



Final Report: TIP-337 Home Battery System for Cybersecure Energy Efficiency and Demand Response

Dane Christensen,¹ Xin Jin,¹ Bethany Sparn,¹ Steve Isley,² Sivasathya Pradha Balamurugan,¹ Scott Carmichael,³ Andrew Michalski,¹ Anuj Sanghvi,¹ Maurice Martin,¹ Kyri Baker,¹ Kaitlyn Garifi,¹ William Gillies,¹ Scott Averitt,⁴ Erdenebat Gantumur,⁵ Brandon Mendrick,⁵ Siddharth Suryanarayanan,⁶ Patricia Aloise-Young,⁶ Rahul Kadavil,⁶ and Salvador Lurbe⁶

1 National Renewable Energy Laboratory

2 Amazon

3 Blockchain

4 Robert Bosch North America

5 ESCRYPT

6 Colorado State University

**NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC**

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Technical Report
NREL/TP-5500-72184
November 2018



Final Report: TIP-337 Home Battery System for Cybersecure Energy Efficiency and Demand Response

Dane Christensen,¹ Xin Jin,¹ Bethany Sparn,¹ Steve Isley,² Sivasathya Pradha Balamurugan,¹ Scott Carmichael,³ Andrew Michalski,¹ Anuj Sanghvi,¹ Maurice Martin,¹ Kyri Baker,¹ Kaitlyn Garifi,¹ William Gillies,¹ Scott Averitt,⁴ Erdenebat Gantumur,⁵ Brandon Mendrick,⁵ Siddharth Suryanarayanan,⁶ Patricia Aloise-Young,⁶ Rahul Kadavil,⁶ and Salvador Lurbe⁶

1 National Renewable Energy Laboratory

2 Amazon

3 Blockchain

4 Robert Bosch North America

5 ESCRYPT

6 Colorado State University

Suggested Citation

Christensen, Dane, Xin Jin, Bethany Sparn, Steve Isley, et al. 2018. *Final Report: TIP-337 Home Battery System for Cybersecure Energy Efficiency and Demand Response*. Golden, CO: National Renewable Energy Laboratory. NREL/TP-5500-72184. <https://www.nrel.gov/docs/fy19osti/72184.pdf>.

**NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC**

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Technical Report
NREL/TP-5500-72184
November 2018

National Renewable Energy Laboratory
15013 Denver West Parkway
Golden, CO 80401
303-275-3000 • www.nrel.gov

NOTICE

This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Buildings Technologies Office and The Bonneville Power Administration. The views expressed herein do not necessarily represent the views of the DOE or the U.S. Government.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free via www.OSTI.gov.

Cover Photos by Dennis Schroeder: (clockwise, left to right) NREL 51934, NREL 45897, NREL 42160, NREL 45891, NREL 48097, NREL 46526.

NREL prints on paper that contains recycled content.

Executive Summary

Many challenges related to energy use and grid participation face the residential building sector and utilities that serve our homes today:

- The opportunity is great...
 - Residential electricity consumption is larger than any other sector
 - Residential buildings drive utility peak load.
- But the challenge is significant
 - Residential comprises 96% of customers, so bulk energy shifts are hard to achieve
 - Homes, homeowner schedules, economics, and preferences are highly diverse
 - Residential energy decisions are less financially driven
 - Emerging connectedness creates new cybersecurity risks.

A solution is needed to reduce energy consumption, increase grid service participation, and improve homeowner benefits. That solution must be able to adapt itself to each home and homeowner/occupant, deliver both building and grid services with reliability and high availability, automate these operations to minimize cost and complexity of deployment, and provide both home data privacy and grid cybersecurity.

We hypothesize that customer-oriented home automation can mutually satisfy home occupant/owner needs, reduce energy consumption, and deliver reliable grid services. This project seeks to develop innovative technology solutions that prove this hypothesis.

The Home Battery System (HBS) is a technology package comprised of connected “smart” appliances, rooftop solar photovoltaics, a home battery, and a coordinating smart controller. This system is envisioned, developed, and demonstrated by the National Renewable Energy Laboratory (NREL), Bosch, ESCRYPT, and Colorado State University (CSU). It is the result of three years of research by our diverse team. NREL developed the home automation controller, and performed simulation and laboratory evaluations of the HBS. Bosch developed and delivered most of the connected appliances used in the project and provided technical and commercialization guidance. ESCRYPT led cybersecurity analysis and developed the cybersecurity layer. CSU provided leadership on preference elicitation methodologies.



Figure ES.1. Conceptual diagram of the Home Battery System.

The HBS concept—automating flexible loads as an active energy-efficiency and demand response (DR) asset for customer and grid benefits—operated successfully in all laboratory scenario experiments. Excepting several clear opportunities where further development and platform maturing is needed, the application of multicriterion decision making results in significant homeowner savings, energy savings, and demand response service outcomes.

- **Energy Efficiency:** 0.4–8.1 kWh/day savings, when no DR called (4%–50% utility cost savings)
- **Demand Response:** 5.0–6.3 kWh over 4-hour DR period (demand response call to decrease load [i.e., increase generation] and demand response call to increase load [i.e., decrease generation])
- **Energy Export:** Hawaii cases all export significantly less with **foresee**. Washington/Oregon cases generally export more.
- **Multicriterion Decision Making:** Clear differences in execution, but room for improvement exists.
 - Cost savings preference did lead to lowest utility costs (and lower battery wear)
 - Negative comfort impacts are apparent in most cases, which motivates further model development to improve thermal predictiveness.

Most notably, energy was saved in every scenario and case where the HBS and **foresee** were able to optimize the home. An example of this outcome, from one of the laboratory scenarios is shown below. Cost was saved (compared to the Baseline) in every Energy-Efficiency case and in many DR cases—despite the fact that we did not include any utility incentive for DR in calculating the daily cost.

When a utility calls a DR event, the HBS is observed to sacrifice approximately 20% of **foresee**-delivered energy savings to provide reliable grid services—the capacity of which varies greatly based on the home, climate, and available appliance resources, but which was observed to be more than 6 kWh in most cases. The actual service level and availability are highly dependent on many factors, including equipment, weather, home construction, tariff, and preferences. Homeowner impact is based on their stated preferences, so in some cases there are positive or negative impacts on thermal comfort, convenience, or cost in order to maximize other objectives.

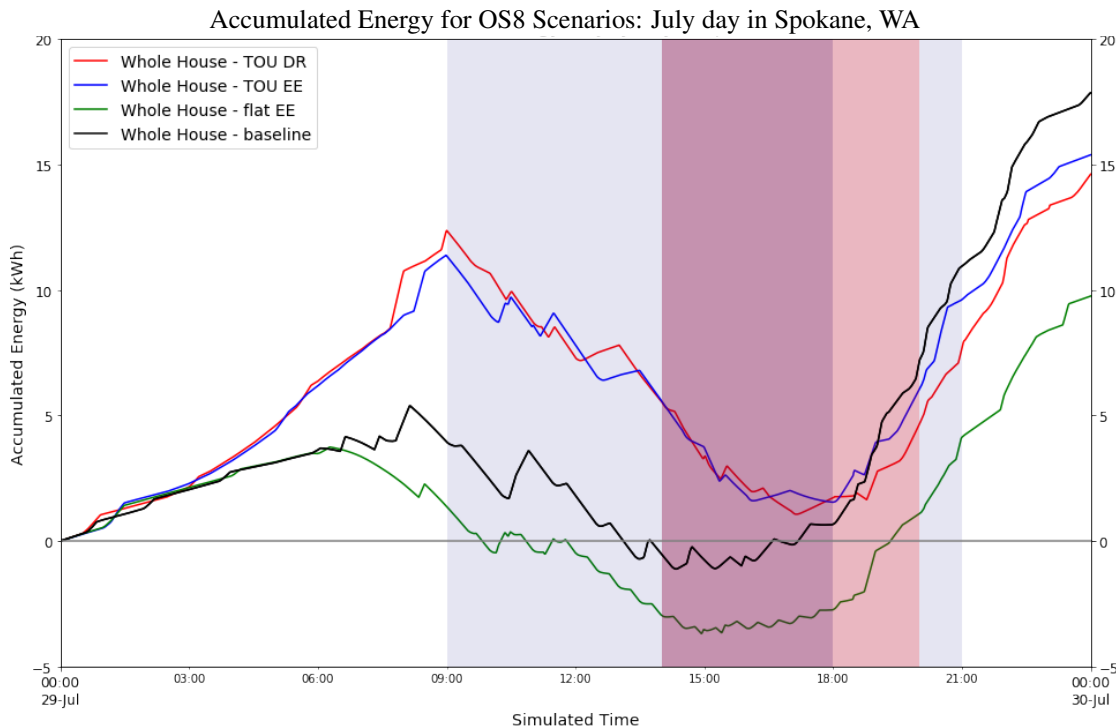


Figure ES.2. Accumulated energy for the OS8 baseline case, OS8 energy efficiency using a flat utility rate, OS8 energy efficiency with a TOU utility rate, and OS8 demand response with TOU pricing. TOU price schedule indicated by blue shading. Load-shed DR period is shaded red.

Acknowledgments

We are deeply grateful for technical contributions, support, and enabling prior and concurrent work done by so many colleagues and partners. It takes a village to raise a smart home! While not all of them worked directly on this project, we would like to acknowledge and thank NREL colleagues without whom this project could not have been completed. Those included Kim Adams, Katherine Avedon, Erin Beaumont, Daniel Beckley, Dylan Cutler, Emiliano Dall’Anese, Harrison Dreves, Lieko Earle, Jennifer Fetzer, Julieta Giraldez, Amy Glickson, Bryan Hannegan, Hilary Hatch Copeland, Caroline Heilbrun, Wayne Hicks, Andrew Hudgins, Erfan Ibrahim, Roderick Jackson, Jennifer Josey, Dheepak Krishnamurthy, Ben Kroposki, Chuck Kutscher, Nick Laws, Jeff Maguire, Killian McKenna, Noel Merket, Anne Miller, Andrew Paulonis, Lucas Phillips, Annabelle Pratt, Mollie Putzig, Emma Raszmann, Brent Rice, David Roberts, Jennifer Schofield, Jean Schulte, Ying Shi, Sandy Steele, Martha Symko-Davies, Kyle Tangler, Sarah Truitt, Linh Truong, Deepthi Vaidhynathan, and Mary Wheeler.

The close partnership with our technical collaborators at Bosch, ETAS/ESCRYPT, and Colorado State University was critical. We appreciate all the many colleagues who contributed to this project.

Funding and Support

Mostly we would like to thank the project sponsors: Bonneville Power Administration’s Technology Innovation group; U.S. Department of Energy, Energy Efficiency and Renewable Energy, Building Technologies Office; and Robert Bosch, N.A. for their incredible commitment to pursue this vision. It was fantastic to work with Judith Estep, Robert Fares, Ryan Fedie, Stephanie Johnson, Kevin Lynn, David Nemptzow, Kari Nordquist, Terry Oliver, Justin Reel, Karma Sawyer, Melanie Smith, Marina Sofos, Stephanie Vasquez, and all their teams.

This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. by Bonneville Power Administration under Technology Innovations Project TIP-337 and by U.S. Department of Energy (DOE) Office of Energy Efficiency and Renewable Energy Building Technologies Office, with substantial in-kind participation by Bosch. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

Acronym List

AHP	Analytic Hierarchy Process	INC	Demand Response Call to Decrease Load (i.e. <u>I</u> ncrease Generation)
AMT	Amazon Mechanical Turk	IQR	Interquartile Rank
API	Application Programming Interface	LAN	Local Area Network
BAU	Business as Usual	MCDM	Multi Criterion Decision Making
BIOS	Basic Input-Output System	MPC	Model Predictive Control
BPA	Bonneville Power Administration	MiM	Man in the Middle
CIP	Critical Infrastructure Protection	NERC	National Electricity Reliability Corporation
CSU	Colorado State University	NMAP	Network Mapper
CTA	Consumer Technology Association	NREL	National Renewable Energy Laboratory
DEC	Demand Response Call to Increase Load (i.e. <u>D</u> ecrease Generation)	NUC	Next Unit of Computing
DER	Distributed Energy Resource	OpenSSL	Open Secure Sockets Layer
DCM	Discrete Choice Model	PV	Photovoltaic
DLC	Direct Load Control	RBSA	Residential Building Stock Assessment
DOE	U.S. Department of Energy	RMSE	Root Mean Squared Error
DoS	Denial of Service	RTP	Real-Time Pricing
DR	Demand Response	SD	Standard Deviation
DSM	Demand-Side Management	SMARTER	Simple Multi-Attribute Rating Technique Exploiting Ranks
EE	Energy Efficiency	SOC	State of Charge
ESIF	Energy Systems Integration Facility	SPL	Systems Performance Laboratory
HB	Hierarchical Bayesian	SYN	Synchronization
HBS	Home Battery System	TCL	Thermostatically-Controlled Load
HEMS	Home Energy Management System	TCP	Transmission Control Protocol
HIL	Hardware In the Loop	TOU	Time of Use
HIT	Human Intelligence Task	UDP	User Datagram Protocol
HSM	Hardware Security Module	UEFI	Unified Extensible Firmware Interface
HTTP	Hypertext Transfer Protocol	USB	Universal Serial Bus
HULK	HTTP Unbearable Load King	UTM	Unified Threat Management
HVAC	Heating, Ventilation and Air Conditioning	WAN	Wide Area Network
ICMP	Internet Control Message Protocol		

Table of Contents

1	Introduction	1
1.1	Research Overview and Objectives	1
1.2	Description of the Home Battery System	2
1.3	foresee Overview and Architecture	3
2	User Preferences	5
2.1	Preference Elicitation Methodologies	5
2.2	Preference Elicitation Survey Development	5
2.2.1	Discrete Choice Modeling	6
2.2.2	SMARTER	7
2.2.3	Analytic Hierarchy Process	8
2.3	Preference Elicitation Survey Results	9
2.3.1	General Comfort Preferences	9
2.3.2	Discrete Choice	10
2.3.3	SMARTER	10
2.3.4	Analytic Hierarchy Process	11
2.4	Method Comparison	12
2.5	User Interface for Preference Initialization	12
3	Machine Learning	16
3.1	Statistical Learning	16
3.2	System Identification	16
3.3	Residential Equipment Models	17
3.3.1	House-HVAC Model	17
3.3.2	Water Heater Model	18
3.3.3	Schedulable Appliance Models	19
3.3.4	PV Model	19
3.3.5	Battery Model	19
3.3.6	House Power Balance	20
3.4	Application to Field Data	20
3.4.1	Estimating Usage Patterns Using Statistical Learning	21
3.4.2	Creating House Models Using System Identification	21
3.5	Comments on Machine Learning	22
4	Optimization	23
4.1	Objective Function	23
4.1.1	Thermal Discomfort	23
4.1.2	Energy Cost	24
4.1.3	Environmental Impacts	24
4.1.4	User Inconvenience	24
4.1.5	Equipment Degradation	25
4.2	Solving the MPC Problem via Quadratic Programming	25
4.3	Chance-Constrained Optimization	25
5	Scenario Experiment Design	27
5.1	Introduction	27
5.2	Experimental Setting	27

5.3	Operational Scenarios	29
5.4	Grid Scenarios	30
5.4.1	GS1: Normal Day	30
5.4.2	GS2: Peak Heat Day	30
5.4.3	GS3: Peak Cold Day	30
5.4.4	GS4: Shoulder Season Oversupply	31
5.4.5	GS7: Solar Self-Consumption	31
5.4.6	GS8: Emergency Loss of Supply	31
5.5	Residence Scenarios	31
5.6	Homeowner Scenarios	31
5.7	Preference Scenarios	32
5.8	Operational Scenario Performance Metrics	32
6	Energy, Cost and Comfort Results	34
6.1	Laboratory Testing Procedure	34
6.1.1	Experiment Initialization	34
6.1.2	Data Analysis	34
6.2	Metric Evaluation	35
6.2.1	Cost Analysis	35
6.2.2	Comfort Analysis	35
6.2.3	Distributed Energy Resource Analysis	36
6.3	Laboratory Experiment Results	36
6.3.1	OS1 Scenario Results - Portland, OR, with Three Different User Profiles	38
6.3.2	OS5 Scenario Results - Honolulu, HI, Under Shoulder Season Oversupply (Load-Add DR)	41
6.3.3	OS8 Scenario Results - Spokane, WA, Under Peak Hot Day (Load-Shed DR)	44
6.4	Simulation Experiment Results	48
6.4.1	OS2 Scenario - Portland, OR, Under Peak Cold Day (Load-Shed DR)	48
6.4.2	OS3 Scenario - Portland, OR, Under Shoulder Season Oversupply (Load-Add DR)	49
6.4.3	OS9 Scenario - Spokane, WA, Under Peak Cold Day (Load-Shed DR)	50
7	Cybersecurity Risk Analysis & Implementation	53
7.1	Cybersecurity Risk Analysis	53
7.2	Recommendation to Mitigate Risks	54
7.2.1	Physical Security	54
7.2.2	Software Security	55
7.3	CIP Compliance	55
7.4	Cybersecurity Layer Implementation	56
8	Cybersecurity Penetration Testing	57
8.1	Objective	57
8.2	Scope	57
8.3	Tests Performed	57
8.4	Penetration Test Results and Recommendations	59
9	Conclusions	60
9.1	Major Findings and Outcomes	60
9.2	Future Work	60
9.2.1	Comfort Delivery Improvements	60
9.2.2	Preferences Improvements	61
9.2.3	Stochastic Optimization and Service Commitment	61
9.2.4	Other Use Cases	62
9.2.5	Impact Analysis	62
Appendix A	Scenario Experiment Details	63
A.1	Grid Scenarios	63

A.2	Climate/Residence Scenarios	65
A.3	Homeowner Scenarios	66
A.3.1	Water Heater Draw Profile	66
A.3.2	Washer, Dryer, and Dishwasher	67
A.3.3	Preferences	68
A.4	Laboratory Procedure	68
A.4.1	Experiment Initialization	69
A.4.2	Data Analysis	69
Appendix B Detailed Laboratory Performance Results		70
B.1	OS1 Scenarios - Portland, OR with three different user profiles	70
B.2	OS2 Scenario - Portland, OR under peak cold day (load shed DR)	73
B.3	OS3 Scenario - Portland, OR, Under Shoulder Season Oversupply (Load-Add DR)	74
B.4	OS4 Scenario - Honolulu, HI, Under Normal Day Operation	75
B.5	OS5 Scenario - Honolulu, HI, Under Shoulder Season Oversupply (Load-Add DR)	79
B.6	OS6 Scenario - Honolulu, HI, Under Solar Self-Consumption	82
B.7	OS7 Scenario - Spokane, WA, Under Normal Operation	85
B.8	OS8 Scenario - Spokane, WA, Under Peak Hot Day (Load-Shed DR)	88
B.9	OS9 Scenario - Spokane, WA, Under Peak Cold Day (Load-Shed DR)	91
B.10	Summary of Results	93
B.10.1	Energy Efficiency Results	93
B.10.2	Demand Response Results	93
B.10.3	DR Service Uncertainty	94
Appendix C Detailed Penetration Test Results		96
C.1	Overview	96
C.2	Objective	96
C.3	Scope	96
C.4	T1 - Foresee Service Disruption via Internal DoS	98
C.4.1	Penetration Testing Tools	98
C.4.2	Results	98
C.4.2.1	ICMP and TCP SYN floods	98
C.4.2.2	HTTP Request Flooding	99
C.4.2.3	Connection Exhaustion	101
C.4.3	T1 Conclusions	102
C.5	T2 - Platform Operation Disruption (via Customer LAN)	102
C.5.1	Penetration Testing Tools	102
C.5.2	Results	102
C.5.3	T2 Conclusions	103
C.6	T3 - Attack foresee Platform from Customer LAN	103
C.6.1	Penetration Testing Tools	104
C.6.2	Results	104
C.6.2.1	Port Scanning	104
C.6.2.2	Vulnerability Scanning	104
C.6.3	T3 Conclusions	105
C.7	T4 - foresee Platform Data Disruption	106
C.7.1	Penetration Testing Tools	106
C.7.2	Results	106
C.7.3	T4 Conclusions	106
C.8	T5 - Attack Communication Channel Outside the UTM	107
C.8.1	Penetration Testing Tools	107
C.8.2	Results	107
C.8.3	T5 Conclusions	107
C.9	T6 - Attack on VOLTTRON	108
C.9.1	Penetration Testing Tools	108

C.9.2 T6 Conclusions	108
C.10 Cybersecurity Recommendations	109

List of Figures

Figure ES.1. Conceptual diagram of the Home Battery System.	iii
Figure ES.2. Accumulated energy for the OS8 baseline case	iv
Figure 1. Conceptual diagram of the Home Battery System	2
Figure 2. Block diagram of foresee software architecture	3
Figure 3. Analytic Hierarchy Process model for Home Energy Management System	9
Figure 4. Discomfort coefficients obtained using the DCM method	10
Figure 5. The distribution of ranks given to the home comfort attributes in the SMARTER method.	11
Figure 6. Intracriterion comparison of the SMARTER method survey ranks.	11
Figure 7. Initialization screen where users input Wi-Fi and location information.	13
Figure 8. Initialization screen where users input air temperature preferences.	13
Figure 9. Initialization screen where users input hot water (shower) preferences.	14
Figure 10. Initialization screen where users rank their service preferences.	14
Figure 11. Device connection screen.	15
Figure 12. Main user interface screen for foresee	15
Figure 13. Comparison of the estimated and actual uncontrollable loads.	21
Figure 14. Model validation of machine-learned thermal house model	22
Figure 15. Appliances used in the laboratory experiments	28
Figure 16. OS1 baseline power consumption	38
Figure 17. OS1 accumulated energy	39
Figure 18. OS1 battery power	39
Figure 19. OS1 accumulated energy	40
Figure 20. OS5 baseline power consumption	41
Figure 21. OS5 accumulated energy	42
Figure 22. OS5 battery power	43
Figure 23. OS5 accumulated discomfort	43
Figure 24. OS8 baseline power consumption	44
Figure 25. OS8 accumulated energy	46
Figure 26. OS8 battery power	47
Figure 27. OS8 accumulated discomfort	47
Figure 28. Accumulated energy from OS2 scenarios	49
Figure 29. Accumulated energy from OS3 scenarios	50
Figure 30. Accumulated energy from OS9 scenarios	51
Figure 31. Flowchart showing high-level steps in the Cybersecurity Risk Analysis process.	53

Figure 32. Flowchart showing the penetration test process.	57
Figure A.1. Appliance Hourly Probability Profiles	68
Figure B.1. OS1 Baseline experiment power	70
Figure B.2. Accumulated energy from OS1 baseline	71
Figure B.3. Battery power consumption for OS1 scenarios	72
Figure B.4. Accumulated Discomfort for OS1 Scenarios	72
Figure B.5. Accumulated energy from the OS2 simulation for baseline, energy efficiency and demand response cases. TOU price schedule indicated by blue shading and red shading indicates load-shed DR period.	74
Figure B.6. Accumulated energy from the OS3 simulation for baseline, energy efficiency and demand response cases. TOU price schedule indicated by blue shading and green shading indicates load-add DR period.	75
Figure B.7. OS4 Baseline test with all loads shown. TOU price schedule indicated by blue shading.	76
Figure B.8. Accumulated air temperature discomfort from the OS4 baseline and experimental cases. TOU price schedule indicated by blue shading.	77
Figure B.9. Battery power consumption for OS4 optimization cases. Battery charging is considered positive power and discharging is considered negative. TOU price schedule indicated by blue shading.	78
Figure B.10. Accumulated energy from the OS4 baseline and experimental case. TOU price schedule indicated by blue shading.	78
Figure B.11. OS5 Baseline test with all loads shown. TOU price schedule indicated by blue shading.	79
Figure B.12. Accumulated energy for the OS5 baseline, OS5 energy efficiency, and OS5 demand response cases. TOU price schedule indicated by blue shading. Load-add DR period shaded green.	80
Figure B.13. Battery power consumption for OS5 optimization cases. Battery charging is considered positive power and discharging is considered negative. TOU price schedule indicated by blue shading, with DR period indicated with the green shading.	81
Figure B.14. Accumulated air temperature discomfort for the OS5 baseline, the energy-efficiency case, and the load-up demand response case. TOU price schedule indicated by blue shading, with DR period indicated with the green shading.	81
Figure B.15. Baseline Power from OS6 Scenarios	82
Figure B.16. Accumulated energy for the OS6 baseline and OS6 energy-efficiency cases, with the results from the Solar Self-Consumption case coming from simulation. TOU price schedule indicated by blue shading.	83
Figure B.17. Battery power consumption for OS6 optimization cases, with the results from the Solar Self-Consumption case coming from simulation. Battery charging is considered positive power and discharging is considered negative. TOU price schedule indicated by blue shading.	84
Figure B.18. Accumulated air temperature discomfort for the OS6 baseline and the energy-efficiency case. TOU price schedule indicated by blue shading.	84
Figure B.19. OS7 Baseline test with all loads shown. TOU price schedule indicated by blue shading.	85
Figure B.20. Accumulated energy for the OS7 baseline and OS7 energy-efficiency cases. TOU price schedule indicated by blue shading.	86
Figure B.21. Battery power consumption for OS7 optimization cases. Battery charging is considered positive power and discharging is considered negative. TOU price schedule indicated by blue shading.	87
Figure B.22. Accumulated air temperature discomfort for the OS7 baseline and the energy-efficiency case. TOU price schedule indicated by blue shading.	87
Figure B.23. OS8 Baseline test with all loads shown. TOU price schedule indicated by blue shading.	88

Figure B.24. Accumulated energy for the OS8 baseline case	89
Figure B.25. Battery power consumption for OS8 optimization cases. Battery charging is considered positive power and discharging is considered negative. TOU price schedule indicated by blue shading, with DR period indicated with the red shading.	90
Figure B.26. Accumulated air temperature discomfort for the OS8 baseline, the energy efficiency with flat rate case, energy efficiency with TOU rates, and the load-shed demand response case. TOU price schedule indicated by blue shading, with DR period indicated with the red shading.	90
Figure B.27. Accumulated energy from the OS9 simulation for baseline, energy efficiency and demand response cases. TOU price schedule indicated by blue shading and red shading indicates load-shed DR period.	92
Figure B.28. Average load reduction during the DR period relative to the EE case for OS9 Scenario Monte Carlo simulations.	95
Figure B.29. Histogram of the average load reduction relative to the EE case during the DR period for the OS9 Scenario Monte Carlo Simulations.	95
Figure C.1. Flowchart showing the penetration test process.	96
Figure C.2. Screenshot showing the test client’s attempt to access the foresee appliance. The left side is the result of a continuous ping showing that the client experienced 54% packet loss over 4 minutes. The right side shows the browser window where the foresee appliance failed to load in its entirety.	99
Figure C.3. Screenshot showing the test client’s attempt to access the foresee appliance. The left side is the result of a continuous ping showing that the client experienced 45% packet loss over 6 minutes. The right side shows the browser window where the foresee user interface failed to load.	100
Figure C.4. Screenshot showing the test client’s attempt to access the foresee appliance. The left side is the result of a continuous ping showing that the client experienced 35% packet loss over 3 minutes. The right side shows the browser window where the foresee appliance failed to load.	100
Figure C.5. Screenshot showing the penetration tester’s laptop as it uses the PyLoris graphical interface to launch an attack against the foresee application.	101
Figure C.6. Screenshot showing the test client’s attempt to access the foresee appliance. The left side is the result of a continuous ping showing that the client experienced 32% packet loss over 1 minute. The right side shows the browser window where the foresee appliance loaded successfully.	101
Figure C.7. Screenshot showing packet capture of communication between the foresee application and an IoT device. A secure connection is established and a secure encryption protocol is in use. The collected communication traffic is obfuscated preventing a penetration tester from spoofing communications to the foresee appliance.	103
Figure C.8. Screenshot from the OpenVAS report alerts the penetration tester that a directory traversal vulnerability exists on the web server hosting the foresee appliance.	105
Figure C.9. Screenshot taken from the foresee appliance itself. The screenshot shows that wireless networking has been disabled from the attack, causing it to continuously attempt to reconnect.	108

List of Tables

Table 1.	The rank-based inferred weights with w_1 being the most important attribute	8
Table 2.	Home air temperature preferences (mean, with interquartile range in parentheses) (°F)	9
Table 3.	Shower preferences in minutes by gender (IQR = interquartile range)	10
Table 4.	Comparison of Preference Elicitation Methods	12
Table 5.	Summary of five homes from the residential building stock assessment data set used in the study	21
Table 6.	System identification results in five RBSA homes	22
Table 7.	Operational Scenario Definitions	29
Table 8.	Homeowner Scenario Definition	32
Table 9.	Homeowner Preference Definition	32
Table 10.	Performance Summary, EE Mode	37
Table 11.	Summary of Demand Response Mode Energy and Comfort Performance	37
Table 12.	Summary from OS1 Scenarios	39
Table 13.	OS1 DER and Comfort Impacts	40
Table 14.	Summary from OS5 Scenarios	41
Table 15.	OS5 DER and Comfort Impacts	44
Table 16.	OS8 Scenario Summary	45
Table 17.	OS8 DER and Comfort Impacts	45
Table 18.	Summary from OS2 Scenarios - vs. Baseline	48
Table 19.	DER and Comfort Impacts from OS2 Scenarios	48
Table 20.	Summary from OS3 Scenarios	50
Table 21.	DER and Comfort Impacts from OS3 Scenarios	50
Table 22.	Summary from OS9 Scenarios	51
Table 23.	DER and Comfort Impacts from OS9 Scenarios	52
Table 24.	Summary of Calculated Risks	53
Table 25.	Cybersecurity Recommendations	54
Table 26.	Penetration Tests Performed	58
Table A.1.	Operational Scenario Definitions	63
Table A.2.	Time-of-Use Utility Rates for each Location	64
Table A.3.	Summary of Grid Scenario details	65
Table A.4.	Climate Scenarios and Days for testing	66
Table A.5.	Description of EnergyPlus models used in scenario experiments	66
Table A.6.	Homeowner Description	66
Table A.7.	Hot water draw profiles for different user profiles	67

Table A.8. Large appliance schedules for different user profiles	68
Table B.1. Summary from OS1 Scenarios	71
Table B.2. DER and Comfort Impacts from OS1 Scenarios	73
Table B.3. Summary from OS2 Scenarios	73
Table B.4. DER and Comfort Impacts from OS2 Scenarios	74
Table B.5. Summary from OS3 Scenarios	75
Table B.6. DER and Comfort Impacts from OS3 Scenarios	76
Table B.7. Summary from OS4 Scenarios	76
Table B.8. DER and Comfort Impacts from OS4 Scenarios	77
Table B.9. Summary from OS5 Scenarios	79
Table B.10. DER and Comfort Impacts from OS5 Scenarios	82
Table B.11. Summary from OS6 Scenarios	83
Table B.12. DER and Comfort Impacts from OS6 Scenarios	85
Table B.13. Summary from OS7 Scenarios	86
Table B.14. DER and Comfort Impacts from OS7 Scenarios	86
Table B.15. Summary from OS8 Scenarios	88
Table B.16. DER and Comfort Impacts from OS8 Scenarios	91
Table B.17. Summary from OS9 Scenarios	91
Table B.18. DER and Comfort Impacts from OS9 Scenarios	92
Table B.19. Summary of Energy-Efficiency Mode Energy and Comfort Performance	93
Table B.20. Summary of Demand Response Mode Energy and Comfort Performance	94
Table C.1. Penetration Tests Performed	97
Table C.8. Cybersecurity Recommendations	109

1 Introduction

Residential buildings present a significant opportunity for improving system-wide energy efficiency and grid operations. In aggregate, homes are both the largest energy-consuming sector and the driving force behind most utilities' peak loads. We understand that residential loads are particularly dominant in Bonneville Power Administration's (BPA's) wintertime morning peak, as a result of coincident electric heating (driven by thermostats setting back up), electric water heating, and lighting (driven by customers waking up early while it's still dark outside). Despite this significant opportunity, it is very challenging to make gains in this sector of electricity customers. Homes comprise more than 90% of a utility's customers. Utilities often approach the residential demand response (DR) problem solely by providing financial incentives. Pricing mechanisms include direct load control (DLC) programs, rebates, and various dynamic pricing schedules, including time-of-use (TOU) and real-time pricing (RTP). Additional demand-side management (DSM) programs usually entail pamphlets distributed with information on recommended behaviors for energy savings. These techniques often produce fewer savings than anticipated and only for a limited duration. To bridge the gap between goals and results, experts in fields outside the typical domain of power system engineering—specifically those in social psychology—can provide valuable insights. That field has proven exceptionally capable of motivating lasting behavior change in other domains, specifically human health. By using similar techniques applied to consumer energy behaviors, significant results can be achieved with residential DR.

Aggregation of distributed energy resources (DERs) promises some relief, but has yet to achieve significant adoption. Specifically, a number of adoption barriers exist—cost of implementation, perception of utilities circumventing homeowner preferences such as comfort, operating cost and technical hurdles in managing many small distributed loads, data privacy, and cybersecurity, to name a few.

Many challenges exist in traditional DR technologies. For utilities, homes are not able to reliably provide dispatchable DR because the DR resources in a home are not always available. Utilities are unable to predict the effectiveness of a DR command as a result of lack of information about the availability of DR resources in homes. Homeowners are not able to specify preferences and their comfort may not always be guaranteed during DR events. These barriers undermine the effectiveness of DR programs and result in high opt-out rate.

We hypothesize that customer-oriented home automation can mutually satisfy home occupant/owner needs, reduce energy consumption, and deliver reliable grid services. This project seeks to identify innovative technology solutions that prove this hypothesis.

1.1 Research Overview and Objectives

The Home Battery System (HBS) project sought to develop and demonstrate a novel technology package that can overcome these adoption barriers and achieve energy savings as well as provide highly flexible demand-side management including DR. The HBS and the integrated advanced control technologies achieved the following technical objectives:

- Guaranteed comfort and improved energy savings for home owners
- Delivery of highly available (more than 90%) and reliable DR capacity from individual homes
- Reliable DR capacity (more than 2 kW/home) prediction from individual homes across multiple look-ahead timeframes (day/hour/minute)
- Optimal scheduling of home appliances based on user preferences and DR requests
- Cybersecure DR delivered by critical infrastructure protection (CIP)-compliant systems.

This report documents the development and laboratory demonstration of this novel technology package. The HBS is comprised of energy-efficient communicating appliances, energy storage, a photovoltaic (PV) inverter, and a cybersecure home energy management system (HEMS) with innovative self-learning and predictive controls. The HBS is designed to address several significant energy efficiency (EE) and DR opportunities, including: a) HEMS system with integrated homeowner preference and comfort guarantee using survey-based input and a multicriteria decision making (MCDM) algorithm; b) delivery of more than 2 kW dispatchable DR per home (both demand response call to decrease load [i.e., increase generation] (INC) and demand response call to increase load [i.e.,

decrease generation] (DEC)), with high (more than 90%) availability and reliability; c) reliable resource predictions across multiple look-ahead timeframes; and d) intelligent control of connected appliances including load shifting and scheduling based on user preferences and DR requests. We also designed the HBS to enable CIP-compliance (CIP-002 through CIP-0014) when aggregated to provide bulk grid services.

1.2 Description of the Home Battery System

The HBS is comprised of hardware and software that are expected to exist in a future “smart home.” The central innovation of this project is the coordinating home automation controller, particularly its decision-making engine. We call this decision engine **foresee**TM. **foresee** is developed as a standalone application which uses the VOLTRONTM platform (Akyol et al. 2012). An overview of **foresee**’s architecture is provided in the next section.

foresee operates the home with guidance from homeowners and occupants on their preferences, which may be complex. In Chapter 2, we document our novel work to develop a methodology to elicit those preferences at setup and inform the relative weighting of multiple connected criteria. The **foresee** platform is flexible to include additional criteria, but the multicriteria objectives are currently:

- Thermal Comfort (air temperature)
- Shower Length (hot water temperature and volume)
- Money (utility bill and operating costs)
- Convenience (having dishes and laundry complete at the earliest possible time)
- Environmental Impact (externalities resulting from the choice among generation sources).

foresee uses machine learning methods to customize itself to a home’s occupant patterns and needs, as well as the physical building parameters such as air conditioning energy and wall insulation parameters—this is described in Chapter 3.

foresee performs optimization to select control actions for the home based on how they will change the state of the home and therefore impact the home occupants. We have implemented a deterministic optimization routine, but are also exploring chance constrained optimization as a future improvement. These are discussed in Chapter 4.

We exercised these energy management capabilities in a laboratory by connecting the **foresee** controller to physical appliances. For this project, the residential appliances and equipment included:



Figure 1. Conceptual diagram of the Home Battery System

- Bosch™ Heat Pump, operated by an ecobee4™ Wi-Fi connected Thermostat
- Bosch HomeConnect-enabled Dishwasher, Refrigerator, Washing Machine, and Electric Dryer
- AO Smith™ CTA-2045-enabled Water Heater
- Eguana™ Residential Battery
- SMA™ Photovoltaic Inverter.

These are shown schematically in Figure 1. The experimental scenarios to which this system was exposed are defined in Chapter 5, and results are provided in Chapter 6.

A cybersecurity layer was developed to protect both the HBS from external intrusion and the homeowner’s data privacy and security. The purpose of the cybersecurity layer is to further support power grid security because a home is not subject to bulk-grid cybersecurity standards but when homes or home appliances are aggregated to provide a grid service, that service must be secure and reliable. As more connected, or “smart,” devices participate in grid operations, this need will grow. We performed a cybersecurity risk assessment to identify potential vulnerabilities and their system impacts, and then designed an implementation plan to address those risks by integrating a cybersecurity layer with **foresee** as summarized in Chapter 7.

An independent team performed penetration testing to confirm that high-risk vulnerabilities had been successfully mitigated by the cybersecurity layer; these experiments and findings are summarized in Chapter 8.

Finally, we provide conclusions and recommendations for future work in Chapter 9.

1.3 foresee Overview and Architecture

Like all VOLTRON applications, **foresee** is developed as a modular, agent-based system. A high-level block diagram of its functional components is shown in Figure 2.

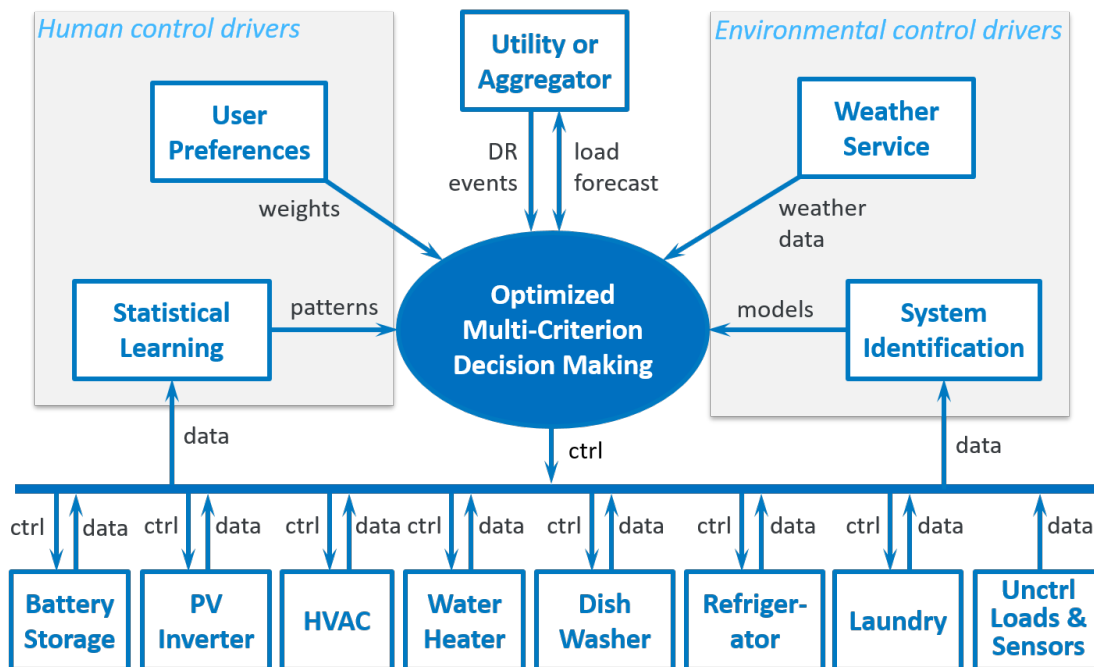


Figure 2. Block diagram of foresee software architecture

Devices in the home, including sensors and other devices that **foresee** is not controlling but which may affect its decisions and outcomes, are shown across the bottom of the figure. These devices provide data that can inform **foresee**’s planning. Those data are used in two ways.

foresee uses Statistical Learning to observe and learn occupants' patterns and needs. For example, when are showers taken and how long are they? This helps **foresee** ensure comfort by planning for sufficient hot water resource at the time it is needed. **foresee** also uses System Identification to establish predictive models of the house and its equipment. For example, what is the wall R-value, and how fast can the air conditioner cool the home down on a hot summer day? This helps **foresee** plan when to run the air conditioner to provide comfort for the occupants by the time they come back home, but reduce energy costs while they are away.

User preferences are interpreted as weighting factors which establish a balance between the potential outcomes (described previously) of implementing different control actions. For example, energy cost savings can result from increasing the air conditioner setpoint, but this may lead to occupant discomfort. How much does that occupant prefer air temperature comfort vs. monetary savings? And how do each of those compare to their preference for environmental benefits or having the dishwasher finished when they get home?

foresee also draws on public web data such as weather services to understand the future need for (for example) air conditioning, and energy cost of operating it. Using all these data and models of the building and occupants, it performs receding-horizon model predictive control (MPC). "Receding Horizon" describes its planning method. At a given time, it will identify an optimal set of control actions at time increments into the future. The first set of those control actions is implemented, after which **foresee** will perform another optimization over the same future planning timeframe. In this way, it is always planning the same distance into the future. MPC is a methodology for using models of a system to predict its future state based on some inputs—in this case our control decisions.

When the local utility, or another entity such as an aggregator, calls a DR event, **foresee** accounts for the offered incentive and forecasts its updated load profile. This establishes a DR service forecast by comparing the new energy use plan to the prior plan (absent the DR event). It then operates the home to achieve that DR service. By planning ahead, it should be able to preheat the water tank, cool the house, and run other appliances in a way that increases DR service and also ensures better delivery of occupant preferences, compared to today's business-as-usual practices.

A more complete description of User Preference, Machine Learning, and Optimization modules are provided in the next three chapters.

2 User Preferences

Considerable research has gone into understanding how to take advantage of connected appliances to enable participation in DR events and improve overall efficiency. However, in order to reach the market share necessary to achieve substantial system benefits, smart homes will need to provide benefits to the homeowner and not just the utility (Balta-Ozkan et al. 2013; Zipperer et al. 2013). Limited research has been conducted to understand how individuals value different aspects of home comfort and the trade-offs they are willing to make between monetary savings, environmental benefits, and the home-services provided by their appliances. For instance, in a critical peak period DR event, some homeowners may prefer to delay their dryer while others would choose to adjust the thermostat. A smart home needs to understand these differences in preferences and control the home accordingly, ideally with little or no user interaction and minimal setup.

