Supervisory Control of Loads and Energy Storage in Next-Generation Zero Energy Buildings

Feitau Kung, Stephen Frank, Jennifer Scheib, Willy Bernal Heredia, and Shanti Pless

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.

Technical Report
NREL/TP-5500-67007
September 2016

Contract No. DE-AC36-08GO28308

Supervisory Control of Loads and Energy Storage in Next-Generation Zero Energy Buildings

Feitau Kung, Stephen Frank, Jennifer Scheib, Willy Bernal Heredia, and Shanti Pless

Prepared under Task No. 0655.1501

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.

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

Technical Report
NREL/TP-5500-67007
September 2016

Contract No. DE-AC36-08GO28308

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

Available electronically at SciTech Connect http:/www.osti.gov/scitech

Available for a processing fee to U.S. Department of Energy and its contractors, in paper, from:

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
OSTI http://www.osti.gov
Phone: 865.576.8401
Fax: 865.576.5728
Email: reports@osti.gov

Available for sale to the public, in paper, from:

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Road
Alexandria, VA 22312
NTIS http://www.ntis.gov
Phone: 800.553.6847 or 703.605.6000
Fax: 703.605.6900
Email: orders@ntis.gov

Cover Photos by Dennis Schroeder: (left to right) NREL 26173, NREL 18302, NREL 19758, NREL 29642, NREL 19795.

NREL prints on paper that contains recycled content.

Acknowledgments

The authors would like to thank the many colleagues who assisted with conceptual design and design review for the thermal systems, lighting systems, controls, and integration strategies proposed for this project. In particular, the authors thank Tony Markel, Eric Kozubal, Andrew Meintz, Mark Martinez, Anna Hoenmans, Sal Sferrazza, Jacob Reynoldson, Michael Sheppy, and Adam Phipps of the National Renewable Energy Laboratory, as well as Denise Dihle and Travis Menard of 360 Engineering.

The authors would also like to acknowledge several colleagues who provided energy modeling advice and assistance during the course of this project. In particular, the authors thank Andrew Parker for creating the initial geometry for a whole-building energy model of the Vehicle Testing and Integration Facility; Kyle Benne for helping with modeling tool selection and workflow recommendations; and Edwin Lee, David Goldwasser, Robert Guglielmetti, and Daniel Studer for assistance with modeling strategies and model review. Finally, the authors thank Ren Anderson and Ron Judkoff for reviewing this report.

Nomenclature

AWHP	air-to-water heat pump
BAS	building automation system
DFT	discrete Fourier transform
DOE	U.S. Department of Energy
EMS	energy management system
FCU	fan coil unit
HVAC	heating, ventilating, and air conditioning
kWh	kilowatt-hour
mHz	millihertz
MWh	megawatt-hour
NREL	National Renewable Energy Laboratory
RSF	Research Support Facility
PV	photovoltaic
SOC	state of charge
TES	thermal energy storage
VTIF	Vehicle Testing and Integration Facility
ZEB	zero energy building

Executive Summary

A zero energy building (ZEB)—also known as a net zero energy or zero net energy building—is a building that exports as much renewable energy as the total energy it imports from other sources on an annual basis (DOE 2015). Large-scale and commercially viable ZEBs are now in the marketplace, and they are expected to become a larger share of the commercial building footprint as government and private sector policies continue to promote the development of buildings that produce more on-site energy than they use.

However, the load profiles of ZEBs are currently perceived by electric utilities to be unfavorable and unpredictable. As shown in Figure ES-1, ZEB load profiles can have abrupt changes in magnitude, at times switching rapidly between exporting and importing electricity. This is a challenge for utilities, which are responsible for constantly balancing electricity supply and demand across the grid. Addressing these concerns will require new strategies and tools.

