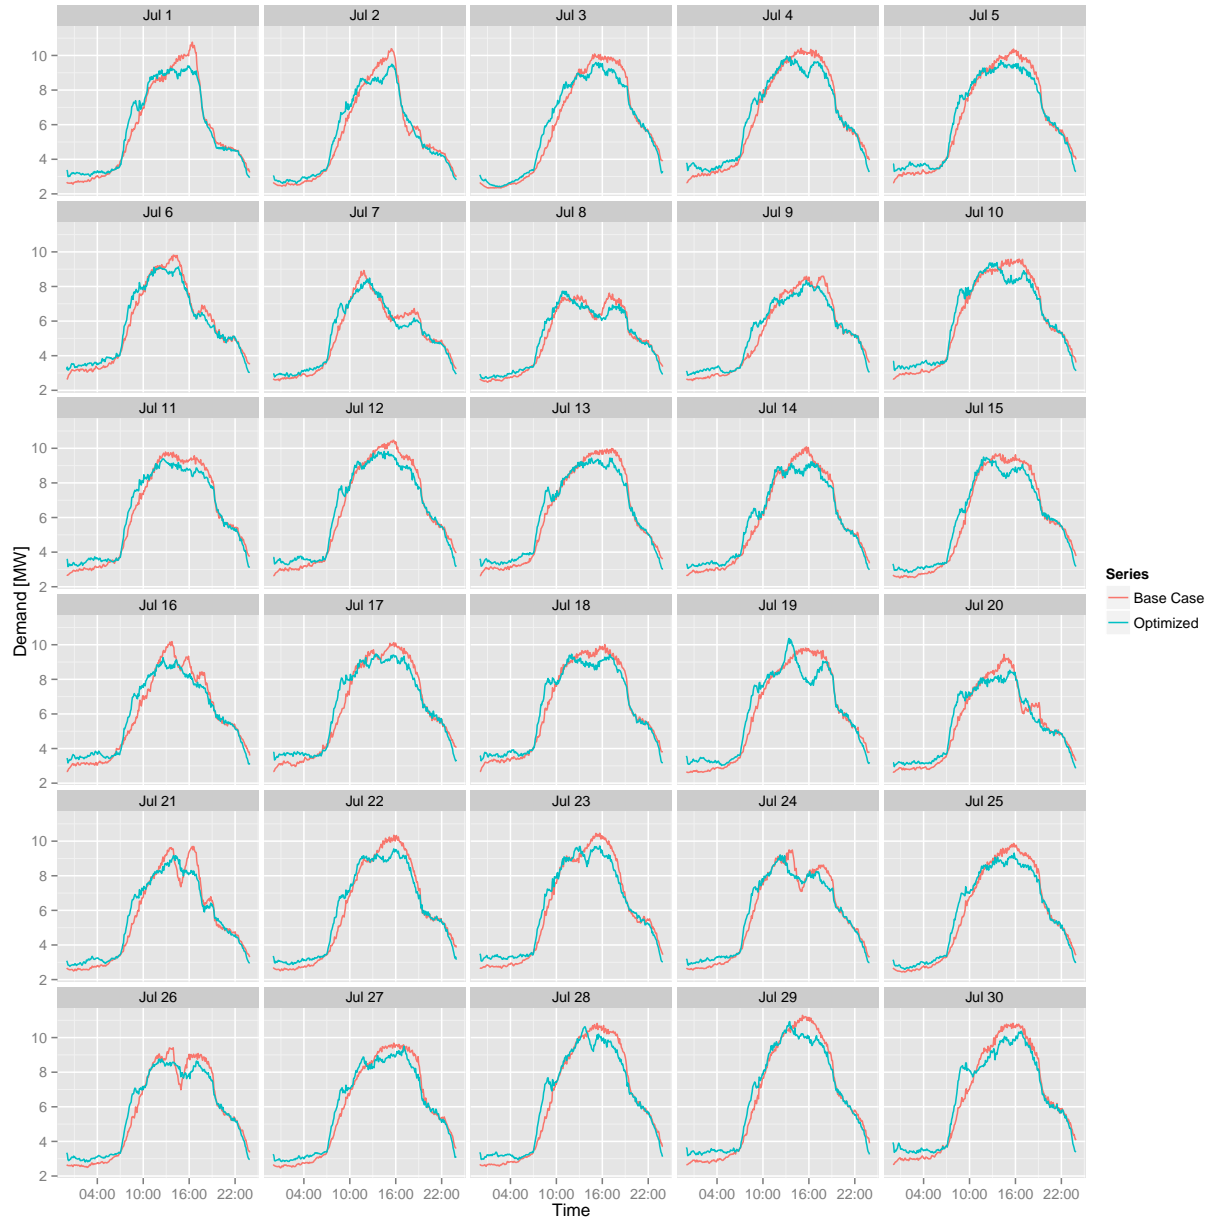


Figure 7.10: Feeder demand profiles for Houston day-ahead price optimization, zero-degree upper boundary case, 70% participation.

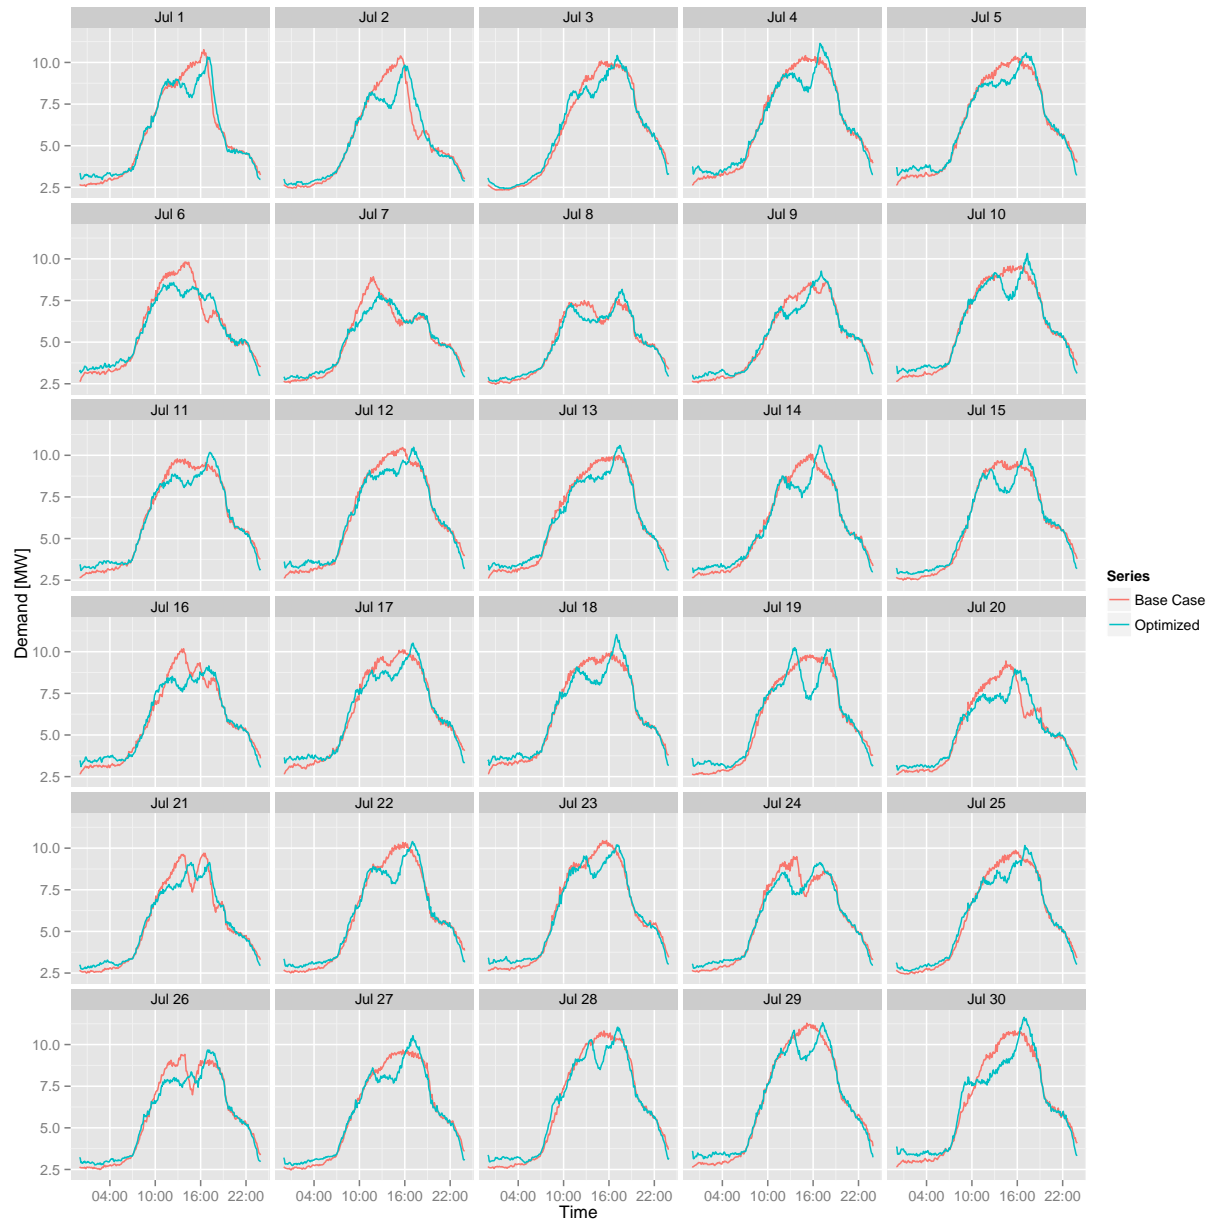


Figure 7.11: Feeder demand profiles for Houston day-ahead price optimization, ramp-return upper boundary case, 70% participation.

eliminate the demand spikes on many days, although the spikes are somewhat attenuated compared to Figure 7.9. Both experiments result in feeder demand curves containing relatively large fluctuations in demand, failing to provide demand smoothing on most days.

7.3 Day-Ahead Price Simulations

Based on the results of the preceding cases, it is not clear which of the upper boundaries should be implemented. Accordingly, three sets of simulations have been performed using the refined methodology: one using the familiar unoccupied cooling boundary set up, and two representing the proposed variations. The results and analysis presented below will reference all three and cases, but charts and tables for Los Angeles and New York reside in Appendix C.

7.3.1 Simulation Results

Referring back to Figures 7.9, 7.10 and 7.11, optimization using the dynamic price approach tends to introduce several features to the feeder demand that are undesirable. Occurrences of increased peak demand, reduced load factor, increased ramping, and increased peak to valley ratio can all be observed in Tables 7.2, 7.3 and 7.4. The metrics tend to be mixed. No variation shows improvements across all metrics, but the zero-degree variation is arguably the best performing of the lot. This variation shows an increase in ramping, but an improved power spectrum distribution (Figure 7.12) over the base case. Power spectrum distributions for the remaining cases (Figures 7.13 and 7.14) indicate degraded performance. At the lower level of participation, (Figure C.2) feeder demand is much less affected and results are generally improved. This is confirmed by Table C.1 and Figure C.1.

Results from the 30% participation rate simulations are consistent with the 70% corollaries, and show features consistent with a smaller participating population, i.e. reduced ability to shape load. These results can be seen in Figures C.2 and C.1 and Table C.1.

Plots of individual days show heavy precooling activity throughout the morning hours. In the zero-degree examples shown by Figures 7.15 and 7.16, optimization succeeds in reducing peak

demand and increasing load factor at the expense of increased ramping. Both examples show a demand spike at approximately 8:00 when the lower cooling set point boundary is reduced. This indicates that in both cases, the controller chooses to depress the set point in preparation of the upcoming high prices later in the afternoon.

Los Angeles and New York results shown in Figures C.8, C.10, C.9, C.27, C.25 and C.26 reflect the findings of the Houston cases. Metrics and spectral densities favor the zero-degree variation over the other two. Consistent with observations in previous chapters, New York and Los Angeles both show little improvement in performance metrics.

Table 7.2: Performance metrics for Houston feeder day-ahead price optimization, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	-2.01	-2.98	-0.82
Peak Demand [MW]	0.36	-1.68	1.74
Peak to Valley [%]	92.95	77.21	116.46
Load Factor [%]	-2.85	-13.32	8.89
Ramp [MW]	3.09	-2.94	8.79

Table 7.3: Performance metrics for Houston feeder day-ahead price optimization, zero-degree upper boundary case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	1.12	0.50	1.91
Peak Demand [MW]	-0.52	-1.26	0.51
Peak to Valley [%]	82.81	75.13	94.49
Load Factor [%]	3.80	-2.32	7.78
Ramp [MW]	0.72	-2.91	4.94

Table 7.4: Performance metrics for Houston feeder day-ahead price optimization, ramp-return upper boundary case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	-0.99	-1.66	0.25
Peak Demand [MW]	0.12	-1.29	1.06
Peak to Valley [%]	88.85	76.81	100.62
Load Factor [%]	-1.05	-6.70	8.25
Ramp [MW]	1.93	-2.89	7.22




Figure 7.12: Total spectral power as a function of frequency bin for Houston feeder day-ahead price optimization, zero-degree upper boundary case, 70% participation.




Figure 7.13: Total spectral power as a function of frequency bin for Houston feeder day-ahead price optimization, 70% participation.




Figure 7.14: Total spectral power as a function of frequency bin for Houston feeder day-ahead price optimization, ramp-return upper boundary case, 70% participation.




Figure 7.15: Feeder demand profiles for Houston, July 1 day-ahead price optimization, zero-degree upper boundary case, 70% participation.




Figure 7.16: Feeder demand profiles for Houston, July 26 day-ahead price optimization, zero-degree upper boundary case, 70% participation.

7.3.2 Summary and Conclusions

Feeder demand plots and metrics vary widely across the three upper boundaries investigated in this chapter. However, one common theme does emerge consistently for many of the days investigated, which can be explained by the shape of the day-ahead price curves shown earlier. Specifically, the relationship between price and demand is highly nonlinear. Day-ahead price curves tend to show very high peak prices in late afternoon that are many times greater than prices earlier in the day. Typically, there is one single hour, and at most two, that is extremely high, surrounded by relatively low prices. By contrast, the shape of the feeder demand more closely approximates a smooth bell curve on most days. When subjected to the extreme prices in the early afternoon, the controller shifts a large portion of demand to the hours immediately before and after the price spike, resulting in a large dip in demand. Across thousands of homes, all responding to the same extreme price during one or two hours, this results in a dip in the feeder demand during the high priced hour(s) and demand spikes on either sides. Despite a relatively extensive set of results and several refinements to methodology, the efficacy of adopting a price responsive control strategy is questionable.

7.4 Synthetic Price Simulations

Introducing a penalty term into the objective function or adopting a multi-objective approach may avoid the spikes observed previously. However, the observation that the controller responds principally to one or two hours of high price suggests a different solution: provide the controller with a synthetic price signal that is both smooth and linearly related to demand. The rationale is that a smooth price reflecting the actual feeder demand will result in an optimized feeder demand that tends to flatten demand without spikes or oscillations.

The next issue to be resolved is the generation of a synthetic price signal. An obvious candidate is the base case feeder demand itself. To test the proposed hypothesis, a price curve is constructed by simply taking the feeder demand (measured in megawatts) and applying a two-sided

moving average to smooth the high frequency oscillations. A 30-minute moving average provides a good degree of smoothing without removing the important features of the demand curve; this is the value used to generate the synthetic price curves used in this experiment.

The following results were generated using the synthetic price signal and the two upper boundary assumptions described in the previous section. Because of the large number of plots and tables, only a handful are presented here. The remaining results are included in Appendix C.

7.4.1 Simulation Results

Like the previous day-ahead results, the synthetic price optimizations show significant fluctuations in demand as the controller avoids high prices. This is particularly evident in the Houston case due to the number of spikes and troughs present in the base case demand curve used to generate the synthetic price curve. These fluctuations can be seen in Figures 7.20 and 7.21. While the zero-degree variation outperforms the ramp-return variation, neither show particularly desirable performance metrics in Tables 7.5 and 7.7. The distribution of spectral power in Figure 7.17 shows the zero-degree variant able to increase the DC power component while the ramp-return variant shows the opposite in Figure 7.18. The zero-degree variant at 30% participation shows the greatest consistency in metrics, with improvements in nearly all areas. Results for this case can be seen in Figures 7.22 and 7.19 and Table 7.7.

Similar to the day-ahead scenarios, observations made regarding the Houston simulations apply to both the Los Angeles and New York simulations, with the zero-degree variant showing the best performance. In the Los Angeles zero-degree case, the size of the fluctuations are much smaller, consistent with the amount of flexible cooling demand available. In the ramp-return cases, controller decisions tend to result in a delay of cooling demand until after the high price periods rather than precooling.

Table 7.5: Performance metrics for Houston feeder synthetic price optimization, zero-degree upper boundary case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	0.23	-0.26	0.86
Peak Demand [MW]	-0.39	-1.03	0.20
Peak to Valley [%]	83.77	75.66	93.36
Load Factor [%]	2.57	-1.62	7.93
Ramp [MW]	1.25	-1.99	2.95

Table 7.6: Performance metrics for Houston feeder synthetic price optimization, ramp-return upper boundary case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	-2.08	-3.17	-0.82
Peak Demand [MW]	0.36	-1.10	1.21
Peak to Valley [%]	91.17	81.85	105.89
Load Factor [%]	-2.94	-8.19	6.81
Ramp [MW]	2.08	-1.19	4.55

Table 7.7: Performance metrics for Houston feeder synthetic price optimization, zero-degree case, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	0.10	-0.10	0.38
Peak Demand [MW]	-0.33	-0.69	-0.07
Peak to Valley [%]	88.00	81.50	96.62
Load Factor [%]	2.16	0.47	4.30
Ramp [MW]	-0.75	-3.99	0.41




Figure 7.17: Total spectral power as a function of frequency bin for Houston feeder synthetic price optimization, zero-degree upper boundary case, 70% participation.




Figure 7.18: Total spectral power as a function of frequency bin for Houston feeder synthetic price optimization, ramp-return upper boundary case, 70% participation.




Figure 7.19: Total spectral power as a function of frequency bin for Houston feeder synthetic price optimization, zero-degree upper boundary case, 30% participation.

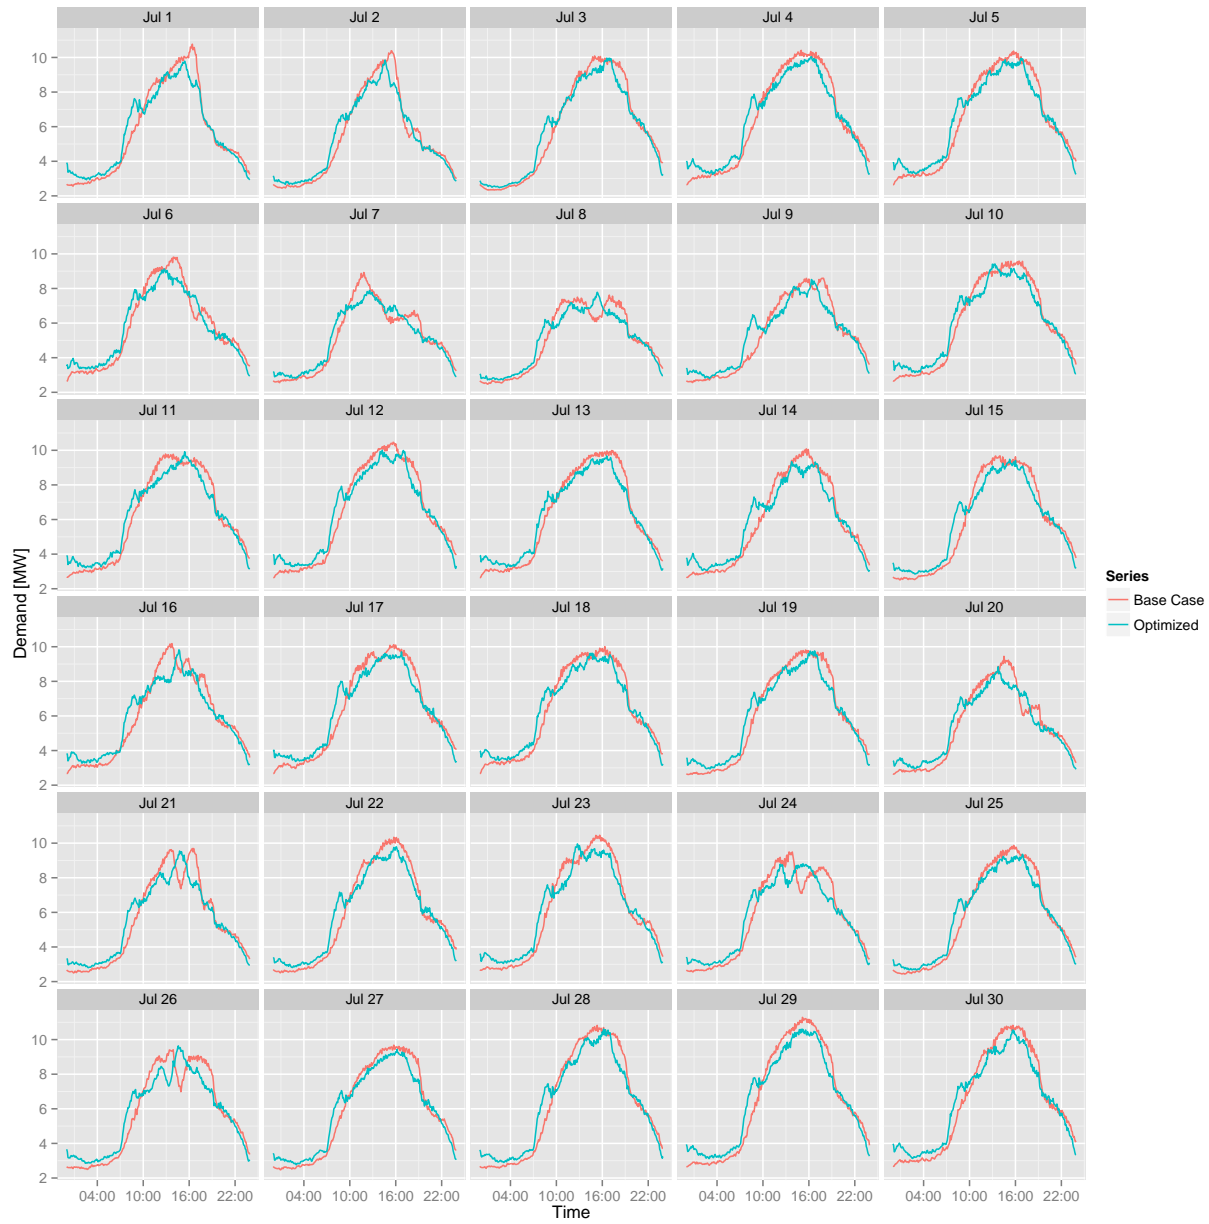


Figure 7.20: Feeder demand profiles for Houston synthetic price optimization, zero-degree upper boundary case, 70% participation.

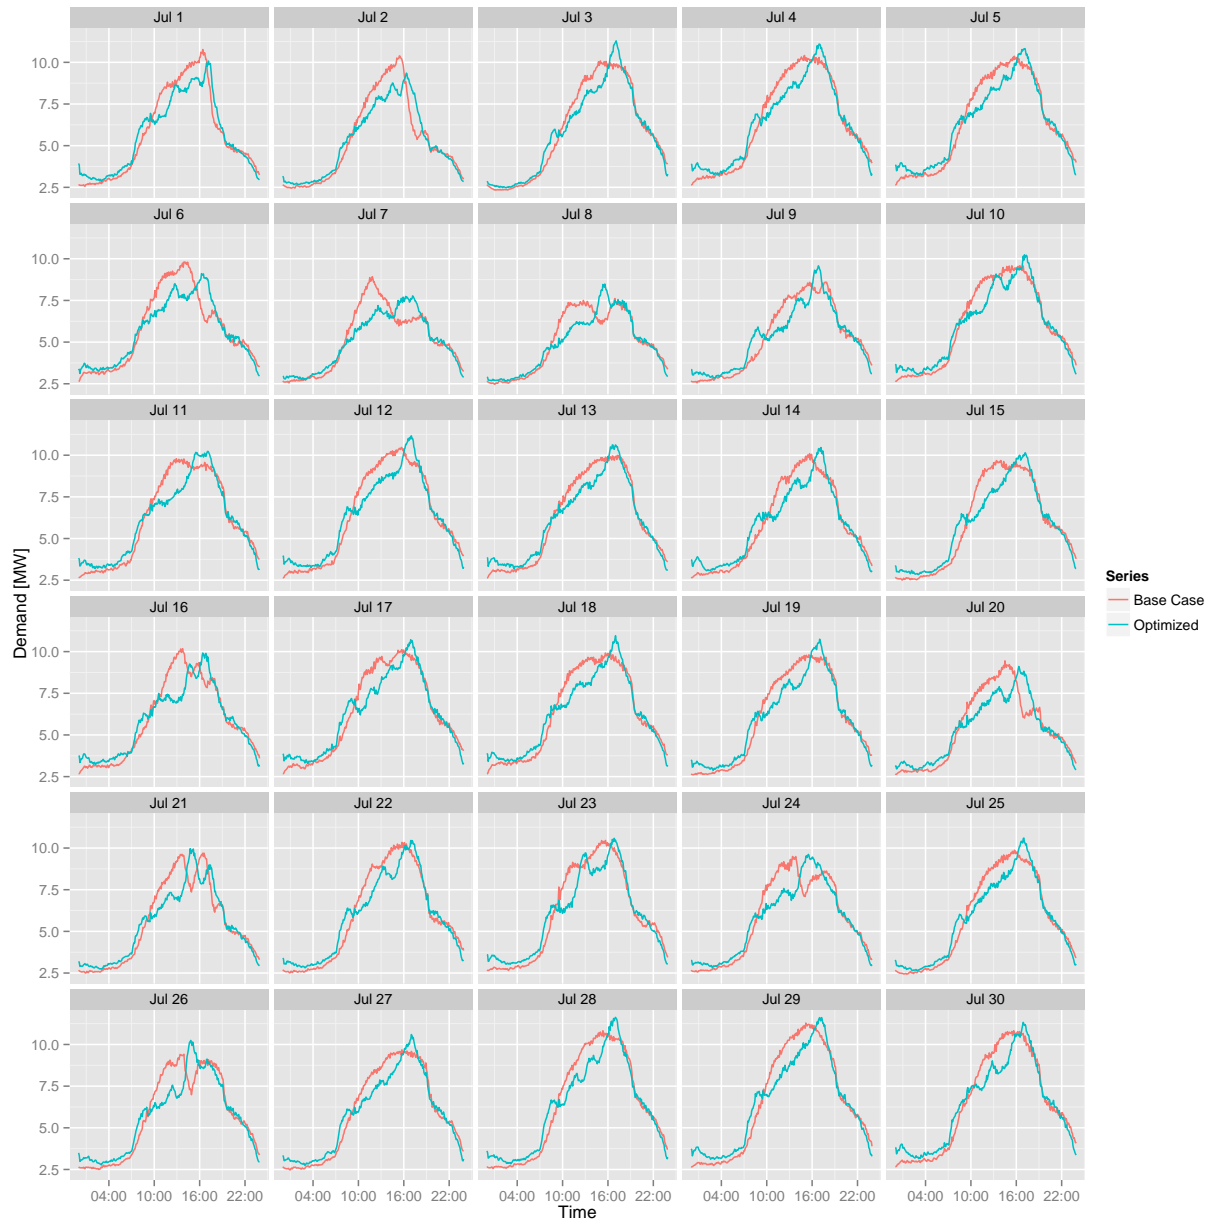


Figure 7.21: Feeder demand profiles for Houston synthetic price optimization, ramp-return upper boundary case, 70% participation.

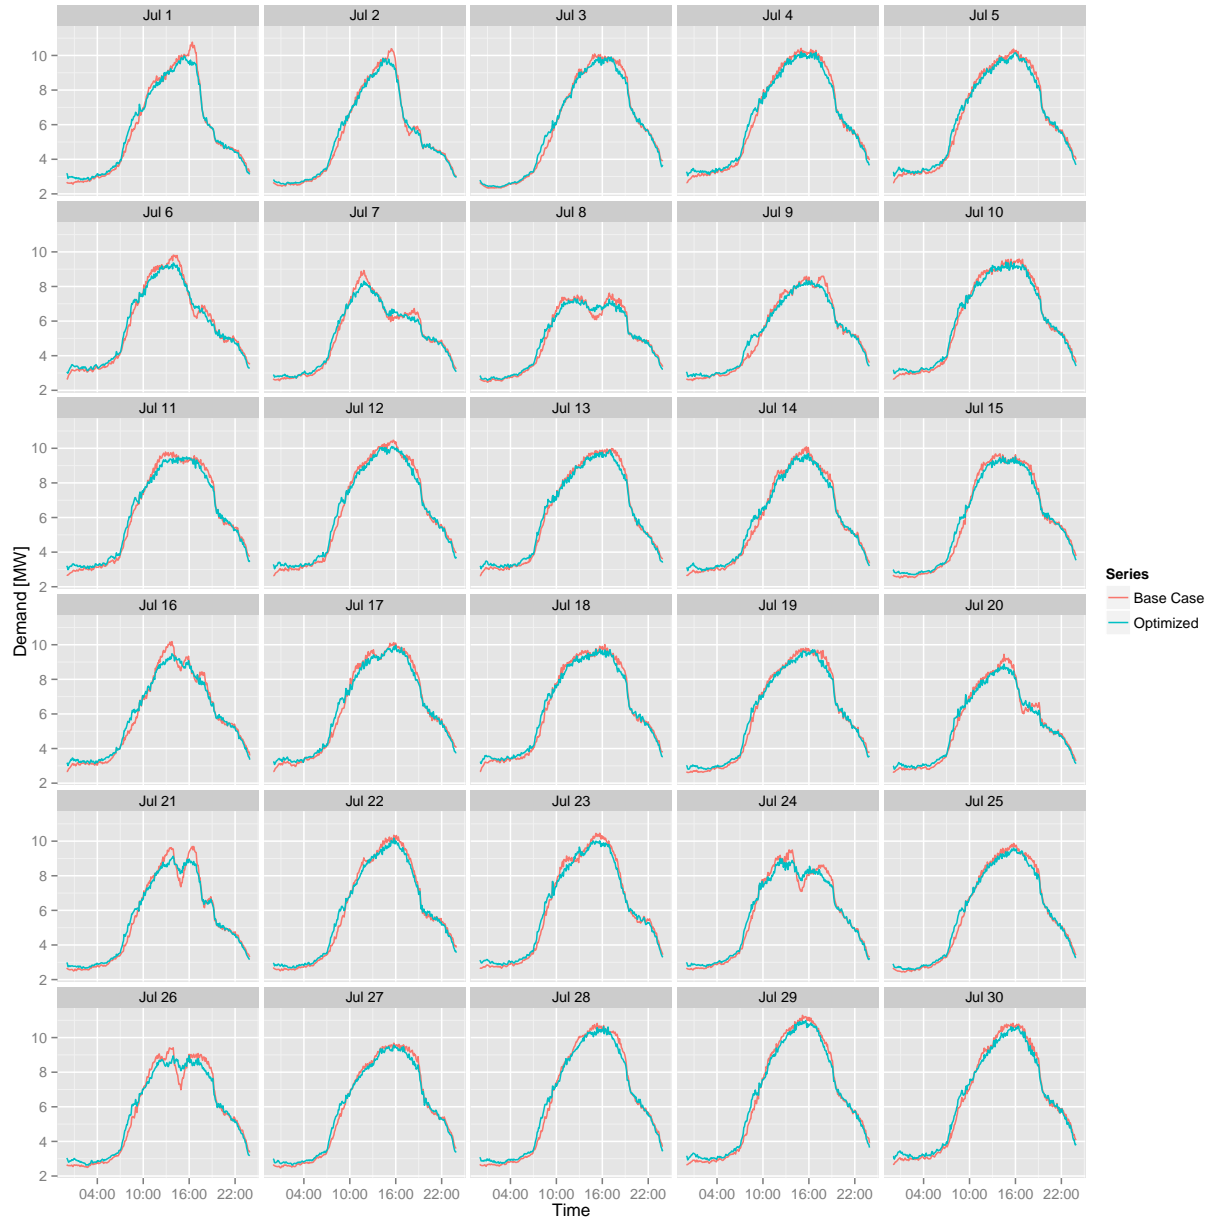


Figure 7.22: Feeder demand profiles for Houston synthetic price optimization, zero-degree upper boundary case, 30% participation.

7.4.2 Summary and Conclusions

Two key observations can be made regarding the results of this experiment. First, despite being provided a smooth price curve, the synthetic price optimization does not necessarily result in a smooth response. Those cases in which the controller is allowed to raise the cooling set point during the unoccupied period above the occupied set point force a return to the occupied set point resulting in a spike; this can not be easily avoided by the method used to formulate the price curve. Those cases in which the upper boundary is the occupied set point for all hours of the day still exhibit demand spikes at the controller attempts to avoid higher price periods.

Second, in those days that contain price spikes not removed by the moving average smoothing, the resulting demand curve contains new spikes preceding or following the original as the controller avoids the high price. This sometimes results in a disproportionate response, i.e. a demand spike larger than the spike being avoided. These days tend to have multiple sequential demand spikes that give the appearance of oscillations. Although it is not clear why these features form in the optimized feeder demand curve, a more controlled experiment can be conducted to observe this phenomenon.

7.5 Demand Oscillations

The appearance of demand spikes in optimized results, even when provided a smooth price curve is problematic. This result has important implications for large scale deployment, as these features would almost certainly result in instabilities in the electric grid. The following experiment examines the formation of these features by iteratively optimizing a single day, July 1, using the base case feeder demand curve from New York. The optimization parameters are the same as in previous cases, using the upper bound defined in Table 6.1.

The experiment begins by performing the same energy cost minimization, using the base case feeder demand as the price curve. The return of the cooling set point upper boundary to the occupied temperature results in a demand spike centered around 18:00 as expected. In the next

iteration, this optimized feeder demand curve is fed back into the optimization as a new synthetic price signal. The result is then fed back into the optimization in the next iteration. This process is repeated several more times. Results of this experiment are shown in Figure 7.23

The demand spike created by the first iteration results in a controller adjustment that avoids it. In iteration two, this appears as a demand spike preceding the spike from iteration one. It is also worth noting that small oscillations form in the beginning of the day as well; their origin can be seen in a small demand bump at 8:00 in iteration one. By iteration three, the original demand spike has returned, but is larger than the spike in iteration one. Also interesting is the series of oscillations seen after 20:00. Features become magnified in each iteration, resulting in very large oscillations after five iterations.

One cause for this behavior is likely related to the thermal efficiency of the building envelope. When the controller shifts cooling to an earlier time, it must generally provide more cooling in total due to gains from the environment. If the price at the earlier time is lower, the controller will be able to provide the additional cooling required without an increase in cost. This can result in demand spikes larger than the spikes being avoided. Similarly, when the controller shifts cooling to a later, less expensive time, it generally results in a larger cooling demand as the cooling system has to provide cooling to cover the cooling avoided previously, the cooling required at the new time, and the losses incurred along the way. Over time, this results in the increasing magnitude of the oscillations seen in the experiment.

Another cause for this behavior is due to the formulation of the objective. The overarching goal of this work is show how MPC can be used to shape load in a beneficial way, e.g. to reduce peak demand, reduce ramping (oscillations), increase load factor, etc. A price-based optimization that only considers energy cost does not consider any of these objectives. The controller is blind to these metrics, and if lowest-cost strategy results in the formation of new demand spikes, it will be adopted regardless of problems this introduces. The original hypothesis set forth in this chapter is that dynamic price can be used to reduce peak demand. Although this particular set of simulations do not prove the hypothesis, they fail to reject it. That said, providing a dynamic price signal,




Figure 7.23: Oscillations in demand profile created by iteratively supplying feeder demand as electricity price to optimization.

optimizing for lowest cost and hoping for a predictable response that satisfies these goals is most likely the wrong approach.

This last observation suggests that perhaps a more fruitful approach would be to attack the problem from the opposite direction. That is, instead of trying to illicit the desired feeder demand shape from an arbitrary signal and objective function combination, why not specify the desired feeder demand as the objective explicitly. This worked relatively well for the demand limiting optimization: the goal was to reduce peak demand and the objective function was formulated to do so. Recognizing the potential need to *increase* demand in a predictable way at certain times requires a more generic methodology than that used previously. The development and demonstration of this new methodology is the subject of the next chapter.

Chapter 8

Load Shape Optimization

Previous chapters introduced several approaches to shaping the demand of a distribution feeder using a decentralized model predictive control methodology. Demand limiting optimization resulted in improvements to key performance metrics, reduced peak demand, reduced ramping, and increased load factor in a consistent and predictable way. In contrast, dynamic pricing optimizations resulted in mixed results and unpredictable demand shapes. The success of the demand limiting optimization is in part due to the objective function, which minimized demand directly at the home level to achieve similar results at the feeder level. This chapter introduces a novel new demand shaping methodology for reducing peak demand and intra-day variability using a similar approach in which the objective of the optimization is to drive demand at each home towards a desired feeder demand shape.

This chapter first describes the general approach, followed by the development of the signal that is provided to the controller during optimization. The chapter then shows an application of the new methodology to the three feeder models to demonstrate its effectiveness. A brief discussion of the methodology's advantages and the applications it enables concludes the chapter and points the way towards the final investigations in Chapters 9 and 10.

8.1 Methodology

The idea behind the methodology proposed in this chapter is very simple: tell the homes how to increase or decrease demand in order to shape the feeder demand curve. As with the demand

limiting optimization, this approach relies on each home making incremental adjustments to their demand to achieve a common goal. The goal, i.e. the specific shape of the feeder demand curve desired, is termed the feeder “reference demand curve”. This curve contains features telling homes where additional demand is needed or where demand should be reduced relative to the base case feeder demand.

This is somewhat different from the previous optimizations methods in concept, as it may be used to create additional demand where needed, rather than simply reduce demand where it is not. Whereas the previous methods simply reduced demand at specific time of day — either by limiting demand, or by using a price-based proxy — this method allows more control over the specific shape of the demand by defining the desired demand curve explicitly. The proposed method follows three steps described in the subsections below. These are:

1. Generate a reference demand curve that represents the desired aggregate feeder demand
2. Transform the feeder reference demand curve into a reference demand curve for each house
3. Minimize the difference between the house demand curve and house reference demand curve

8.1.1 Generation of Feeder Reference Demand Curve

To illicit a predictable response, the author proposes that the feeder reference demand curve should reflect the limitations of the underlying physical processes in the system. That is, a reference curve should not contain features that are unachievable by the system. For example, with 2MW of flexible demand at any given moment, the optimization, no matter how effective, may not be able to eliminate all demand on the feeder if the demand is greater than 2MW. Nor may it eliminate all 2MW for an extended period of time because of comfort requirements. Therefor, the feeder reference demand curve must reflect two important constraints of the system.

The first constraint that should satisfied is the first law of thermodynamics. The interpretation adopted here is somewhat less formal, but guides the construction of the reference demand curve: the assumption is that in both the base case and the optimized case, the total area under the

demand curve should be equivalent. Obviously this can not strictly be true due to losses related to thermal storage efficiency, part-load efficiency, and the partial deferral of load to subsequent hours outside of the planning horizon. Nevertheless, the areas should be approximately the same, which is a condition easily addressed by normalization. Normalization results in a reference demand curve with approximately the same amount of cooling related demand as the base case. Note that through normalization at the house level, the same condition can be enforced.

Second, the amount of load reduced or increased at any given time is limited by the available flexible cooling demand. That is, cooling demand may only be reduced by the amount that is active. Similarly, demand may only be increased as much as the amount of cooling demand not already active. Although not designed explicitly to do so, the experiments in Chapter 6 give an indication of the size of the flexible cooling demand. Additional experiments can be conducted measure flexible cooling demand directly. However, because optimization enforces upper and lower set point boundaries, the exact amount of flexible demand is much more difficult to determine, since the flexible demand available at a given time step is a function of HVAC operation in preceding hours. Enforcing this constraint at the feeder level does not enforce the constraint at each home. Therefore, this constraint is treated as a soft constraint that merely guides the formation of the feeder reference demand curve.

It is not entirely clear from the constraints what method should be chosen to generate the feeder demand curve. The author proposes that the desired outcome should inform the choice. In practice, the choice of technique may be specific to the needs of the utility or demand aggregator on any given day.¹ Here, the goal of the optimization is to shape the feeder demand curve to have a lower peak and reduced variability, effectively shaving peak load in order to fill-in troughs, and removing oscillations in demand. This implies that a smoothing algorithm applied to the base case feeder demand curve should be used to generate the reference demand curve.

This approach is advantageous for several reasons. First, generating a reference demand

¹ Exploring a set of different possible techniques is beyond the scope of this work.

curve with a smoothing algorithm is relatively simple and easily automated. Second a smoothing algorithm such as a simple moving average addresses both constraints automatically,² without requiring them to be explicitly enforced. Finally, and perhaps most importantly, the method can be used to study a wide range of base case scenarios in which both demand and supply could be highly variable, for example, in systems with wind and solar generation.

For the studies presented in this and following chapters, a simple moving average is used to generate the reference feeder demand curve. The choice of period was chosen to be four hours in order to reduce short-term variation in feeder demand, while still creating realistic demand targets. This determination is somewhat subjective. That said, initial experiments performed with several values suggested that this value results in an achievable demand for many of the cases considered.

The method adopted here is one of many, and likely not the ‘best’ for all scenarios. Rather, it is simply one alternative that attempts to address the two constraints, while driving feeder demand towards lower peaks and reduced variability. An interesting question that arises is whether or not the shape of this curve really matters at all. An experiment, testing an alternative “zero demand” reference demand curve intended to explore this question, is discussed in Appendix H.

8.1.2 Generation of House Reference Demand Curve

A new curve representing the deviation between base case and reference feeder demands can be useful for guiding optimization at each home. This curve is referred to as the ‘reference signal’. This curve is defined as the normalized feeder base case demand curve, minus the normalized feeder reference demand curve. Curves are normalized by dividing the demand at each time step by the sum of demand over the cost horizon. The reference signal can contain both positive and negative values, corresponding with increased and decreased reference demands. An example of the feeder base case demand, reference demand and reference signal curves for the Houston feeder on July

² This is a matter of interpretation and not strictly true all of the time due to end-effects related smoothing. However, it is fairly easy to show that 1) the area under a smoothed demand curve is very nearly that of the original, and 2) a smoothed curve contains less variability than the original, implying that the new demand can be met if flexible cooling load is available.

26th is shown in Figure 8.1.




Figure 8.1: Example of feeder reference demand and reference signal created from simple moving average of base case feeder demand profile.

At the beginning of optimization, each home is provided the reference signal. To generate the *house* reference demand curve, house demand is first smoothed using a simple moving average.³ The smoothed house demand curve is then normalized by the sum of demand over the cost horizon. The reference signal is added to the normalized house demand to produce the normalized house reference demand. This is the demand curve which the optimization attempts to produce.

8.1.3 Objective Function Definition

The house reference demand curve is generated once for each home at the beginning of the optimization. At each iteration, house demand is smoothed and normalized using the same moving average period. The optimization objective can be formulated as a minimization of deviation

³ Recall that cooling system staging results in discrete cooling capacities that are either on or off. The time-averaging process results in a much smoother demand curve compared to the square-wave curve produced by HVAC cycling. This allows the controller to try to match the relatively smooth reference curve.

between candidate and house reference demand curves. Specifically, the objective function to be minimized is the sum of squared error between the house reference and candidate demand curves:

$$\min \left(\sum_{i=j}^k (p'_i - p_i)^2 \right) \quad (8.1)$$

where:

i is the time step index

j is the cost horizon time step start index

k is the cost horizon time step end index

p'_i is the normalized reference demand at time step i

p_i is the normalized electric demand at time step i

Similar to the experiment in Section 6, the effect of simple moving average period length is tested to arrive at the appropriate value. In this case, only 120 minutes is tested to confirm that oscillating behavior observed in that test is still present. Based on this result, the decision is made to keep the 60 minute moving average period for all simulations moving forward. Results of the 120 minute moving average period are included in Appendix D, Figures D.2, D.4 and D.7.

8.2 Houston

Metrics from Houston optimizations, listed in Table 8.1, show fairly consistent improvement across all dimensions. Compared to the results in Chapter 6, improvements over the base case are not quite as large. It should be noted as well that in these results, peak demand does increase in a small number of cases, e.g. July 3. The reason for this is unclear as the reference curve for this day (not shown) is lower than the base case during the peak hours. Still, the approach shows promise in achieving the intended result.

As with previous investigations, overall electricity consumption increases as a result of optimization. Although the reference demand curve in theory maintains the total area under the original, thus enforcing the same cooling related electricity consumption, this is not necessarily

achievable. This is consistent with findings regarding thermal storage efficiency which suggest that precooling incurs a penalty due to losses resulting in a total cooling energy increase. Changes to the reference demand curve may minimize the additional energy required, but are beyond the scope of this investigation.

The power spectrum distribution (Figure 8.2) shows the greatest reduction of power at all frequencies in all results seen thus far. This result is intuitive: given a smooth reference demand curve, the controller will attempt to generate a smooth response. It is somewhat amazing, however, that this emerges from the individual actions of thousands of homes which are not coordinated, but instead directed to a common goal.

Table 8.1: Performance metrics for Houston feeder load shaping optimization, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	4.87	4.37	5.25
Peak Demand [MW]	-0.15	-0.87	0.14
Peak to Valley [%]	84.26	75.70	95.57
Load Factor [%]	3.02	1.06	6.87
Ramp [MW]	-0.23	-3.46	1.84

Table 8.2: Performance metrics for Houston feeder load shaping optimization, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	2.07	1.79	2.27
Peak Demand [MW]	-0.09	-0.54	0.06
Peak to Valley [%]	91.47	88.07	98.16
Load Factor [%]	1.40	0.52	3.77
Ramp [MW]	-0.59	-2.03	0.14

Viewed across all days in July (Figure 8.4), the optimized feeder demand appears to be smoother and contain fewer oscillations than all previous results. The 30% participation case (Figure 8.5) shows a limited ability to affect feeder demand as expected, but metrics and spectral power analysis still show improvements over the base case.

Plots of individual days such as those shown in Figures 8.6 and 8.7 illustrate the smoothing




Figure 8.2: Total spectral power as a function of frequency bin for Houston feeder load shape optimization, 70% participation.




Figure 8.3: Total spectral power as a function of frequency bin for Houston feeder load shape optimization, 30% participation.

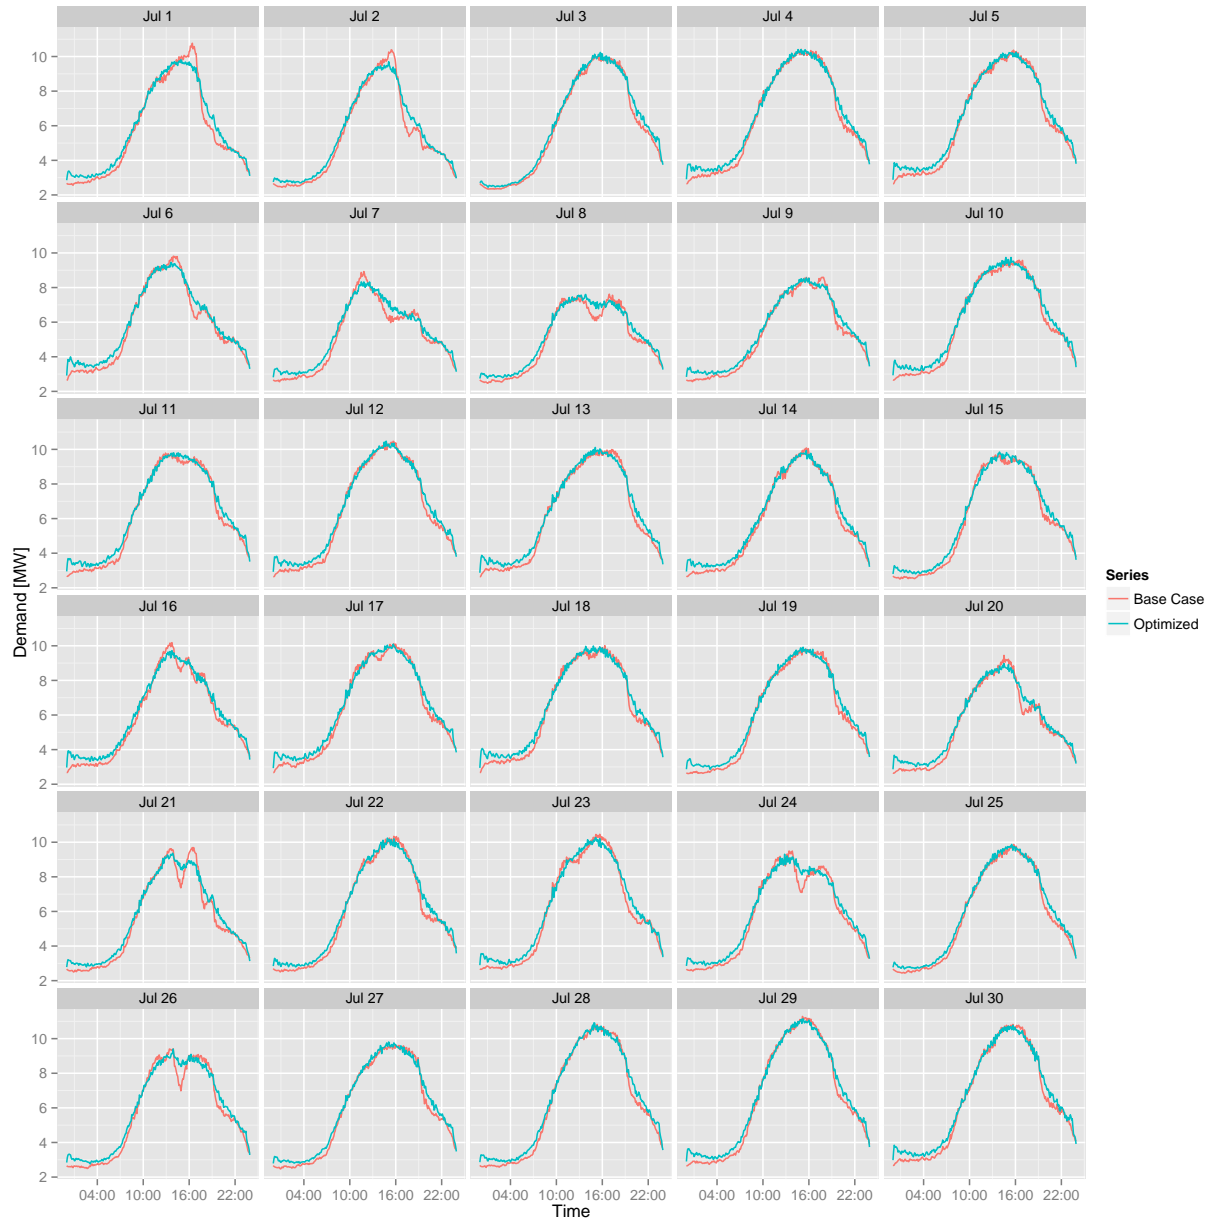


Figure 8.4: Feeder demand profiles for Houston load shape optimization, 70% participation.