Prior research has focused on solving the load scheduling problem and usually assumes a generic function for comfort (Tsui and Chan 2012) or user-specified weights and operating windows for schedulable appliances (Rastegar, Fotuhi-Firuzabad, and Aminifar 2012; Anvari-Moghaddam, Monsef, and Rahimi-Kian 2015). These methods do not have a basis in behavioral science, and we could find very little research on how to elicit the parameters necessary to personalize any real-world implementation, though see Suryanarayanan, Devadass, and Hansen 2015 for an exception, and none that test their elicitation method with human subjects. This research seeks to fill this gap in the literature by testing several methods of quantifying comfort and eliciting home-service preferences. The end result is a top-level objective function that can be used by an optimization algorithm to determine an appropriate set of control signals over a given planning horizon.

To that end, we describe three different models of human comfort and how they were tested using online sample populations. We evaluate each method along three dimensions, a subjective user evaluation of the elicitation process, the time required to complete the process, and a measure of each methods' predictive ability. For each method, an online experiment was created to elicit the necessary parameters. All three methods shared the same starting and ending section involving comfort preferences and demographic questions respectively. Details and screenshots for each method can be found in the supplemental information.

Each experiment was fielded online using approximately 1,000 participants drawn from Amazon's Mechanical Turk population for a total of 3,000 participants. After a delay of at least five days, a random sample of 200 participants for each method was invited back to complete a validation experiment used to estimate each method's predictive ability. The methods and validation experiment will be described first, followed by a discussion of results.

2.1 Preference Elicitation Methodologies

For this study, the aspects of comfort considered included home air temperature, shower temperature and length, status of laundry and dishes, and financial and environmental costs. For convenience, we are defining comfort broadly to include both benefits, like comfortable air temperature, a nice shower, etc., and costs including monetary, convenience and environmental. Each of the methods is extensible to include other home services, such as an electric vehicle or window control for ventilation.

Three different models of human comfort, each inspired by a different field of study, were explored. The methods are Discrete Choice Modeling (DCM), Simplified Multi Attribute Rating Technique Exploiting Ranks (SMARTER) (Edwards and Barron 1994), and the Analytic Hierarchy Process (AHP). Each is capable of evaluating the quality of a home outcome and generating a unique, personalized score that can be used in an optimization algorithm to arrive at the best outcome, as evaluated by the system owner. Importantly, the functions are not tightly coupled to any particular optimization algorithm or type of DR and can work in both incentive- and priced-based programs (Albadi and El-Saadany 2008).

2.2 Preference Elicitation Survey Development

To explore these models, we conducted human subjects research using Amazon's Mechanical Turk (AMT), an online platform providing a low-cost pool of subjects that can be used for research purposes. AMT is an online platform where individuals from anywhere in the United States can sign up to perform 'Human Intelligence Tasks,' or HITs.

Common HITs include categorizing images, writing summaries, or transcribing images. However, researchers have been using the platform for social science experiments since at least 2009 (Mason and Suri 2012). The AMT population is more educated, younger, less wealthy, and more female than the overall adult U.S. population (Ipeirotis 2010). Samples drawn from AMT are more representative than in-person conveniences, but less representative than specially designed internet panels (Berinsky, Huber, and Lenz 2012). Many traditional findings have been replicated by Paolacci, Chandler, and Ipeirotis 2010 using AMT. Details of the experiments associated with each method can be found in the supplemental information.

We employed a two stage process to evaluate our three methods. In stage one, a survey for each method was designed, piloted, and completed by approximately 1,000 AMT workers (for a total of 3,000 responses). Individual level comfort models were created, and in stage two, 200 random participants for each method were invited back to take part in a validation experiment wherein their choices could be compared to the predictions made by the models. The methods will be described first, followed by a discussion of the results.

2.2.1 Discrete Choice Modeling

Discrete choice modeling (DCM) is a technique that is well established and used extensively in transportation (Ben-Akiva and Lerman 1985), health care (Reed Johnson et al. 2013), environmental valuation (Hoyos 2010), and others. DCM uses responses to a series of choice situations to estimate coefficients in a presumed utility model. Each choice situation includes a finite set of alternatives, and individuals either choose their favorite from among them or rank them from most desirable to least desirable depending on the experimental design. In short, individuals are assumed to make a deterministic choice among alternatives in a choice situation, but the researcher cannot observe all the relevant factors in the decision and so can only make probabilistic estimates about future choices. The typical use-case for DCM is to understand how individuals evaluate a product or service composed of multiple attributes. For instance, some of the earlier work using DCM was in the field of transportation (Ben-Akiva and Lerman 1985) to predict consumer response to new public transit options. In our case, we're using DCM to understand the comfort services that a home provides. Rather than the service being transportation, we're conceptualizing home comfort as a service. Thus, analogous to the characteristics of a travel option, we can treat the daily outcome in a home as composed of comfort related attributes, money, and environmental impact.

For that purpose, we assume an individual i has a goal to maximize utility. Therefore the home automation system optimizes using the following problem formulation, where our problem is formulated as equation 2.1 and utility costs are represented by equation 2.2:

$$\min \sum_{t=1}^H U_i(t) \quad (2.1)$$

$$U_i(t) = \beta_{i,m}(t)M + \beta_{i,c}(t)C + \beta_{i,d}(t)D + \beta_{i,l}(t)L + \beta_{i,sl}(t)S_l + \beta_{i,T}(t)\Delta T(t)^2 + \beta_{i,Tn}(t)I_{\Delta T < 0}(t)\Delta T(t)^2 + \varepsilon_i(t) \quad (2.2)$$

where:

H = the number of timesteps in the planning horizon

M = the monetary cost of operating the home over the planning horizon

C = the environmental impact of operating the home over the planning horizon

D = indicator variable if the dishes are done when needed

L = indicator variable if the laundry is done when needed

S_l = the shower length in minutes less than an individual's preferred shower length

$\Delta T(t)$ = the air temperature in the home at time t , relative to an individual's preferred set point (e.g. a value of 2 means the indoor air is 2°F higher than their desired air temperature)

$I_{\Delta T < 0}(t)$ = an indicator equal to one if the air temperature at time t is below the homeowner's preferred value

$\varepsilon_i(t)$ = the random error term

and the β 's are parameters which personalize the system's outcomes

The air temperature coefficients ($\beta_{i,T}$ and $\beta_{i,Tn}$) allow for different curvatures above and below the desired temperature. We have assumed that people respond to temperature changes in a quadratic manner. This ensures that small deviations from the preferred temperature count less than subsequent deviations. However, initial pilot testing revealed that people have asymmetric preferences between being too cold or too hot. Some people are more sensitive to cold air than hot and vice versa. To capture this asymmetry, the $\beta_{i,Tn}$ coefficient modifies the magnitude of the quadratic coefficient for temperatures below the preferred value. A positive value will increase the slope at lower temperatures, indicating an individual that is more sensitive to being too cold than hot, and a negative value will decrease it, indicating the opposite.

We note that the β 's are likely time- and/or state-dependent for most homeowners. For example, the sensitivity to air temperature discomfort is much lower when the home is unoccupied than when it is occupied, and to water temperature variation is much higher when an occupant is showering than when washing dishes. There are also expected to be seasonal variations in preferred HVAC setpoints and sensitivities to the air being warmer or cooler than the deadband. The objective function in 2.1 is written such that the β 's are time-dependent, and this could easily be generalized to include other relevant parameters. For this project the β 's were treated as constants and we leave other formulations for future work.

The problem is formulated such that all preference weights can be time-dependent. For example, the value of $\beta_{i,T}$, or weight on thermal discomfort, is expected to be higher when people are at home - and therefore impacted by discomfort - than when they are away from home. The purpose of this section is to develop methods to learn these β 's, and in the remainder of our work the β parameters are assumed for each home, and used to execute model predictive control.

Methods for estimating DCM utility functions can be found in numerous textbooks (Train 2009; Louviere, Hensher, and Swait 2000). However, traditional DCM cannot generally be used to estimate individual-specific models because of the dozens or even hundreds of decisions that an individual would have to make to attain adequate accuracy (Louviere et al. 2008). However, hierarchical Bayesian (HB) techniques partially overcome this limitation by combining aggregate information on preferences with individual specific data (Lenk et al. 1996). Given these limitations, our DCM method excluded the shower temperature attribute after pilot testing found that individuals consistently valued this highly. Rather than increase the complexity of the design (which would result in a reduction in statistical power or increase in choice situations) we opted instead to use a constant multiple of each individual's shower length coefficient. This is discussed more thoroughly in the SI, along with a description of the choice experiment and screenshots of the actual implementation. Future implementations can safely assume that showers should be the desired temperature and any required shower curtailment would come from reduced length.

HB requires one to assume a distribution for each coefficient to be estimated. Log-normal and negative log-normal distributions were used throughout as most attributes have definite signs (e.g. all else being equal people should prefer spending less money, having clean laundry, not having to curtail their shower, etc.). The exception is the coefficient for the cool temperature modifier, $\beta_{i,m}$, it was given a normal distribution as it can reasonably be either positive or negative. For carbon, there are in theory individuals that prefer an increased environmental impact to reducing it. These individuals will instead end up with coefficients close to zero (indicating no preference for reducing environmental impact).

While we obtained a ranking of alternatives for each choice set, we have only used the most preferred attribute in estimating HB coefficients, rather than use an exploded logit approach. HB models were estimated using the RSGHB software library (Dumont, Keller, and Carpenter 2015).

2.2.2 SMARTER

SMARTER stands for Simple Multi-attribute Rating Technique Exploiting Ranks and is a technique developed by Edwards and Barron 1994 to quickly create a decision model for an individual. It works by breaking a problem up into attribute utilities and asking individuals questions to determine their shapes. Those attribute utility functions are normalized and linearly combining into an overall utility function via weights inferred from a user-provided ranking of the attributes. Individuals rank the attributes by considering the best and worst outcomes for each attribute. For the shower and air temperature, the best and worst values are user-supplied. For money and carbon, average values of \$2.00 and 12 lbs. of CO₂ more than usual were used. For further details and screenshots see the SI. Edwards and Barron found that SMARTER enabled users to avoid the most cognitively difficult task of weighting attributes

Table 1. The rank-based inferred weights with w_1 being the most important attribute

	Attributes	
	7	6
w_1	0.2590	0.2966
w_2	0.2174	0.2410
w_3	0.1781	0.1884
w_4	0.1406	0.1387
w_5	0.1038	0.0908
w_6	0.0679	0.0445
w_7	0.0334	-

relative to each other as required by many other preference elicitation processes. The authors claim that, “In short, when [inferred] weights don’t pick the best option, the one they do pick isn’t too bad.” (Edwards and Barron 1994).

While simple, evidence suggests that SMARTER is a robust method. (Srivastava and Connolly 1995) found that SMARTER resulted in recommendations that correlated well with actual judgments by test subjects. Pöyhönen and Hämmäläinen 2001 found that SMARTER estimated subjects’ top alternative identically to three other methods 76%–78% of the time, and subjects were equally satisfied with the overall scores produced by each method. Huang 2011 found that participants rated SMARTER the most useful and most trustworthy of the three methods they evaluated. Equation 2.3 summarizes the final SMARTER utility function.

$$U_i = w_{r_{i,m}}(-M/2) + w_{r_{i,c}}(-C/12) + w_{r_{i,d}}D + w_{r_{i,l}}L + w_{r_{i,s_l}}(1 - S_l/S_{l,min}) + w_{r_{i,s_T}}(S_l - S_{T,min}/S_{T,p} - S_{T,min}) + w_{r_{i,T}}(1/T_m a x^2 T^2 + 1/T_m i n^2 - 1/T_m a x^2 I_{T < 0} T^2) \quad (2.3)$$

Where variables in common with Eq. 2.2 are defined the same, min and max subscripts refer to an individual’s stated minimum or maximum value on the associated variable, and the w ’s are the rank-inferred weights (Roberts and Goodwin 2002). The r subscripts in the weights are the user-provided ranks for each attribute. For example, if individual ‘ i ’ ranked money as third most important, then $w_{r_{i,m}}$ would be w_3 from Table 1, 0.1781. Note that while equation 2.3 is provided in this section, the SMARTER-elicited weights can be converted for use in equation 2.2 and this was the method implemented in this project’s controller, which is discussed in section 4.

Eq. 2.3 also includes an explicit shower temperature term, S_l . Conveying shower temperature to respondents is difficult. The approach we took was to describe the shower as ranging from ‘Too Cold’ which was defined as the hot water being nearly gone $S_{l,min}$, to ‘Perfect Temperature’ $S_{l,p}$. These values would then need to be mapped to the corresponding actual water temperatures as measured by the smart home. For this reason and the importance of shower temperature generally, in future work we recommend assuming shower temperature is always as desired and focusing on shower length as the relevant control variable. In this case, the SMARTER equation drops the shower temperature term and the right-most column in Table 1 is used for the weights.

With some manipulation, the SMARTER utility function described previously can be transformed to the same utility function as used in DCM method. The major difference is that instead of inferring the weights using participants’ choices, SMARTER directly asks participants to specify the shape of the attribute utility curve, and then the remaining coefficients are inferred from the attribute ranking step. Thus, every individual that specifies the same rank will have the same weights. This makes the method simpler, while acquiring more information about the attribute utility curves with fewer questions than DCM, but in a less realistic fashion.

2.2.3 Analytic Hierarchy Process

Due to the team’s prior work with Analytic Hierarchy Process (AHP), especially at Colorado State University, a survey was developed to elicit preferences using this methodology. AHP is a popular MCDM tool that decomposes the decision problem into a hierarchical structure containing the goal and the decision elements—alternatives, and

the criteria. AHP presents the user with a series of questions to make a pairwise comparison between each element at every level in the hierarchy.

AHP fundamentally works by assigning:

- global priorities (normalized weights) for the criteria in terms of their importance, and
- local priorities (normalized weights) for the alternatives for each criterion.

The local and global priorities are then subsumed to form the final judgment matrix to arrive at the decision ratings (rankings) for the alternatives that contribute towards the goal.

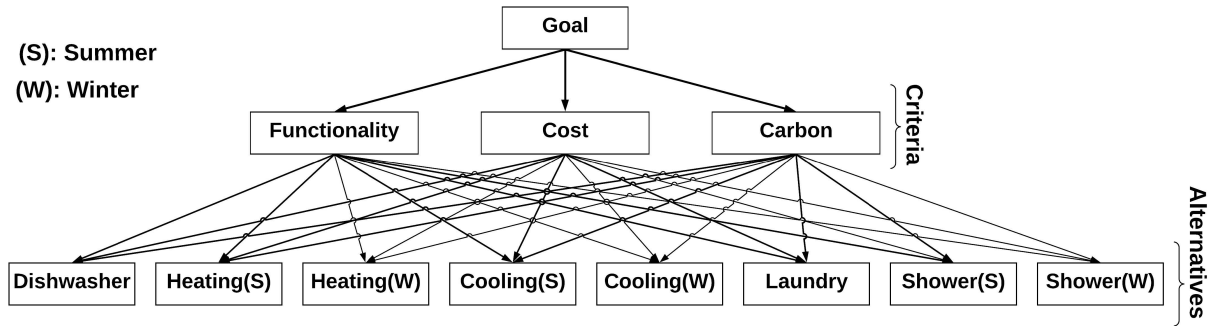


Figure 3. Analytic Hierarchy Process model for Home Energy Management System

The formulation of AHP preference weights and their use in an objective function is beyond the scope of this document. Because a paper providing this formulation is expected to be published very soon (Kadavil et al. 2018) and we did not select it for use in this project’s MCDM implementation, we refer you to the publication for the full detail.

2.3 Preference Elicitation Survey Results

2.3.1 General Comfort Preferences

Participants first answered questions about their home air temperature including their preferred, minimum, and maximum allowable values for summer and winter, and their preferred night time temperatures for summer and winter. These responses are summarized in table 2.

Table 2. Home air temperature preferences (mean, with interquartile range in parentheses) (°F)

Season	Min	Preferred	Max	Night
Summer	65.2 (60–70)	70.5 (68–74)	76.8 (74–80)	68.8 (65–72)
Winter	64.1 (60–68)	69.5 (68–73)	75.4 (72–78)	68.9 (65–72)

The average preferred temperature, both during the day and at night, did not change between summer and winter. However, in the summer, participants on average indicated a slight, 1°F shift towards higher acceptable temperatures. These averages mask puzzling underlying heterogeneity. About one-third of participants indicated no difference in preferred summer and winter temperatures, another one-third indicated a preference for higher summer than winter temperatures by an average of 5.6°F, and the final one-third actually preferred higher temperatures in the winter than the summer, by 5.0°F. The size of this final group was a surprise to the research team, and further cognitive testing should be performed on the data entry form to ensure that participants were in fact interpreting the questions correctly.

After answering home air temperature questions, participants answered questions about their shower preferences. Interestingly, 68.8% of respondents believe their shower is slightly hotter or much hotter than average, and only 5.6% think it is slightly or much colder than average. For reference, in an experiment involving women who were allowed to set their own shower temperature (Ohnaka, Tochihara, and Watanabe 1994), the average temperature chosen was

104°F. Table 3 summarizes shower preferences by gender from our surveys. We also asked participants about what appliances they owned and how they used them.

Table 3. Shower preferences in minutes by gender (IQR = interquartile range)

Gender	Count	Average Length			Minimum Length	
		Mean	SD	IQR	Mean	SD
Male	1,246	12.0	6.3	8-15	5.2	3.3
Female	1,658	14.1	6.6	10-15	6.6	3.5

23 individuals did not report a gender; those results are excluded

2.3.2 Discrete Choice

The responses to the choice situations in the discrete choice experiment were used to estimate individual level coefficients using hierarchical Bayesian methods. Figure 4 shows the distribution of individual level coefficients.

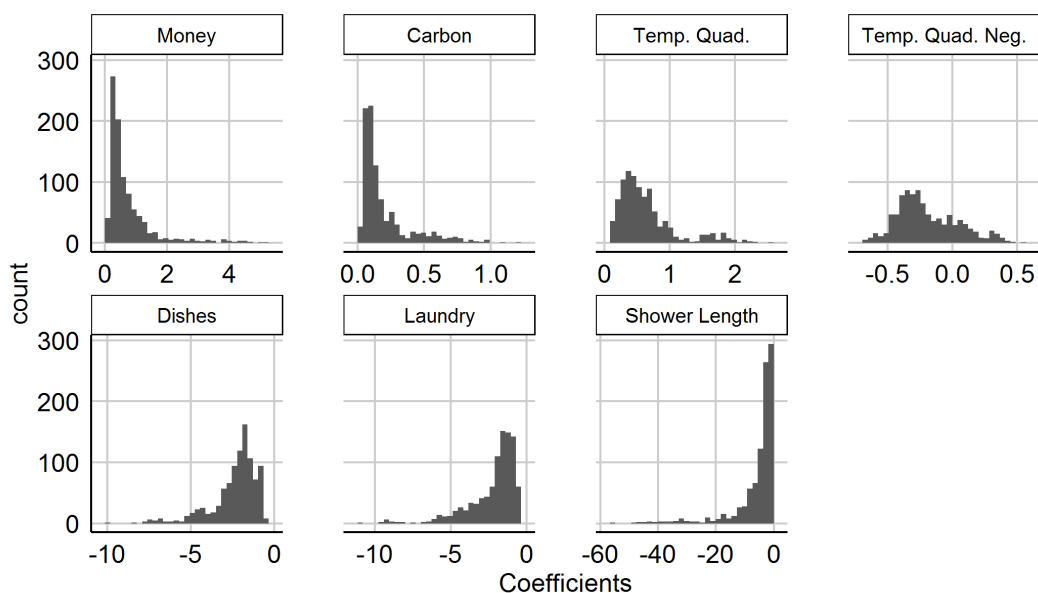


Figure 4. Discomfort coefficients obtained using the DCM method

From the figure, “Temp. Quad. Neg” refers to the $\beta_{i,ln}$ coefficient from Eq. (1). The results show more individuals with a negative value, which matches the results from the SMARTER method that indicated more people disliked being too hot than too cold.

2.3.3 SMARTER

A total of 956 individuals successfully completed the SMARTER preference elicitation experiment. Across all SMARTER participants, 70% reported that being too hot was worse than being too cold.

One of the most important outcomes in SMARTER is the ranking of attributes. The full distribution of ranks is shown Figure 5. Carbon is very often the least important item to individuals, but then somewhat uniformly distributed throughout the remaining ranks. Money is the opposite, most often ranked first, it is then distributed throughout the remaining ranks with a heavy tail compared to other attributes. Taking the average rank given to each attribute by all SMARTER participants, the average rank from most important to least important is air temperature, money, shower temperature, shower length, having clean clothes, carbon, and finally having clean dishes.

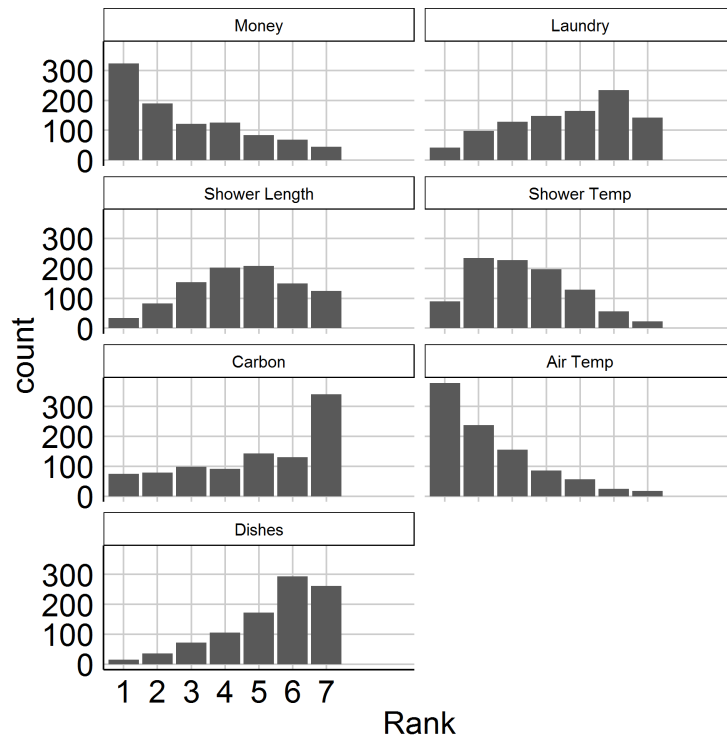


Figure 5. The distribution of ranks given to the home comfort attributes in the SMARTER method.




Figure 6. Intracriterion comparison of the SMARTER method survey ranks.

2.3.4 Analytic Hierarchy Process

As indicated in the next section, the initial predictive measure comparison did not indicate AHP was compelling for use in **foresee**. Thus we elected to implement another preference elicitation methodology in the software. The AHP method and findings are expected to be published soon (Kadavil et al. 2018); but ongoing analysis indicates a stronger potential for AHP than initially understood.

2.4 Method Comparison

After the method specific experiments, approximately 200 randomly chosen participants from each method were invited back for a follow-up experiment. Similarly to the discrete choice experiment, they were presented with a series of choices situations, each with two options. First, they made 12 choices related to full-day comfort outcomes from an orthogonal fractional factorial design. Next, they answered four questions that were custom-made to represent demand response scenarios. The twelve choice situations were created using an orthogonal fractional factorial design using the Ngene software package (Collins, Rose, and Bliemer 2007). Other designs were possible, but orthogonal fractional factorials are well established in the literature, simple to generate, and allow the unbiased estimation of main effects (Louviere, Hensher, and Swait 2000). Future research will consider altering the validation experiment. An ideal experiment would provide more tests around those attributes that are considered more important. Currently, the pairwise comparisons are methodologically similar to the ranking exercise in the DCM method. This may unfairly advantage this method over the others. Further details can be found in the supplemental information, and in (Aloise-Young et al. 2018), (Christensen et al. 2016), and (Kadavil et al. 2018).

Results from the longitudinal surveys were compared against predictions made from the preference survey data from the same users' initial survey inputs. These were compared to each other, as were the survey participants' report of usability, and the average time they took to complete the surveys. Those results are shown in Table 4.

Table 4. Comparison of Preference Elicitation Methods

Survey Method	Percent Correctly Predicted	Average Usability Score (scale of 1-7)	Average Completion Time (min)
AHP	49	2.48	9.0
DCM	68	2.57	5.7
SMARTER	72	2.53	5.5

Based on the preliminary findings in the project's first year, and the need to implement a method to perform model-predictive control to achieve those preferences, the SMARTER method was chosen. It appeared to have the highest predictive ability, a short survey time, and its objective function was sufficiently similar to the DCE method that DCE could be leveraged in the future for continuous (and active) reinforcement learning of preferences.

We performed a sensitivity analysis and found that a subset of the SMARTER survey questions appeared sufficient to initiate the **foresee** multicriterion decision making objective function. These few questions would be used when **foresee** was initialized, to establish an initial set of control objectives.

2.5 User Interface for Preference Initialization

The SMARTER-subset questions were implemented in a simple graphical user interface (GUI) in our prototype user interface, to enable the system to be manually initialized for our tests. (Of course, we can also set these automatically to enable laboratory experiments.)

First, the user must connect **foresee** to their Wi-Fi network to enable access to the Internet and to networked appliances, as shown in figure 7. Then, the user is asked to provide information about their air temperature preferences, as shown in figure 8. Then, the user is asked to provide information about their hot water preferences, as shown in figure 9. In the final initialization step, the user is asked to rank the relative importance of different outcomes by dragging them up and down the screen, as shown in figure 10. This completes the initial setup.

Of course, **foresee** needs other devices in order learn about the home and help manage the home. The user can set up and manage those devices in the user interface as well, as shown in figure 11. The homeowner can monitor their current house state and benefits received on the main homescreen of **foresee**, as shown in figure 12.

This GUI was developed primarily for our use in monitoring and demonstrating the capabilities of the early prototype **foresee** platform, but it also serves as an example for others who may commercialize **foresee** or seek to apply the preference elicitation in other ways.

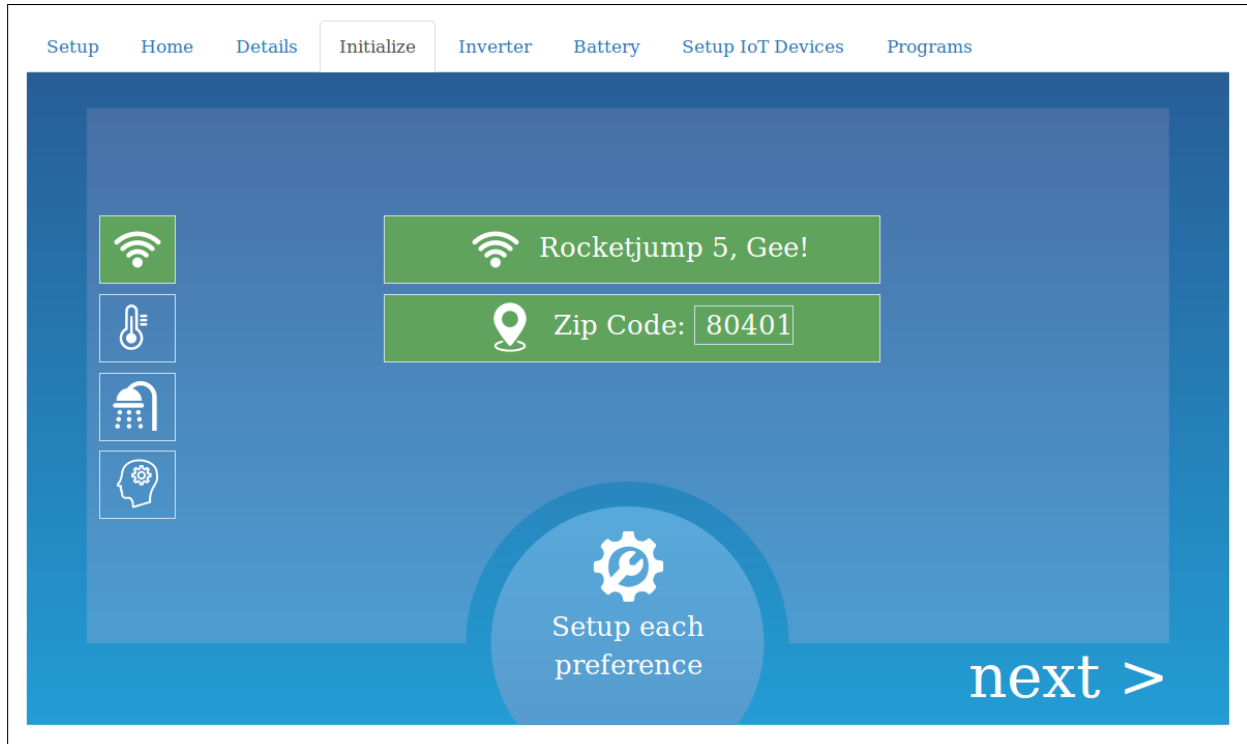


Figure 7. Initialization screen where users input Wi-Fi and location information.

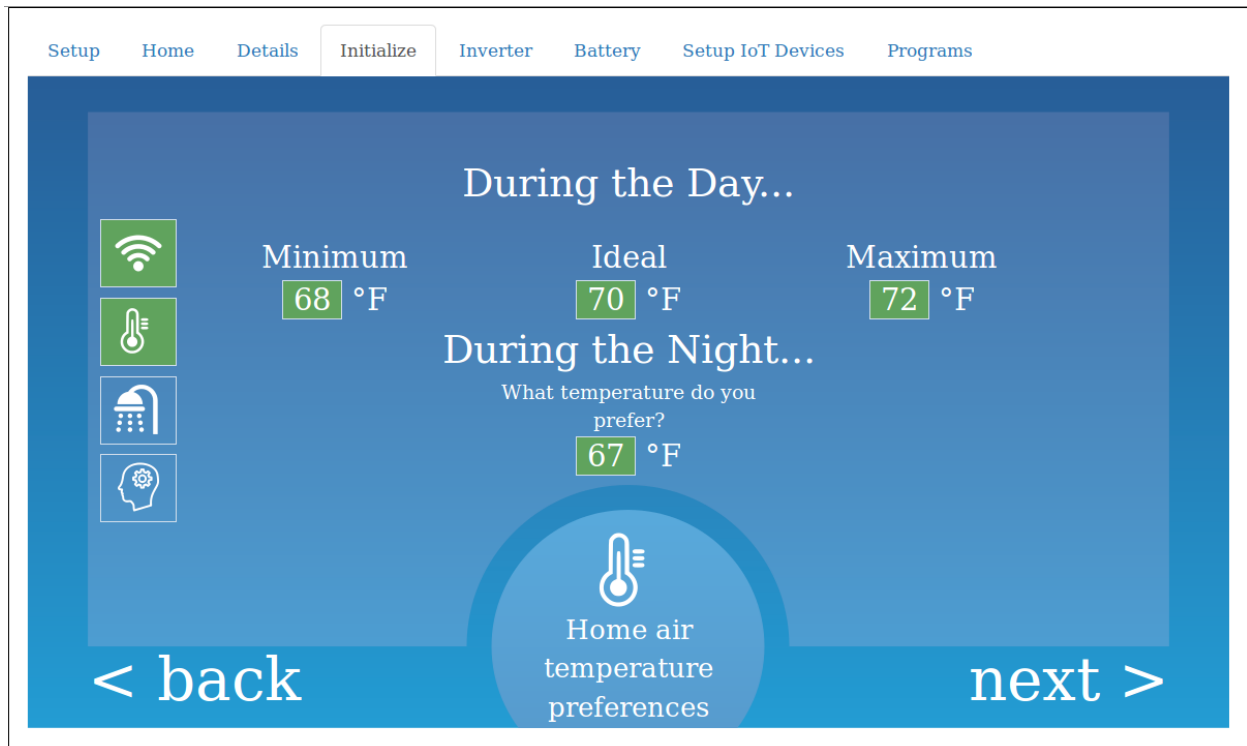


Figure 8. Initialization screen where users input air temperature preferences.

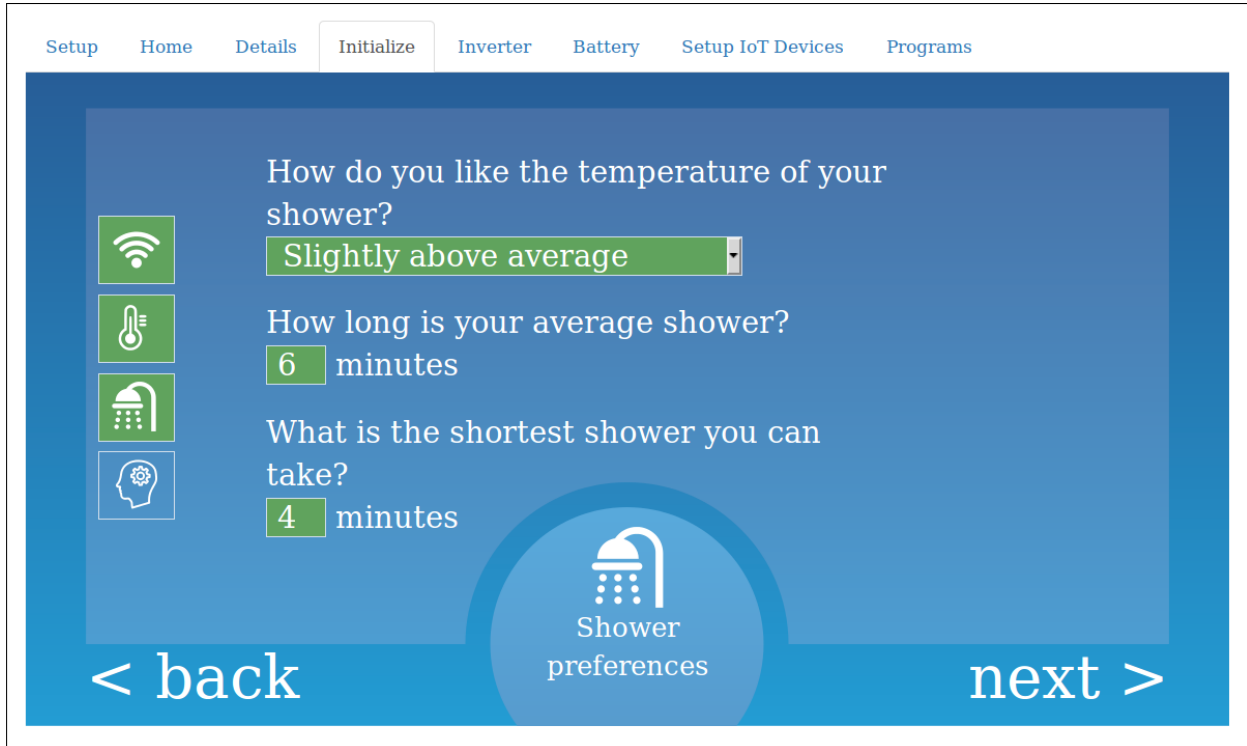


Figure 9. Initialization screen where users input hot water (shower) preferences.

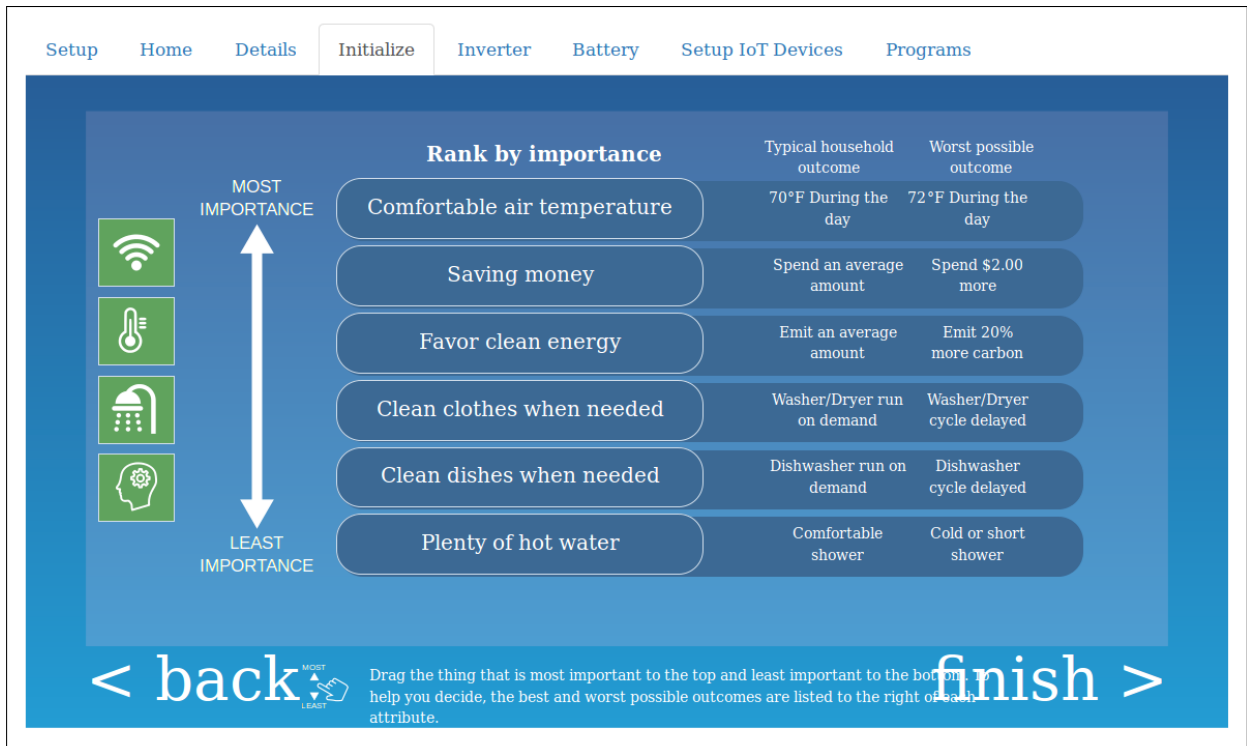


Figure 10. Initialization screen where users rank their service preferences.

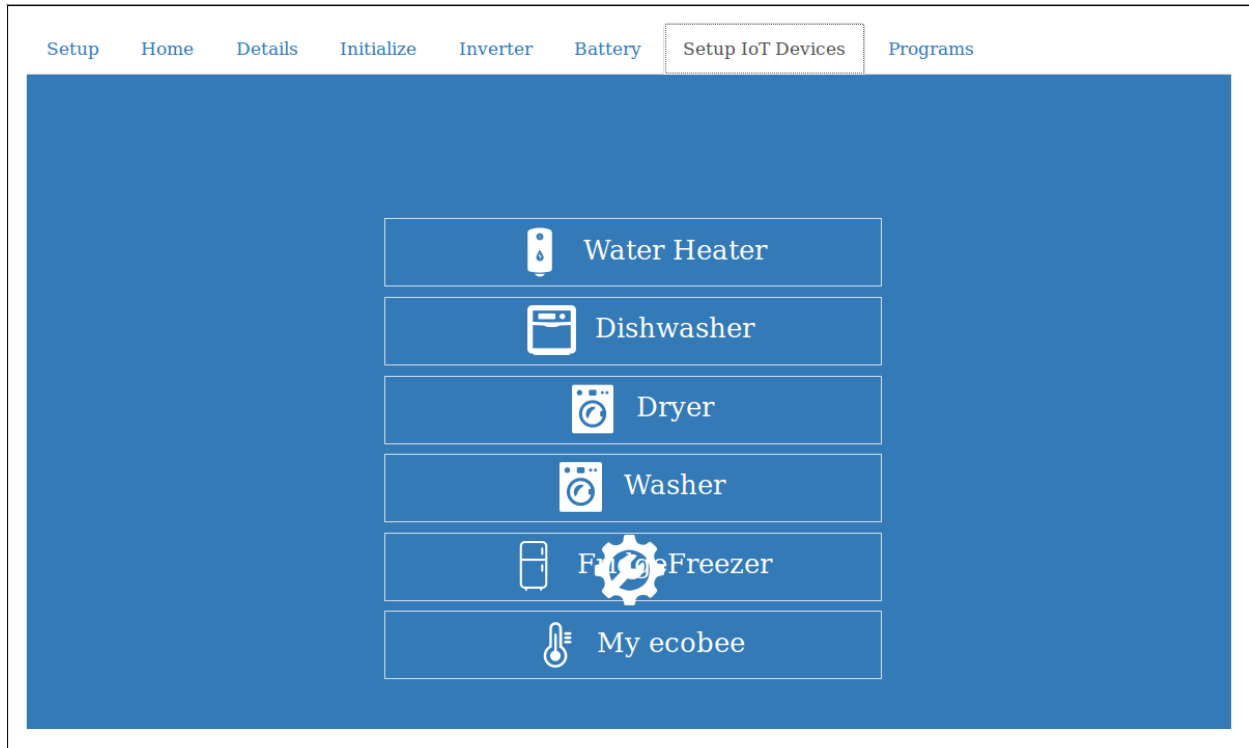


Figure 11. Device connection screen.

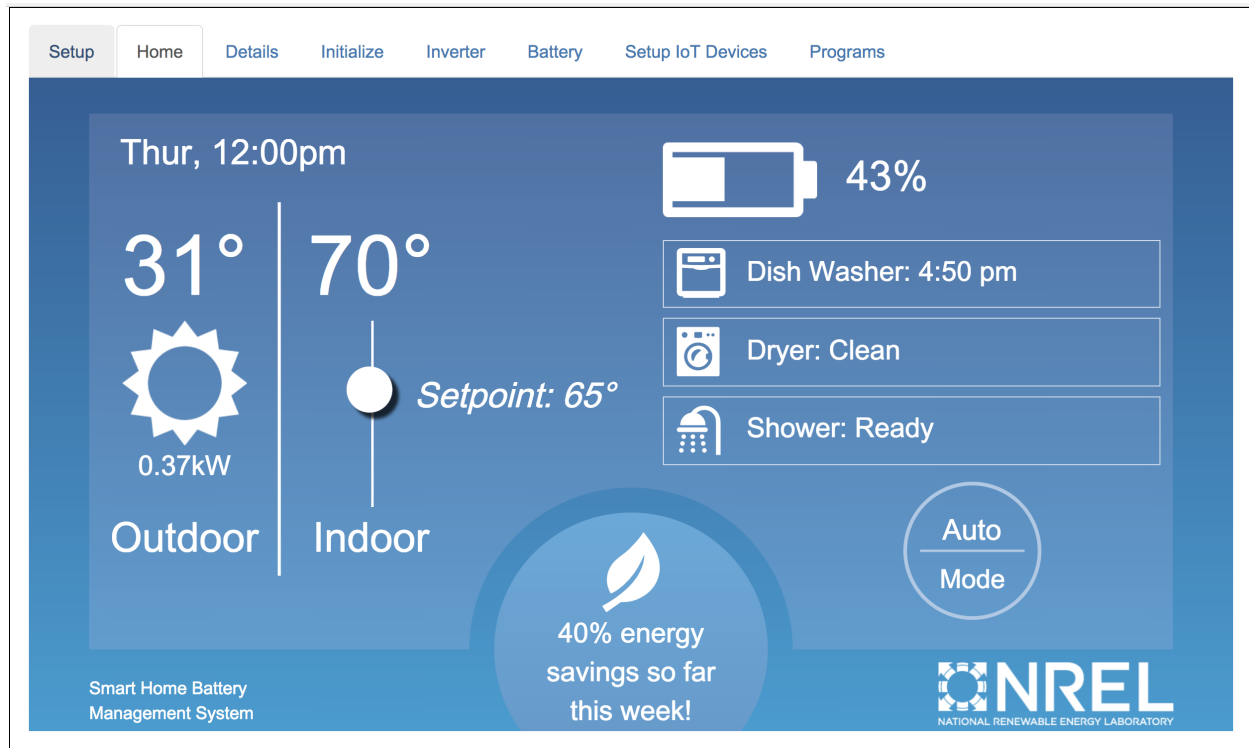


Figure 12. Main user interface screen for foresee.