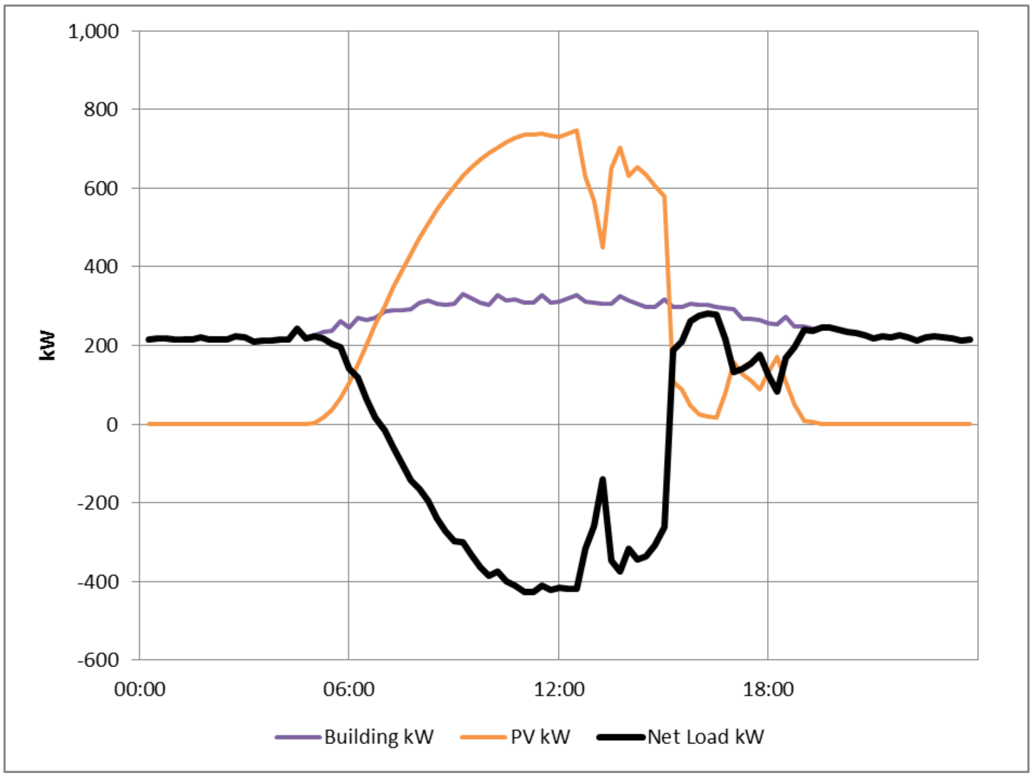

Figure ES-1. Sample daily load profile for a "first-generation" zero energy building with rapid changes in net load

Integrating advanced controls and on-site energy storage will allow future ZEBs to actively manage their load profiles and provide new value to utility operators and building owners. These efforts will enable greater deployment of renewable generation while facilitating grid stability. To advance capabilities in this area, the project team:

- Developed a prototype system design and issued a request for proposals for adding a ZEB load profile controller and flexible building loads to an existing building. Working through this process enabled the team to document real-world integration challenges and identify considerations that may be extended to future projects.

- Created and demonstrated novel capabilities for using whole-building simulation tools to design and test control strategies for systems that integrate on-site generation, on-site storage, and flexible building loads.
- Developed and explored potential methods for characterizing and synthesizing minute-scale variability in building loads. These methods could be refined through future research to enable modeling of control strategies for buildings that lack 1-minute-interval metered data.

Key findings from these real-world and simulation-based efforts include:

- ZEB supervisory controller performance is sensitive to multiple design parameters. These can range from control choices (e.g., net load targets and battery constraints) to system characteristics (e.g., battery capacity).
- Different control objectives can work in opposition to each other, such as when peak demand reduction is improved at the expense of an increase in imported energy and battery cycling. Developers of supervisory controls for next-generation ZEBs will therefore require tools that can navigate complex parameter spaces.
- A simulation-based approach can help modelers draft, test, compare, and improve ZEB supervisory control strategies based on a variety of metrics.
- Carefully designed systems combining batteries, thermal energy storage, flexible loads, and advanced supervisory controls can sometimes outperform systems that rely solely on batteries for ZEB load profile management.

Table of Contents

List of Figures

List of Tables

1 Introduction

A zero energy building (ZEB)—also known as a net zero energy or zero net energy building—is a building that exports as much renewable energy as the total energy it imports from other sources on an annual basis (DOE 2015).[1] Large-scale and commercially viable ZEBs are now in the marketplace, and they are expected to become a larger share of the commercial building market as government and private sector policies continue to promote the development of buildings that produce more on-site energy than they use.[2]

However, the load profiles of ZEBs are currently perceived by electric utilities to be unfavorable and unpredictable, as shown in Figure 1. At the same time, building owners are increasingly interested in enhancing energy surety for ZEBs. Addressing these concerns will require new strategies and tools for building stakeholders, utilities, and technology developers.

Integrating advanced controls and on-site storage will allow future ZEBs to actively manage their load profiles and provide new value to utility operators and building owners. This report discusses control challenges, solutions, and newly developed capabilities for designing and testing supervisory control strategies for buildings with on-site generation and storage.