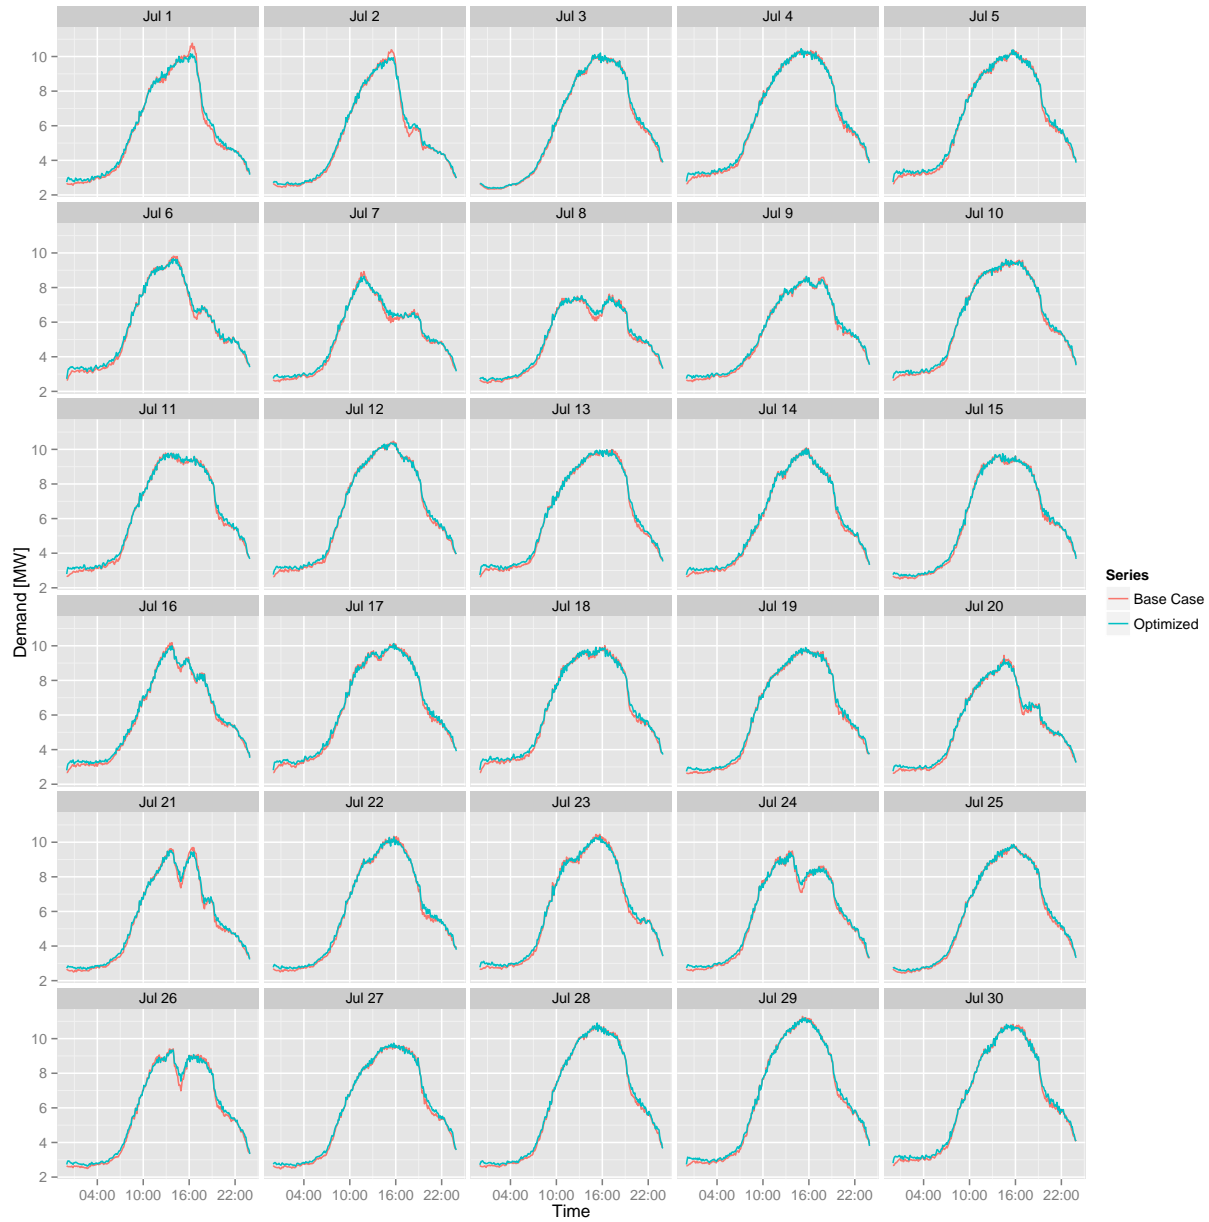


Figure 8.5: Feeder demand profiles for Houston load shape optimization, 30% participation.

effect of the optimization in greater detail. Both days begin with higher demand relative to the base case as the controller depresses the cooling set point. The set point is then raised to avoid the peak demand of the base case, then raised to fill-in the trough that follows. In both days, the controller manipulates the cooling set point further to shave a secondary peak around 19:00.




Figure 8.6: Feeder demand profiles for Houston, July 2 load shape optimization, 70% participation.




Figure 8.7: Feeder demand profiles for Houston, July 7 load shape optimization, 70% participation.

8.3 Los Angeles

As underwhelming as the Los Angeles results appear to be from a visual inspection of the demand curves in Figure 8.9 and the metrics in Table 8.3, the power spectrum analysis in Figure 8.8 tells a slightly different story. That is, despite the controller’s inability to significantly improve these metrics, there appears to be some benefit provided by the controller with regards to smoothing the feeder load shape. The DC power spectrum component is increased, while all other frequency bins are decreased. The 30% participation results included in Appendix D show improvements in this regard as well.

Table 8.3: Performance metrics for Los Angeles feeder load shaping optimization, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	0.70	0.47	0.94
Peak Demand [MW]	0.00	-0.04	0.03
Peak to Valley [%]	99.38	96.77	100.60
Load Factor [%]	0.94	0.34	1.77
Ramp [MW]	0.05	-0.27	0.29

Looking at several days in closer detail helps illustrate this point. In Figures 8.9, 8.10 and 8.11, small bumps and depressions in the base case feeder demand curve are partially or fully eliminated by the actions of the controllers. The net result is a feeder demand curve with a smaller rate of change, and therefore a lower ramp rate required by generation to match the load. Although the difference in this case is marginal, this can be a valuable feature to the load serving entity.




Figure 8.8: Total spectral power as a function of frequency bin for Los Angeles feeder load shape optimization, 70% participation.

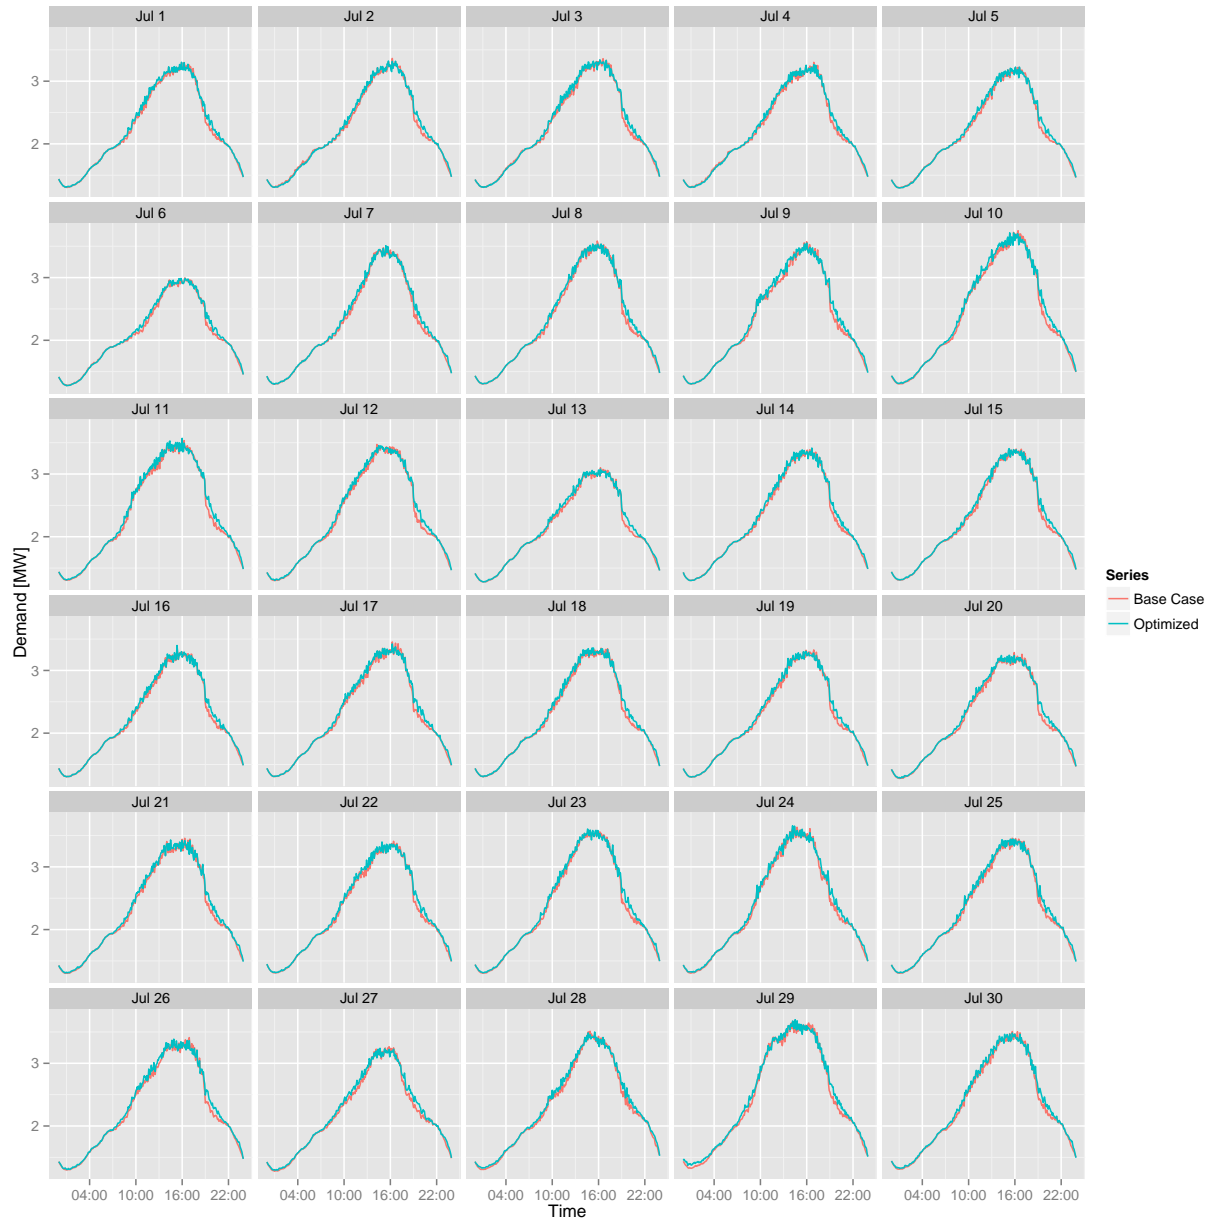


Figure 8.9: Feeder demand profiles for Los Angeles load shape optimization, 70% participation.




Figure 8.10: Feeder demand profiles for Los Angeles, July 9 load shape optimization, 70% participation.




Figure 8.11: Feeder demand profiles for Los Angeles, July 11 load shape optimization, 70% participation.




Figure 8.12: Feeder demand profiles for Los Angeles, July 29 load shape optimization, 70% participation.

8.4 New York

Similar to the Los Angeles case, the utility of optimization in this case is measured in the reduction of ramping requirements. While the effect is only slightly more noticeable in the monthly plots of Figure 8.14, plots of individual days in Figures 8.15 and 8.16 show the smoothing effects more clearly. The metrics Table 8.4 and power spectrum analysis Figure 8.13 indicate that performance of most other metrics is improved over the base case with few exceptions. Again, 30% participation results included in Appendix D show similar trends.

Table 8.4: Performance metrics for New York feeder load shaping optimization, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	2.38	1.91	2.92
Peak Demand [MW]	-0.03	-0.15	0.09
Peak to Valley [%]	91.19	79.93	99.38
Load Factor [%]	2.13	0.87	4.24
Ramp [MW]	-0.01	-0.93	0.86




Figure 8.13: Total spectral power as a function of frequency bin for New York feeder load shape optimization, 70% participation.

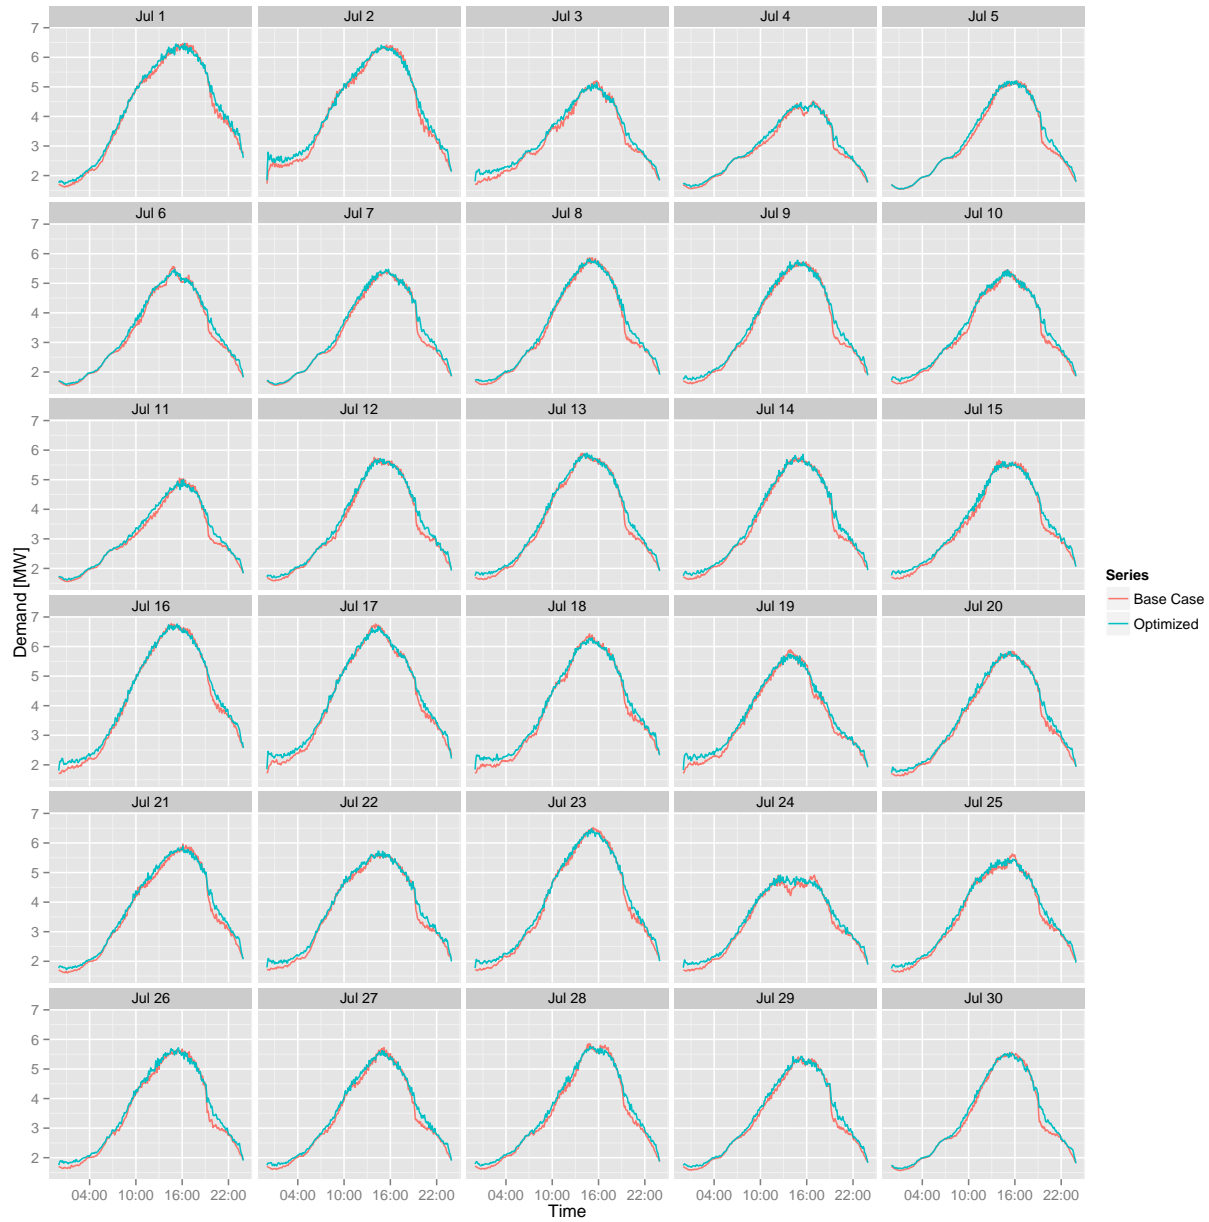


Figure 8.14: Feeder demand profiles for New York load shape optimization, 70% participation.




Figure 8.15: Feeder demand profiles for New York, July 4 load shape optimization, 70% participation.




Figure 8.16: Feeder demand profiles for New York, July 24 load shape optimization, 70% participation.

8.5 Summary and Conclusions

The results for the three feeders show the proposed methodology to be effective in shaping demand, resulting in a relatively smooth demand curve like that of the reference curve provided to the controller. Although occurrences of increased peak demand are observed in a few cases, improved demand reduction could likely be achieved through further refinements to the reference curve, or by the introduction of a weighting term to periods of particular importance. Further refinement is not considered in this work.

The proposed methodology is well suited to a wide range of applications because of the generic nature of the approach. In these studies, demand shaping utilized only the HVAC systems of the home. While these systems tend to be the largest loads in the home, other loads such as pool pumps and electric hot water heaters could be manipulated in a similar manner, provided that they are equipped with the appropriate controls.

The distributed and decentralized nature of the approach allows the entity directing the load shaping effort, whether at the feeder level or higher, to be agnostic to the type of equipment participating in load shaping. In the near future, these could include electric vehicle chargers, building-sited battery storage and cogeneration plants. The extension of this approach to micro-grid management is a logical next step and interesting area of future research.

The studies presented in this chapter illustrate what this methodology can do to shape demand in response to the loads present on existing distribution networks. The next two chapters will investigate how this methodology can be used to shape anticipated feeder load as new generation is added to the electric system.

Chapter 9

Rooftop Solar Electric Generation

The adoption rate of residential rooftop photovoltaic (PV) systems has steadily risen over the last decade due to strong incentives and falling prices [83]. In California, the high penetration of PV is expected to cause conflicts with traditional generation, resulting in periods of rapid ramping and over-generation. This phenomenon, popularized by the so-called “Duck Curve” (see Figure 9.1) presents a real and present problem for utilities — who have little control over the production of electricity from these distributed rooftop solar resources — where high penetration rates are occurring.

This chapter evaluates the application of the methodology developed in Chapter 8 as a potential solution to this problem. The chapter first describes minor modifications to the methodology necessary for this study, then introduces the solar models used to generate electricity in the distribution feeder. Next, the selection of PV penetration levels is discussed, followed by a description of the method developed for sizing and distributing the PV systems to homes in the feeder model. The chapter then explores two levels of PV system penetration for each of the three feeders.




Figure 9.1: Example of the Duck Curve produced in systems with high PV penetration levels leads to times of high ramping and over generation. The moniker is derived from the area between the two curves which resembles a duck on some days [19].

9.1 Methodology

At a high level, the methodology employed in this chapter is identical to that described in Chapter 8: 1) a reference demand curve is generated from the base case feeder demand, 2) this reference demand is translated into a reference signal and provided to the controller in each simulated home, 3) the controller optimizes house electric demand to match a house reference demand created from the reference signal using the same objective function described previously.

Differences between the investigations in this chapter and those studied earlier lie in the distribution feeder demand which, in these investigations, contains electric generation from rooftop solar systems. The generation by these systems has a significant effect on the shape of the feeder demand curve which results in the Duck Curve characteristics described above. The only adjustment required to the methodology is simply the introduction of solar electric generation in the feeder model. Annual electricity plots of feeder demand with solar can be seen in Appendix E, Figures E.2, E.6 and E.12. Additional examples of the “Duck Curve” can be seen in the Houston and New York feeders in Figures E.1 and E.13.

9.2 Model Development and Validation

Estimating the production of rooftop solar electricity in the distribution feeder requires the addition of a solar collector model to the building model described in Chapter 3. While this level of detail is not strictly required, the simulation of solar at the individual residence allows the controller to see the effect of solar generation on whole house demand, thus yielding a more realistic representation of the distributed optimization. The model is composed of two components: a solar collector producing direct current, and an inverter which converts the direct current to alternating current.

The direct current portion of the photovoltaic system is a simple flat-plate collector model. Cell efficiency is assumed to vary with temperature, which is assumed to be uniform across all cells in the collector. Temperature is calculated using the simple relationship developed by Ross [81].

Although the collector here is assumed to be directly roof mounted rather than rack mounted, this relationship is sufficient for the purpose of this research. The calculation of direct current collector efficiency, η_c , and power, Q_{DC} , is shown in Equations 9.1 and 9.2. Parameter values assumed for the model are shown in Table 9.1.

$$\eta_c = \eta_0[1 - \beta(T_c - T_0)] \quad (9.1)$$

$$Q_{DC} = \eta_c I_t A_c \quad (9.2)$$

where:

η_c is the collector efficiency at temperature T_c

η_0 is the collector efficiency at standard test conditions

β is the efficiency temperature coefficient

T_c is the collector operating temperature

T_0 is the collector operating temperature at standard test conditions

I_t is the total collector insolation calculated by Equation 3.6

A_c is the total collector area

The amount AC power, Q_{AC} , produced by the system can be calculated using a simple inverter model with part-load efficiency developed by Yewdall [99]:

$$Q_{AC} = \eta_i \eta_s Q_{DC} \quad (9.3)$$

where:

η_i is the inverter efficiency at specified input power

η_s is the balance of system nominal efficiency

To ensure that the model produces reasonable values, it is cross-validated against results produced by the PVWatts Performance Calculator for Grid-Connected PV Systems [62]. For this comparison, both PVWatts and the model above are simulated using a year of typical Houston

Table 9.1: Default photovoltaic system parameters.

Parameter	Value	Units
η_0	15	%
β	0.004	K ⁻¹
T_0	47	°C
η_s	95	%

weather with the parameters in Table 9.1. Note that the default system derating factor of 0.77 was simulated for PVWatts.

Table 9.2: Photovoltaic system parameters used for model cross-validation.

Parameter	Value
Collector tilt	22.6°
Collector azimuth	0°
Nominal size	5kW

Agreement between the models is good, with an annual difference in AC power of -1.7%. A normalized root mean square error of 8.6% indicates reasonable agreement at each time step. An example of AC power time series produced by both models is shown in Figure 9.2. Visual inspection of the time series plot also indicates that the model produces values comparable to PVWatts.




Figure 9.2: Example of PV model output compared to PVWatts, July 14–21.

9.3 Desired PV Penetration

Penetration levels are adopted from the National Renewable Energy Laboratory's Western Wind and Solar Integration Study Phase 2 report (WWSIS-2) [59]. The WWSIS-2 report outlines several scenarios for solar penetration levels. This chapter adopts the values from the "High Solar" and "TEPPC" scenarios, corresponding to 25% and 3.6% of annual regional electricity consumption, respectively. Although the WWSIS-2 report is specific to the western United States, these values are adopted for the New York feeder as well.

Following the assumptions in the WWSIS-2, 60% of the installed solar generation is assumed to be solar electric. The WWSIS-2 report further divides this between rooftop solar and utility scale solar at 40% and 60% distributions. Here, this research departs from the WWSIS-2 and assumes all of the PV generation is provided by rooftop solar, distributed equally between the residential sector and all others. That is, 50% of all solar electricity generated on an annual basis is provided by residential rooftop solar.

The feeder models considered in this work represent feeders heavily populated with residential buildings which dominate the annual electricity consumed. It is assumed that the distribution of rooftop solar is consistent with a feeder having a composition matching the regional average as measured by annual electricity consumption. The distribution of residential solar must be corrected by the fraction of residential electricity consumption in the feeder so that the distribution reflects the difference between region and feeder. For example, if the 50% of rooftop solar is generated by residential buildings where residences represent 38% of the annual electricity consumed in a region, and 80% of the electricity consumed in the feeder, then a higher percentage residential rooftop solar would need to be modeled in the feeder to account for the larger residential consumption. Residential consumption on an annual basis is easily calculated from the hybrid model; values for electricity consumption by sector are available from the EIA [1], making this correction fairly easy.

The corrected rooftop solar penetration values are summarized in Table 9.3. Note that the values reported for both sectors have been also been adjusted according to the 60% solar electric

assumption in the WWSIS-2 report. In the section that follows, the penetration level shown in the residential row will be the value used to distribute PV to homes in the feeder. The contribution of solar from other sources is not modeled.

Table 9.3: PV Penetration for three feeder models.

		GridMPC Load	EIA Load	3.6% Penetration	25% Penetration
Houston	Residential	80%	38%	1.87%	13.01%
	Other	20%	62%	0.29%	1.99%
New York	Residential	78%	35%	1.88%	13.04%
	Other	22%	65%	0.28%	1.96%
Los Angeles	Residential	86%	35%	1.99%	13.81%
	Other	14%	65%	0.17%	1.19%

9.4 Sizing and Distribution of PV Systems

Simulating the desired residential penetration level requires distributing the correct number of residential PV system to homes on the distribution feeder. This is accomplished with a two stage process that begins with the sizing of PV systems for each home on the feeder. This first stage is as follows:

1. A prototypical photovoltaic system is simulated for the entire year using the parameters specified in Table 9.2.

2. The ratio of electricity produced to nominal system size is calculated, giving

$$kWh_{produced}/kW_{nominal}/year.$$

3. For each home, annual electricity consumption is summed, giving

$$kWh_{used}/year.$$

4. The nominal system size for each home is calculated using the ratio found in step 2, assuming that each system offsets 80% of the value found in step 3.

The second stage distributes PV systems across the population of homes. This process need only take place once for each feeder and desired level of penetration. Were every home in the

feeder to be given a PV system, the total contribution of all distributed generation in the feeder would equal 80% of the residential feeder electricity consumption. Therefore, to arrive at the desired residential contribution, only a fraction of the population need be simulated with a PV system. Choosing the population is easily accomplished by:

1. Assigning a random number to each home between zero and one.
2. Choosing an arbitrary threshold value between zero and one. If a home has a random number less than this threshold, it is assigned a PV system.
3. Summing the total contribution of distributed generation for those homes under the threshold value, and calculating the ratio of this sum to the total annual feeder consumption, giving the penetration level.
4. Adjusting the threshold value up or down until the calculated penetration measures the desired penetration shown in Table 9.3.

9.5 Houston

The introduction of rooftop solar creates significant variation in the feeder demand curve. In the monthly plots, Figures 9.3, 9.4, 9.5 and 9.6, this can be seen in the high solar penetration cases in particular, and to a lesser extent in the low solar penetration cases. In all cases, the controller appears to reduce demand variations by smoothing the demand curve. Plots for July 1 and July 21 in Figures 9.11 and 9.12 illustrate the smoothing effect for the high solar and penetration case more clearly. In 30% cases the ability is limited, consistent with previous studies, but optimization appears to be beneficial nonetheless.

Metrics shown in Tables 9.4, 9.5, 9.6 and 9.7 are generally improved across all dimensions, with a few exceptions. As with previous studies, total energy consumption increases as a result of the load shifting activities. Trends in the improvements can be seen across the four cases: improvements over the base case are greatest for the high participation and high penetration,

followed by high participation and low penetration, then low participation and high penetration and finally low participation and low penetration.

Power spectrum analysis reveals the same trends in Figures 9.7, 9.8, 9.9 and 9.10. Plots for the high solar case reveal a larger amount of the power spectrum concentrated in the twelve to six hour frequency range compared to previous studies. While there appears to be a very marginal increase in power concentrated in the twenty-four to twelve hour range in the optimized case, the reduction in the twelve to six hour range is significant. This trend and the increase in the DC power component show the benefits of the optimization, similar to that observed in previous chapters.

Table 9.4: Performance metrics for Houston feeder load shaping optimization, high solar penetration case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	4.82	4.03	5.36
Peak Demand [MW]	-0.48	-0.92	0.02
Peak to Valley [%]	80.04	72.91	88.88
Load Factor [%]	5.85	3.01	8.95
Ramp [MW]	-0.70	-2.56	1.03

Table 9.5: Performance metrics for Houston feeder load shaping optimization, low solar penetration case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	4.85	4.17	5.26
Peak Demand [MW]	-0.16	-0.92	0.22
Peak to Valley [%]	84.01	76.39	96.20
Load Factor [%]	3.20	0.66	7.45
Ramp [MW]	-0.31	-2.81	1.49

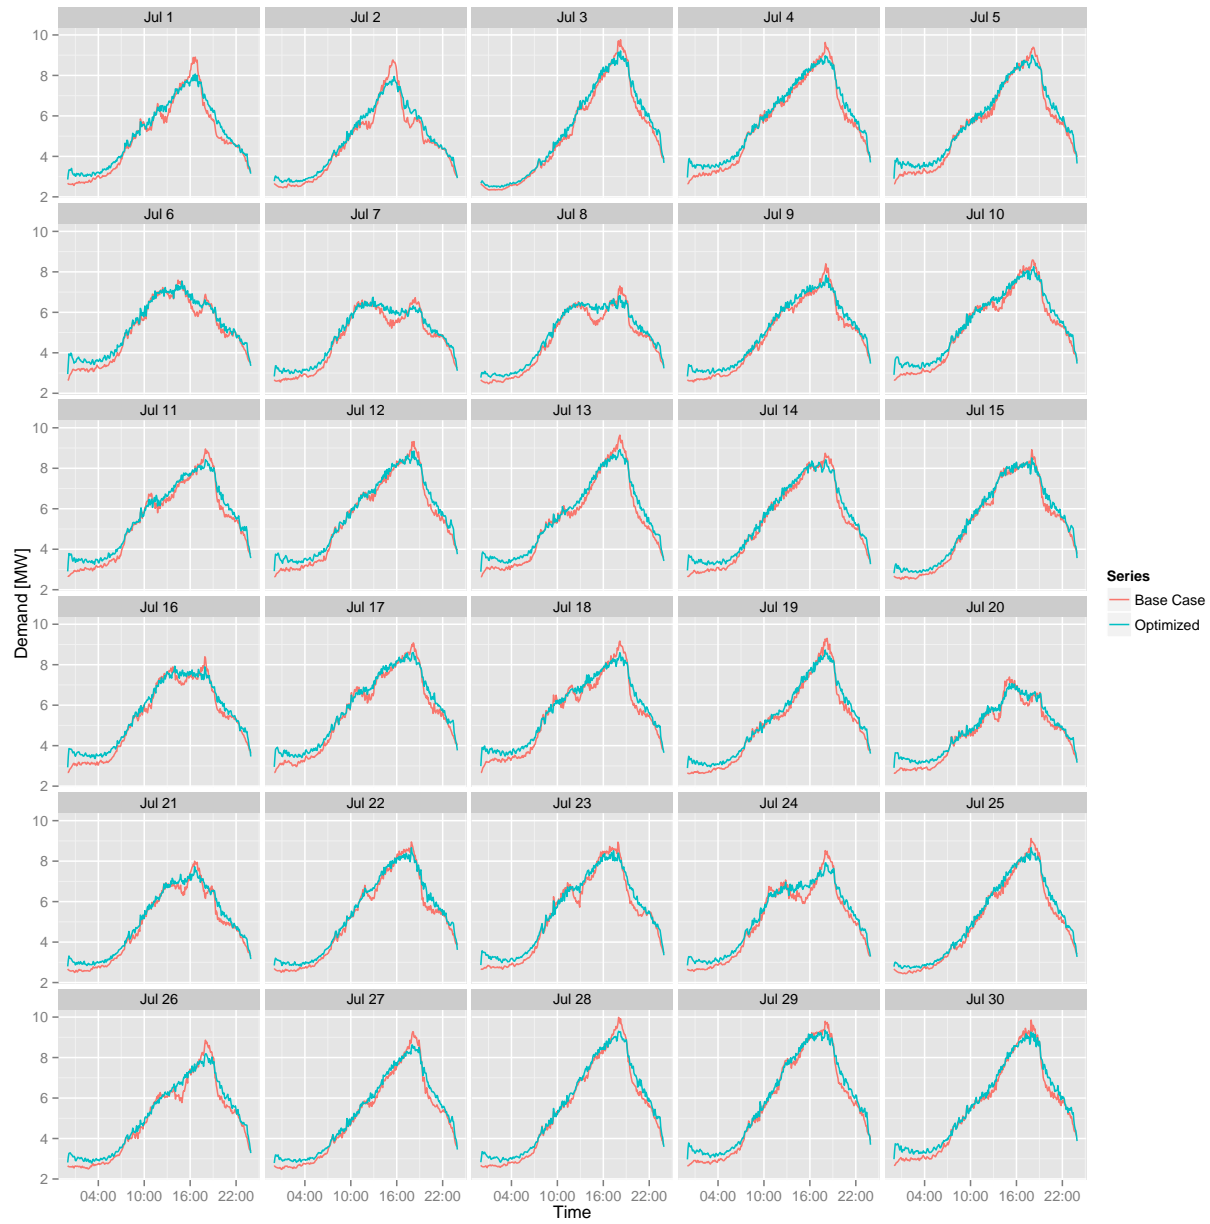


Figure 9.3: Feeder demand profiles for Houston load shape optimization, high solar penetration case, 70% participation.

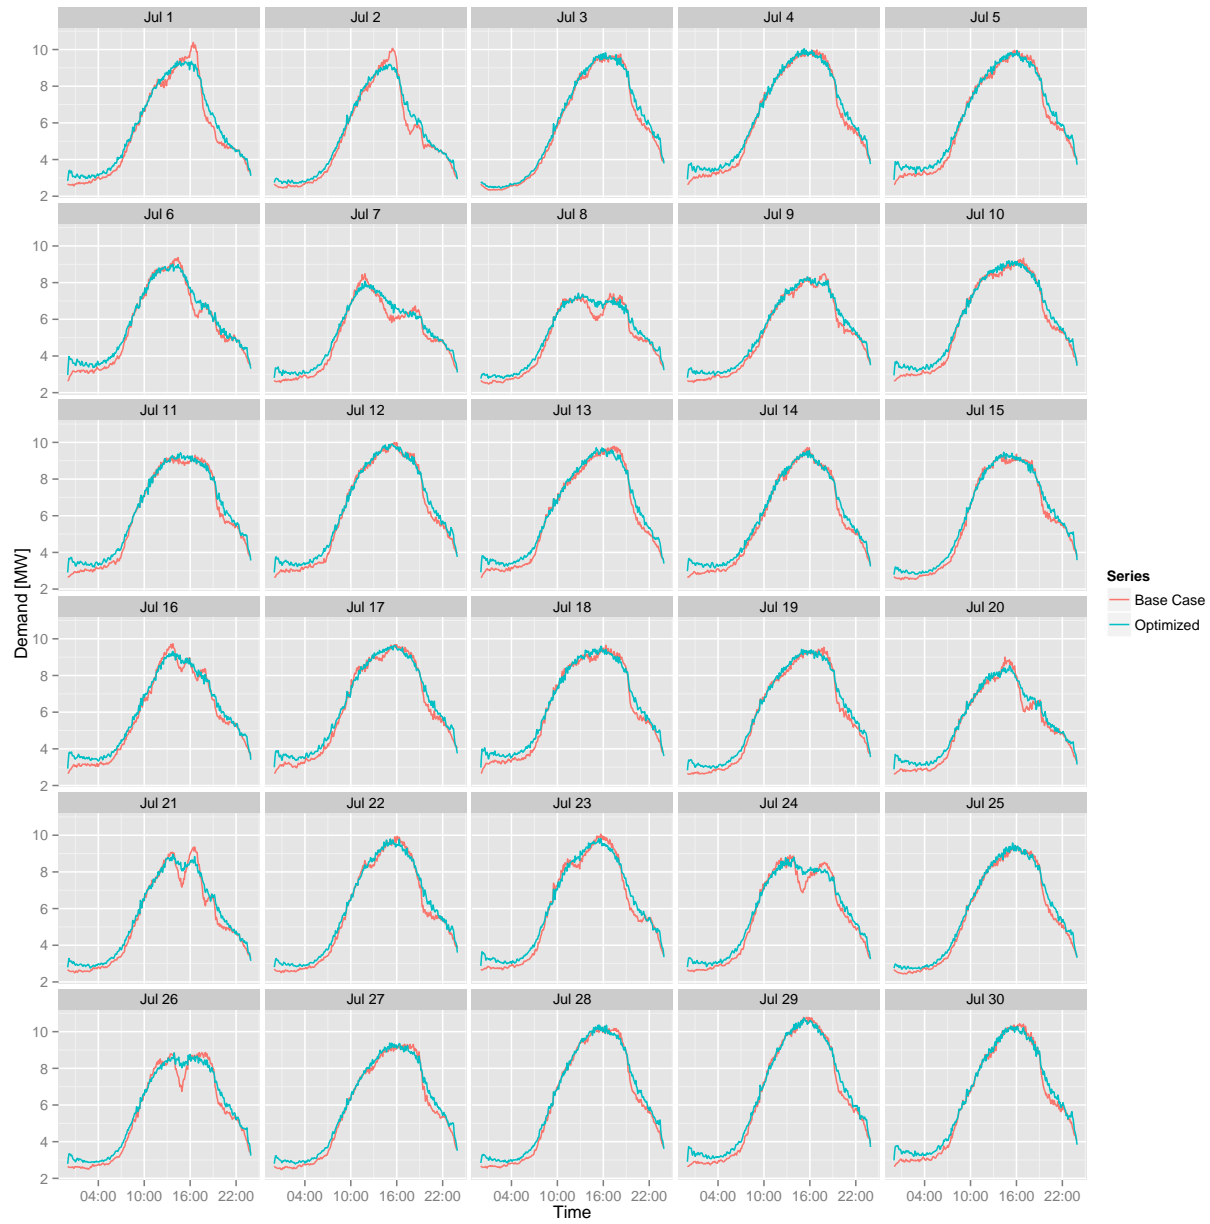


Figure 9.4: Feeder demand profiles for Houston load shape optimization, low solar penetration case, 70% participation.

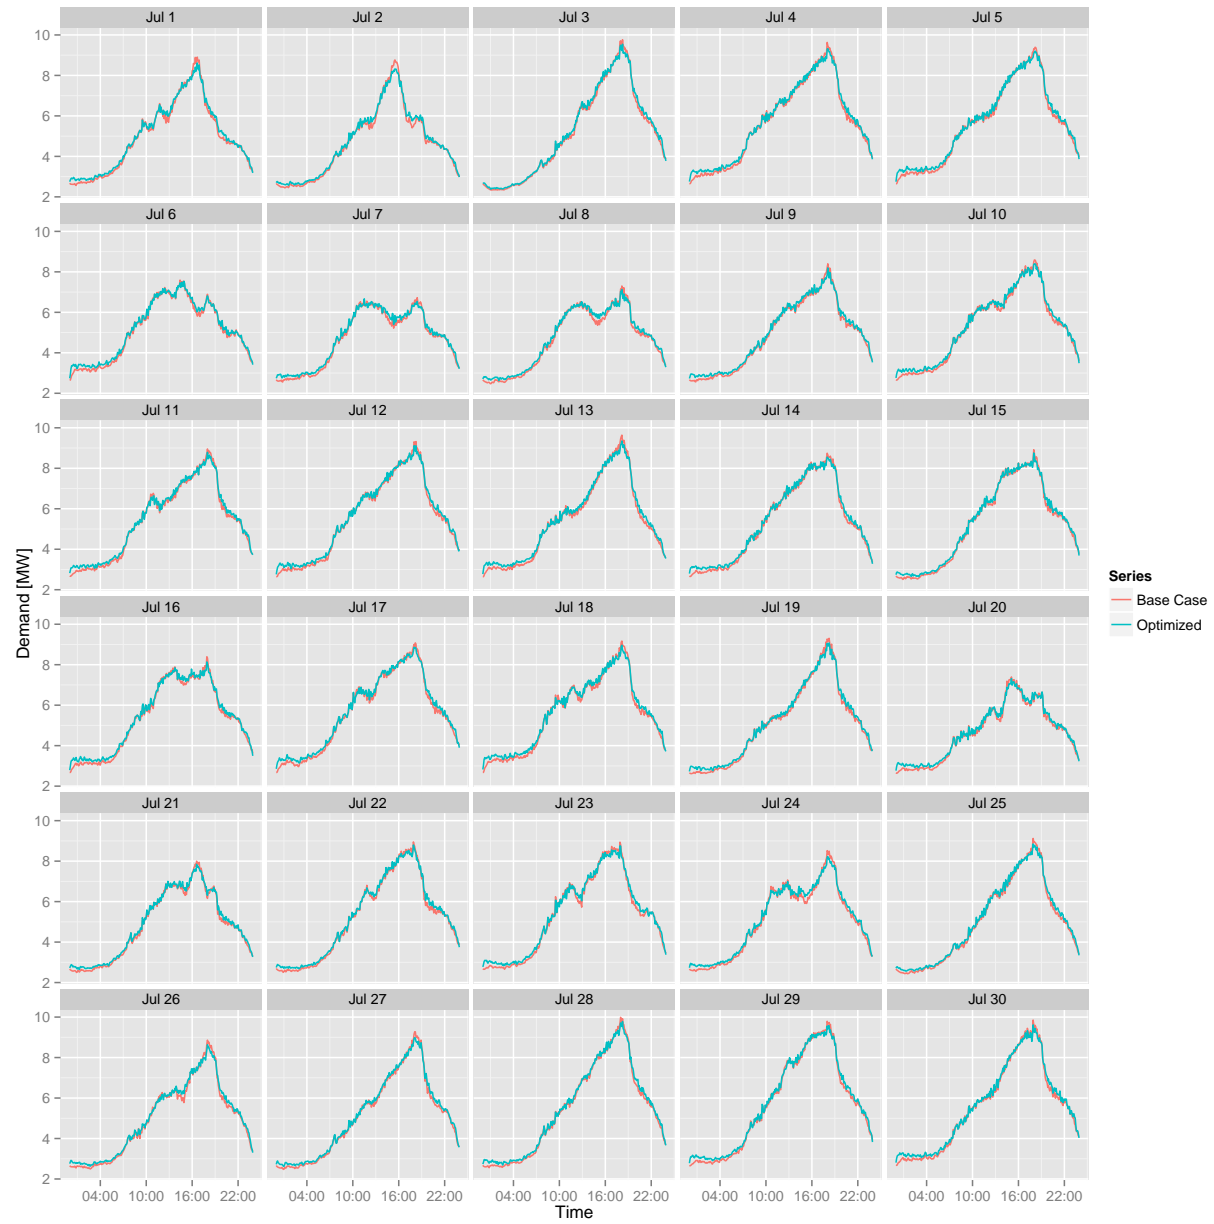


Figure 9.5: Feeder demand profiles for Houston load shape optimization, high solar penetration case, 30% participation.

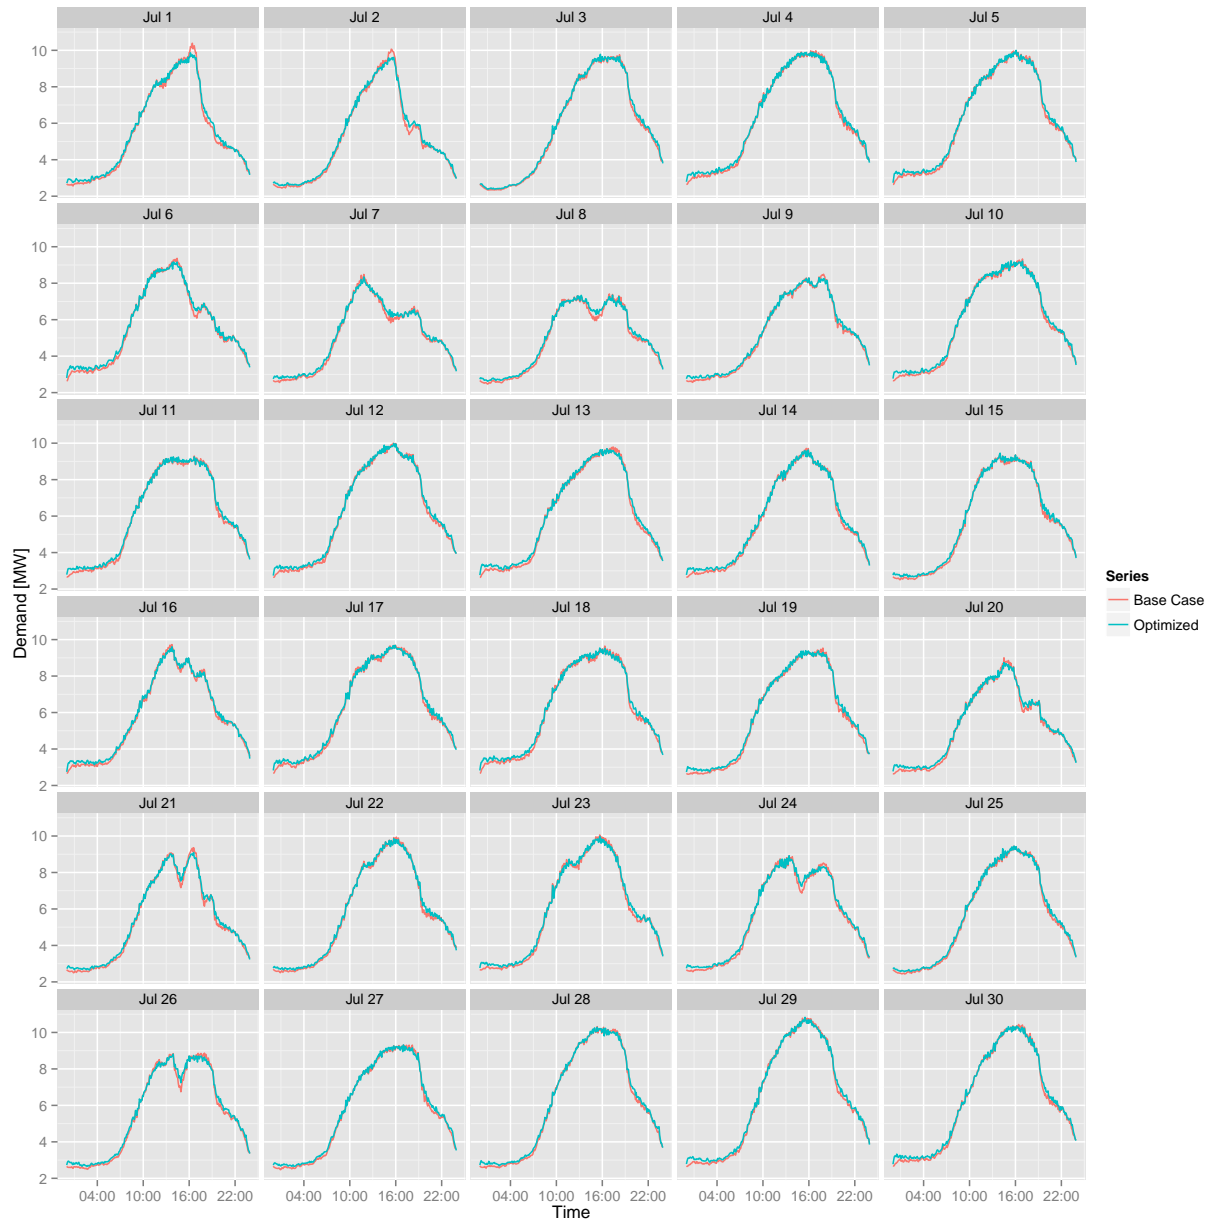


Figure 9.6: Feeder demand profiles for Houston load shape optimization, low solar penetration case, 30% participation.