3 Machine Learning

foresee use machine learning methods to use data available from the smart home systems in two ways:

- a) to develop predictive models of the occupants, their schedules and their demands for building services; and
- b) to develop predictive models of the building and other physical assets to predict the energy requirements and human impacts which arise from executing available control actions.

We used statistical learning and system identification methods to derive these models, respectively. The models then become inputs to the decision-making optimization routine, which is described in the next chapter.¹

3.1 Statistical Learning

Statistical learning is used to understand the patterns in uncontrollable loads and user behaviors, such as hot water draws, both of which are important for **foresee** to accurately predict future loads and plan the resources for upcoming DR events. Statistical methods, such as a simple averaging method, or more sophisticated machine-learning methods, such as Expectation Maximization of Gaussian Mixture Models (Bishop 2006), can be used to learn the usage patterns. In this initial development, a simple averaging method is used to represent the load patterns at the each of the intervals equal to the MPC's prediction time step, which is 30 min in our study. MPC is performed over multiple time steps, such as 8-hr look-ahead comprising 16 time steps. The daily load patterns are represented by 48 bins, and the value of each bin is the average of the training data at the corresponding 30-min interval. This time discretization is flexible and can be changed for future use case development.

Uncontrollable loads are relatively small compared to major building loads, but still consume a significant amount of energy. The estimated uncontrollable loads are included in the MPC to predict the whole-home load. Hot water draws are the driving force of water heater energy consumption and are critical to the accurate prediction of water heating loads. The estimated water draws are used as an input to the water heater model.

3.2 System Identification

To accurately predict the behavior (i.e., the output) of a controllable device with a given control command (i.e., the input), it is critical to obtain a dynamic model to represent the input-output relationship of the device. The house envelope model and the associated heating, ventilation, and air conditioning (HVAC) model are of particular interest not only because space cooling and heating is the largest load in the residential building sector, but also because each home is unique and no default model can accurately predict the HVAC power and indoor temperature without performing model calibration in the actual home.

Many system identification techniques have been developed for buildings applications and a review of the techniques can be found in (Bloem 1994). In this project, an equivalent circuit model of one resistor-capacitor pair is used to represent a single-zone house model and the coefficients can be obtained using multivariate linear regression. We use the learned house model to update the indoor temperature and power consumption based on the outdoor temperature, solar irradiance, and thermostat set point. The indoor temperature is measured at or near the thermostat. The outdoor temperature and solar irradiance forecast can be obtained from weather service providers. Other appliances are also modeled in a simplified way using the identified parameters. A water heater is modeled as a two-node stratified tank with nominal power ratings learned from historical data (Jin, Maguire, and Christensen 2014).

Schedulable appliances such as dishwashers and clothes dryers are represented by an average runtime and power consumption. A categorization of the appliances and the appliance models are discussed in detail in the following.

Quadratic programming is used to solve the optimization problem, which means the optimal solutions for the appliances are duty cycle commands in fraction form instead of binary variables (i.e., on/off controls). Appliances with variable-speed or variable-power control capabilities are able to directly use the duty cycle commands as the input.

¹Note that much of this chapter has been previously published in (Jin et al. 2017), (Jin et al. 2016), and (Raszmann et al. 2017), and we duplicate that prior work here for the reader's convenience.

Appliances with on/off characteristics may convert the duty cycle commands to on-off control signals using time slot techniques that consider minimum on and off time.

User-initiated loads, such as those generated by the dishwasher and clothes dryer, are also important for predictions in the optimization process. Simplified models are derived from the data to represent the appliances using average runtime and power consumption.

3.3 Residential Equipment Models

The residential equipment considered in the paper includes both building loads and DERs. Residential building loads can be generally categorized into thermostatically control loads (TCLs), schedulable loads, and uncontrollable loads. DERs, such as PV and battery storage, are gaining ground in the residential sector thanks to the decreasing cost of PV panels and lithium-ion batteries.

TCLs include space heating, space cooling, water heating, and refrigeration. Schedulable loads refer to the loads that operate based on recurring schedules or one-time schedules. Examples include pool pumps and wet cleaning devices, such as dishwashers and clothes washers and dryers. TCLs and schedulable loads can be operated at low-energy states to improve EE or shifted to enable DR. They are considered controllable loads and make up about 50% of the total electricity consumption in the sector. Other loads, such as lighting, cooking, electronics, computers, and miscellaneous loads, are unsuitable for providing EE or DR applications for various reasons. These loads are considered uncontrollable loads because they are not directly controlled by the HEMS.

Among the controllable loads, TCLs, such as HVAC and water heaters, have inherent thermal storage and can provide large load reduction if precooling or preheating is implemented to shift the loads. These devices can also be operated at low-energy states when they do not need to provide primary services. For example, temperature setback can be performed in the HVAC system when the home is not occupied. Note that refrigeration is not considered in this paper because of food safety and user inconvenience concerns.

Schedulable loads are usually initiated by users with anticipated completion time. The cycles can be shifted to accommodate excessive PV generation or reduce load during DR events as long as the cycles end by the anticipated completion time. In this paper, we consider dishwashers and clothes dryers because they have heating elements and tend to consume more energy. Clothes washers and pool pumps are not considered here.

Although uncontrollable loads cannot be used to provide DR because of a lack of controllability, they make significant contributions to the base load, which needs to be estimated accurately for **foresee** to coordinate with DERs and maximize the load reduction during DR periods.

Building loads' energy consumption is not always coincidental with generation from rooftop PV. Under net metering policies, any excessive PV generation can backfeed to the grid; however, under high penetrations of distributed generation, this may cause local overvoltage problems in the distribution grid. With **foresee**, both the home battery system and controllable loads can be used to maximize the on-site consumption of the electricity from PV, thereby improving the PV hosting capacity in the grid.

Submetered data are required by **foresee** to track the appliance status in the prediction models or estimate the energy consumption using data-driven approaches. We assume that the submetered utility data are either self-reported by the appliances or provided by low-cost submetering solutions, such as nonintrusive load monitoring techniques (Christensen, Earle, and Sparr 2012; Jin 2017).

3.3.1 House-HVAC Model

The indoor air temperature, T_{air}^{in} , of a house is updated using a house model coupled with an HVAC model:

$$T_{air}^{in}(t+1) = T_{air}^{in}(t) + \gamma_1(T_{air}^{out}(t) - T_{air}^{in}(t)) + \gamma_2[\eta_{hvac}^h P_{hvac}^h(t) - \eta_{hvac}^c P_{hvac}^c(t)] + \gamma_3 E_e(t) \quad (3.1)$$

where T_{air}^{in} is the indoor air temperature, T_{air}^{out} is the outdoor air temperature, and E_e is the solar irradiance. γ_1 , γ_2 , and γ_3 are the coefficients learned from system identification and correspond to the building envelope, heating or cooling gain, and solar heat gain, respectively.

Simplifications are made to assume the HVAC system has constant efficiency (η_{hvac}^h and η_{hvac}^c) and nominal power consumption ($P_{hvac}^{h,nom}$ and $P_{hvac}^{c,nom}$) when operating in both heating and cooling modes. In a real building, these values will need to be machine-learned and nominal/constant values may not be sufficiently representative of actual system performance.

We denote HVAC runtime fraction as the control variable:

$$\begin{aligned} U_{hvac}^h &= P_{HVAC}^h / P_{hvac}^{h,nom} \\ U_{hvac}^c &= P_{HVAC}^c / P_{hvac}^{c,nom} \end{aligned} \quad (3.2)$$

Then the model in Eq. 3.1 is subject to the following constraints:

$$T_{air}^{min} \leq T_{air}^{in} \leq T_{air}^{max} \quad (3.3)$$

$$0 \leq U_{hvac}^h, U_{hvac}^c \leq 1 \quad (3.4)$$

where Eq. 3.3 defines the range of the allowable indoor air temperature, and Eq. 3.4 indicates the range of HVAC control variables that can be interpreted as the desired runtime fractions. The HVAC cannot be operated in two modes during the same time interval.

We note here, and discuss later, that our laboratory results indicate this fairly naïve building and HVAC model may be insufficient to enable accurate model-predictive control of the air conditioner to ensure comfort of the occupants. Future work is warranted to examine and develop more predictive and customer-satisfactory building models. It is important for **foresee**'s future commercial use cases that the model used in MPC be compatible with machine learning methods and data available from the home and its equipment, to minimize or avoid any barriers from requiring to user inputs or a home audit.

3.3.2 Water Heater Model

A two-node electric resistance water heater model based on Jin, Maguire, and Christensen 2014 is used to create a balance between prediction accuracy and computation speed. Compared to the one-node model used in most literature, the two-node model can capture the stratification phenomenon in the tank, thus representing the outlet temperature more accurately:

$$T_{wh}^{low}(t+1) = T_{wh}^{low}(t) + \frac{1}{C_{wh}^{low}} \left[UA_{wh}^{low} (T_{air}^{in}(t) - T_{wh}^{low}(t)) + \Delta m(t) C_p (T_{mains} - T_{wh}^{low}(t)) + \eta_{wh}^{low} P_{wh}^{nom,low} U_{wh}^{low}(t) \right] \quad (3.5)$$

$$T_{wh}^{up}(t+1) = T_{wh}^{up}(t) + \frac{1}{C_{wh}^{up}} \left[UA_{wh}^{up} (T_{air}^{in}(t) - T_{wh}^{up}(t)) + \Delta m(t) C_p (T_{wh}^{low} - T_{wh}^{up}(t)) + \eta_{wh}^{up} P_{wh}^{nom,up} U_{wh}^{up}(t) \right] \quad (3.6)$$

$$T_{dhw}(t+1) = T_{wh}^{up}(t+1) \quad (3.7)$$

where superscripts *low* and *up* represent the lower node and upper node of the tank. Eqs. 3.5 and 3.6 indicate how the water temperature is updated by a number of input variables. T_{wh}^i is the water temperature and U_{wh}^i is the control signal of tank node i in terms of duty cycle. T_{air}^{in} is the indoor air temperature, and T_{mains} is the mains water temperature entering the tank. UA_{wh}^i is the product of heat loss coefficient and surface area of node i , Δm is the flow rate of hot water draws, C_p is the heat capacity of water, C_{wh}^i is the thermal capacitance of tank node i , and $P_{wh}^{nom,i}$ and η_{wh}^i are the rated power and efficiency of the resistive element in node i , respectively.

Eq. 3.7 indicates the temperature at which the hot water is delivered by the top node of the water heater. The water heater model is subject to the following constraints:

$$T_{wh}^{min} \leq T_{wh}^{up} \leq T_{wh}^{max} \quad (3.8)$$

$$T_{wh}^{low} \leq T_{wh}^{up} \quad (3.9)$$

$$U_{wh}^{low} + U_{wh}^{up} \leq 1 \quad (3.10)$$

$$0 \leq U_{wh}^{low}, U_{wh}^{up} \leq 1 \quad (3.11)$$

where Eq. 3.8 dictates the constraints of the hot water temperature for safety reasons as the hot water exits the tank from the upper node. Eq. 3.9 enforces the thermal stability in the tank such that the lower node should not be hotter than the upper node because of buoyancy. Eqs. 3.10 and 3.11 are the constraints of control signals U_{wh}^{low} and U_{wh}^{up} , which are continuous variables between 0 and 1 and can be interpreted as duty cycles.

Machine learning can, over time, be used to identify the appliance variables and coefficients: C_{wh}^i , UA_{wh}^i , η_{wh}^i , and P_{wh}^i . Because residential electric resistance water heaters are, generally, commodity products, common models can be developed for standard appliances and tank sizes. These can be more efficiently learned with application of active learning methods; we plan to implement this as future work.

3.3.3 Schedulable Appliance Models

Schedulable appliances operate based on user-specified schedules and the cycles can be shifted to meet other high-priority needs. These appliances are distinct from TCLs, such as HVAC and water heaters, because the only control available is to shift the scheduled cycles. Dishwashers and clothes dryers are considered in this project² with the following constraints:

$$P_{app}(t) = U_{app}(t)P_{app}^{nom} \quad (3.12)$$

$$0 \leq U_{app} \leq 1 \quad (3.13)$$

$$\sum_{h=1}^H U_{app}(t) = RT_{app}^{nom} \quad (3.14)$$

where P_{app}^{nom} is the nominal power of the appliance, U_{app} is the control variable indicating the duty cycle of that appliance, and RT_{app}^{nom} is the average runtime of the appliance in terms of prediction steps. P_{app}^{nom} and RT_{app}^{nom} can either be learned via system identification or provided by the appliances.

3.3.4 PV Model

The PV model and the operating constraint are as follows:

$$P_{pv}(t) = (1 - U_{pv}(t))P_{pv}^{max}(t) \quad (3.15)$$

$$0 \leq U_{pv} \leq 1 \quad (3.16)$$

where Eq. 3.15 indicates that the available PV power is a function of the maximum PV generation P_{pv}^{max} and PV curtailment U_{pv} . P_{pv}^{max} is proportional to the solar irradiance, E_e , and the area of the PV panel. Note $P_{pv}^{max} \geq 0$ and $P_{pv} \leq 0$. As shown in Eq. 3.16, U_{pv} is a continuous variable and represents the amount of real power curtailment, with 0 denoting no curtailment and 1 denoting a curtailment of 100%. Reactive power control was not considered in this project, and incorporation of this increasingly-common feature of PV inverters is left for future work.

3.3.5 Battery Model

The battery state of charge SOC , the charging power, P_{bat}^{ch} , and the discharging power, P_{bat}^{dis} , are related as follows:

$$SOC(t+1) = SOC(t) + \frac{\eta_{bat}^{ch} \Delta t}{Q_{bat}} P_{bat}^{ch}(t) + \frac{\Delta t}{\eta_{bat}^{ch} Q_{bat}} P_{bat}^{dis}(t) \quad (3.17)$$

²Note that clothes washers are not considered strong candidates for load deferral in this project because the energy consumption of this appliance is very low compared to the dryer. We also did not include the refrigerator's ice making and defrost operating states. While these are potentially good candidates for load deferral, the connected refrigerator we received did not have an ice maker, and lacked the capability for external defrost control.

where: $P_{bat}^{ch} \geq 0$, $P_{bat}^{dis} \leq 0$, $\eta_{bat}^{ch}, \eta_{bat}^{dis} \in (0, 1]$ are the discharging and charging efficiency of the battery system, Δt is the length of the prediction step, and Q_{bat} is the capacity of the battery.

Define the battery control variables $U_{bat}^{ch} = P_{bat}^{ch} / P_{bat}^{ch,max}$ and $U_{bat}^{dis} = P_{bat}^{dis} / P_{bat}^{dis,max}$, which represent the percentage of maximum charging power, $P_{bat}^{ch,max}$, and maximum discharging power, $P_{bat}^{dis,max}$, respectively.

The battery system model in Eq. 3.17 is subject to the following constraints:

$$SOC_{min} \leq SOC(t+1) \leq SOC_{max} \quad (3.18)$$

$$0 \leq U_{bat}^{ch}(t), U_{bat}^{dis}(t) \leq 1 \quad (3.19)$$

where Eq. 3.18 defines the operable SOC range for reducing battery degradation and Eq. 3.19 indicates the range of the normalized battery control variables. The battery cannot charge and discharge at the same time.

We are working, but have not completed adaptation of the internal battery model from (Raszmann et al. 2017) into a usable model in the MPC formulation, which will enable better battery degradation representations in the objective (cost) function. This is left as future work.

3.3.6 House Power Balance

Denote the total load in a home as:

$$\begin{aligned} P_{load}(t) &= P_{hvac}(t) + P_{wh}(t) + P_{dw}(t) + P_{cd}(t) + P_{unctrl}(t) \\ &= U_{hvac}^c(t)P_{hvac}^{nom,c} + U_{hvac}^h(t)P_{hvac}^{nom,h} + U_{wh}^{low}(t)P_{wh}^{nom,low} + U_{wh}^{up}(t)P_{wh}^{nom,up} + \\ &\quad U_{dw}(t)P_{dw}^{nom} + U_{cd}(t)P_{cd}^{nom} + P_{unctrl}(t) \end{aligned} \quad (3.20)$$

where variable P_{unctrl} denotes the estimate of uncontrollable electrical loads (i.e., lighting, television, and plug loads). Then the overall power balance of the home is as follows:

$$-P_{grid}(t) + U_{bat}^{ch}(t)P_{bat}^{ch,max} - U_{bat}^{dis}(t)P_{bat}^{dis,max} - (1 - U_{pv}(t))P_{pv}^{max}(t) + P_{load}(t) = 0 \quad (3.21)$$

$\forall t \in [0, H - 1]$. P_{grid} denotes the power drawn from the grid and is positive when the power flows from the grid to the home.

There is no explicit constraint on PV backfeeding (energy export to the grid) although backfeeding may be penalized during optimization to account for utility rate disincentives (feed-in tariff) and/or loss of environmental benefits. Battery backfeeding to the grid is undesirable due to the non-negligible round-trip energy loss and unnecessary degradation caused by cycling fade. Therefore we further impose a constraint on battery power such that the power discharged from the battery must be consumed behind the meter:

$$U_{bat}^{ch}(t)P_{bat}^{ch,max} - U_{bat}^{dis}(t)P_{bat}^{dis,max} + P_{load}(t) \geq 0 \quad (3.22)$$

Note that this formulation assumes a constant efficiency and capacity for the operating equipment. We leave the generalization of this formulation to state-dependent efficiency and capacity to be completed in future work.

3.4 Application to Field Data

Field data from the Residential Building Stock Assessment (RBSA) data set (Ecotope 2014) were chosen for this study. RBSA was a field survey performed by the Northwest Energy Efficiency Alliance on more than 1,850 residential sites, including 100 submetered single family homes. In the submetered homes, temperature readings and energy metering were recorded every 15 min except for the indoor temperature, which was recorded at 1-hr intervals.

Five representative homes out of the 100 submetered homes were selected for numerical evaluation in this study. These homes were selected because they are all electric, distributed across the entire Pacific Northwest, and equipped with the most common space heating and cooling technologies in this region. Table 5 provides a brief summary of the homes selected for use in our study.

Table 5. Summary of five homes from the residential building stock assessment data set used in the study

ID	Location	Climate	Controllable building loads			
			Space heating	Space cooling	Water heating	Wet cleaning
1	Tenino, WA	Marine	Heat pump	Heat pump	Electric resistive	Clothes dryer & dishwasher
2	Eugene, OR	Marine	Heat pump	Heat pump	Electric resistive	Clothes dryer & dishwasher
3	Seattle, WA	Marine	Forced air furnace	N/A	Electric resistive	Clothes dryer
4	Emmett, ID	Cold	Baseboard heater	N/A	Electric resistive	Clothes dryer
5	Tacoma, WA	Marine	Baseboard heater	N/A	Electric resistive	Clothes dryer

3.4.1 Estimating Usage Patterns Using Statistical Learning

Statistical learning methods extracted the load patterns of uncontrollable loads and water draw patterns, which were used in the MPC framework to facilitate decision-making. Two weeks’ data prior to the testing days were used to learn the usage patterns and predict the uncontrollable loads and water draws on the testing days. Fig. 13 shows the result of the uncontrollable load estimation in one of the RBSA submetered homes using the statistical learning method.




Figure 13. Comparison of the estimated and actual uncontrollable loads.

The root mean squared error (RMSE) of the averaging-based estimation method is less than 0.2 kW, whereas the daily average of the uncontrollable load is about 1 kW for this particular RBSA home. The spikes in the actual uncontrollable load have significant contributions to the RMSE because the estimated load profile is much smoother.

Similarly, the water draw patterns were derived from the historical data and used as an input to the water heater model in section 3.3.2. Water draw data collected from a different field test (Jin, Maguire, and Christensen 2014) was used to augment the RBSA data set that has no recorded water draw data. The same draw profile was used in all five RBSA homes to avoid additional complexity. Water draw patterns were difficult to predict because of day-to-day variation in user behavior. The statistical learning method accurately predicted the water draws in most parts of the day when large draws occurred and the daily RMSE of draw volume prediction was 4.43 gallons for a field test site with total average daily draw of 90 gallons.

3.4.2 Creating House Models Using System Identification

System identification was performed on the RBSA data to learn the house and HVAC model from the historical data. Three weeks’ data were exported from each RBSA house: the first 2 weeks’ data were used as the training data and the last week’s data were held back as the validation data. A linear regression was performed on the training data to extract the coefficients of the house model in Eq. 3.1. The house model was then used to predict the indoor air temperature and compared with the validation data. Fig. 14 shows a comparison of the indoor air temperature recorded in the RBSA data set and the estimated value derived from the learned house model for a RBSA house in Tenino, Washington. The estimated temperature follows the actual value with relatively high accuracy; statistical results are shown in Table 6.

Two metrics were used to evaluate the performance of the learned house model from system identification, namely, root mean square error (RMSE) and coefficient of determination (or R^2). RMSE is a commonly used measure of

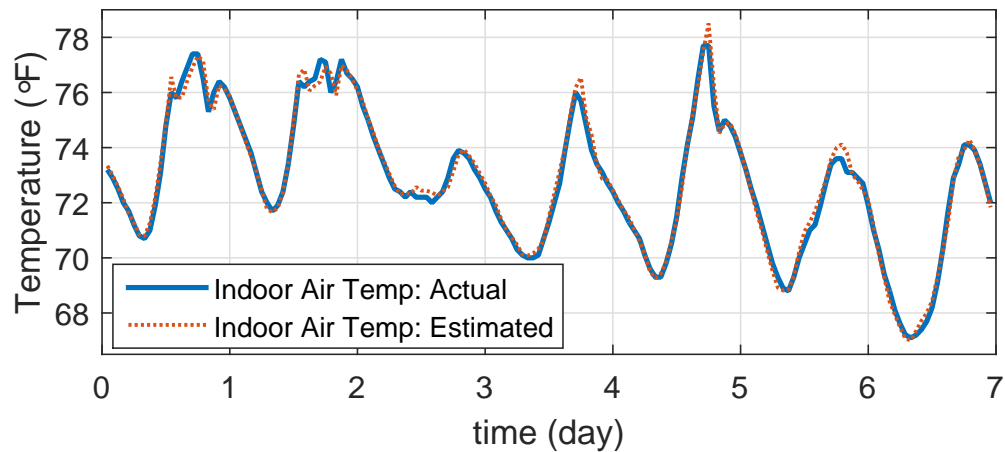


Figure 14. Comparison of actual and estimated indoor air temperature in the RBSA home in Tenino, Washington. The house model was trained using 2 weeks' historical data and evaluated using another week's testing data.

Table 6. System identification results in the five RBSA homes using 2 weeks' training data and 1 week's testing data.

ID	Location	Season	RMSE (°F)	R ²
1	Tenino	Cooling	0.27	0.99
		Heating	1.32	0.94
2	Eugene	Cooling	0.39	0.86
		Heating	0.51	0.89
3	Seattle	Heating	0.69	0.94
4	Emmett	Heating	0.49	0.87
5	Tacoma	Heating	0.58	0.97

the differences between values predicted by a model and the actual values. Smaller RMSE values indicate higher accuracy of the prediction model. R² is a dimensionless statistical measure of goodness-of-fit for a linear model. R² is always between 0 and 1; in general, the higher R² is, the better the model fits the data. Table 6 indicates that most house models have RSME values smaller than 1°F and R² values greater than 0.85.

3.5 Comments on Machine Learning

In this work, the data available to us had some limitations as identified above. These included the relatively-lower temporal resolution of measurements needed for energy efficiency evaluations, as compared to the higher-resolution measurements that are typically available from IoT devices, the inconsistency of data collected across sites, and some important data that were missing from many houses' data.

Secondly, the historical and laboratory data was limited in terms of its ability to reflect the fully-complex real-world of an actual building. This is particularly true in terms of calibrating models against the expected state changes that occur when control is applied (for historical data), the homeowner/occupants' actual operation of end uses (in both cases) and the homeowners/occupants' perceived comfort and other benefits (in both cases).

In the next phase of development we expect to take **foresee** into pilots in real buildings, where the full richness and complexity of data that reflects real buildings, real weather, and real occupants will permit substantial improvement in machine learning implementations.

4 Optimization

Scheduling and coordination of residential equipment is formulated as a multiobjective MPC problem. This section first introduces the objective functions in the optimization problem, using the models and constraints described in Chapter 3.³

4.1 Objective Function

The goal of a MPC is to find the optimal control strategy U_{opt} that minimizes the overall cost:

$$U_{opt} = \underset{U}{\operatorname{argmin}} J(x_0, U) \quad (4.1)$$

where the overall objective function

$$J(x_0, U) = \sum_{t=0}^{H-1} B\Phi^T(x(t), U(t)) \quad (4.2)$$

is a linear combination of multiple objectives over the prediction horizon H . J is a function of the initial equipment status x_0 and control actions U taken over the entire prediction horizon. $B = [\beta_{cmft}, \beta_{cost}, \beta_{carbon}, \beta_{conv}, \beta_{dgrd}]$ is a set of weighting factors derived from user preferences for individual objectives, $\Phi = [\Phi_{cmft}, \Phi_{cost}, \Phi_{carbon}, \Phi_{conv}, \Phi_{dgrd}]$ is a set of objective functions for calculating the cost of individual objectives, and $x(t)$ is a set of time-dependent variables representing the equipment status and is updated at each prediction step by the control actions $U(t)$ applied to the equipment.

The optimization problem is subject to the following constraints:

- Battery system dynamics,
- Single-zone house model heat balance,
- Water heater dynamics,
- Schedulable appliance constraints,
- Limits on PV generation and curtailment, and
- Overall home energy balance.

These constraints and related appliance models are explained in greater detail in Chapter 3. The subscripts in the weighting factors β and objective functions Φ indicate the five types of objectives that are considered in the paper. The objectives are represented, for the purpose of optimization, as costs or "disutilities" so the overall weighted cost can be minimized.

4.1.1 Thermal Discomfort

Thermal discomfort is represented by a function for temperature deviation between the actual and desired value. Indoor air temperature and hot water temperature during showers are considered in the study:

$$\beta_{cmft} = [\beta_{air}, \beta_{water}] \quad (4.3)$$

$$\Phi_{cmft} = [\Phi_{air}, \Phi_{water}] \quad (4.4)$$

and the thermal discomfort is defined as follows:

$$\beta_{air}\Phi_{air}(T_{air}^{in}) = \beta_a^+ [T_{air}^{in} - T_{air}^{max}]_+^2 + \beta_a^- [T_{air}^{min} - T_{air}^{in}]_+^2 \quad (4.5)$$

$$\beta_{dhw}\Phi_{dhw}(T_{dhw}) = \beta_w^+ [T_{dhw} - T_{dhw}^{max}]_+^2 + \beta_w^- [T_{dhw}^{min} - T_{dhw}]_+^2 \quad (4.6)$$

³Note that much of this chapter has been previously published in (Jin et al. 2017) and (Garifi et al. 2018b), and we duplicate that prior work here for the reader's convenience.

where T is the temperature of the variable specified in the subscripts and $[\]_+$ indicates $\max(0, \cdot)$, so the costs of Φ_{air} and Φ_{dhw} are zero unless the indoor air temperature T_{air} in or the hot water temperature T_{dhw} is outside the desirable bounds.

User preferences on thermal comfort may be different if the temperature becomes higher or lower than the desired value. For example, the weighting factor for air temperature disutility, β_{air} , is equal to β_a^+ when T_{air}^{in} exceeds the upper limit T_{air}^{max} and is equal to β_a^- when T_{air}^{in} falls below the lower limit T_{air}^{min} . Similarly, the weighting factor for the water temperature disutility β_{dhw} is equal to β_w^- and β_w^+ when the water temperature T_{dhw} is too cold or too hot, respectively.

4.1.2 Energy Cost

Energy cost represents the cost of electricity from the utility grid:

$$\Phi_{cost}(P_{grid}) = c_e P_{grid} \quad (4.7)$$

where c_e is the unit cost of electricity and could accommodate various pricing schemes, such as fixed pricing, time-of-use pricing, and real-time pricing. P_{grid} denotes the power drawn from the grid and is positive when the power flows from the grid to the home. Demand charge, DR incentives, and feed-in tariffs may also be incorporated in Φ_{cost} with minor modifications in the objective function.

In this project, the energy cost relates solely to electricity. In the future, this could be expanded to include other utility costs, such as natural gas, water, fuel oil, propane, or purchased heat.

4.1.3 Environmental Impacts

In this project, environmental impacts are restricted to the negative outcome resulting from curtailing rooftop PV in order to achieve other goals. Environmental impacts could be expanded to include other societal factors in the future, such as grid resiliency and water consumption by the power generation sector.

$$\Phi_{pv} = U_{pv}(t)P_{pv}^{max} + [P_{grid}]_+ \quad (4.8)$$

where P_{pv}^{max} is the maximum PV generation when there is no curtailment, and U_{pv} is the curtailment of PV generation that should be minimized whenever possible to maximize environmental benefits. $[P_{grid}]_+$ is the amount of power backfeeding to the grid and is equal to the surplus PV power because battery backfeeding is disabled by Eq. 3.22. Minimizing this objective would encourage the generation and self-consumption of clean PV energy.

4.1.4 User Inconvenience

User inconvenience is represented by the delayed cycles of schedulable appliances. Dishwashers and clothes dryers, which have the highest DR potential among all intermittently schedulable appliances, are considered in the study:

$$\beta_{conv} = [\beta_{dw}, \beta_{cd}] \quad (9)$$

$$\Phi_{conv} = [\Phi_{dw}, \Phi_{cd}] \quad (10)$$

$$\Phi_{dw} = I_{dw} ET_{dw}^t U_{dw}(t) \quad (4.11)$$

$$\Phi_{cd} = I_{cd} ET_{cd}^t U_{cd}(t) \quad (4.12)$$

where I_{app} is a binary variable indicating if a cycle is scheduled, ET_0^h defines the elapsed time from the scheduled start time of the cycle, and U_{app} is the control variable representing duty cycles of the appliance.

In other words, the cost increases as a cycle is delayed for a longer period of time. For simplicity, $ET_{app}^t = t/H$ is used in this project, so the inconvenience is proportional to the time delay of a cycle.

4.1.5 Equipment Degradation

Equipment degradation is additional wear and tear caused by the control actions. This objective is not explicitly evaluated in the user preference survey, but it is included in the overall objective function because a HEMS should not cause excessive equipment degradation.

The battery is vulnerable to degradation if not controlled properly; cycling fade is one of the main aging factors of lithium-ion batteries. The Ampere-hour throughput (Ah), which is the energy delivered or stored by the battery in terms of the current throughput, is used to represent the cycling fade. Other aging factors such as the calendar fade are not considered in this paper. Battery degradation is modeled as:

$$\Phi_{dgrd} = \alpha_{bat}^{Ah} (|P_{bat}^{dis}| + |P_{bat}^{ch}|) \quad (4.13)$$

where α_{bat}^{Ah} is the coefficient of the battery cycling fading, P_{bat}^{dis} is the battery discharging power, and P_{bat}^{ch} is the battery charging power.⁴

4.2 Solving the MPC Problem via Quadratic Programming

The optimization problem formulated in Section 4.1 is a quadratic programming problem in which the objective function is quadratic and the constraints are linear. Unlike mixed-integer programming or dynamic programming problems that often require commercial solvers and/or high-performance computers, the optimization problem in our approach can be solved by lightweight convex programming solvers capable of being implemented on an embedded or resource-constrained platform. CVX, a MATLAB-based modeling system designed for specifying and solving disciplined convex programs (Grant and Boyd 2014), was used in this paper. The prediction time step, the time in between control decisions, was 30 min and therefore the **foresee** controls were updated every 30 min using the CVX solver. The prediction horizon used in this paper was 16 time steps, or 8 hours, though we have used timesteps and horizons of both shorter and longer durations in related work. The appliance simulation was implemented at 1-min intervals to interpret the load control from MPC and update the appliance status; this process is done in real time in actual appliances.

The solution to the optimization problem defined in Eq. 4.1 is a set of control variables that minimize the objective function over the prediction horizon $[0, H - 1]$:

$$U_{opt} = U(t)^* \quad (4.14)$$

where $U(t) = [U_{bat}^{ch}, U_{bat}^{dis}, U_{pv}, U_{hvac}^h, U_{hvac}^c, U_{wh}^{low}, U_{wh}^{up}, U_{dw}, U_{cd}]$ is a set of control variables that updates the equipment status $x(t) = [SOC, P_{pv}, T_{air}^{in}, T_{dhw}, P_{dw}, P_{cd}]$ at each step within the horizon, and * denotes the optimal solutions.

The battery inverter and PV inverter are able to perform variable output power control, whereas water heaters and HVAC are simple on/off appliances. For these on/off appliances, the continuous control variables are interpreted as the duty-cycle signals, which can be converted into time-dependent on/off control signals over the prediction step. Appliances with compressors such as HVAC also have constraints on the minimum off time or maximum number of daily cycles to protect the equipment against abuse. These constraints are also implemented when converting the continuous control variables to real-time control commands.

Results from performing this optimization on simulated buildings are presented in (Jin et al. 2017). The Home Battery System was also exercised in the laboratory under a set of scenario tests, as described in the next two chapters.

4.3 Chance-Constrained Optimization

Incorporating uncertainty into the MPC problem can be performed by implementing chance, or probabilistic, constraints. Chance constraints seek to satisfy a constraint with a prescribed probability; for example, the upper limit on the deterministic constraint in (3.3) can be formulated as a function of random variables (for example, outdoor air temperature, solar irradiance, etc.) and thus rewritten as a probabilistic constraint:

⁴Minimizing this objective helps to avoid simultaneous charging and discharging in the battery model. Additional work on this specific outcome can be found in (Garifi et al. 2018a) which is currently under journal review for publication.

$$Pr(T_{air}^{in} \leq T_{air}^{max}) \leq 1 - \varepsilon \quad (4.15)$$

where ε denotes the maximum violation probability; i.e., the air temperature must not exceed its upper limit with probability $1 - \varepsilon$. If the random variable within the chance constraint is Normally distributed, an exact analytical reformulation of the chance constraint can be used; otherwise, another approach must be taken, such as scenario-based approaches or distributionally robust approaches (Dall’Anese, Baker, and Summers 2017). The work on integrating chance constraints into the **foresee** code base is currently underway; preliminary results can be seen in (Garifi et al. 2018b) where the HEMS aims to satisfy a DR request with a certain probability as well as maintain comfort bounds with a high probability.

5 Scenario Experiment Design

5.1 Introduction

The Scenario Test Plan defines realistic and reasonable scenarios in which the HBS would be expected to deliver meaningful benefits to the homeowner/occupant, to the utility, and/or to society through infrastructure support or environmental impact reduction. Because the HBS operates in the cyberphysical domain - at the intersection of information technology and the physical world - the scenarios include both a set of operational scenarios and a set of cybersecurity scenarios. The operational scenarios are a combination of grid scenarios, weather (location) scenarios, home scenarios, and occupant scenarios. Cybersecurity scenarios are independent of these physical instantiation scenarios, and will be executed separately from the operational scenarios.

The primary purpose of the scenario experiments is to validate the Home Battery System's controller, comprised of the **foresee** automation software and the cybersecurity layer developed by ETAS/ESCRYPT. These software components are integrated onto a single microcomputer that will serve as the experimental test article.

The HBS hardware is comprised of this controller, communicating equipment includes a Bosch heat pump, refrigerator, dishwasher, washing machine and dryer; an ecobee™ thermostat, a Fronius™ photovoltaic inverter, an A.O. Smith™ water heater, and an Eguana residential battery storage system. The HBS is operated in the laboratory to use this set of equipment, but could be easily adapted to include other generation, storage, sensors, more or less appliances/equipment, and/or different control objectives. A diagram of the HBS is shown in Figure 1.

All experimental equipment is physically installed in the Systems Performance Laboratory (SPL) in the Energy Systems Integration Facility (ESIF) at the National Renewable Energy Laboratory (NREL). All experiments will be performed using established methods for driving the equipment in realistic ways, using laboratory capabilities such as a PV simulator and a hardware-in-the-loop system for operating the heat pump (Pratt et al. 2017).

In addition to the different operational scenarios that will be evaluated, the laboratory phase of this project will also allow us to explore additional research questions that will be valuable to the later phases of field deployment. Research questions guiding the laboratory work described here are given in the following section.

5.2 Experimental Setting

NREL has established an experimental test bed for the Home Battery System scenario tests in Systems Performance Lab (SPL) within the Energy Systems Integration Facility on its campus in Golden, Colorado. The project's Stage Gate 2 milestone demonstrated laboratory readiness of the test articles, which include a residential heat pump with ecobee4 thermostat, an electric water heater controlled via CTA-2045 module, simulated solar PV array with a SunSpec-enabled PV inverter, 12kWh/5kW SunSpec-enabled home battery, and finally a set of HomeConnect-enabled appliances: dishwasher, refrigerator, washing machine, and dryer. NREL has also installed a Neurio residential power meter. These test articles are all operated from a microcomputer which runs the cybersecure foresee control software.

A residential Wi-Fi router is installed to provide standard home networking, and it operates a virtual network within ESIF which connects to external sites with lower cybersecurity restrictions than NREL's typical enterprise network. This is representative of a normal house's internet connection.

SPL hosts an embedded data acquisition system based on National Instruments hardware, specifically the EtherCAT RIO deterministic platform. All appliances listed previously have dedicated power metering, with real and reactive power sampled simultaneously at 5Hz with uncertainty below 0.2% of reading.

Similarly, the data acquisition system monitors water and air temperatures and flowrates at 5Hz, using special limits of error T-type thermocouples and high-accuracy pitot tubes. All sensors and sampling intervals far exceed the required accuracy for the experiments in the following because the experiments are intended to showcase representative operation in example use cases rather than provide highly-accurate predictions which might be appropriate for regulatory purposes.

Additional data sources including the appliances and equipment under test, weather data, and weather forecasts from web resources, typically lack stated uncertainties.



Figure 15. Appliances used in the laboratory experiments. Clockwise from top left: Neurio power meter (Wi-Fi), AO Smith electric water heater (CTA-2045), Bosch air-source heat pump, ecobee4 thermostat (ecobee cloud API), Eguana home battery (Modbus), SMA SunnyBoy inverter (SunSpec), Bosch clothes washer and dryer (HomeConnect cloud API), Bosch refrigerator (HomeConnect cloud API), Intel NUC minicomputer running foresee, and Bosch dishwasher (HomeConnect cloud API).

5.3 Operational Scenarios

Operational scenarios describe the physical services provided to the utility/aggregator, to the homeowner, and to society. Any operational scenario under which **foresee** is operated can be described as a combination of one scenario from each of the following categories: Grid, Climate, Residence, Homeowner, and Preference. Eight operational scenarios were completed experimentally, and three via simulation.

Table 7. Operational Scenario Definitions

Operational Scenario	Grid Scenario	Climate/Residence Scenario	Homeowner Scenario	Preference Scenario	Baseline Studies
OS1a	GS1-T4	CS1v	HS4	PS4	BAU
OS1b	GS1-T4	CS1v	HS4	PS5	
OS1c	GS1-T4	CS1v	HS4	PS6	
OS2*	GS3	CS1v	HS2	PS1	BAU, EE
OS3*	GS4	CS1v	HS2	PS1	BAU, EE
OS4	GS1-T4	CS2n	HS1	PS1	BAU
OS5	GS4	CS2n	HS1	PS1	BAU, EE
OS6	GS7	CS2n	HS1	PS1	BAU, EE
OS7	GS1-T4	CS5v	HS1	PS1	BAU
OS8	GS2	CS5v	HS1	PS1	BAU, EE
OS8-2	GS1-T2	CS5v	HS1	PS1	BAU
OS9*	GS4	CS5v	HS1	PS1	BAU, EE

*OS2, OS3, OS9 were performed in simulation due to budget and schedule constraints.

Each operational scenario in Table 7 is comprised of one Grid, Climate/Residence, Homeowner and Preference scenario. Those are each defined in the next sections, and the operational scenario was selected from the broader list of scenarios that could potentially be explored. This was a large parametric space, and the Technology Innovations project is not designed to exhaustively explore it. Instead we attempted to define operational scenarios which were representative of current or future use cases relevant to BPA, as well as to technology adoption trends underway today. The operational scenarios resulted from discussions by all project partners, and approved by BPA and DOE prior to experiments beginning.

- OS1a-c are designed to demonstrate how the Home Battery System meets a variation in homeowner preferences/goals under normal operation.
- OS2-3 are designed to demonstrate Portland winter load-shed and shoulder-season load-add demand response service.
- OS4-6 are designed to demonstrate response to a time of use rate, renewable oversupply, and solar self-consumption for a temperate climate (Honolulu).
- OS7-9 are designed to demonstrate cold-climate summer and winter load-shed demand response service.

Each operating scenario has one or two relevant baseline experiments against which the Home Battery System performance is compared.

- BAU A “Business as Usual” baseline study will be performed to establish how the home would operate without the **foresee** controller - all equipment will operate on its own internal controls and the solar inverter will backfeed power where possible, or curtail it if required by the utility contract.
- EE For cases where a grid service is called, an energy efficiency (EE) baseline study will also be performed to establish how the **foresee** controller would operate the home absent any grid service call.

5.4 Grid Scenarios

We have defined nine grid scenarios to provide a broad sample of potential use cases in which the HBS would be expected to operate.

- GS1 Normal Day: No demand participation is called and no other constraints are imposed by the utility.
- GS2 Peak Heat Day and Energy Supply*: The grid is severely strained in capacity due to extreme afternoon hot weather, and requires load reduction.
- GS3 Peak Cold Day and Energy Supply: The grid is severely strained in capacity due to extreme morning cold weather, and requires load reduction.
- GS4 Shoulder Season Oversupply: The grid experiences oversupply and requires addition of load to avoid shutting down base load or other negative supply-side impacts.
- GS7 Solar Self-Consumption: Utility interconnection agreement requires the home to avoid backfeeding photovoltaic generation.

** As originally defined in the SGIP Transactive Energy Application Landscape Scenarios whitepaper (SGIP 2016), and summarized in the following.*

Under each of the Grid Scenarios, one or more relevant referenceable tariff structures or participation incentives will be used to define a financial basis for residential load participation.