1.1 Background: First-Generation Zero Energy Buildings

A ZEB is required to balance energy use and renewable generation on an annual basis, but it is otherwise free to export and import energy. This means a ZEB is not required to have any on-site energy storage. A ZEB without storage relies on the grid to provide power when on-site generation is low, and it relies on other grid customers to use ZEB exports when on-site generation is high. Such buildings can be thought of as "first-generation" ZEBs.

One obstacle to greater ZEB adoption is that the load profiles of first-generation ZEBs can be unfavorable and unpredictable from the perspective of grid operation. For example, Figure 1 depicts a daily load profile from the U.S. Department of Energy's (DOE's) National Renewable Energy Laboratory (NREL) Research Support Facility I (RSF I), a zero energy office building in Golden, Colorado (DOE 2012). The black line in the figure is the building's electrical "net load," which is defined as follows:

$$\textit{Net Load} = \textit{Imported Electricity} - \textit{Exported Electricity}.$$

"Imported electricity" is electricity delivered to the building from the grid, and "exported electricity" is electricity sent from the building to the grid. For a building without energy storage, net load is positive whenever the building needs more electricity that it can produce, and net load is negative whenever on-site generation exceeds the building's requirements.

[1] For a more detailed definition and guidance for ZEB calculations, see *A Common Definition for Zero Energy Buildings,* produced by the U.S. Department of Energy and the National Institute of Building Sciences (DOE 2015).

[2] Example policies include Title 24 in California (State of California 2016) and Executive Order 13693 (Office of Federal Sustainability 2015).

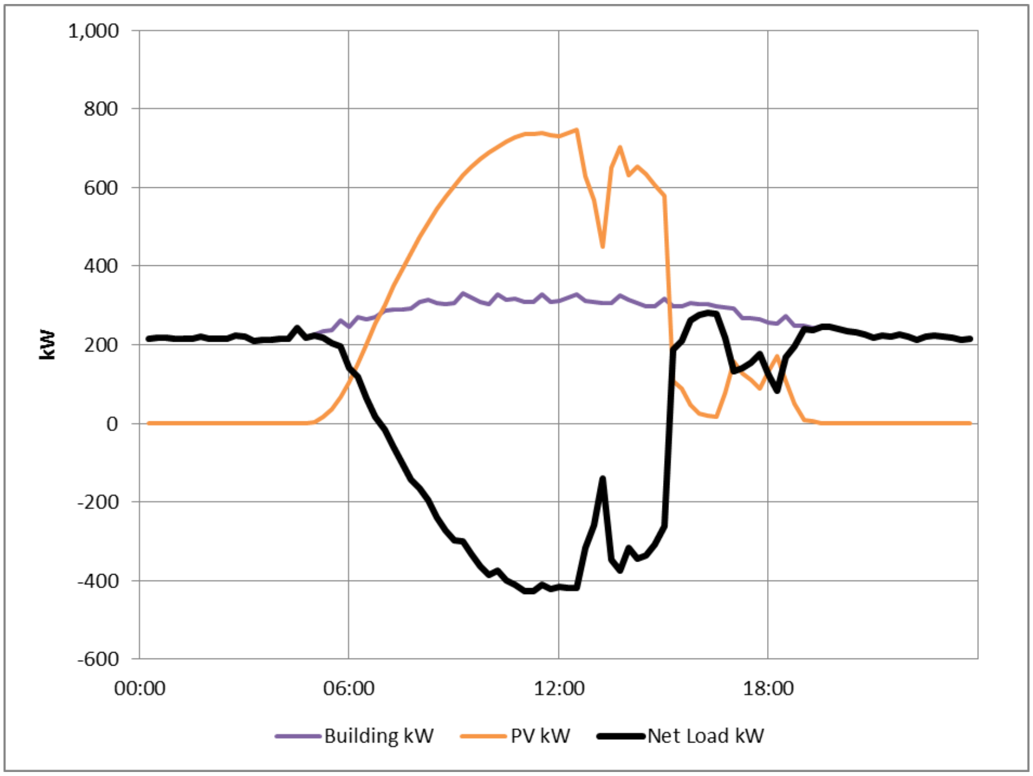

Figure 1. Sample daily load profile for a "first-generation" zero energy building with rapid changes in net load

At this site, the renewable generation is from a photovoltaic (PV) system. On this particular day, cloud cover increased in the late afternoon, causing dips in PV output and corresponding spikes in net load. This kind of variability is a concern for utilities, which are responsible for balancing electricity supply and demand.

Such load profile characteristics may be less of an issue in the near term, because ZEBs currently represent a small fraction of the total building stock. But as the adoption of ZEBs and distributed renewable generation grows, so too will the demand for ZEBs to have more consistent or controllable load profiles.