Table 9.6: Performance metrics for Houston feeder load shaping optimization, high solar penetration case, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	2.07	1.73	2.35
Peak Demand [MW]	-0.20	-0.41	-0.07
Peak to Valley [%]	89.86	86.36	95.39
Load Factor [%]	2.43	1.84	3.54
Ramp [MW]	-0.77	-1.72	-0.09

Table 9.7: Performance metrics for Houston feeder load shaping optimization, low solar penetration case, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	2.07	1.76	2.40
Peak Demand [MW]	-0.10	-0.49	0.07
Peak to Valley [%]	91.05	87.08	97.64
Load Factor [%]	1.60	0.44	3.60
Ramp [MW]	-0.61	-1.72	0.22




Figure 9.7: Total spectral power as a function of frequency bin for Houston feeder load shape optimization, high solar penetration case, 70% participation.




Figure 9.8: Total spectral power as a function of frequency bin for Houston feeder load shape optimization, low solar penetration case, 70% participation.




Figure 9.9: Total spectral power as a function of frequency bin for Houston feeder load shape optimization, high solar penetration case, 30% participation.




Figure 9.10: Total spectral power as a function of frequency bin for Houston feeder load shape optimization, low solar penetration case, 30% participation.




Figure 9.11: Feeder demand profiles for Houston, July 1 load shape optimization, high solar penetration case, 70% participation.




Figure 9.12: Feeder demand profiles for Houston, July 21 load shape optimization, high solar penetration case, 70% participation.

9.6 Los Angeles

High penetrations of solar in the Los Angeles feeder results in base case demand curves that exhibit the Duck Curve characteristics very clearly. Unfortunately, the ability of homes to shift large portions of electricity demand is limited. This affects the ability of the controller to remove these characteristics. Figures 9.13, 9.14 illustrate this for the high and low solar cases. One may notice that the low solar case resembles the no solar case very closely. Given the low level of penetration, this is not terribly surprising.

Nevertheless, in both the high and low solar cases, the optimization does show improved performance in demand reduction, peak to valley ratio, load factor, ramping and spectral power distribution. The trends regarding participation and penetration seen in the Houston feeder can be observed here. High participation results are shown in Tables 9.8 and 9.9, and Figures 9.15 and 9.16. Low participation metrics as well as monthly demand plots are included in Appendix E.

Similar to the Houston feeder, the Los Angeles feeder exhibits a large amount of power concentrated in the twenty-four to twelve hour frequency; this is consistent with the shape of the demand curve seen in the daily demand profiles. Optimization reduces the power concentration in this frequency range only slightly; box-whisker plots for the high solar cases show a wide range of results.

Plots of July 13 and 15 in Figures 9.17 and 9.18 show very little change in demand between base and optimal cases during early morning hours. This has been previously observed in other cases; the absence of flexible cooling demand during these hours is the cause.

Table 9.8: Performance metrics for Los Angeles feeder load shaping optimization, high solar penetration case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	0.64	0.43	0.91
Peak Demand [MW]	-0.10	-0.14	-0.03
Peak to Valley [%]	95.05	85.04	99.09
Load Factor [%]	2.82	1.54	3.43
Ramp [MW]	-0.31	-0.67	0.03




Figure 9.13: Feeder demand profiles for Los Angeles load shape optimization, high solar penetration case, 70% participation.

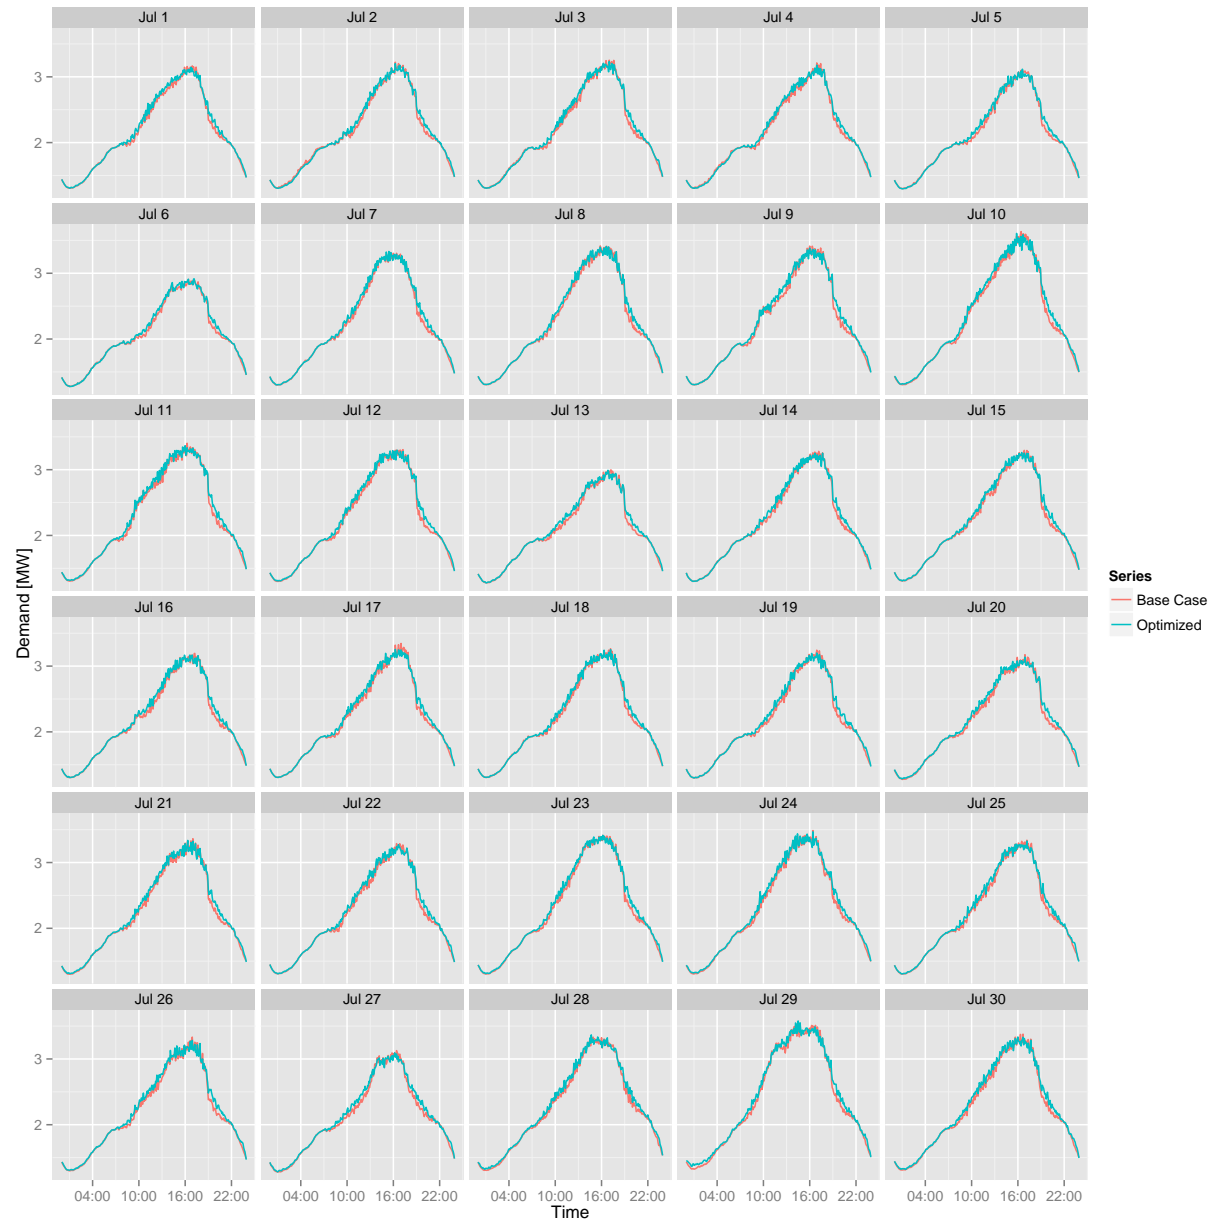


Figure 9.14: Feeder demand profiles for Los Angeles load shape optimization, low solar penetration case, 70% participation.

Table 9.9: Performance metrics for Los Angeles feeder load shaping optimization, low solar penetration case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	0.68	0.43	0.91
Peak Demand [MW]	-0.02	-0.06	0.02
Peak to Valley [%]	98.93	97.19	100.57
Load Factor [%]	1.31	0.41	2.05
Ramp [MW]	-0.01	-0.25	0.24




Figure 9.15: Total spectral power as a function of frequency bin for Los Angeles feeder load shape optimization, high solar penetration case, 70% participation.




Figure 9.16: Total spectral power as a function of frequency bin for Los Angeles feeder load shape optimization, low solar penetration case, 70% participation.




Figure 9.17: Feeder demand profiles for Los Angeles, July 13 load shape optimization, high solar penetration case, 70% participation.




Figure 9.18: Feeder demand profiles for Los Angeles, July 15 load shape optimization, high solar penetration case, 70% participation.

9.7 New York

The introduction of solar electric generation into the New York feeder results in demand profiles resembling both Houston and Los Angeles profiles. For example, in the high solar case shown in Figure 9.19, July 5, 7, 29 and others exhibit Duck Curve characteristics of Los Angeles, while other days such as July 13, 18 and 19 contain relatively large amounts of variation from hour to hour similar to Houston. The low solar case (Figure 9.20), similar to previous results, shows demand profiles with much less variation.

The mix of characteristics in the same feeder results in a key observation: the controller appear to be less able to smooth demand in days with Duck Curve characteristics, and more effective at smoothing demand in days with high hour to hour variation in demand. This indicates that in addition to the amount of flexible cooling demand available, the ability of the controller to shift demand is related to the length of time it must be shifted. Specifically, the controller is able to move demand forwards or backwards by timescales measured in one or two hours, not four, six, or twelve. This makes sense intuitively: because of losses in the system related to building envelope integrity and the limited amount of thermal mass present in a residence, the controller finds no benefit to moving demands by longer timescales. Doing so would be inefficient because the benefit of shifting diminishes with time. This has been observed by the short amount of time it takes for the demand curve to recover in the demand response cases shown in Chapter 5 and in the precooling experiment in Chapter 6.

Optimization results in improvements over the base case similar to those observed in the Houston and Los Angeles feeders. High participation results are summarized in Tables 9.10 and 9.11, and in Figures 9.21 and 9.22. Low participation results follow the same trends observed previously. Those results can be seen in Tables E.3 and E.4 and Figures E.10 and E.11 in Appendix E, along with plots of individual days showing the smoothing effects of optimization.

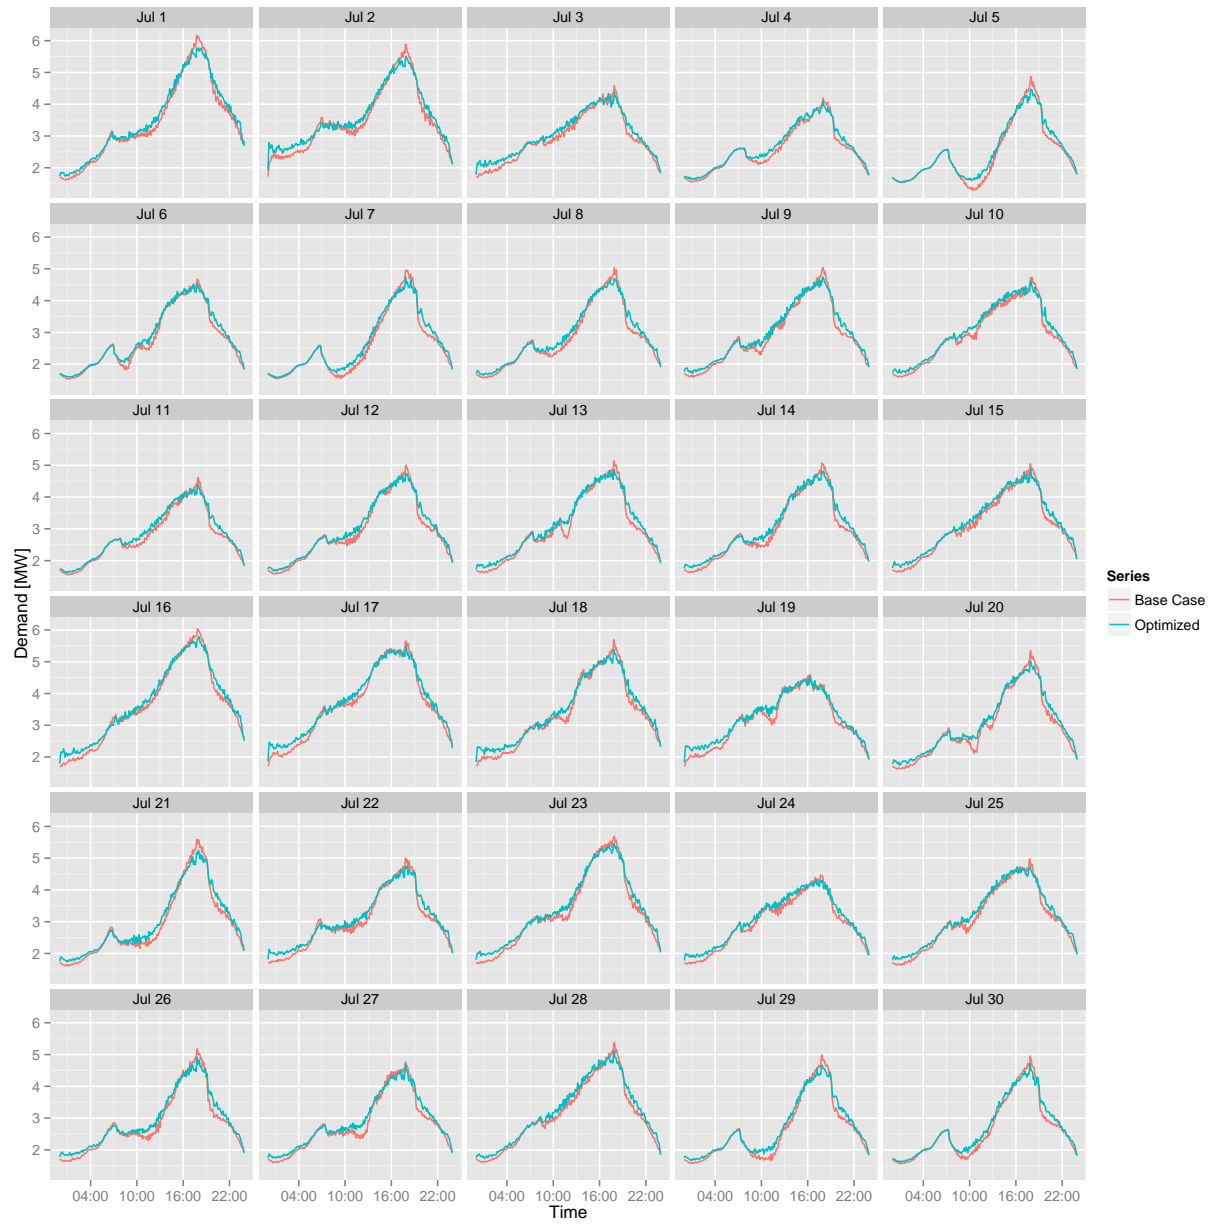


Figure 9.19: Feeder demand profiles for New York load shape optimization, high solar penetration case, 70% participation.

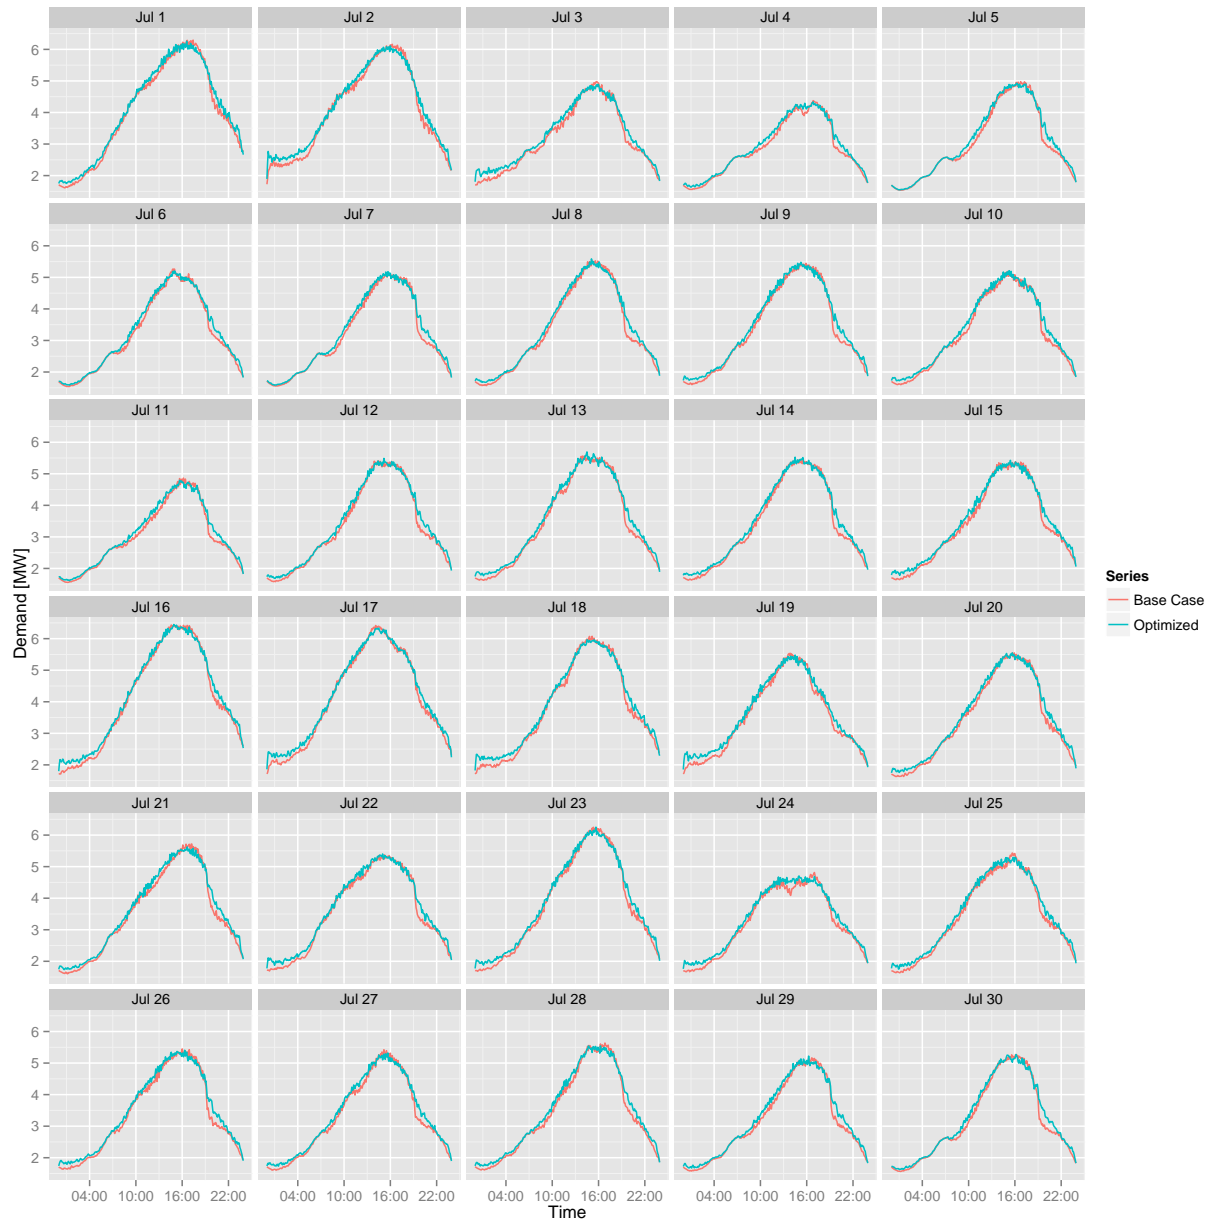


Figure 9.20: Feeder demand profiles for New York load shape optimization, low solar penetration case, 70% participation.

Table 9.10: Performance metrics for New York feeder load shaping optimization, high solar penetration case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	2.32	1.71	3.01
Peak Demand [MW]	-0.25	-0.39	-0.04
Peak to Valley [%]	86.62	78.55	93.07
Load Factor [%]	5.00	3.23	6.04
Ramp [MW]	-0.47	-1.87	0.14

Table 9.11: Performance metrics for New York feeder load shaping optimization, low solar penetration case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	2.36	1.99	2.94
Peak Demand [MW]	-0.03	-0.12	0.08
Peak to Valley [%]	91.11	81.61	98.68
Load Factor [%]	2.17	0.89	3.95
Ramp [MW]	0.04	-0.95	0.59




Figure 9.21: Total spectral power as a function of frequency bin for New York feeder load shape optimization, high solar penetration case, 70% participation.




Figure 9.22: Total spectral power as a function of frequency bin for New York feeder load shape optimization, low solar penetration case, 70% participation.

9.8 Summary and Conclusions

The introduction of rooftop solar electric to residences at high penetration levels creates significant variability in the feeder demand at several timescales. Short term variability — on the order of one or two hours — creates fluctuations in demand over short periods of time, requiring traditional generation to respond quickly and repeatedly. At longer timescales, i.e. six to twelve hours, the contribution of solar electric generation results in large troughs and steep, sustained increases in demand. Both conditions create potential problems for load serving entities who must manage generation assets to follow demand.

This chapter has investigated whether or not the methodology developed in Chapter 8 can be used to shape demand in anticipation of variations in supply, shifting control of the system from a load-following to supply-following paradigm. Results suggest that this can in theory be realized with high participation rates and the availability of flexible demand. In the cases studied here, the methodology shows the greatest promise in reducing short-term demand variations by smoothing the feeder demand profile. The ability of the approach to address longer term variations spanning multiple hours is limited, but still provides some improvement in terms of load factor, peak demand, peak to valley, and ramping metrics.

Chapter 10

Utility Scale Wind Generation

Previous investigations in Chapters 8 and 9 considered how load a shaping methodology could be utilized to address the intrinsic variability of the base case feeder electric demand. This chapter explores the ability of the methodology to shape demand based on some extrinsic variability defined outside of the distribution feeder, by introducing wind generator output into the construction of the feeder reference demand curve. The hypothesis tested here is: although the production of wind generated electricity does not shape the demand characteristics of the feeder directly, the methodology can be used to offset the feeder demand in order to absorb the variations wind generated electricity introduces.

The chapter first discusses the inclusion of wind generation into the reference demand curve. Next, the selection and sizing of wind turbines is described. Simulation results from six combinations of feeders and wind penetration levels are then presented with interpretation and conclusions.

10.1 Methodology

The methodology adopted in this chapter is identical to that described in Chapter 8 with one exception: the feeder demand profile used to generate the reference demand curve is a composite of base case feeder demand and utility scale wind generation. Generation of electricity from wind turbines is assumed to occur outside of the distribution feeder and is therefore not normally reflected in base case feeder demand. To allow the controller to see and respond to the presence of wind generation outside of the feeder, wind generated electricity is subtracted from the base case feeder

demand forecast to create a *composite* reference demand curve. Models used to simulate wind generated electricity are described in Section 10.2. Sections 10.3 and 10.4 describe the selection of the desired wind penetration levels and how they are simulated.

With the introduction of wind generation into the reference demand curve, the curve reflects the variability of generation outside of the distribution feeder. This curve is constructed using a 4-hour moving average window in the manner outlined in Chapter 8. The smoothed composite demand curve is then converted into the reference signal sent to each building, and controllers optimize house electric demand to offset the variations introduced in the composite curve. In all of the results presented in this chapter and Appendix F, the base case and optimized curves shown are composite curves. It is important to keep in mind that the optimized feeder demand curve is shaped to offset the variability in generation; presenting the results as a composite allows the smoothing effect of optimization to be observed and measured easily.

It is also important to remember that this study does not attempt to model a utility scale wind farm as accurately as possible. Rather, the intent is to show how MPC may allow residential buildings to smooth the demand curve when wind is present, removing the need for traditional generation to perform this task. This study is in essence a prototype for future studies, where generation variability can be introduced by other source.

10.2 Model Selection and Description

GridLAB-D contains models of utility scale wind turbines, making the task of modeling the contribution from utility scale wind relatively simple. The two utility scale models available in GridLAB-D are developed from the GE 2.5 and Vestas V82 turbines. Coefficients of performance, cut-in and cut-out speeds vary between the two, resulting in somewhat different power output under the same wind conditions. To illustrate the differences, the normalized power output from both turbines during a single week in July is plotted in Figure 10.1.

Given the intent of this work, there is no turbine model better suited to this study than another. Therefore, to roughly account for the diversity of generation that may be present in a real




Figure 10.1: Normalized wind turbine output for two turbine models showing the difference in output characteristics.

system, e.g. contributions from several wind farms, a blend of output from the two turbines is used, producing a more diverse set of conditions for optimization. The assumption made is that half of the wind generated electricity, on an annual basis, will be provided by each of the two types of turbines.

10.3 Desired Wind Penetration

Wind penetration levels are adopted from the National Renewable Energy Laboratory Western Wind and Solar Integration Study Phase 2 report (WWSIS-2) [59]. The WWSIS-2 report outlines several scenarios for wind penetration levels. This chapter adopts the values from the “High Wind” and “TEPPC” scenarios, corresponding to 25% and 9.4% of annual regional electricity consumption, respectively. Although a similar study has been performed for the eastern United States with slightly different penetration assumptions [29], values from WWSIS-2 report are adopted for the New York feeder for convenience and ease of comparison.

10.4 Sizing Wind Systems

Because the production of wind is assumed to occur outside of the distribution feeder and has no direct relationship to the number of homes or composition of the feeder load, the sizing of the wind systems is rather trivial:

1. Simulate each turbine for an entire year in each feeder location.
2. Normalize the electricity production from each turbine by dividing by the sum of total production by the turbine.
3. Add the two normalized curves for each location and divide by two (blending assumption).
4. Scale the new curve by the desired penetration level for the location.

The last step is simply a matter of multiplying the normalized curve by the number of megawatt hours required to meet the desired penetration level. For the three feeders, this results

in the scaling factors shown in Table 10.1 for each of the penetration levels. This process need only take place once for each feeder and desired level of penetration.

Table 10.1: Scaling factors used for scaling wind turbine output to desired penetration levels.

	Houston	Los Angeles	New York
High Penetration [MWh]	10,633	6,425	4,914
Low Penetration [MWh]	4,253	2,570	1,966

10.5 Houston

The most striking feature of the demand profiles shown in Figure 10.3 is the lack of variation in the optimized result despite the hour to hour variation that wind generation introduces. The effect is much less pronounced in the low wind scenario shown in Figure 10.4, but present nonetheless. In both scenarios, the controller exhibits an impressive ability to smooth the composite demand curve. It bears repeating that the optimized demand curve includes the contribution from the wind turbines. Were the wind generated electricity removed from the optimized result presented, one could fully appreciate the amount of demand the controllers shift in order to achieve the result shown. This can be seen in Figure 10.2 which shows the wind turbine output plotted with the pre-composited optimized feeder demand profile.

According to the metrics in Tables 10.2 and 10.3, the performance of the optimized results is undoubtedly the best seen in terms of peak to valley ratio, load factor, and ramping. Peak demand is very respectable on average, with the largest single reduction observed in the minimum column. Total electric consumption, however, continues to show an increase as seen in the vast majority of previous cases. Low participation cases show similar results, but are obviously limited by the number of homes optimized and therefor limited in flexible cooling demand. Low participation results can be seen in Appendix F.

Power spectrum analysis shown in Figures 10.5 and 10.6 is consistent with the qualitative assessment regarding smoothing. Examples showing more detail can be seen in plots of July 6, 14

and 20, shown in Figures 10.7, 10.8 and 10.9. In the July 6 plot, a large increase in feeder demand resulting from a rapid decrease in wind turbine output results in a spike which is effectively removed by optimization. In the July 14 plot, oscillations occurring roughly on a two hour frequency are completely removed. Frequent step changes in feeder demand as a result of turbine cut-outs and cut-ins shown in the July 20 plot are also removed by the optimization.

Table 10.2: Performance metrics for Houston feeder load shaping optimization, high wind penetration case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	4.80	4.05	5.69
Peak Demand [MW]	-0.50	-1.84	0.10
Peak to Valley [%]	63.32	39.40	88.18
Load Factor [%]	5.63	1.59	13.86
Ramp [MW]	-2.52	-7.10	2.38

Table 10.3: Performance metrics for Houston feeder load shaping optimization, low wind penetration case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	4.85	4.20	5.37
Peak Demand [MW]	-0.25	-1.08	0.20
Peak to Valley [%]	79.75	65.78	93.54
Load Factor [%]	3.78	0.82	9.08
Ramp [MW]	-0.92	-4.13	2.06

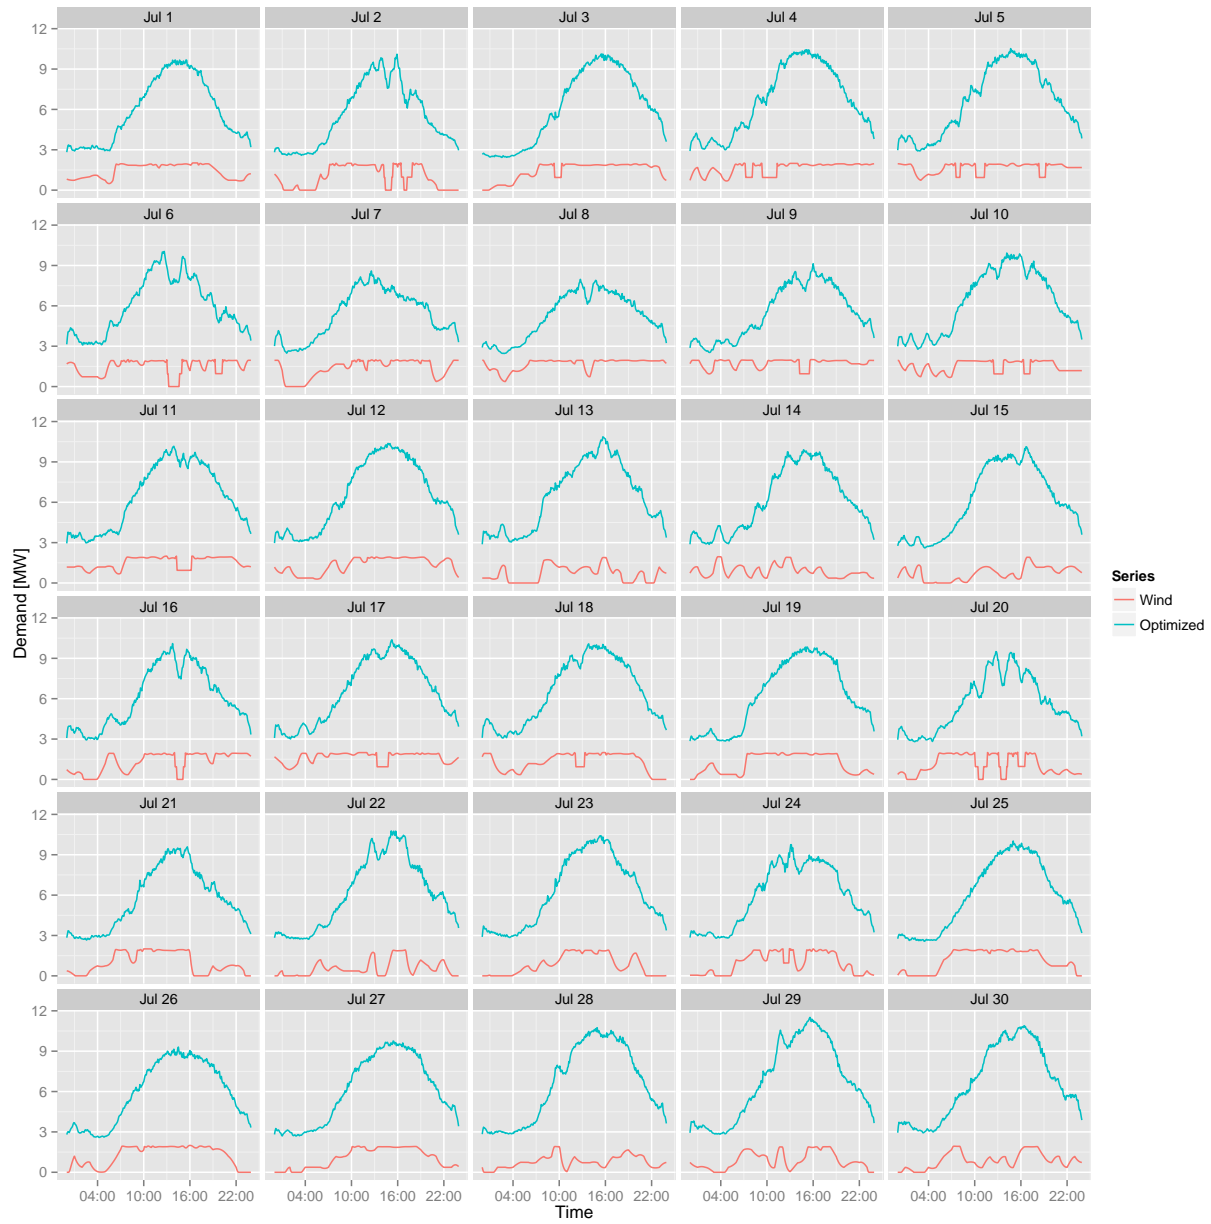


Figure 10.2: Pre-composited optimized feeder demand and wind production profile for Houston load shape optimization, high wind penetration case, 70% participation.




Figure 10.3: Feeder demand profiles for Houston load shape optimization, high wind penetration case, 70% participation.

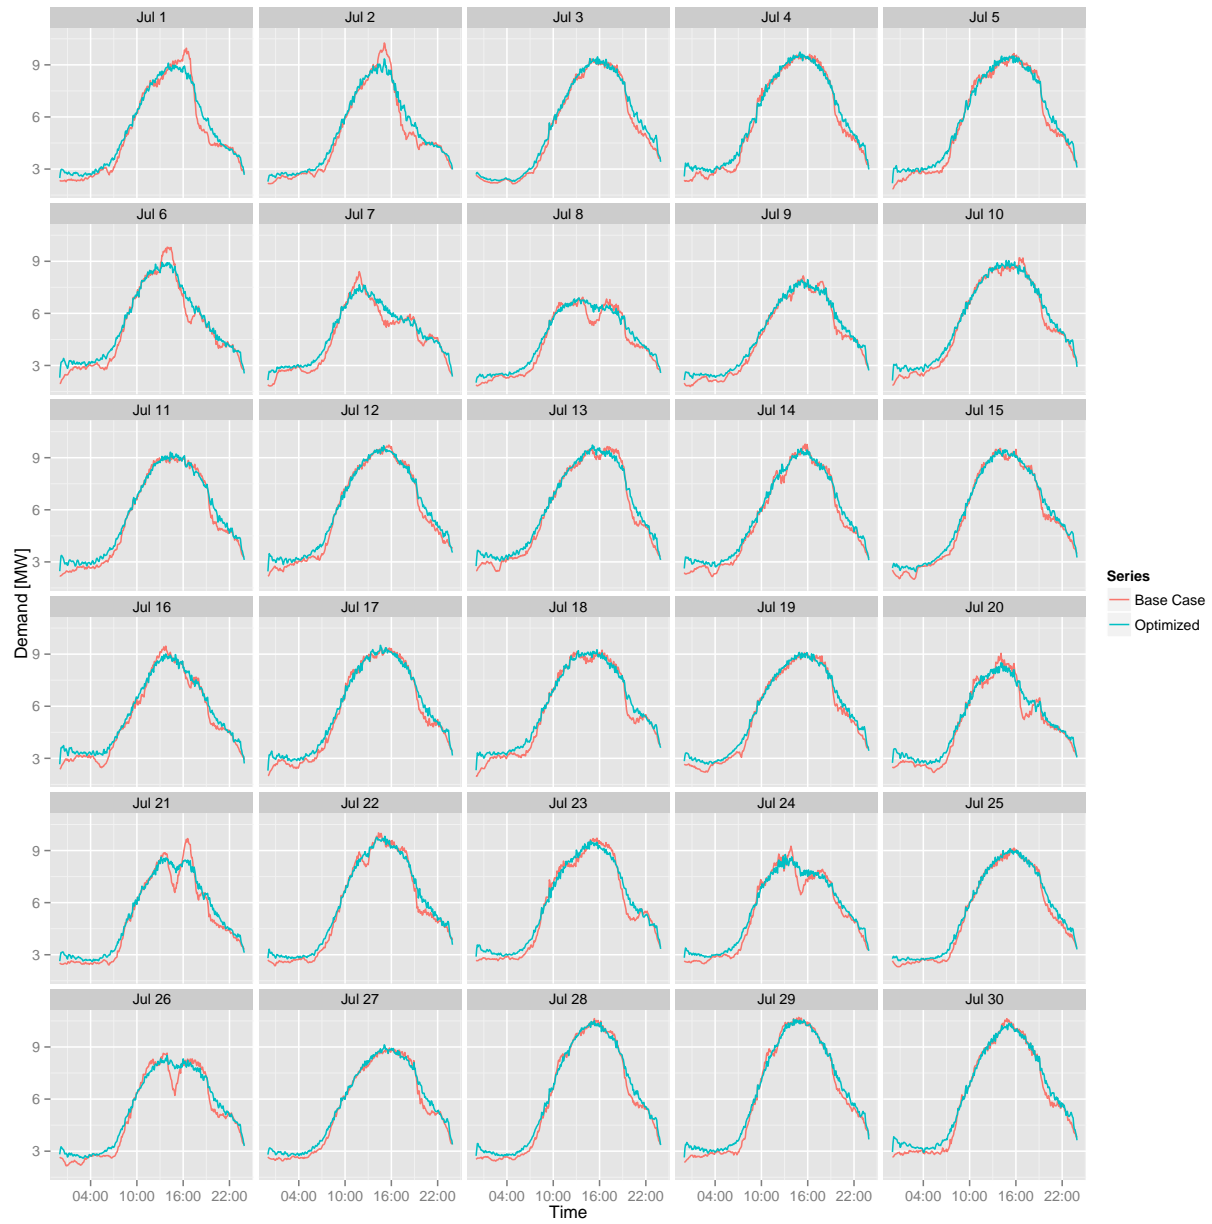


Figure 10.4: Feeder demand profiles for Houston load shape optimization, low wind penetration case, 70% participation.




Figure 10.5: Total spectral power as a function of frequency bin for Houston feeder load shape optimization, high wind penetration case, 70% participation.




Figure 10.6: Total spectral power as a function of frequency bin for Houston feeder load shape optimization, low wind penetration case, 70% participation.




Figure 10.7: Feeder demand profiles for Houston, July 6 load shape optimization, high wind penetration case, 70% participation.




Figure 10.8: Feeder demand profiles for Houston, July 14 load shape optimization, high wind penetration case, 70% participation.




Figure 10.9: Feeder demand profiles for Houston, July 20 load shape optimization, high wind penetration case, 70% participation.

10.6 Los Angeles

Los Angeles results indicate a limited ability of the controller to shape demand in any significant way. Nevertheless, the performance metrics and spectral power analysis show improvements across all dimensions, but with an increase in total energy consumption. Results for the 70% participation case can be seen in Table 10.4 and Figure 10.12. Power spectral analysis shows a large portion of power concentrated in the twenty-four to six hour frequency range. Optimization reduces this, shifting more to into the DC component, but variation across the observations is quite high.

Composite feeder demand curves, both base case and optimized, exhibit a large amount of variation as a result of the variation in wind generation (Figure 10.10). The effect appears to be much more marked in the Los Angeles feeder because a large percentage of the annual wind generation happens to be coincident with the time frame studied, and generation is not consistent. This, in combination with the lack of available flexible cooling demand presents a much higher hurdle for the controller.

Composite feeder demand in the low wind penetration case (Figure 10.11) is much less affected by the variability of generation. This does not change the amount of load the controller can shift, obviously, but does result in marginally smoother optimized demand curves. Performance for this case is relatively good. Improvements are shown in Table 10.5 and Figure 10.13).

Results for 30% participation cases follow the same trends as their 70% correlaries. These can be seen in Appendix F. Plots of individual days show that the controller is unable to shift demand in the early morning and late evening hours, although the actions do result in minimal shaving of peak demand, and filling of troughs. Figures 10.14 and 10.15 both show good examples of this behavior.

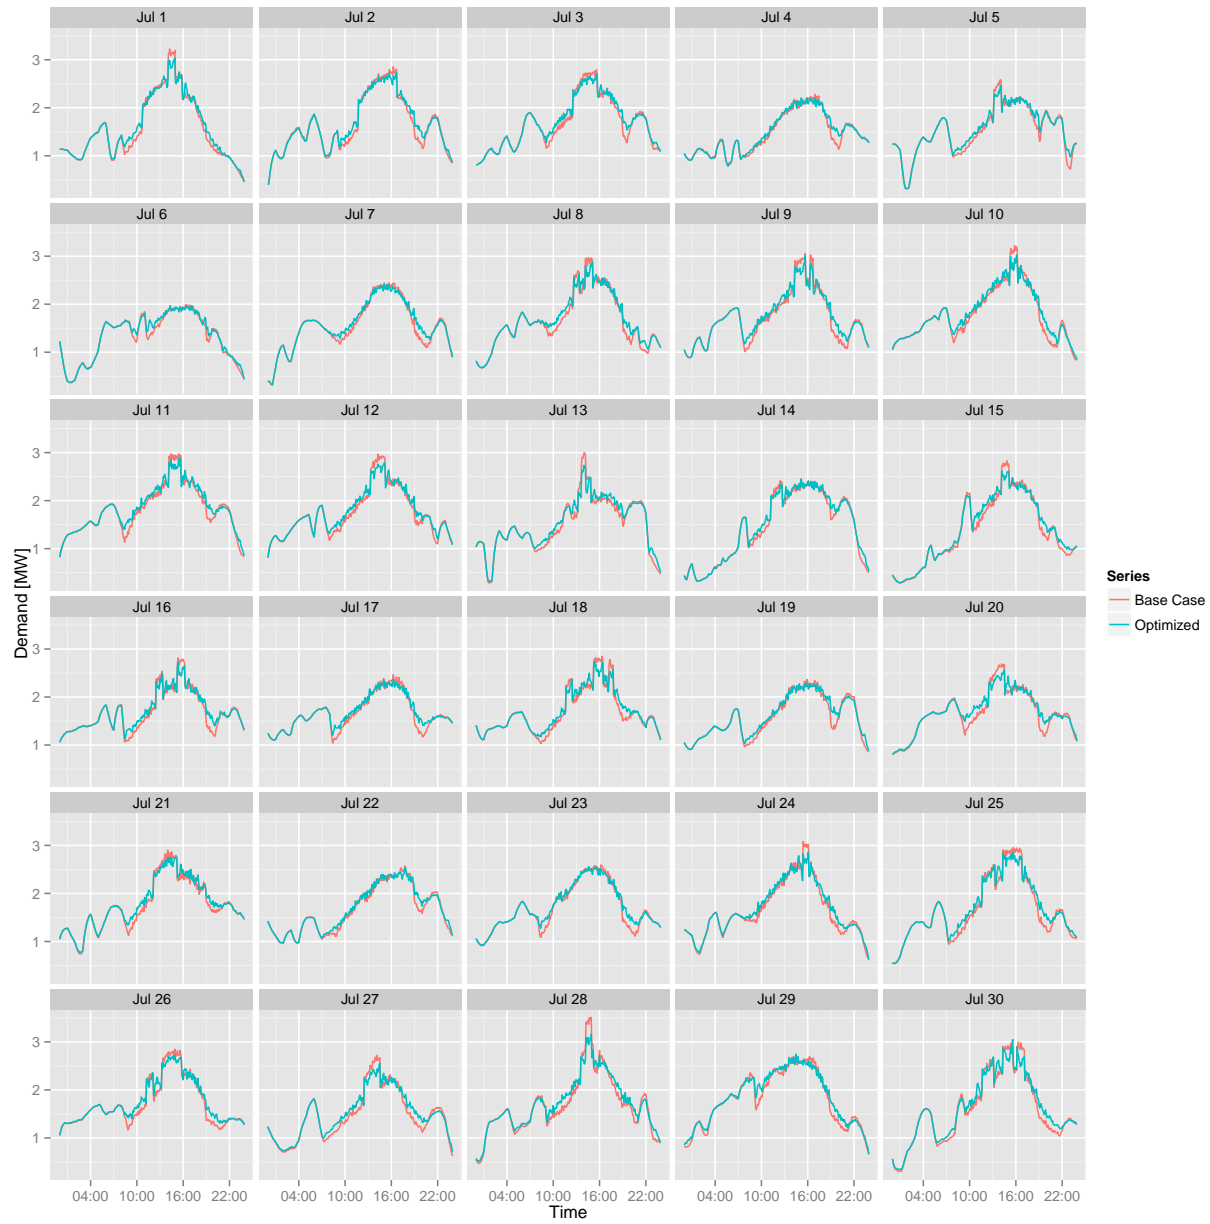


Figure 10.10: Feeder demand profiles for Los Angeles load shape optimization, high wind penetration case, 70% participation.

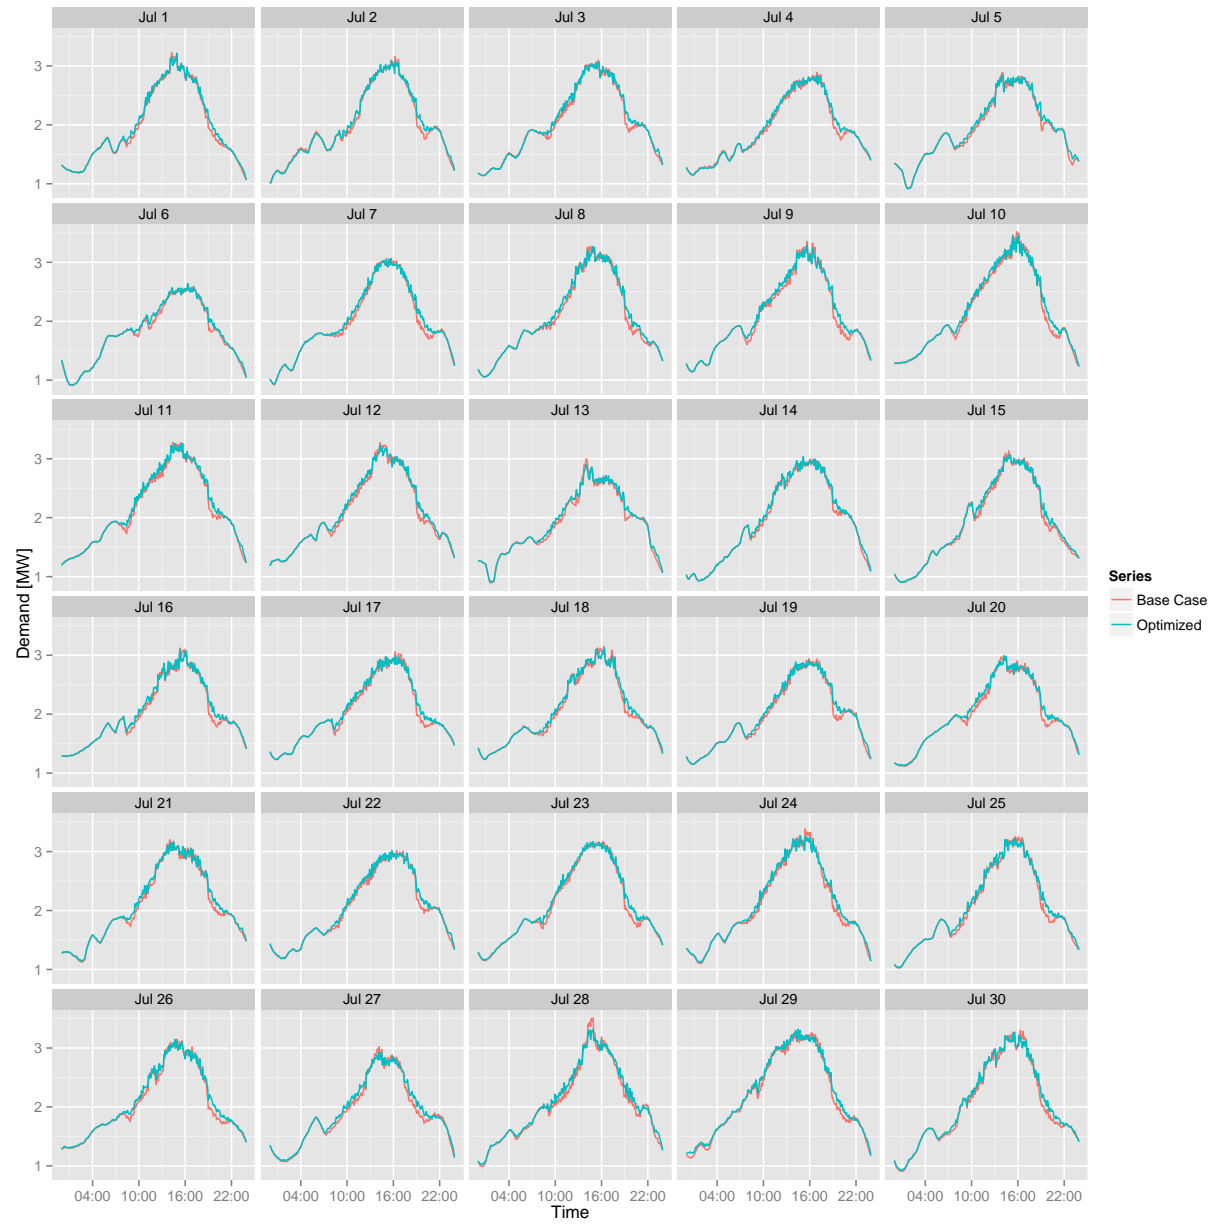


Figure 10.11: Feeder demand profiles for Los Angeles load shape optimization, low wind penetration case, 70% participation.