5.4.1 GS1: Normal Day

The GS1 grid scenario is an uneventful early summer day. No demand-side participation is called and no constraints, such as limits on energy backfeed, are imposed by the utility.

Five tariff rates are defined:

- GS1-T1 Flat rate, net metered - at the location's 2016 average cost per kWh (EIA 2017).
- GS1-T2 Flat rate, feed-in-tariff - at location's 2016 average cost per kWh (ibid) and with a 75%-reduced feed-in tariff paid for any exported energy.
- GS1-T3 Time of Use rate, net metered - using the location's largest serving utility rates, if available, or if not Southern California Edison's TOU-D-A rate (SCE 2017) scaled to better match the location's annual flat rate.
- GS1-T4 Time of Use rate, feed-in-tariff - TOU as defined in GS1-T3, with feed-in-tariff as defined in GS1-T2.
- GS1-T5 Residential Demand Rate - kW and kWh rate, patterned on Xcel Energy's residential rate (Xcel 2018) currently being piloted.

5.4.2 GS2: Peak Heat Day

The GS2 grid scenario occurs on a day of extreme heat. Because both transmission/distribution and generation capacity are constrained due to the afternoon heat, additional demand shedding is required to avoid a brownout or rolling blackouts (INC event). The Home Battery System is subscribed to an aggregator service, and is called to participate in an acute event, 6 hours in duration, with the participation called only 4 hours in advance. An incentive is provided for each kWh delivered during the event, and the aggregator charges a penalty of 250% of the incentive for underdelivering on the forecasted unit service.

5.4.3 GS3: Peak Cold Day

The GS3 grid scenario occurs on a day of extreme cold. Because transmission/distribution capacity is constrained due to the cold winter morning, additional demand shedding is required to avoid a brownout or rolling blackouts (INC event). The Home Battery System is subscribed to an aggregator service, and is called to participate in an acute

event, 4 hours in duration (6 - 10 a.m.), with the participation called 12 hours in advance. An incentive is provided for each kWh delivered during the event, and the aggregator charges a penalty of 250% for underdelivering on the forecasted unit service.

5.4.4 GS4: Shoulder Season Oversupply

The GS4 grid scenario occurs on a spring day. Demand is low due to the mild weather, so renewable hydropower and wind generation resources result in an oversupply (DEC event). The Home Battery System is subscribed to an aggregator service, and is called to add load. This event is 6 hours in duration, and occurs across mid-day (10 a.m.–4 p.m.) or nighttime (midnight–6 a.m.) hours. An incentive is provided for each kWh consumed during the event, taking the customer’s marginal consumption price near zero during this time period.

5.4.5 GS7: Solar Self-Consumption

The GS7 grid scenario occurs when a customer’s utility is unable to support energy export from new distributed energy resources. The Home Battery System is charged with managing the home’s resources to avoid exporting energy according to the interconnection agreement and rate structure. We plan to use Hawaii Electric Company’s self-supply tariff as a model for this scenario.

5.4.6 GS8: Emergency Loss of Supply

The GS8 grid scenario occurs when a utility incurs a sudden loss of supply, and commands emergency, unscheduled demand response. This could occur from a grid fault, such as a downed transmission line, or due to unscheduled outage at a generating facility.

CS1 Portland, Oregon. Calendar year 2014 weather was be used.

CS2 Honolulu, Hawaii. Calendar year 2011 weather was be used.

CS5 Spokane, Washington. Calendar year 2014 weather was be used.

5.5 Residence Scenarios

The physical construction of a residence defines, and constrains, what services HBS can deliver to homeowner and grid alike. We have leveraged prior work in selecting regionally-appropriate constructions—a single new home and a single vintage home—for each climate (Wilson et al. 2016).

For each residence scenario, we will first simulate the house in EnergyPlus. Equipment parameters will be used to closely-match what is installed in the laboratory. The EnergyPlus 6-minute output data will be used as training data for **foresee** to learn the equipment capabilities and the building’s physical response.

To initialize each experiment, the learned house state (temperature, water heater temperature, battery state of charge, etc.) will be aligned between the simulation and the physical laboratory as closely as possible. i.e. the Home Battery System will be initialized in abstract time before each experiment is conducted in real time.

For each climate zone, a designation of “n” is used to represent a new home built with efficient construction practice by today’s builders, and “v” is used to represent a vintage home sampled from the ResStock housing stock characterization for that region. For any location where recent utility-provided building surveys are available, they will be used instead of the ResStock data set.

5.6 Homeowner Scenarios

Homeowner archetypes will be developed based on the Building America House Simulation Protocol Wilson et al. 2014 and the Building America Hot Water Draw Schedule Generator (Hendron, Burch, and Barker 2014). A detailed description of these occupant profiles and schedules is provided in Appendix A.

Table 8. Homeowner Scenario Definition

	HS1	HS2	HS4	
Number of Occupants	2	2	4	
Occupancy Schedule	Low	High	Low	
Hot Water Use	High	Low	Med	
Cooling Setpoint	72	75	75	°F
Heating Setpoint	69	65	69	°F

5.7 Preference Scenarios

Initialization preference weights follow those used in prior Stage Gate, as shown in Table 9. PS1 was selected from the set of responses elicited using the SMARTER method developed in the project’s AMT surveys, and PS4-6 result from synthetic preference inputs. Preference weighting values, and service importance rankings from which they were derived are both shown. The lower half of the table provides the subjectively-ranked preferences, as a user would have provided them. The numerical weightings in the top half of the table are calculated from those subjective inputs. In the case of PS1, which is an actual input sample from the AMT surveys, the ranked inputs were not recorded.

Table 9. Homeowner Preference Definition

Criterion	PS1	PS4	PS5	PS6
Money	-0.180	-0.148	-0.0693	-0.0222
Carbon	-0.049	-0.0037	-0.0037	-0.0247
Dishes	0.590	0.0908	0.188	0.139
Laundry	0.630	0.241	0.0908	0.0908
Shower	1.060	0.0236	0.0301	0.0301
AirA	-0.180	0.0445	0.0445	0.297
AirB	0.082	-0.0154	-0.330	-0.0209
Most importance Least	Representative actual input sample from Y1 preference elicitation surveys	Saving Money	Air Comfort	Less Pollution
		Clean Clothes	Hot Water	Hot Water
		Hot Water	Clean Dishes	Air Comfort
		Air Comfort	Saving Money	Clean Dishes
		Clean Dishes	Clean Clothes	Clean Clothes
		Less Pollution	Less Pollution	Saving Money

5.8 Operational Scenario Performance Metrics

Each scenario is expected to operate for twenty-four hours. Daily summaries of total energy, discomfort, carbon, cost and other relevant summary metrics will be developed for comparison against baselines and other scenarios.

Performance metrics that apply to all operational scenarios include: Cost/financial benefit, Comfort impacts, Environmental costs, and Energy use in EE vs. Baseline. Some performance metrics are scenario specific and are only used to evaluate the relevant scenarios, such as:

- Daily cost of operating the home in EE or DR mode, in US \$, compared to the baseline cost;
- Solar firming and/or backfeed reduction, measured in % vs. the baseline; and
- DR prediction, measured in % error as defined in the following.

$$\epsilon_{DR} = \begin{cases} \frac{E_{actual} - E_{predict}}{E_{actual} - E_{baseline}}, & \text{when } E_{predict} > E_{actual} \\ 0, & \text{otherwise} \end{cases} \quad (5.1)$$

where ϵ_{DR} represents the committed, but undelivered, DR energy normalized by the actual energy reduction during the DR event interval, in %; E_{actual} is the measured energy consumption during the DR event interval, in kWh; $E_{predict}$ is the energy consumption **foresee** predicted it would use during the DR event interval, in kWh; and $E_{baseline}$ is the measured energy consumption of the home during the DR event interval, in kWh, without any whole-home controls

Note that $\epsilon_{DR} \geq 0$ indicates underdelivered DR services. Overdelivered DR service are not considered as an error, so corresponding ϵ_{DR} were set to 0 to not penalize overdelivery. The goal at project conclusion is for the HBS to provide the forecasted DR service with 90% confidence.

6 Energy, Cost and Comfort Results

6.1 Laboratory Testing Procedure

Of the 10 operational scenarios that were outlined in Section 5, three could not be performed in the lab due to constraints on time and budget. The three scenarios that were not evaluated in the laboratory were simulated and those results will be presented in Section 6.4.

Each scenario consisted of at least two different tests: a baseline test and an experimental test with **foresee** optimizing the controls for all the connected devices based on the user preferences. These tests are referred to as "EE cases" because energy efficiency is one of the main goals of the **foresee** controller. In some cases, the scenario included other tests that introduced demand response (DR) events on top of the overall optimization. All experiments (baseline, EE and DR) were run for 24 hours, beginning at midnight on the day indicated in Table A.4. Each case was configured with the same start up procedure to ensure that the appliances began in the same state.

The HVAC HIL system is run by starting an EnergyPlus simulation that will determine the initial conditions for the 'Outdoor' chamber and the 'Indoor' environmental chamber housing the thermostat (Pratt et al. 2017). The HVAC system was given time to reach its initial conditions before the test was started. The water heater was set to 125°F for all tests. Before each test began, water was drawn from the water heater until a heating element turned on. The water heater was allowed to fully reheat the tank before the test officially started. The washer and dryer were preloaded with a load of dry and wet clothes, respectively, before a test began. The dryer needs to be loaded with wet clothes to run in typical fashion. Once the loads of laundry were preloaded, the "remote start enabled" button was pressed on each device. The dishwasher did not need any special treatment, but its "remote start enabled" button was pressed for each test. The battery was charged to 2.4 kWh before each test. Generally, **foresee** would discharge the battery to below 2.4 kWh by the end of each test day, so charging was usually required before each new test.

6.1.1 Experiment Initialization

Once all the systems were at their initial conditions, different simulators were started to begin the test. A building was simulated using EnergyPlus to determine thermal state of the building, which drove load and actuation of the HVAC system. The appropriate 24-hour schedule for the hot water draw profile was started for the water heater. The PV simulator used to drive the PV inverter with a 24-hour insolation profile based on data from the test location was turned on. Lastly, **foresee** was started in whatever mode was needed. We used an agent in **foresee** to schedule the appliances for the baseline tests, even though **foresee** was not optimizing the control of appliances in the baseline tests.

6.1.2 Data Analysis

Most of the data for the laboratory tests was collected using a National Instruments-based data acquisition system at 1-second intervals. The 1-second data were averaged over a minute before analyzing the results. Individual appliance power data were collected for all the devices that were part of these experiments. Additional temperature and flow data were collected for the HVAC HIL systems and the water heater. The data acquisition system also has the ability to send control signals and was used to run the schedule of hot water draws and maintain the HVAC HIL set point temperatures. After each test was completed, data were exported from the data acquisition system and formatted for analysis.

A few corrections were applied to the data to ensure that the results are more realistic. The first correction involved correcting some unrealistic performance of loads in the lab and accounting for uncontrolled loads that did not exist in the laboratory. We used a bank of 30 light bulbs to represent the uncontrolled base load in a home, which added a constant power draw of about 500W at all times. We also ran the air handler continuously to ensure well-mixed air in the duct loop portion of the HVAC HIL system; normally the air handler would only run when cooling was called for. The laboratory experiments included all the major appliances in a home, but there are additional loads that we did not have in the lab, such as miscellaneous electric loads and cooking loads. We ran an annual simulation using the EnergyPlus models used to drive the HVAC HIL system and generated disaggregated simulated load data for each house in each climate. After we removed the appliance loads that we were controlling in the lab, we were left

with the total uncontrolled load, which was comprised of the range hood, bath exhaust fan, miscellaneous electric loads, lights, and the electric range, all based on data in the Building America Home Simulation protocol (Wilson et al. 2014). This simulated load varies hourly, in contrast to the base load from the experimental tests that never changed. We removed the lighting load from the lab data and added in the simulated uncontrolled load. We also modified the air handler power data to be on only during time when the condenser is on, as it would typically be controlled. This correction was applied to all laboratory data sets.

The second correction involves the battery and so does not affect the baseline cases. We started each test with the battery charged to 2.4 kWh, which is equal to 20% of the battery's energy capacity. However, **foresee** always discharged the battery beyond its initial level, ending each day at about 1.8 kWh. That difference of 0.6 kWh would look like additional energy savings because more energy was drawn from the battery than was added. To account for this discrepancy, we adjusted the total energy consumption for all tests that used by battery by 0.69 kWh, the amount of energy needed to charge the battery back to its initial state of 2.4 kWh. This correction is larger than 0.6 kWh because the battery does not charge perfectly efficiently (the measured charging efficiency was 0.875).

6.2 Metric Evaluation

Data from each test were collected and analyzed to evaluate different aspects of the test. The total energy consumed by the home was one key parameter, but other metrics were also used to more fully understand the impact of the **foresee** control.

6.2.1 Cost Analysis

As described in Section 5, each location used a different TOU rate based on existing utility rates for the location (or a similar location) and those rates were used to calculate the electricity rates for that day. The other consideration that affected cost was the feed-in tariff. For all cases, a feed-in tariff of 25% of the current electricity rate was applied to any energy that was back fed to the grid. Because **foresee** operated every half hour, we analyzed the energy use on a half-hourly basis to determine costs. For every half hour, if the net energy use was positive (drawing power from utility), the cost of that energy was calculated based on the TOU rate. If the net energy during a half hour was negative (backfeeding power to the utility), the cost credit of the generation was calculated using 25% of the current TOU rate. The total cost for each day is the sum of all the half-hourly costs and credits. For scenario OS8-2, a flat utility rate was used.

6.2.2 Comfort Analysis

In some cases, energy was saved at the expense of comfort. The different metrics to determine the impact on comfort are described as follows.

Air temperature comfort: For these scenarios, the set point temperature was always constant throughout the day. If the indoor temperature stayed within the bounds of $T_{setpoint} \pm 1^\circ\text{F}$, no discomfort was experienced. If the temperature strayed outside of those bounds, we calculated a discomfort parameter equal to the difference between the current temperature and the top or bottom of the comfort bounds (which ever was closer), multiplied by the duration. Discomfort was calculated on a minutely basis so the duration for each calculation was one minute. We calculated the accumulated discomfort for each day in terms of °F-hr outside of the comfort bounds, which gives a measure of how large the temperature deviation was and for how long. In some cases, cooling was only needed for part of the day, so discomfort was only calculated during the time that HVAC was active. In Table 10, discomfort was calculated for the full 24-hour period unless otherwise noted.

Water temperature comfort: The default setpoint for the water heater was chosen to be 125°F, based on the average setpoint listed in the Building America Home Simulation Protocol (Wilson et al. 2014). People generally do not take showers with 125°F water though—the hot water is mixed with cold water to achieve an average shower temperature of 110°F. Therefore, the minimum outlet temperature for the water heater is 110°F. If the outlet temperature dropped below 110°F, we calculated an energy penalty equal to the energy needed to raise the temperature of the too-cold water above 110°F.

6.2.3 Distributed Energy Resource Analysis

This home included a 5 kW PV system and a 12 kWh/5 kW battery, both of which are considered DERs. Controls provided by **foresee** can ensure that more of the energy generated by the PV system is used by the house—either by changing with different appliances run or by storing the energy in the battery. The battery can also be used to store energy from periods of time when the cost of utility power is cheaper or to provide grid services during certain times of day.

Reduced solar backfeeding: Because PV panels generate the most power when most people are not home, it is common for homes with PV to backfeed a significant amount of locally generated energy to the grid. This costs homeowners money because the energy backfed to the utility is bought at a rate that is 25% of the cost of electricity. The baseline case is used to establish the total energy backfed from the home without control. The associated experimental cases will be evaluated to see if **foresee** reduced the amount of energy backfed to the utility.

Battery throughput: Batteries are valuable energy storage devices, but there are efficiency losses when charging and discharging the battery and the battery degrades slowly overtime. Therefore, charging and discharging the battery has an associated cost, which was accounted for in cases with a home battery by adding a term in the objective function, $\beta_{bat} P_{bat}$. We used the coefficient $\beta_{bat} = 8.8 \times 10^{-5}/\text{kW}$. In the future, we expect to develop this model further to include a state- and rate-dependent (i.e. nonconstant) coefficient. Initial work on that effort was published in Raszmann et al. 2017, but completion of that formulation is left as future work. The amount of energy charged and discharged over each experimental cases will be calculated and in cases where the final state of charge of the battery is lower than the initial value, we will adjust the total energy use for that case by the amount of energy needed to return the battery to its initial state.

6.3 Laboratory Experiment Results

Of the 10 operational scenarios that were outlined in Section 5, three could not be performed in the lab due to constraints on time and budget. The results of the remaining scenario tests are summarized in the following. Full details are given in Appendix B. The three scenarios that were not evaluated in the laboratory were simulated and those results will be presented in Section 6.4.

Each scenario consisted of at least two different tests: a baseline test and an experimental test with **foresee** running to optimize the controls for all the connected devices based on the user preferences. These tests are referred to as "EE cases" because energy efficiency is one of the main goals of the **foresee** controller. In some cases, the scenario included other tests that introduced demand response events.

The laboratory results of the scenario tests are summarized in Table 10 and Table 11. Table 10 the energy, costs, and comfort impacts for all the energy-efficiency mode tests, relative to the baseline tests. Results from the demand-response tests, relative to the energy-efficiency cases and the baseline cases, are shown in Table 11. Some illustrative results from the scenarios are given in Sections 6.3.1 through 6.3.3. Detailed results for all the scenarios can be found in Appendix B.

In all the energy-efficiency mode tests, energy and cost were saved relative to the baseline case. For some tests, the energy savings came at the expense of comfort delivered by the air conditioner. The baseline cases incurred some level of discomfort, but the **foresee** control always increased the level of discomfort to save energy. Cost savings were generally higher than the relative energy savings because **foresee** was able to shift load out of the higher price periods of the TOU schedules and into the lower cost periods. The battery enabled much of the load shifting, as the battery can be easily charged at any time of day and discharged later to supply power to appliances that are not as flexible in their operation. The battery could also be used to absorb more of the local solar generation, but that was only done in cases when the lowest energy cost period coincided with the times that the PV system was generating power. See Appendix B for more details on the results from the energy-efficiency mode tests.

Of the five demand response cases outlined in Section 5, two of them were evaluated during the laboratory tests: a load-shed DR case on a summer day in Spokane, WA and a load day on a spring day in Honolulu, HI. For both scenarios, the EE case provided some load add or load shed because the TOU rate naturally shifted energy into or out of the demand-response period, but the DR case provided additional load shifting. Both DR cases resulted in energy savings relative to the baseline case, but the load-add DR case resulted in a small increase in energy consumption

relative to the EE case. Cost savings were achieved in each DR case relative to the baseline case and EE case, despite the fact that no additional incentives were provided to motivate the additional load shifting.

Table 10. Summary of Energy-Efficiency Mode Daily Energy and Comfort Performance

HBS Scenario	Baseline Case			Energy-Efficiency Case‡				
	Energy Use (kWh)	Cost (\$)	Discomfort (Air, °F-hr)	Energy Use (kWh)	Cost (\$)	Discomfort (Air, °F-hr)	Savings vs. Baseline	
							E (kWh)	Cost (\$)
OS1a ^a	0.36	0.93	1.6	-3.56	0.68	3.1	3.33	0.25 (24%)
OS1b ^a				-2.07	0.75	5.8	1.87	0.18 (19%)
OS1c ^a				-2.13	0.75	1.8	1.87	0.18 (19%)
OS2*	39.58	3.83	0.02	39.41	2.89	0.44	0.17	0.95 (25%)
OS3*	4.98	1.16	0.00	3.23	0.51	2.61	1.76	0.65 (56%)
OS4 ^b	12.75	7.04	6.6	8.70	4.11	16.0	3.37	2.93 (42%)
OS5	11.57	7.10	5.5	9.40	4.56	12.4	1.43	2.54 (36%)
OS6	14.10	7.56	7.0	10.52	4.37	12.1	2.90	3.20 (42%)
OS7	3.40	1.09	3.3	2.33	1.04	5.7	0.37	0.05 (4%)
OS8	17.85	1.91	11.6	15.38	1.71	29.7	1.70	0.21 (11%)
OS8-2†	17.85	2.29	11.6	9.76	1.56	46.4	8.10	0.73 (32%)
OS9*	60.30	5.20	0.08	59.31	4.69	0.25	0.99	0.50 (10%)

^a Accumulated discomfort was calculated between 3 p.m. and 12 a.m.

^b Accumulated discomfort was calculated between 12 p.m. and 12 a.m.

* OS2, OS3, OS9 were performed in simulation due to budget and schedule constraints.

† OS8 was repeated with flat rate to study the impact of tariff on control outcomes.

‡ Note that values may not directly equate, due to rounding and correcting for building state from experiment start to end.

Table 11. Summary of Demand Response Mode Energy and Comfort Performance, versus the Energy-Efficiency Case

Scenario	Daily DR Case			DR vs. Baseline			DR vs. Energy Efficiency		
	Energy (kWh)	Cost‡ (\$)	Discomfort (°F-hr)	Service (kWh)	Daily Savings		Service (kWh)	Daily Savings	
					E (kWh)	Cost (\$)		E (kWh)	Cost (\$)
OS2*	39.42	2.90	0.44	7.57	0.16	0.94 (25%)	0.03	-0.01	-0.01 (0%)
OS3*	19.29	1.57	2.61	24.79†	-14.30	-0.41 (-35%)	19.26†	-16.06	-1.06 (-207%)
OS5	10.13	4.30	13.14	10.15†	0.70	2.80 (40%)	6.26†	-0.73	0.27 (6%)
OS8	14.59	1.57	28.16	8.53	2.49	0.34 (18%)	1.28	0.80	0.14 (8%)
OS9*	60.48	4.91	0.29	8.74	-0.18	0.29 (6%)	11.08	-1.17	-0.22 (-5%)

* OS2, OS3, & OS9 were performed in simulation due to budget and schedule constraints.

† These values are load-add (DEC) DR service; all others are load-shed (INC).

‡ The incentives for participating in the DR program were not included in the cost calculation.

6.3.1 OS1 Scenario Results - Portland, OR, with Three Different User Profiles

The OS1 scenarios were designed to show the difference in performance for three different user preference profiles in **foresee**. All four tests, a baseline day and three days using **foresee** optimization, were run using weather for June 5, 2014. Homeowner profile HS4 defined the thermostat set point, the hot water draw profile, and the schedule for the washer, dryer and dishwasher (See Tables A.6, A.8). The three experimental cases for OS1 used three different homeowner preference profiles: PS4, PS5, and PS6. These profiles are intended to represent people that prioritize saving money, improved comfort, and maximizing use of renewable energy, respectively. In the following results, OS1a refers to profile PS4, OS1b corresponds to PS5, and OS1c uses profile PS6.

The results from the baseline day are shown in Figure 16, broken down by the different end loads. The blue shading indicates the relative price of the time-of-use price schedule, with the darkest shading representing the most expensive period of time. For more detail on the TOU schedule, see Table A.2.




Figure 16. OS1 Baseline experiment with all loads shown. TOU price schedule indicated by blue shading.

The experimental cases of OS1a, OS1b and OS1c have similar daily load profiles as shown in Figure 16, but with changes to when the appliances ran and the use of battery as needed. The accumulated energy for all the cases are shown in Figure 17. Due to the amount of solar energy generated by the simulated solar panels and the low cooling needs on a mild early summer day, the baseline case came very close to being net zero (generating enough energy to offset the energy drawn from the grid). The experimental cases used less energy than the baseline case, meaning that each case exported excess solar energy to the grid. A summary of the net energy consumption of each case and the impact on daily cost and solar generation are given in table 12.

The battery was not operated during the baseline case, but was controlled by **foresee** in the experimental cases. The Figure 18 shows how the battery was operated in cases OS1a, OS1b and OS1c. OS1a used the battery the least, while OS1b and OS1c used the battery in similar fashions. Battery throughput is listed in Table 13 for each case and the baseline case resulted in a nonzero throughput because the battery continues to draw power even when it is not being used.

Accumulated Energy for OS1 Scenarios: June day in Portland, OR




Figure 17. Accumulated energy from the OS1 baseline and the three scenarios with different preferences. TOU price schedule indicated by blue shading.

Table 12. Summary from OS1 Scenarios

Scenario	Total Energy (kWh)	Energy Savings (kWh)	Daily Cost (\$)	Cost Savings (%)
Baseline	0.4	-	0.93	-
OS1a (Cost)	-3.6	3.3	0.68	23.9
OS1b (Comfort)	-2.1	1.9	0.75	18.7
OS1c (Carbon)	-2.1	1.9	0.75	19.4

Battery Power Consumption for OS1 Scenarios: June day in Portland, OR




Figure 18. Battery power consumption for OS1 optimization cases. Battery charging is considered positive power and discharging is considered negative. TOU price schedule indicated by blue shading.

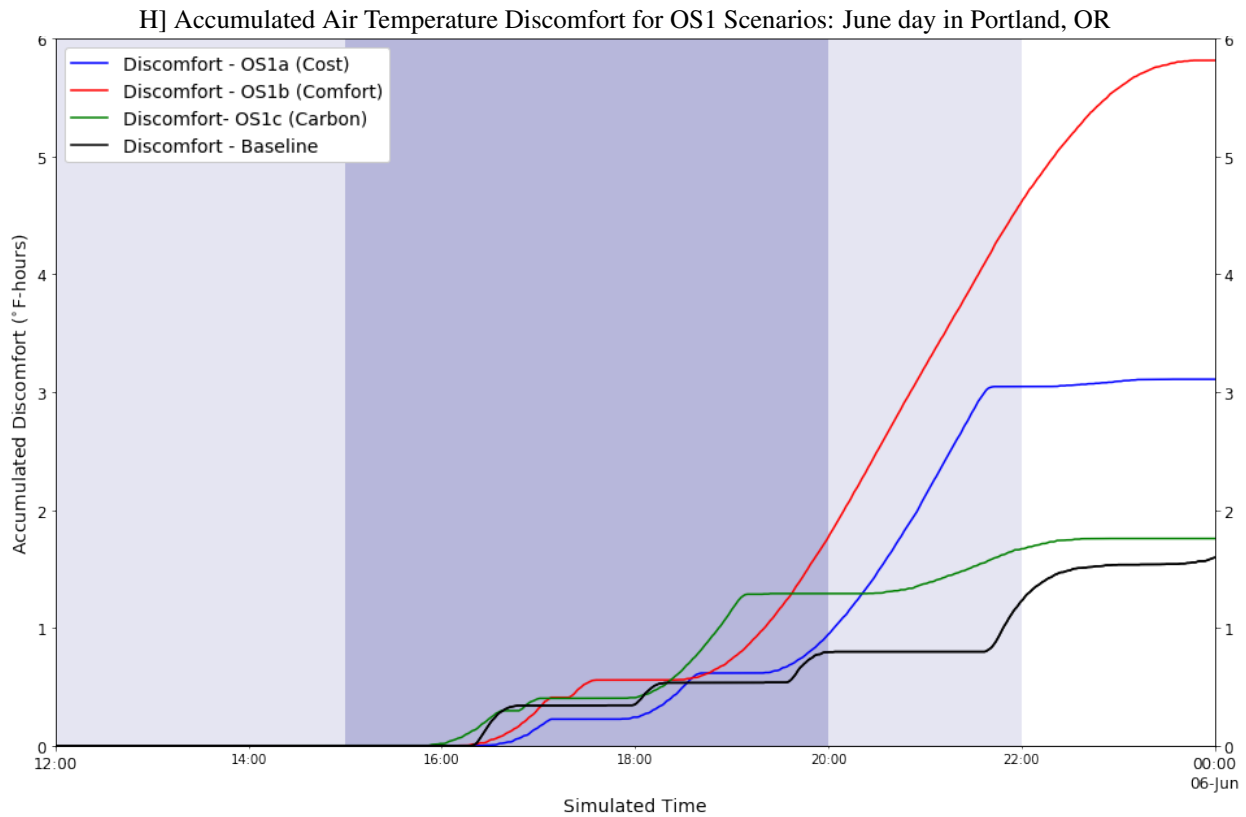


Figure 19. Accumulated Discomfort for OS1 cases. Cooling was not needed until 15:00. TOU price schedule indicated by blue shading.

Use of the battery in OS1b and OS1c enabled some cost savings because the battery charged during the lowest price period in the morning and then discharged during the higher cost periods later in the day. The battery did not facilitate more utilization of solar energy generation. In fact, in all experimental cases, the amount of energy exported from the house went up relative to the baseline because the PV-equipped home consumed negligible electricity in the baseline case, and **foresee** operated it to use less energy in all other cases.

The energy savings in each experimental case came primarily from a reduction in air conditioner energy use and a reduction in water heater energy use. In this scenario, all cases, including the baseline, experienced some level of discomfort due to the air temperature deviating from the comfort band, as shown in Figure 19. In this case, the discomfort metric was only calculated between 3 p.m. and midnight because cooling was not needed until the later afternoon. The house cools off significantly at night, but that deviation outside of the comfort bound was not counted in the discomfort metric because it occurred naturally. On the other hand, there was no impact to the delivered hot water temperature and the outlet water was always above the minimum temperature of 110°F.

Table 13. DER and Comfort Impacts from OS1 Scenarios

Scenario	Total Air Temp Discomfort (°F-hrs)	Water Heater Comfort Impact (kWh)	Solar Energy Backfeed (kWh)	Battery Throughput (kWh)
Baseline	1.6	0	18.2	0.3
OS1a (Cost)	3.1	0	18.6	3.1
OS1b (Comfort)	5.8	0	23.1	13.7
OS1c (Carbon)	1.8	0	22.5	13.4

A summary of the DER and comfort impacts is presented in Table 13. It is unclear why **foresee** selected control actions that resulted in the highest degree of discomfort in the OS1b case, where users' preferences indicated com-

fort as the highest priority. Because the other two use cases achieved significant energy and greater cost savings with lower comfort impacts, it appears there are sensitivities that are not represented in **foresee**'s model-predictive building and HVAC model. Thus there is room for improvement; this merits future research.

6.3.2 OS5 Scenario Results - Honolulu, HI, Under Shoulder Season Oversupply (Load-Add DR)

The OS5 scenarios demonstrate the ability of **foresee** to respond to a load-add demand response event on a spring day in Honolulu, HI with excess solar energy generation. The three tests, a baseline case, a case using **foresee** optimization, and a case using **foresee** with a load-add event between 10 a.m. and 4 p.m., were run using weather for March 20, 2011. Homeowner profile HS1 defined the thermostat set point, the hot water draw profile, and the schedule for the washer, dryer and dishwasher (See Tables A.6, A.8). Preference profile PS1 was used to drive the **foresee** controls.

The results from the baseline day are shown in the Figure 20, broken down by the different end loads. The blue shading indicates the relative price of the time-of-use price schedule, with the darkest shading representing the most expensive period of time. For more detail on the TOU schedule, see Table A.2.

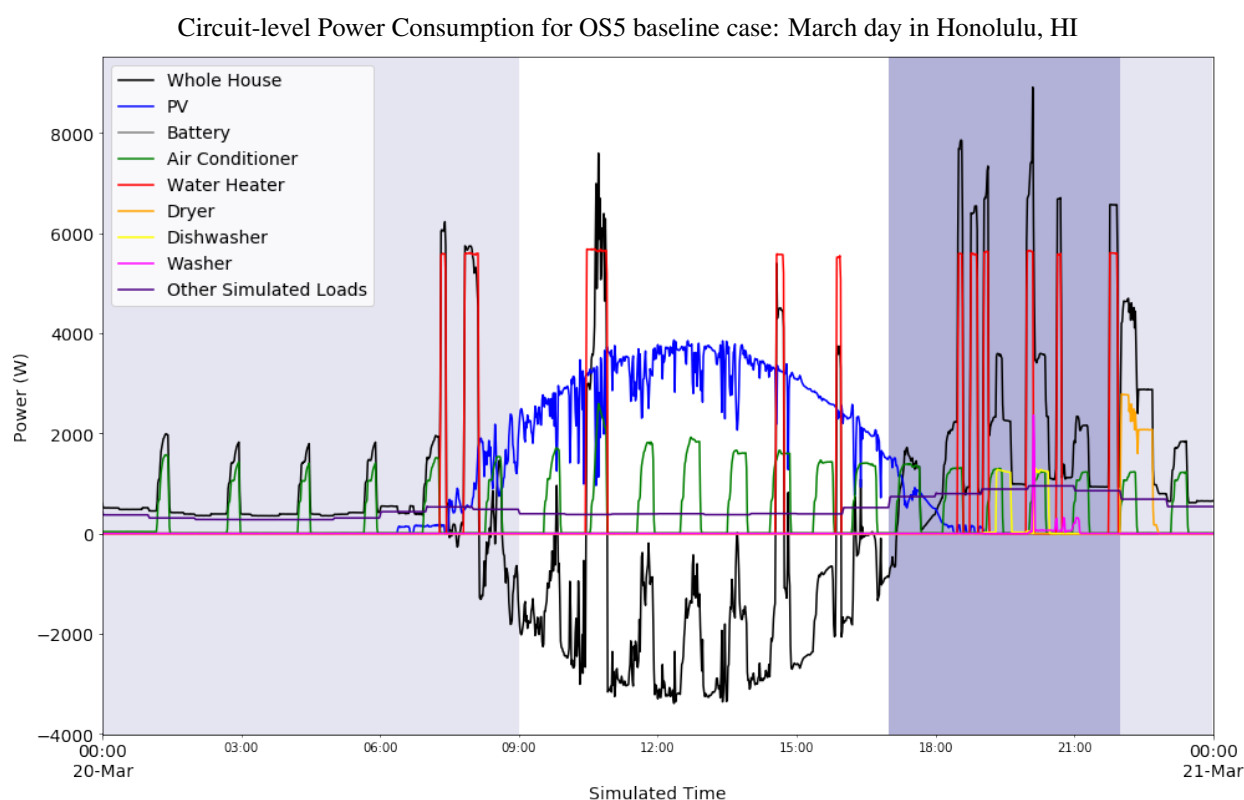


Figure 20. OS5 Baseline experiment with all loads shown. TOU price schedule indicated by blue shading.

Table 14. Summary from OS5 Scenarios

Scenario	Daily Energy (kWh)		DR Period (kWh)		Daily Cost (\$)	Cost Savings (%)
	Total	Saved	Energy Use	DR Service		
Baseline	11.6	-	-8.4	-	7.10	-
OS5 EE	10.1	1.4	-4.6	3.9	4.56	35.8
OS5 DR (load-add)	10.9	0.7	1.7	10.1	4.29	39.5

Weather in Honolulu is fairly constant throughout the year, but this day in March used less energy than the early June day from the OS4 cases. This scenario was designed to look at the ability of **foresee** to add load during a demand-

response event through the middle of the day. In Hawaii, the need for residents to use more energy during the middle of the day is driven by the high levels of distributed solar panels on residential buildings. The load-add demand response event overlapped with the lowest price period of the TOU schedule so the EE case saw an increase of energy use during the DR period, with the DR case further increasing the amount of energy used during the event period, as shown in Table 14. Despite the large increase in energy used during the DR period, both cases that used **foresee** control, decreased the overall energy use relative to the baseline case. Significant cost savings were seen in both cases with the EE case saving 35% and the DR case saving nearly 40% relative to baseline costs. The costs for all three cases were evaluated using the same TOU rate schedule, despite the fact that there is likely an additional financial incentive for customers to participate in demand-response event, so the cost savings for the DR case may be higher in real life. The DR case saw a small increase in energy use over the EE case, but resulted in a small decrease in cost savings.

The accumulated energy for all cases is shown in Figure 21, which shows the effect that **foresee** had, especially during the DR period, denoted by the green shading. The accumulated energy of the EE case increases towards the end of the lowest price period, while the DR case was able to maintain a nearly flat accumulated energy during the DR period, resulting in an increase of 10 kWh during that 6-hour window relative to the baseline case. Most of the increase in energy use for both EE and DR cases during the middle of the day came from the operation of the battery, shown in 22. In the EE case, the battery charging ramped up towards the end of the lowest price period and then discharged during the highest price period. The DR case charged the battery at a higher power level during the DR period and then discharged during the high price period. For both cases, the use of the battery took advantage of the lowest cost electricity from the grid as well as using more of the on-site solar generation, which is likely the main way that the **foresee** cases saved a significant fraction of daily costs. As in other cases, the EE and DR cases resulted in higher accumulated air temperature discomfort than the baseline case. A summary of the net energy consumption of each case and the impact on daily cost and solar generation are given in Table 14.

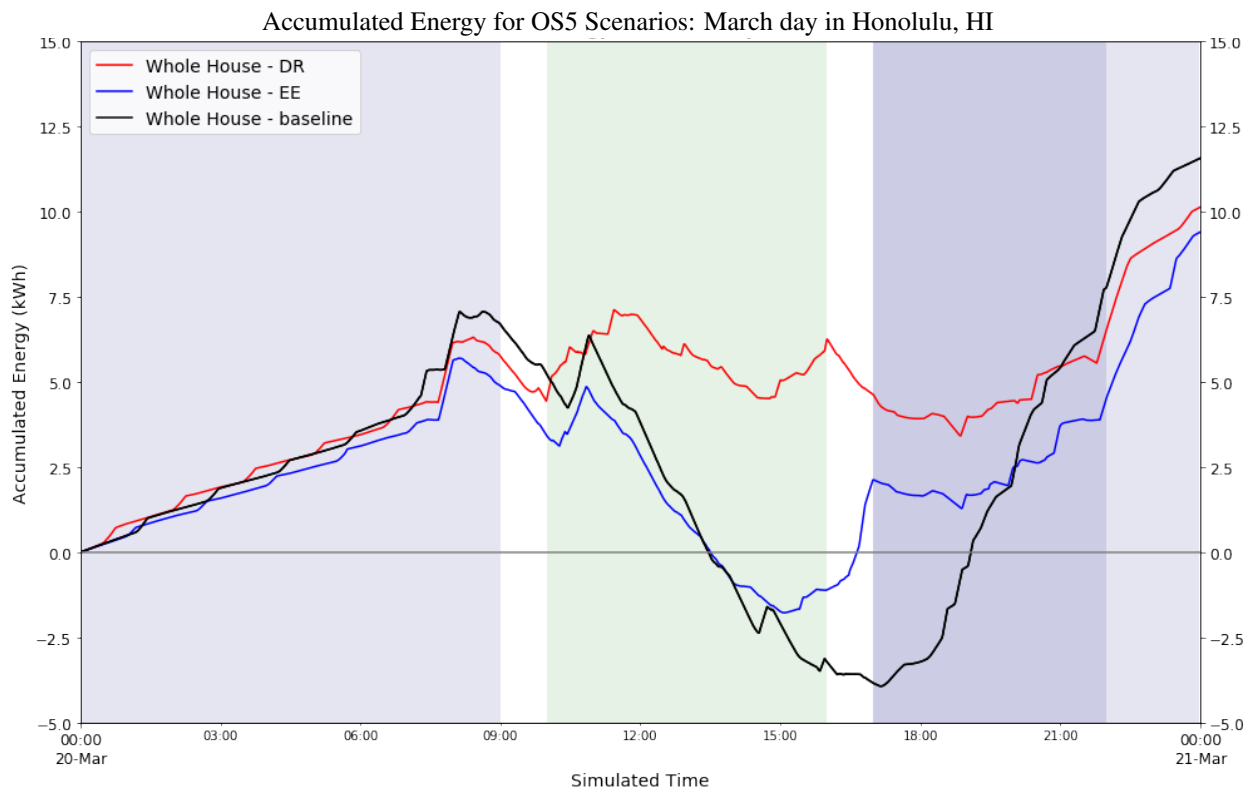


Figure 21. Accumulated energy for the OS5 baseline, OS5 energy efficiency, and OS5 demand response cases. TOU price schedule indicated by blue shading. Load-add DR period shaded green.

Battery Power Consumption for OS5 Scenarios: March day in Honolulu, HI




Figure 22. Battery power consumption for OS5 optimization cases. Battery charging is considered positive power and discharging is considered negative. TOU price schedule indicated by blue shading, with DR period indicated with the green shading.




Figure 23. Accumulated air temperature discomfort for the OS5 baseline, the energy efficiency case, and the load-up demand response case. TOU price schedule indicated by blue shading, with DR period indicated with the green shading.

Table 15. DER and Comfort Impacts from OS5 Scenarios

Scenario	Total Air Temp Discomfort (°F-hrs)	Water Heater Comfort Impact (kWh)	Solar Energy Backfeed (kWh)	Battery Throughput (kWh)
Baseline	5.5	0	14.5	0.3
OS5 EE	12.4	0	10.6	15.6
OS5 DR (load-up)	13.1	0	8.9	18.1

6.3.3 OS8 Scenario Results - Spokane, WA, Under Peak Hot Day (Load-Shed DR)

The OS8 scenarios demonstrate the ability of **foresee** to respond to a load-shed demand response event on a hot summer day in Spokane, WA. A total of four tests were conducted on this simulated day: a baseline case, a case using **foresee** optimization with a TOU utility rate, a case using **foresee** optimization with a flat utility rate, and a case using **foresee** with a load-shed event between 2 and 8 p.m. on top of a TOU rate. Each test case was run weather for July 29, 2014. Homeowner profile HS1 defined the thermostat set point, the hot water draw profile, and the schedule for the washer, dryer and dishwasher (See Tables A.6, A.8). Preference profile PS1 was used to drive the **foresee** controls.

The results from the baseline day are shown in the Figure 24, broken down by the different end loads. The blue shading indicates the relative price of the time-of-use price schedule, with the darkest shading representing the most expensive period of time. For more detail on the TOU schedule, see Table A.2.

Circuit-level Power Consumption for OS8 baseline case: July day in Spokane, WA




Figure 24. OS8 Baseline experiment with all loads shown. TOU price schedule indicated by blue shading.

This scenario used weather from the hottest day in Spokane in 2014, with a maximum outdoor temperature of 99°F. Hot summer days stress the electric grid because air conditioning loads are much higher and concentrated in the late afternoon and evening. This scenario was used to demonstrate **foresee**'s ability to provide load-shed during the hottest part of the day: 2–8 p.m. The demand response period overlaps with the highest price period of the TOU period so one additional test was run to see how the the energy-efficiency optimization varied when a flat utility rate

Table 16. Summary from OS8 Scenarios - vs. Baseline

Scenario	Daily Energy (kWh)		DR Period (kWh)		Daily Cost (\$)	Cost Savings (%)
	Total	Saved	Energy Use	DR Service		
Baseline - flat	17.9	-	7.6	-	2.29	-
Baseline -TOU					1.91	-
OS8 EE - flat	9.8	8.1	4.0	3.6	1.56	31.6
OS8 EE - TOU	16.2	1.7	0.3	7.2	1.71	10.8
OS8 TOU + DR (load-shed)	15.4	2.5	-0.9	8.5	1.57	17.8

is used versus a TOU utility rate. In all cases, energy was saved relative to the baseline case and cost was also saved in all cases. Also, all **foresee**-controlled cases led to a reduction of energy use during the DR period, with the DR case producing the largest reduction in DR period energy use, as shown in Table 16. The EE flat rate case led to the largest energy and cost savings. The DR TOU case resulted in nearly 18% cost savings, followed by the EE TOU case that produced over 10% cost savings. The costs for all three cases were evaluated using the same TOU rate schedule, despite the fact that there is likely an additional financial incentive for customers to participate in demand-response event, so the cost savings for the DR case may be higher in real life. The DR case delivered additional energy and cost savings relative to the EE TOU case (the EE flat rate case should not be compared to the DR case because the TOU rate drives the **foresee** controls).