1.2 Next-Generation Zero Energy Buildings

"Next-generation" ZEBs could address these concerns by actively controlling electricity demand at various timescales. This could be accomplished with two critical additional elements: on-site energy storage and advanced supervisory controls.

- Energy storage systems make it possible to change the timing of electrical demand without impacting the building's primary mission of serving building occupants. The energy storage could take a variety of forms, such as electrochemical batteries, thermal energy storage (TES) integrated with a building's heating, ventilating, and air-conditioning (HVAC) system, or the thermal capacitance of the building itself.

- New supervisory control strategies must also be developed to manage combinations of flexible building loads, on-site renewable generation systems, and on-site energy storage.

In many commercial buildings, particularly medium to large ones, a wide range of device-level controllers for equipment like chillers and pumps are supervised by a building automation system (BAS). However, BAS products are not typically designed to supervise load shaping with energy storage. Next-generation ZEB supervisory controllers could either be incorporated into a BAS product or packaged as a separate product that can communicate with a BAS and other device-level controllers.

By combining flexible loads, on-site generation, on-site storage, and advanced supervisory controls, next-generation ZEBs have the potential to provide building owners and utilities with new value streams, such as:

- Maximizing on-site renewable generation while avoiding exports to the grid
- Minimizing unexpected variability in building net load profiles
- Minimizing or shifting net demand based on incentives in utility rate structures
- Managing loads, generation, and storage strategically during demand response events
- Managing loads, generation, and storage strategically and providing energy surety during islanding events
- Responding to utility control signals to provide ancillary grid services.

To advance capabilities for designing next-generation ZEB systems and controls, the project team:

- Developed a prototype system design and issued a request for proposals for adding a ZEB load profile controller and flexible building loads to an existing building. This integration work is described in Section 2. Working through this process enabled the team to document real-world integration challenges and identify considerations that may be extended to future projects.
- Created and demonstrated novel capabilities for using whole-building energy simulation tools to design and test control strategies for systems that integrate on-site generation, on-site storage, and flexible building loads. These capabilities are described in Section 3.
- Developed and explored potential methods for characterizing and synthesizing minute-scale variability in building loads. These methods are discussed in Section 4 and could be refined through future research to enable modeling of control strategies for buildings that lack 1-minute-interval metered data.

2 Integration Considerations

The team used a building on the NREL campus to explore, address, and document real-world integration requirements for enabling ZEB load profile control in an operational building. This initial integration effort yielded two important results:

1. The system specifications selected in this effort went on to form the basis of simulations described later in this report.

2. The initial integration effort revealed multiple considerations that can be generalized to inform future installations of ZEB systems and load profile controllers.

The following sections describe the case study location and key findings from the integration effort.

2.1 Building Description

The subject of the study was NREL's Vehicle Testing and Integration Facility (VTIF) in Golden, Colorado. The building has about 3,000 ft^2 of floor space divided between two floors and is used for research activities. The first floor primarily consists of a large high-bay research area with a few smaller rooms. The second floor includes a mezzanine with a few office cubicles and rooms. Building systems and spaces are detailed in Section 3.

Figure 2. The Vehicle Testing and Integration Facility. ***Photo by Dennis Schroeder, NREL 31843***

This building was selected for this study in part because its existing infrastructure included a PV system and major energy end-use categories (lighting, heating, cooling, plug and process loads, etc.) that are found in the commercial building sector. Additionally, its small size and research-oriented space types made it easier to propose system changes without disrupting occupants.

The VTIF is not a ZEB, but it does have a substantial amount of on-site PV generation relative to its electricity use. Extrapolating from data collected during parts of 2014 and 2015, the team estimated that the VTIF's PV generation would equal about 85% of its electricity use in a typical year. This generation-to-demand ratio is high enough to create net load profile characteristics that are also seen in ZEBs. This is illustrated in Figure 3, which depicts the electrical net load profiles of RSF I (a ZEB) and the VTIF. RSF I is a zero energy office building located on the same campus as the VTIF.[3]