Table 10.4: Performance metrics for Los Angeles feeder load shaping optimization, high wind penetration case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	0.79	0.54	0.96
Peak Demand [MW]	-0.11	-0.38	0.00
Peak to Valley [%]	93.09	81.38	99.65
Load Factor [%]	3.65	1.38	7.24
Ramp [MW]	-0.44	-1.95	0.57

Table 10.5: Performance metrics for Los Angeles feeder load shaping optimization, low wind penetration case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	0.71	0.46	0.92
Peak Demand [MW]	-0.04	-0.17	0.02
Peak to Valley [%]	97.87	91.83	100.19
Load Factor [%]	1.74	0.74	4.17
Ramp [MW]	-0.23	-0.79	0.10




Figure 10.12: Total spectral power as a function of frequency bin for Los Angeles feeder load shape optimization, high wind penetration case, 70% participation.




Figure 10.13: Total spectral power as a function of frequency bin for Los Angeles feeder load shape optimization, low wind penetration case, 70% participation.




Figure 10.14: Feeder demand profiles for Los Angeles, July 9 load shape optimization, high wind penetration case, 70% participation.




Figure 10.15: Feeder demand profiles for Los Angeles, July 13 load shape optimization, high wind penetration case, 70% participation.

10.7 New York

In contrast to the Los Angeles feeder demand curves, New York exhibits much less variability due to a more consistent output by the wind turbines in both high and low wind scenarios (Figures 10.16 and 10.17). During many of the days, output from the generators is nearly constant at approximately 1MW. Bumps and dips in the base case demand curve relative to the no wind cases considered previously are much less prevalent than in Houston or Los Angeles. Still, some variation is introduced and the controller responds accordingly.

According to metrics in Tables 10.6 and 10.7, performance generally improves over the base case, but shows additional ramping. In contrast, power spectrum analysis suggests the opposite, with a shift towards greater DC power and reduced power in the twenty-four to six hour frequency range. A very marginal increase can be seen in the six to four hour range, but much smaller than any change in other ranges. Low participation results show the ramping metric to decrease compared to the base case, with similar power spectrum characteristics.

Smoothing of the composite base case demand can be seen in greater detail in Figures 10.20 and 10.21. Peak demand is reduced and oscillations in demand between midnight and 10:00 are generally eliminated in the July 3 case. The step increase and decrease seen in the July 10 plot is the result of cut-out as wind speed exceeds the cut-out speed of one of the two turbines. The controller is not able to remove this feature entirely, but provides some smoothing in the hours before and after to help reduce the magnitude of the step. Similarly, the controller reduces the steep change in demand at 19:00 by filling-in the hours between 19:00 and 23:00.

Table 10.6: Performance metrics for New York feeder load shaping optimization, high wind penetration case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	2.48	2.06	2.99
Peak Demand [MW]	-0.13	-0.44	0.03
Peak to Valley [%]	79.98	61.92	98.39
Load Factor [%]	3.59	1.76	6.99
Ramp [MW]	0.26	-1.27	1.77

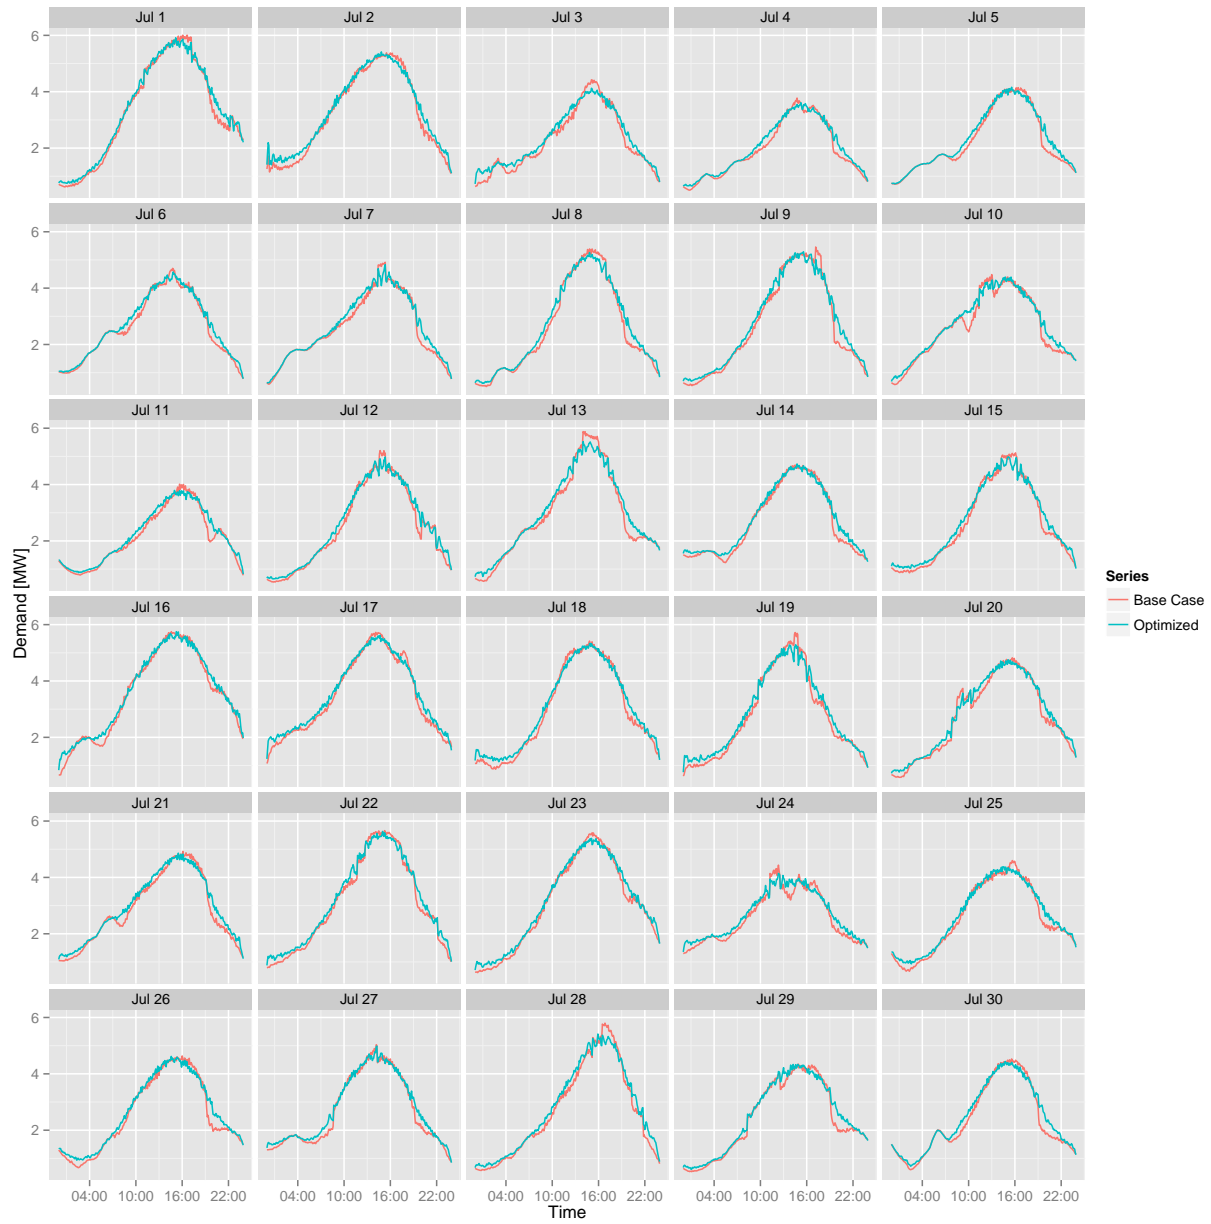


Figure 10.16: Feeder demand profiles for New York load shape optimization, high wind penetration case, 70% participation.

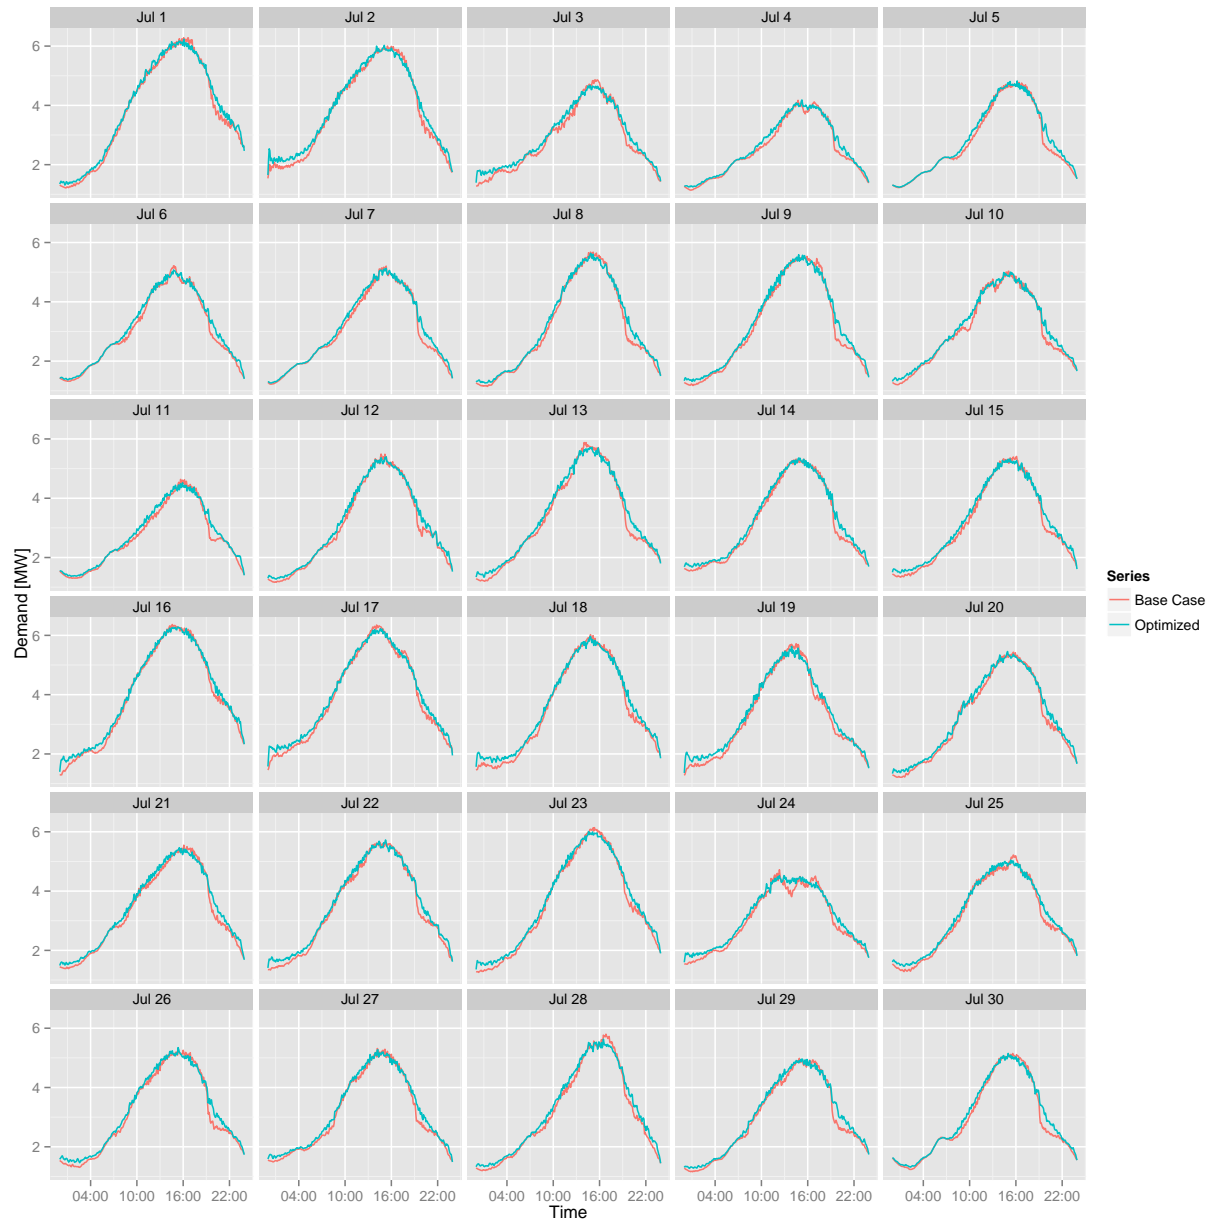


Figure 10.17: Feeder demand profiles for New York load shape optimization, low wind penetration case, 70% participation.

Table 10.7: Performance metrics for New York feeder load shaping optimization, low wind penetration case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	2.44	2.04	2.98
Peak Demand [MW]	-0.06	-0.23	0.08
Peak to Valley [%]	88.74	77.50	98.83
Load Factor [%]	2.58	1.15	5.38
Ramp [MW]	0.05	-0.96	0.73




Figure 10.18: Total spectral power as a function of frequency bin for New York feeder load shape optimization, high wind penetration case, 70% participation.




Figure 10.19: Total spectral power as a function of frequency bin for New York feeder load shape optimization, low wind penetration case, 70% participation.




Figure 10.20: Feeder demand profiles for New York, July 3 load shape optimization, high wind penetration case, 70% participation.




Figure 10.21: Feeder demand profiles for New York, July 10 load shape optimization, high wind penetration case, 70% participation.

10.8 Summary and Conclusions

The effectiveness of the controller in reducing wind-induced variation in the Houston case is impressive. Previous results suggested that the methodology may be well suited to eliminating demand fluctuations on the one to two hour time scale, provided that enough flexible cooling demand is available. These results provides further confirmation of those findings. Also consistent with previous findings, the effectiveness of the optimization varies across feeders, resulting in demand shaping limitations in the Los Angeles¹ feeder in particular.

It should be pointed out that the limitations encountered are not necessarily the fault of the method. Rather, as discussed in previous chapters, the limitations seem to be related to the size of the demand being shifted relative to the demand available, and the effectiveness of buildings to store cooling investments made to shift load. This implies that with larger flexible demands and more efficient storage, the method can be applied to demand shaping more effectively. Applications of this method to battery storage, plug-in hybrid vehicles and electric vehicles would confirm this hypothesis; this is a very interesting area of future research for which this methodology seems perfectly suited.

¹ That is not to say that the method would never work in Los Angeles. Rather, the effectiveness in Los Angeles is probably quite specific to the population of houses participating. This suggests an area of future research.

Chapter 11

Conclusions and Future Work

This work evaluates the effectiveness of model predictive control as a mechanism for shaping electric demand through residential HVAC set point manipulation. A decentralized MPC methodology is developed that allows the study of optimal control strategies applied to thousands of homes simultaneously. Using distribution network power flow simulation, the results of the controllers' actions are evaluated at the distribution feeder level to evaluate the aggregate impact on feeder demand. Three combinations of feeders and climates, representing several typical of those found in the United States, are studied. The work further considers the ability of the methodology to address feeder demand variability introduced when large amounts of solar and wind generation are present. The following chapter briefly summarizes the results and conclusions drawn from this research, followed by a discussion of potential areas of future work.

11.1 Summary of Results & Conclusions

Residential demand management is still very much a nascent idea. Pilot projects and small-scale demonstrations are the norm, due in part to the investment required to install the required hardware and implement the communication framework. The recent introduction and popularity of so-called “smart thermostats” has allowed much deeper penetration of technologies that enable demand management on a large scale. Yet, it is not clear what impacts large-scale implementation of demand management would yield. In this work, a very simple idealized demonstration shows that a large scale deployment of the current demand response methodology could result in excessive

demand spikes exceeding the capacity of the distribution network. These demand rebounds are known to utilities as a result of pilot projects, and present a challenge to implementation.

More intelligent solutions to the residential demand management problem exist. Through model predictive control, it is believed that demand rebounds may be mitigated or eliminated altogether. Further, it is believed that MPC may allow large groups of buildings to shape electric demand according to the needs of the utility: to level demand, remove variability, and potentially mitigate the issues associated with large amounts of renewable energy generation. The decentralized MPC methodology developed in this work allows investigations into demand management strategies to be studied at the distribution feeder level, incorporating the demand behavior of thousands of individual buildings.

Initial experiments using a simple demand limiting objective show the approach to be effective in reducing peak demand. Results vary according to the location, as the air conditioning load, i.e. the flexible cooling demand, differs from one building population to another. Although peak demand reductions are not as large as those measured in the demand response cases, the effect is more sustained, and the rebound eliminated. The smoothing of the feeder demand curve, a side effect of the demand limiting objective, shows the potential of the methodology for reducing demand variability. Several metrics, including peak to valley ratio, load factor and ramping, quantify the additional benefits to the demand limiting approach.

The work then evaluates the methodology applied to dynamic pricing scenarios. Hourly day-ahead pricing forecasts generated from historical price and weather data provide the driving signal for the model predictive controller, which attempts to minimize daily energy costs at each home. Day-ahead pricing has been proposed as a solution to supply-demand imbalances by encouraging the shift of electricity use to low priced portions of the day. The objective considered in this research unfortunately does not consider the larger ramifications of price-based optimization, resulting in aggregate demand behavior that introduces large fluctuations to the feeder demand. To the best of the author's knowledge, this is the first study of this type to be conducted using day-ahead price optimization without a price-feedback mechanism. An experiment where feedback is introduced

shows that oscillations in demand can be created and magnified using a price-based objective, suggesting that alternatives be explored to address demand shaping directly, rather than through a price-based proxy.

By directing the optimizer towards a demand shaping objective explicitly, the problems encountered in price-based optimization is avoided. The new approach employs a reference signal supplied to each home telling the home when to increase or decrease electric demand relative to normal operation. This load shaping optimization results in incremental adjustments by each home that in aggregate create the desired demand profile at the distribution feeder level. The approach shows an impressive ability to smooth the feeder demand profile on days containing demand variability. Peak demand is lowered in many cases, but because demand limiting is not the explicit objective, the reductions are not as great as those seen in the demand limiting optimization. Further refinement to the reference demand curve, i.e. the desired feeder demand shape, may result in additional demand reduction benefits. Other metrics such as peak to valley ratio, load factor and ramping show improvements consistent with the demand smoothing effects of the optimization. The new optimization approach, while still decentralized is considered to be “directed”. That is, although there is no centralized controller making control decisions for all homes, the formulation of the objective allows indirect coordination of the building population by virtue of the reference signal. This provides a template for real-world implementation that can be realized with today’s smart thermostats.

This methodology extends well to removing the variations in demand created when large amounts of rooftop solar electric generation is present. Traditionally, the effect of distributed solar generation has not been problematic for electric utilities due to low penetration levels. In parts of the United States, however, the rapid adoption of rooftop solar is beginning to create problems [21]. These include sustained, steep ramping and over-generation periods that are difficult for traditional thermal generators to accommodate. Results of simulations suggest that the methodology shows the greatest benefit in removing short term variations in demand on the order of one or two hours. The ability to remove the six to twelve hour demand features that solar generation can often produce

is limited, but the methodology does show some benefit. The limitations of the methodology are related to the amount of flexible cooling demand.

The final cases considered in this work investigate the ability of the methodology to shape demand according to needs beyond the distribution feeder. All previous investigations have considered shaping demand within the feeder assuming that the shape of demand at the distribution transformer is a proxy for system demand. In many scenarios, this is not a valid assumption as demand at different parts of the distribution system is a function of many factors, not all residential. The introduction of utility-scale wind generation at high penetration levels is such a case. Over-production by wind farms can result in curtailment, and changing weather conditions can wreak havoc [58] on traditional generation which must be rapidly ramped to adjust system output to offset the peaks and troughs created by the variation introduced. Similar to the previous results, the methodology demonstrates an ability to shape demand at the one and two hour time frame, virtually removing all short term variations in wind turbine output.

The use of model predictive control as a mechanism for shaping residential HVAC electrical demand shows promise in reducing peak demand, increasing load factor and smoothing the variability of renewable energy output when aggregated across a large number of buildings. Furthermore, the decentralized approach studied here provides a framework for a real-world approach which is both computationally tractable and implementable using current technologies. The choice of optimization objective is an important consideration, as unexpected consequences can result when the objective pursued by individual homes does not align well with the needs of the distribution system. The effectiveness of the approach is largely limited by the amount of flexible demand present in the system, a function of building population and characteristics.

11.2 Contributions & Future Work

Chapter 1 introduced a number of research questions that have directed this research. Each is summarized and given separate treatment below to highlight the important contributions of this work, limitations and suggestions for future research.

How would distributed HVAC MPC be performed and can this be a template for real-world implementation?

A significant contribution of this work has been the development and demonstration of a methodology that can be used to study the large-scale impacts of distributed residential MPC on the electric distribution system. The method studied in this work avoids the inherent difficulty of performing a centralized optimization with thousands of individual HVAC controllers by distributing the details of optimization to each controller instead. The method allows for global objectives to be realized by directing the actions of the individual controllers towards a common goal. Not only does this allow for more realistic models of system and controller, it provides a realistic solution to a very real problem that would be encountered in actual implementation.

It is not clear from this work that simulating and optimizing each building separately is truly required. For example, the entire residential building stock may be easily classified by half a dozen or so prototypical buildings. Considerable computation may be avoided by simply optimizing a few prototypical buildings and applying the optimized result to a larger population in simulation. That said, the diversity of responses from individual buildings may be the critical feature that yields the results seen in this work. In practice, the diversity of response due to occupant behavior may undermine any attempts at model reduction. Identifying how the methodology explored in this work could be simplified is a potential avenue for future research.

The method studied here can be considered a type of agent-based approach in that the controllers all make their own decisions about how best to operate. That is not to say that this approach results in the global best solution; proving that is not trivial. Indeed, there are likely opportunities to coordinate operation that this method may overlook. Studies comparing this method to a centralized method using the same level of modeling detail may help illuminate areas of improvement.

There are of course other alternatives using a more formal agent-based approach. These methods represent a hybrid of the two above: not entirely centralized, but not entirely distributed either. Comparison of classic agent-based methods with the current would illustrate the value in

limited coordination between agents. It must be noted, however, that these methods assume some inter-agent communication that complicates implementation considerably in any real-world application — privacy and security concerns notwithstanding. The latter are already a source of debate in the larger smart grid discussion.

What, is the signal or driving function that must be fed to individual homes such that the aggregate response meets the objectives of the utility?

This work investigates three driving functions that can be used by residential buildings to shape electric demand, evaluated using a set of performance metrics assumed to represent the needs common to load serving entities. These entities include electric utilities (regulated, deregulated, municipal and cooperative), electricity retailers and aggregators. A contribution of this work has been to show how these driving functions result in very different performance under a variety of conditions.

In the case of the dynamic-price investigations, this work highlights a very important finding which provides guidance for future work. The use of a proxy signal, price, to drive a particular aspect of grid operation makes the formulation of an objective very difficult. Day-ahead price optimization at the home level results in increased demand variability and large fluctuations, possibly destabilizing operation. In retrospect this is an obvious result because the homes are incentivized to adjust operation to avoid high price, regardless of the effect on absolute demand. Therefore any future investigations need to consider how utility goals align with those of the individual homes so that one does not subvert the other.

Demand limiting and load shaping investigations consider primarily the needs of the load serving entities, and do not address individual homes explicitly. Although performance metrics show a great potential for lowering peak demand, increasing load factor, and eliminating intra-day variability, the increase in total energy consumption suggests that costs to residents may increase under these control schemes. The increase in consumption implies a possible increase in emissions as well. Future work should consider the economic and environmental impacts of this approach to

provide a larger context to the results.

Can distributed MPC be used for day-ahead resource planning, given a set of installed traditional generation?

The work here assumes perfect forecast knowledge of building loads, weather and a whole host of other factors affecting the operation of the electrical grid. In theory, this method could be applied to multiple feeders to allow better day-ahead planning and eliminate the need for peaking plants on selected days. This work, however, does not provide an answer to this question. Rather, it provides a first step towards investigations which could. Future investigations which consider the uncertainty inherent in many of these assumptions are required before any definite conclusions can be drawn.

What impact would distributed MPC have on electric grid operation in the presence of distributed generation, and would distributed MPC allow a higher penetration of renewable energy to be utilized?

This work is the first to consider the use of residential MPC to address the inherent variability that large scale renewable energy integration introduces. This work has contributed an understanding of the limitations of residential HVAC demand shaping as it relates to variability occurring on different time scales. The cases investigated here indicate that variations on shorter time scales can be almost completely removed by the reference demand approach. Variations on longer time scales cannot be eliminated due to physical limitations associated with shifting cooling loads for long periods of time.

The separating line between short and long time frames is difficult to quantify in a single value. In fact, this line likely changes from day to day as a function of the amount of flexible cooling demand. Referring back to Chapter 5, it can be seen on the highest demand days that the system returns back to normal demand between 1-3 hours depending on the location. The distinction between short and long time frames has a relationship with this settling time. That

is, the length of time the controller is able to shift demand is only as long as this settling time. Thus, a short time frame is limited to this settling time. On mild days, this time may be longer; on hot days, shorter. Of course, this settling time is also function of the optimization boundaries and many other factors. Given this definition, anything greater than this settling time would be considered a long time frame.

The investigations studied here use day-ahead information for optimization with a full 24 hour planning horizon. Because the method studied here appears to work well at short time frames, future work should consider the use of this method as a means for short term curtailment similar to the investigations by Olivieri et al. [73]. Beyond curtailment, additional opportunities to accommodate much higher frequency variations in demand, both up and down, may exist using this method. For example, this work considered solar variations on the order of hours, but higher frequency variations are observed in practice. Shortening the planning horizon, simulation time step and decision variable length may allow extremely quick response to rapidly changing conditions caused by solar and wind variations, or outages. This may allow participation in ancillary electricity markets, e.g. spinning reserve, frequency regulation and voltage control.

With improved storage efficiency and capacity, it is likely that this decentralized MPC approach could be used to smooth diurnal variations in electric system demand. In fact, the generic nature of the method lends itself to many applications, not just HVAC. For example, the extension of the methodology to electric vehicle (EV) charging or residential battery storage is quite obvious. There is a large amount of research being performed on EVs as load shaping resources; to the best of the author's knowledge, the reference demand optimization developed here is not represented in the literature and would provide a valuable contribution to the body of work that already exists in that area. Similarly, the use of the technique as a micro-grid management strategy represents another fruitful area of application.

Bibliography

- [1] Annual Energy Review. United States Energy Information Administration, Washington, D.C., 2011.
- [2] Annual Energy Outlook 2012. United States Energy Information Administration, Washington, D.C., 2012.
- [3] EnergyPlus Engineering Reference: The Reference to EnergyPlus Calculations. United States Department of Energy, 2012.
- [4] Adaptive Computing, Inc. TORQUE Resource Manager. Retrieved from <http://www.adaptivecomputing.com/products/open-source/torque/>, December 2013.
- [5] U.S. Energy Information Administration. Residential Energy Consumption Survey (RECS). Retrieved from <http://www.eia.gov/consumption/residential/>, 2009.
- [6] American Society of Heating, Refrigerating and Air-Conditioning Engineers. ASHRAE Handbook: Fundamentals. American Society of Heating, Refrigerating and Air-Conditioning Engineers, Atlanta, Georgia, 2009.
- [7] Baltimore Gas and Electric Company. BGE Smart Energy Rewards. Retrieved from <http://www.bge.com/smartenergy/smart-energy-rewards/pages/>, January 2014.
- [8] Michael J. Brandemuehl. HVAC 2 Toolkit: A Toolkit for Secondary HVAC System Energy Calculations. American Society of Heating, Refrigerating and Air-Conditioning Engineers, Atlanta, Georgia, 1993.
- [9] James E. Braun. Reducing energy costs and peak electrical demand through optimal control of building thermal mass. ASHRAE Transactions, 96(2):264–273, 1990.
- [10] James E. Braun. Load Control Using Building Thermal Mass. Journal of Solar Energy Engineering, 125(3):292, 2003.
- [11] James E. Braun and Nitin Chaturvedi. An Inverse Gray-Box Model for Transient Building Load Prediction. HVAC&R Research, 8(1):73–99, 2002.
- [12] James E. Braun and Kyoung-ho Lee. Assessment of Demand Limiting Using Building Thermal Mass in Small Commercial Buildings. ASHRAE Transactions, 112(1):547–558, 2006.
- [13] James E. Braun, Kent W. Montgomery, and Nitin Chaturvedi. Evaluating the Performance of Building Thermal Mass Control Strategies. HVAC&R Research, 7(3):403–428, 2001.

- [14] Robert S Briggs, RG Lucas, and ZT Taylor. Climate Classification for Building Energy Codes and Standards: Part 2-Zone Definitions, Maps, and Comparisons. Transactions-American Society of Heating Refrigerating and Air Conditioning Engineers, 109(1):122–130, 2003.
- [15] Robert S Briggs, Robert G Lucas, and Z Todd Taylor. Climate Classification for Building Energy Codes and Standards: Part 1-Development Process. Transactions-American Society of Heating Refrigerating and Air Conditioning Engineers, 109(1):109–121, 2003.
- [16] Torsten Broeer, Jason Fuller, Francis Tuffner, David Chassin, and Ned Djilali. Modeling framework and validation of a smart grid and demand response system for wind power integration. Applied Energy, 113(C):199–207, 2014.
- [17] Tim Brom. Microwulf Software and Network Configuration Notes. Retrieved from <http://www.calvin.edu/~adams/research/microwulf/>, January 2011.
- [18] BuildingIQ. BuildingIQ - Next Generation Energy Management. Retrieved from <http://www.buildingiq.com/>, November 2012.
- [19] California Independent System Operator. Demand Response and Energy Efficiency Roadmap. Retrieved from <http://www.caiso.com/Documents/DR-EERoadmap.pdf>, December 2013.
- [20] California Independent System Operator. California ISO Open Access Same-time Information System (OASIS). Retrieved from <http://oasis.caiso.com/mrioasis/>, January 2014.
- [21] California Independent System Operator. Flexible Resource Adequacy Criteria and Must-Offer Obligation. Retrieved from <https://www.caiso.com/Documents/RevisedDraftFinalProposal-FlexibleRACriteriaMustOfferObligation-Clean.pdf>, March 2014.
- [22] Carrier Corporation. Controls & Thermostats : Carrier Home Comfort. Retrieved from <http://www.carrier.com/homecomfort/en/us/products/controls-and-thermostats/>, January 2014.
- [23] C. Cecati, C. Citro, A. Piccolo, and P. Siano. Smart Grids Operation with Distributed Generation and Demand Side Management. In Lingfeng Wang, editor, Modeling and Control of Sustainable Power Systems, Green Energy and Technology, pages 27–46. Springer Berlin Heidelberg, 2012.
- [24] D. P. Chassin, K. Schneider, and C. Gerkenmeyer. GridLAB-D: An open-source power systems modeling and simulation environment. In Transmission and Distribution Conference and Exposition, pages 1–5. IEEE, 2008.
- [25] DP Chassin, P Du, and JC Fuller. The potential and limits of residential demand response control strategies. In Power and Energy Society General Meeting, pages 1–6. IEEE, 2011.
- [26] Clean Urban Energy. Clean Urban Energy. Retrieved from <http://www.cleanurbanenergy.com/>, November 2012.
- [27] Brian Coffey, Fariborz Haghighat, Edward Morofsky, and Edward Kutrowski. A software framework for model predictive control with GenOpt. Energy and Buildings, 42(7):1084–1092, July 2010.

- [28] Charles D. Corbin, Gregor P. Henze, and Peter May-Ostendorp. A model predictive control optimization environment for real-time commercial building application. Journal of Building Performance Simulation, 6(3):159–174, 2013.
- [29] David Corbus, J King, T Mousseau, R Zavadil, B Heath, L Hecker, J Lawhorn, D Osborn, J Smit, R Hunt, et al. Eastern wind integration and transmission study. Technical report, National Renewable Energy Laboratory, Golden, Colorado, 2010.
- [30] D. B. Crawley, L. K. Lawrie, F. C. Winkelmann, W. F. Buhl, Y. J. Huang, C. O. Pedersen, R. K. Strand, R. J. Liesen, D. E. Fisher, M. J. Witte, and J. Glazer. EnergyPlus: creating a new-generation building energy simulation program. Energy & Buildings, 33(4):319–331, 2001.
- [31] T. Dewson, B. Day, and A.D. Irving. Least squares parameter estimation of a reduced order thermal model of an experimental building. Building and Environment, 28(2):127–137, April 1993.
- [32] James Druitt and Wolf-Gerrit Früh. Simulation of demand management and grid balancing with electric vehicles. Journal of Power Sources, 216(0):104–116, 2012.
- [33] R. Eberhart and J. Kennedy. A new optimizer using particle swarm theory. Proceedings of the Sixth International Symposium on Micro Machine and Human Science, pages 39–43, 1995.
- [34] EcoFactor. EcoFactor - A personalized solution for managing residential energy use. Retrieved from <http://www.ecofactor.com/>, November 2012.
- [35] Electric Reliability Council of Texas. Day-Ahead Market. Retrieved from <http://www.ercot.com/mktinfo/dam/>, January 2014.
- [36] Executive Office of the President. The President’s Climate Action Plan. Retrieved from <http://www.whitehouse.gov/sites/default/files/image/president27sclimateactionplan.pdf>, June 2013.
- [37] C. Felsmann, G. Knabe, and H. Werdin. Simulation of domestic heating systems by integration of TRNSYS in a MATLAB/SIMULINK model. In Proc. 6th Int. Conference System Simulation in Buildings, pages 16–18, Liege, Belgium, 2002.
- [38] J. C. Fuller, K. P. Schneider, and D. Chassin. Analysis of Residential Demand Response and Double-Auction Markets. In Power and Energy Society General Meeting, pages 1–7, San Diego, California, 2011. IEEE.
- [39] M.M. Gouda, S. Danaher, and C.P. Underwood. Building thermal model reduction using nonlinear constrained optimization. Building and Environment, 37(12):1255–1265, December 2002.
- [40] Gregor P. Henze. Energy and Cost Minimal Control of Active and Passive Building Thermal Storage Inventory. Journal of Solar Energy Engineering, 127(3):343, 2005.

- [41] Gregor P. Henze, Michael J. Brandemuehl, and Clemens Felsmann. ASHRAE Research Project 1313-RP: Evaluation of Building Thermal Mass Savings. Technical report, American Society of Heating, Refrigerating and Air Conditioning Engineers, Atlanta, Georgia, 2007.
- [42] Gregor P. Henze, Clemens Felsmann, and Gottfried Knabe. Evaluation of optimal control for active and passive building thermal storage. International Journal of Thermal Sciences, 43(2):173–183, February 2004.
- [43] Gregor P. Henze, Anthony R. Florita, Michael J. Brandemuehl, Clemens Felsmann, and Hwakong Cheng. Advances in Near-Optimal Control of Passive Building Thermal Storage. Journal of Solar Energy Engineering, 132(2):021009, 2010.
- [44] Gregor P. Henze, Doreen Kalz, Simeng Liu, and Clemens Felsmann. Experimental Analysis of Model-Based Predictive Optimal Control for Active and Passive Building Thermal Storage Inventory. HVAC&R Research, 11(2):189–213, April 2005.
- [45] Haitham Hindi, Daniel Greene, and Caitlin Laventall. Coordinating regulation and demand response in electric power grids using multirate model predictive control. In Innovative Smart Grid Technologies, pages 1–8, Atlanta, Georgia, 2011. IEEE.
- [46] Bri-Mathias S Hodge, Shisheng Huang, Aviral Shukla, Joseph F Pekny, and Gintaras V Reklaitis. The effects of vehicle-to-grid systems on wind power integration in california. Computer Aided Chemical Engineering, 28:1039–1044, 2010.
- [47] Tanguy Hubert and Santiago Grijalva. Realizing smart grid benefits requires energy optimization algorithms at residential level. In Innovative Smart Grid Technologies, pages 1–8, Atlanta, Georgia, January 2011. IEEE.
- [48] Honeywell International Inc. Honeywell Thermostats. Retrieved from <http://yourhome.honeywell.com/home/Products/Thermostats/>, January 2014.
- [49] Ron Judkoff and Joel Neymark. International Energy Agency building energy simulation test (BESTEST) and diagnostic method. Technical report, National Renewable Energy Laboratory, Golden, Colorado, 1995.
- [50] Ron Judkoff, Ben Polly, Marcus Bianchi, and Joel Neymark. Building Energy Simulation Test for Existing Homes (BESTEST-EX) Phase 1 Test Procedure : Building Thermal Fabric Cases. Technical report, National Renewable Energy Laboratory, Golden, Colorado, 2010.
- [51] Willett Kempton and Jasna Tomić. Vehicle-to-grid power implementation: From stabilizing the grid to supporting large-scale renewable energy. Journal of Power Sources, 144(1):280–294, 2005.
- [52] James Kennedy and Russell C. Eberhart. A discrete binary version of the particle swarm algorithm. In IEEE International Conference on Systems, Man, and Cybernetics., pages 4–8. IEEE, 1997.
- [53] Hamid Khayyam, Hassan Ranjbarzadeh, and Vincenzo Marano. Intelligent control of vehicle to grid power. Journal of Power Sources, 201:1–9, 2012.

- [54] D.N. 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, pages 1–8. IEEE, July 2008.
- [55] K. Lee and James E. Braun. Development and Application of an Inverse Building Model for Demand Response in Small Commercial Buildings. In Proceedings of SimBuild, pages 1–12, Boulder, Colorado, 2004. IBPSA-USA.
- [56] Kyoung-ho Lee and James E. Braun. Reducing Peak Cooling Loads through Model-Based Control of Zone Temperature Setpoints. In American Control Conference, pages 5070–5075, Seoul, Korea, July 2007. IEEE.
- [57] Kyoung-ho Lee and James E. Braun. Model-based demand-limiting control of building thermal mass. Building and Environment, 43(10):1633–1646, October 2008.
- [58] Debra Lew, Lori Bird, Michael Milligan, Bethany Speer, Xi Wang, Enrico Maria Carlini, Ana Estanqueiro, Damian Flynn, Emilio Gomez-Lazaro, Nickie Menemenlis, et al. Wind and Solar Curtailment. Technical report, National Renewable Energy Laboratory, Golden, Colorado, 2013.
- [59] Debra Lew, Greg Brinkman, E Ibanez, BM Hodge, and J King. The western wind and solar integration study phase 2. Technical report, National Renewable Energy Laboratory, Golden, Colorado, 2013.
- [60] Na Li, Lijun Chen, and Steven H. Low. Optimal demand response based on utility maximization in power networks. In Power and Energy Society General Meeting, pages 1–8, Pasadena, California, July 2011. IEEE.
- [61] B. Y. H. Liu and R. C. Jordan. The interrelationship and characteristic distribution of direct, diffuse and total solar radiation. Solar Energy, 4(3):1–19, 1960.
- [62] Bill Marion and Mary Anderberg. PVWATTS - an online performance calculator for grid-connected PV systems. In Proceedings of the Solar Conference, pages 119–124. American Solar Energy Society; American Institute of Architects, 2000.
- [63] E. H. Mathews, P. G. Richards, and C. Lombard. A first-order thermal model for building design. Energy and Buildings, 21(2):133–145, 1994.
- [64] Peter T. May-Ostendorp. Offline Model Predictive Control of Mixed Mode Buildings for Near-Optimal Supervisory Control Strategy Development. PhD thesis, University of Colorado at Boulder, 2012.
- [65] Peter T. May-Ostendorp, Gregor P. Henze, Charles D. Corbin, Balaji Rajagopalan, and Clemens Felsmann. Model-predictive control of mixed-mode buildings with rule extraction. Building and Environment, 46(2):428–437, February 2011.
- [66] Peter T. May-Ostendorp, Gregor P. Henze, Balaji Rajagopalan, and Charles D. Corbin. Extraction of supervisory building control rules from model predictive control of windows in a mixed mode building. Journal of Building Performance Simulation, 6(3):199–219, 2013.

- [67] Marc-André Moffet, Frédéric Sirois, and David Beauvais. Review of Open-Source Code Power Grid Simulation Tools for Long-Term Parametric Simulation. Technical report, CanmetENERGY, 2011.
- [68] F. B. Morris, James E. Braun, and S. J. Treado. Experimental and Simulated Performance of Optimal Control of Building Thermal Storage. *ASHRAE Transactions*, 100(1):402–414, 1994.
- [69] Nest Labs. Nest - The Learning Thermostat. Retrieved from <http://www.nest.com/>, November 2012.
- [70] New York Independent System Operator. NYISO (Markets & Operations - Market Data - Pricing Data). Retrieved from http://www.nyiso.com/public/markets_operations/market_data/pricing_data/, January 2014.
- [71] NV Energy. mPowered: Home Energy Management for Northern Nevada. Retrieved from <https://www.nvenergy.com/home/saveenergy/rebates/mpowered/>, January 2014.
- [72] Office of the Press Secretary. President to Attend Copenhagen Climate Talks. Retrieved from <http://www.whitehouse.gov/the-press-office/president-attend-copenhagen-climate-talks>, November 2009.
- [73] Simon J. Olivieri, Gregor P. Henze, Chad D. Corbin, and Michael J. Brandemuehl. Evaluation of commercial building demand response potential using optimal short-term curtailment of heating, ventilation, and air-conditioning loads. *Journal of Building Performance Simulation*, 7(2):100–118, 2014.
- [74] Zhang O'Neill, Satish Narayanan, and Rohini Brahme. Model-based Thermal Load Estimation in Buildings. In *Proceedings of SimBuild*, pages 474–481, New York, New York, 2010. IBPSA-USA.
- [75] OpenDSS. OpenDSS - Electric Power Distribution System Simulator. Retrieved from <http://sourceforge.net/projects/electricdss/>, November 2012.
- [76] Philippe Junior Ossoucah. APREM—Analyse paramétrique des réseaux électriques avec Matlab. Master's thesis, École Polytechnique de Montréal, 2010.
- [77] Raul J. Martinez Oviedo, Zhong Fan, Sedat Gormus, Parag Kulkarni, and Dritan Kaleshi. Residential energy demand management in smart grids. In *Transmission and Distribution Conference and Exposition*, pages 1–8, Bristol, UK, May 2012. IEEE.
- [78] Pecan Street Research Institute. Pecan Street Inc. Retrieved from <http://www.pecanstreet.org/>, January 2014.
- [79] Therese Pepper, Marco Pritoni, Alan Meier, Cecilia Aragon, and Daniel Perry. How people use thermostats in homes: A review. *Building and Environment*, 46(12):2529–2541, 2011.
- [80] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2013.

- [81] RG Ross Jr. Flat-plate photovoltaic array design optimization. In 14th Photovoltaic Specialists Conference, volume 1, pages 1126–1132, 1980.
- [82] KP Schneider, Y Chen, DP Chassin, RG Pratt, DW Engel, and SE Thompson. Modern grid initiative-distribution taxonomy. Technical report, Pacific Northwest National Laboratory, Richland, Washington, 2008.
- [83] Solar Energy Industries Association. Solar Market Insight Report 2013 Year in Review. Technical report, Solar Energy Industries Association, March 2014.
- [84] Solar Energy Laboratory. TRNSYS 16–A TRaNsient SYstem Simulation program. University of Wisconsin, Madison, Wisconsin, 2004.
- [85] Robert C Sonderegger. Dynamic models of house heating based on equivalent thermal parameters. PhD thesis, Princeton University, 1978.
- [86] Ryan A. Tanner and Gregor P. Henze. Stochastic control optimization for a mixed mode building considering occupant window opening behaviour. Journal of Building Performance Simulation, 7(6):427–444, 2014.
- [87] The MathWorks, Inc. MathWorks - MATLAB and Simulink for Technical Computing. Retrieved from <http://www.mathworks.com/>, November 2012.
- [88] Jianhui Wang, Cong Liu, Dan Ton, Yan Zhou, Jinho Kim, and Anantray Vyas. Impact of plug-in hybrid electric vehicles on power systems with demand response and wind power. Energy Policy, 39(7):4016–4021, 2011.
- [89] Shengwei Wang and Xinhua Xu. Parameter estimation of internal thermal mass of building dynamic models using genetic algorithm. Energy Conversion and Management, 47(13-14):1927–1941, August 2006.
- [90] Xing Wang and Avnaesh Jayantilal. Dispatch wind generation and demand response. In Power and Energy Society General Meeting, San Diego, California, 2011. IEEE.
- [91] Weather Underground. Weather Forecast & Reports - Long Range & Local. Retrieved from <http://www.wunderground.com/>, January 2014.
- [92] Claire Weiller. Plug-in hybrid electric vehicle impacts on hourly electricity demand in the United States. Energy Policy, 39(6):3766–3778, 2011.
- [93] Michael Wetter. GenOpt, Generic Optimization Program. In Seventh International IBPSA Conference, pages 601–608, Rio de Janeiro, Brazil, 2001. IBPSA.
- [94] Michael Wetter and Philip Haves. A modular building controls virtual test bed for the integration of heterogeneous systems. In Proceedings of SimBuild, pages 69–76, Berkeley, California, 2008. IBPSA-USA.
- [95] S. Widergren, K. Subbarao, D. Chassin, J. C. Fuller, and Robert Pratt. Residential real-time price response simulation. In Power and Energy Society General Meeting, pages 1–5, San Diego, California, 2011. IEEE.

- [96] Xcel Energy Inc. Xcel Energy - Saver's Switch for Residences - CO. Retrieved from https://www.xcelenergy.com/Save_Money_&_Energy/Residential/Heating_&_Cooling/Saver's_Switch_for_Residences_-_CO, January 2014.
- [97] Xinhua Xu and Shengwei Wang. Optimal simplified thermal models of building envelope based on frequency domain regression using genetic algorithm. *Energy and Buildings*, 39(5):525–536, May 2007.
- [98] Xinhua Xu and Shengwei Wang. A simplified dynamic model for existing buildings using CTF and thermal network models. *International Journal of Thermal Sciences*, 47(9):1249–1262, September 2008.
- [99] Zeke Yewdall. Active Cooling and Heat Recovery for Building Integrated Photovoltaics. Master's thesis, University of Colorado at Boulder, 2003.
- [100] Wei Zhang, Karanjit Kalsi, and Jason Fuller. Aggregate Model for Heterogeneous Thermostatically Controlled Loads with Demand Response. In *Power and Energy Society General Meeting*, pages 1–8, Columbus, Ohio, 2012. IEEE.
- [101] G. Zhou, P. Ihm, M. Krarti, S. Liu, and G. P. Henze. Integration of an internal optimization module within EnergyPlus. In *Eighth International IBPSA Conference*, pages 1475–1482, Eindhoven, Netherlands, 2003. IBPSA.

Appendix A

Supplement to Demand Response

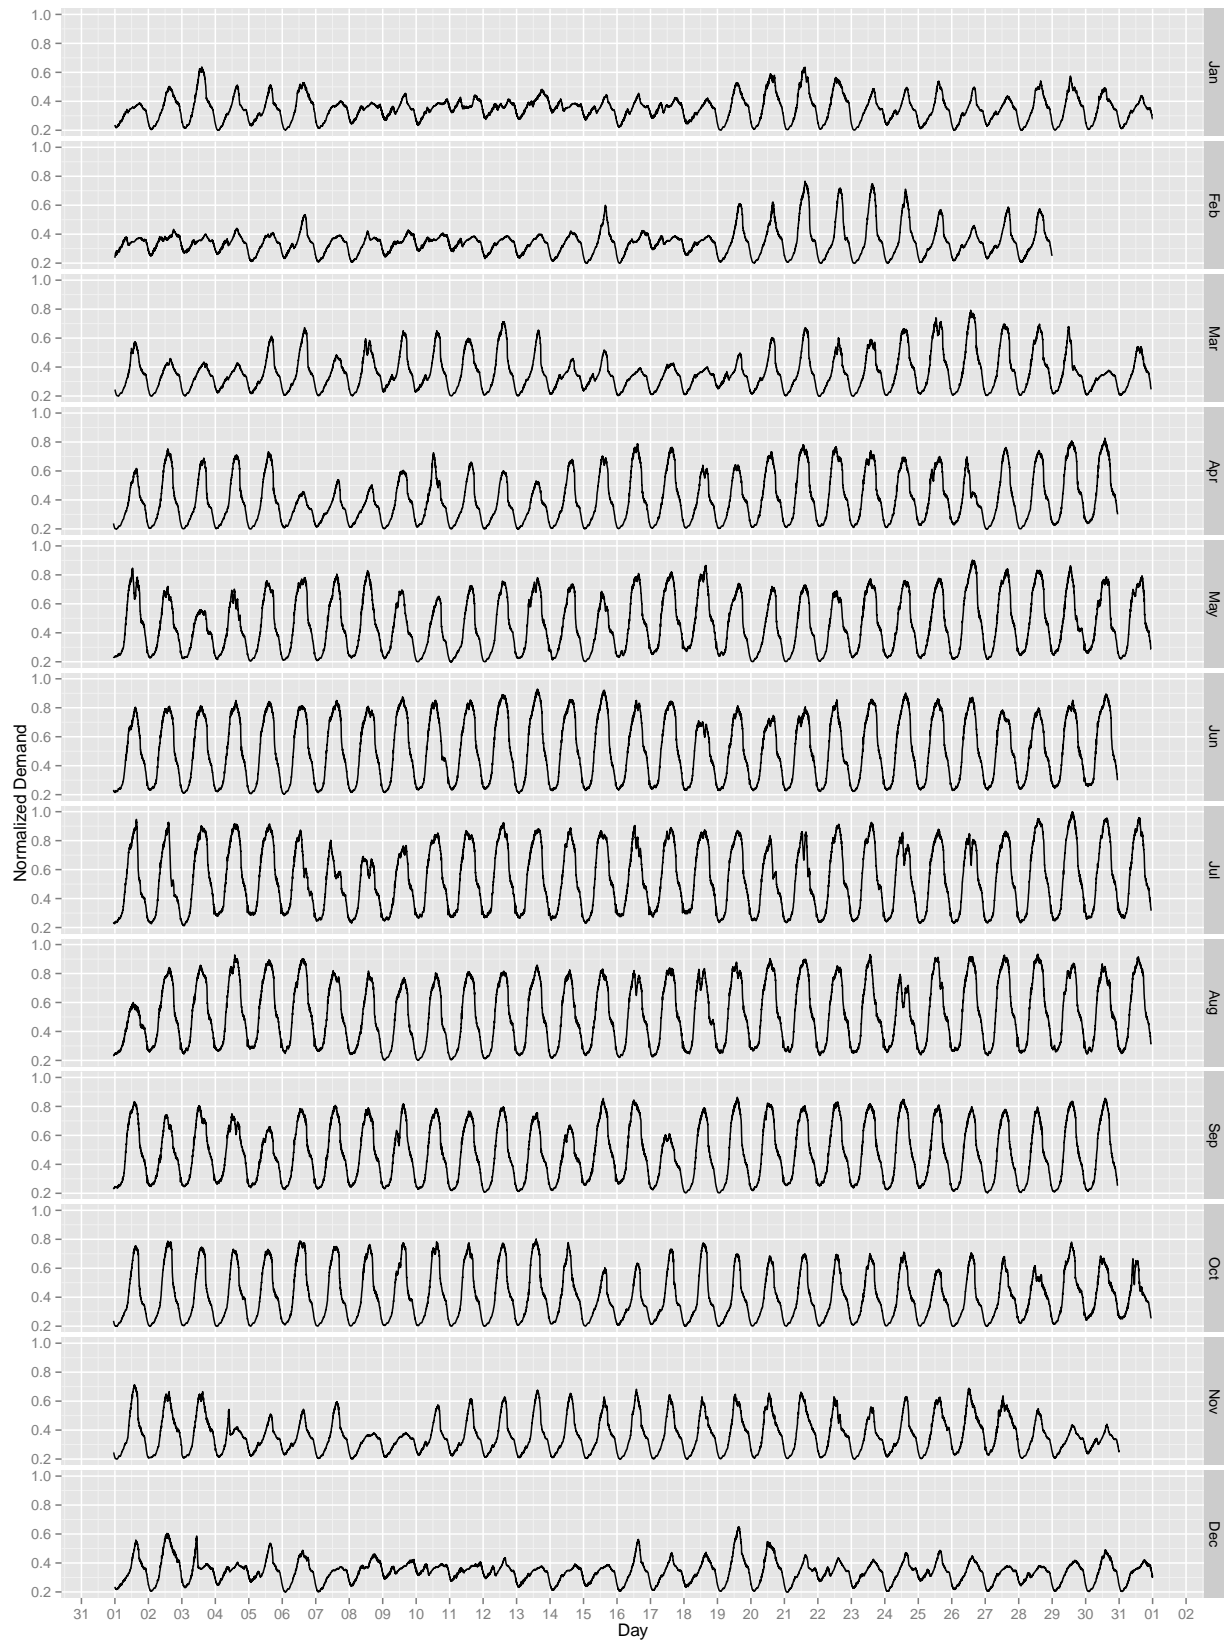


Figure A.1: Normalized annual Houston feeder demand profile.