The accumulated energy for all cases is shown in Figure 25. The two cases that used the TOU rate (EE and DR) both charged the battery before the medium price period began and discharged during the medium and high price periods (see Figure 26). Neither the baseline case or the EE flat rate case used the battery at all. The EE flat rate case was able to achieve significant energy savings by reducing the amount of air conditioning cycling, which also contributed to a high level of accumulated air temperature discomfort, as shown in Figure 27. The EE TOU and DR TOU cases also experienced higher levels of air temperature discomfort than the baseline case, but the air conditioner actually used slightly more energy. Both cases tended to let the temperature drift above the upper comfort bound and then would cool beyond the lower comfort bound, which led to additional cooling energy use. Most of the energy savings in the EE TOU and DR TOU cases came from energy savings from the water heater. The battery usage in those cases contributed to higher cost savings relative to the energy savings. See Table 16 for the summary of additional impacts in each case.

Table 17. DER and Comfort Impacts from OS8 Scenarios

Scenario	Total Air Temp Discomfort (°F-hrs)	Water Heater Comfort Impact (kWh)	Solar Energy Backfeed (kWh)	Battery Throughput (kWh)
Baseline	11.6	0	11.8	0.3
OS8 EE - flat	46.4	0	10.7	0.3
OS8 EE - TOU	29.7	0	13.0	16.9
OS8 TOU + DR (load-shed)	28.2	0	13.7	17.8

Accumulated Energy for OS8 Scenarios: July day in Spokane, WA

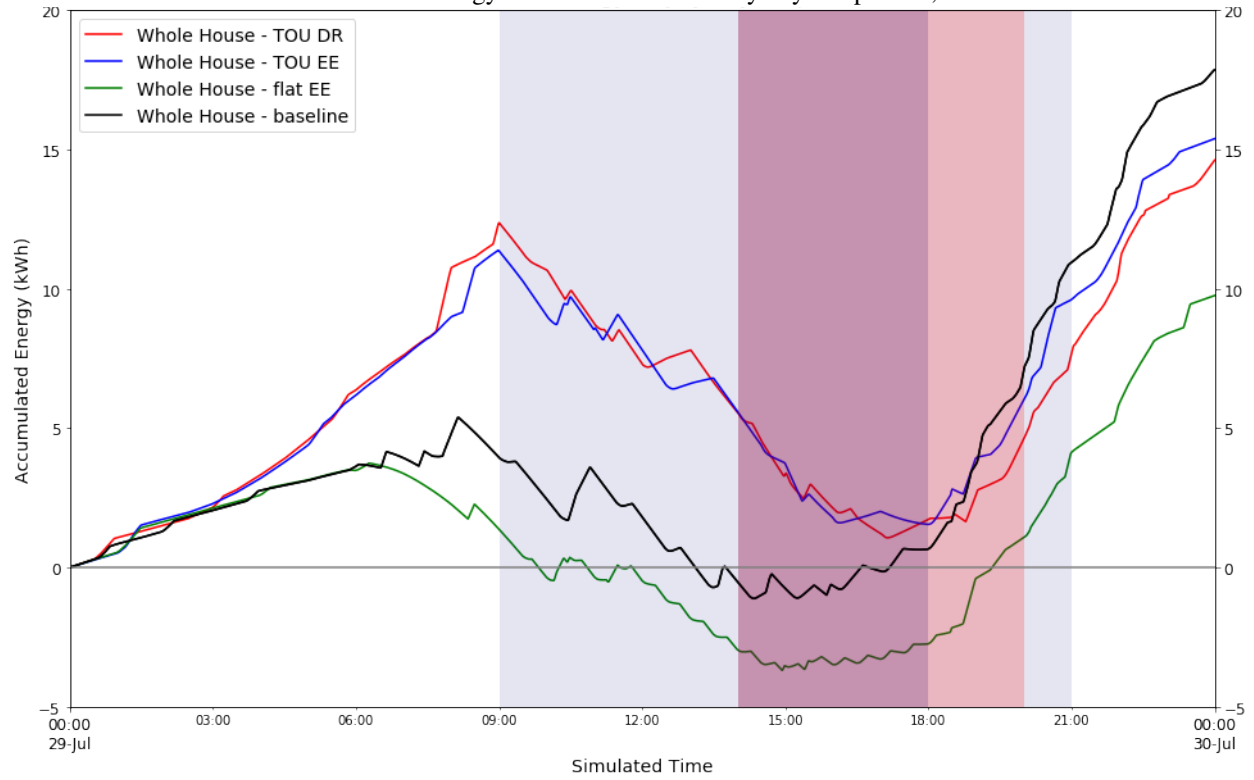


Figure 25. Accumulated energy for the OS8 baseline case, OS8 energy efficiency using a flat utility rate, OS8 energy efficiency with a TOU utility rate, and OS8 demand response with TOU pricing. TOU price schedule indicated by blue shading. Load-shed DR period shaded red.

Battery Power Consumption for OS8 Scenarios: July day in Spokane, WA




Figure 26. Battery power consumption for OS8 optimization cases. Battery charging is considered positive power and discharging is considered negative. TOU price schedule indicated by blue shading, with DR period indicated with the red shading.

Accumulated Air Temperature Discomfort for OS8 Scenarios: July day in Spokane, WA

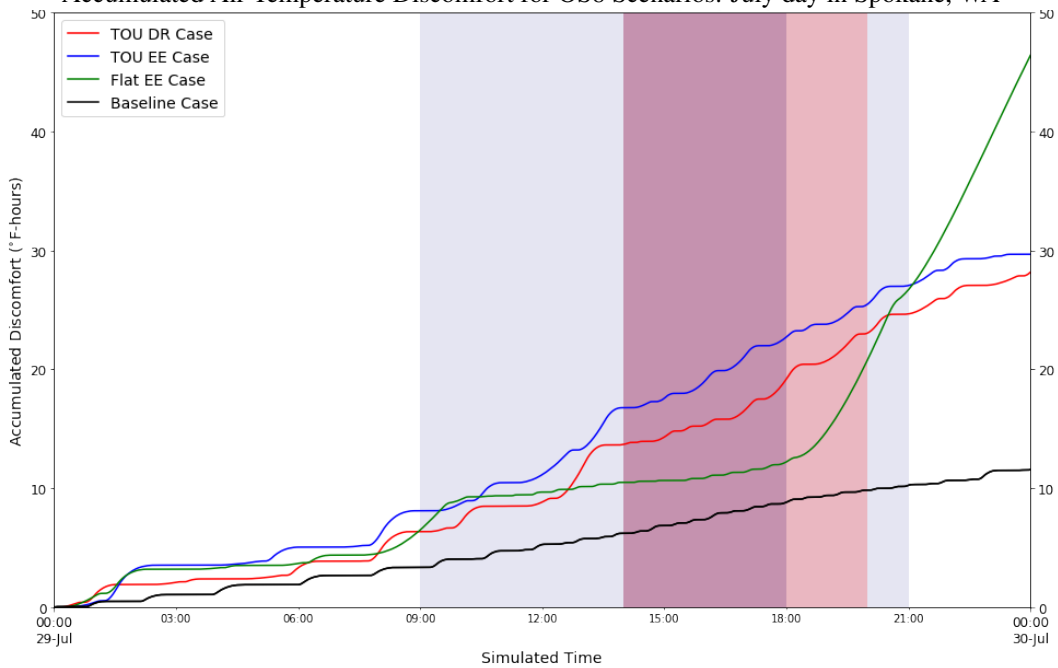


Figure 27. Accumulated air temperature discomfort for the OS8 baseline, EE with flat rate, EE with TOU, and loadshed DR cases. TOU price schedule indicated by blue shading, with DR period indicated with the red shading.

6.4 Simulation Experiment Results

6.4.1 OS2 Scenario - Portland, OR, Under Peak Cold Day (Load-Shed DR)

The OS2 scenario demonstrate the ability of **foresee** respond to a load-shed demand response event on a cold winter day in Portland, OR. A total of three tests were conducted on this simulated day: a baseline case, a case using **foresee** optimization with a TOU utility rate, a case using **foresee** optimization with a load-shed event between 6 a.m. and 10 a.m. on top of a TOU rate. Each test case was run using weather for February 5, 2014. Homeowner profile HS2 defined the thermostat set point, the hot water draw profile, and the schedule for the washer, dryer and dishwasher (See Tables A.6, A.8). Preference profile PS1 was used to derive the **foresee** controls.

This scenario used weather from one of the coldest day in Portland in 2014, with a lowest outdoor temperature of 21°F. Cold winter days stress the electric grid because space heating loads are much higher and concentrated in the early morning. This scenario was used to demonstrate **foresee**'s ability to provide load-shed during time period when heating loads are the highest. The demand response period overlaps with the medium price period of the TOU period. As shown in Table 18, the EE and DR cases have small energy savings compared to the baseline case, but the cost savings in the EE and DR cases are as high as 25%. It shows **foresee** responded to the TOU rates by shifting the loads to the time periods with lower energy prices. The daily energy savings and energy savings over the DR time periods for the EE and DR cases are very close because the TOU rates dominated the **foresee** operation and not much additional loads can be reduced during the DR time period.

The accumulated energy is shown in Figure 28. The accumulated energy of the EE and DR cases stayed almost constant during the DR time period (6–10 a.m.) whereas the baseline case has lower accumulated energy at 6 a.m. and ended up with higher accumulated energy at than the EE and DR cases at 10 a.m. Without **foresee**, the home has much higher solar energy export in the afternoon compared to the EE and DR cases.

As shown in Table 19, the EE and DR use cases also have similar DER and comfort impacts. Both cases have slightly higher air temperature discomfort but much lower energy backfeed than the baseline case. Battery was heavily used in the EE and DR cases; more than 26 kWh battery throughput occurred in both cases, which was about three times of the available battery capacity [i.e., $12 \text{ kWh} \times (85\% - 15\%) = 8.4 \text{ kWh}$].

Summary from OS2 Scenarios

Table 18. Summary from OS2 Scenarios - vs. Baseline

Scenario	Energy (kWh)		DR Period (kWh)		Daily Cost (\$)	Cost Savings (%)
	Total	Saved	Energy Use	DR Service		
Baseline	39.58	-	7.61	-	3.84	-
OS2 EE	39.41	0.17	0.06	7.54	2.89	24.8%
OS2 DR (load-shed)	39.42	0.16	0.04	7.57	2.90	24.6%

Table 19. DER and Comfort Impacts from OS2 Scenarios

Scenario	Total Air Temp Discomfort (°F-hrs)	Water Heater Comfort Impact (kWh)	Solar Energy Backfeed (kWh)	Battery Throughput (kWh)
Baseline	0.02	0.00	4.64	0.00
OS2 EE	0.44	0.00	0.16	26.21
OS2 DR (load-shed)	0.44	0.00	0.05	26.27

Accumulated Energy for OS2 Scenarios: February day in Portland, OR




Figure 28. Accumulated energy from the OS2 simulation for baseline, energy efficiency and demand response cases. TOU price schedule indicated by blue shading and red shading indicates load-shed DR period.

6.4.2 OS3 Scenario - Portland, OR, Under Shoulder Season Oversupply (Load-Add DR)

The OS3 scenarios demonstrate the ability of **foresee** to respond to a load-add demand response event on a spring day in Portland, OR with excess hydro energy generation. The three tests, a baseline case, a case using **foresee** optimization, and a case using **foresee** with a load-add event between 10 a.m.–4 p.m., were run using weather for March 20, 2014. Homeowner profile HS2 defined the thermostat set point, the hot water draw profile, and the schedule for the washer, dryer and dishwasher (See Tables A.6, A.8). Preference profile PS1 was used to drive the **foresee** controls.

The weather of this test day is milder compared to that in OS2 and only a moderate amount of heating was required in the early morning. This scenario was designed to look at the ability of **foresee** to add load during a demand-response event through the middle of the day. The load-add demand response event overlapped with the lowest price period of the TOU schedule so the EE case saw an increase of energy use during the DR period, with the DR case further increasing the amount of energy used during the event period, as shown in Table 20. Despite the large increase in energy used during the DR period, **foresee** in the EE case decreased the overall energy use relative to the baseline case. In the DR case, **foresee** used 24.8 kWh more energy compared to the baseline case, which contributed to the significant decrease in daily energy savings and cost savings.

Figure 29 shows the accumulated energy of the three cases. The DR case has very similar accumulated energy profile with the EE case until 10 a.m. when the load add DR period started. 16.83-kWh energy was consumed by the home during the 4-hour DR period whereas the home exported energy during the same period in the baseline and EE cases.

As shown in Table 21, the EE and DR cases have the same amount of air temperature discomfort. **foresee** successfully decreased the energy backfeed by self-consuming the PV energy in the EE and DR cases. Battery was used in the EE case to store the excess PV and was used more aggressively in the DR case to store the free energy from the grid during the load add DR period.

Table 20. Summary from OS3 Scenarios

Scenario	Energy (kWh)		DR Period (kWh)		Daily Cost (\$)	Cost Savings (%)
	Total	Saved	Energy Use	DR Service		
Baseline	4.98	-	-7.97	-	1.16	-
OS3 EE	3.22	1.76	-2.44	5.53	0.51	55.9%
OS3 DR (load-add)	19.3	-14.30	16.83	24.8	1.57	-34.9%

Table 21. DER and Comfort Impacts from OS3 Scenarios

Scenario	Total Air Temp Discomfort (°F-hrs)	Water Heater Comfort Impact (kWh)	Solar Energy Backfeed (kWh)	Battery Throughput (kWh)
Baseline	0.00	0.00	12.88	0.00
OS3 EE	2.61	0.00	7.32	9.71
OS3 DR (load-up)	2.61	0.00	6.13	16.70




Figure 29. Accumulated energy from the OS3 simulation for baseline, energy efficiency and demand response cases. TOU price schedule indicated by blue shading and green shading indicates load-add DR period.

6.4.3 OS9 Scenario - Spokane, WA, Under Peak Cold Day (Load-Shed DR)

The OS9 scenario demonstrate the ability of **foresee** respond to a load-shed demand response event on a cold winter day in Spokane, WA. A total of three tests were conducted on this simulated day: a baseline case, a case using **foresee** optimization with a TOU utility rate, a case using **foresee** optimization with a load-shed event between 6 and 10 a.m. on top of a TOU rate. Each test case was run using weather for November 14, 2014. Homeowner profile HS1 defined the thermostat set point, the hot water draw profile, and the schedule for the washer, dryer and dishwasher (See Tables A.6 and A.8). Preference profile PS1 was used to derive the **foresee** controls.

This scenario used weather from a cold day in Spokane in 2014, with a lowest outdoor temperature of 16°F. Cold

winter days stress the electric grid because space heating loads are much higher and concentrated in the early morning. Similar to OS1, this scenario was used to demonstrate **foresee**'s ability to provide load-shed during time period when heating loads are the highest. The demand response period overlaps with the medium price period of the TOU period. As shown in Table 22, the EE and DR cases have small energy savings compared to the baseline case, but both cases yielded moderate cost savings by shifting the loads to the time periods with lower energy prices. Compared to the EE case, the DR case delivered over 11 kWh energy reduction during the DR period at the cost of lower energy savings and cost savings.

Figure 30 shows the accumulated energy for all three cases. The accumulated energy in the DR case increased much more quickly than the other cases between 4 and 6 a.m. after **foresee** received the notification of the DR event and stayed almost constant throughout the entire DR period. Both EE and DR cases also used much less energy than the baseline case during the high-price period between 2 p.m. and 6 p.m. to save energy cost.

As shown in Table 23, the EE and DR use cases also have similar DER and comfort impacts. Both cases have slightly higher air temperature discomfort compared to the baseline case but have near zero energy backfeed. Battery was heavily used in the EE and DR cases. The battery throughput in the EE case was 15 kWh, which was almost equivalent to a full charging and discharging cycle. The battery was used more heavily in the DR case as its throughput doubles the EE case.

Table 22. Summary from OS9 Scenarios

Scenario	Energy (kWh)		DR Period (kWh)		Daily Cost (\$)	Cost Savings (%)
	Total	Saved	Energy Use	DR Service		
Baseline	60.30	-	9.08	-	5.20	-
OS9 EE	59.31	0.99	11.42	-2.34	4.69	9.67%
OS9 DR (load-shed)	60.48	-0.18	0.34	8.74	4.91	5.51%

Accumulated Energy for OS9 Scenarios: November day in Spokane, WA

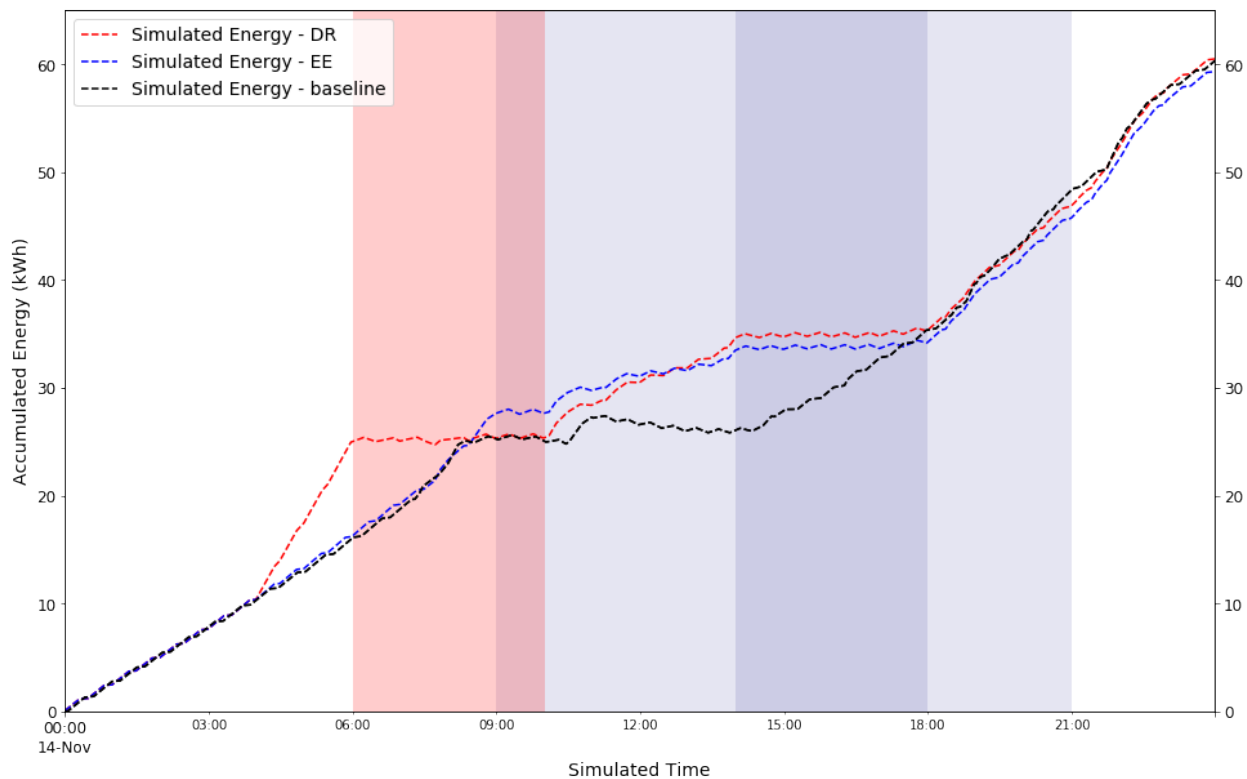


Figure 30. Accumulated energy from the OS9 simulation for baseline, energy efficiency and demand response cases. TOU price schedule indicated by blue shading and red shading indicates load-shed DR period.

Table 23. DER and Comfort Impacts from OS9 Scenarios

Scenario	Total Air Temp Discomfort (°F-hrs)	Water Heater Comfort Impact (kWh)	Solar Energy Backfeed (kWh)	Battery Throughput (kWh)
Baseline	0.08	0.00	1.75	0.00
OS9 EE	0.25	0.00	0.09	15.11
OS9 DR (load-shed)	0.29	0.00	0.09	31.43

7 Cybersecurity Risk Analysis & Implementation

7.1 Cybersecurity Risk Analysis

A detailed cybersecurity risk analysis was performed, using a process developed by ETAS/ESCRYPT, which was minimally adapted from established standards (ISO 2008). It also drew on methods developed for the electric power sector, compiled in the National Electric Sector Cybersecurity Organization Resource failure scenarios (EPRI). It is beyond the scope of this document to include the full derivation of vulnerabilities and their impacts; we provide summary findings as follows. A nonproprietary discussion of the method used, applied to connected outlet products, can be found in Langner and Christensen 2018.

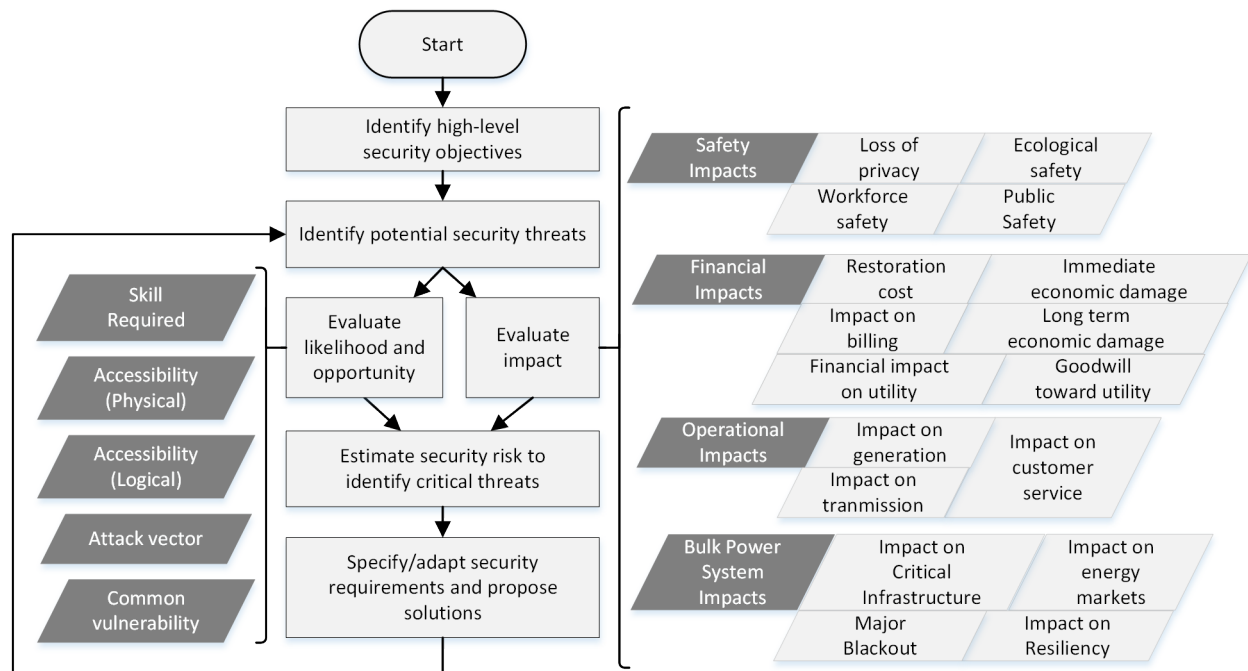


Figure 31. Flowchart showing high-level steps in the Cybersecurity Risk Analysis process.

Table 24. Summary of Calculated Risks

Attack	Attack Description	Impact	Difficulty	Risk
C-ATK-1	Extract IP from software update package.	Medium	Enhanced Basic	Undesirable
C-ATK-2	Retrieve critical data over the network	Medium	Enhanced Basic	Undesirable
C-ATK-3	Retrieve critical data from the HEMS hardware, using physical access	Medium	Enhanced Basic	Undesirable
AU-ATK-1	Install malware through software update	Critical	Enhanced Basic	Unacceptable
AU-ATK-2	Attack other devices over internal home network	Medium	Enhanced Basic	Undesirable
AU-ATK-3	Steal, disrupt or damage critical data	Critical	Enhanced Basic	Unacceptable
AU-ATK-4	Steal or modify cryptographic credentials	Critical	High	Unacceptable
AU-ATK-5	Compromise foresee software	Medium	High	Tolerable
AV-ATK-1	Disrupt normal operation of foresee , leading to loss of grid services	Catastrophic	Enhanced Basic	Unacceptable
F-ATK-1	Rollback attack using software update mechanism	Critical	Moderate	Unacceptable

As shown in Table 24, 10 attacks were identified and assessed. One was found to be of low (tolerable) risk. The other nine were deemed worthy of further evaluation and protection at this time. Because of the research/prototype nature of the **foresee** platform, several of these vulnerabilities could not be evaluated. For example, our research scope does not include develop a method for remote software updating. This is a feature required for products, but not for fundamental research. Therefore, C-ATK-1, AU-ATK-1 and F-ATK-1 could not be considered in this project. Similarly, this project uses a virtualized hardware security module (HSM) whereas a future product is expected to use a significantly more mature physical HSM to store cryptographic keys. Therefore software cybersecurity to address AU-ATK-4 is not needed. In the end, this project required a software cybersecurity layer to protect against attacks C-ATK-2, AU-ATK-3, AU-ATK-5, and AV-ATK-1. We defer addressing AU-ATK-2 until the solution is closer to product level because it is as much a challenge to the other smart home products as it is to **foresee**.

7.2 Recommendation to Mitigate Risks

Risk mitigation recommendations were generated to inform future cybersecurity layer development.

Table 25. Cybersecurity Recommendations

Recommendation	Attack
There are multiple risks to software updates. Adapt industry best practices to securely software update. It is common that such a solution verifies authenticity of the software update package through digital signature and has its own version numbering system which is checked during the software update.	C-ATK-1, AU-ATK-1, F-ATK-1
For any system that employs cryptography, it is crucial to handle cryptographic materials securely to preserve integrity of the overall system. This is out of scope for the current project. Secure key and certificate management practice, with proper key provisioning, should be used to keep the VOLTTRON platform intact.	AU-ATK-4
Utilize supporting security backend infrastructure that compliments various use cases. In addition, such infrastructure makes management and maintenance of multiple security-relevant use cases easier with better logical separation, i.e. using different cryptographic materials for individual use cases.	AU-ATK-2, AU-ATK-3, AU-ATK-4
Employ secure storage or HSM for storing cryptographic credentials. Such hardware guarantees the integrity of cryptographic materials and serves as a root of trust.	AU-ATK-4
To mitigate the risk of DoS, use mutual authentication in SSL/TLS communication. In this way, entities which are involved in secure network communication can verify each other, so communication can be terminated if found illegitimate or inauthentic.	AV-ATK-1
Protect all critical data with encryption, so confidentiality of the data is preserved even if the data are lost or stolen. Use the hardware root of trust to preserve key materials.	AU-ATK-3
Use secure boot to ensure that only authentic software from the manufacturer is permitted to run on the system.	AU-ATK-1, AU-ATK-5

Finally, recommendations were made to harden the platform on which the **foresee** software executes. Those recommendations include physical security recommendations - for future hardware implementations—and software recommendations.

7.2.1 Physical Security

- The HEMS should be installed in a place that is inside of the home and in a location that is not readily accessible to potential attackers. Ideally it would be kept in a secure area.

- Homeowners should be advised to practice basic home safety rules, such as not to leave the residence without locking the doors and windows.
- All unused network and communication (USB, etc.) ports on the HEMS product should be secured using port locks.
- The HEMS product should utilize tamper-proofing methods on the physical devices. This may include tamper-proof tape, screws, or other labels that state the device is not to be tampered with.
- A log should be kept of persons with physical access to the HBS, where appropriate.
- A HSM should be present, and used to store all cryptographic materials, including symmetric and asymmetric keys and certificates. The HSM should be used for performing secure cryptographic operations.

7.2.2 Software Security

- All firmware should be kept up-to-date at all times. BIOS should be password-protected, and all devices and ports should be disabled.
- Disable auto-mounting external devices. Restrict boot devices and disable network booting. Do not enable kernel parameters to be edited.
- UEFI secure boot must be enabled and the relevant key set for UEFI secure boot installed. Avoid using traditional boot loaders; instead use kernel, initramfs, and boot command as an EFI application.
- There should only be a single user account on the Linux platform, and it should have a strong password. Root login should be disabled, especially for remote users. Restrict “sudo” privilege escalation as much as possible.
- It is recommended to use SELinux, a kernel extension which implements a Mandatory Access Control architecture to provide fine-grain control over files, processes, users and applications.
- All unnecessary software services should be disabled.
- All software shall be kept up-to-date especially regarding security updates. The **foresee** software shall be updated using secure software update only. All software updates which could disrupt system functionality or require a system restart should be scheduled at the same time, and in conjunction with an Actuator Agent so no control actions can happen during update and restart.
- A Unified Threat Management (UTM) system should be used between the HEMS device and the homeowner’s network. It achieves several goals: network isolation, firewall, intrusion detection and intrusion prevention. All of these are required for compliance with aggregated resource standards, as described in Section 7.3.
 - Network Isolation: UTM device between the HEMS and the homeowner’s network creates a natural network separation to ensure the HEMS is only accessible through them UTM, in other words, a trusted device.
 - Firewall: UTM uses a firewall for basic packet filtering. All unused ports should be closed on this firewall to eliminate unnecessary or malicious packets from entering the HEMS.
 - Intrusion Detection and Prevention: These enable the UTM to search the incoming and outgoing network traffic for signs of malicious attacks, and stop those attacks early.
- All communications to external services should use well-known secure cryptographic algorithms. These channels should be well-documented and use no more than one port each.

7.3 CIP Compliance

One goal of this project was to enable the Home Battery System to support aggregation by a third-party aggregator or utility for participation in bulk grid markets. Currently any resources participating in those markets are required to

comply with North American Electric Reliability Corporation (NERC) Critical Infrastructure Protection (CIP) cybersecurity standards. (NERC 2017). At the time of this work, CIP version 5 was the most current. It was determined that the HBS system needs to be compliant with CIP-002-5.1 and CIP-003-6 to enable such aggregation.

A plan was developed to enable future commercialized versions of **foresee** to enable CIP compliance in aggregate. Its inclusion in this document is unnecessary, but future CIP compliance was one aspect of the cybersecurity layer design and implementation. A paper is being developed to document key aspects of this analysis (Martin et al. 2018).

7.4 Cybersecurity Layer Implementation

An implementation plan, software design and software requirements specification were developed, reviewed by all parties, and approved. ESCRYPT developed the cybersecurity layer based on that plan. NREL installed that cybersecurity layer alongside the **foresee** prototype hardware, on which the hardening steps described previously had been performed.

This system underwent penetration testing to evaluate the cyber hardness against the previously mentioned attacks; these test results are provided in the next section.

8 Cybersecurity Penetration Testing

To confirm the cybersecurity risk assessment and implementation, NREL performed a series of penetration tests on the hardened **foresee** platform described in Section 7. The **foresee** application and cybersecurity layer were installed on a Linux NUC which had been hardened according to the recommendations above.

8.1 Objective

The **foresee** platform is a prototype, not an operational system. On an operational system, the goal of a penetration test is to identify vulnerabilities in the system and develop mitigations to protect operational assets from compromise. (For instance, if weak or default passwords are in use, switch to strong passwords.)

On a prototype system, the goal of penetration testing should be to explore the security characteristics of the system under development to develop a list of recommended controls for production versions of the system. There are no operational assets to protect, but there is a need to understand what is needed to protect future operational assets. This subtle difference impacts both the assumptions made going into the test and the design of the test itself.

8.2 Scope

The scope of this test is to evaluate the prototype cybersecurity layer's ability to withstand attacks identified by the Risk Assessment, and to generate recommendations for future development of the **foresee** platform if it becomes commercialized. The tests performed have resulted in a detailed list of actionable items that can be implemented on the **foresee** platform in the future.

8.3 Tests Performed

Cybersecurity penetration tests were described in the Scenario Experiment Plan, shared with our customers and stakeholders, and approved by all parties. Among the nine potential cybersecurity penetration tests discussed there, six were found to be relevant at this stage of prototype development. A more complete discussion of these tests, including tools and methods used, along with code snippets and screenshots, is provided in Appendix C.

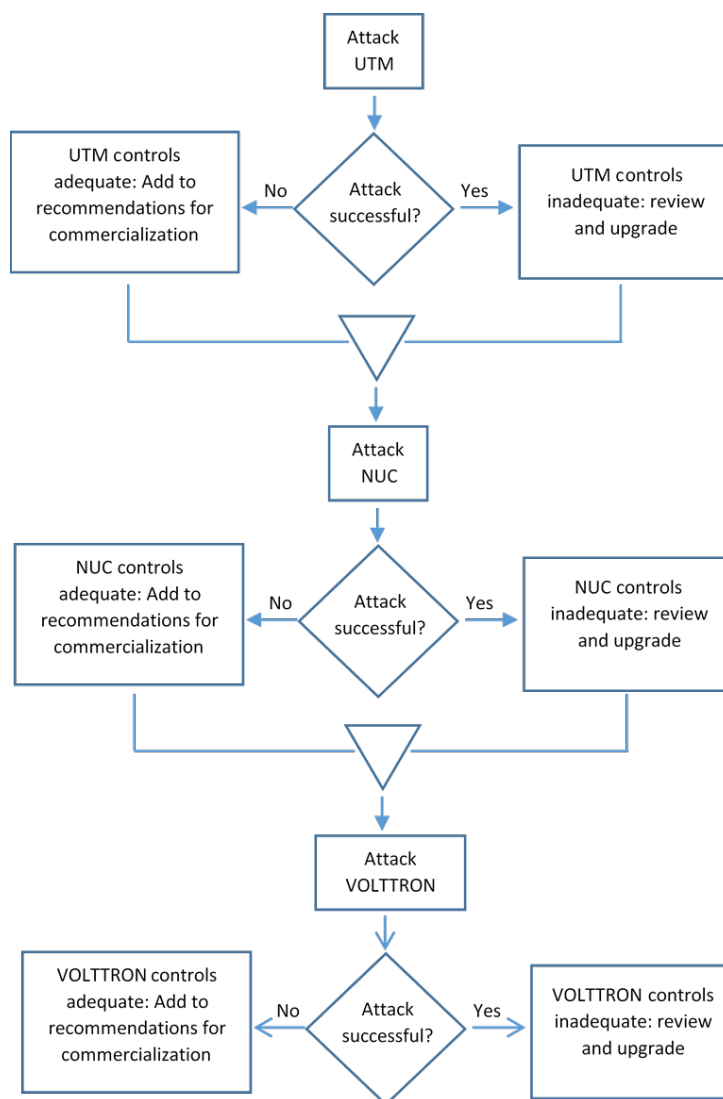


Figure 32. Flowchart showing the penetration test process.

Table 26. Penetration Tests Performed

#	Scenario Name	Description	Threat Analysis Reference	Penetration Test Activities	Test Success Metrics (i.e. Home Battery System Failure Criteria)
T1	foresee service disruption via internal DoS	Disrupt foresee platform. Assumes access to the customer LAN.	AV-ATK-1	On LAN: - DOS	Good: Slow operation of foresee platform Best: Cause foresee platform to crash
T2	Platform Operation Disruption (via customer LAN)	Alter behavior of the foresee platform. Assumes access to the customer LAN.	AU-ATK-5	On LAN: - Packet capture and spoofing (MiM)	Good: spoof data from appliances Better: spoof commands
T3	Attack foresee platform from customer LAN	Get access to the foresee platform from the customer LAN.	AU-ATK-5	On foresee platform: - Port scanning - Password guessing - Vulnerability scans	Good: gather quality data about foresee platform Better: retrieve foresee credentials
T4	foresee platform Data Disruption	Steal, disrupt, or damage foresee platform data. Assumes foresee platform access (CS2).	C-ATK-2 AU-ATK-3	On foresee platform: - Escalate privileges	Good: log onto foresee platform, access data Better: log onto foresee platform, alter data
T5	Attack communication channel outside the UTM	Attempt to deny communications from the public Internet to the customer LAN	AU-ATK-3	Outside UTM: - Physical disconnect of communication - DoS on network	Good: Disrupt communications via DoS
T6	Attack on VOLTTRON	Probe for gaps in VOLTTRON's security implementation		From inside the foresee platform: - Check for outdated versions of OpenSSL. - Attempt to connect via depreciated connection protocol - Check for key misconfiguration on OMQ and CurveZMQ	Good: Find outdated or misconfigured system software Better: Exploit outdated or misconfigured system software

8.4 Penetration Test Results and Recommendations

Of the vulnerabilities discovered the chart below organizes each discovered vulnerability by level of criticality and the test case it was discovered in. Each vulnerability includes a mitigation strategy to assist in the **foresee** application's resilience against the discovered vulnerabilities when the application is deployed. A more complete discussion of these results is provided in Appendix C.

Test #	Vulnerability	Criticality
T3	Directory Traversal	HIGH
<p>Findings: API keys used to integrate devices with foresee were able to be recovered in a useable format (unencrypted plain-text), so the metric for high risk has been met.</p> <p>Recommendations: Services accessible to users should be updated and patched promptly as security updates are released. Additionally, file permissions and permissions of applications should be limited so that critical files require higher levels of permissions to access.⁵</p>		

Test #	Vulnerability	Criticality
T1	DoS From Flooding Attacks	MODERATE
<p>Findings: The application was made inaccessible to legitimate users for as long as the flood attacks were running, but whenever the attack stopped the system returned to a working state with no administrative maintenance or repairs required. As the foresee platform never crashed the metric for high risk was not met. As the operation of the foresee platform was brought to a halt during the attacks the metric for moderate risk was exceeded.</p> <p>Recommendations: Some Denial of Service attacks may be reduced by deploying foresee with a UTM. TCP SYN floods can be reduced by using SYN cookies, micro blocks, or other mitigation techniques to reduce the likelihood of resources being exhausted during a TCP SYN flood. ICMP flooding can likely be reduced with firewalls or Intrusion Prevention Systems that are configured specifically to prevent ICMP flooding.</p>		

Test #	Vulnerability	Criticality
T4	Lack of User Authentication	MODERATE
<p>Findings: With assumed access to the foresee platform, data could be accessed. Therefore, the metric for moderate risk has been met. Because no attempts to alter data beyond the basic functionality foresee provides to a typical user was achieved, the metric for high risk has not been met.</p> <p>Recommendations: Currently, users accessing the foresee interface do not undergo any authentication process. Any form of authentication to the foresee interface can further protect it against malicious users. Htaccess Authentication can be used to restrict access to certain directories with basic HTTP authentication.</p>		

Test #	Vulnerability	Criticality
T5	WiFi Deauthentication Attack	LOW
<p>Findings: Although communications were disrupted via a Denial of Service from outside the customer LAN, the attack requires close proximity, which makes it unlikely to come from the public Internet. Since the Denial of Service attack is unlikely to come from the public Internet the metric for moderate risk has not been met.</p> <p>Recommendations: 802.11 can encrypt management frames presuming that the AP and clients support it. Other wireless network protection features can block this type of attack but will most likely require implementing security on the customer's LAN network infrastructure.</p>		

⁵To avoid such attacks, a Input validation of the requests to the web server should be implemented and the access rights of the files by the server reduced to only the necessary files. This was not yet implemented due to the prototype nature of the **foresee** platform; the web server is used for convenience in this project and an alternate method would certainly be used in production software.

9 Conclusions

9.1 Major Findings and Outcomes

The Home Battery System concept—automating flexible loads as an active EE and DR asset for customer and grid benefits—actually works. Excepting several clear opportunities where further development and platform maturing is needed, the application of multicriterion decision making results in significant homeowner savings, energy savings, and DR service outcomes.

- **Energy Efficiency:** 0.4–8.1 kWh/day savings, when no DR called (4%–50% utility cost savings)
- **Demand Response:** 5.0–6.3 kWh over 4-hour DR period (INC and DEC)
- **Energy Export:** Hawaii cases all export significantly less with **foresee**. Washington/Oregon cases generally export slightly more.
- **Multicriterion Decision Making:** Clear differences in execution, but room for improvement exists.
 - Cost savings preference did lead to lowest utility costs (and lower battery wear)
 - Negative comfort impacts are apparent in most cases, which motivates further model development to improve thermal predictiveness.

Most notably, energy was saved in every scenario and case where the HBS and **foresee** were able to optimize the home. Cost was saved in every EE case, and (compared to the Baseline) in many DR cases—despite the fact that we did not account for any utility incentive for DR in calculating the daily cost.

foresee was demonstrated to operate successfully on a microcomputer, which shows promise for a future transition to an embedded application and/or operation of numerous simultaneous instances in a cloud setting. The platform’s cybersecurity layer proved capable of handling the advanced penetration testing successfully, aside from a vulnerability that existed to facilitate early prototyping, and would not exist in a final product instance.

In parallel projects, and a DOE-funded extension, we have seen that application of **foresee** to other operational scenarios is possible (e.g. Cutler et al. 2018; Laws et al. 2018); extending the HBS to other equipment configurations is possible (e.g. Baker et al. 2016); and studying annual outcomes and establishing financial metrics to justify its adoption will be valuable (e.g. Jin, Maguire, and Christensen. 2018).

9.2 Future Work

A key opportunity to improve the performance of the HBS is in leveraging real-world data. This will allow a significantly-greater sophistication of machine learning algorithms, since the full diversity of homeowners, building constructions, appliances and weather is impossible to replicate in simulation or laboratory settings. To gather real-world data, a pilot project is anticipated. This pilot will also provide the opportunity to develop, implement, and validate the performance findings summarized earlier in this section as well as to confirm that the improvements described immediately below are implemented and ready for real-world use.

9.2.1 Comfort Delivery Improvements

Our laboratory results showed an unsatisfactory level of comfort delivery. There are a number of potential causes, and as a result a number of future work opportunities to improve the delivery of comfort. These will require additional effort to research and develop, which should include a field “alpha” pilot to validate the methods in real buildings. We also strongly recommend that the comfort improvements be performed in concert with additional work on human preferences as discussed in the next section.

- Building thermal model fidelity—increase and explore computation/accuracy tradeoffs to improve accuracy of model predictive control algorithms
- Dispatch controller—evaluate possible need to improve how this manages the appliances

- Review to ensure that preference weight is correctly computed
- Determine whether a network or implementation issue prevented **foresee** from receiving data to update house state at the beginning of some/all optimization periods
- Develop reinforcement learning methods for continuously/periodically updating the building model to better estimate building state
- Renormalize and/or reinforcement learning of preference weights to improve delivery of value vs. the initialization.

9.2.2 Preferences Improvements

This project was able to develop an initialization method for preferences, which is only a starting point for human-machine interaction in the automated home. Significant additional work is needed for **foresee** to become a proven solution for homeowners.