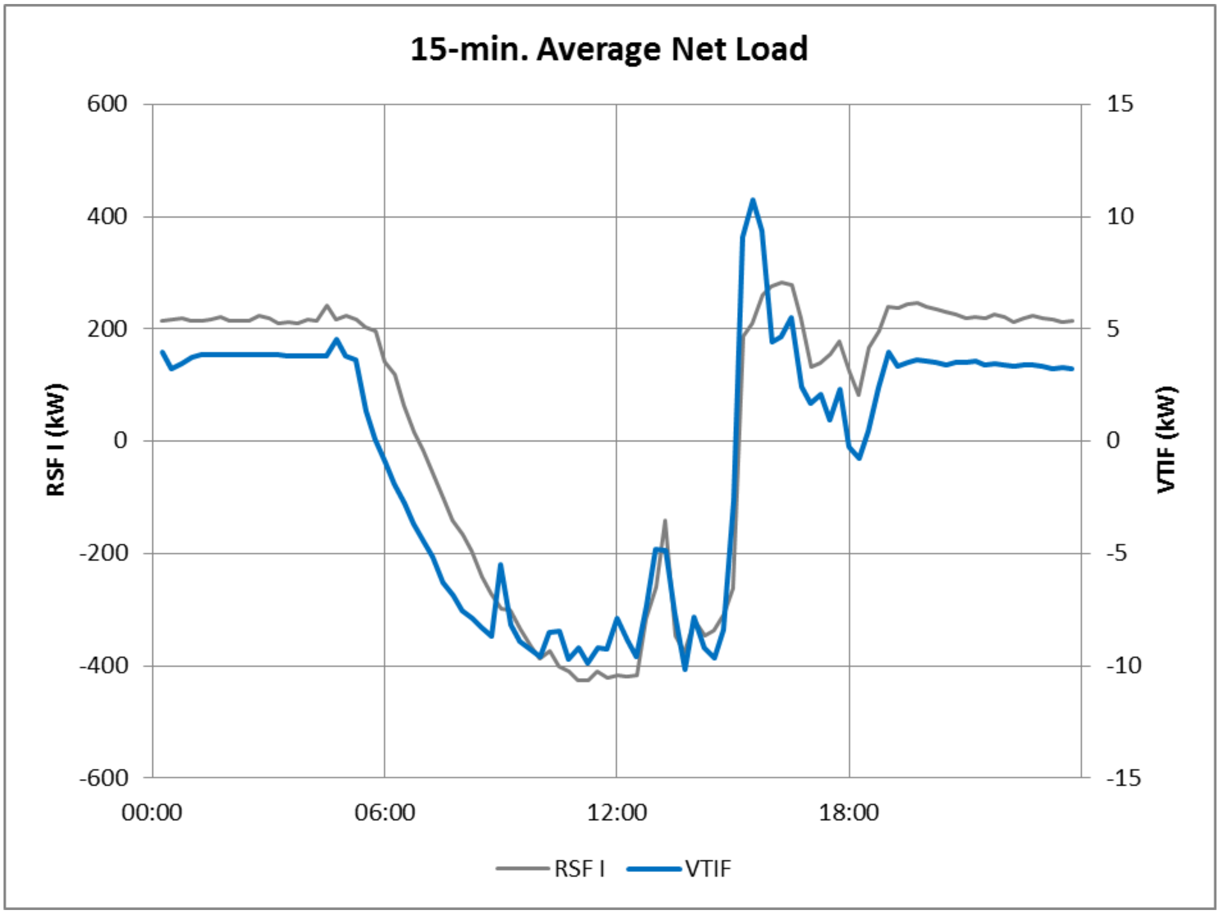

Figure 3. Sample daily net load profile for Research Support Facility I and the Vehicle Testing and Integration Facility on the same day

Like its neighbor RSF I, the VTIF has a net load profile with rapid changes that may concern utilities. Indeed, utilities' concerns about ZEBs also extend more broadly to customer-sited renewable generation in general. Similarly, the implementation strategies, new design capabilities, and research findings described in this report are applicable not just to ZEBs, but also to other buildings with a high ratio of energy generation to energy demand.

2.2 Considerations for Designing a Control Strategy

2.2.1 Level of Complexity and Data Requirements

A ZEB load profile controller could be designed to be simpler or more complex, depending on project-specific needs and priorities. Related considerations include:

[3] RSF I is the original 222,000-ft^2 ZEB completed in 2010. A subsequent ZEB addition called RSF II expanded the facility's total floor space to 360,000 ft^2. For more information on the RSF, see *The Design-Build Process for the Research Support Facility* (DOE 2012).

- Will the control approach be purely heuristic, or will it include real-time optimization?
- Will there be a learning component to the control algorithm?
- Will there be a predictive component to the control system?
- Will the supervisory controller be part of or separate from the BAS?
- Will the supervisory controller send many specific commands to many low-level actuators? Or will it send fewer commands and allow device-level controllers to be more autonomous?
- Will the control approach involve any communication with utility systems, and if so, will the communication be one-way or two-way?