Figure A.2: Normalized annual Los Angeles feeder demand profile.




Figure A.3: Normalized annual New York feeder demand profile.

Appendix B

Supplement to Demand Limiting Optimization

B.1 Houston

Table B.1: Performance metrics for Houston feeder demand limiting optimization, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	1.60	1.26	1.96
Peak Demand [MW]	-0.26	-0.49	-0.14
Peak to Valley [%]	88.36	79.84	95.63
Load Factor [%]	2.30	1.61	3.49
Ramp [MW]	-0.53	-1.75	0.55




Figure B.1: Total spectral power as a function of frequency bin for Houston feeder demand limiting optimization, 30% participation.

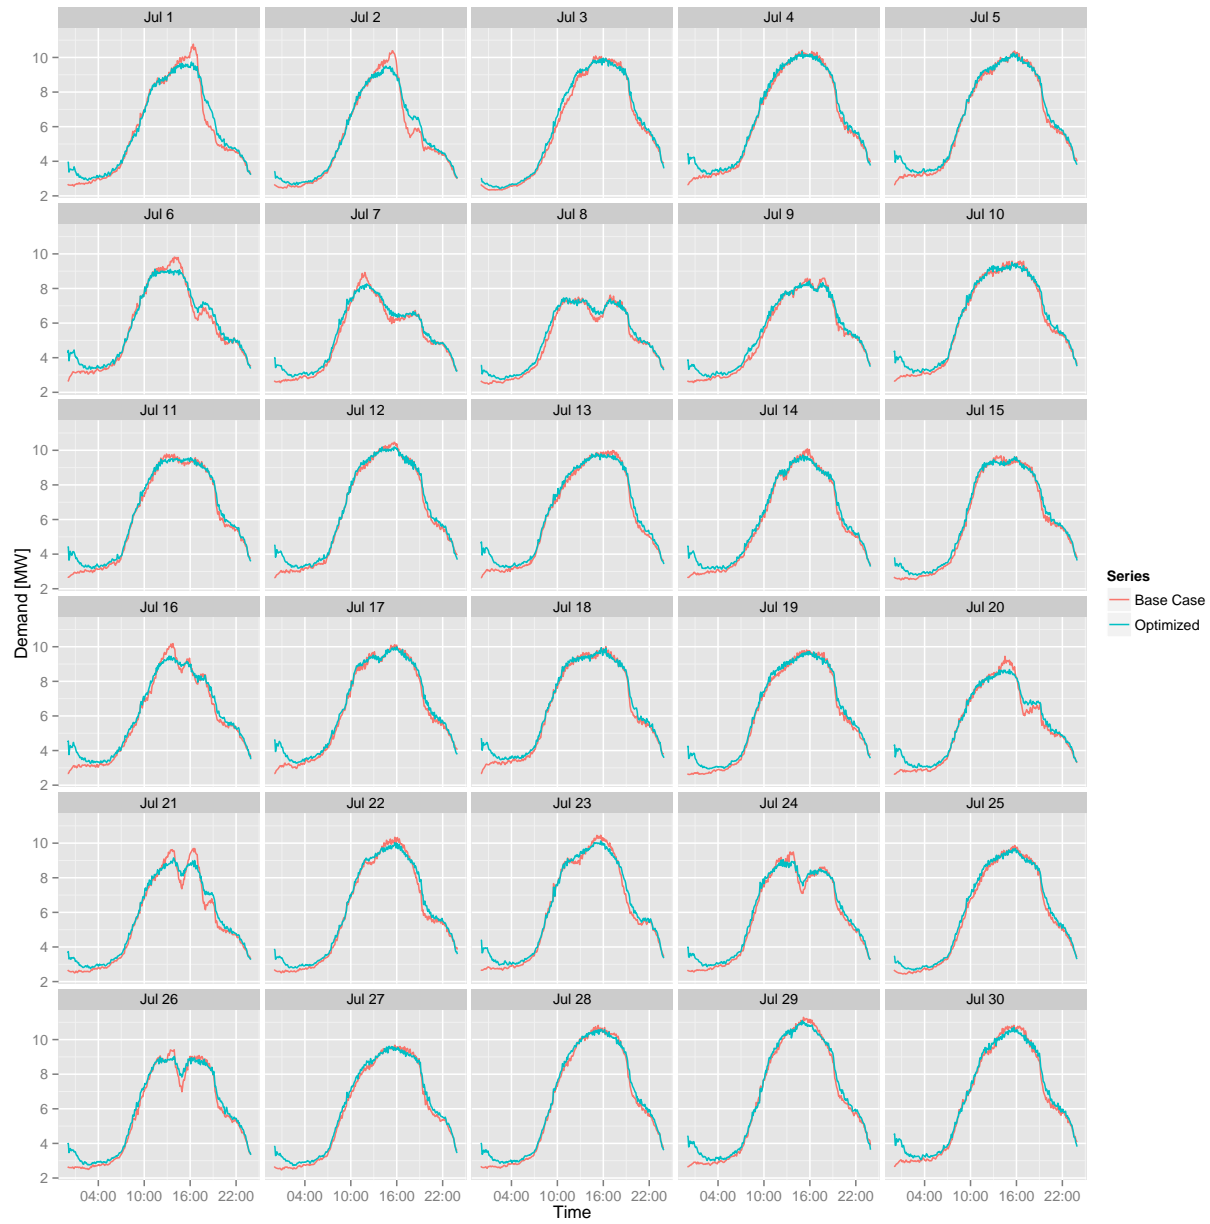


Figure B.2: Feeder demand profiles for Houston demand limiting optimization, 15 minute moving average, 70% participation.

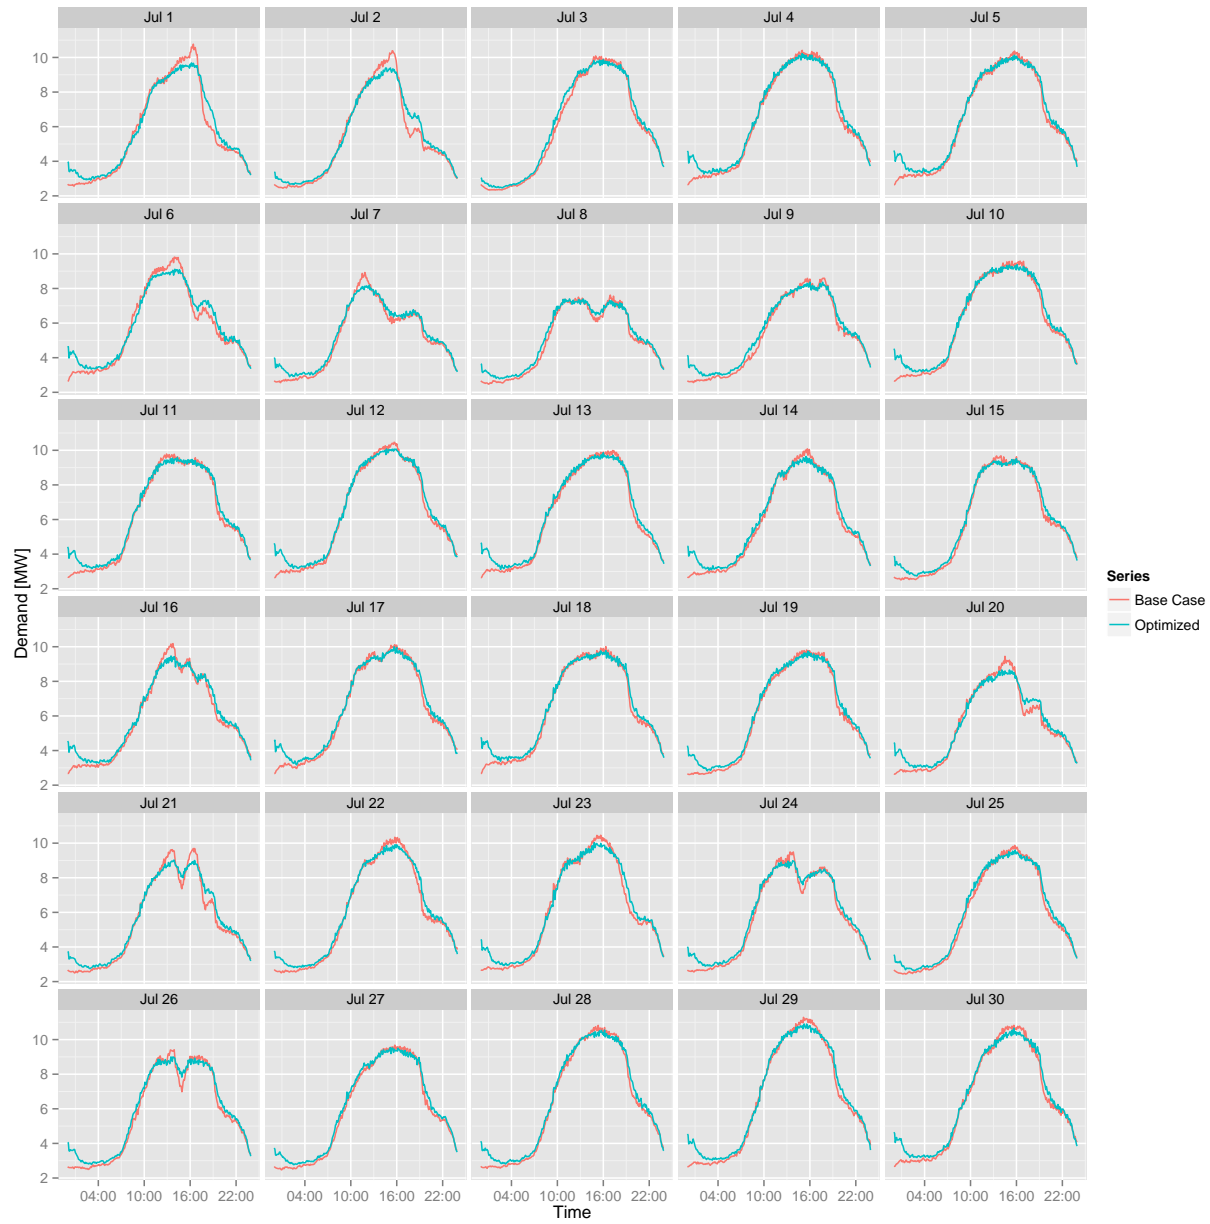


Figure B.3: Feeder demand profiles for Houston demand limiting optimization, 30 minute moving average, 70% participation.

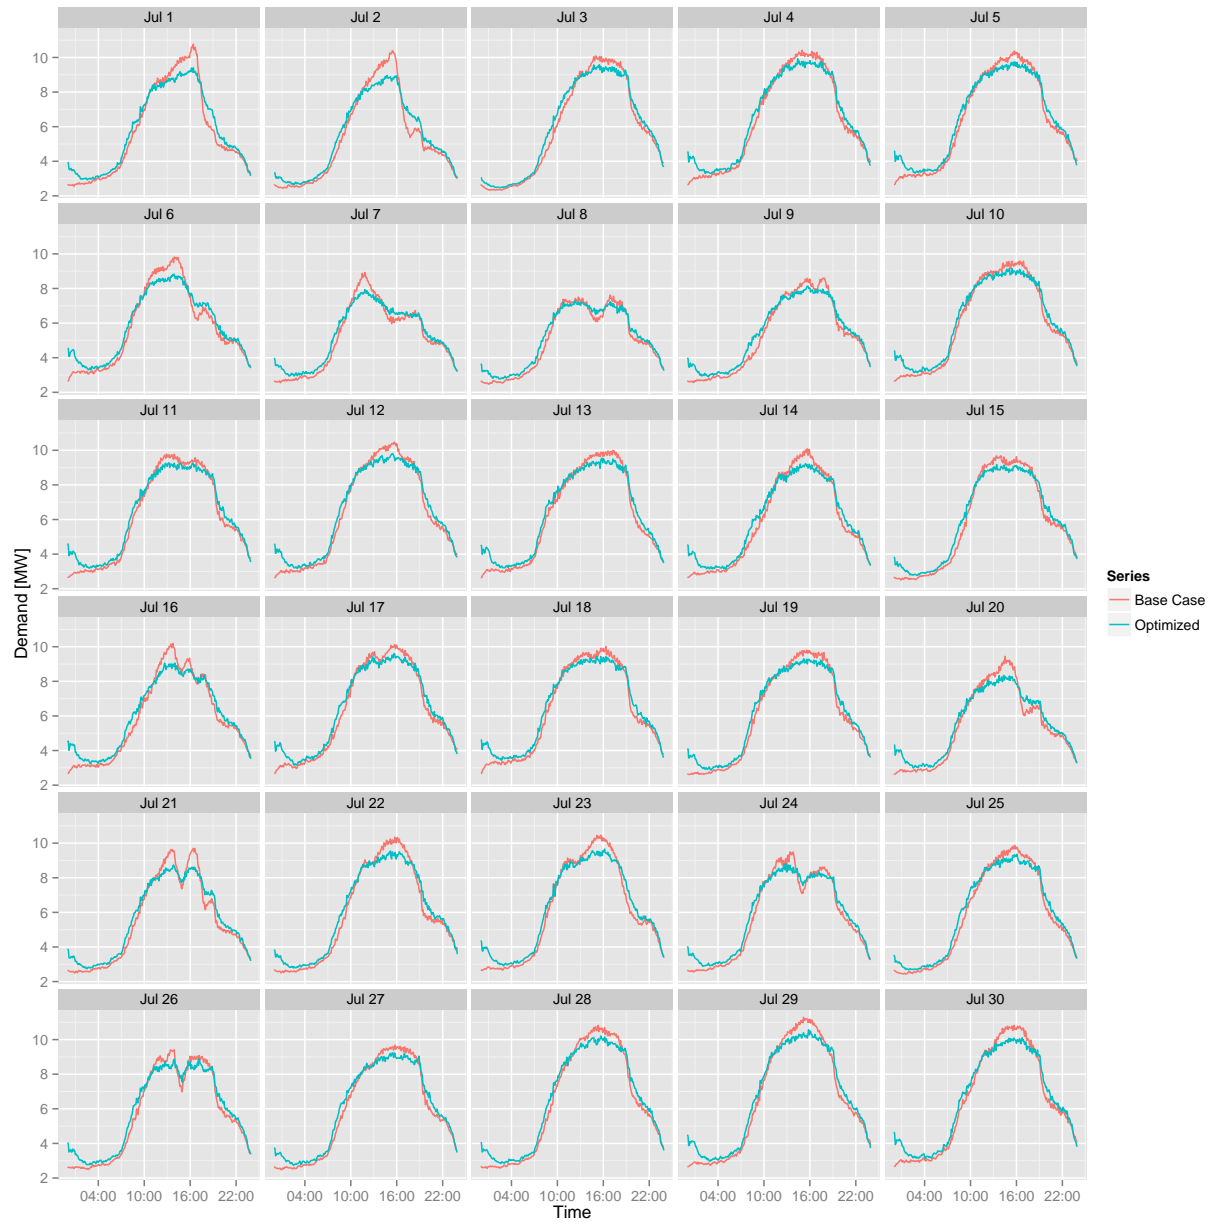


Figure B.4: Feeder demand profiles for Houston demand limiting optimization, 90 minute moving average, 70% participation.

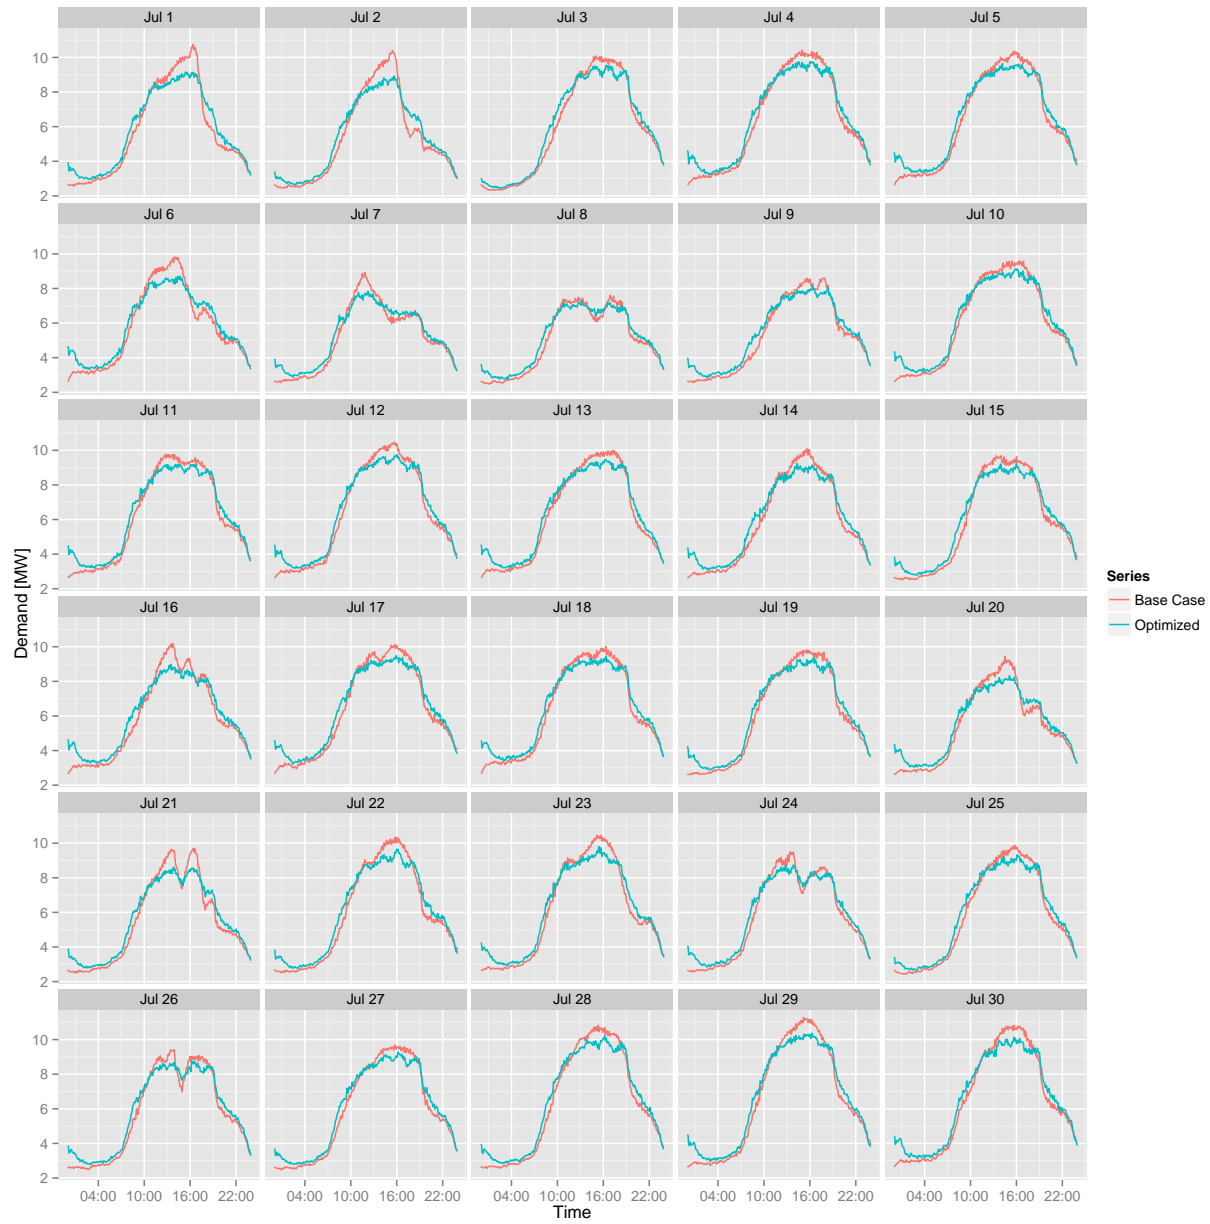


Figure B.5: Feeder demand profiles for Houston demand limiting optimization, 120 minute moving average, 70% participation.

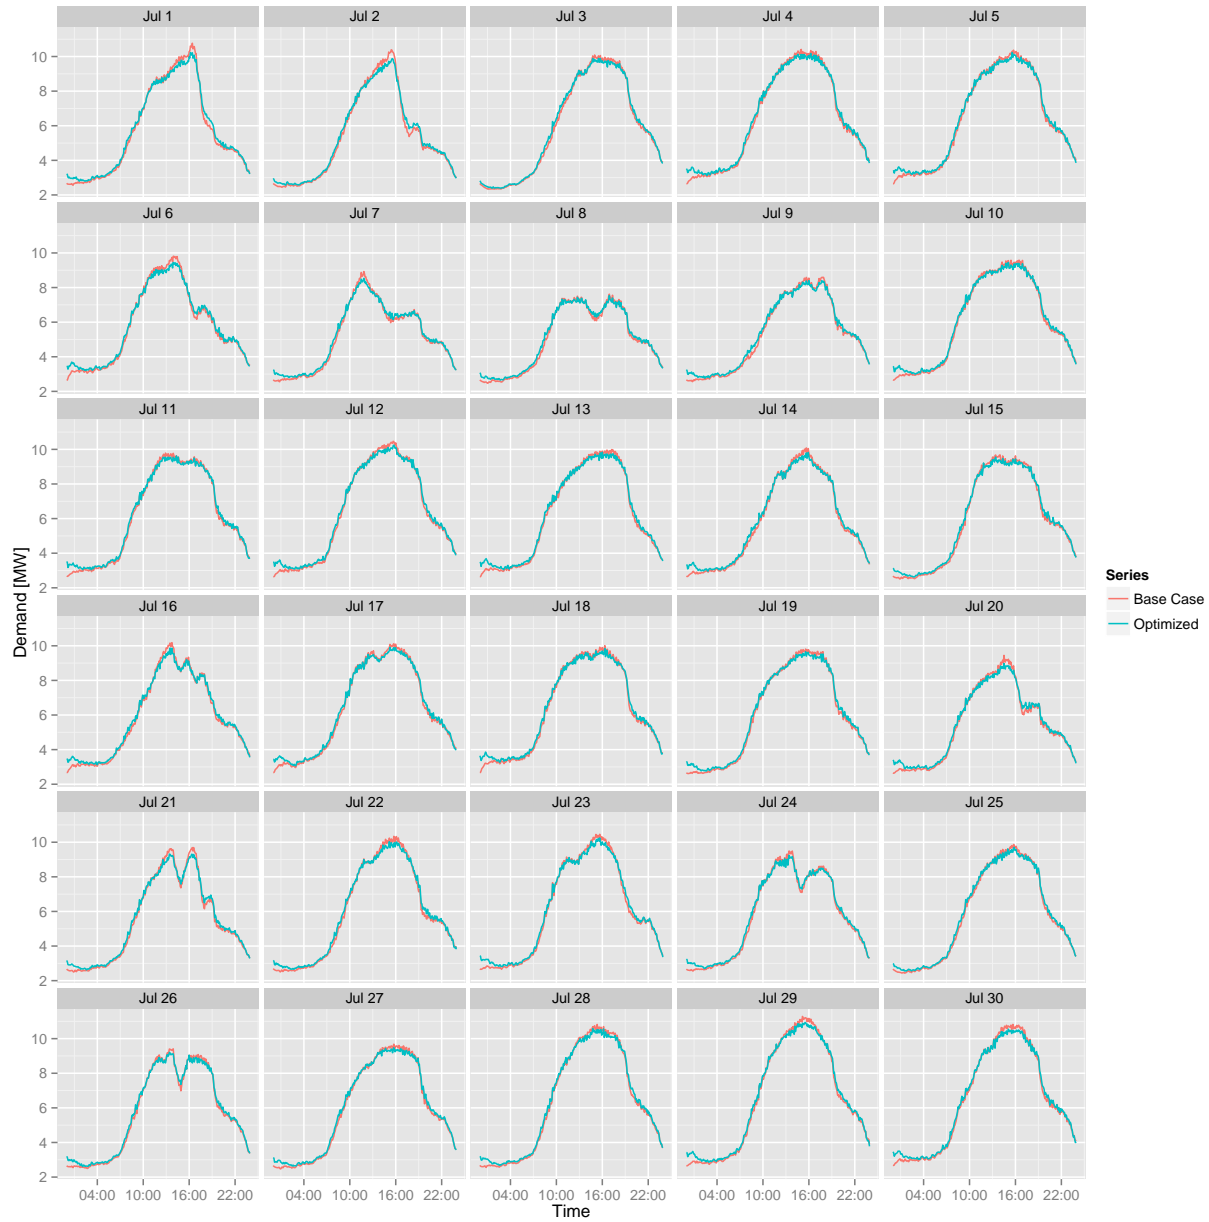


Figure B.6: Feeder demand profiles for Houston demand limiting optimization, 30% participation.




Figure B.7: Feeder demand profiles for Houston, July 29 demand limiting optimization, 70% participation.




Figure B.8: Experiment to measure precooling investment and flexible demand in Houston, July 29 feeder, 70% participation.

B.1.1 Demand Limiting as a Function of Mass Level

Section 3.1.11 showed the reduced order model to be somewhat insensitive to the range of internal mass levels observed in the building population. One question that arises is whether or not this affects the optimization results. In theory, a very massive building should be able to shift a larger amount of the cooling load, thus reducing peak demand. Although the range of mass levels simulated in this work is relatively small, an experiment can be conducted to show how mass level affects demand limiting at an aggregate level.

The analysis involves performing a demand limiting optimization on a single day using the Houston feeder model. Participants are ranked and grouped into quartiles by mass level. Mass level is expressed as mass to floor area ratio and values range between 2.5 and 4.0. 2100 homes in the Houston feeder have air conditioning, yielding 525 homes in each quartile. The remaining homes are not included in the analysis. The optimization is performed four times, i.e. once per quartile, assuming that the entire quartile group participates.

Figure B.9 shows the resulting aggregate demand profiles. The Base Case series represents the aggregate demand of the 2100 homes before optimization. The Fourth Quartile series represents optimized aggregate demand assuming only the highest mass level quartile participates in the optimization. The First Quartile series assumes only the lowest mass level quartile participates.

Results of the experiment support the previous observations regarding sensitivity to internal mass level. Whether or not the range of mass levels is truly representative of the building stock is of some considerable debate. Further exploration of how this and other building characteristics affect aggregate feeder demand is left for future consideration.




Figure B.9: Feeder demand profiles by mass level for Houston, July 29 demand limiting optimization.

B.2 Los Angeles

Table B.2: Performance metrics for Los Angeles feeder demand limiting optimization, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	0.26	0.18	0.35
Peak Demand [MW]	-0.03	-0.08	-0.01
Peak to Valley [%]	98.91	97.15	99.70
Load Factor [%]	0.89	0.43	1.84
Ramp [MW]	-0.02	-0.16	0.11




Figure B.10: Total spectral power as a function of frequency bin for Los Angeles feeder demand limiting optimization, 30% participation.




Figure B.11: Feeder demand profiles for Los Angeles demand limiting optimization, 30% participation.




Figure B.12: Feeder demand profiles for Los Angeles, July 24 demand limiting optimization, 70% participation.

B.3 New York

Table B.3: Performance metrics for New York feeder demand limiting optimization, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	0.80	0.64	1.22
Peak Demand [MW]	-0.11	-0.19	-0.06
Peak to Valley [%]	92.33	80.91	97.96
Load Factor [%]	1.82	1.11	3.00
Ramp [MW]	-0.21	-0.70	0.17




Figure B.13: Total spectral power as a function of frequency bin for New York feeder demand limiting optimization, 30% participation.

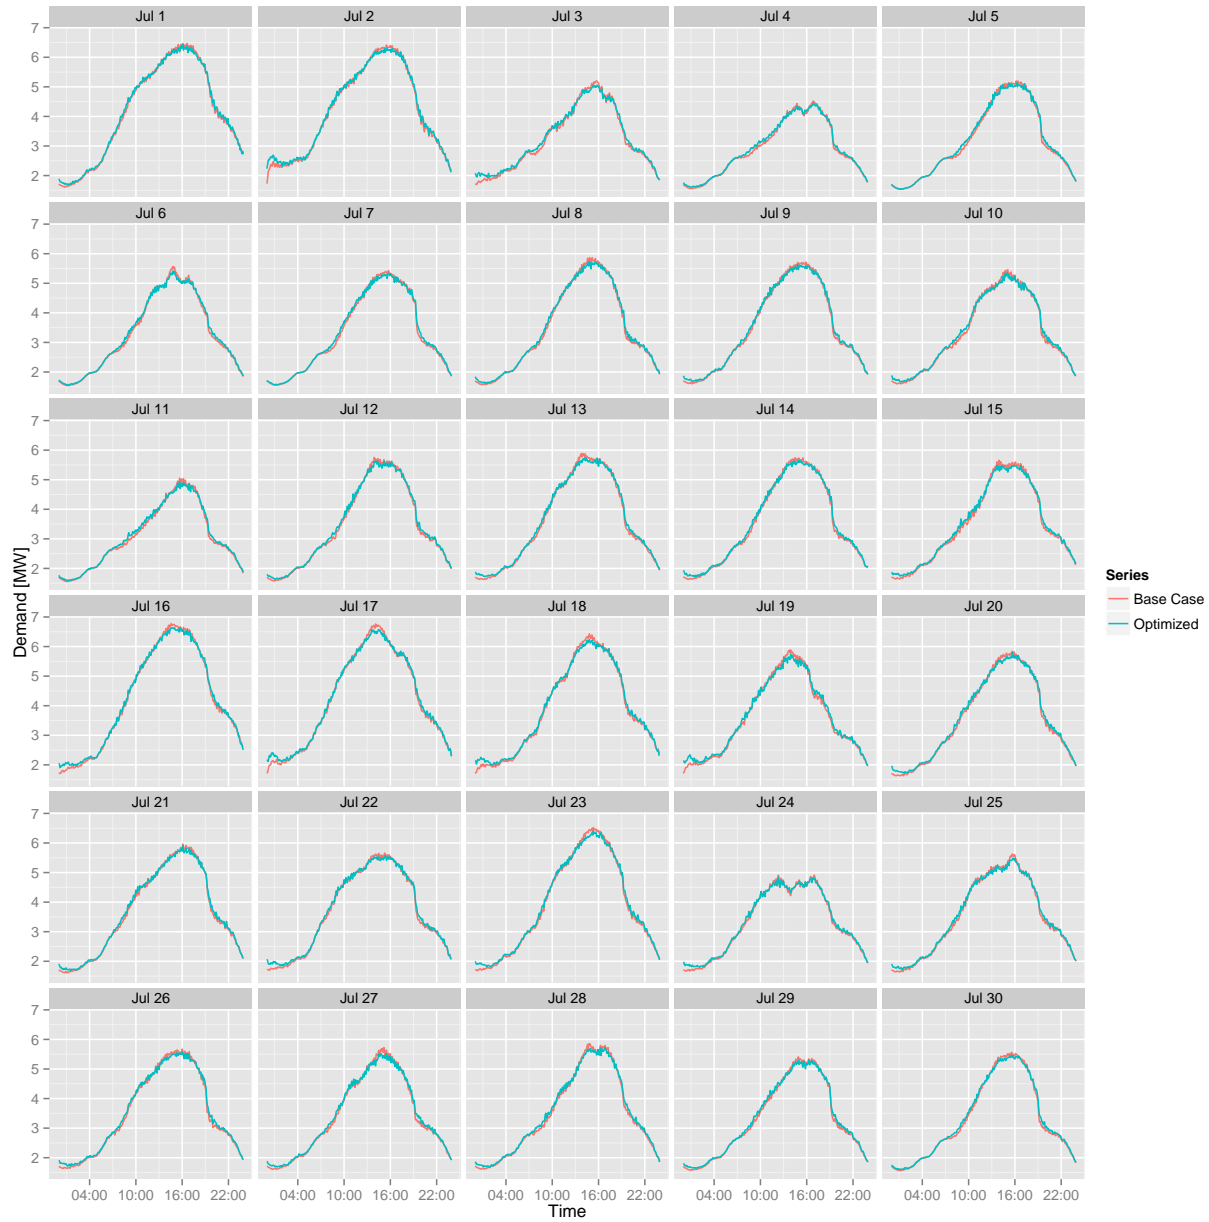


Figure B.14: Feeder demand profiles for New York demand limiting optimization, 30% participation.




Figure B.15: Feeder demand profiles for New York, July 23 demand limiting optimization, 70% participation.




Figure B.16: Experiment to measure precooling investment and flexible demand in New York, July 16 feeder, 70% participation.

Appendix C

Supplement to Dynamic Price Optimization

C.1 Houston – Day-Ahead Price

Table C.1: Performance metrics for Houston feeder day-ahead price optimization, ramp-return case, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	-0.85	-1.14	-0.30
Peak Demand [MW]	-0.24	-1.08	0.36
Peak to Valley [%]	90.94	79.77	99.74
Load Factor [%]	1.11	-2.67	6.65
Ramp [MW]	-0.19	-3.31	1.82




Figure C.1: Total spectral power as a function of frequency bin for Houston feeder day-ahead price optimization, ramp-return upper boundary case, 30% participation.

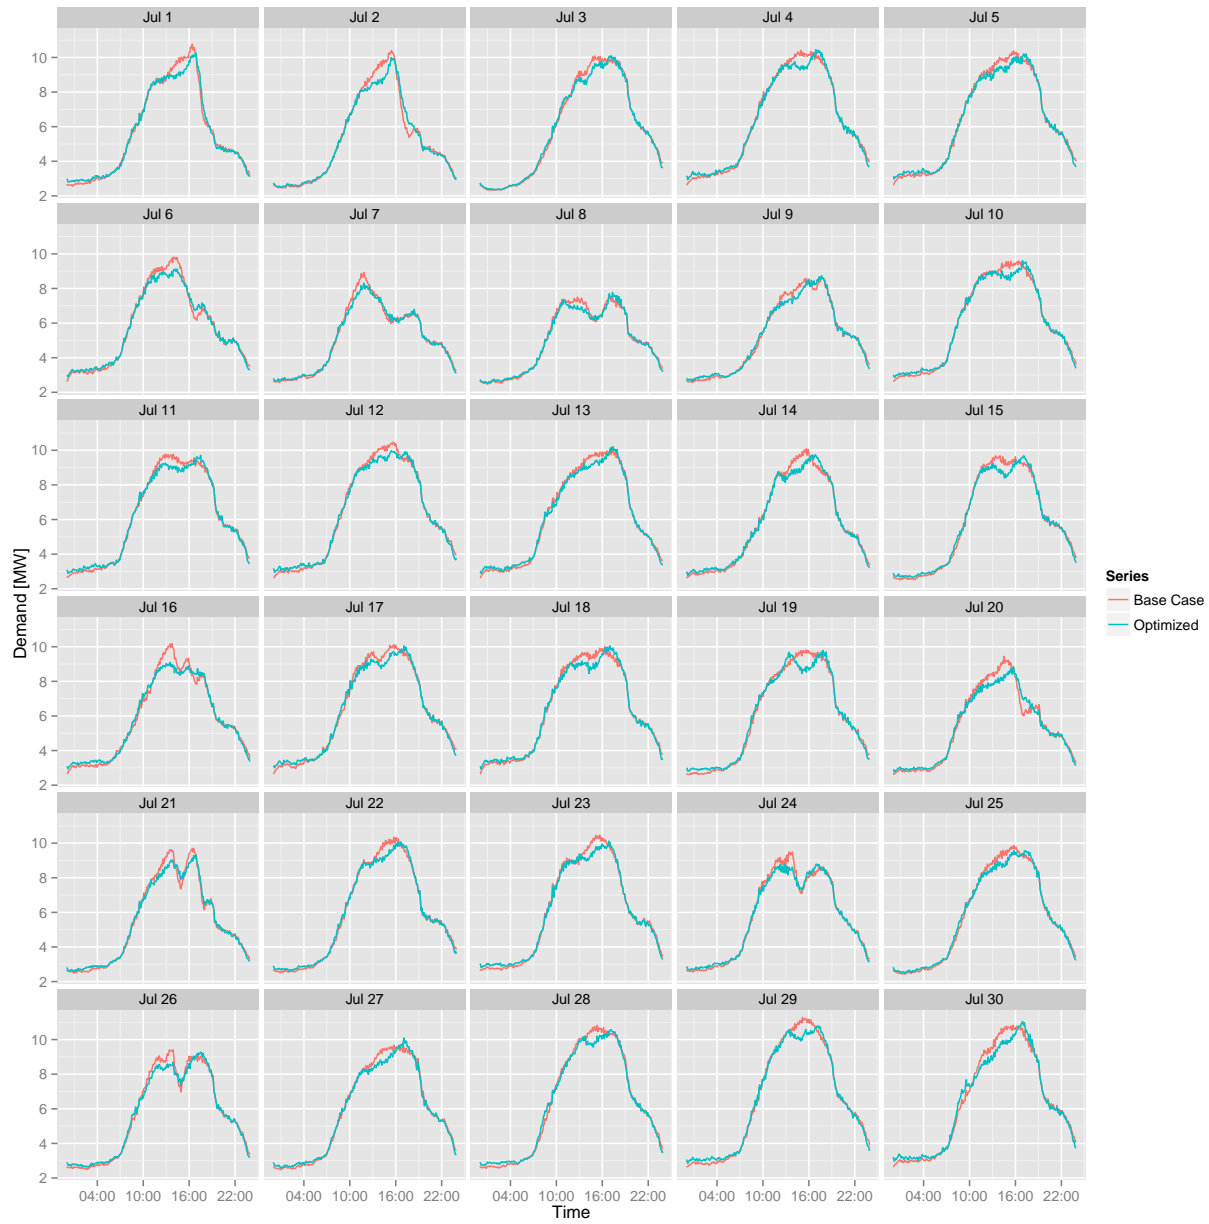


Figure C.2: Feeder demand profiles for Houston day-ahead price optimization, ramp-return upper boundary case, 30% participation.




Figure C.3: Feeder demand profiles for Houston, July 29 day-ahead price optimization, zero-degree upper boundary case, 70% participation.

C.2 Los Angeles – Day-Ahead Price

Table C.2: Performance metrics for Los Angeles feeder day-ahead price optimization, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	-0.64	-0.87	-0.50
Peak Demand [MW]	-0.02	-0.17	0.09
Peak to Valley [%]	99.54	95.04	102.65
Load Factor [%]	-0.45	-2.81	2.86
Ramp [MW]	0.18	-0.18	0.46

Table C.3: Performance metrics for Los Angeles feeder day-ahead price optimization, zero-degree upper boundary case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	-0.08	-0.18	-0.03
Peak Demand [MW]	0.02	-0.06	0.13
Peak to Valley [%]	100.70	98.29	104.50
Load Factor [%]	-0.58	-3.06	0.99
Ramp [MW]	0.15	-0.16	0.46

Table C.4: Performance metrics for Los Angeles feeder day-ahead price optimization, ramp-return upper boundary case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	-0.48	-0.65	-0.34
Peak Demand [MW]	0.19	0.01	0.32
Peak to Valley [%]	105.58	100.31	109.44
Load Factor [%]	-4.14	-6.34	-0.66
Ramp [MW]	0.29	0.04	0.56

Table C.5: Performance metrics for Los Angeles feeder day-ahead price optimization, ramp-return case, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	-0.21	-0.32	-0.15
Peak Demand [MW]	0.07	-0.03	0.13
Peak to Valley [%]	102.02	98.97	104.00
Load Factor [%]	-1.59	-2.88	0.40
Ramp [MW]	0.10	-0.17	0.35




Figure C.4: Total spectral power as a function of frequency bin for Los Angeles feeder day-ahead price optimization, 70% participation.




Figure C.5: Total spectral power as a function of frequency bin for Los Angeles feeder day-ahead price optimization, zero-degree upper boundary case, 70% participation.




Figure C.6: Total spectral power as a function of frequency bin for Los Angeles feeder day-ahead price optimization, ramp-return upper boundary case, 70% participation.




Figure C.7: Total spectral power as a function of frequency bin for Los Angeles feeder day-ahead price optimization, ramp-return upper boundary case, 30% participation.

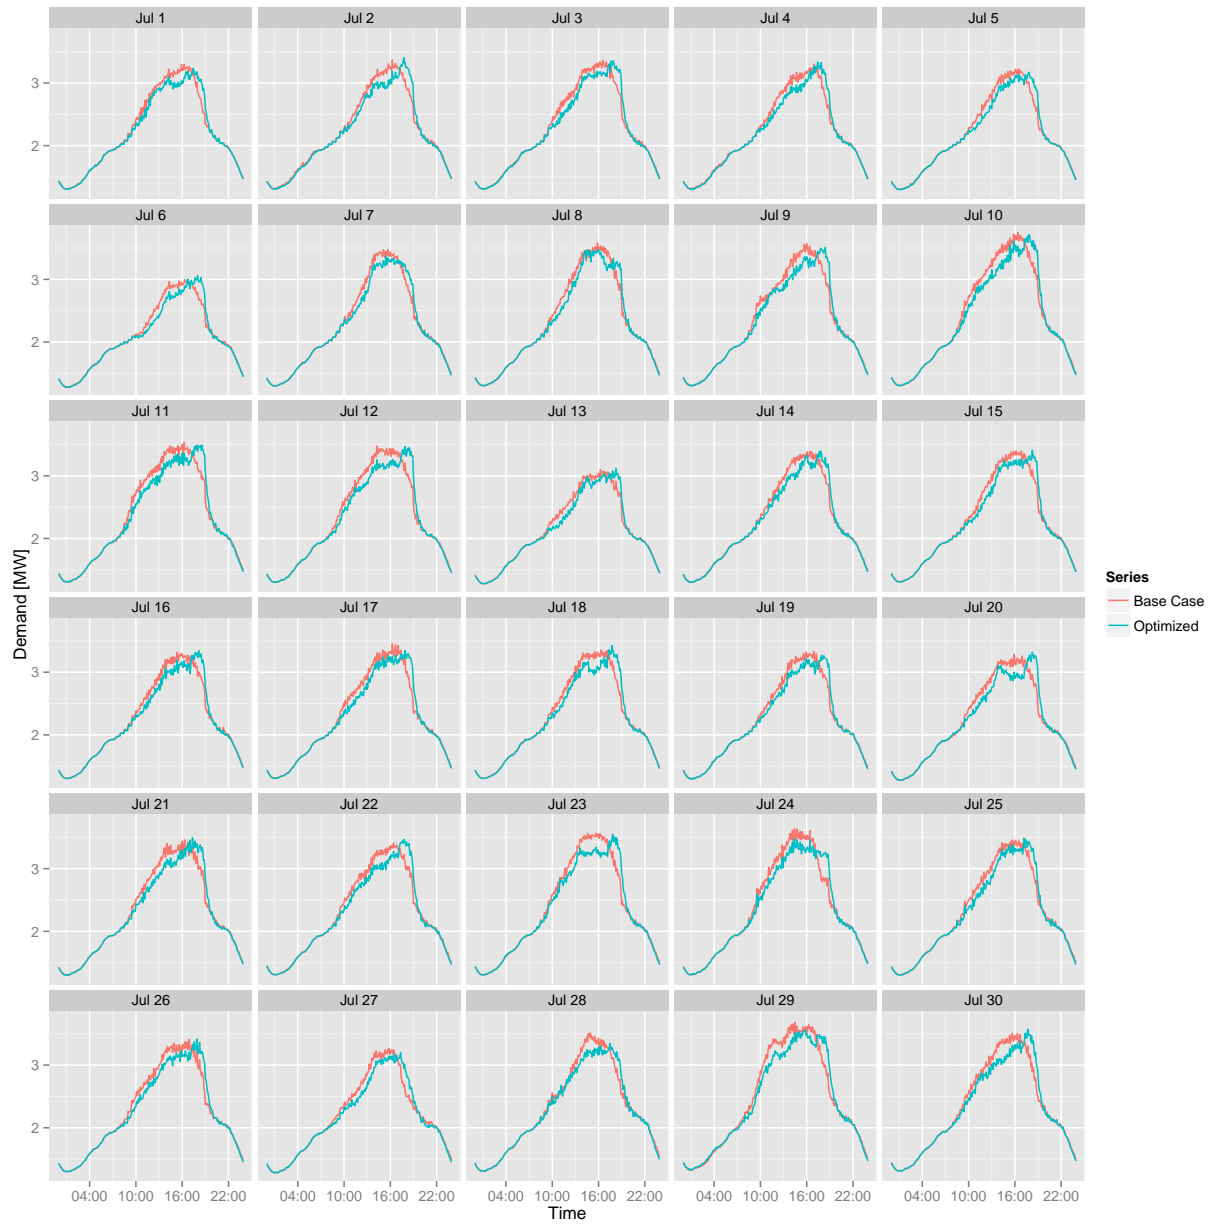


Figure C.8: Feeder demand profiles for Los Angeles day-ahead price optimization, 70% participation.




Figure C.9: Feeder demand profiles for Los Angeles day-ahead price optimization, ramp-return upper boundary case, 70% participation.