- Study compensatory behaviors and include these to manage secondary risks which could result from control actions
- Explore how to improve preference initialization methodology
- Validate preference methodology and MCDM
- Develop and integrate continuous/active/reinforcement learning methods for preferences
- Develop and implement time-dependent preferences
- Research hybridizing SMARTER/DCE and AHP to leverage best of each
- Study community preferences
- Evaluate methods to estimate or machine-learn preferences to replace or reduce initialization effort.

9.2.3 Stochastic Optimization and Service Commitment

We feel there is ample opportunity to improve the home's optimization under uncertainty. Uncertainty exists in all manners of inputs to the **foresee** system, including weather forecasts, occupancy forecasts, machine-learned building and appliance models, sensor-reported data, network latency, occupancy schedules, and many variable utility tariffs or incentives. Accounting for these uncertainties to deliver the highest-reliability outcomes for homeowners and utilities/aggregators has numerous potential benefits. The research performed in this project is a mere start.

- Develop probability distribution functions for inputs to the HBS
- Evaluate other stochastic optimization techniques besides chance-constrained optimization
- Identify approaches to minimize risk and uncertainty of outcomes
- Determine the impact of internal building models—which are inherently reduced-order compared to the real systems because of the need to machine-learn them from reasonably-achievable data streams—on the uncertainty of outcomes
- Determine the impact of input uncertainties on uncertainty of outcomes
- Study how aggregators could leverage uncertainty intervals to improve the confidence of committed aggregated resources into the bulk grid market
- Study how utilities could reduce reserve margins based on improved uncertainties in load forecasts or ancillary service commitment.

9.2.4 Other Use Cases

Numerous other potential use cases exist. We were unable to explore those in this project, and look forward to pursuing some of these in future work.

- Generation-following load (as opposed to demand-following generation) to enable higher penetrations of renewable energy
- Resiliency, such as enabling a home to operate islanded with higher levels of service compared to today's state of the art
- Enabling high penetrations of electric vehicles
- Supporting feeder voltage regulation and other feeder-level services
- Participating in community-scale energy solutions to reduce costs vs. today's zero-energy construction practice
- Participating in novel and emerging markets, such as energy blockchain and transactive controls
- Enabling nonenergy benefits such as bulk water conservation, carbon emission reduction, and outdoor air quality management

9.2.5 Impact Analysis

Several requests and suggestions have been made to characterize and/or analyze the potential impact that deployment of **foresee** could have in US homes. This was out of scope, so must be left for future work. Those efforts could include:

- Rigorously evaluate the maximum technical potential for different smart home technologies and packages, including the incremental benefits of whole-home automation such as **foresee**
- Determine the regional variability of energy and cost savings in order to help market adoption in the highest-benefit areas proceed quickly
- Establish operational strategies for aggregating and managing distributed homes across a utility power network

A Scenario Experiment Details

This appendix provides additional detail we developed to operate the laboratory experiments.

The Home Battery System will be evaluated under operational scenarios designed to exercise different features of the control software and equipment, to demonstrate their usefulness in different locations and with different users.

With input from DOE and BPA, NREL defined the operational scenario tests, shown in the following table. This table describes the major differences between the different tests, but a number of additional details were needed before testing could begin.

Table A.1. Operational Scenario Definitions

Operational Scenario	Grid Scenario	Climate/Residence Scenario	Homeowner Scenario	Preference Scenario	Baseline Studies
OS1a	GS1-T4 (<i>TOU+feed-in</i>)	CS1v	HS4	PS4	BAU
OS1b	GS1-T4 (<i>TOU+feed-in</i>)	CS1v	HS4	PS5	
OS1c	GS1-T4 (<i>TOU+feed-in</i>)	CS1v	HS4	PS6	
OS2*	GS3 (<i>Peak Cold INC DR</i>)	CS1v	HS2	PS1	BAU, EE
OS3*	GS4 (<i>Shoulder Season DEC DR</i>)	CS1v	HS2	PS1	BAU, EE
OS4	GS1-T4 (<i>TOU+feed-in</i>)	CS2n	HS1	PS1	BAU
OS5	GS4 (<i>Shoulder Season DEC DR</i>)	CS2n	HS1	PS1	BAU, EE
OS6	GS7 (<i>Self-Consume PV</i>)	CS2n	HS1	PS1	BAU, EE
OS7*	GS1-T4 (<i>Shoulder Season DEC DR</i>)	CS5v	HS1	PS1	BAU
OS8	GS2 (<i>Peak Hot INC DR</i>)	CS5v	HS1	PS1	BAU, EE
OS8-2	GS1-T2 (<i>flat rate+feed-in</i>)	CS5v	HS1	PS1	BAU
OS9*	GS4 (<i>Peak Cold INC DR</i>)	CS5v	HS1	PS1	BAU, EE

*OS2, OS3, and OS9 were performed in simulation due to budget and schedule constraints.

A.1 Grid Scenarios

The grid scenarios describe the utility rate and additional grid services that are requested for each scenario, when applicable. The base case for all locations (GS1-T4) is a Time-of-Use (TOU) rate with a feed-in tariff for any excess generation. Of the three locations we will be evaluating, Portland and Hawaii both have optional TOU rates available to residential customers. Those TOU schedules will be used for testing. Spokane does not have a TOU rate available, so the Xcel Energy optional TOU rate available in Denver (Xcel 2018) was scaled by 0.75 because the base rate in Spokane is roughly 75% that of Denver's base rate. We will also apply a feed-in tariff during the base case that sets the feed-in tariff (compensation the homeowner receives for exporting energy back to the grid) to 25% of the current rate of energy purchases. An additional scenario (OS8-2) was run with a flat rate and feed-in tariff in Spokane, WA. The flat rate was chosen to be the middle cost from Spokane's TOU rate: \$0.0975/kWh.

The other grid scenarios will evaluate different conditions that would benefit from residential demand-side management. We will evaluate the need to reduce load (also referred to as demand response) during the late afternoon on peak hot days and during the morning on peak cold days. We will also look at the need to increase load due to shoulder season oversupply (a phenomenon that is common in regions with a lot of hydroelectric power generation). Lastly, we will evaluate a case found in Hawaii where the utility will not pay for any excess generation returned to the grid, so the customer has an incentive to self-consume all their on-site generation.

In each case, a relevant monetary incentive will be applied to encourage the desired behavior. For the demand response cases, a 5x multiplier will be applied to the cost of electricity during the period of time when load reduction

Table A.2. Time-of-Use Utility Rates for each Location

Time	Portland		Honolulu (\$/kWh)	Spokane (\$/kWh)
	Summer ^a (\$/kWh)	Winter ^b (\$/kWh)		
12:00 a.m.	0.04399	0.04399	0.263	0.06
1:00 a.m.	0.04399	0.04399	0.263	0.06
2:00 a.m.	0.04399	0.04399	0.263	0.06
3:00 a.m.	0.04399	0.04399	0.263	0.06
4:00 a.m.	0.04399	0.04399	0.263	0.06
5:00 a.m.	0.04399	0.04399	0.263	0.06
6:00 a.m.	0.07572	0.13197	0.263	0.06
7:00 a.m.	0.07572	0.13197	0.263	0.06
8:00 a.m.	0.07572	0.13197	0.263	0.06
9:00 a.m.	0.07572	0.13197	0.175	0.0975
10:00 a.m.	0.07572	0.07572	0.175	0.0975
11:00 a.m.	0.07572	0.07572	0.175	0.0975
12:00 a.m.	0.07572	0.07572	0.175	0.0975
1:00 p.m.	0.07572	0.07572	0.175	0.0975
2:00 p.m.	0.07572	0.07572	0.175	0.135
3:00 p.m.	0.13197	0.07572	0.175	0.135
4:00 p.m.	0.13197	0.07572	0.175	0.135
5:00 p.m.	0.13197	0.13197	0.399	0.135
6:00 p.m.	0.13197	0.13197	0.399	0.0975
7:00 p.m.	0.13197	0.13197	0.399	0.0975
8:00 p.m.	0.07572	0.07572	0.399	0.0975
9:00 p.m.	0.07572	0.07572	0.399	0.06
10:00 p.m.	0.04399	0.04399	0.263	0.06
11:00 p.m.	0.04399	0.04399	0.263	0.06

^a May–October

^b November–April

is desired, and notification of the event will be send several hours in advance. Similarly, a 1/5 multiplier will be applied to the cost of energy during the ‘load add’ period for the shoulder season case. Lastly, self-consuming local generation will be encouraged by removing the feed-in tariff. Back-feeding energy to the grid will not be penalized.

For each operational scenario, a baseline test (without foresee control) was run first, followed by an optimization test where foresee was allowed to manage the controlled appliances based on the preferences shown in Table 7. For some scenarios, a third test was completed to demonstrate demand response or some other grid service request. The Baseline Studies column of Table 7 indicates whether that scenario will be compared to just a Baseline case (“BAU”) and often an optimization case (“EE” for energy efficiency).

Table A.3. Summary of Grid Scenario details

	Description	Event Duration	Incentive	Notification window
GS1-T4	Time of Use rate + 75% multiplier feed-in-tariff	-	-	-
GS2	Peak Heat Day	6 hours (2–8 p.m.)	5x multiplier on energy cost	4 hours prior
GS3	Peak Cold Day	4 hours (6–10 a.m.)	5x multiplier on energy cost	12 hours prior
GS4	Shoulder Season Oversupply	6 hours (10 a.m–4p.m.)	1/5x multiplier on energy cost	??
GS7	Solar Self-Consumption	No set duration	No Feed-in tariff	None

A.2 Climate/Residence Scenarios

Each region has different grid service needs, mainly driven by climate and other unique factors (such as Hawaii’s high PV penetration driven by their high electricity costs). We will evaluate three different locations in the laboratory tests: Portland, OR; Honolulu, HI; and Spokane, WA. Portland is in a marine climate, Spokane is in a colder climate with more extreme temperatures in both summer and winter, and Honolulu has unique challenges related to local generation.

Weather drives several of the residential appliances in the laboratory—the heat pump and PV inverter in particular. There are seasonal impacts on the use of other appliances that will not be included here. An EnergyPlus (*EnergyPlus*) simulation is used in a hardware-in-the-loop configuration to drive the operation of the heat pump, which requires a standard hourly weather file that includes outdoor temperature, humidity, wind speed and solar insolation. The PV inverter is driven by a power supply, known as a PV simulator, that requires second-level solar radiation data. Hourly weather data are available for many locations around the country, but higher resolution data were harder to obtain for all past years. One-second data solar radiation data were only available in Honolulu for 2011, so we are using weather data from 2011 for the Hawaii cases. We had 2014 hourly weather files for Portland and Spokane and one-minute solar radiation data were available for both locations (the Spokane solar data were actually collected in Cheney, WA, a town about 15 miles southwest of Spokane) through the University of Oregon Solar Radiation Monitoring Laboratory for all years since 2004. We interpolated the 1-minute solar data to create 1-second data files.

The following table shows the dates in those weather files that will be used to test the different grid scenarios. The individual days were chosen based on the ambient temperature and solar radiation, depending on the grid scenario to be tested. GS1-T4 is described in the scenario test plan as “an uneventful early summer day” so the actual day of data were chosen to be representative of a typical early summer day—not too hot, an average amount of solar radiation. GS2 occurs on a peak heat day, so the hottest day in 2014 in Spokane was chosen. GS3 occurs on a peak cold day, so cold days were chosen for Portland and Spokane. The days chosen were not the coldest day of the year but were close. The coldest days had strangely shaped temperature profiles so days with more typical profiles were chosen for testing. GS4 occurs on mild spring days so days in Portland and Hawaii were chosen that are expected to have minimal heating or cooling demands and some solar radiation, but not a lot. Lastly, GS7 concerns the backfeed of excess solar generation, so a day in Hawaii was found with the most solar radiation on a hot day, but not the hottest day available.

The other climate-dependent aspect of the laboratory testing is the building simulation used to drive the heat pump hardware-in-the-loop system. Construction methods and materials vary depending on the region of the country, so the building models were modified to match the regional variation. Drawing from data from a Pacific Northwest-specific version of ResStock that uses BPA’s climate zone definitions and RBSA data to determine building characteristics, building models with the most frequently seen characteristics for Portland and Spokane were created (Wilson et al. 2016). Per a request from BPA, the Portland and Spokane homes are not new homes, but of the most common vintage (1970s for both cities). There is less data available for Hawaii and it is not included in the general

Table A.4. Climate Scenarios and Days for testing

	Location	Year of Data	Weather Day for Scenario Experiment				
			GS1-T4	GS2	GS3	GS4	GS7
CS1v	Portland, OR	2014*	06/05/14	-	02/05/14	04/18/14	-
CS2n	Honolulu, HI	2011‡	06/19/11	-	-	03/20/11	08/07/11
CS5v	Spokane, WA	2014*†	06/04/14	07/29/14	11/14/14	-	-

* Solar data from Univ of Oregon SRML (Oregon 2018), hourly weather data from Weather Analytics

‡ Solar data from Oahu Solar Measurement Grid (Sengupta and Andreas 2010), hourly weather data from NREL NSRD (NREL 2018).

† Solar data for Spokane was actually collected in Cheney, WA.

version of ResStock, so we recreated construction information from a 2011 Oahu project where NREL provided building energy modeling used to design the homes (P. Norton 2013). The homes were designed to be affordable zero energy homes, which seemed appropriate for simulating a new home in Hawaii today.

Table A.5. Description of EnergyPlus models used in scenario experiments

	Location	Size (ft ²)	Vintage	Walls	Siding	Attic	Foundation
CS1v	Portland, OR	2016	1970s	R-13, 2x4 walls	Vinyl siding	Vented, R-19 insulation	Crawlspace, R- 19 insulation, vented
CS2n	Honolulu, HI	1341	New	Uninsulated, 2x4 walls	Aluminum siding	Vented, R-38 insulation	Uninsulated Slab
CS5v	Spokane, WA	2016	1970s	R-13, 2x4 walls	Vinyl siding	Vented, R-19 insulation	Uninsulated Basement

A.3 Homeowner Scenarios

While the operation of the heat pump and PV inverter were largely driven by weather and building construction, the rest of the appliances in SPL are driven mostly by the behavior of the people in home. Three different homeowner profiles were created for lab testing, with a high-level description of each profile given in the following. The profiles were used to determine tangible schedules for the remaining large appliances.

Table A.6. Homeowner Description

	HS1	HS2	HS4
Number of Occupants	2	2	4
Hot Water Use (gal)	24.3	44.4	59.9
Cooling Setpoint (°F)	72	75	75
Heating Setpoint (°F)	69	65	69

A.3.1 Water Heater Draw Profile

Hot water consumption, not tank thermal loss, is the main drive for water heater energy use. A programmable solenoid valve on the hot water outlet from the water heater is used to create a realistic hot water draw profile for each of the homeowners. The hot water use given in Table 6 is interpreted as being relative to households of similar sizes, rather than relative to the average across all home sizes. We are assuming HS1 and HS2 are composed of two

occupants each, while HS4 has four occupants. Based on a recent report (Parker, Fairey, and Lutz 2015), the average daily hot water use for two occupants is about 35 gallons and about 60 gallons for four occupants.

We then used the Domestic Hot Water Event Schedule Generator (Hendron, Burch, and Barker 2014) for a 3-bedroom house in Denver to generate a full year of daily schedules. We selected specific days from that year to find daily hot water usage volumes that matched our targets that also had reasonable timing for draws (for example, we excluded days that had large shower or bath draws in the middle of the night). The three draw profiles selected are shown in the following. HS2 is the lowest volume profile, with 24.3 gallons drawn, HS2 is the middle with 44.4 gallons used daily and HS4 is the highest draw volume with 59.9 gallons per day. The draw schedule is shown in Table A.7.

Table A.7. Hot water draw profiles for different user profiles

Low Volume		Medium Volume		High Volume	
Draw Time	Draw (Gal)	Draw Time	Draw (Gal)	Draw Time	Draw (Gal)
5:48 a.m.	12.2	1:04 a.m.	0.7	7:16 a.m.	3.2
7:45 a.m.	0.5	1:06 a.m.	1.1	7:46 a.m.	11.6
9:29 a.m.	0.6	7:48 a.m.	0.8	10:25 a.m.	14.3
9:36 a.m.	1.4	7:50 a.m.	1.2	10:45 a.m.	1.4
11:58 a.m.	1.7	9:16 a.m.	21.1	12:06 p.m.	0.6
5:31 p.m.	2.6	9:46 a.m.	8.2	2:32 p.m.	3.3
6:19 p.m.	0.5	2:18 p.m.	1.1	3:47 p.m.	1.5
6:29 p.m.	2.5	4:51 p.m.	2.1	6:02 p.m.	0.5
6:36 p.m.	2.2	4:53 p.m.	1.3	6:26 p.m.	2.0
Total:	24.3	6:25 p.m.	0.5	6:36 p.m.	1.6
		8:15 p.m.	3.2	6:46 p.m.	2.1
		8:21 p.m.	0.8	6:50 p.m.	1.2
		11:16 p.m.	2.3	6:56 p.m.	2.4
		Total:	44.4	7:39 p.m.	1.0
				8:00 p.m.	1.4
				8:03 p.m.	0.8
				8:33 p.m.	1.8
				9:40 p.m.	3.6
				9:45 p.m.	2.8
				9:58 p.m.	2.8
				Total:	59.9

A.3.2 Washer, Dryer, and Dishwasher

The remaining large appliances are all devices that run with finite cycles. Because these devices traditionally require people to be home to start them, we had to define hours that the home was occupied. Based on the occupancy metrics given in Table 8, we assumed that HS1 is unoccupied between 9 a.m. and 5 p.m., HS2 is occupied all day, and HS4 is unoccupied between 9 a.m. and 4 p.m. This information was paired with data from the Home Simulation protocol about when these three appliances are mostly likely to be run. The HSP profiles are averages across all types of households are shown in Figure A.1.

The following schedules were created for each user profile based on the assumed times that people were home, the HSP hourly profiles, and by requiring that the clothes washer complete its cycle before the dryer could start.

In all baseline cases, we scheduled the appliances to run on the schedule associated with that scenario. In all experimental cases that used foresee, the appliance schedule was used by default, but foresee was allowed to shift the start time for any of the appliances to a later time, as long as the cycle finished by midnight.




Figure A.1. Hourly Probability Profiles for the washer, dryer and dishwasher from the Home Simulation Protocol (Wilson et al. 2014).

Table A.8. Large appliance schedules for different user profiles

HS1			HS2			HS4		
Washer	Dryer	Dishwasher	Washer	Dryer	Dishwasher	Washer	Dryer	Dishwasher
8:00 p.m.	10:00 p.m.	7:00 p.m.	9:00 a.m.	11:00 a.m.	10:00 a.m.	4:00 p.m.	6:00 p.m.	8:00 p.m.

A.3.3 Preferences

Four different preference profiles were used during the **foresee** laboratory tests. These profiles are generated using the foresee setup process that asks the user to rank the priority of different outcomes, such as saving money or having plenty of hot water. Operational Scenario 1 compares the outcomes for three different preference profiles: PS4, PS5 and PS6. PS4 describes a user whose highest priority is saving money, PS5 prioritizes comfort over all other outcomes, and PS6 wants to make choices that maximize renewable energy usage. The rest of the scenarios were run using the same PS1 profile, a profile from the preference surveys collected in the first year of this project.

A.4 Laboratory Procedure

All experiments (baseline, EE and DR) were run for 24 hours, beginning at midnight on the day indicated in Table 4. Each case was configured with the same start up procedure to ensure that the appliances began in the same state.

The HVAC HIL system is run by starting an EnergyPlus simulation that will determine the initial conditions for the ‘Outdoor’ chamber and the ‘Indoor’ environmental chamber housing the thermostat (Pratt et al. 2017). After 1 minute, EnergyPlus has determined those initial conditions and sent setpoint values to the control systems managing the temperature of the two chambers. The simulation was paused at this point to allow the chambers to get to their initial conditions.

The water heater was set to 125°F for all tests. Before each test began, water was drawn from the water heater until a heating element turned on. The water heater was allowed to fully reheat the tank before the test officially started. The washer and dryer were preloaded with a load of dry and wet clothes, respectively, before a test began. The dryer

needs to be loaded with wet clothes to run in typical fashion. By using separate loads in the washer and dryer, each appliance could be run on demand, without needing a researcher to move the wet load from the washer into the dryer. Once the loads of laundry were preloaded, the “remote start enabled” button was pressed on each device. The dishwasher did not need any special treatment, but its “remote start enabled” button was pressed for each test.

The battery was charged to 2.4 kWh before each test. Generally, **foresee** would discharge the battery to below 2.4 kWh by the end of each test day, so charging was usually required before each new test.

A.4.1 Experiment Initialization

Once all the systems were at their initial conditions, different simulators were started to begin the test. The paused EnergyPlus simulation was unpaused and allowed to run. The appropriate 24-hour schedule for the hot water draw profile was started for the water heater. The PV simulator used to drive the PV inverter with a 24-hour insolation profile based on data from the test location was turned on. Lastly, **foresee** was started in whatever mode was needed. We used an agent in **foresee** to schedule the appliances for the baseline tests, even though **foresee** was not optimizing the control of appliances in the baseline tests.

A.4.2 Data Analysis

Most of the data for the laboratory tests was collected using a National Instruments-based data acquisition system at 1-second intervals. The 1-second data were averaged over a minute before analyzing the results. Individual appliance power data were collected for all the devices that were part of these experiments. Additional temperature and flow data were collected for the HVAC HIL systems and the water heater. The data acquisition system also has the ability to send control signals and was used to run the schedule of hot water draws and maintain the HVAC HIL set point temperatures. After each test was completed, data were exported from the data acquisition system and formatted for analysis.

A few corrections were applied to the data to ensure that the results are more realistic. The first correction involved correcting some unrealistic performance of loads in the lab and accounting for uncontrolled loads that did not exist in the laboratory. We used a bank of 30 light bulbs to represent the uncontrolled base load in a home, which added a constant power draw of about 500W at all times. We also ran the air handler continuously to ensure well-mixed air in the duct loop portion of the HVAC HIL system; normally the air handler would only run when cooling was called for. The laboratory experiments included all the major appliances in a home, but there are additional loads that we did not have in the lab, such as miscellaneous electric loads and cooking loads. We ran an annual simulation using the EnergyPlus models used to drive the HVAC HIL system and generated disaggregated simulated load data for each house in each climate. After we removed the appliance loads that we were controlling in the lab, we were left with the total uncontrolled load, which was comprised of the range hood, bath exhaust fan, miscellaneous electric loads, lights, and the electric range, all based on data in the Building America Home Simulation protocol (Wilson et al. 2014). This simulated load varies hourly, in contrast to the base load from the experimental tests that never changed. We removed the lighting load from the lab data and added in the simulated uncontrolled load. We also modified the air handler power data to be on only during time when the condenser is on, as it would typically be controlled. This correction was applied to all laboratory data sets.

The second correction involves the battery and so does not affect the baseline cases. We started each test with the battery charged to 2.4 kWh, which is equal to 20% of the battery’s energy capacity. However, **foresee** always discharged the battery beyond its initial level, ending each day at about 1.8 kWh. That difference of 0.6 kWh would look like additional energy savings because more energy was drawn from the battery than was added. To account for this discrepancy, we adjusted the total energy consumption for all tests that used by battery by 0.69 kWh, the amount of energy needed to charge the battery back to its initial state of 2.4 kWh. This correction is larger than 0.6 kWh because the battery does not charge perfectly efficiently (the measured charging efficiency was 0.875).

B Detailed Laboratory Performance Results

B.1 OS1 Scenarios - Portland, OR with three different user profiles

The OS1 scenarios were designed to show the difference in performance for three different user preference profiles in **foresee**. All four tests, a baseline day and three days using **foresee** optimization, were run using weather for June 5, 2014. Homeowner profile HS4 defined the thermostat set point, the hot water draw profile, and the schedule for the washer, dryer and dishwasher (See Tables A.6, A.8). The three experimental cases for OS1 used three different homeowner preference profiles: PS4, PS5, and PS6. These profiles are intended to represent people that prioritize saving money, improved comfort, and maximizing use of renewable energy, respectively. In the following results, OS1a refers to profile PS4, OS1b corresponds to PS5, and OS1c uses profile PS6.

The results from the baseline day are shown in the following figure, broken down by the different end loads. The blue shading indicates the relative price of the time-of-use price schedule, with the darkest shading representing the most expensive period of time. For more detail on the TOU schedule, see Table A.2.




Figure B.1. OS1 Baseline experiment with all loads shown. TOU price schedule indicated by blue shading.

The experimental cases of OS1a, OS1b and OS1c have similar daily load profiles as shown in Figure B.1, but with changes to when the appliances ran and the use of battery as needed. The accumulated energy for all the cases are shown in Figure B.2. Due to the amount of solar energy generated by the simulated solar panels and the low cooling needs on a mild early summer day, the baseline case came very close to being net zero (generating enough energy to offset the energy drawn from the grid). The experimental cases used less energy than the baseline case, meaning that each case exported excess solar energy to the grid. A summary of the net energy consumption of each case and the impact on daily cost and solar generation are given in Table B.1.

The battery was not operated during the baseline case, but was controlled by **foresee** in the experimental cases. The following figure shows how the battery was operated in cases OS1a, OS1b and OS1c. OS1a used the battery the least, while OS1b and OS1c used the battery in similar fashions. Battery throughput is listed in Table B.2 for each




Figure B.2. Accumulated energy from the OS1 baseline and the three scenarios with different preferences. TOU price schedule indicated by blue shading.

Table B.1. Summary from OS1 Scenarios

Scenario	Total Energy (kWh)	Energy Savings (kWh)	Daily Cost (\$)	Cost Savings (%)
Baseline	0.4	-	0.93	-
OS1a (Cost)	-3.6	3.3	0.68	23.9
OS1b (Comfort)	-2.1	1.9	0.75	18.7
OS1c (Carbon)	-2.1	1.9	0.75	19.4

case and the baseline case resulted in a nonzero throughput because the battery continues to draw power even when it is not being used.

Use of the battery in OS1b and OS1c enabled some cost savings because the battery charged during the lowest price period in the morning and then discharged during the higher cost periods later in the day. The battery did not facilitate more utilization of solar energy generation. In fact, in all experimental cases, the amount of energy exported from the house went up relative to the baseline because the PV-equipped home consumed negligible electricity in the baseline case, and **foresee** operated it to use less energy in all other cases.

The energy savings in each experimental case came primarily from a reduction in air conditioner energy use and a reduction in water heater energy use. In this scenario, all cases, including the baseline, experienced some level of discomfort due to the air temperature deviating from the comfort band, as shown in Figure B.4. In this case, the discomfort metric was only calculated between 3 p.m. and midnight because cooling was not needed until the later afternoon. The house cools off significantly at night, but that deviation outside of the comfort bound was not counted in the discomfort metric because it occurred naturally. On the other hand, there was no impact to the delivered hot water temperature and the outlet water was always above the minimum temperature of 110°F.

Battery Power Consumption for OS1 Scenarios: June day in Portland, OR




Figure B.3. Battery power consumption for OS1 optimization cases. Battery charging is considered positive power and discharging is considered negative. TOU price schedule indicated by blue shading.

Accumulated Air Temperature Discomfort for OS1 Scenarios: June day in Portland, OR

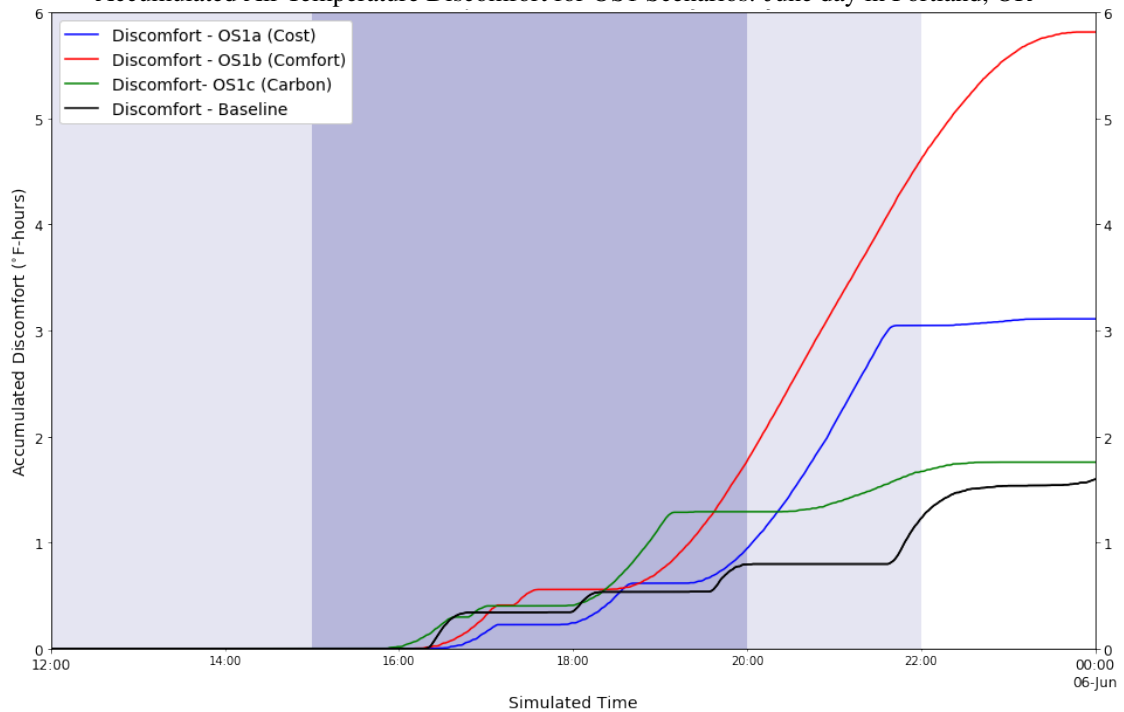


Figure B.4. Accumulated Discomfort for OS1 cases. Cooling was not needed until 15:00. TOU price schedule indicated by blue shading.

Table B.2. DER and Comfort Impacts from OS1 Scenarios

Scenario	Total Air Temp Discomfort (°F-hrs)	Water Heater Comfort Impact (kWh)	Solar Energy Backfeed (kWh)	Battery Throughput (kWh)
Baseline	1.6	0	18.2	0.3
OS1a (Cost)	3.1	0	18.6	3.1
OS1b (Comfort)	5.8	0	23.1	13.7
OS1c (Carbon)	1.8	0	22.5	13.4

A summary of the DER and comfort impacts is presented in Table B.2. It is unclear why **foresee** selected control actions that resulted in the highest degree of discomfort in the OS1b case, where users’ preferences indicated comfort as the highest priority. Because the other two use cases achieved significant energy and greater cost savings with lower comfort impacts, it appears there are sensitivities that are not represented in **foresee**’s model-predictive planning model. Thus there is room for improvement; this merits future research.

B.2 OS2 Scenario - Portland, OR under peak cold day (load shed DR)

The OS2 scenario demonstrate the ability of **foresee** respond to a load-shed demand response event on a cold winter day in Portland, OR. A total of three tests were conducted on this simulated day: a baseline case, a case using **foresee** optimization with a TOU utility rate, a case using **foresee** optimization with a load-shed event between 6 and 10 a.m. on top of a TOU rate. Each test case was run using weather for February 5, 2014. Homeowner profile HS2 defined the thermostat set point, the hot water draw profile, and the schedule for the washer, dryer and dishwasher (See Tables A.6, A.8). Preference profile PS1 was used to derive the **foresee** controls.

This scenario used weather from one of the coldest day in Portland in 2014, with a lowest outdoor temperature of 21°F. Cold winter days stress the electric grid because space heating loads are much higher and concentrated in the early morning. This scenario was used to demonstrate **foresee**’s ability to provide load-shed during time period when heating loads are the highest. The demand response period overlaps with the medium price period of the TOU period. As shown in Table B.3, the EE and DR cases have small energy savings compared to the baseline case, but the cost savings in the EE and DR cases are as high as 25%. It shows **foresee** responded to the TOU rates by shifting the loads to the time periods with lower energy prices. The daily energy savings and energy savings over the DR time periods for the EE and DR cases are very close because the TOU rates dominated the **foresee** operation and not much additional loads can be reduced during the DR time period.

The accumulated energy is shown in Figure B.5. The accumulated energy of the EE and DR cases stayed almost constant during the DR time period (6 - 10 a.m.) whereas the baseline case has lower accumulated energy at 6 a.m. and ended up with higher accumulated energy at than the EE and DR cases at 10 a.m. Without **foresee**, the home has much higher solar energy export in the afternoon compared to the EE and DR cases.

The EE and DR use cases also have similar DER and comfort impacts. Both cases have slightly higher air temperature discomfort but much lower energy backfeed than the baseline case. Battery was heavily used in the EE and DR cases; more than 26 kWh battery throughput occurred in both cases, which was about three times of the available battery capacity [i.e., 12 kWh × (85%-15%) = 8.4 kWh].

Table B.3. Summary from OS2 Scenarios

Scenario	Daily Energy (kWh)		DR Period (kWh)		Daily Cost (\$)	Cost Savings (%)
	Total	Saved	Energy Use	DR Service		
Baseline	39.58	-	7.61	-	3.84	-
OS2 EE	39.41	0.17	0.06	7.54	2.89	24.8%
OS2 DR (load-shed)	39.42	0.16	0.04	7.57	2.90	24.6%

Accumulated Energy for OS2 Scenarios: February day in Portland, OR




Figure B.5. Accumulated energy from the OS2 simulation for baseline, energy efficiency and demand response cases. TOU price schedule indicated by blue shading and red shading indicates load-shed DR period.

Table B.4. DER and Comfort Impacts from OS2 Scenarios

Scenario	Total Air Temp Discomfort (°F-hrs)	Water Heater Comfort Impact (kWh)	Solar Energy Backfeed (kWh)	Battery Throughput (kWh)
Baseline	0.02	0.00	4.64	0.00
OS2 EE	0.44	0.00	0.16	26.21
OS2 DR (load-shed)	0.44	0.00	0.05	26.27

B.3 OS3 Scenario - Portland, OR, Under Shoulder Season Oversupply (Load-Add DR)

The OS3 scenarios demonstrate the ability of **foresee** to respond to a load-add demand response event on a spring day in Portland, OR with excess hydro energy generation. The three tests, a baseline case, a case using **foresee** optimization, and a case using **foresee** with a load-add event between 10 a.m. and 4 p.m., were run using weather for March 20, 2014. Homeowner profile HS2 defined the thermostat set point, the hot water draw profile, and the schedule for the washer, dryer and dishwasher (See Tables A.6, A.8). Preference profile PS1 was used to drive the **foresee** controls.

The weather of this test day is milder compared to that in OS2 and only a moderate amount of heating was required in the early morning. This scenario was designed to look at the ability of **foresee** to add load during a demand-response event through the middle of the day. The load-add demand response event overlapped with the lowest price period of the TOU schedule so the EE case saw an increase of energy use during the DR period, with the DR case further increasing the amount of energy used during the event period, as shown in Table B.5. Despite the large increase in energy used during the DR period, **foresee** in the EE case decreased the overall energy use relative to the baseline case. In the DR case, **foresee** used 24.8 kWh more energy compared to the baseline case, which contributed

to the significant decrease in daily energy savings and cost savings.

Figure B.6 shows the accumulated energy of the three cases. The DR case has very similar accumulated energy profile with the EE case until 10 a.m. when the load add DR period started. 16.83-kWh energy was consumed by the home during the 4-hour DR period whereas the home exported energy during the same period in the baseline and EE cases.

The EE and DR cases have the same amount of air temperature discomfort. **foresee** successfully decreased the energy backfeed by self-consuming the PV energy in the EE and DR cases. Battery was used in the EE case to store the excess PV and was used more aggressively in the DR case to store the free energy from the grid during the load add DR period.

Table B.5. Summary from OS3 Scenarios

Scenario	Total Energy (kWh)	Energy Savings (kWh)	DR Period Energy (kWh)	DR Energy Add (kWh)	Daily Cost (\$)	Cost Savings (%)
Baseline	4.98	-	-7.97	-	1.16	-
OS3 EE	3.22	1.76	-2.44	5.53	0.51	55.9%
OS3 DR (load-add)	19.3	-14.30	16.83	24.8	1.57	-34.9%

Accumulated Energy for OS3 Scenarios: April day in Portland, OR

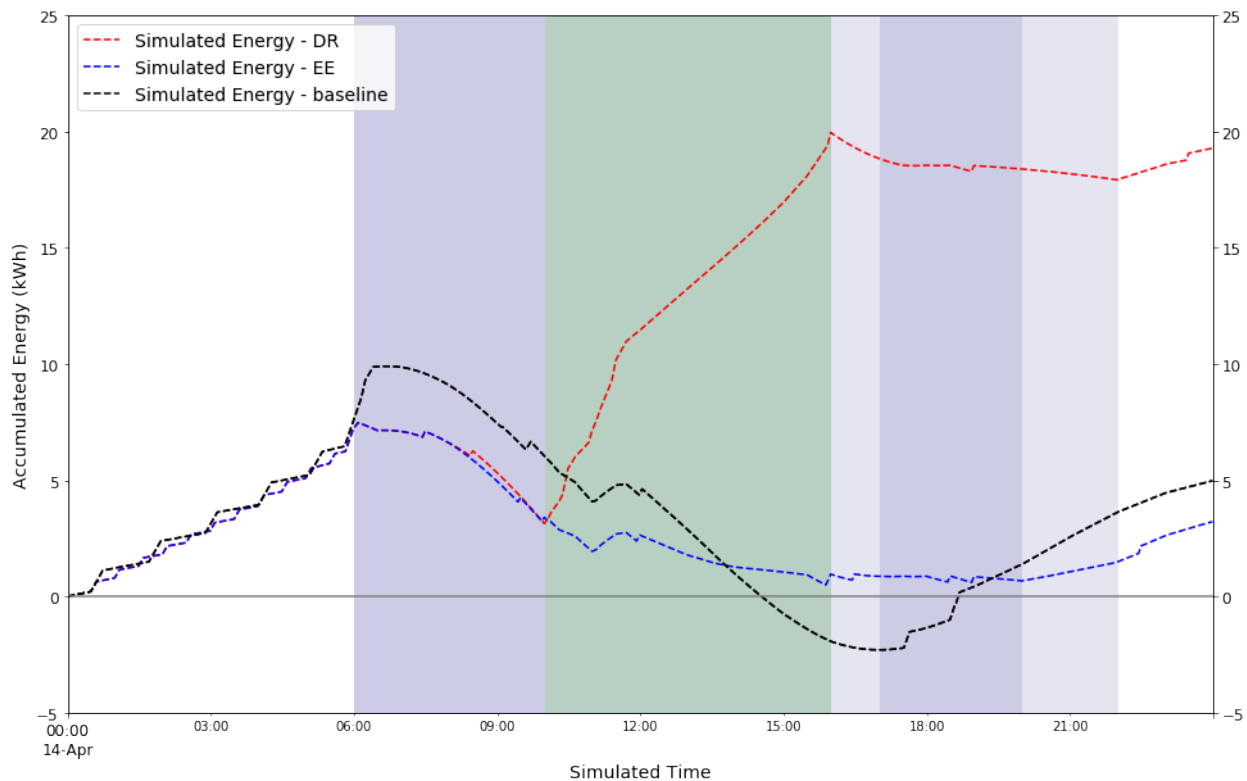


Figure B.6. Accumulated energy from the OS3 simulation for baseline, energy efficiency and demand response cases. TOU price schedule indicated by blue shading and green shading indicates load-add DR period.

B.4 OS4 Scenario - Honolulu, HI, Under Normal Day Operation

The OS4 scenarios were designed to show the impact of **foresee** control on an early summer day in Honolulu, HI. The two days of tests, a baseline day and a day using **foresee** optimization, were run using weather for June 19,

Table B.6. DER and Comfort Impacts from OS3 Scenarios

Scenario	Total Air Temp Discomfort (°F-hrs)	Water Heater Comfort Impact (kWh)	Solar Energy Backfeed (kWh)	Battery Throughput (kWh)
Baseline	0.00	0.00	12.88	0.00
OS3 EE	2.61	0.00	7.32	9.71
OS3 DR (load-up)	2.61	0.00	6.13	16.70

2011. Homeowner profile HS1 defined the thermostat set point, the hot water draw profile, and the schedule for the washer, dryer and dishwasher (See Tables A.6, A.8). Preference profile PS1 was used to drive the **foresee** controls.

The results from the baseline day are shown in the Figure B.7, broken down by the different end loads. The blue shading indicates the relative price of the time-of-use price schedule, with the darkest shading representing the most expensive period of time. For more detail on the TOU schedule, see Table A.2.

Circuit-level Power Consumption for OS4 baseline case: June day in Honolulu, HI




Figure B.7. OS4 Baseline test with all loads shown. TOU price schedule indicated by blue shading.

Table B.7. Summary from OS4 Scenarios

Scenario	Total Energy (kWh)	Energy Savings (kWh)	Daily Cost (\$)	Cost Savings (%)
Baseline	12.8	-	7.04	-
OS4 EE	9.4	3.4	4.11	41.6

Despite being an early summer day, the energy use over the day in Honolulu is much higher than the early summer day in Portland. The **foresee** controls implemented during the energy-efficiency case resulted in over 3 kWh of energy savings (over 25%) and more than 40% cost savings. Cost savings are typically higher than energy savings as

a percentage because **foresee** can shift energy use into lower cost periods of the day, which increases the cost savings relative to the energy savings. The accumulated energy for both cases is shown in Figure B.8, which illustrates that **foresee** was able to keep the accumulated energy nearly flat during the highest price period for the EE case. The operation of the battery in the OS4 scenario is shown in Figure B.9. The battery charged during the middle of the day when the electricity price was lowest and there was excess solar energy available. In addition to energy savings due to changes in the water heater and air conditioner set points, charging the battery in the middle of the day and discharging later in the day when electricity was more expensive resulted in significant cost savings and reduced the amount of energy that was backfed to the grid. As in other cases, the EE case resulted in higher accumulated air temperature discomfort than the baseline case. A summary of the net energy consumption of each case and the impact on daily cost and solar generation are given in Table B.7.




Figure B.8. Accumulated air temperature discomfort from the OS4 baseline and experimental cases. TOU price schedule indicated by blue shading.

Table B.8. DER and Comfort Impacts from OS4 Scenarios

Scenario	Total Air Temp Discomfort (°F-hrs)	Water Heater Comfort Impact (kWh)	Solar Energy Backfeed (kWh)	Battery Throughput (kWh)
Baseline	6.6	0	13.7	0.3
OS4 EE	16.0	0	10.9	17.9

Battery Power Consumption for OS4 Scenario: June day in Portland, OR




Figure B.9. Battery power consumption for OS4 optimization cases. Battery charging is considered positive power and discharging is considered negative. TOU price schedule indicated by blue shading.

Accumulated Air Temperature Discomfort for OS4 Scenarios: June day in Honolulu, HI




Figure B.10. Accumulated energy from the OS4 baseline and experimental case. TOU price schedule indicated by blue shading.