For example, if someone wants to implement ZEB load profile control within an existing BAS, a heuristic, non-predictive approach may be a good fit. The algorithm could be designed and tested in advance through simulation, laboratory testing, and field testing, but the intent would be for the algorithm to be inexpensive to maintain once implemented.

On the other hand, if someone wants to improve results by accounting for changing weather forecasts, then a model-predictive control approach may be beneficial. Such an approach would likely require more inputs, computing resources, and maintenance costs, but it would also be able to reevaluate and optimize solutions on an ongoing basis.

Future incentive programs may also impact choices about controller complexity. The most sophisticated controllers might communicate directly with utility systems that transmit real-time pricing or other control signals to participating customers. Such programs may be designed to assist utilities with advanced load management or ancillary services, such as frequency regulation. Alternatively, if utilities develop rate structures that incentivize simpler load management outcomes (e.g., maintaining net load within a certain range), customers may be able to participate with simpler, more autonomous control systems.

For the VTIF project, the team decided to pursue a simpler, heuristic control approach so that the algorithm could be implemented in the VTIF's existing BAS. In this case, the team opted to sacrifice some level of control precision and adaptability in favor of lower cost and greater ease of implementation.

Any ZEB load profile controller will involve inputs and outputs, but initial decisions about the architecture of a control system can have interesting downstream effects on system integration costs. Considerations of this type include:

- How many inputs and outputs will be required?
- How will the number of inputs and outputs affect installation costs?
- What communication protocols will be required to complete the system?

- Will communication across network firewalls be required?
- How will the supervisory controller respond to short-term power outages or anomalies in metered data?
- Is external data (from outside the building) required? If so, will the supervisory controller still function properly if communication with external networks is disrupted?

2.2.2 Defining Control Stages

The team also chose to design a supervisory controller that sends as few commands as necessary to lower-level devices, allowing device-level controllers to be more autonomous. To accomplish this, the team defined “supervisory control stages” ranging from -4 to +4. The intent was for the supervisory controller to communicate the control stage to lower-level controllers, which would then respond by implementing specific device-level actions. The supervisory control stages were defined as follows:

- A positive stage number (+1 to +4) indicates that load-reducing actions are needed to keep the net load within the deadband for net load control. The supervisory controller will limit controllable end uses and allow energy storage systems to respond if specified conditions are met. Stages farther from zero expand the conditions under which TES discharging is allowed, taking into account variables such as outside air conditions and storage state of charge (SOC).
- “Stage 0” indicates that the supervisory controller is not executing any load management actions. Default device-level control deadbands will apply.
- A negative stage number (−1 to −4) indicates that load-increasing actions are needed to keep the net load within the deadband for net load control. The supervisory controller will allow energy storage systems to respond if specified conditions are met. Stages farther from zero expand the conditions under which TES charging is allowed.

An example of a supervisory control staging strategy is summarized in Table 1. Tables such as these can be used to communicate the order in which different load-shaping resources are dispatched. Additional conditional statements can be defined to further constrain action. For example, the supervisory controller could restrict TES charging if outdoor air temperature conditions are unfavorable from an HVAC efficiency perspective or utility cost perspective. Device-level control actions are discussed in more detail in Section 3.5.

Table 1. Summary of Supervisory Control Stages and Device-Level Responses

<table>
<tr><th>Stage</th><th>Description</th><th>Battery</th><th>Cold TES</th><th>Hot TES</th><th>Lighting</th></tr>
<tr><td>+4</td><td rowspan="4">Load-reducing actions needed</td><td rowspan="3">Discharge if conditions allow</td><td rowspan="4">Raise the tank temperature trigger for charging the tank</td><td rowspan="4">Lower the tank temperature trigger for charging the tank</td><td rowspan="4">Reduce levels while maintaining minimum required level of service</td></tr>
<tr><td>+3</td></tr>
<tr><td>+2</td></tr>
<tr><td>+1</td><td rowspan="3">Idle</td></tr>
<tr><td>0</td><td>Default control actions apply</td><td>Default action</td><td>Default action</td><td rowspan="5">Default action</td></tr>
<tr><td>-1</td><td rowspan="4">Load-increasing actions needed</td><td rowspan="4">Allow charging under more conditions</td><td rowspan="4">Allow charging under more conditions</td></tr>
<tr><td>-2</td><td rowspan="3">Charge if conditions allow</td></tr>
<tr><td>-3</td></tr>
<tr><td>-4</td></tr>
</table>

2.2.3 Grid-Connected versus Microgrid Operation

In designing a supervisory control strategy, the team tried to avoid including actions that would negatively impact occupant productivity. This was particularly true when defining control logic for a "grid-connected" mode of operation during which it is assumed that the building is functioning normally with little or no justification for reductions in occupant service from lighting or HVAC systems.