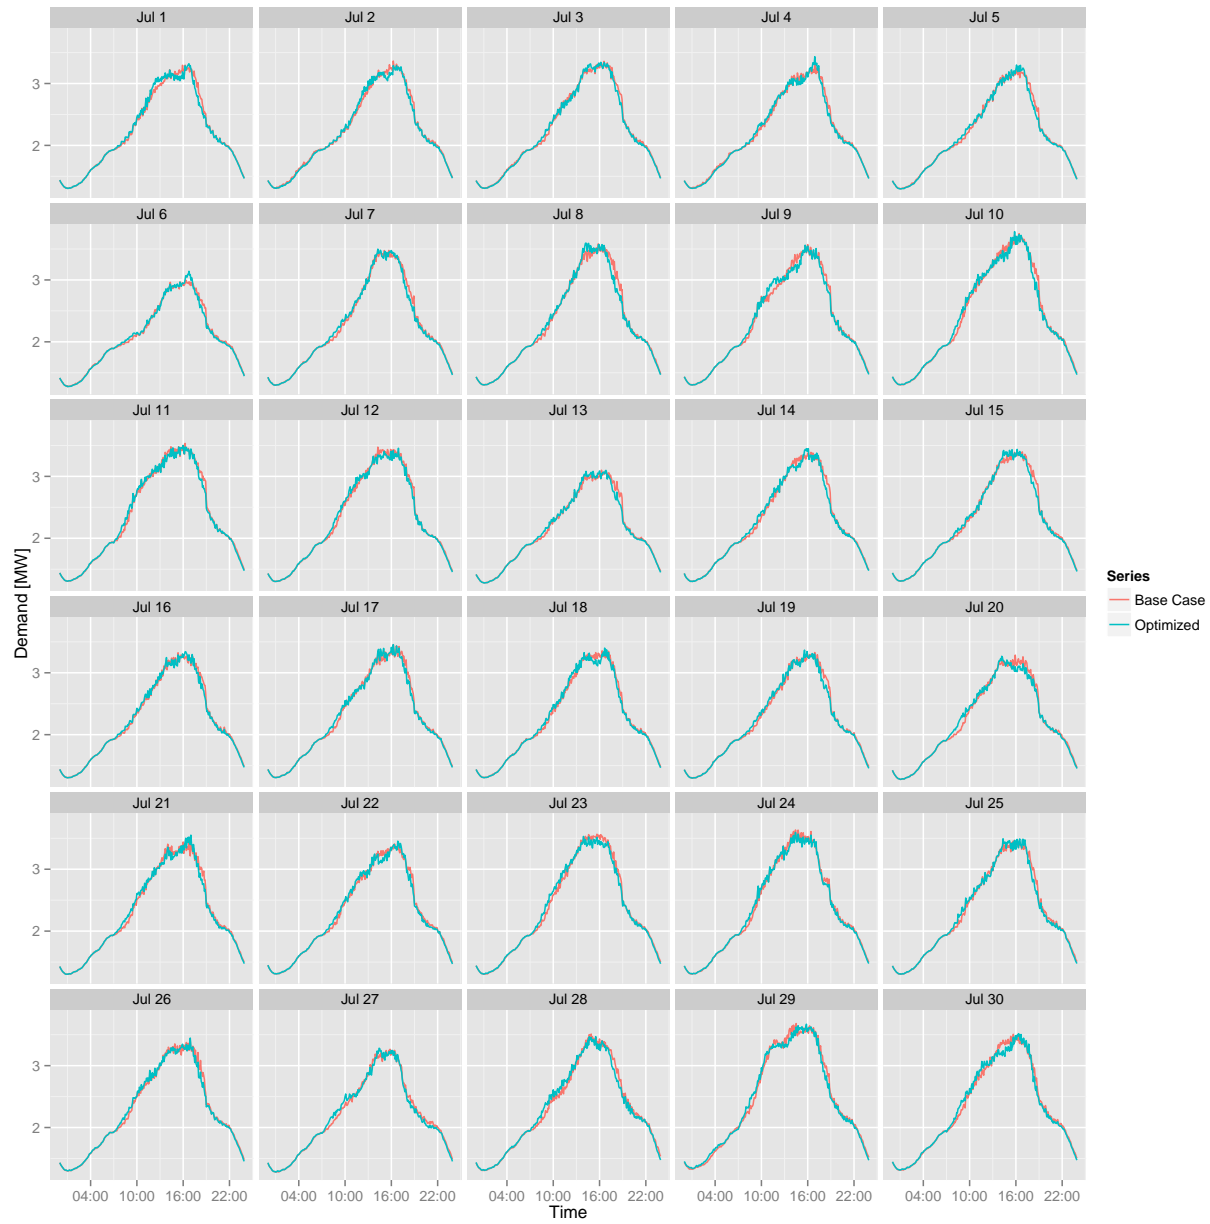


Figure C.10: Feeder demand profiles for Los Angeles day-ahead price optimization, zero-degree upper boundary case, 70% participation.

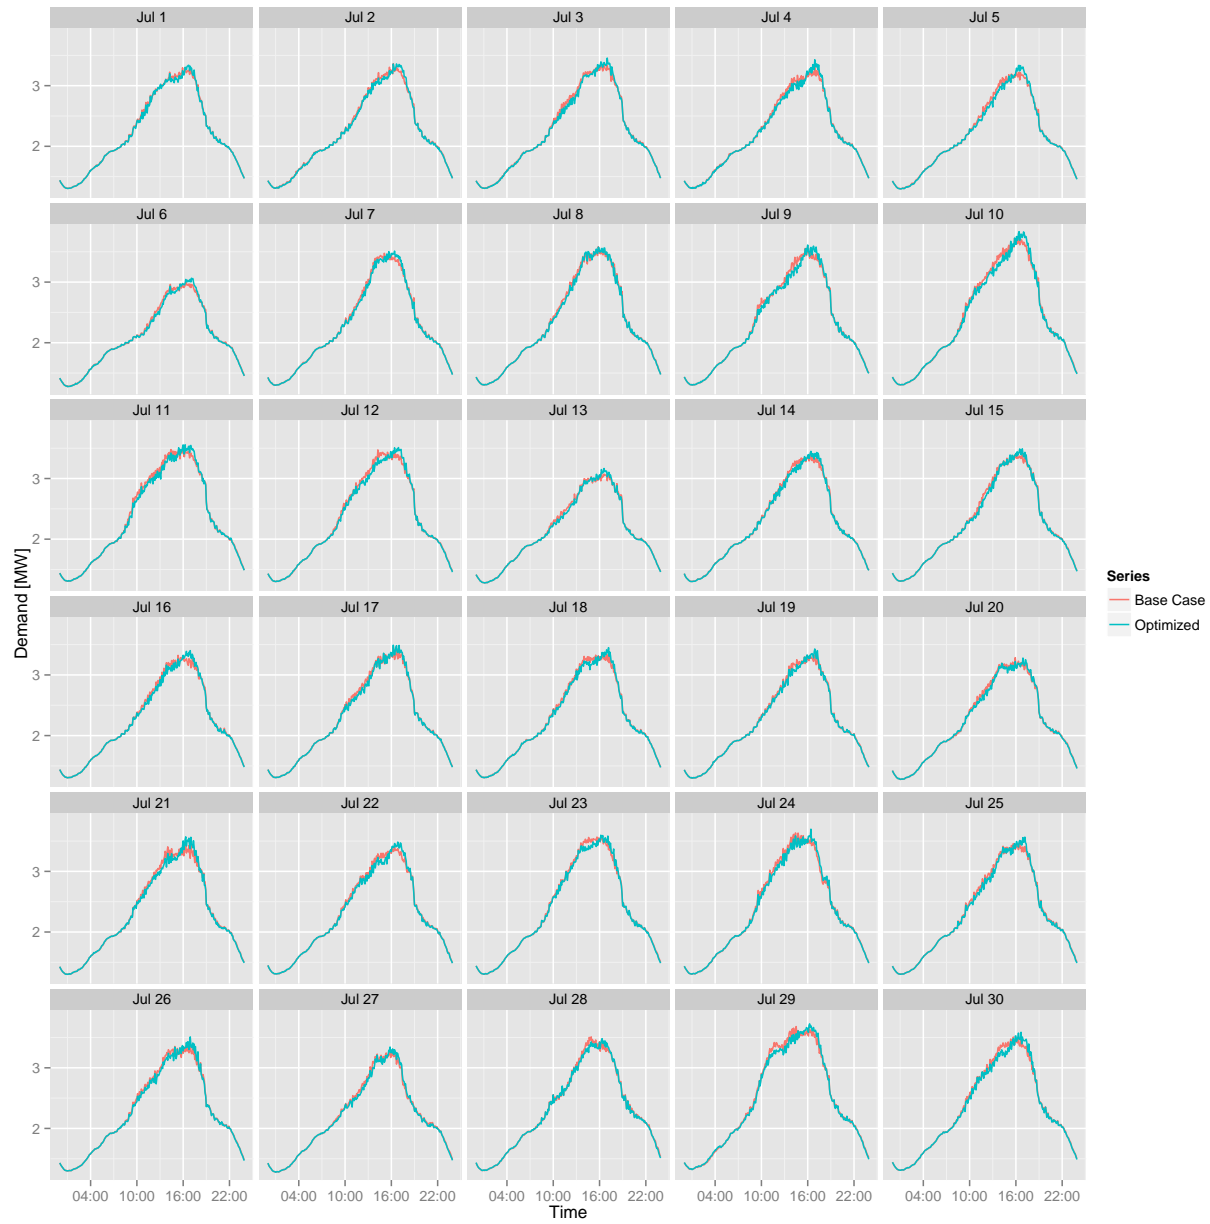


Figure C.11: Feeder demand profiles for Los Angeles day-ahead price optimization, ramp-return upper boundary case, 30% participation.




Figure C.12: Feeder demand profiles for Los Angeles, July 7 day-ahead price optimization, 70% participation.




Figure C.13: Feeder demand profiles for Los Angeles, July 24 day-ahead price optimization, 70% participation.




Figure C.14: Feeder demand profiles for Los Angeles, July 28 day-ahead price optimization, 70% participation.

C.3 Los Angeles – Synthetic Price

Table C.6: Performance metrics for Los Angeles feeder synthetic price optimization, zero-degree upper boundary case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	-0.08	-0.19	-0.04
Peak Demand [MW]	-0.03	-0.11	0.00
Peak to Valley [%]	98.90	96.49	100.18
Load Factor [%]	0.46	-0.14	1.82
Ramp [MW]	-0.02	-0.29	0.21

Table C.7: Performance metrics for Los Angeles feeder synthetic price optimization, ramp-return upper boundary case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	-0.49	-0.66	-0.37
Peak Demand [MW]	0.13	0.06	0.22
Peak to Valley [%]	103.59	100.83	106.54
Load Factor [%]	-3.02	-5.23	-1.61
Ramp [MW]	0.18	-0.04	0.52

Table C.8: Performance metrics for Los Angeles feeder synthetic price optimization, zero-degree case, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	-0.03	-0.08	-0.01
Peak Demand [MW]	-0.02	-0.06	0.03
Peak to Valley [%]	99.44	98.24	100.64
Load Factor [%]	0.28	-0.60	1.01
Ramp [MW]	0.00	-0.16	0.17




Figure C.15: Total spectral power as a function of frequency bin for Los Angeles feeder synthetic price optimization, zero-degree upper boundary case, 70% participation.




Figure C.16: Total spectral power as a function of frequency bin for Los Angeles feeder synthetic price optimization, ramp-return upper boundary case, 70% participation.




Figure C.17: Total spectral power as a function of frequency bin for Los Angeles feeder synthetic price optimization, zero-degree upper boundary case, 30% participation.

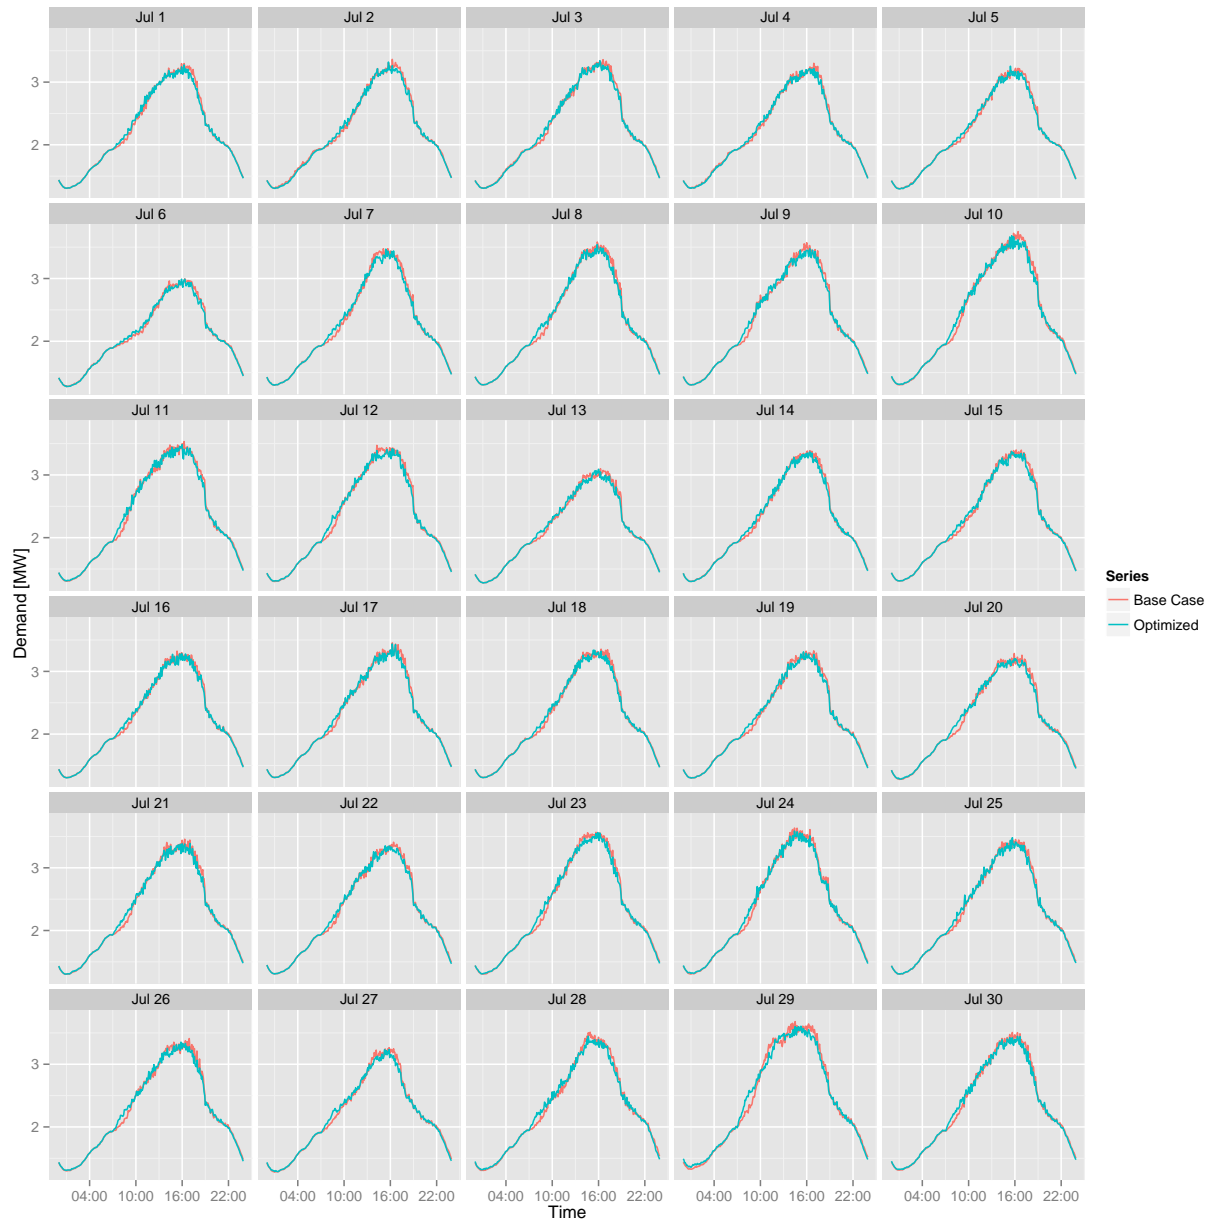


Figure C.18: Feeder demand profiles for Los Angeles synthetic price optimization, zero-degree upper boundary case, 70% participation.

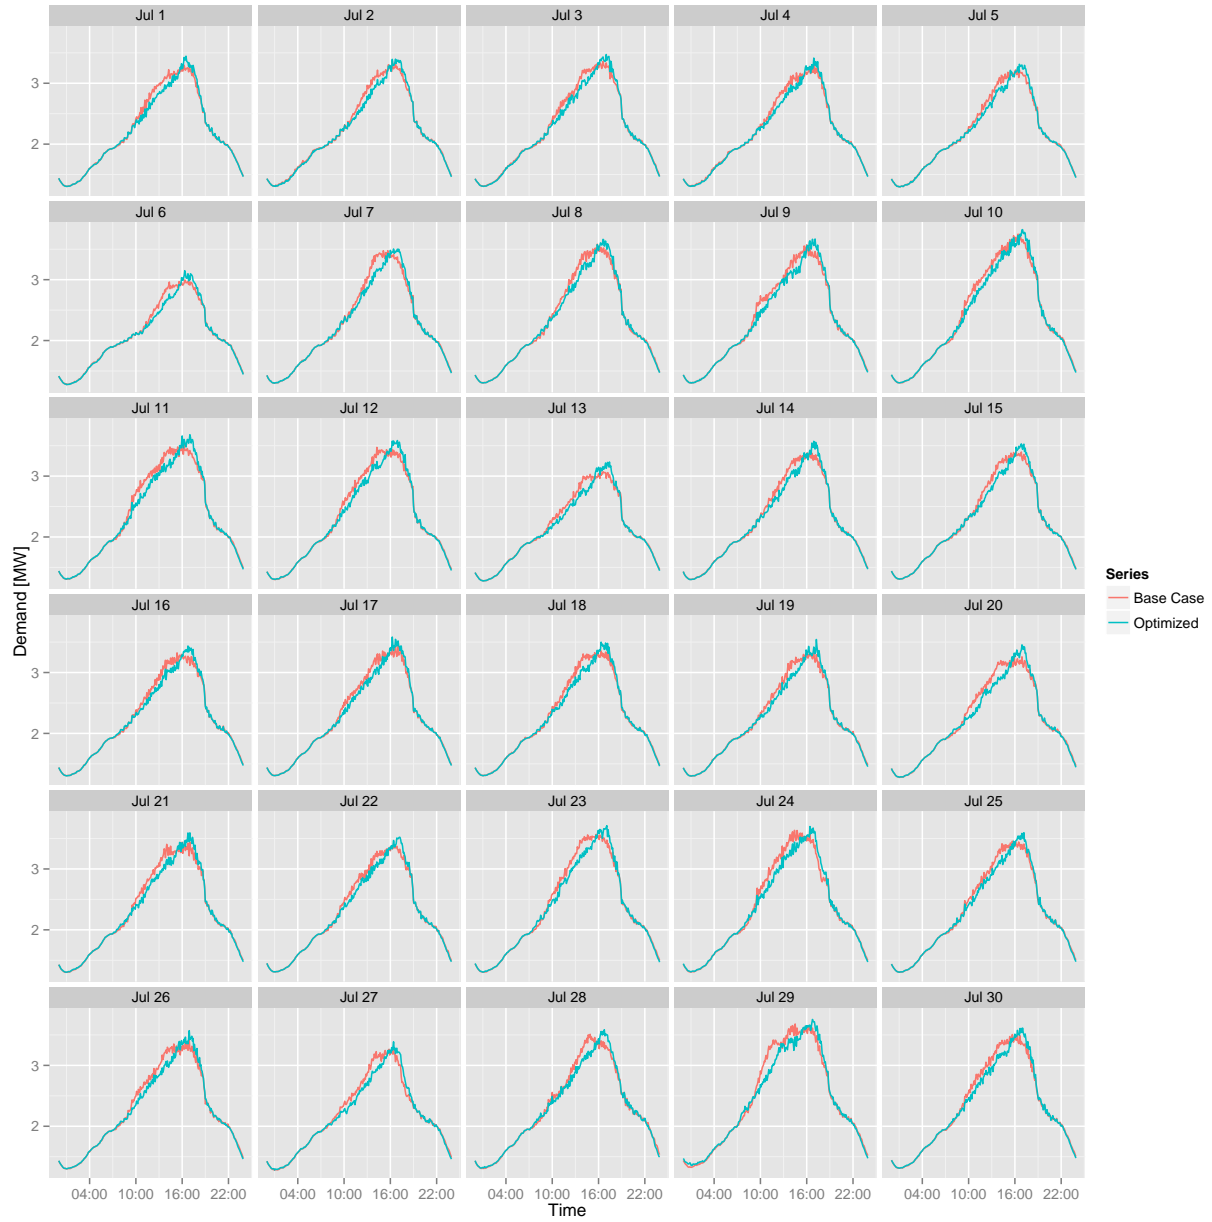


Figure C.19: Feeder demand profiles for Los Angeles synthetic price optimization, ramp-return upper boundary case, 70% participation.

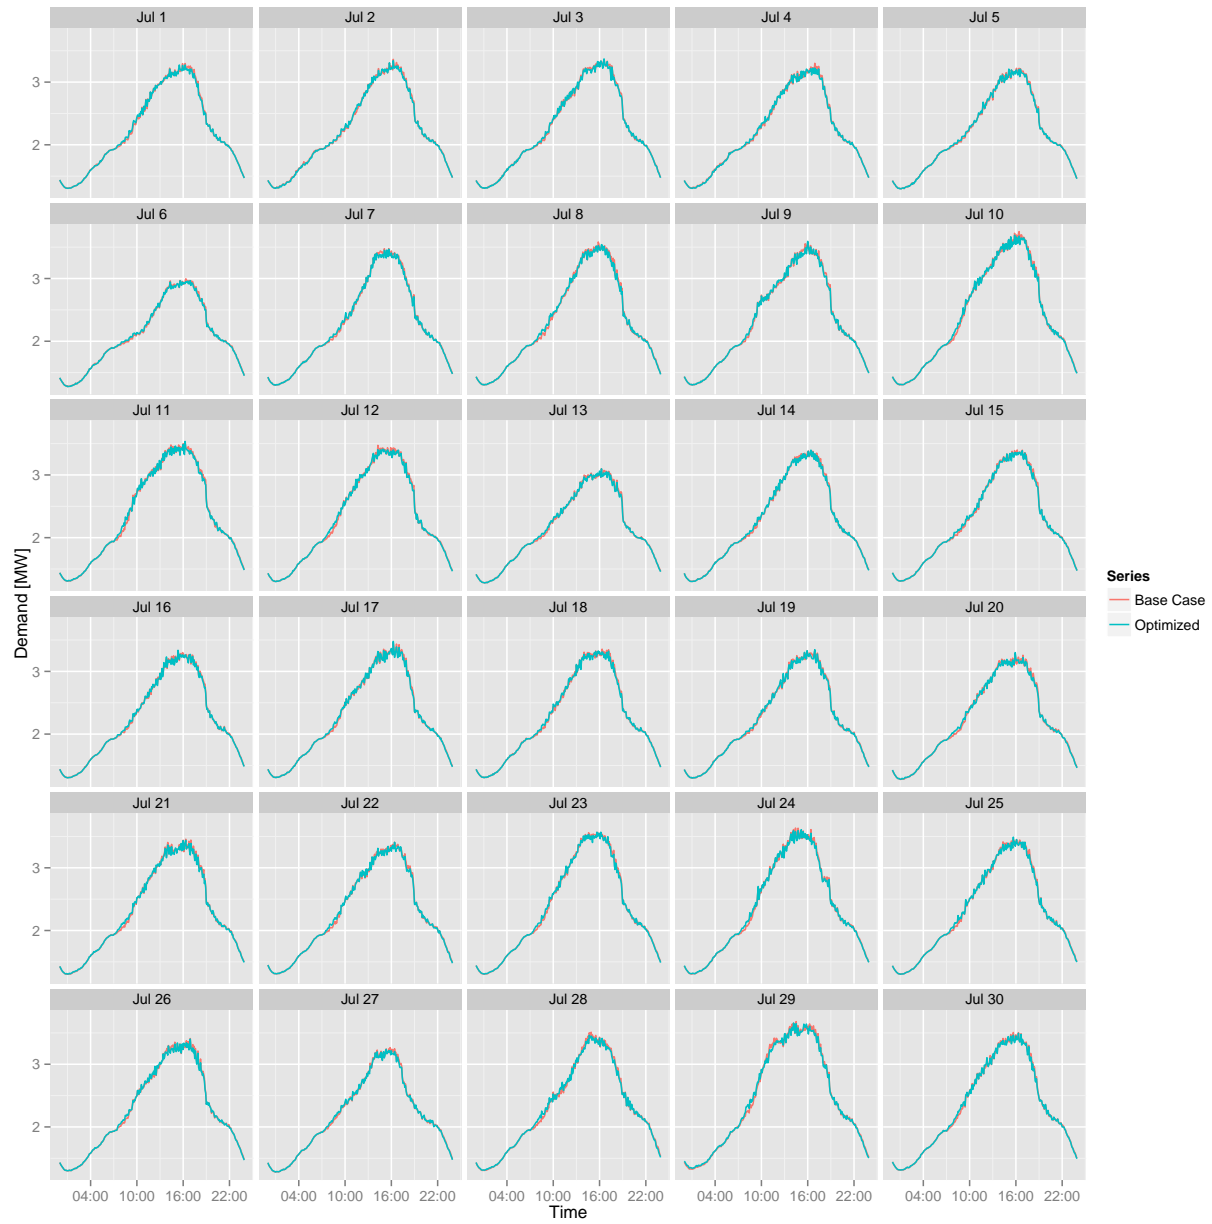


Figure C.20: Feeder demand profiles for Los Angeles synthetic price optimization, zero-degree upper boundary case, 30% participation.

C.4 New York – Day-Ahead Price

Table C.9: Performance metrics for New York feeder day-ahead price optimization, zero-degree upper boundary case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	3.24	1.89	4.24
Peak Demand [MW]	0.14	-0.32	0.43
Peak to Valley [%]	91.81	81.30	98.67
Load Factor [%]	0.82	-2.52	4.88
Ramp [MW]	0.72	-0.45	2.82

Table C.10: Performance metrics for New York feeder day-ahead price optimization, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	1.54	0.20	2.90
Peak Demand [MW]	0.15	-0.46	0.52
Peak to Valley [%]	92.62	79.44	102.04
Load Factor [%]	-0.53	-5.44	6.31
Ramp [MW]	1.32	-0.09	4.50

Table C.11: Performance metrics for New York feeder day-ahead price optimization, ramp-return upper boundary case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	2.11	1.14	3.29
Peak Demand [MW]	0.41	-0.09	0.79
Peak to Valley [%]	96.20	82.51	103.39
Load Factor [%]	-2.81	-6.81	2.17
Ramp [MW]	1.30	-0.01	4.09

Table C.12: Performance metrics for New York feeder day-ahead price optimization, ramp-return case, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	-0.63	-0.93	-0.34
Peak Demand [MW]	0.02	-0.31	0.31
Peak to Valley [%]	98.37	84.92	103.87
Load Factor [%]	-0.79	-3.77	2.81
Ramp [MW]	0.28	-0.28	1.02




Figure C.21: Total spectral power as a function of frequency bin for New York feeder day-ahead price optimization, zero-degree upper boundary case, 70% participation.




Figure C.22: Total spectral power as a function of frequency bin for New York feeder day-ahead price optimization, 70% participation.




Figure C.23: Total spectral power as a function of frequency bin for New York feeder day-ahead price optimization, ramp-return upper boundary case, 70% participation.




Figure C.24: Total spectral power as a function of frequency bin for New York feeder day-ahead price optimization, ramp-return upper boundary case, 30% participation.

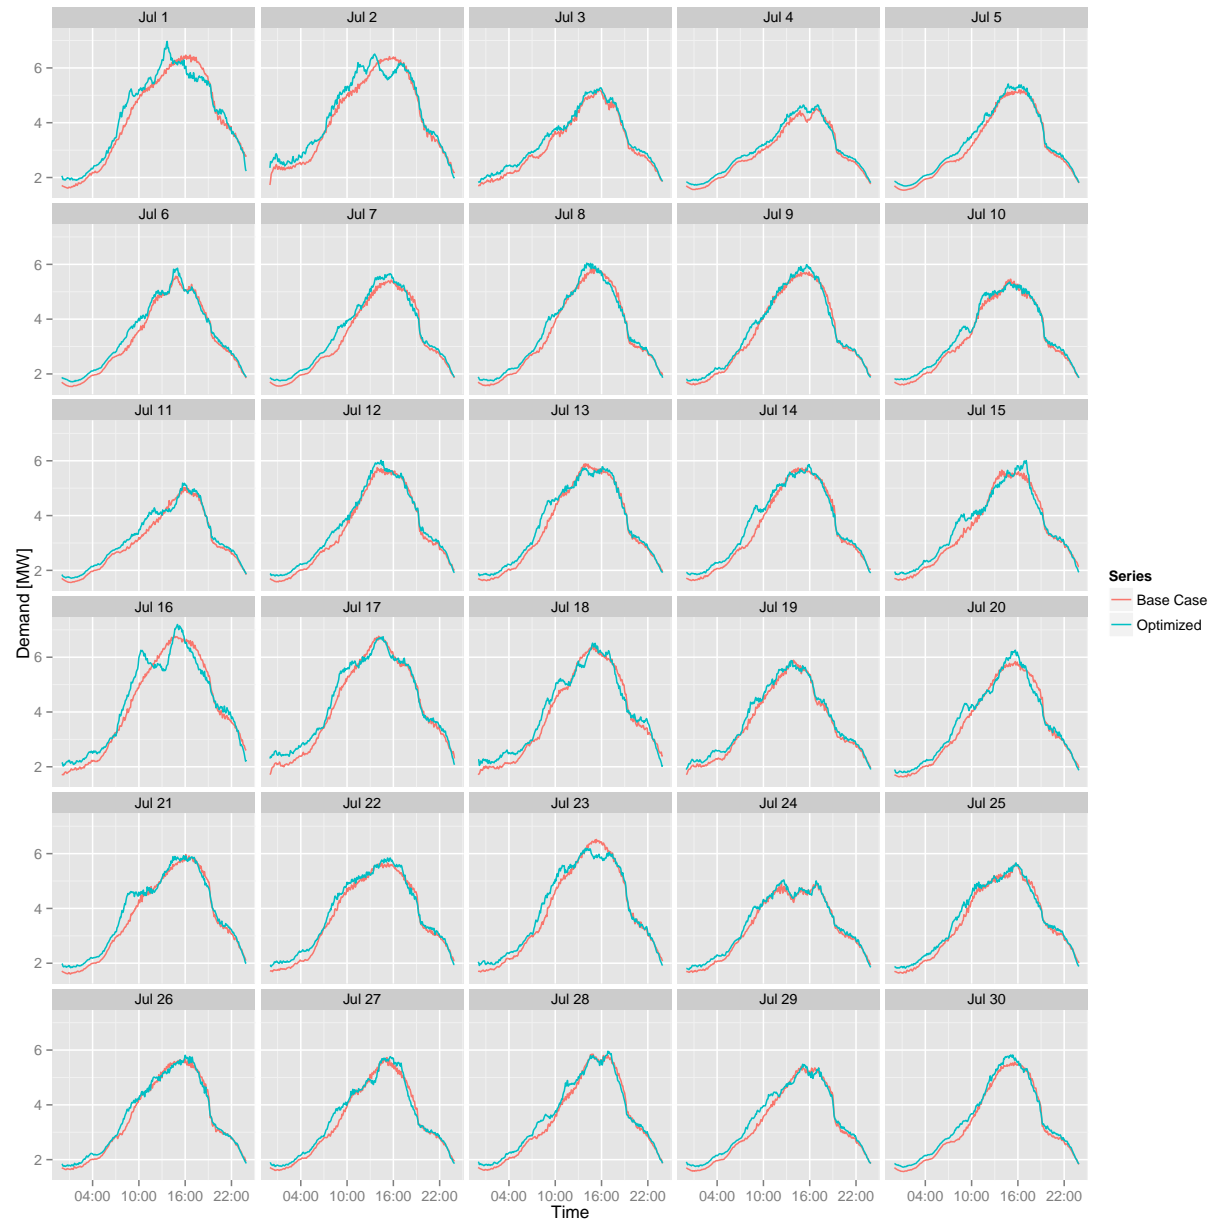


Figure C.25: Feeder demand profiles for New York day-ahead price optimization, zero-degree upper boundary case, 70% participation.

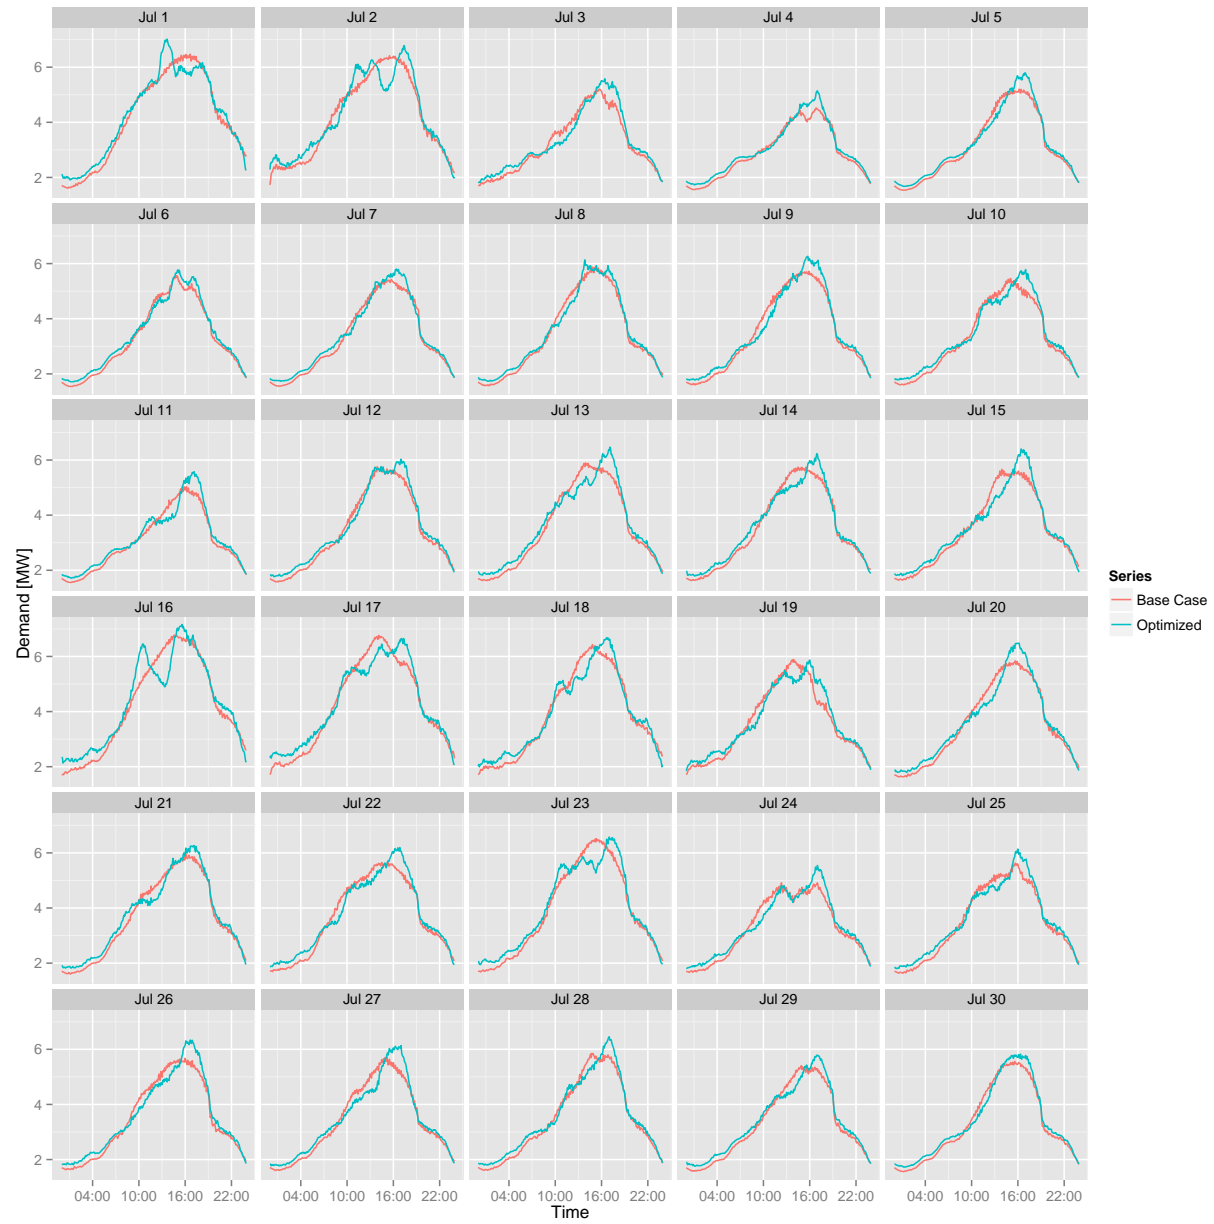


Figure C.26: Feeder demand profiles for New York day-ahead price optimization, ramp-return upper boundary case, 70% participation.

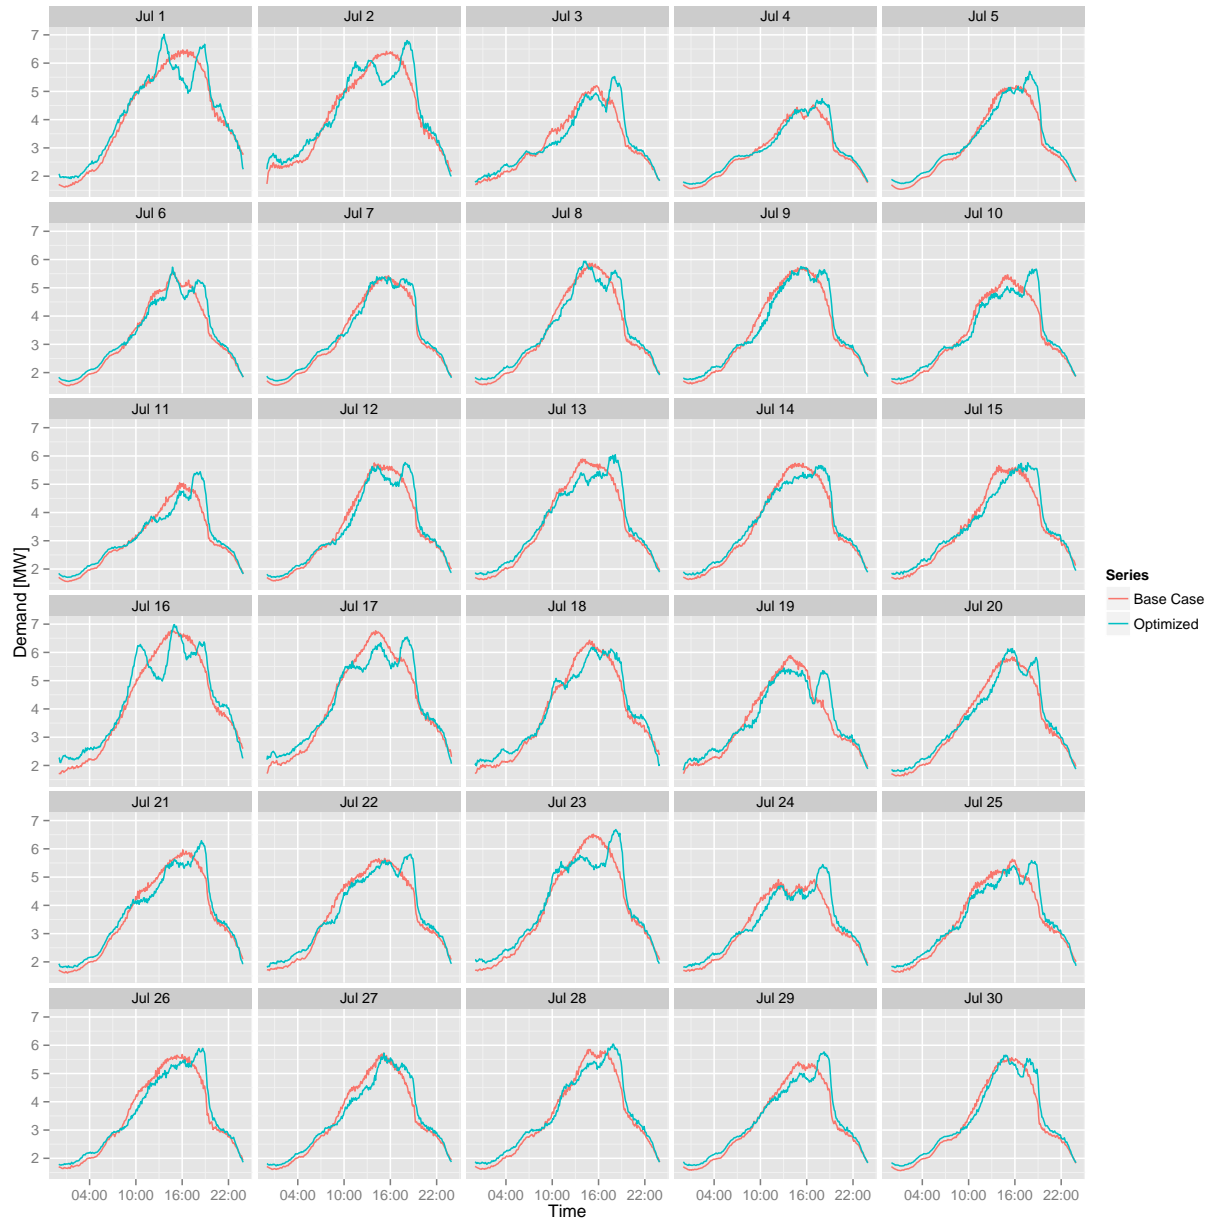


Figure C.27: Feeder demand profiles for New York day-ahead price optimization, 70% participation.

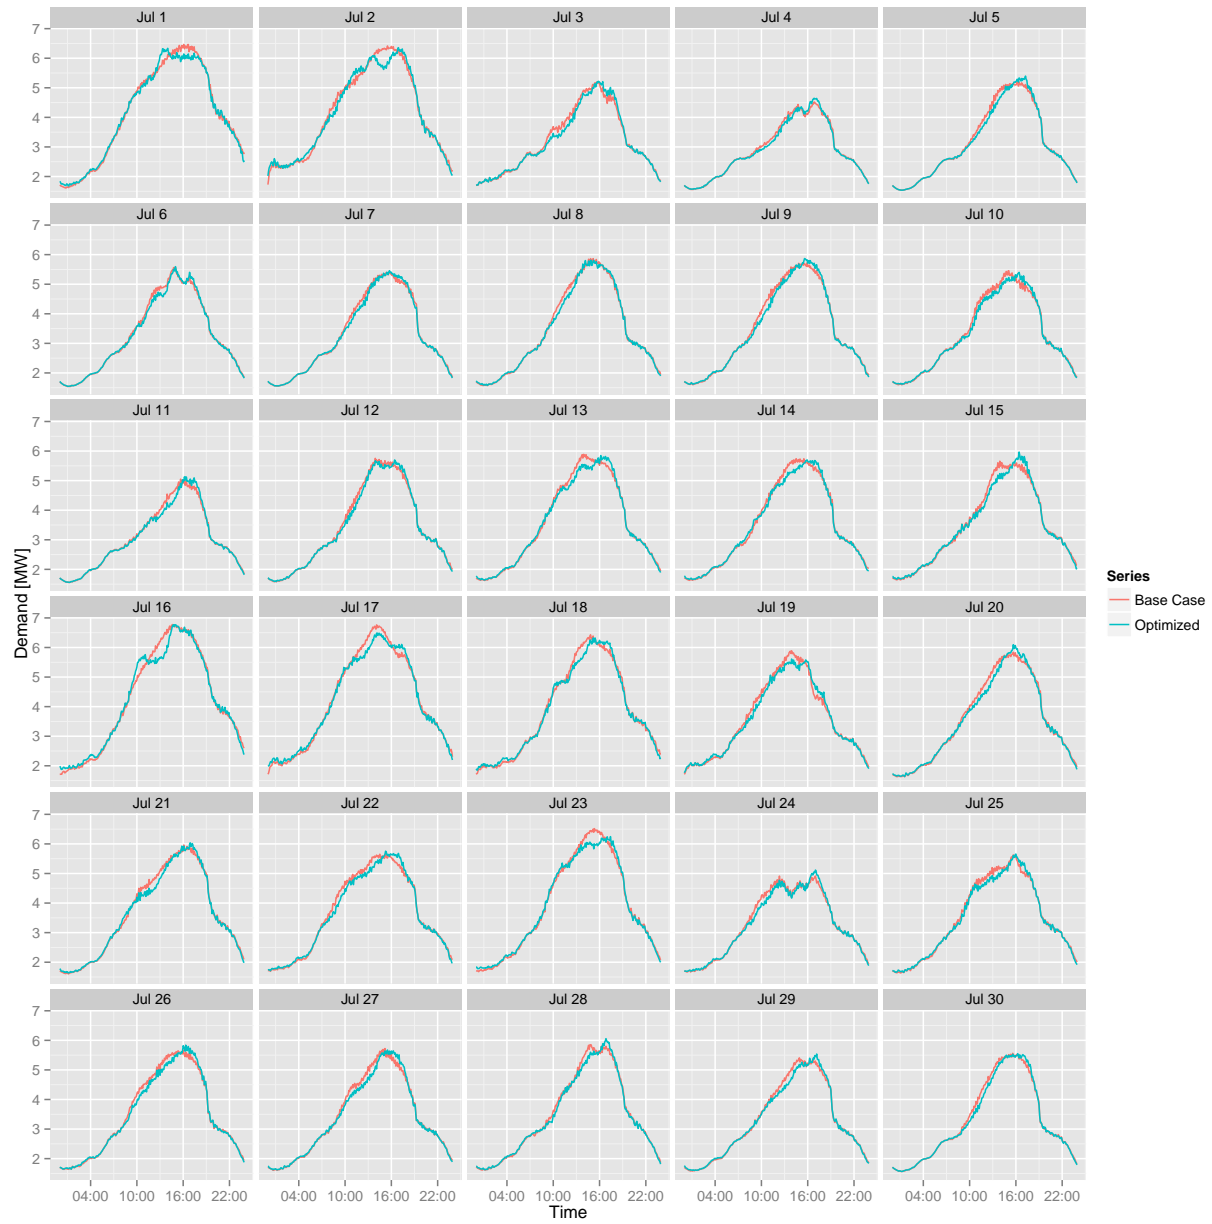


Figure C.28: Feeder demand profiles for New York day-ahead price optimization, ramp-return upper boundary case, 30% participation.



Figure C.29: Feeder demand profiles for New York, July 14 day-ahead price optimization, 70% participation.




Figure C.30: Feeder demand profiles for New York, July 17 day-ahead price optimization, 70% participation.




Figure C.31: Feeder demand profiles for New York, July 19 day-ahead price optimization, 70% participation.

C.5 New York – Synthetic Price

Table C.13: Performance metrics for New York feeder synthetic price optimization, zero-degree upper boundary case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	-0.05	-0.22	0.19
Peak Demand [MW]	-0.16	-0.29	-0.04
Peak to Valley [%]	90.32	80.17	98.19
Load Factor [%]	1.72	0.40	3.90
Ramp [MW]	0.32	-1.00	1.30

Table C.14: Performance metrics for New York feeder synthetic price optimization, ramp-return upper boundary case, 70% participation.

	Mean	Min	Max
Electric Consumption [MWh]	-1.48	-1.97	-0.95
Peak Demand [MW]	0.21	-0.32	0.39
Peak to Valley [%]	97.15	84.77	107.00
Load Factor [%]	-3.24	-5.28	2.36
Ramp [MW]	0.64	-0.77	1.41

Table C.15: Performance metrics for New York feeder synthetic price optimization, zero-degree case, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	-0.02	-0.09	0.08
Peak Demand [MW]	-0.08	-0.16	0.00
Peak to Valley [%]	94.69	88.11	98.69
Load Factor [%]	0.89	-0.06	2.25
Ramp [MW]	-0.06	-0.78	0.41




Figure C.32: Total spectral power as a function of frequency bin for New York feeder synthetic price optimization, zero-degree upper boundary case, 70% participation.




Figure C.33: Total spectral power as a function of frequency bin for New York feeder synthetic price optimization, ramp-return upper boundary case, 70% participation.




Figure C.34: Total spectral power as a function of frequency bin for New York feeder synthetic price optimization, zero-degree upper boundary case, 30% participation.

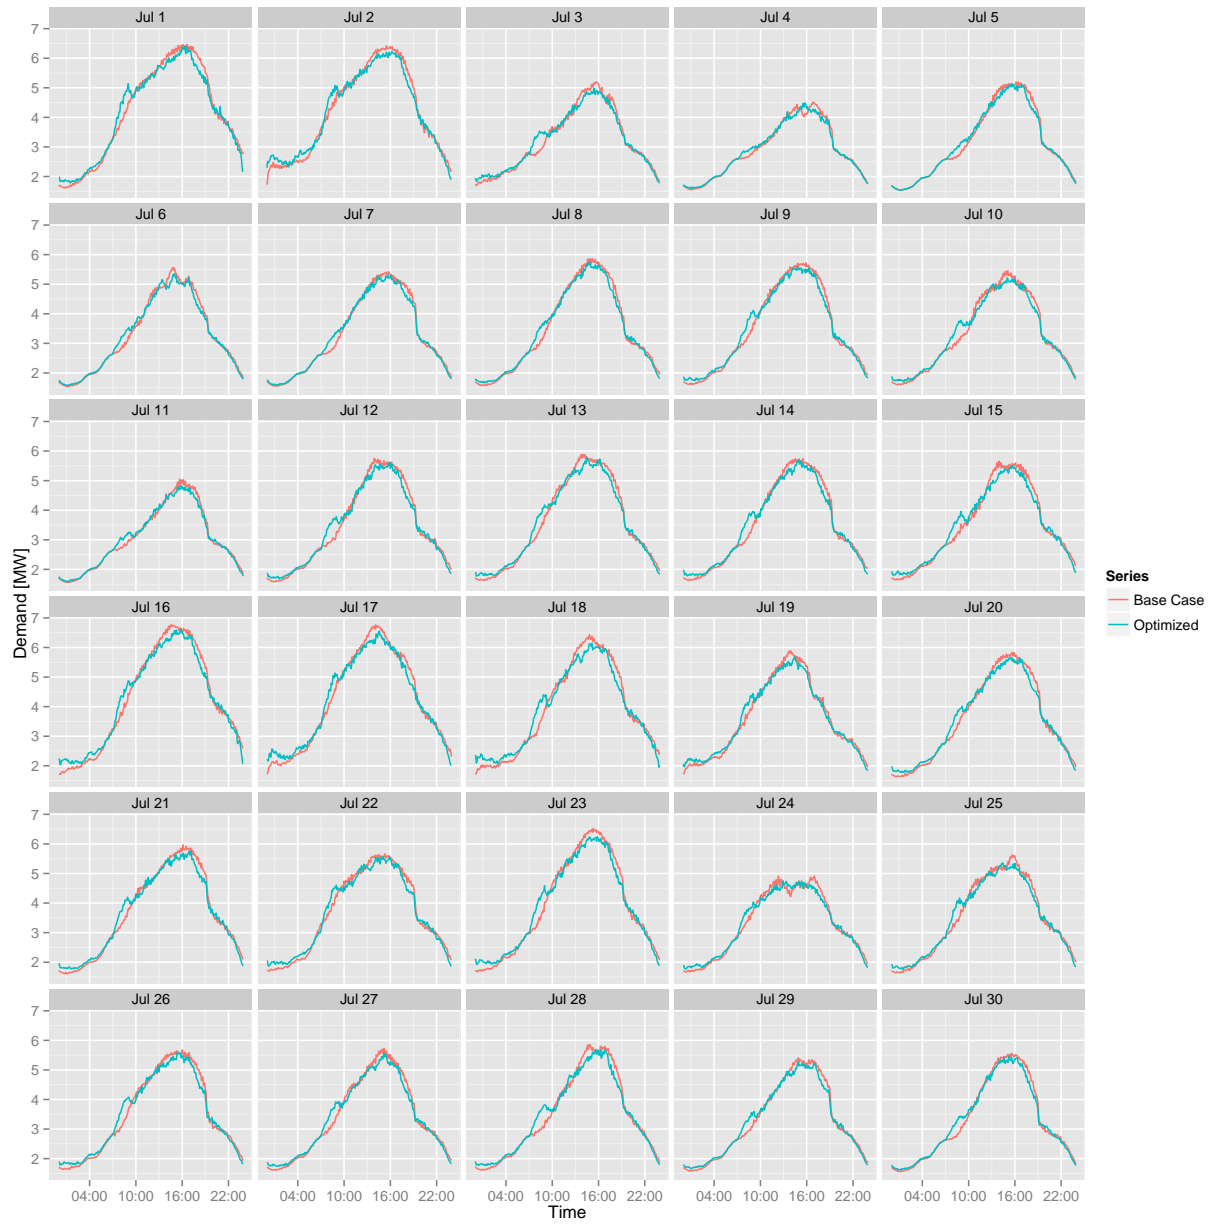


Figure C.35: Feeder demand profiles for New York synthetic price optimization, zero-degree upper boundary case, 70% participation.