B.5 OS5 Scenario - Honolulu, HI, Under Shoulder Season Oversupply (Load-Add DR)

The OS5 scenarios demonstrate the ability of **foresee** to respond to a load-add demand response event on a spring day in Honolulu, HI with excess solar energy generation. The three tests, a baseline case, a case using **foresee** optimization, and a case using **foresee** with a load-add event between 10 a.m. and 4 p.m., were run using weather for March 20, 2011. Homeowner profile HS1 defined the thermostat set point, the hot water draw profile, and the schedule for the washer, dryer and dishwasher (See Tables A.6, A.8). Preference profile PS1 informed **foresee**'s controls.

The results from the baseline day are shown in the Figure B.11, broken down by the different end loads. The blue shading indicates the relative price of the time-of-use price schedule, with the darkest shading representing the most expensive period of time. For more detail on the TOU schedule, see Table A.2.




Figure B.11. OS5 Baseline test with all loads shown. TOU price schedule indicated by blue shading.

Table B.9. Summary from OS5 Scenarios

Scenario	Daily Energy (kWh)		DR Period (kWh)		Daily Cost	Cost
	Total	Saved	Energy Use	DR Service	(\$)	Savings (%)
Baseline	11.6	-	-8.4	-	7.10	-
OS5 EE	10.1	1.4	-4.6	3.9	4.56	35.8
OS5 DR (load-add)	10.9	0.7	1.7	10.1	4.29	39.5

Weather in Honolulu is fairly constant throughout the year, but this day in March used less energy than the early June day from the OS4 cases. This scenario was designed to look at the ability of **foresee** to add load during a demand-response event through the middle of the day. In Hawaii, the need for residents to use more energy during the middle of the day is driven by the high levels of distributed solar panels on residential buildings. The load-add demand response event overlapped with the lowest price period of the TOU schedule so the EE case saw an increase of energy use during the DR period, with the DR case further increasing the amount of energy used during the event

period, as shown in Table B.9. Despite the large increase in energy used during the DR period, both cases that used **foresee** control, decreased the overall energy use relative to the baseline case. Significant cost savings were seen in both cases with the EE case saving 35% and the DR case saving nearly 40% relative to baseline costs. The costs for all three cases were evaluated using the same TOU rate schedule, despite the fact that there is likely an additional financial incentive for customers to participate in demand-response event, so the cost savings for the DR case may be higher in real life. The DR case saw a small increase in energy use over the EE case, but resulted in a small decrease in cost savings.

The accumulated energy for all cases is shown in Figure B.12, which shows the effect that **foresee** had, especially during the DR period, denoted by the green shading. The accumulated energy of the EE case increases towards the end of the lowest price period, while the DR case was able to maintain a nearly flat accumulated energy during the DR period, resulting in an increase of 10 kWh during that 6-hour window relative to the baseline case. Most of the increase in energy use for both EE and DR cases during the middle of the day came from the operation of the battery, shown in Figure B.13. In the EE case, the battery charging ramped up towards the end of the lowest price period and then discharged during the highest price period. The DR case charged the battery at a higher power level during the DR period and then discharged during the high price period. For both cases, the use of the battery took advantage of the lowest cost electricity from the grid as well as using more of the on-site solar generation, which is likely the main way that the **foresee** cases saved a significant fraction of daily costs. As in other cases, the EE and DR cases resulted in higher accumulated air temperature discomfort than the baseline case. A summary of the net energy consumption of each case and the impact on daily cost and solar generation are given in Table B.9.




Figure B.12. Accumulated energy for the OS5 baseline, OS5 energy efficiency, and OS5 demand response cases. TOU price schedule indicated by blue shading. Load-add DR period shaded green.

Battery Power Consumption for OS5 Scenarios: March day in Honolulu, HI

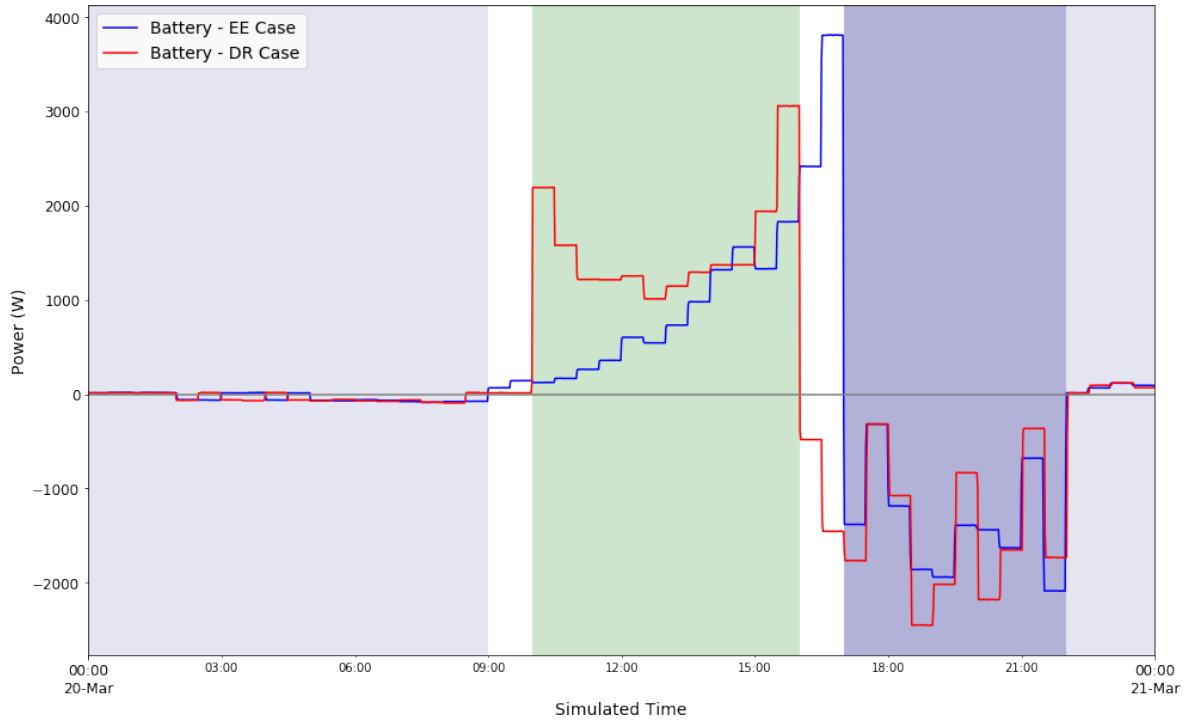


Figure B.13. Battery power consumption for OS5 optimization cases. Battery charging is considered positive power and discharging is considered negative. TOU price schedule indicated by blue shading, with DR period indicated with the green shading.




Figure B.14. Accumulated air temperature discomfort for the OS5 baseline, the energy-efficiency case, and the load-up demand response case. TOU price schedule indicated by blue shading, with DR period indicated with the green shading.

Table B.10. DER and Comfort Impacts from OS5 Scenarios

Scenario	Total Air Temp Discomfort (°F-hrs)	Water Heater Comfort Impact (kWh)	Solar Energy Backfeed (kWh)	Battery Throughput (kWh)
Baseline	5.5	0	14.5	0.3
OS5 EE	12.4	0	10.6	15.6
OS5 DR (load-up)	13.1	0	8.9	18.1

B.6 OS6 Scenario - Honolulu, HI, Under Solar Self-Consumption

The OS6 scenarios were designed to demonstrate a capability of **foresee** that meets a need that is currently unique to Hawaii—self-consumption of solar energy generation. Hawaii has a large number of distributed PV systems installed on homes, to the point that the utility has difficulty maintaining the required voltage levels during the day when solar generation is high. Currently, one solution to high levels local generation is to curtail the PV inverters if the feeder voltage gets too high (Nelson et al. 2016). The use of **foresee** control with a stationary battery and controllable loads could enable the home to self-consume the locally-generated power, thereby avoiding curtailing the solar generation. For this scenario, two cases—the baseline case and the EE case—were implemented with laboratory tests while the self-consumption case was performed in simulation. Each case used weather data from August 7, 2011 from Honolulu, HI. Homeowner profile HS1 defined the thermostat set point, the hot water draw profile, and the schedule for the washer, dryer and dishwasher (See Tables A.6, A.8). Preference profile PS1 was used to drive the **foresee** controls.

The results from the baseline day are shown in the Figure B.15, broken down by the different end loads. The blue shading indicates the relative price of the time-of-use price schedule, with the darkest shading representing the most expensive period of time. For more detail on the TOU schedule, see Table A.2.




Figure B.15. OS6 Baseline test with all loads shown. TOU price schedule indicated by blue shading.

Table B.11. Summary from OS6 Scenarios, with the results from the Solar Self-Consumption case coming from simulation.

Scenario	Total Energy (kWh)	Energy Savings (kWh)	Daily Cost (\$)	Cost Savings (%)
Baseline	14.1	-	7.56	-
OS6 EE	11.2	2.9	4.37	42.3
OS6 Solar Self-Consumption*	11.5	2.6	4.50	40.5

This scenario was designed to look at the ability of **foresee** to self-consume as much of the solar generation as possible. The baseline and EE case were performed through laboratory testing and the self-consumption case was executed in simulation. The EE case led to nearly 3 kWh of energy savings and over 40% cost savings, as shown in Table B.11. The self-consumption case has similar energy and cost savings compared to the EE case.

The accumulated energy for all cases is shown in Figure B.16, which shows the effect that **foresee** had. The accumulated energy of the EE case increases towards the end of the lowest price period due to increasing power consumption of the battery as shown in Figure B.17. The accumulated energy of the baseline case increases steadily in the morning and then decreases significantly during the day due to the solar energy backfeeding. In contrast, the accumulated energy of the self consumption case stayed almost constant during the day, illustrating the effect of **foresee** in solar self-consumption.

In the EE case, the battery charging ramped up towards the end of the lowest price period and then discharged during the highest price period. As in other cases, the EE case resulted in higher accumulated air temperature discomfort than the baseline case. A summary of the net energy consumption of each case and the impact on daily cost and solar generation are given in Table B.11.




Figure B.16. Accumulated energy for the OS6 baseline and OS6 energy-efficiency cases, with the results from the Solar Self-Consumption case coming from simulation. TOU price schedule indicated by blue shading.

Battery Power Consumption for OS6 Scenario: August day in Honolulu, HI




Figure B.17. Battery power consumption for OS6 optimization cases, with the results from the Solar Self-Consumption case coming from simulation. Battery charging is considered positive power and discharging is considered negative. TOU price schedule indicated by blue shading.

Accumulated Air Temperature Discomfort for OS6 Scenarios: August day in Honolulu, HI

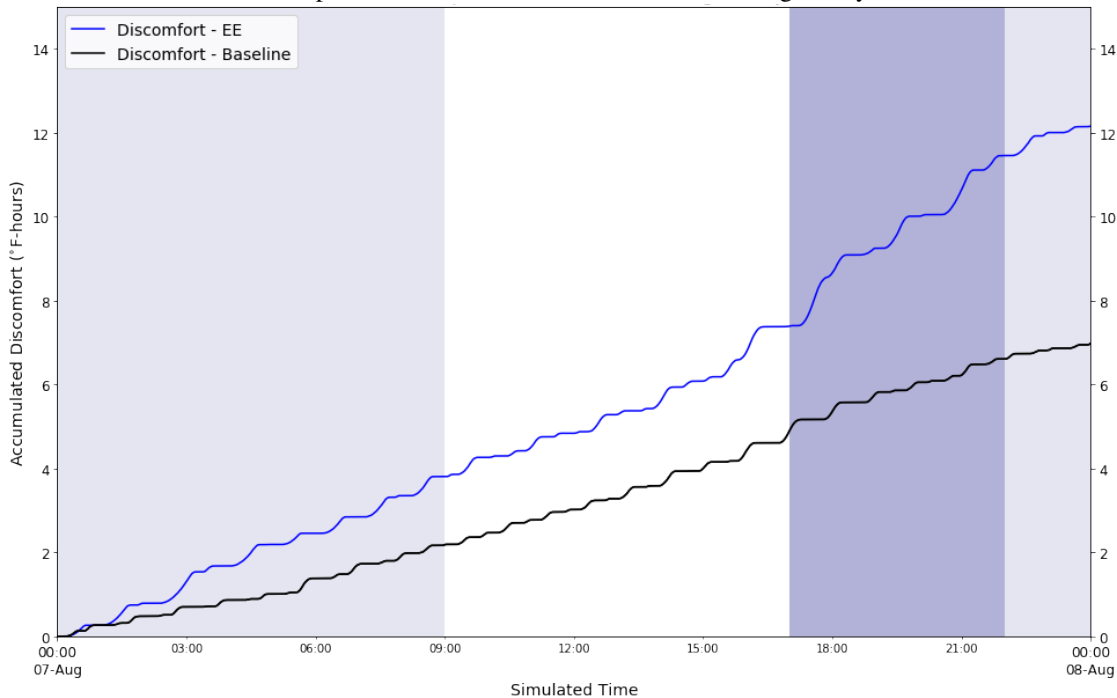


Figure B.18. Accumulated air temperature discomfort for the OS6 baseline and the energy-efficiency case. TOU price schedule indicated by blue shading.

Table B.12. Daily Comfort and DER Impacts from OS6 Scenarios, with the results from the Solar Self-Consumption case coming from simulation.

Scenario	Air Temp Discomfort (°F-hrs)	Water Heater Comfort Impact (kWh)	Solar Energy Backfeed (kWh)	Battery Throughput (kWh)
Baseline	7.0	0	13.6	0.3
OS6 EE	12.2	0	10.6	18.1
OS6 Solar Self-Consumption*	2.0	0	0.7	18.8

B.7 OS7 Scenario - Spokane, WA, Under Normal Operation

The OS7 scenarios were designed to show the impact of **foresee** control on an early summer day in Spokane, WA. The two days of tests, a baseline day and a day using **foresee** optimization, were run using weather for June 4, 2014. Homeowner profile HS1 defined the thermostat set point, the hot water draw profile, and the schedule for the washer, dryer and dishwasher (See Tables A.6, A.8). Preference profile PS1 was used to drive the **foresee** controls.

The results from the baseline day are shown in the Figure B.19, broken down by the different end loads. The blue shading indicates the relative price of the time-of-use price schedule, with the darkest shading representing the most expensive period of time. For more detail on the TOU schedule, see Table A.2.

Circuit-level Power Consumption for OS7 baseline case: June day in Spokane, WA




Figure B.19. OS7 Baseline test with all loads shown. TOU price schedule indicated by blue shading.

The OS7 scenarios demonstration normal operation for **foresee** on an early summer day, with some need for cooling. The baseline energy use is relatively low and there was energy saved by **foresee** in the EE Case, but both the energy and cost savings are small. The accumulated energy for both cases is shown in Figure B.20 and the operation of the battery in the EE case is shown in Figure B.21. The battery charged before the TOU price increased at 9 a.m.

Table B.13. Summary from OS7 Scenarios

Scenario	Total Energy (kWh)	Energy Savings (kWh)	Daily Cost (\$)	Cost Savings (%)
Baseline	3.4	-	1.08	-
OS7 EE	3.0	0.4	1.03	4.4

and discharged during the middle and highest cost periods of the day. The modest energy savings mostly came from changes in the water heater and air conditioner set points, charging the battery in the middle of the day and discharging later in the day when electricity was more expensive resulted in significant cost savings and reduced the amount of energy that was backfed to the grid. Both the baseline and EE case resulted in a small amount of air temperature discomfort, as shown in Figure B.22. In this case, the discomfort calculation was only performed between 12 p.m. and midnight as cooling was not needed before noon. A summary of the net energy consumption of each case and the impact on daily cost and solar generation are given in Table B.13.




Figure B.20. Accumulated energy for the OS7 baseline and OS7 energy-efficiency cases. TOU price schedule indicated by blue shading.

Table B.14. DER and Comfort Impacts from OS7 Scenarios

Scenario	Total Air Temp Discomfort (°F-hrs)	Water Heater Comfort Impact (kWh)	Solar Energy Backfeed (kWh)	Battery Throughput (kWh)
Baseline	3.3	0	17.7	0.3
OS5 EE	5.9	0	18.7	10.0

Battery Power Consumption for OS7 Scenario: June day in Spokane, WA




Figure B.21. Battery power consumption for OS7 optimization cases. Battery charging is considered positive power and discharging is considered negative. TOU price schedule indicated by blue shading.

Accumulated Air Temperature Discomfort for OS7 Scenarios: June day in Spokane, WA




Figure B.22. Accumulated air temperature discomfort for the OS7 baseline and the energy-efficiency case. TOU price schedule indicated by blue shading.

B.8 OS8 Scenario - Spokane, WA, Under Peak Hot Day (Load-Shed DR)

The OS8 scenarios demonstrate the ability of **foresee** to respond to a load-shed demand response event on a hot summer day in Spokane, WA. A total of four tests were conducted on this simulated day: a baseline case, a case using **foresee** optimization with a TOU utility rate, a case using **foresee** optimization with a flat utility rate, and a case using **foresee** with a load-shed event between 2 and 8 p.m. on top of a TOU rate. Each test case was run weather for July 29, 2014. Homeowner profile HS1 defined the thermostat set point, the hot water draw profile, and the schedule for the washer, dryer and dishwasher (See Tables A.6, A.8). Preference profile PS1 was used to drive the **foresee** controls.

The results from the baseline day are shown in the Figure B.23, broken down by the different end loads. The blue shading indicates the relative price of the time-of-use price schedule, with the darkest shading representing the most expensive period of time. For more detail on the TOU schedule, see Table A.2.




Figure B.23. OS8 Baseline test with all loads shown. TOU price schedule indicated by blue shading.

Table B.15. Summary from OS8 Scenarios

Scenario	Daily Energy (kWh)		DR Period (kWh)		Daily Cost (\$)	Cost Savings (%)
	Total	Saved	Energy Use	DR Service		
Baseline - flat	17.9	-	7.6	-	2.29	-
Baseline -TOU					1.91	-
OS8 EE - flat	9.8	8.1	4.0	3.6	1.56	31.6
OS8 EE - TOU	16.2	1.7	0.3	7.2	1.71	10.8
OS8 TOU + DR (load-shed)	15.4	2.5	-0.9	8.5	1.57	17.8

This scenario used weather from the hottest day in Spokane in 2014, with a maximum outdoor temperature of 99°F. Hot summer days stress the electric grid because air conditioning loads are much higher and concentrated in the

late afternoon and evening. This scenario was used to demonstrate **foresee**’s ability to provide load-shed during the hottest part of the day: 2–8 p.m. The demand response period overlaps with the highest price period of the TOU period so one additional test was run to see how the the energy-efficiency optimization varied when a flat utility rate is used versus a TOU utility rate. In all cases, energy was saved relative to the baseline case and cost was also saved in all cases. Also, all **foresee**-controlled cases led to a reduction of energy use during the DR period, with the DR case producing the largest reduction in DR period energy use, as shown in Table B.15. The EE flat rate case led to the largest energy and cost savings. The DR TOU case resulted in nearly 18% cost savings, followed by the EE TOU case that produced over 10% cost savings. The costs for all three cases were evaluated using the same TOU rate schedule, despite the fact that there is likely an additional financial incentive for customers to participate in demand-response event, so the cost savings for the DR case may be higher in real life. The DR case delivered additional energy and cost savings relative to the EE TOU case (the EE flat rate case should not be compared to the DR case because the TOU rate drives the **foresee** controls).

The accumulated energy for all cases is shown in Figure B.24. The two cases that used the TOU rate (EE and DR) both charged the battery before the medium price period began and discharged during the medium and high price periods (see Figure B.25). Neither the baseline case or the EE flat rate case used the battery at all. The EE flat rate case was able to acheive significant energy savings by reducing the amount of air conditioning cycling, which also contributed to a high level of accumulated air temperature discomfort, as shown in Figure B.26. The EE TOU and DR TOU cases also experienced higher levels of air temperature discomfort than the baseline case, but the air conditioner actually used slightly more energy. Both cases tended to let the temperature drift above the upper comfort bound and then would cool beyond the lower comfort bound, which led to additional cooling energy use. Most of the energy savings in the EE TOU and DR TOU cases came from energy savings from the water heater. The battery usage in those cases contributed to higher cost savings relative to the energy savings. See Table B.15 for the summary of additional impacts in each case.




Figure B.24. Accumulated energy for the OS8 baseline case, OS8 energy efficiency using a flat utility rate, OS8 energy efficiency with a TOU utility rate, and OS8 demand response with TOU pricing. TOU price schedule indicated by blue shading. Load-shed DR period shaded red.

Battery Power Consumption for OS8 Scenarios: July day in Spokane, WA




Figure B.25. Battery power consumption for OS8 optimization cases. Battery charging is considered positive power and discharging is considered negative. TOU price schedule indicated by blue shading, with DR period indicated with the red shading.

Accumulated Air Temperature Discomfort for OS8 Scenarios: July day in Spokane, WA




Figure B.26. Accumulated air temperature discomfort for the OS8 baseline, the energy efficiency with flat rate case, energy efficiency with TOU rates, and the load-shed demand response case. TOU price schedule indicated by blue shading, with DR period indicated with the red shading.

Table B.16. DER and Comfort Impacts from OS8 Scenarios

Scenario	Total Air Temp Discomfort (°F-hrs)	Water Heater Comfort Impact (kWh)	Solar Energy Backfeed (kWh)	Battery Throughput (kWh)
Baseline	11.6	0	11.8	0.3
OS8 EE - flat	46.4	0	10.7	0.3
OS8 EE - TOU	29.7	0	13.0	16.9
OS8 TOU + DR (load-shed)	28.2	0	13.7	17.8

B.9 OS9 Scenario - Spokane, WA, Under Peak Cold Day (Load-Shed DR)

The OS9 scenario demonstrate the ability of **foresee** respond to a load-shed demand response event on a cold winter day in Spokane, WA. A total of three tests were conducted on this simulated day: a baseline case, a case using **foresee** optimization with a TOU utility rate, a case using **foresee**'s optimization with a load-shed event between 6 and 10 a.m. on top of a TOU rate. Each test case was run using weather for November 14, 2014. Homeowner profile HS1 defined the thermostat set point, the hot water draw profile, and the schedule for the washer, dryer and dishwasher (See Tables A.6, A.8). Preference profile PS1 was used to derive the **foresee** controls.

This scenario used weather from a cold day in Spokane in 2014, with a lowest outdoor temperature of 16°F. Cold winter days stress the electric grid because space heating loads are much higher and concentrated in the early morning. Similar to OS1, this scenario was used to demonstrate **foresee**'s ability to provide load-shed during time period when heating loads are the highest. The demand response period overlaps with the medium price period of the TOU period. As shown in Table B.17, the EE and DR cases have small energy savings compared to the baseline case, but both cases yielded moderate cost savings by shifting the loads to the time periods with lower energy prices. Compared to the EE case, the DR case delivered over 11 kWh energy reduction during the DR period at the cost of lower energy savings and cost savings.

Figure B.27 shows the accumulated energy for all three cases. The accumulated energy in the DR case increased much more quickly than the other cases between 4 and 6 a.m. after **foresee** received the notification of the DR event and stayed almost constant throughout the entire DR period. Both EE and DR cases also used much less energy than the baseline case during the high-price period between 2 and 6 p.m. to save energy cost.

The EE and DR use cases also have similar DER and comfort impacts. Both cases have slightly higher air temperature discomfort compared to the baseline case but have near zero energy backfeed. Battery was heavily used in the EE and DR cases. The battery throughput in the EE case was 15 kWh, which was almost equivalent to a full charging and discharging cycle. The battery was used more heavily in the DR case as its throughput doubles the EE case.

Table B.17. Summary from OS9 Scenarios

Scenario	Daily Energy (kWh)		DR Period (kWh)		Daily Cost (\$)	Cost Savings (%)
	Total	Saved	Energy Use	DR Service		
Baseline	60.30	-	9.08	-	5.20	-
OS9 EE	59.31	0.99	11.42	-2.34	4.69	9.67%
OS9 DR (load-shed)	60.48	-0.18	0.34	8.74	4.91	5.51%

Accumulated Energy for OS9 Scenarios: November day in Spokane, WA




Figure B.27. Accumulated energy from the OS9 simulation for baseline, energy efficiency and demand response cases. TOU price schedule indicated by blue shading and red shading indicates load-shed DR period.

Table B.18. DER and Comfort Impacts from OS9 Scenarios

Scenario	Total Air Temp Discomfort (°F-hrs)	Water Heater Comfort Impact (kWh)	Solar Energy Backfeed (kWh)	Battery Throughput (kWh)
Baseline	0.08	0.00	1.75	0.00
OS9 EE	0.25	0.00	0.09	15.11
OS9 DR (load-shed)	0.29	0.00	0.09	31.43

B.10 Summary of Results

B.10.1 Energy Efficiency Results

A summary of measured and simulated results from the baseline experiments and the energy-efficiency use case experiments is provided in Table B.19. Key observations are that large energy and cost savings were achieved, partially at a comfort cost. It is unclear why the comfort disutility is underweighted, resulting in this apparently negative outcome. Additional research should be performed to correct this weakness. The comfort metrics shown in the following are accumulated over the entire day, unless otherwise noted.

Table B.19. Summary of Energy-Efficiency Mode Energy and Comfort Performance

HBS Scenario Case	Baseline			Energy Efficiency (savings vs. Baseline) [‡]				
	Energy Use (kWh)	Cost (\$)	Discomfort (Air, °F-hr)	Energy Use (kWh)	Cost (\$)	Discomfort (Air, °F-hr)	Savings	
							E (kWh)	Cost (\$)
OS1a ^a	0.36	0.93	1.6	-3.56	0.68	3.1	3.33	0.25 (24%)
OS1b ^a				-2.07	0.75	5.8	1.87	0.18 (19%)
OS1c ^a				-2.13	0.75	1.8	1.87	0.18 (19%)
OS2*	39.58	3.84	0.02	39.41	2.89	0.44	0.17	0.95 (25%)
OS3*	4.98	1.16	0.00	3.23	0.51	2.61	1.76	0.65 (56%)
OS4 ^b	12.75	7.04	6.6	8.70	4.11	16.0	3.37	2.93 (42%)
OS5	11.57	7.10	5.5	9.40	4.56	12.4	1.43	2.54 (36%)
OS6	14.10	7.56	7.0	10.52	4.37	12.1	2.90	3.20 (42%)
OS7	3.40	1.09	3.3	2.33	1.04	5.7	0.37	0.05 (4%)
OS8	17.85	1.91	11.6	15.38	1.71	29.7	1.70	0.21 (11%)
OS8-2 [†]	17.85	2.29	11.6	9.76	1.56	46.4	8.10	0.73 (32%)
OS9*	60.30	5.20	0.08	59.31	4.69	0.25	0.99	0.50 (10%)

^a Accumulated discomfort was calculated between 3 p.m. and 12 a.m.

^b Accumulated discomfort was calculated between 12 p.m. and 12 a.m.

* OS2, OS3, OS9 were performed in simulation due to budget and schedule constraints.

[†] OS8 was repeated with flat rate to study the impact of tariff on control outcomes.

[‡] Note that values may not directly equate, due to rounding and correcting for building state from experiment start to end.

B.10.2 Demand Response Results

BPA considers technologies as “efficiency first” which means that demand response service is calculated with a baseline of the **foresee**-controlled home as opposed to the home absent **foresee**. Because other utilities or entities may use a different baseline, we provide DR performance results for both baselines in Table B.20. DR performance against the no-**foresee** baseline are shown on the left, and DR performance against the **foresee** EE baseline are both provided.

It is clear that a modest overall energy takeback may occur on days when DR is called, versus the Energy-Efficiency case, to deliver the demand response service. However, that is not uniformly true, and is dependent on the homeowner’s preferences, the equipment and home, the weather, and the utility rate. It is likely that rate design or incentive could be tailored to elicit demand response service without sacrificing any energy efficiency, but we were unable to develop such rate guidelines in this work, and leave it for future research.

Table B.20. Summary of Demand Response Mode Energy and Comfort Performance, versus the Energy-Efficiency Case

Scenario	Daily DR Case			DR vs. Baseline			DR vs. Energy Efficiency		
	Energy (kWh)	Cost‡ (\$)	Discomfort (°F-hr)	Service (kWh)	Daily Savings		Service (kWh)	Daily Savings	
					E (kWh)	Cost (\$)		E (kWh)	Cost (\$)
OS2*	39.42	2.90	0.44	7.57	0.16	0.94 (25%)	0.03	-0.01	-0.01 (0%)
OS3*	19.29	1.57	2.61	24.79†	-14.30	-0.41 (-35%)	19.26†	-16.06	-1.06 (-207%)
OS5	10.13	4.30	13.14	10.15†	0.70	2.80 (40%)	6.26†	-0.73	0.27 (6%)
OS8	14.59	1.57	28.16	8.53	2.49	0.34 (18%)	1.28	0.80	0.14 (8%)
OS9*	60.48	4.91	0.29	8.74	-0.18	0.29 (6%)	11.08	-1.17	-0.22 (-5%)

* OS2, OS3, & OS9 were performed in simulation due to budget and schedule constraints.

† These values are load-add (DEC) DR service; all others are load-shed (INC).

‡ The incentives for participating in the DR program were not included in the cost calculation.

B.10.3 DR Service Uncertainty

The DR results shown earlier assume perfect knowledge of the weather conditions (outdoor temperature and solar irradiance) and occupancy (hot water draw schedules and uncontrollable load forecasts), and the sample size of laboratory experiments is too small to confirm that **foresee**'s DR service forecast was significant. The DR forecasts from a **foresee**-controlled home were validated by running 50 Monte Carlo simulations for Scenario OS9 where the uncertainty on weather forecasts and occupancy were assumed to have the following probability distributions:

- Outdoor Temperature Forecast Error (°C): $N(0, 0.5)$,
- Solar Irradiance Forecast Error (kW/m^2): $N(0, 0.1)$,
- Hot Water Draw Forecast Error (kg/sec): $N(0, 0.0025)$,
- Uncontrollable Load Forecast Error (kW): $N(0, 0.1)$

where $N(\mu, \sigma)$ denotes a normal distribution with mean μ and standard deviation σ . In other work with MPC-based home energy management systems, similar to **foresee**, Gaussian distributions have been used to describe uncertainty in weather forecasts (Oldewurtel et al. 2012), and are used almost universally in uncertainty representation for home energy management research and development.*

The average load reduction relative to the EE case during the DR period is shown in Figure B.28 for each of the Monte Carlo simulations and the corresponding DR case with no uncertainty in the weather forecast or occupancy. A histogram of the of average load reduction during the DR period is provided in Figure B.29. The load reduction under perfect forecast was 2.77 kW over the 4-hour DR period, and the delivered load reduction among the Monte Carlo simulations was between 2.70 kW and 2.86 kW, which is within 10% of the case with perfect forecast.

*We feel there is a significant opportunity for future research in development of large-scale data sets which can be used to improve on the engineering-judgment assumptions such as on occupancy profiles and occupancy uncertainty.




Figure B.28. Average load reduction during the DR period relative to the EE case for OS9 Scenario Monte Carlo simulations.




Figure B.29. Histogram of the average load reduction relative to the EE case during the DR period for the OS9 Scenario Monte Carlo Simulations.

C Detailed Penetration Test Results

C.1 Overview

To confirm the cybersecurity risk assessment and implementation, NREL performed a series of penetration tests on the hardened **foresee** platform described in Section 7. The **foresee** application and cybersecurity layer were installed on a Linux NUC which had been hardened according to the previously mentioned recommendations.

C.2 Objective

The **foresee** platform is a prototype, not an operational system. On an operational system, the goal of a penetration test is to identify vulnerabilities in the system and develop mitigations to protect operational assets from compromise. (For instance, if weak or default passwords are in use, switch to strong passwords.)

On a prototype system, the goal of penetration testing should be to explore the security characteristics of the system under development to develop a list of recommended controls for production versions of the system. There are no operational assets to protect, but there is a need to understand what is needed to protect future operational assets. This subtle difference impacts both the assumptions made going into the test and the design of the test itself.

C.3 Scope

The scope of this test is to evaluate the prototype cybersecurity layer's ability to withstand attacks identified by the Risk Assessment, and to generate recommendations for future development of the **foresee** platform if it becomes commercialized. The tests performed have resulted in a detailed list of actionable items that can be implemented on the **foresee** platform in the future.

We repeat table 26 from section 8, page 58, here for convenience, followed by penetration test findings for each of the cybersecurity penetration tests.

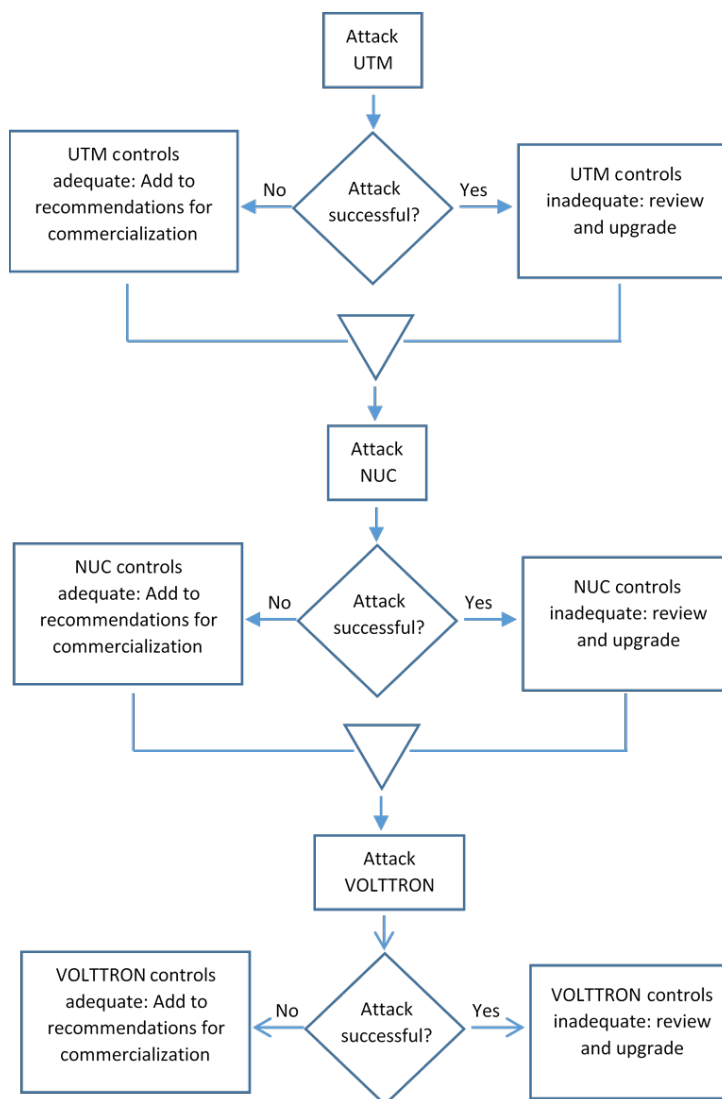


Figure C.1. Flowchart showing the penetration test process.

Table C.1. Penetration Tests Performed

#	Scenario Name	Description	Threat Analysis Reference	Pen Test Activities	Test Success Metrics (i.e. Home Battery System Failure Criteria)
T1	Disrupt foresee services via internal Denial of Service (DoS)	Disrupt foresee platform. Assumes access to the customer LAN.	AV-ATK-1	On LAN: - DoS	Good: Slow operation of foresee platform Best: Cause foresee platform to crash
T2	Platform Operation Disruption (via customer LAN)	Alter behavior of the foresee platform. Assumes access to the customer LAN.	AU-ATK-5	On LAN: - Packet capture and spoofing (MiM)	Good: spoof data from appliances Better: spoof commands
T3	Attack foresee platform from customer LAN	Get access to the foresee platform from the customer LAN.	AU-ATK-5	On foresee platform: - Port scanning - Password guessing - Vulnerability scans	Good: gather quality data about foresee platform Better: retrieve foresee credentials
T4	foresee platform Data Disruption	Steal, disrupt, or damage foresee platform data. Assumes foresee platform access (CS2).	C-ATK-2 AU-ATK-3	On foresee platform: - Escalate privileges	Good: log onto foresee platform, access data Better: log onto foresee platform, alter data
T5	Attack communication channel outside the UTM	Attempt to deny communications from the public Internet to the customer LAN	AU-ATK-3	Outside UTM: - Physical disconnect of communication - DoS on network	Good: Disrupt communications via DoS
T6	Attack on VOLTTRON	Probe for gaps in VOLTTRON's security implementation		From inside the foresee platform: - Check for outdated versions of OpenSSL. - Attempt to connect via depreciated connection protocol - Check for key misconfiguration on OMQ and CurveZMQ	Good: Find outdated or misconfigured system software Better: Exploit outdated or misconfigured system software

C.4 T1 - Foresee Service Disruption via Internal DoS

#	Scenario Name	Description	Risk Analysis Reference	Pen Test Activities	Test Success Metrics (i.e. Home Battery System Failure Criteria)
T1	Disrupt foresee services via internal Denial of Service (DoS)	Disrupt foresee platform. Assumes access to the customer LAN.	AV-ATK-1	On LAN: - DoS	Good: Slow operation of foresee platform Best: Cause foresee platform to crash

The intention of T1 was to cause a denial of service to impair the availability of the **foresee** application or to interrupt **foresee**'s service. To perform this task various pen testing tools were utilized to attempt to impact **foresee**'s performance by overwhelming the application with extensive network traffic. The test was performed from within the LAN network containing the **foresee** application. Multiple tools were able to cause the operation of the **foresee** platform to slow down or become inaccessible altogether.

The attacks that were attempted against **foresee** for this test included:

- ICMP flooding
- TCP SYN flood
- HTTP request flooding

C.4.1 Penetration Testing Tools

Hping3 is a tool designed for firewall testing, network testing and more. For this test Hping was used to perform a TCP and an ICMP network flood. A TCP network flood aims to interrupt the remote service by creating more active sessions than the service is capable of handling. In this situation a remote application can crash or be unable to sustain legitimate sessions from users attempting to access **foresee**.

The ICMP network flood aims to send a significant amount of ICMP requests which may cause the receiving device to exhaust itself from trying to reply to each request. This can make the receiving device crash or make its services unavailable.

HULK The HTTP Unbearable Load King (HULK) is a tool to test web applications using the HTTP protocol. HULK was used to perform a flood of HTTP traffic to the **foresee** appliance. Because the **foresee** application answers HTTP requests sent to it, this could exhaust the application's ability to further communicate to legitimate requests. HULK ensures each request is uniquely generated and mimics a legitimate HTTP request in an effort to bypass firewalls, Intrusion Prevention Systems (IPS) and caching engines in order to hit the web application's direct resource pool.

PyLoris is designed to test a remote service's vulnerability to connection exhaustion. PyLoris opens a connection the remote application and sends slow reply messages to the server. When multiplied across many concurrent connections this attack can cause a remote service to crash or become unavailable due to keeping the server in a listening state as it waits for the clients to send their slow response.

C.4.2 Results

C.4.2.1 ICMP and TCP SYN floods

When the attack is launched at **foresee** with no intermediary devices both the ICMP flood and the TCP flood both seemed equally effective against disabling the user interface for **foresee**, but only while the flooding attacks were taking place.

During either flooding attack a legitimate user will not be able to adjust settings to the **foresee** application or be able to review the **foresee** application to inspect if it is working as desired by the end user.

Procedure: The flood attack took place from within the same Local Area Network (LAN) as the **foresee** appliance. A test client started a continuous ping against the **foresee** appliance and verified that prior to the attack it was able to access the **foresee** interface. The penetration tester's laptop (running Kali Linux) ran the following command to execute an ICMP flooding attack:

```
» hping3 -flood -q -1 -I eth0 -d 10240 192.168.1.101
```

The test client attempted to access the **foresee** appliance again. After several minutes the **foresee** page never loaded in its entirety in the test client's browser, as shown in figure C.2.

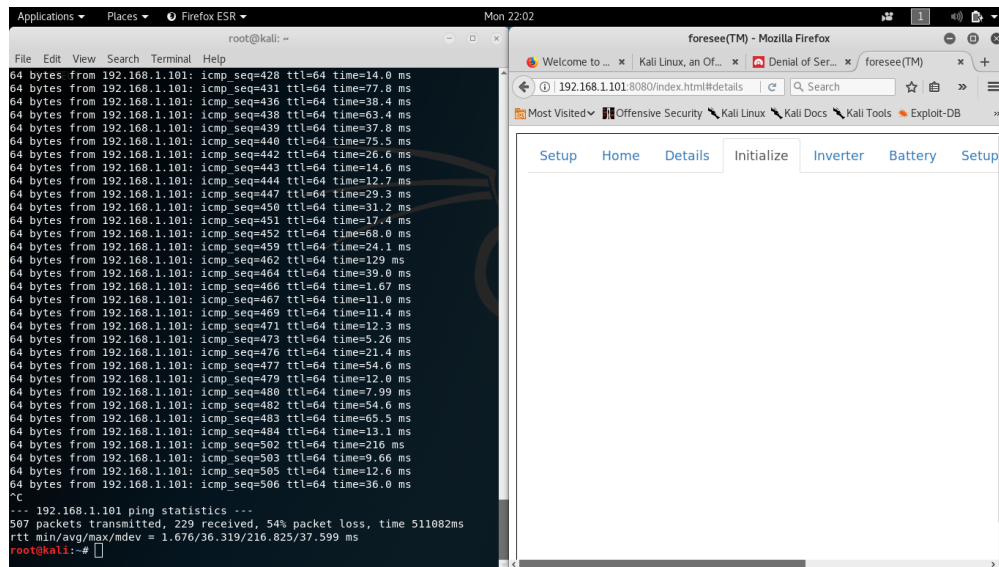


Figure C.2. Screenshot showing the test client's attempt to access the **foresee** appliance. The left side is the result of a continuous ping showing that the client experienced 54% packet loss over 4 minutes. The right side shows the browser window where the **foresee** appliance failed to load in its entirety.

The flood attack was halted on the penetration tester's laptop, after which the test client resumed full functionality to the **foresee** appliance.

The penetration tester's laptop (running Kali Linux) ran the following command to execute a TCP SYN flood attack:

```
» hping3 -d 120 -S -w 64 -p 8080 -flood -rand-source 192.168.1.101
```

The test client attempted to access the **foresee** appliance again. After several minutes the **foresee** page failed to present any content to the user, as shown in figure C.3

C.4.2.2 HTTP Request Flooding

When the attack is launched at **foresee** with no intermediary devices the HTTP Request flooding attack generated by HULK was effective against disabling the user interface for **foresee**, but only while the flooding attacks were taking place.

During the flooding attack a legitimate user will not be able to adjust settings to the **foresee** application or be able to review the **foresee** application to inspect if it is working as desired by the end user.

HULK seemed to be at least equal to the other Denial of Service (DoS) attacks mentioned under this test. Unlike the other attacks under this test, the HTTP Request flood had the lowest interference with the network itself as it relied on exhausting resources of the application rather than the network.