In an "islanded-microgrid" mode of operation, however, it is possible that some reduction in occupant service might be tolerated, depending on assumptions about the conditions that would trigger islanding. For example, if islanding is only intended to occur during an emergency weather event that lasts a certain number of hours or days, then control stages could be designed that reduce service to predefined non-critical zones or spaces. In practice, such control actions would have to be accompanied by a clear communication plan so that occupants understand what to expect and how to respond during an emergency-driven islanding event.

An additional consideration for "islandable" building-scale microgrids is that their energy storage systems may be required to maintain a certain SOC during grid-connected mode. This would ensure that the storage systems are ready to provide backup power at a moment's notice in case the building switches to islanded-microgrid mode. Alternatively, a designer could specify a hybrid system in which certain batteries are dedicated to backup power, while other batteries are allowed to support load shaping when the building is in grid-connected mode.

2.3 Considerations for Selecting Compatible Building Systems

2.3.1 Efficiency

One well-established ZEB design principle is that energy efficiency should precede the design of renewable energy systems. The primary benefit of this approach is that designers can reduce the size—and costs—of on-site renewable generation systems if there is less demand for energy to begin with.

This principle also extends to next-generation ZEBs. For example, by reducing generation requirements, energy efficiency measures may also reduce the amount of energy storage required to reduce peak demand or exports down to a certain target. Unfortunately, next-generation ZEBs also come with a few energy efficiency twists that can add a layer of complexity to design decisions:

- Not all energy efficiency measures reduce peak demand. In fact, some can increase building load variability.
- Energy storage systems come with energy losses.

Because energy efficiency is one of several important metrics to ZEB owners and designers, the above caveats mean that care must be taken when evaluating alternative technologies and control approaches for load-shaping applications. Interactive effects between generation, storage, lighting, HVAC, and environmental conditions can be captured through the use of appropriate simulation tools. Designers can use these tools to ensure that load-shaping solutions also are energy-efficient and effective at maintaining satisfactory space conditions for occupant needs. Methods for conducting such evaluations are discussed further in Section 3.

2.3.2 Availability and Responsiveness of Resources

As a load management resource, dedicated energy storage systems are advantageous in that they are not constrained by other operational objectives. By contrast, lighting and HVAC systems are required to provide occupants with other services, and this restricts the timing and magnitude of their responses to calls to increase or decrease demand.

Initially, this observation led the team to suspect that the supervisory control algorithm should enable lighting and HVAC load management actions before dispatching a battery. In this way, constrained resources would be given time to act if available—after that, if the desired change in building net load had not yet been achieved, the battery could then act to make up the difference.

However, batteries can also respond more quickly than lighting and HVAC-integrated TES systems. This suggests potential tradeoffs between dispatch order and different metrics. For example, if controlling net load within a tight deadband is paramount for a particular application, then the ideal control algorithm might rely heavily on dedicated batteries. On the other hand, if a wider deadband is tolerable, and if dispatching other resources can potentially reduce battery size or extend battery life, then the ideal control algorithm might leverage multiple resources. Tradeoffs such as these can be evaluated with simulation tools as discussed in Section 3. Inputs to such simulations might come from a combination of design decisions, manufacturer specifications, and independent laboratory testing.

3 Using Simulation to Develop and Test Supervisory Controls

Whole-building energy simulation will play a critical role in the design and evaluation of next-generation ZEBs. In this project, the team demonstrated new capabilities for 1-minute-time step modeling of supervisory control strategies that link generation, storage, and flexible building loads.

3.1 Selected Tools

For this project, the team used EnergyPlus and OpenStudio for simulation-based analyses. EnergyPlus is a physics-based, whole-building energy simulation engine developed by DOE. OpenStudio is a DOE-sponsored cross-platform collection of software tools that can interface with EnergyPlus models to provide added usability and functionality (DOE 2016a, DOE 2016b).

These tools were selected for their ability to dynamically model building systems and their interactions, such as how lighting decisions can affect sizing and operation of HVAC equipment. It was also necessary to model these interactions at sufficiently granular timescales. The goal in this study was to manage net load profiles at a 15-minute timescale, which is commonly used by utilities to assess demand charges. To effectively control a 15-minute-interval output, the team needed to model building control actions and system responses at even shorter intervals. EnergyPlus met this requirement, allowing simulation time steps as short as 1 minute.