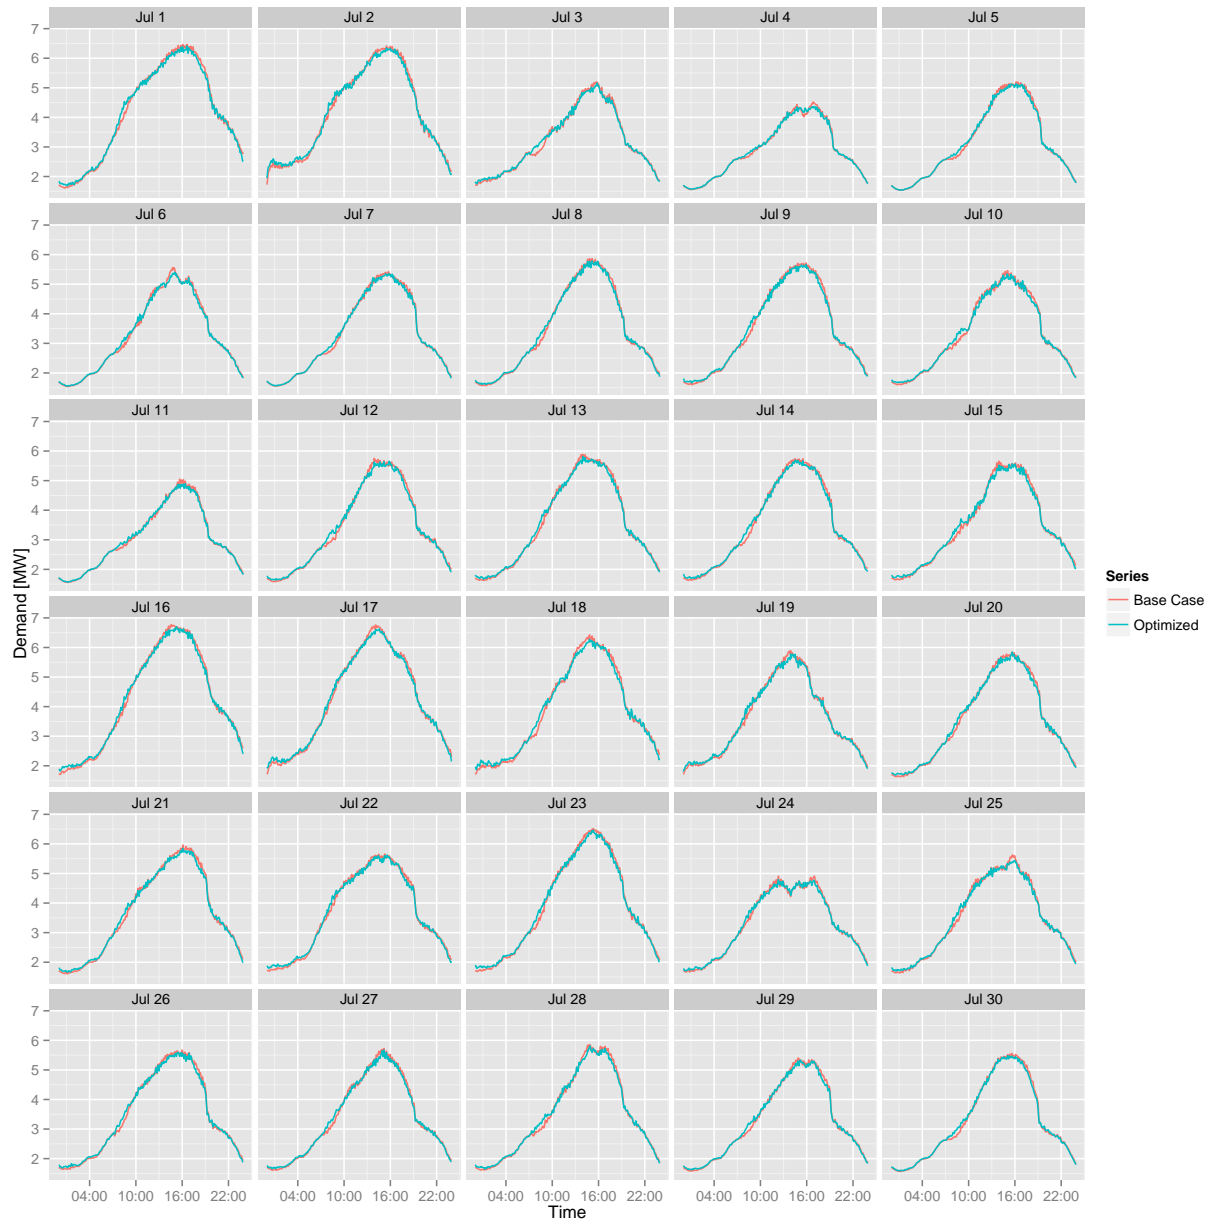


Figure C.36: Feeder demand profiles for New York synthetic price optimization, zero-degree upper boundary case, 30% participation.

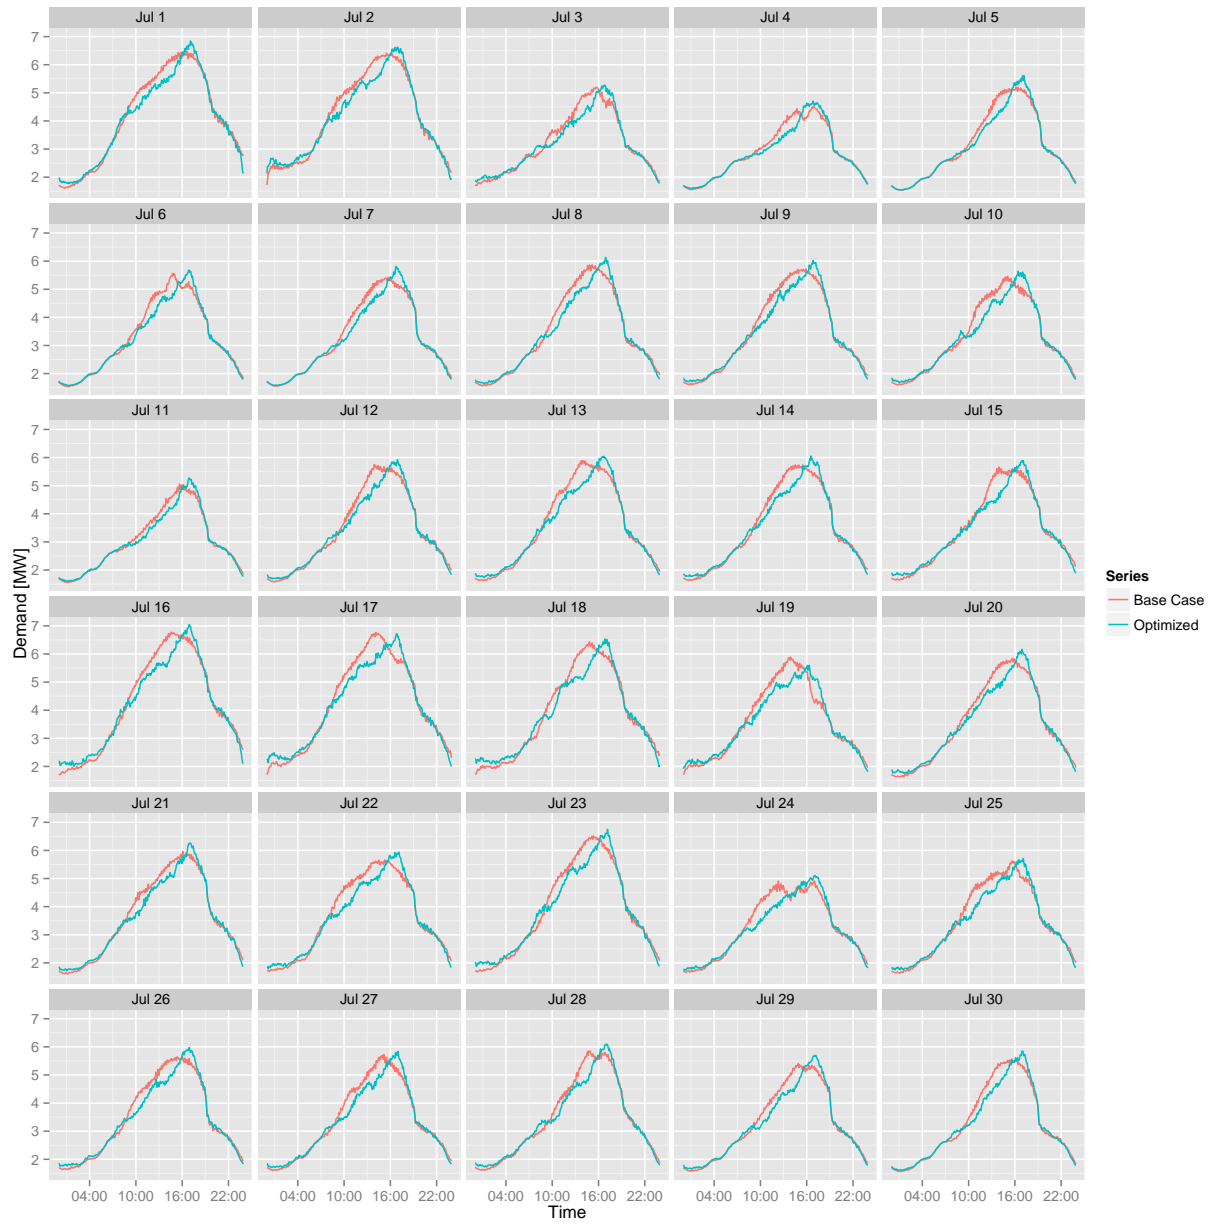


Figure C.37: Feeder demand profiles for New York synthetic price optimization, ramp-return upper boundary case, 70% participation.

Appendix D

Supplement to Load Shape Optimization

D.1 Houston




Figure D.1: Feeder demand profiles for Houston, July 21 load shape optimization, 70% participation.

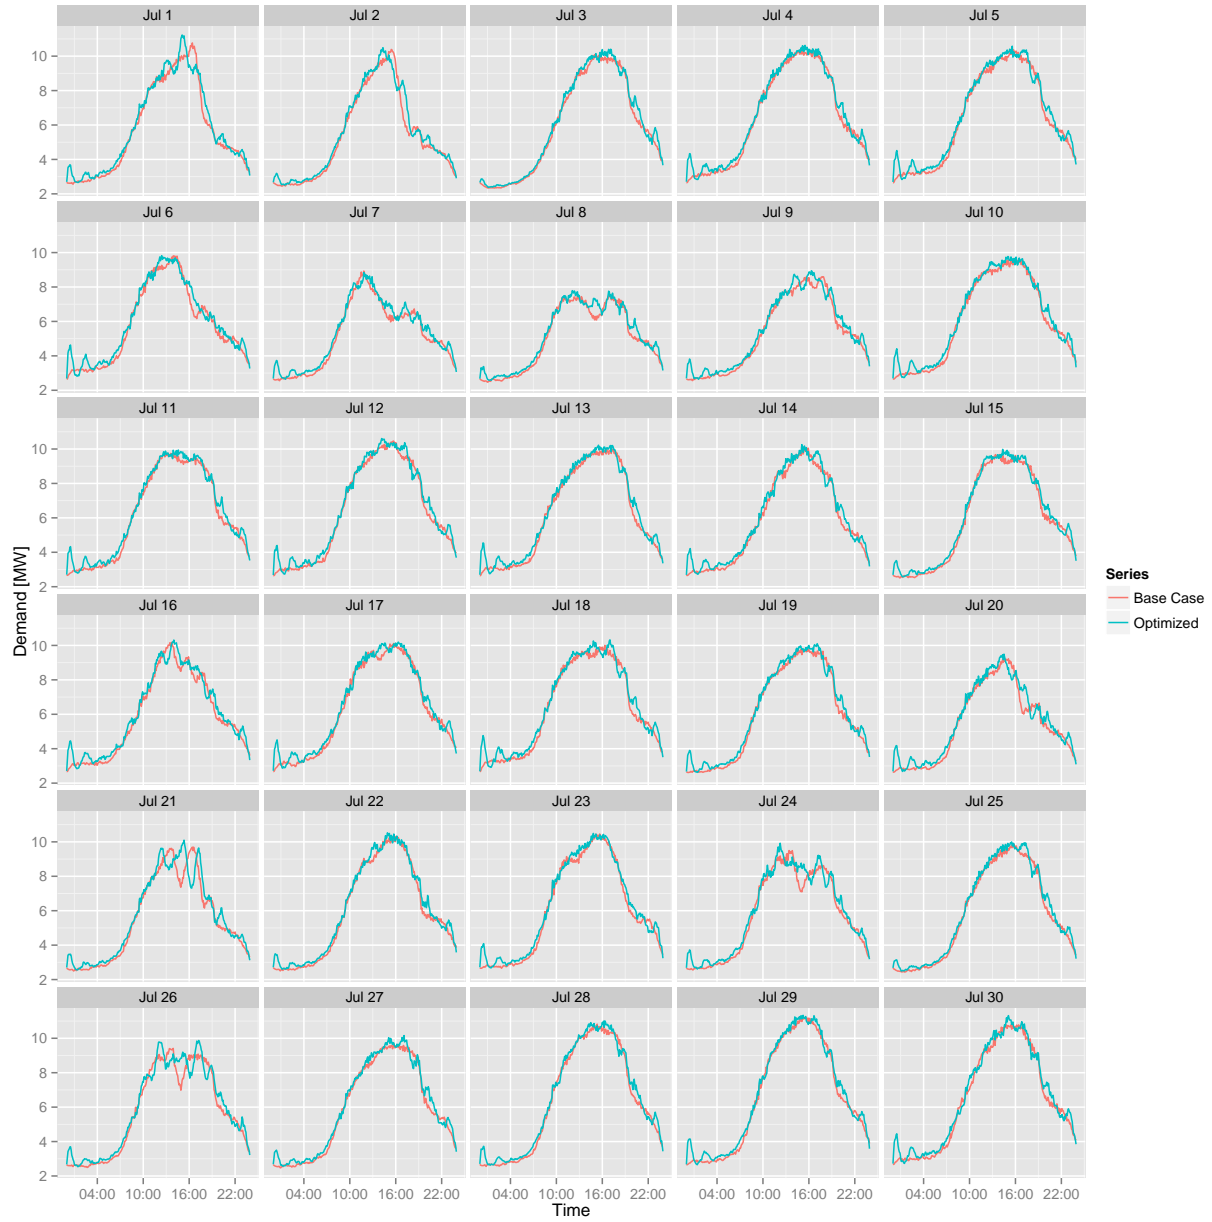


Figure D.2: Feeder demand profiles for Houston load shape optimization, 120 minute moving average, 70% participation.

D.2 Los Angeles

Table D.1: Performance metrics for Los Angeles feeder load shaping optimization, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	0.31	0.20	0.46
Peak Demand [MW]	0.00	-0.03	0.01
Peak to Valley [%]	99.62	98.15	100.43
Load Factor [%]	0.48	0.01	1.06
Ramp [MW]	-0.02	-0.19	0.11




Figure D.3: Total spectral power as a function of frequency bin for Los Angeles feeder load shape optimization, 30% participation.

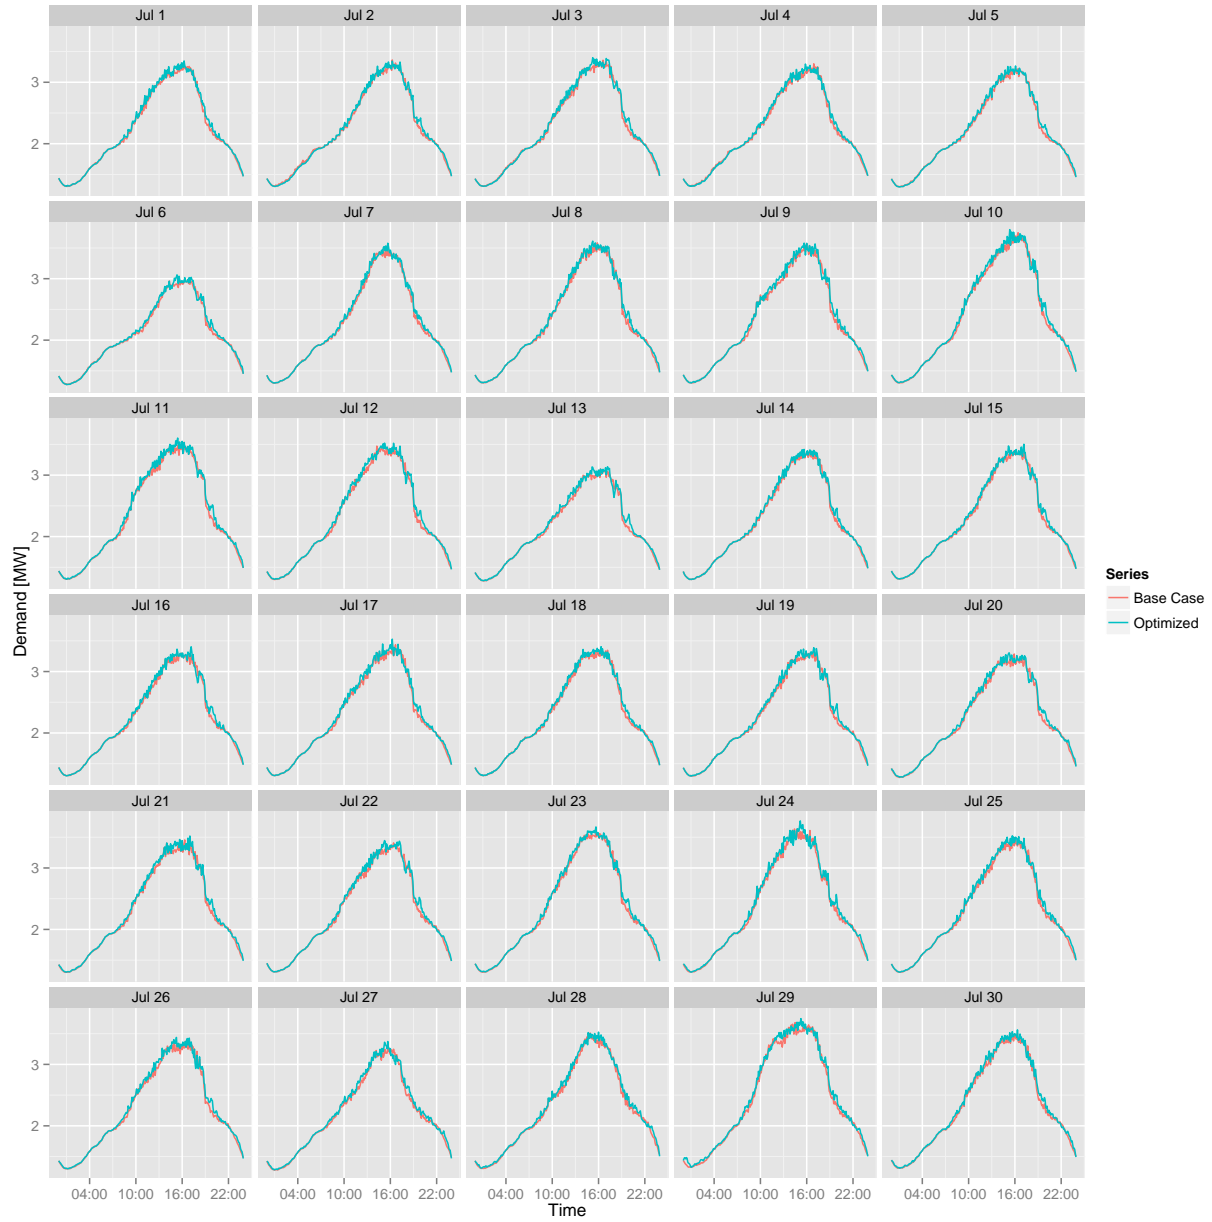


Figure D.4: Feeder demand profiles for Los Angeles load shape optimization, 120 minute moving average, 70% participation.

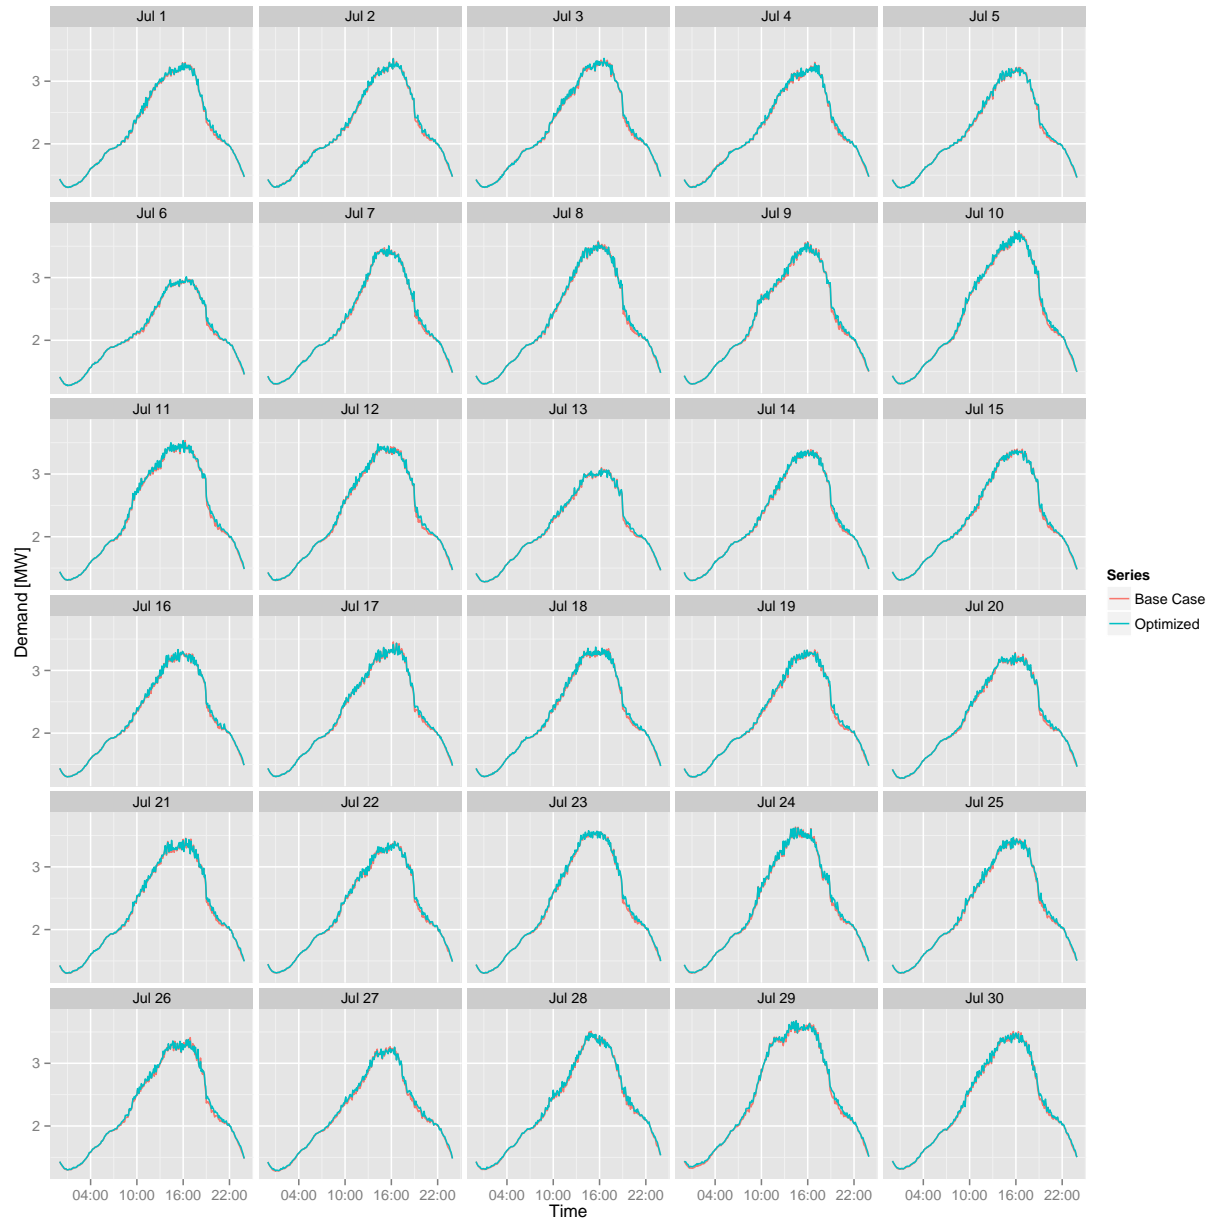


Figure D.5: Feeder demand profiles for Los Angeles load shape optimization, 30% participation.

D.3 New York

Table D.2: Performance metrics for New York feeder load shaping optimization, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	1.03	0.86	1.25
Peak Demand [MW]	-0.02	-0.08	0.03
Peak to Valley [%]	95.70	90.14	99.75
Load Factor [%]	1.00	0.39	2.01
Ramp [MW]	-0.11	-0.58	0.26




Figure D.6: Total spectral power as a function of frequency bin for New York feeder load shape optimization, 30% participation.

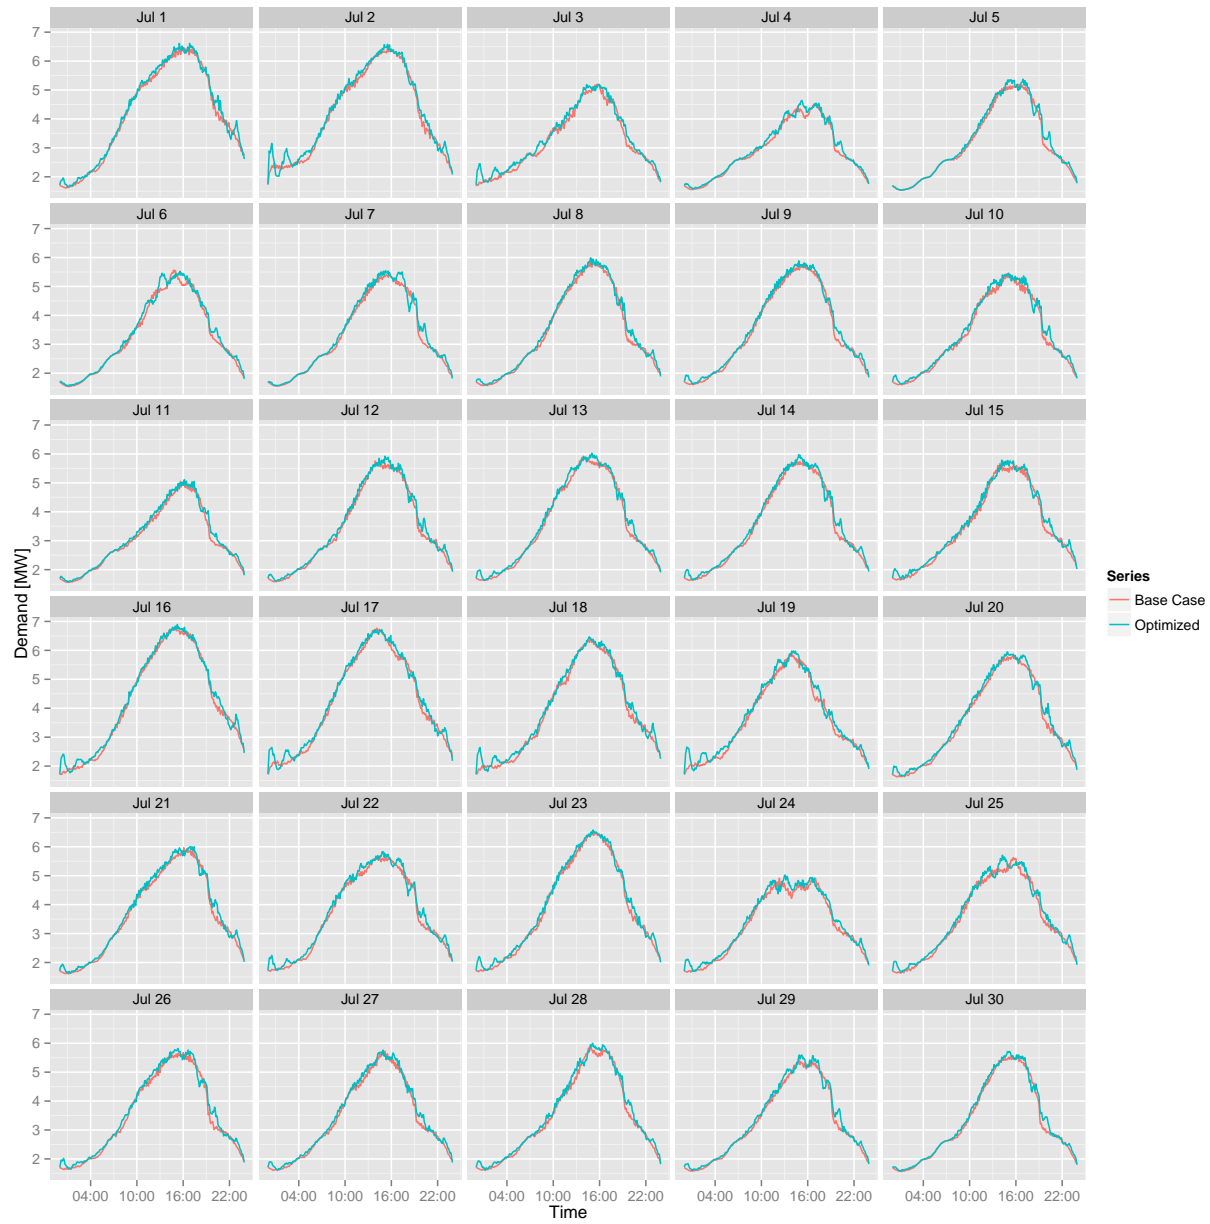


Figure D.7: Feeder demand profiles for New York load shape optimization, 120 minute moving average, 70% participation.

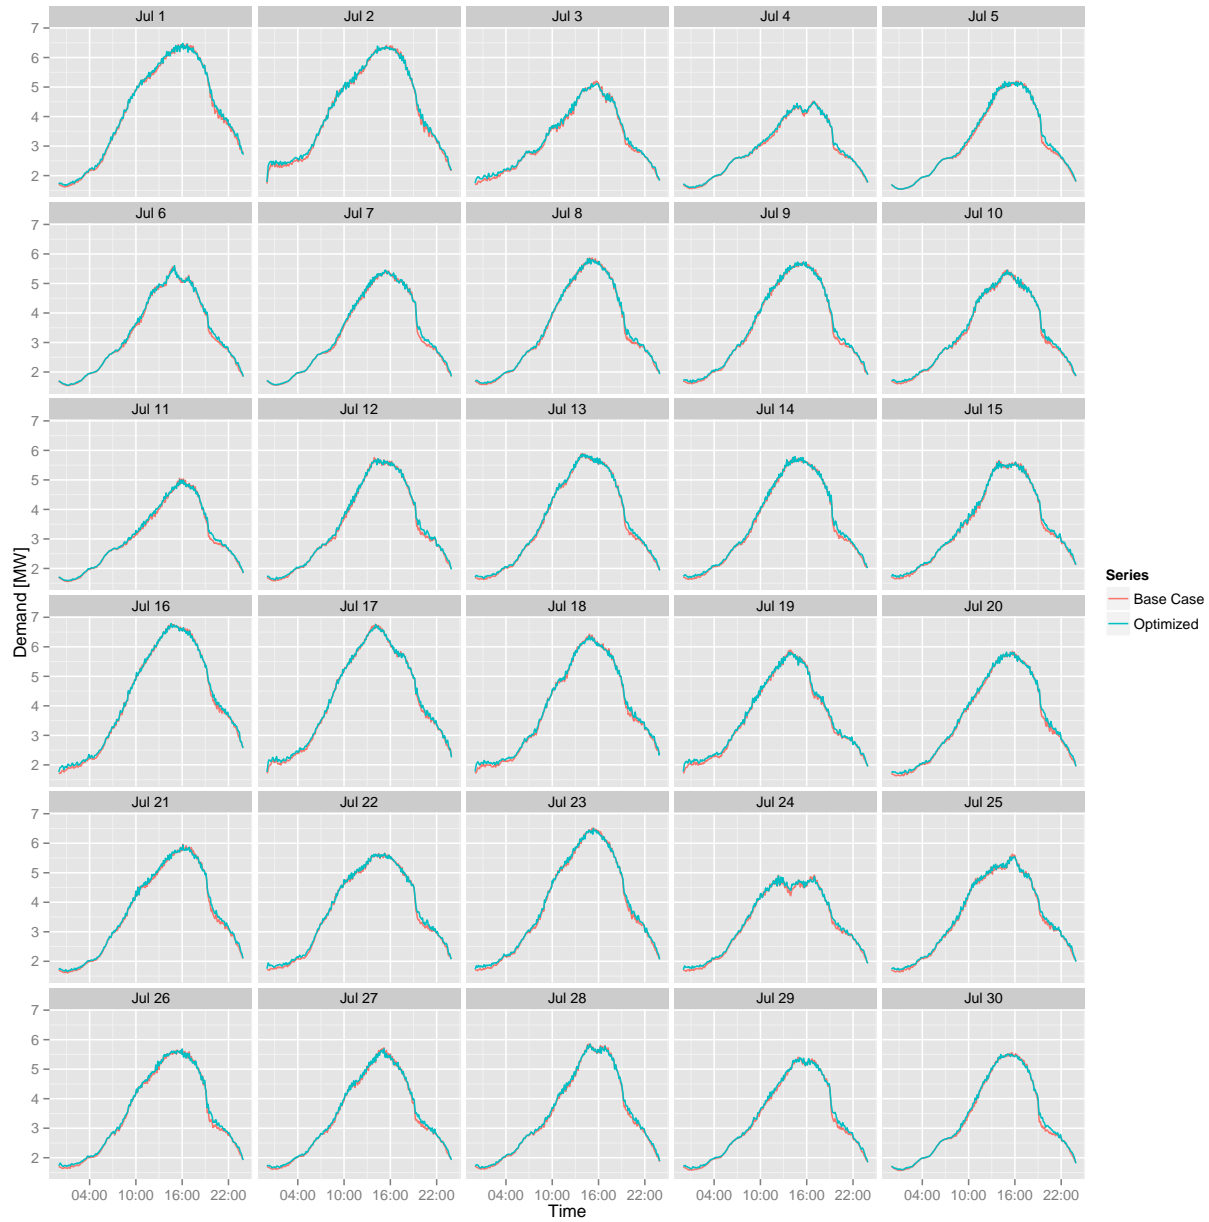


Figure D.8: Feeder demand profiles for New York load shape optimization, 30% participation.




Figure D.9: Feeder demand profiles for New York, July 15 load shape optimization, 70% participation.

Appendix E

Supplement to Rooftop Solar Electric Generation

E.1 Houston




Figure E.1: Houston feeder demand profile for the week of July 1–7 with high rooftop solar penetration.




Figure E.2: Annual Houston feeder demand profile with high rooftop solar penetration.




Figure E.3: Feeder demand profiles for Houston, July 6 load shape optimization, high solar penetration case, 70% participation.

E.2 Los Angeles

Table E.1: Performance metrics for Los Angeles feeder load shaping optimization, high solar penetration case, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	0.28	0.18	0.42
Peak Demand [MW]	-0.04	-0.06	-0.01
Peak to Valley [%]	97.34	93.33	99.40
Load Factor [%]	1.20	0.71	1.61
Ramp [MW]	-0.14	-0.31	0.18

Table E.2: Performance metrics for Los Angeles feeder load shaping optimization, low solar penetration case, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	0.29	0.19	0.39
Peak Demand [MW]	-0.01	-0.03	0.01
Peak to Valley [%]	99.50	98.21	100.11
Load Factor [%]	0.59	0.23	0.99
Ramp [MW]	-0.02	-0.17	0.13




Figure E.4: Total spectral power as a function of frequency bin for Los Angeles feeder load shape optimization, high solar penetration case, 30% participation.




Figure E.5: Total spectral power as a function of frequency bin for Los Angeles feeder load shape optimization, low solar penetration case, 30% participation.




Figure E.6: Annual Los Angeles feeder demand profile with high rooftop solar penetration.




Figure E.7: Feeder demand profiles for Los Angeles load shape optimization, high solar penetration case, 30% participation.

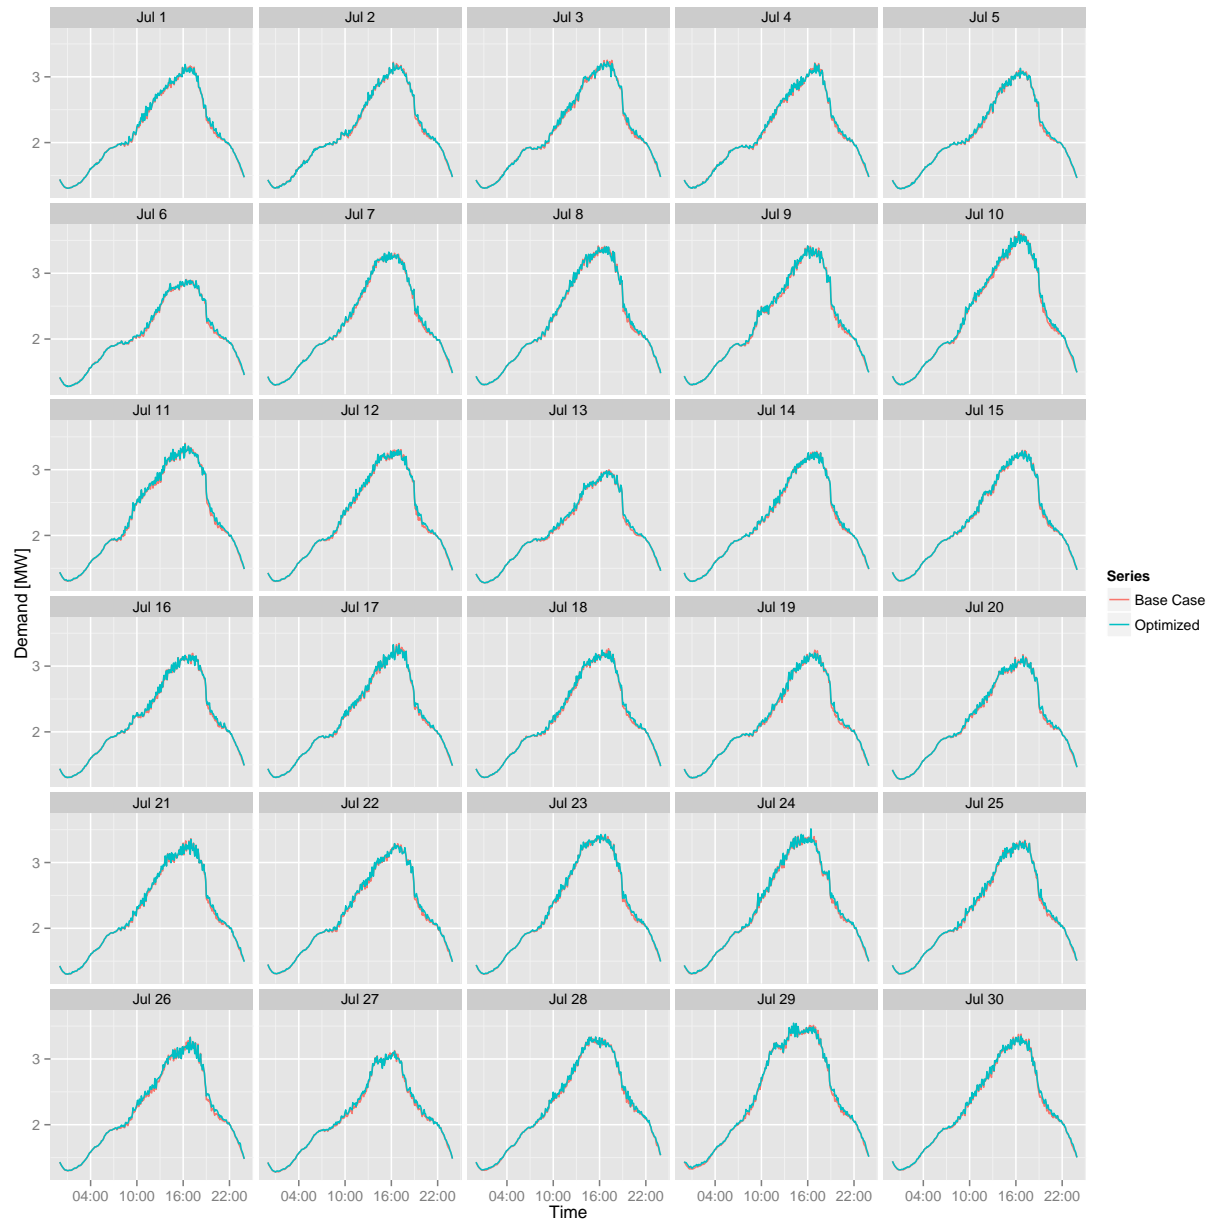


Figure E.8: Feeder demand profiles for Los Angeles load shape optimization, low solar penetration case, 30% participation.




Figure E.9: Feeder demand profiles for Los Angeles, July 3 load shape optimization, high solar penetration case, 70% participation.

E.3 New York

Table E.3: Performance metrics for New York feeder load shaping optimization, high solar penetration case, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	0.98	0.72	1.25
Peak Demand [MW]	-0.10	-0.20	-0.02
Peak to Valley [%]	93.61	86.46	97.91
Load Factor [%]	2.05	1.33	3.22
Ramp [MW]	-0.36	-0.67	-0.03

Table E.4: Performance metrics for New York feeder load shaping optimization, low solar penetration case, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	1.01	0.80	1.27
Peak Demand [MW]	-0.02	-0.07	0.07
Peak to Valley [%]	95.66	89.60	99.35
Load Factor [%]	0.98	-0.01	1.97
Ramp [MW]	-0.18	-0.74	0.23




Figure E.10: Total spectral power as a function of frequency bin for New York feeder load shape optimization, high solar penetration case, 30% participation.




Figure E.11: Total spectral power as a function of frequency bin for New York feeder load shape optimization, low solar penetration case, 30% participation.




Figure E.12: Annual New York feeder demand profile with high rooftop solar penetration.




Figure E.13: New York feeder demand profile for the week of July 1–7 with high rooftop solar penetration.




Figure E.14: Feeder demand profiles for New York load shape optimization, low solar penetration case, 30% participation.

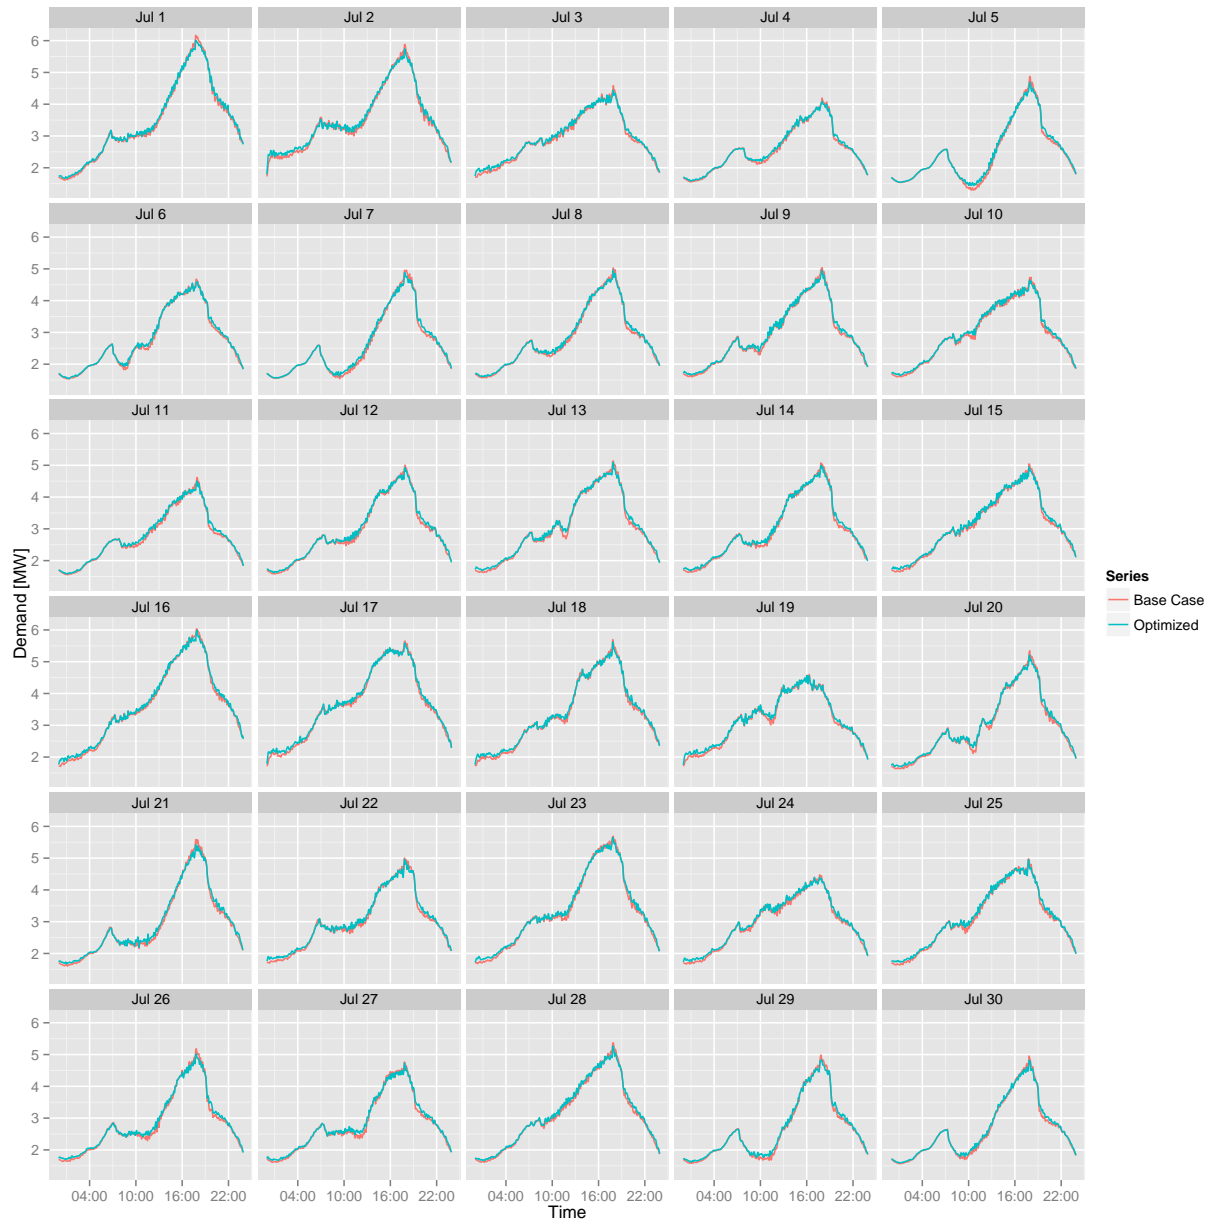


Figure E.15: Feeder demand profiles for New York load shape optimization, high solar penetration case, 30% participation.




Figure E.16: Feeder demand profiles for New York, July 2 load shape optimization, high solar penetration case, 70% participation.




Figure E.17: Feeder demand profiles for New York, July 13 load shape optimization, high solar penetration case, 70% participation.




Figure E.18: Feeder demand profiles for New York, July 20 load shape optimization, high solar penetration case, 70% participation.

Appendix F

Supplement to Utility Scale Wind Generation

F.1 Houston

Table F.1: Performance metrics for Houston feeder load shaping optimization, high wind penetration case, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	2.07	1.77	2.44
Peak Demand [MW]	-0.28	-0.80	0.05
Peak to Valley [%]	77.38	62.83	95.33
Load Factor [%]	2.72	0.68	5.27
Ramp [MW]	-2.61	-4.94	-0.43

Table F.2: Performance metrics for Houston feeder load shaping optimization, low wind penetration case, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	2.09	1.84	2.37
Peak Demand [MW]	-0.13	-0.53	0.07
Peak to Valley [%]	89.55	82.68	97.20
Load Factor [%]	1.75	0.50	4.08
Ramp [MW]	-1.11	-2.88	0.22




Figure F.1: Total spectral power as a function of frequency bin for Houston feeder load shape optimization, high wind penetration case, 30% participation.