Figure C.3. Screenshot showing the test client’s attempt to access the foresee appliance. The left side is the result of a continuous ping showing that the client experienced 45% packet loss over 6 minutes. The right side shows the browser window where the foresee user interface failed to load.

Procedure: The attack took place from within the same Local Area Network (LAN) as the **foresee** appliance. A test client started a continuous ping against the **foresee** appliance and verified that prior to the attack it was able to access the **foresee** interface.

The penetration test laptop (running Kali Linux) used the HULK tool from the following location: <https://github.com/grafov/hulk>. The following command executes an HTTP Request flooding attack:

```
» python hulk.py http://192.168.1.101:8080
```

The test client attempted to access the **foresee** appliance again. After several minutes the **foresee** page never loaded in the test client’s browser, as shown in figure C.4.

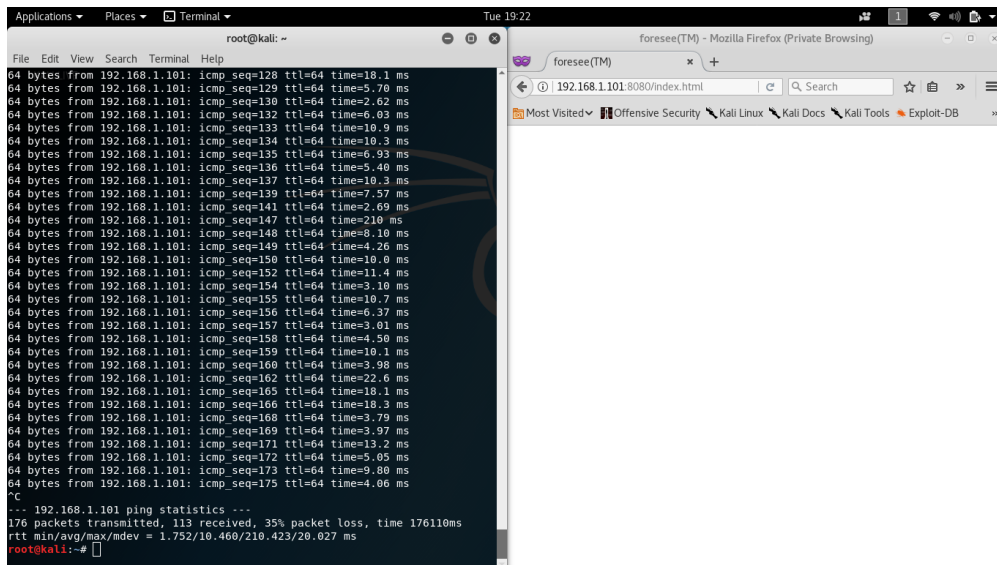


Figure C.4. Screenshot showing the test client’s attempt to access the foresee appliance. The left side is the result of a continuous ping showing that the client experienced 35% packet loss over 3 minutes. The right side shows the browser window where the foresee appliance failed to load.

C.4.2.3 Connection Exhaustion

PyLoris deemed to be ineffective against the **foresee** application. The test client was able to access the **foresee** application without any apparent differences.

Procedure: The attack took place from within the same LAN as the **foresee** appliance. A test client started a continuous ping against the **foresee** appliance and verified that prior to the attack it was able to access the **foresee** interface. The penetration tester's laptop downloaded the PyLoris application from the following source: <https://sourceforge.net/p/pyloris/wiki/Home/>. The penetration tester's laptop leveraged an attack through the graphical interface, as shown in figures C.5 and C.6.

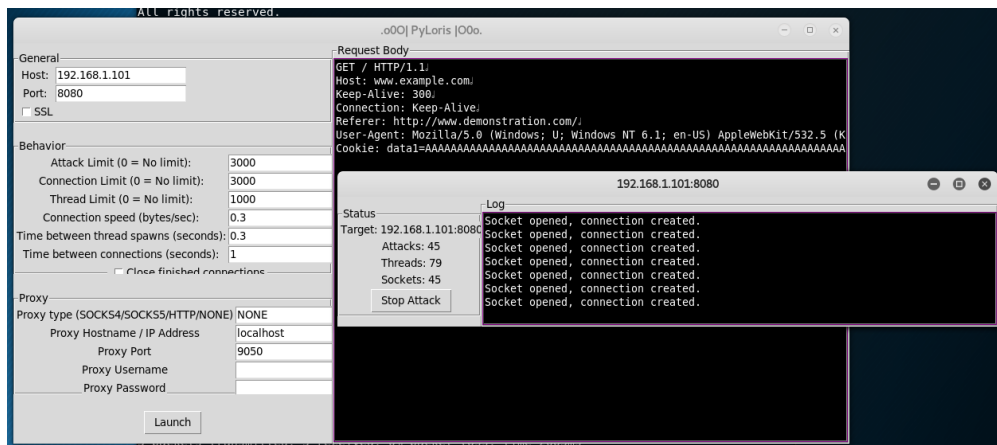


Figure C.5. Screenshot showing the penetration tester's laptop as it uses the PyLoris graphical interface to launch an attack against the foresee application.

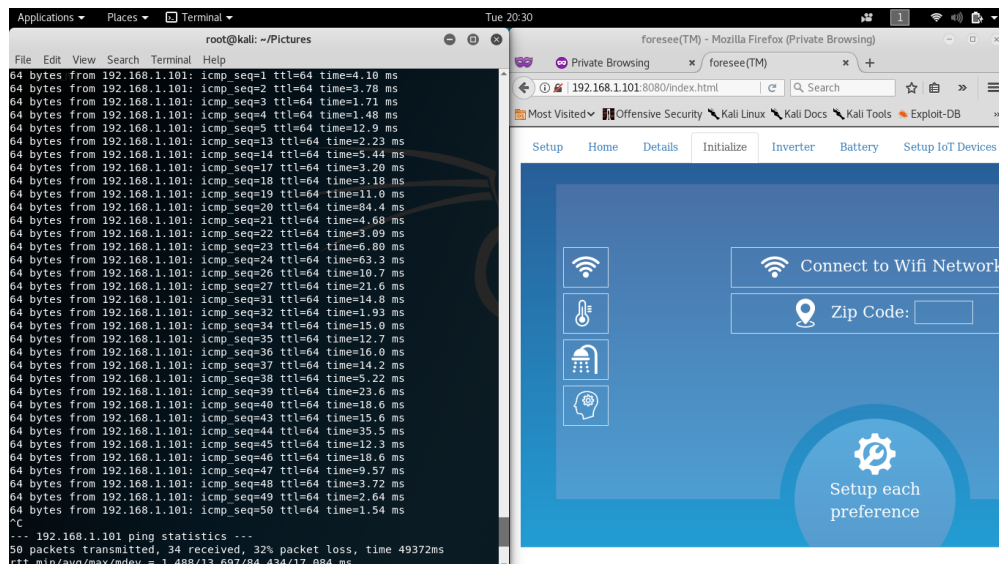


Figure C.6. Screenshot showing the test client's attempt to access the foresee appliance. The left side is the result of a continuous ping showing that the client experienced 32% packet loss over 1 minute. The right side shows the browser window where the foresee appliance loaded successfully.

C.4.3 T1 Conclusions

Several DoS attacks appeared to be effective at making the **foresee** application inaccessible to a legitimate user while the flood was running. However, none of the attacks caused the **foresee** application to crash indefinitely. Whenever the attack was stopped, **foresee** returned to its normal working state without requiring any administrative maintenance or repairs. All these attacks ran within the same LAN as the **foresee** appliance. Because the ideal implementation of **foresee** will include a Universal Threat Management (UTM) device, it is possible that some of these attacks may not be as effective.

The HTTP Request flood attack generated by HULK may be the most likely to remain effective if there were an intermediary UTM on the network due to it relying more so on exhausting application resources rather than network resources.

RISK = MODERATE

The application was made inaccessible to legitimate users for as long as the flood attacks were running, but whenever the attack stopped the system returned to a working state with no administrative maintenance or repairs required. As the **foresee** platform never crashed the metric for high risk was not met. As the operation of the **foresee** platform was brought to a halt during the attacks the metric for moderate risk was exceeded.

C.5 T2 - Platform Operation Disruption (via Customer LAN)

#	Scenario Name	Description	Threat Analysis Reference	Pen Test Activities	Test Success Metrics (i.e. Home Battery System Failure Criteria)
T2	Platform Operation Disruption (via customer LAN)	Alter behavior of the foresee platform. Assumes access to the customer LAN.	AU-ATK-5	On LAN: - Packet capture and spoofing (MiM)	Good: spoof data from appliances Better: spoof commands

The intention of T2 was to identify vulnerabilities in spoofing communication between the **foresee** application and other devices. If communication could be spoofed by an attacker, a Man in the Middle (MiM) attack could send false communication to the **foresee** application.

C.5.1 Penetration Testing Tools

TCP Dump was used to collect network traffic communicating with the **foresee** appliance.

Wireshark was used to read and dissect the packet captures generated from TCPDump.

C.5.2 Results

Of the communication identified between the **foresee** application and the IoT devices, all communication was encrypted with a strong encryption protocol that the penetration tester was unable to compromise.

Procedure: The task involved collecting a packet capture from the **foresee** device as well as from a penetration testing laptop inside the customer LAN. The packet capture ran for multiple days to ensure as much critical communication as possible was recorded. Upon collecting the network traffic, a thorough inspection was performed to identify unique communications taking place with the **foresee** application. This is shown in figure C.7.

At the time the test was performed, no vulnerabilities in the encryption protocol (TLSv1.2) could be found.

C.5.3 T2 Conclusions

The encryption used by the **foresee** application to communicate with IoT devices was considered secure. No communication between the **foresee** application and an IoT device could be compromised.

RISK = UNDISCOVERED

Communication from the **foresee** application could not be compromised by the penetration test.

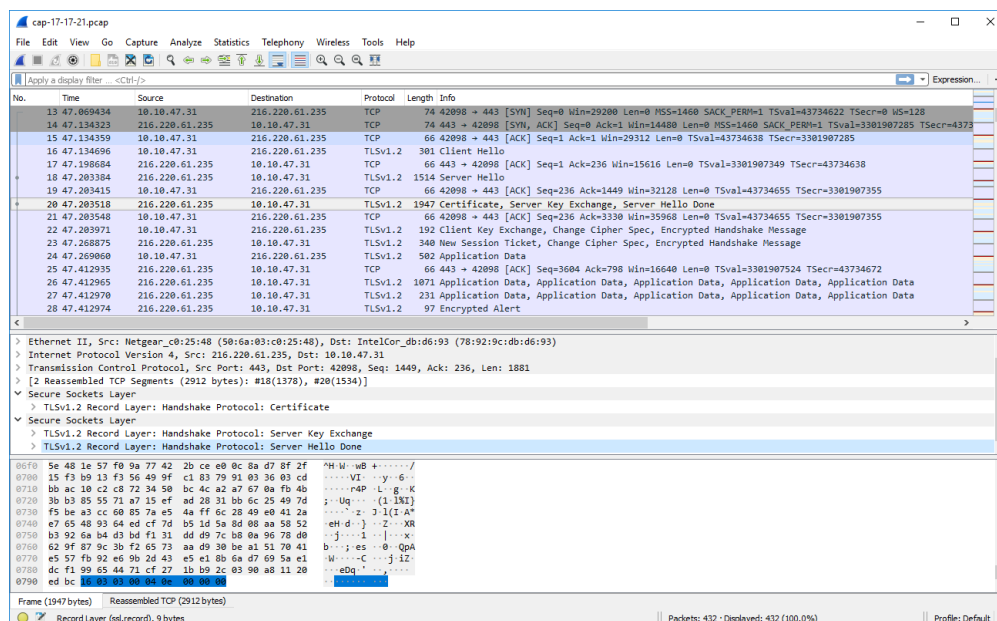


Figure C.7. Screenshot showing packet capture of communication between the **foresee** application and an IoT device. A secure connection is established and a secure encryption protocol is in use. The collected communication traffic is obfuscated preventing a penetration tester from spoofing communications to the **foresee** appliance.

C.6 T3 - Attack **foresee** Platform from Customer LAN

#	Scenario Name	Description	Threat Analysis Reference	Pen Test Activities	Test Success Metrics (i.e. Home Battery System Failure Criteria)
T3	Attack foresee platform from customer LAN	Get access to the foresee platform from the customer LAN.	AU-ATK-5	On foresee platform: - Port scanning - Password guessing - Vulnerability scans	Good: gather quality data about foresee platform Better: retrieve foresee credentials

This test intended to identify vulnerabilities in the **foresee** appliance either from the application or the way the application was deployed. The goal of this test was to attempt to gain a heightened level of access to the **foresee** application.

C.6.1 Penetration Testing Tools

NMAP is an extensive network port scanning tool. NMAP was used to identify open network ports on the **foresee** platform to identify services or communications that the **foresee** application may respond to from the penetration test.

OpenVAS is a vulnerability scanner that can automate the process of testing vulnerabilities that exist. It can check for vulnerabilities with a device, or its services.

DirBuster is a directory enumeration tool designed to crawl webservers in attempt to find content that might not be linked to from accessible web pages. It uses a wordlist to guess the name of the file.

C.6.2 Results

C.6.2.1 Port Scanning

Port scanning revealed open TCP ports on port 8080 and 8086. These ports were occupied by the **foresee** application and the database that the application communicates to. UDP scans revealed port 123 which was occupied by the Network Time Protocol.

C.6.2.2 Vulnerability Scanning

The vulnerability scanner has identified a version of Allaire JRun web server hosting the **foresee** application which appeared vulnerable to a directory traversal attack. This vulnerability was confirmed to be effective in the **foresee** application. This allowed the penetration testing laptop to attempt to expose the content of files so long as the file paths could be guessed.

User accounts, API keys for **foresee** devices, and password hashes from the grub configuration were all readable by the penetration testing laptop by leveraging the directory traversal attack.

Procedure: NMAP was used to identify available services listening on the **foresee** appliance.

Relevant results of the NMAP scan which identified the running services on the **foresee** appliance are shown below.

```
» nmap -sT -T2 -Pn -v -reason -p- -n -oA pre-nmap-tcp 192.168.1.101
...
Nmap scan report for 192.168.1.101
Host is up, received user-set (0.059s latency).
Not shown: 65518 closed ports
Reason: 65518 conn-refused
PORT STATE SERVICE REASON
8080/tcp open http-proxy syn-ack
8086/tcp open d-s-n syn-ack
...
```

```
» nmap -sU -T2 -Pn -v -reason -p- -n -oA pre-nmap-udp 192.168.1.101
...
Nmap scan report for 192.168.1.101
Host is up, received arp-response (0.026s latency).
Not shown: 34640 open|filtered ports, 30894 closed ports
Reason: 34640 no-responses and 30894 port-unreaches
PORT STATE SERVICE REASON
123/udp open ntp udp-response ttl 64
MAC Address: 78:92:9C:DD:4F:23 (Intel Corporate)
...
```


2.1.1 High 8080/tcp

High (CVSS: 7.8) NVT: Generic HTTP Directory Traversal
Summary Generic check for HTTP directory traversal vulnerabilities.
Vulnerability Detection Result The following traversal URL(s) where found: Vulnerable url: <code>http://192.168.1.101:8080/../../../../../../../../etc/passwd</code> ...continues on next page ...

Figure C.8. Screenshot from the OpenVAS report alerts the penetration tester that a directory traversal vulnerability exists on the web server hosting the foresee appliance.

The vulnerability was further exploited to retrieve important information. The code snippet shows the penetration tester's laptop retrieving API keys from Ecobee by passing a malformed request to the remote server.

```
» ncat 192.168.1.101 8080

GET ../../../../../../../../../../home/foresee/foresee/Ecobee/ecobee.config HTTP/1.1

HTTP/1.1 200 OK
Content-type: text/plain
Date: Tue, 08 May 2018 21:29:40 GMT
Transfer-Encoding: chunked

140
{
  "api_key": "uRktn3b1ZsbTMV2q116Ao9VKEVdCYQAa",
  "agentid": "Ecobee",
  "poll_time": 10,
  "on_request_only": false,
  "token":{"access_token": "HCI1GqayYwsy417Tb10BskLoEVigKmtS", "token_type":
"Bearer", "expires_in": 3599, "refresh_token": "DTD03Ymr7NOK5QqPymqanNpY7xnrb1Ak",
"scope": "smartWrite"}
}

0
```

C.6.3 T3 Conclusions

The **foresee** application was sitting on a web service vulnerable to directory traversal attacks. The vulnerability allowed the penetration tester to view many files on the system that exposed sensitive configuration files of the application, and the Operating System.

Under real-world deployment **foresee** will most likely be deployed with the capability of retrieving security updates, which will reduce the likelihood of one of its services being vulnerable to directory traversal. However, this vulnerability was able to be far more impactful due to permissions granted to the web service. Critical files on the **foresee** application should have more strict permissions set to prevent vulnerabilities like this in the future.

RISK = HIGH

Because the API keys used to integrate devices with **foresee** were able to be recovered in a usable format (unencrypted plain-text), the metric for high risk has been met. To avoid such attacks, a Input validation of the requests to

the web server should be implemented and the access rights of the files by the server reduced to only the necessary files. This was not yet implemented due to the prototype nature of the **foresee** platform; the web server is used for convenience in this project and an alternate method would certainly be used in production software.

C.7 T4 - foresee Platform Data Disruption

#	Scenario Name	Description	Threat Analysis Reference	Pen Test Activities	Test Success Metrics (i.e. Home Battery System Failure Criteria)
T4	foresee platform Data Disruption	Steal, disrupt, or damage foresee platform data. Assumes foresee platform access (CS2).	C-ATK-2 AU-ATK-3	On foresee platform: - Escalate privileges	Good: log onto foresee platform, access data Better: log onto foresee platform, alter data

This test intended to see if a penetration tester could access or alter data on the **foresee** platform with assumed access to the **foresee** platform. This identifies security risk to the **foresee** platform if the customer LAN is compromised, or a malicious user was given access to the customer LAN.

C.7.1 Penetration Testing Tools

OWASP-ZAP is a web proxy used to intercept web requests and can modify web requests. It was used to modify web requests to test how **foresee** handled unexpected requests in hopes that a malicious user can get more access to data or more control of the platform.

OWASP-Mantra is a web-browser-based suite of penetration testing tools specific to testing web applications.

C.7.2 Results

In the way **foresee** was deployed in the test environment, no authentication took place that would prevent a malicious user from within the customer LAN to the **foresee** application. From the customer LAN it was possible to access the **foresee** interface along with all its features the same as a customer would be able to.

The **foresee** Web Application did not pass any web requests that could be tampered. Access to the application did not reveal any place for the penetration tester to increase level of access to data or control of the platform.

C.7.3 T4 Conclusions

foresee did not require any authentication from within the customer LAN in the way it was currently deployed. This allowed anyone with access to the customer LAN to have the same permissions of the customer. Real-world deployments should mitigate this risk with some form of authenticating the customer.

Access to the **foresee** application did not allow the penetration tester to preform any task beyond what a customer is permitted to. All attempts to intercept and tamper data shown to a legitimate customer failed.

RISK = MODERATE

With assumed access to the **foresee** platform data could be accessed. Therefore, the metric for moderate risk has been met. Since, no attempts to alter data beyond the basic functionality **foresee** provides to a typical user was achieved, the metric for high risk has not been met.

C.8 T5 - Attack Communication Channel Outside the UTM

Because the test environment was safely sandboxed from the public Internet, this test intended to identify the possibility of any attacks from outside the customer LAN.

#	Scenario Name	Description	Threat Analysis Reference	Pen Test Activities	Test Success Metrics (i.e. Home Battery System Failure Criteria)
T5	Attack communication channel outside the UTM	Attempt to deny communications from the public Internet to the customer LAN	AU-ATK-3	Outside UTM: - Physical disconnect of communication - DoS on network	Good: Disrupt communications via DoS

C.8.1 Penetration Testing Tools

Aircrack-NG is a toolset allowing various attacks through Wi-Fi networks. Aircrack-NG was used to attempt a DoS attack to disrupt communications of the **foresee** platform.

C.8.2 Results

The **foresee** platform was protected by a NAT offered by a basic home router that prevented remote users from establishing a connection with the **foresee** application unless the connection was initiated by the **foresee** platform.

Aircrack-NG was able to successfully DoS the **foresee** appliance from its Wi-Fi network making the hardware of the **foresee** platform unable to sustain a connection to its Wi-Fi network. This disrupted all communication going to or from the **foresee** platform. Although this attack was successful without any type of access to the customer LAN, it is important to note that this attack is only possible if launched from within Wi-Fi signal range and can only disrupt the communication for as long as the attack runs.

Procedure: The penetration tester identified the MAC address of the NUC hosting the **foresee** platform with passive packet sniffing. The penetration tester used Aircrack-NG to launch a Denial of Service attack.

```
» aireplay-ng -0 0 -a 00:14:6C:80:1a:00 -c EE:FA:BC:0C:C2:CD wlan0mon
```

C.8.3 T5 Conclusions

An attacker within Wi-Fi signal range of the **foresee** platform can disrupt communications of the **foresee** appliance for as long as the attacker wishes to run the attack for. This will disrupt all wireless communications to the **foresee** appliance which will prevent any new data from entering or leaving the **foresee** appliance. This attack can be performed without any level of access to the customer LAN other than being within its wireless signal range.

RISK = LOW

Although communications were disrupted via a Denial of Service from outside the customer LAN, the attack requires close proximity, which makes it unlikely to come from the public Internet. Because the Denial of Service attack is unlikely to come from the public Internet the metric for moderate risk has not been met.

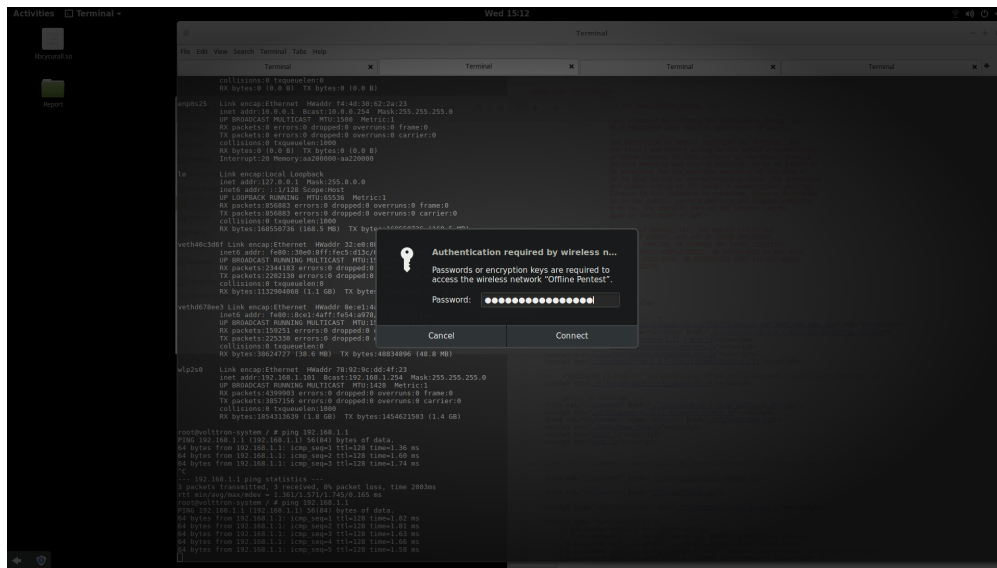


Figure C.9. Screenshot taken from the foresee appliance itself. The screenshot shows that wireless networking has been disabled from the attack, causing it to continuously attempt to reconnect.

C.9 T6 - Attack on VOLTRON

#	Scenario Name	Description	Threat Analysis Reference	Pen Test Activities	Test Success Metrics (i.e. Home Battery System Failure Criteria)
T6	Attack on VOLTRON	Probe for gaps in VOLTRON's security implementation	-	From inside the foresee platform: <ul style="list-style-type: none"> - Check for outdated versions of OpenSSL. - Attempt to connect via deprecated connection protocol - Check for key misconfiguration on 0MQ and CurveZMQ 	Good: Find outdated or misconfigured system software Better: Exploit outdated or misconfigured system software

This test sought to identify and exploit vulnerabilities of the VOLTRON platform.

C.9.1 Penetration Testing Tools

OpenVAS is a vulnerability scanner that can automate the process of testing vulnerabilities that exist. It can identify misconfigured OpenSSL keys and unsecured encryption protocols.

C.9.2 T6 Conclusions

No misconfigured keys or out of date software was found beyond the vulnerabilities discovered from other tests. The penetration testing tools did not find any additional vulnerabilities that will allow an attack from outside the UTM.

RISK = UNDISCOVERED

Versions of SSL and key configurations appeared to be secure by current industry standards.

C.10 Cybersecurity Recommendations

Of the vulnerabilities discovered the chart below organizes each discovered vulnerability by level of criticality and the test case it was discovered in. Each vulnerability includes a mitigation strategy to assist in the **foresee** application's resilience against the discovered vulnerabilities when the application is deployed.

Table C.8. Cybersecurity Recommendations

Test #	Vulnerability	Criticality
T3	Directory Traversal	HIGH
Services accessible to users should be updated and patched promptly as security updates are released. Additionally, file permissions and permissions of applications should be limited so that critical files require higher levels of permissions to access.*		

Test #	Vulnerability	Criticality
T1	DoS From Flooding Attacks	MODERATE
Some Denial of Service attacks may be reduced by deploying foresee with a UTM. TCP SYN floods can be reduced by using SYN cookies, micro blocks, or other mitigation techniques to reduce the likelihood of resources being exhausted during a TCP SYN flood. ICMP flooding can likely be reduced with firewalls or Intrusion Prevention Systems that are configured specifically to prevent ICMP flooding.		

Test #	Vulnerability	Criticality
T4	Lack of User Authentication	MODERATE
Currently, users accessing the foresee interface do not undergo any authentication process. Any form of authentication to the foresee interface can further protect it against malicious users. Htaccess Authentication can be used to restrict access to certain directories with basic HTTP authentication.		

Test #	Vulnerability	Criticality
T5	Wi-Fi Deauthentication Attack	LOW
802.11 can encrypt management frames presuming that the AP and clients support it. Other wireless network protection features can block this type of attack but will most likely require implementing security on the customer's LAN network infrastructure.		

*To avoid such attacks, a Input validation of the requests to the web server should be implemented and the access rights of the files by the server reduced to only the necessary files. This was not yet implemented due to the prototype nature of the **foresee** platform; the web server is used for convenience in this project and an alternate method would certainly be used in production software.

References

- Akyol, Bora, Jereme Haack, Selim Ciraci, Brandon Carpenter, Maria Vlachopoulou, and C. W. Tews. 2012. "VOLTTRON: An Agent Execution Platform for the Electric Power System." *Third International Workshop on Agent Technologies for Energy Systems (ATES 2012). A workshop of the Eleventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)* (Valencia, Spain). VOLTTRON open-source code and documentation can be found here: <http://github.com/volttron/VOLTTRON/>.
- Albadi, M. H., and E. F. El-Saadany. 2008. "A Summary of Demand Response in Electricity Markets." *Electric Power Systems Research* 78 (11). doi:10.1016/j.epsr.2008.04.002.
- Aloise-Young, Patricia, Salvador Lurbe, Steven Isley, Xin Jin, Rahul Kadavil, Siddharth Suryanarayanan, and Dane Christensen. 2018. "Demand Response Load Curtailment Preferences: Strategies for Multi-Criterion Decision Making in a Smart Home." *In development*.
- Anvari-Moghaddam, A., H. Monsef, and A. Rahimi-Kian. 2015. "Optimal Smart Home Energy Management Considering Energy Saving and a Comfortable Lifestyle." *IEEE Transactions on Smart Grid* 6 (1).
- Baker, Kyri, Xin Jin, Deepthi Vaidhynathan, Wesley Jones, Dane Christensen, Bethany Sparn, Jason Woods, Harry Sorensen, and Monte Lunacek. 2016. "Frequency Regulation Services from Connected Residential Devices." *Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments* (Palo Alto, CA).
- Balta-Ozkan, Nazmiye, Rosemary Davidson, Martha Bicket, and Lorraine Whitmarsh. 2013. "Social Barriers to the Adoption of Smart Homes." *Energy Policy* 63. doi:10.1016/j.enpol.2013.08.043.
- Ben-Akiva, Moshe E., and Steven R. Lerman. 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand*. MIT Press.
- Berinsky, Adam J., Gregory A. Huber, and Gabriel S. Lenz. 2012. "Evaluating Online Labor Markets for Experimental Research: Amazon.com's Mechanical Turk." *Political Analysis* 20 (3). doi:10.1093/pan/mpr057.
- Bishop, C. 2006. "Pattern recognition and machine learning." Springer-Verlag.
- Bloem, J. 1994. *System identification applied to building performance data*. Joint Research Centre, European Commission.
- Christensen, Dane, Lieko Earle, and Bethany Sparn. 2012. "NILM Applications for the Energy-Efficient Home." *1st International Workshop on Non-Intrusive Load Monitoring* (Pittsburgh, PA).
- Christensen, Dane, Steve Isley, Kyri Baker, Xin Jin, Patricia Aloise-Young, Rahul Kadavil, and Siddharth Suryanarayanan. 2016. "Homeowner Preference Elicitation: A Multi-Method Comparison." *Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments (BuildSys '16)* (Palo Alto, CA).
- Collins, A. T., J. M. Rose, and M. C. J. Bliemer. 2007. *Ngene-The Cutting Edge in Experimental Design for Stated Choice Experiments*. Institute of Transport / Logistics Studies, The University of Sydney, Australia.
- Cutler, Dylan, Ted Kwasnik, Sivasathya Balamurugan, Sam Booth, Bethany Sparn, and Karen Hsu. 2018. "A Demonstration of Blockchain-based Energy Transactions between Laboratory Test Homes." *ACEEE 2018 Summer Study on Energy Efficiency in Buildings* (Pacific Grove, CA).
- Dall'Anese, E., K. Baker, and T. Summers. 2017. "Chance-Constrained AC Optimal Power Flow for Distribution Systems With Renewables." *IEEE Transactions on Power Systems* 32 (5): 3427–3438.
- Dumont, Jeff, Jeff Keller, and Chase Carpenter. 2015. "RSGHB: Functions for Hierarchical Bayesian Estimation: A Flexible Approach (version 1.1.2)." <https://github.com/RSGInc/RSGHB>.
- Ecotope. 2014. *Residential building stock assessment: metering study*. Tech. rep. E 14-283. Northwest Energy Efficiency Alliance.
- Edwards, Ward, and F.Hutton Barron. 1994. "SMARTS and SMARTER: Improved Simple Methods for Multiattribute Utility Measurement." *Organizational Behavior and Human Decision Processes* (3). doi:10.1006/obhd.1994.1087.
- EIA. 2017. *Electric Power Monthly: Average Price of Electricity to Ultimate Customers by End-Use Sector*. Accessed September 23, 2017, https://www.eia.gov/electricity/monthly/epm_table_grapher.php?t=epmt_5_6_a. *EnergyPlus*. <https://energyplus.net/>.
- EPRI. *National Electric Sector Cybersecurity Organization Resource (NESCOR) Electric Sector Failure Scenarios and Impact Analysis*. <http://smartgrid.epri.com/NESCOR.aspx>.

- Garifi, Kaitlyn, Kyri Baker, Dane Christensen, and Behrouz Touri. 2018a. "Non-Simultaneous Charging and Discharging Guarantees in Energy Storage System Models for Home Energy Management Systems." *in review*. <http://arxiv.org/abs/1805.00100>.
- Garifi, Kaitlyn, Kyri Baker, Behrouz Touri, and Dane Christensen. 2018b. "Stochastic Model Predictive Control for Demand Response in a Home Energy Management System." *IEEE Power & Energy Society General Meeting* (Portland, OR).
- Grant, M, and S Boyd. 2014. *CVX: matlab software for disciplined convex programming, version 2.1*. <http://cvxr.com/cvx>.
- Hendron, Robert, Jay Burch, and Greg Barker. 2014. "Tool for Generating Realistic Residential Hot Water Event Schedules," no. NREL/CP-550-47685. The event generator spreadsheet can be found here: <http://energy.gov/eere/buildings/building-america-analysis-spreadsheets>, <http://www.nrel.gov/docs/fy10osti/47685.pdf>.
- Hoyos, David. 2010. "The State of the Art of Environmental Valuation with Discrete Choice Experiments." *Ecological Economics* 69 (8). doi:10.1016/j.ecolecon.2010.04.011.
- Huang, Shiu-li. 2011. "Designing Utility-Based Recommender Systems for E-Commerce: Evaluation of Preference-Elicitation Methods." *Electronic Commerce Research and Applications* 10 (4). doi:10.1016/j.elerap.2010.11.003..
- Ipeirotis, Panagiotis G. 2010. *Demographics of Mechanical Turk*. Tech. rep. SSRN Scholarly Paper ID 1585030. Rochester, NY: Social Science Research Network. <http://papers.ssrn.com/abstract=1585030>.
- ISO. 2008. *ISO-18045: Common Methodology for Information Technology Security Evaluation*. Tech. rep. ISO-18045. International Standards Organization. <https://www.iso.org/standard/46412.html>.
- Jin, Xin. 2017. "An active learning framework for non-intrusive load monitoring." *3rd International Workshop on Non-Intrusive Load Monitoring*.
- Jin, Xin, Jeff Maguire, and Dane Christensen. 2014. "Model predictive control of heat pump water heaters for energy efficiency." *2014 ACEEE Summer Study on Energy Efficiency in Buildings*.
- Jin, Xin, Jeff Maguire, and Dane Christensen. 2018. "Economic Sizing of Batteries for the Smart Home." *5th International High Performance Building Conference* (Lafayette, IN).
- Jin, Xin, Kyri Baker, Steven Isley, and Dane Christensen. 2016. "User-Preference-Driven Multi-Objective Model Predictive Control of Residential Building Loads and Battery Storage for Demand Response." *Proceedings of 2017 American Control Conference* (Seattle, WA). <https://www.nrel.gov/docs/fy17osti/67809.pdf>.
- Jin, Xin, Kyri Baker, Dane Christensen, and Steven Isley. 2017. "Foresee: A user-centric home energy management system for energy efficiency and demand response." *Applied Energy* 205.
- Kadavil, Rahul, Siddharth Suryanarayanan, Patricia Aloise-Young, Steve Isley, and Dane Christensen. 2018. "An Application of the Analytic Hierarchy Process for Prioritizing User Preferences in the Design of a Home Energy Management System." *Sustainable Energy, Grids and Networks* 16. doi:10.1016/j.segan.2018.07.009.
- Langner, M. Rois, and Dane Christensen. 2018. "Navigating Cyber Security Implications of Smart Outlets." *ACEEE 2018 Summer Study on Energy Efficiency in Buildings* (Pacific Grove, CA).
- Laws, Nicholas, Xin Jin, Killian McKenna, Julieta Giraldez, and Dane Christensen. 2018. "Development of Machine Learning for Predicting Distribution Feeder Voltage Violations." *in preparation*.
- Lenk, Peter J., Wayne S. DeSarbo, Paul E. Green, and Martin R. Young. 1996. "Hierarchical Bayes Conjoint Analysis: Recovery of Partworth Heterogeneity from Reduced Experimental Designs." *Marketing Science* 15 (2). doi:10.1287/mksc.15.2.173.
- Louviere, Jordan J., David A. Hensher, and Joffre D. Swait. 2000. *Stated Choice Methods: Analysis and Applications*. New York: Cambridge University Press.
- Louviere, Jordan J., Deborah Street, Leonie Burgess, Nada Wasi, Towhidul Islam, and Anthony A. J. Marley. 2008. "Modeling the Choices of Individual Decision-Makers by Combining Efficient Choice Experiment Designs with Extra Preference Information." *Journal of Choice Modelling* 1 (1). doi:10.1016/S1755-5345(13)70025-3.
- Martin, Maurice, Erdenebat Gantumur, Brandon Mendrick, and Dane Christensen. 2018. "Risk Assessment at the Edge: Applying NERC CIP to Aggregated Grid-Edge Resources." *in development*.
- Mason, Winter, and Siddharth Suri. 2012. "Conducting Behavioral Research on Amazon's Mechanical Turk." *Behavior Research Methods* 44 (1). doi:10.3758/s13428-011-0124-6.
- Nelson, Austin, et al. 2016. "Hawaiian Electric Advanced Inverter Grid Support Function Laboratory Validation and Analysis," no. NREL/TP-5D00-67485. <https://www.nrel.gov/docs/fy17osti/67485.pdf>.
- NERC. 2017. *Critical Infrastructure Protection Standards*. Version 5. <https://www.nerc.com/pa/Stand/Pages/CIPStandards.aspx>.
- NREL. 2018. *National Solar Radiation Database*. Accessed January 11, 2018, <https://nsrdb.nrel.gov/>.

- Ohnaka, Tadakatsu, Yutaka Tochiara, and Yumiko Watanabe. 1994. "The Effects of Variation in Body Temperature on the Preferred Water Temperature and Flow Rate during Showering." *Ergonomics* 37 (3). doi:10.1080/00140139408963669.
- Oldewurtel, Frauke, Alessandra Parisio, Colin N. Jones, Dimitrios Gyalistras, Markus Gwerder, Vanessa Stauch, Beat Lehmann, and Manfred Morari. 2012. "Use of model predictive control and weather forecasts for energy efficient building climate control." *Energy and Buildings* 45:15–27.
- Oregon, University of. 2018. *Solar Radiation Monitoring Laboratory Solar Data*. Accessed January 10, 2018, <http://solardat.uoregon.edu/SolarData.html>.
- P. Norton, K. J. Kelly, K. Kiatreungwattana. 2013. "Evaluation of Model Results and Measured Performance of Net-Zero Energy Homes in Hawaii: Preprint." *2013 ASHRAE Annual Conference* (Denver, CO), no. NREL/CP-7A30-58000. <https://www.nrel.gov/docs/fy13osti/58000.pdf>.
- Paolacci, Gabriele, Jesse Chandler, and Panagiotis G. Ipeirotis. 2010. "Running Experiments on Amazon Mechanical Turk." *Judgment and Decision Making* 5 (5).
- Parker, Danny, Philip Fahey, and James Lutz. 2015. "Estimating Daily Domestic Hot Water Use in North American Homes." *Proceedings of the 2015 ASHRAE Annual Conference*, no. FSEC-PF-464-15. <http://www.fsec.ucf.edu/en/publications/pdf/FSEC-PF-464-15.pdf>.
- Pöyhönen, Mari, and Raimo P Hämäläinen. 2001. "On the Convergence of Multiattribute Weighting Methods." *European Journal of Operational Research* 129 (3). doi:10.1016/S0377-2217(99)00467-1.
- Pratt, Annabelle, Mark Ruth, Dheepak Krishnamurthy, Bethany Sparr, Monte Lunacek, and Wesley Jones. 2017. "Hardware-in-the-Loop Simulation of a Distribution System with Air Conditioners under Model Predictive Control," no. NREL/CP-5D00-67392. <https://www.nrel.gov/docs/fy17osti/67392.pdf>.
- Rastegar, Mohammad, Mahmud Fotuhi-Firuzabad, and Farrokh Aminifar. 2012. "Load Commitment in a Smart Home." *Applied Energy* 96. doi:10.1016/j.apenergy.2012.01.056.
- Raszmann, Emma, Kyri Baker, Dane Christensen, and Ying Shi. 2017. "Modeling Stationary Lithium-ion Batteries for Optimization and Predictive Control." *IEEE Power and Energy Conference at Illinois (PECI)* (Urbana-Champaign, IL). (Best Paper Award), <https://www.nrel.gov/docs/fy17osti/67809.pdf>.
- Reed Johnson, F., Emily Lancsar, Deborah Marshall, Vikram Kilambi, Axel Mühlbacher, Dean A. Regier, Brian W. Bresnahan, Barbara Kanninen, and John F. P. Bridges. 2013. "Constructing Experimental Designs for Discrete-Choice Experiments: Report of the ISPOR Conjoint Analysis Experimental Design Good Research Practices Task Force." *Value in Health* 16 (1). doi:10.1016/j.jval.2012.08.2223.
- Roberts, Ron, and Paul Goodwin. 2002. "Weight Approximations in Multi-Attribute Decision Models." *Journal of Multi-Criteria Decision Analysis* 11 (6). doi:10.1002/mcda.320.
- SCE. 2017. *Southern California Edison Time-Of-Use Residential Rate Plans*. , last accessed September 23, 2017, <https://www.sce.com/wps/portal/home/residential/rates/Time-Of-Use-Residential-Rate-Plans>.
- Sengupta, M., and A. Andreas. 2010. "Oahu Solar Measurement Grid (1-Year Archive): 1-Second Solar Irradiance." (Oahu, HI), no. NREL/DA-5500-56506. <http://dx.doi.org/10.5439/1052451>.
- SGIP. 2016. *Transactive Energy Application Landscape Scenarios*. Tech. rep. SGIP. http://www.sgip.org/wp-content/uploads/SGIP_White_Paper_TE_Application_Landscape_Scenarios_12-15-2016_FINAL.pdf.
- Srivastava, Joydeep, and Terry Connolly. 1995. "Do Ranks Suffice? A Comparison of Alternative Weighting Approaches in Value Elicitation." *Organizational Behavior and Human Decision Processes* 63 (1).
- Suryanarayanan, Siddharth, Marvin Antony Devadass, and Timothy M Hansen. 2015. "A Load Scheduling Algorithm for the Smart Home Using Customer Preferences and Real Time Residential Prices." *IFAC-PapersOnLine* 48 (30): 126–131.
- Train, Kenneth E. 2009. *Discrete Choice Methods with Simulation*. Cambridge University Press.
- Tsui, K. M., and S. C. Chan. 2012. "Demand Response Optimization for Smart Home Scheduling Under Real-Time Pricing." *IEEE Transactions on Smart Grid* 3 (4). doi:10.1109/TSG.2012.2218835.
- Wilson, Eric, Cheryn Engebrecht-Metzger, Scott Horowitz, and Robert Hendron. 2014. "2014 Building America House Simulation Protocols," no. NREL DOE/GO-102013-4336. <https://www.nrel.gov/docs/fy14osti/60988.pdf>.
- Wilson, Eric, Craig Christensen, Scott Horowitz, and Henry Horsey. 2016. "A High-Granularity Approach to Modeling Energy Consumption and Savings Potential in the U.S. Residential Building Stock." *Proceedings of the ASHRAE and IBPSA-USA SimBuild 2016 Building Performance Modeling Conference* (Salt Lake City, Utah). <http://ibpsa-usa.org/index.php/ibpsa/issue/view/35>.
- Xcel. 2018. *Xcel Energy Colorado Residential Electric and natural gas rate schedule summaries*. Last accessed July 10, 2018, <https://www.xcelenergy.com/staticfiles/xcel/Regulatory/COResRates.pdf>.

Zipperer, A., P. A. Aloise-Young, S. Suryanarayanan, R. Roche, L. Earle, D. Christensen, P. Bauleo, and D. Zimmerle. 2013. "Electric Energy Management in the Smart Home: Perspectives on Enabling Technologies and Consumer Behavior." *Proceedings of the IEEE* 101 (11): 2397–2408. doi:10.1109/JPROC.2013.2270172.