Figure F.2: Total spectral power as a function of frequency bin for Houston feeder load shape optimization, low wind penetration case, 30% participation.

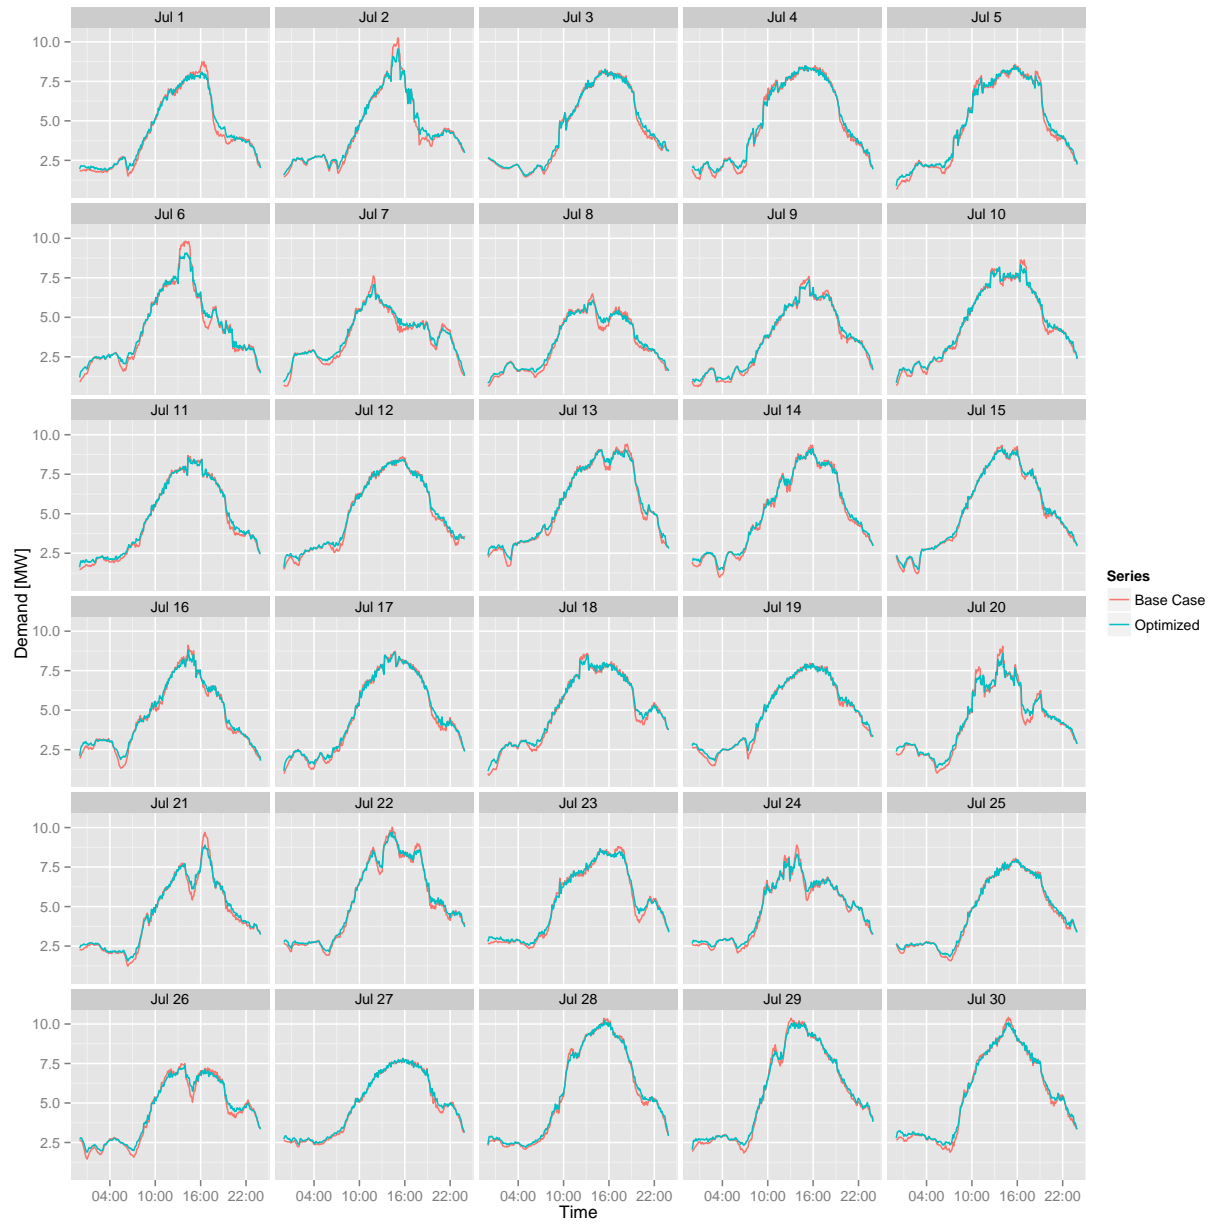


Figure F.3: Feeder demand profiles for Houston load shape optimization, high wind penetration case, 30% participation.

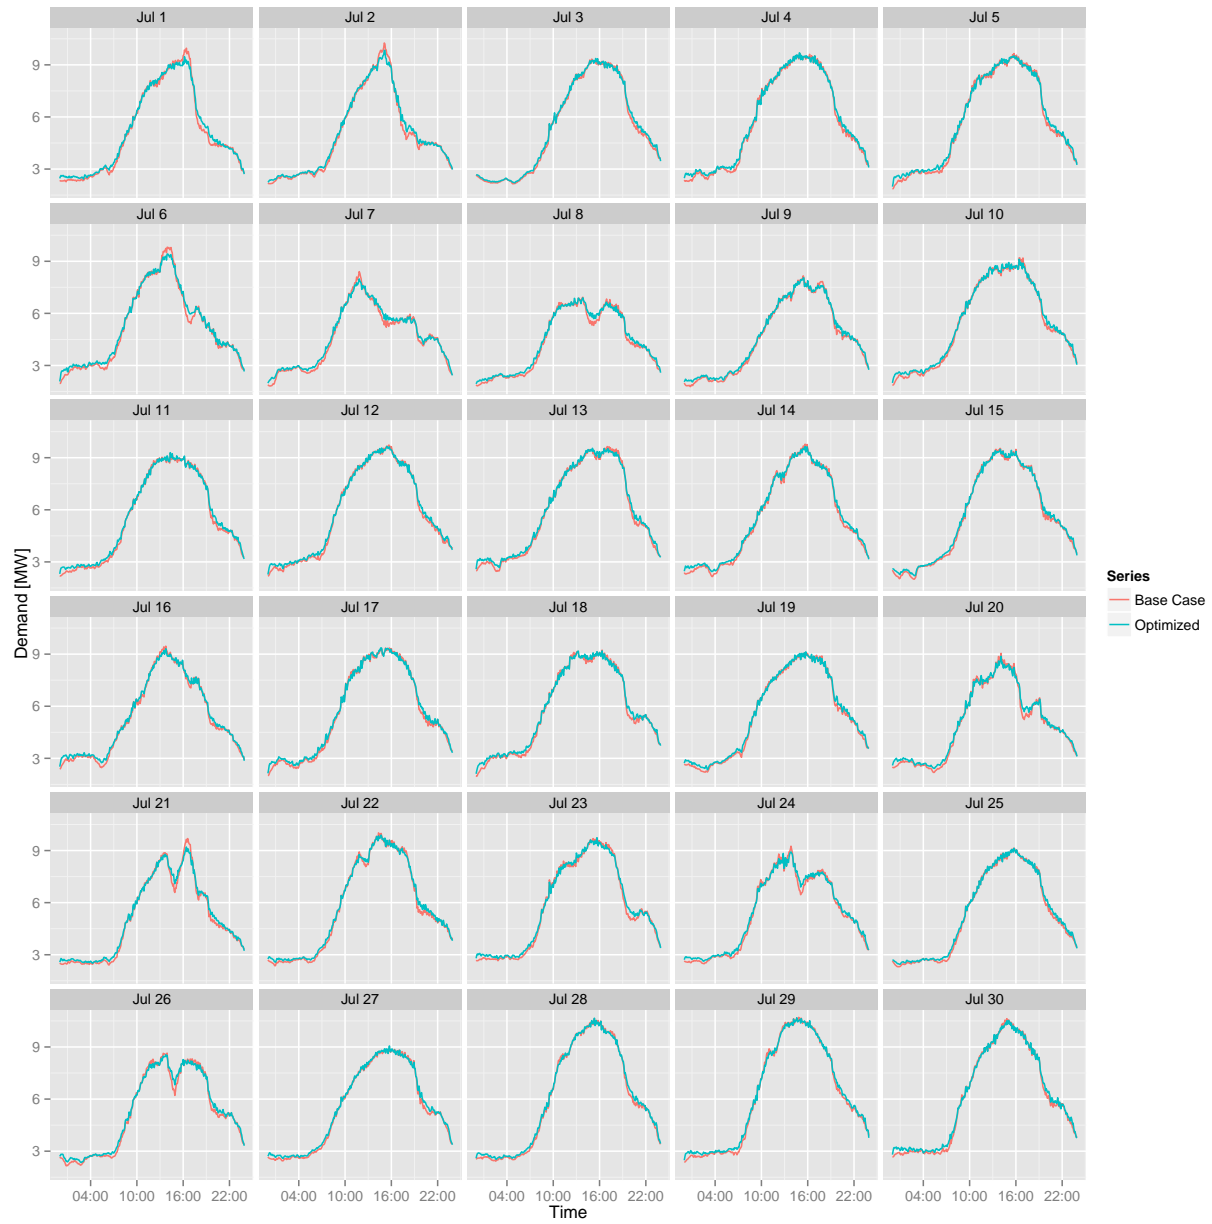


Figure F.4: Feeder demand profiles for Houston load shape optimization, low wind penetration case, 30% participation.

F.2 Los Angeles

Table F.3: Performance metrics for Los Angeles feeder load shaping optimization, high wind penetration case, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	0.34	0.25	0.43
Peak Demand [MW]	-0.05	-0.18	-0.01
Peak to Valley [%]	96.64	90.18	99.97
Load Factor [%]	1.61	0.79	3.12
Ramp [MW]	-0.33	-1.01	0.09

Table F.4: Performance metrics for Los Angeles feeder load shaping optimization, low wind penetration case, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	0.30	0.21	0.40
Peak Demand [MW]	-0.02	-0.07	0.00
Peak to Valley [%]	99.02	96.77	100.48
Load Factor [%]	0.80	0.24	1.67
Ramp [MW]	-0.13	-0.45	0.15




Figure F.5: Total spectral power as a function of frequency bin for Los Angeles feeder load shape optimization, high wind penetration case, 30% participation.




Figure F.6: Total spectral power as a function of frequency bin for Los Angeles feeder load shape optimization, low wind penetration case, 30% participation.

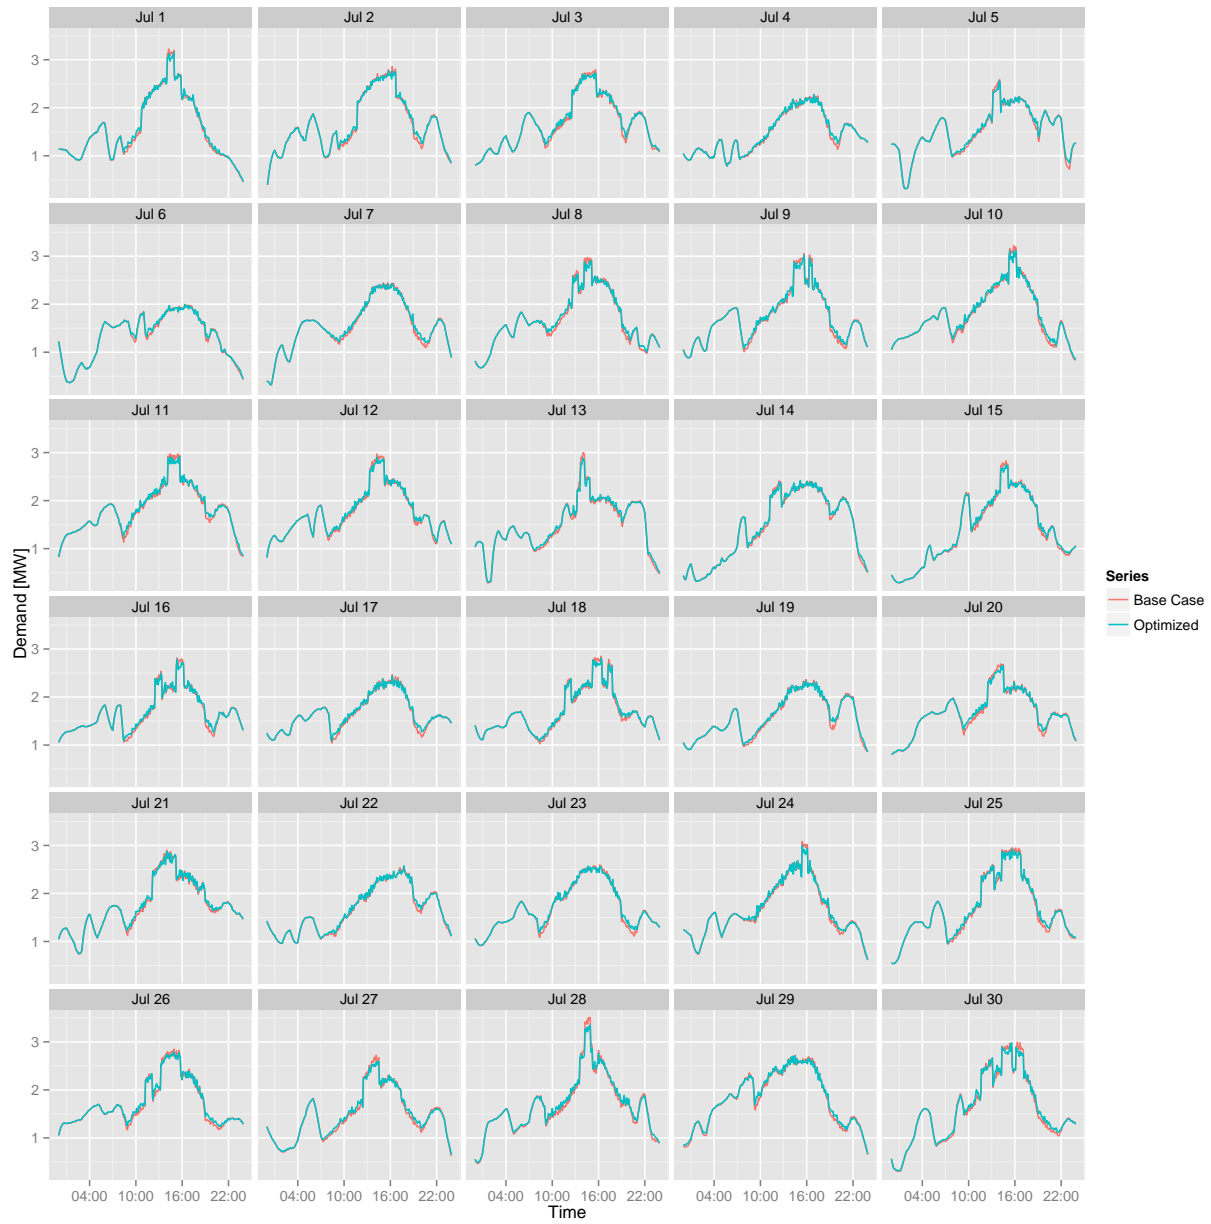


Figure F.7: Feeder demand profiles for Los Angeles load shape optimization, high wind penetration case, 30% participation.

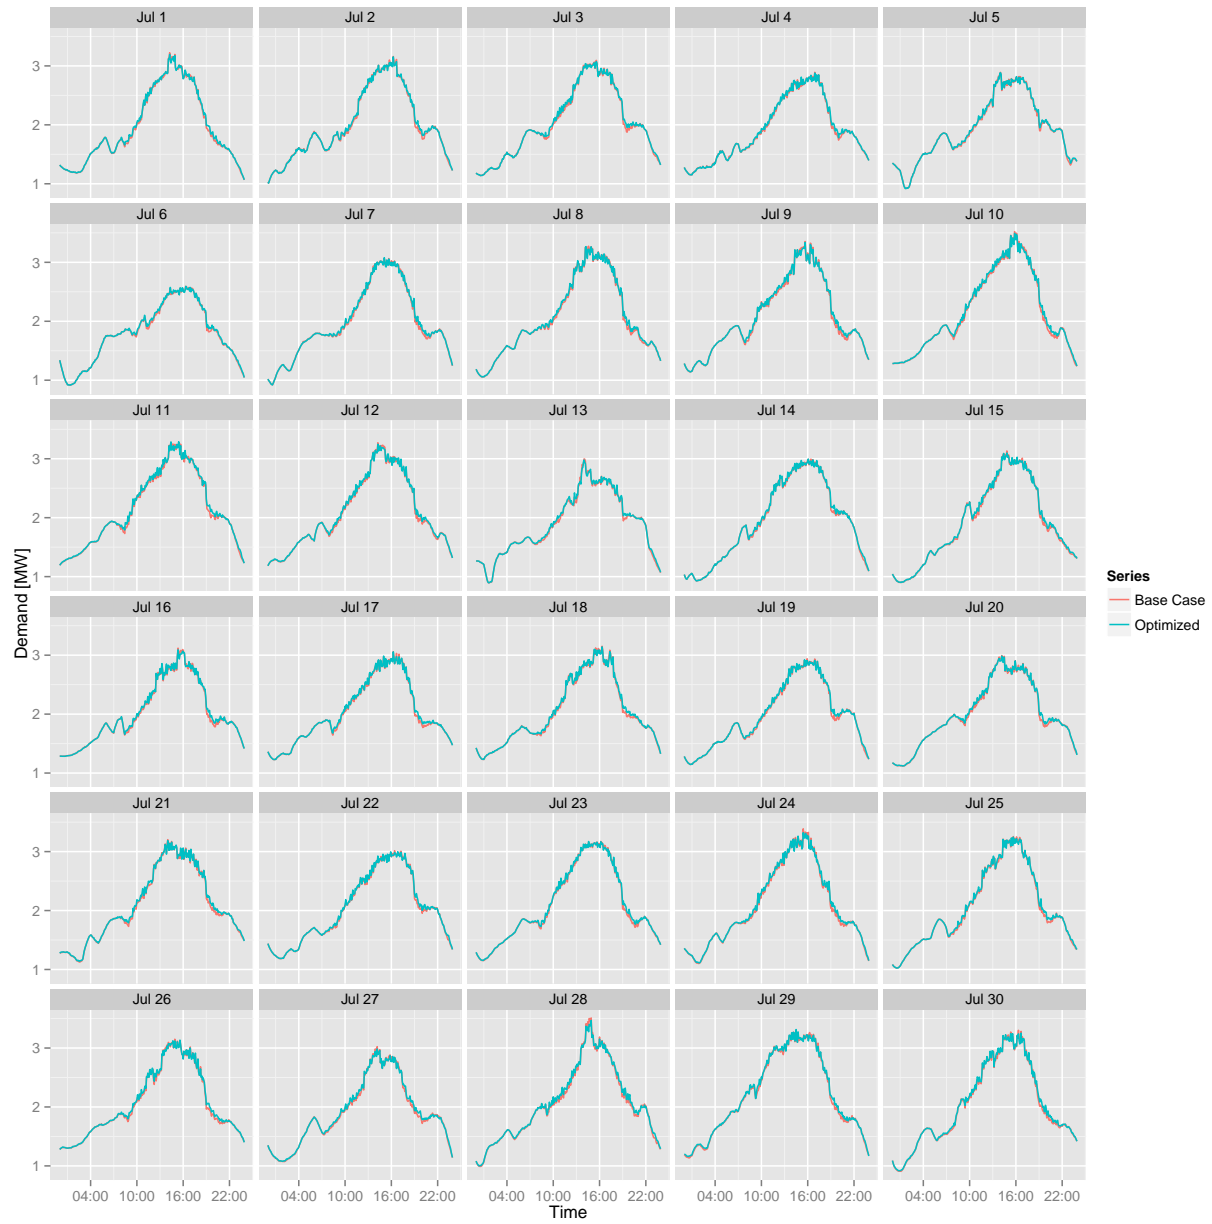


Figure F.8: Feeder demand profiles for Los Angeles load shape optimization, low wind penetration case, 30% participation.




Figure F.9: Feeder demand profiles for Los Angeles, July 23 load shape optimization, high wind penetration case, 70% participation.

F.3 New York

Table F.5: Performance metrics for New York feeder load shaping optimization, high wind penetration case, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	1.07	0.89	1.38
Peak Demand [MW]	-0.07	-0.21	0.02
Peak to Valley [%]	89.05	76.48	98.64
Load Factor [%]	1.72	0.71	3.21
Ramp [MW]	-0.26	-1.13	0.38

Table F.6: Performance metrics for New York feeder load shaping optimization, low wind penetration case, 30% participation.

	Mean	Min	Max
Electric Consumption [MWh]	1.04	0.85	1.31
Peak Demand [MW]	-0.03	-0.09	0.01
Peak to Valley [%]	94.17	88.02	99.27
Load Factor [%]	1.18	0.55	2.32
Ramp [MW]	-0.17	-0.75	0.31




Figure F.10: Total spectral power as a function of frequency bin for New York feeder load shape optimization, high wind penetration case, 30% participation.




Figure F.11: Total spectral power as a function of frequency bin for New York feeder load shape optimization, low wind penetration case, 30% participation.

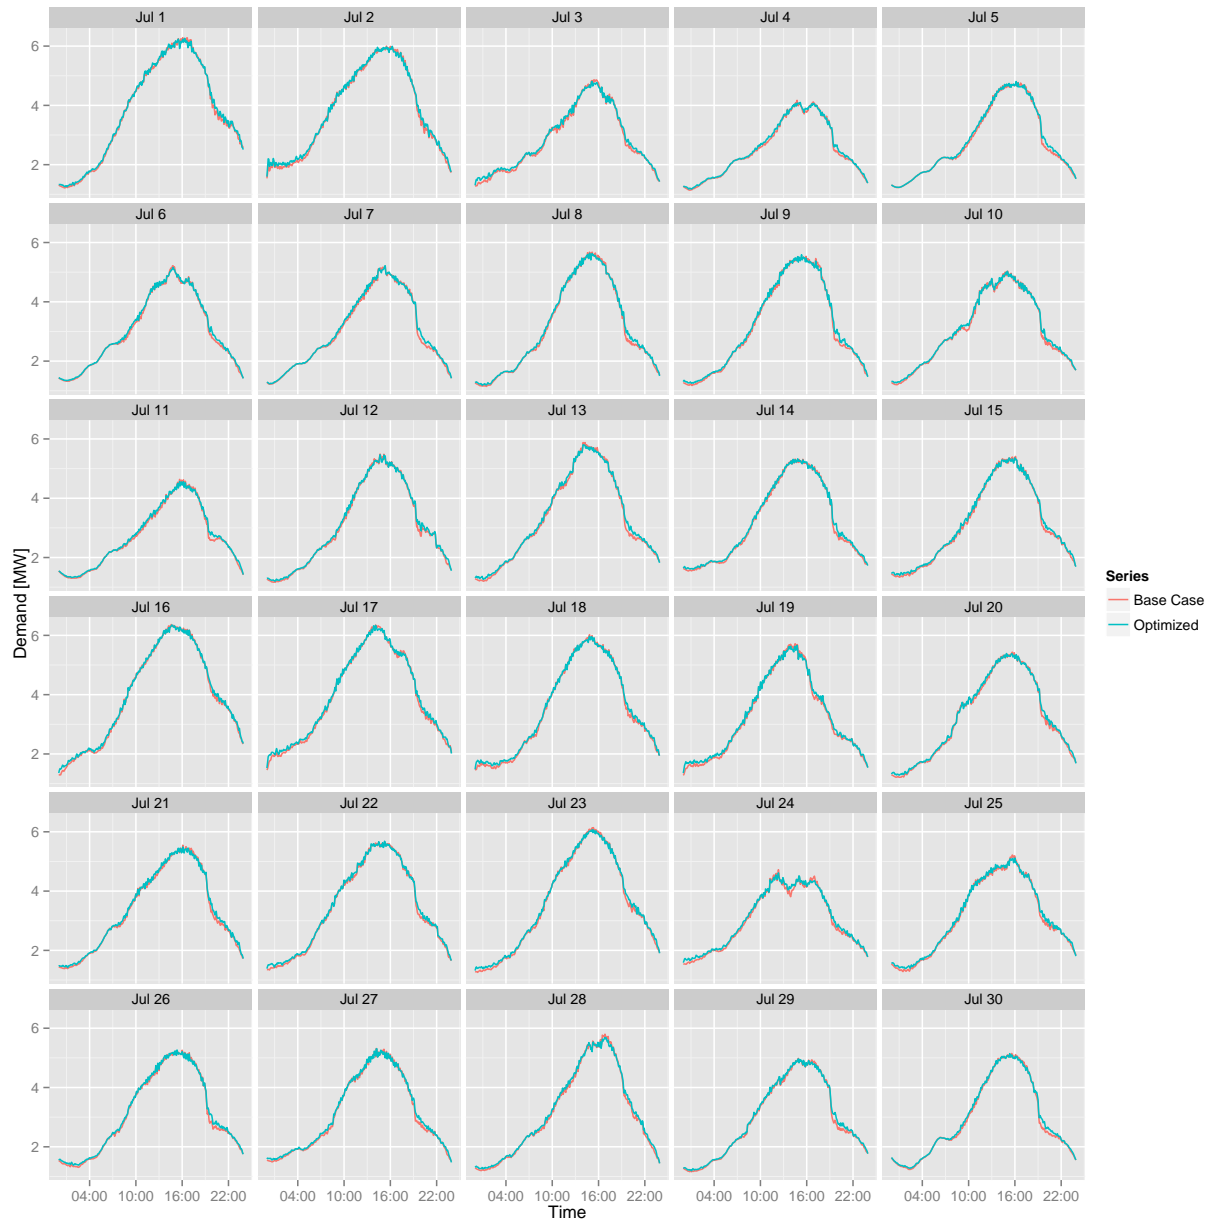


Figure F.12: Feeder demand profiles for New York load shape optimization, low wind penetration case, 30% participation.

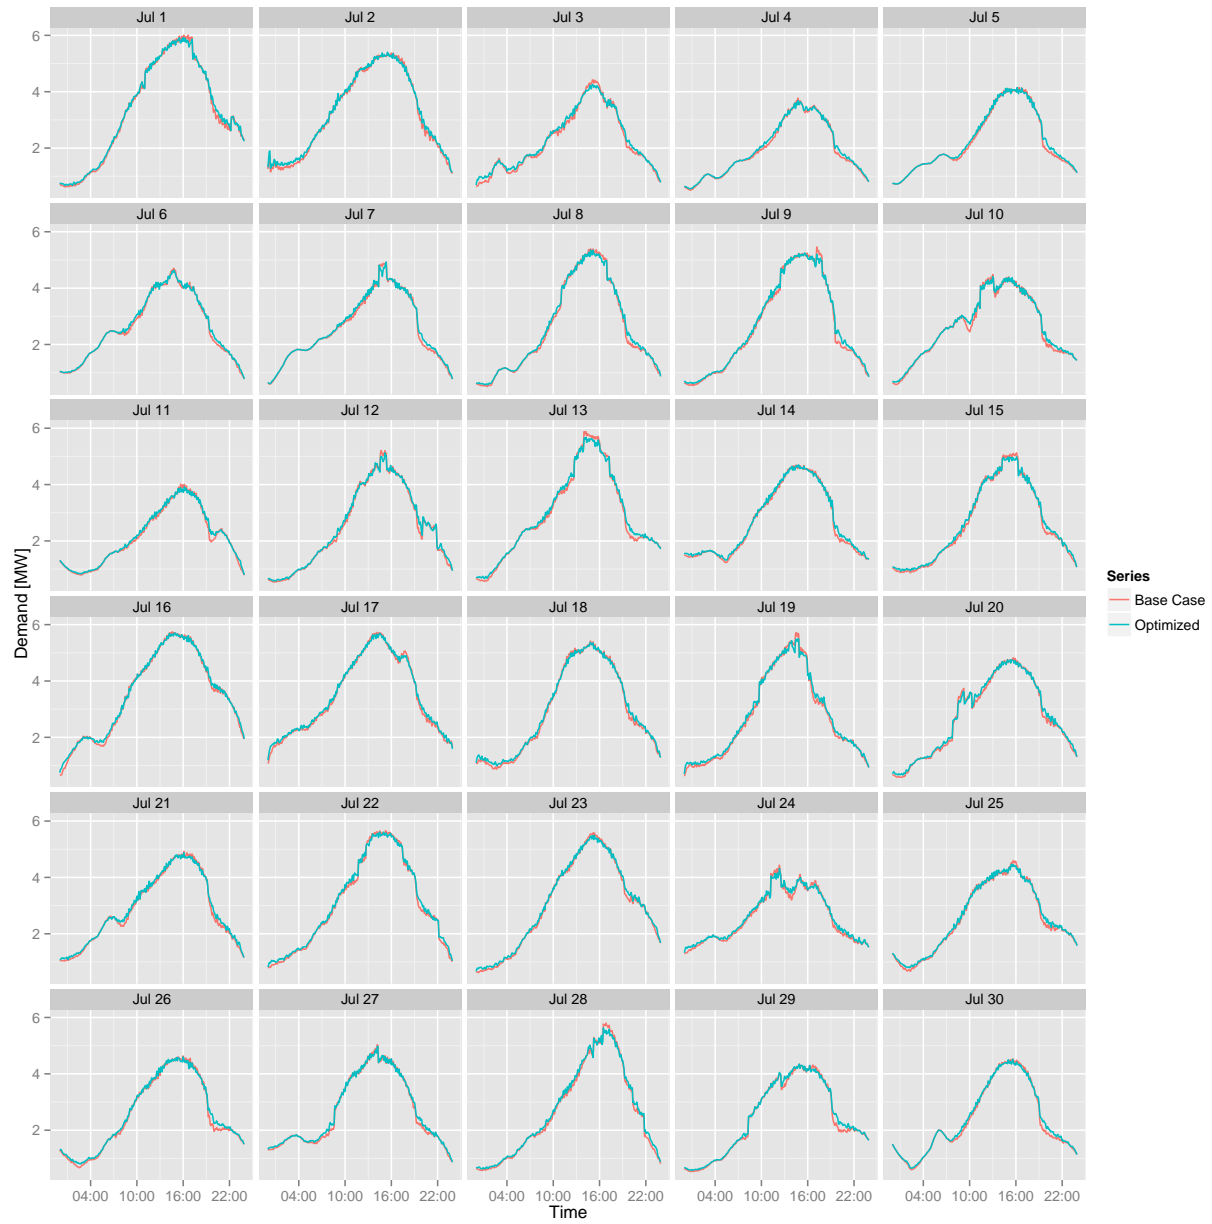


Figure F.13: Feeder demand profiles for New York load shape optimization, high wind penetration case, 30% participation.




Figure F.14: Feeder demand profiles for New York, July 24 load shape optimization, high wind penetration case, 70% participation.

Appendix G

Assessment of Building Thermal Storage Efficiency

The ability of the controller to shift demand from one time of day to another is related to the round-trip storage efficiency of the building envelope. Energy may be invested in the building envelope in order to precool the building and reduce cooling energy at a later time. Some of the precooling effect is lost to the environment and internal gains. The amount of cooling energy avoided at some later time due to precooling is in theory the amount of precooling invested minus the amount lost. The ratio of the amount of cooling energy avoided to the amount of precooling invested is defined as the storage efficiency. High storage efficiency implies that precooling may occur early in the day, and that the effects of precooling may be utilized late into the afternoon.

To test the storage efficiency of the buildings in the Los Angeles feeder, an experiment is constructed in which the feeder is subjected to a demand limiting optimization similar to that described in Chapter 6. In this case, however, the controller is only allowed to make set point adjustments below the base case cooling set point. This forces the controller to use precooling to limit peak demand, rather than a combination of precooling and cooling set point set-ups. The controller is given the same lower boundary as the cases in Chapter 6. For this experiment, all 100% of the homes on the feeder are assumed to participate in the optimization.

Following optimization, the storage efficiency is calculated by taking the difference at each time step between a house's base case and optimized cooling demands. Where optimized cooling demand is higher, the difference is considered an 'investment'; where base case cooling demand is higher, the difference is considered a 'return'. The ratio of the sum of returns to the sum of

investments in a daily basis is the daily storage efficiency. This calculation occurs for each house and each day of the month of July, giving 41,106 observations of storage efficiency.

Storage efficiency calculated in this way is shown as a histogram in Figure G.1. It should be noted that values greater than and equal to one were removed. The existence of these high efficiencies is of no concern. Inspection of the data showed that many of these were the result of rounding error. Efficiencies of unity occurred when base case and optimized case were nearly identical except for small offsets in HVAC cycling.

The distribution mean for all observations is 0.70 or 70% storage efficiency. The distribution shows two modes clearly, centered around 62% and 81%, although it is not obvious why. As it relates to the optimizations presented in this work, the storage efficiency plot suggests that the controller will be limited in the amount of precooling it will utilize. That is, precooling at times that result in low efficiency will be avoided in favor of times that have higher efficiency, or eliminated altogether.

This is, however, only one distribution of results. For control strategies in which the objective function differs, the controller will make different trade-offs that result in the best objective function value subject to the storage efficiency of the envelope. This does not necessarily mean that storage efficiency distributions would change under different objective functions, but it stands to reason that if the control strategy results in additional precooling, the distribution of efficiencies would reflect that decision. Perhaps more precisely, the distribution of efficiencies would inform the decision.

The interaction between storage efficiency, objective function, and the control decisions that result is very complex. Unfortunately, the results below do not provide enough insight to say with confidence what effect storage efficiency has on the differences seen in precooling between Houston, Los Angeles and New York feeders. Further investigation would be required to address this issue. It should be noted that the quantification of a building's thermal storage efficiency does not present itself in the literature; this is an interesting area for recommended future research.




Figure G.1: Histogram showing the thermal storage efficiency of the building envelope from demand limiting optimization of the Los Angeles feeder in July with occupied cooling set point upper boundary.

Appendix H

Assessment of Zero Demand Optimization

Chapter 8 proposes a feeder reference demand curve construction based on a simple moving average of base case feeder demand. The construction is intended to illicit a controller response that in aggregate leads to a smooth feeder demand curve. As seen in Chapter 6, some smoothing is observed when the controller is asked to simply minimize demand over the entire planning horizon. Although smoothing in that case is a side effect of demand minimization, it suggests that demand minimization may be accomplished using the load shaping method and an alternative formulation of reference demand curve. This observation also suggests that the shape of the reference curve may not matter at all and that the constraints discussed in Subsection 8.1.1 may be ignored.

To test this hypothesis, a “zero demand” house reference demand curve is employed in the load shaping method. Recall that the objective of optimization is simply the minimization of the sum of squared error between the house demand curve and the house reference demand curve. In this case, the house reference demand curve is flat and has a value of zero for all time steps. Although zero demand is not physically achievable by the controller, this reference curve will encourage the controller to flatten demand as much as possible. In theory, it should encourage the controller to minimize energy consumption over the entire planning horizon provided that comfort constraints. i.e. temperature boundaries, are observed. Compared to the load shaping results, this experiment illustrates the extent of demand reductions and energy savings possible using the load shaping methodology. In practical terms, this construction represents a default control strategy in which there is no communication from the utility, either because of communication network failure

or owner opt-out.

The experiment is conducted on the Houston feeder model for the month of July, using a participation rate of 70%. It is repeated for the high wind case to provide a contrast to the results of Chapter 10. Metrics for both cases are shown in Tables H.1 and H.2, and in Figures H.1 and H.2. Feeder demand curves can be seen in Figures H.3 and H.4.

Generally speaking, the monthly results show a more consistent demand at the beginning of the day compared to the demand limiting optimization. That is, instead of a high demand that quickly diminishes, the demand stays relatively flat in the very early morning hours. Evening hours show increased demand compared to the demand limiting optimization. It is assumed that this increased demand is the result of demand deferral from earlier in the day. Metrics suggest that the zero demand optimization is better able to reduce peak demand and peak to valley ratio, and improve load factor compared with most cases investigated.

By contrast, energy consumption is increased compared to all cases investigated, presumably due to the increased amount of demand shifting. Ramping is somewhat inconsistent, showing both positive and negative differences from the base case; this is not terribly surprising since reduced ramping is not an objective of the optimization. While the constant power portion of the frequency spectrum is marginally increased, so too are the twelve to four hour ranges

Because the controller has no awareness of feeder demand, it is unable to remove any variations introduced by wind. This is most obvious in the high wind plots which show large hour to hour variations. Although minimizing demand does not address wind produced variations directly, some benefit can still be observed. As with the smoothing seen in previous cases, this is simply a side effect of shifting demand. Metrics are consistent with the previous case, showing improvements in the same areas. It is assumed that the improvements are due solely to the reduction in demand, and not related to any synergies between wind generation and the zero demand optimization.

This result further highlights the utility of the load shaping methodology when information about feeder demand is available to drive controller decisions. If the variations are from sources not related to the demand being manipulated, then the demand limiting and zero demand methods will

not be as effective as the load shaping method; neither are able to eliminate (or create) demand during a specific time to offset demand elsewhere in the system. The high wind case illustrates this very well. Yet the wind case is but one example. It can be argued that the purpose of the high wind case is not to show how the method can be used to manage demand variations introduced by wind, rather its purpose is to demonstrate how the load shaping method can be used to change the shape of the demand in an arbitrary way.

In Section 8.1 it was stated that the desired outcome of the optimization may inform the formation of the reference demand curve. For example, if the utility is concerned about reducing peak demand, then this experiment shows that a zero demand reference curve can produce an outcome similar to the demand limiting case if no wind or solar is present. However, if the utility desires to smooth the demand curve, the load shaping optimization does produce a smoother response compared to both demand limiting and zero demand methods under the same conditions. July 21 and July 26 illustrate this well; peaks and troughs in the optimized results of Chapter 8 are smaller compared to those seen in Chapter 6 and those shown here. This is further supported by the ramping metrics and frequency analysis.

This experiment suggests that the selection of reference curve is important to the objective one is trying to achieve. Indeed, the demand limiting, zero demand and load shaping results are all very similar when no wind is present. This result does not prove that the choice of curve is immaterial; rather, it indicates that the load shaping method is a generalization of the demand limiting method. That is, the load shaping method using a zero demand reference curve can produce the demand limiting result. Based on the results of the experiment, the converse is not true.

When demand reduction is the desired objective, it appears that the constraints proposed for the reference curve can be violated. Whether or not this is strictly true for a load shape intended to smooth the demand curve remains an open question. However, it stands to reason that the greater the violation of the constraint, the less certain the outcome. When constraints are satisfied perfectly, the optimized result will match the reference curve exactly. When constraints are violated, the optimized result will likely lie somewhere between the base case curve and the

reference curve, but the exact result will be difficult to predict. Admittedly, knowing the limitations a priori is not a trivial task, suggesting an area for future work.

Table H.1: Performance metrics for Houston feeder zero demand optimization, 70% participation.

	Average	Min	Max
Electric Consumption [MWh]	5.51	4.72	6.13
Peak Demand [MW]	-0.71	-1.41	-0.34
Peak to Valley [%]	76.78	69.24	86.09
Load Factor [%]	7.11	5.02	10.48
Ramp [MW]	0.24	-3.87	2.65

Table H.2: Performance metrics for Houston feeder zero demand optimization, high wind penetration case, 70% participation.

	Average	Min	Max
Electric Consumption [MWh]	5.51	4.68	5.97
Peak Demand [MW]	-0.62	-1.26	-0.11
Peak to Valley [%]	67.70	47.82	85.49
Load Factor [%]	6.86	4.53	11.33
Ramp [MW]	0.34	-4.02	2.99




Figure H.1: Total spectral power as a function of frequency bin for Houston feeder zero demand optimization, 70% participation.




Figure H.2: Total spectral power as a function of frequency bin for Houston feeder zero demand optimization, high wind penetration case, 70% participation.

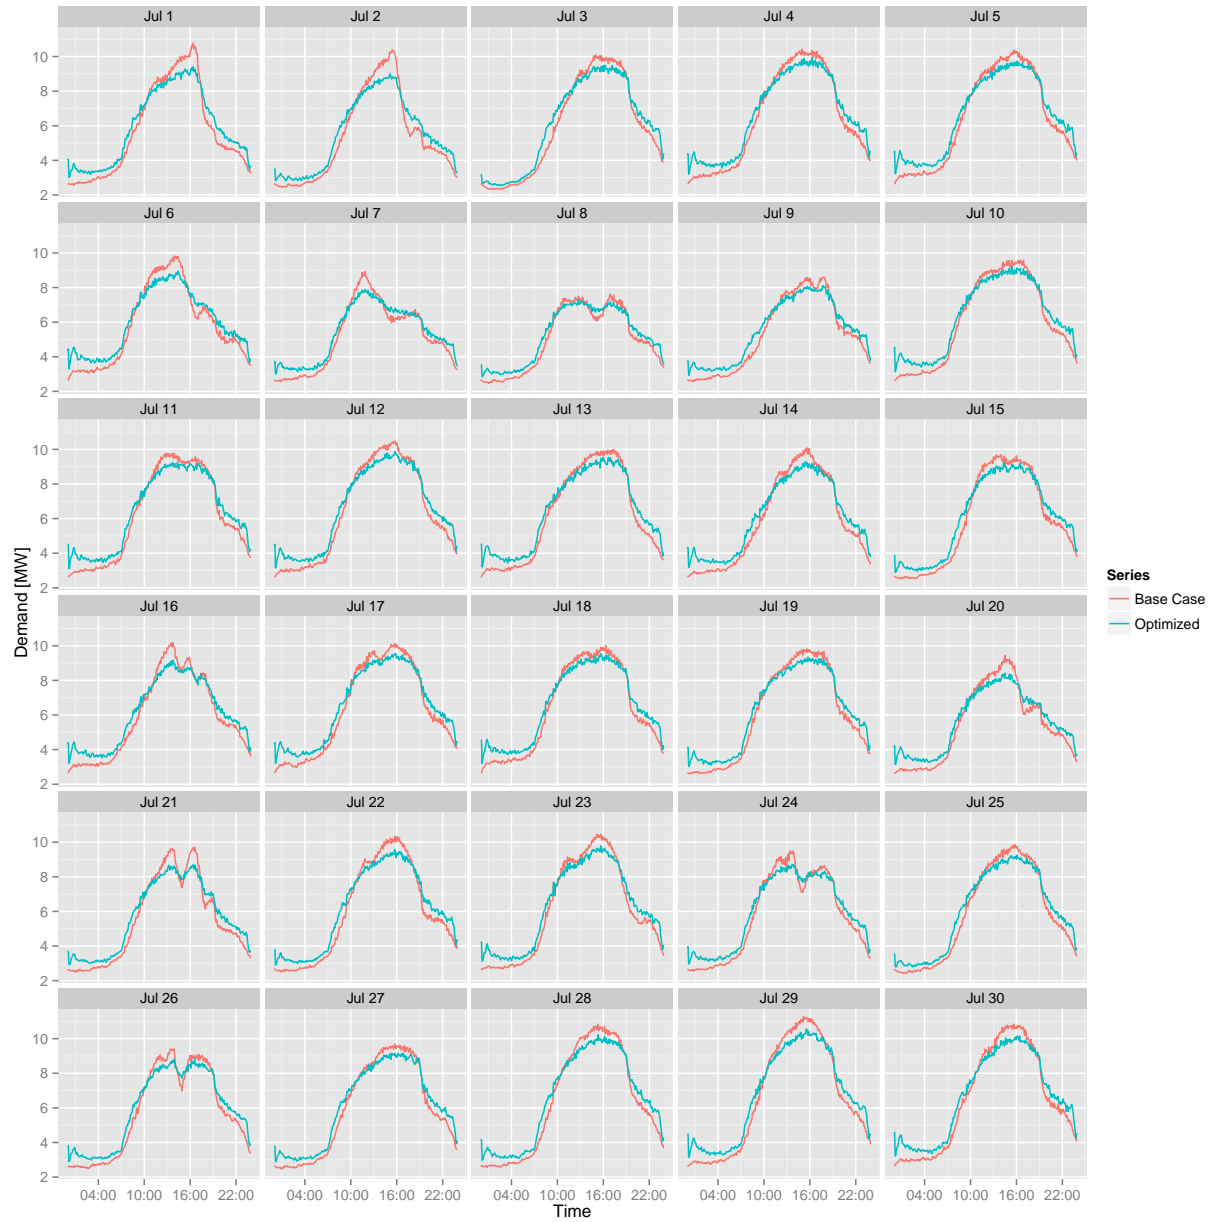


Figure H.3: Feeder demand profiles for Houston zero demand optimization, 70% participation.

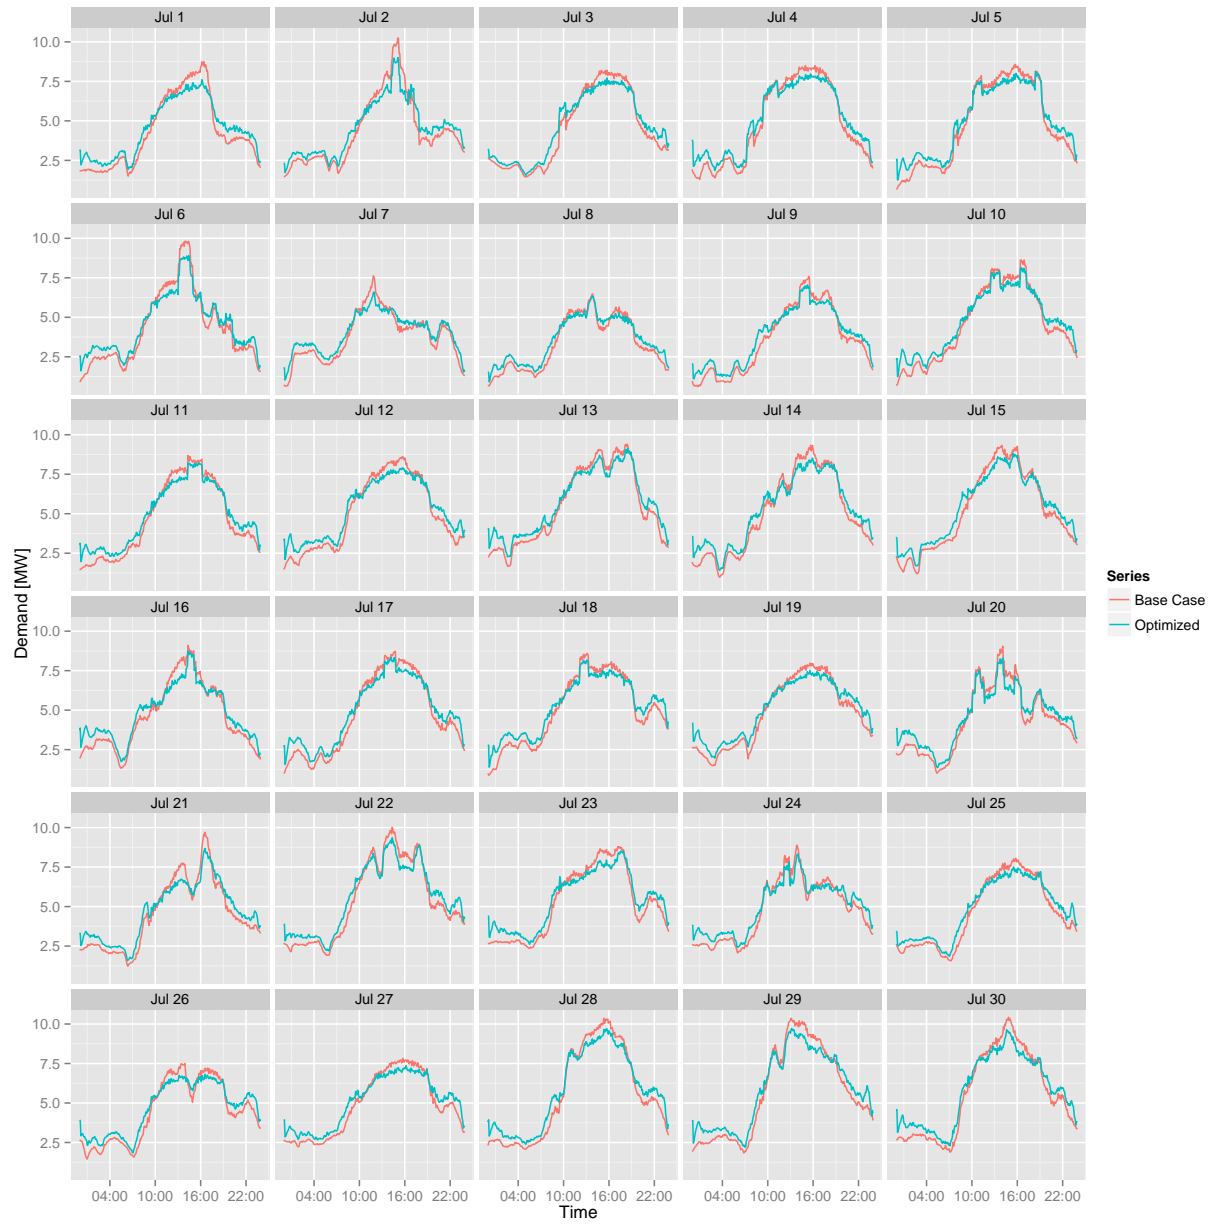


Figure H.4: Feeder demand profiles for Houston zero demand optimization, high wind penetration case, 70% participation.