Within the EnergyPlus application, the team used the EnergyPlus Energy Management System (EMS) feature[4] to write and test algorithms for the ZEB supervisory controller. EnergyPlus EMS allows modelers to write programs to override default EnergyPlus controls with more sophisticated, customized supervisory control logic.

EnergyPlus EMS also enabled the team to model device-level controls more realistically. For example, the team replaced EnergyPlus's default simplified temperature set point managers with more realistic deadband controls, which allowed the direct-expansion HVAC equipment to cycle in a more realistic manner. The result was a more realistic HVAC load profile with 1-minute variability.

Another advantage of using the EnergyPlus/OpenStudio platform to develop these capabilities is that, in the future, it will be easier to improve solutions by using other OpenStudio features, such as the Parametric Analysis Tool. These features will allow future designers and researchers to rapidly vary system sizes, set points, control triggers, and control actions and more easily compare alternative approaches under a variety of conditions.

3.2 Initial Model of Existing Conditions

The first step in the modeling process was to create a whole-building energy model that approximates the existing VTIF building. The team constructed the initial model with information from three sources: (1) construction drawings, which provided envelope dimensions

[4] More details about EnergyPlus EMS are available in the *Application Guide for EMS* (DOE 2016c).

and materials; (2) information from building management; and (3) on-site inspection of building conditions. The model geometry is depicted in Figure 4.

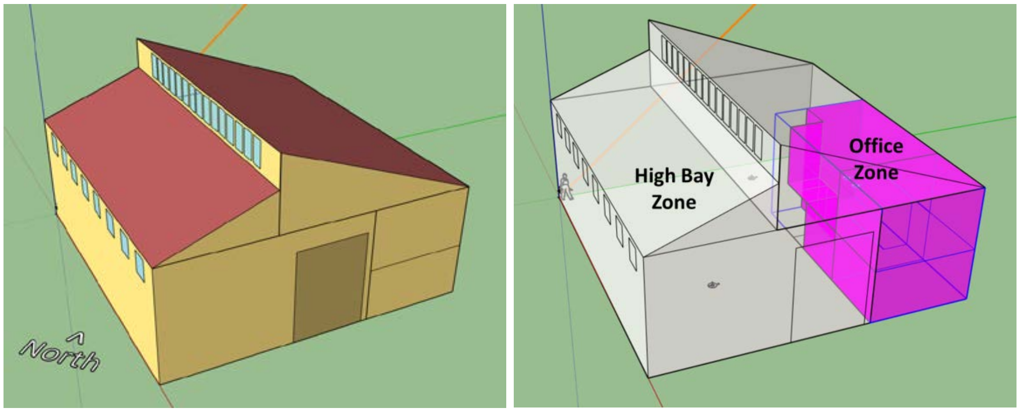

Figure 4. Initial model geometry and thermal zones

The team divided the building model into two major thermal zones: a "high bay zone" and an "office zone." The high bay zone is the larger zone, consisting primarily of an open area used for research activities, including tests of vehicle systems. It also includes a few small closet spaces. The office zone includes a shop room, storage room, and restroom on the first floor and a mezzanine office area, laboratory room, and mechanical closet on the second floor. The team included clerestory glass as a daylight source in the high bay zone but excluded the transfer of this light into the mezzanine space and omitted two tubular daylighting devices that were small contributors to office zone lighting.

3.3 Zero Energy Building Model Types

After modeling the building envelope and thermal zones of the VTIF, the team transformed the initial model into ZEB models by modifying and replacing several systems. Compared to the existing conditions at the VTIF, the ZEB models have more PV generation capacity, more efficient lighting fixtures, more sophisticated lighting sensors and controls, and an all-electric HVAC system. Changes to specific building systems are discussed in more detail in the following sections.

To enable comparisons of ZEB performance, the team generated three types of ZEB models with different load management resources:

- A baseline ZEB model without batteries
- A baseline ZEB model with batteries
- Alternative ZEB models with multiple load-shaping resources and supervisory control.

The first baseline model represents a first-generation ZEB, which lacks energy storage. Using an annual simulation with a 1-minute time step, the model's 15-minute average net load ranged from a high of 30 kilowatts (kW) to a low of -49 kW. These values provided points of reference for the metrics of peak demand and peak export rate.

Because a major focus of this research was on developing supervisory control strategies, the team also created a second baseline model that included batteries but lacked other load-shaping resources. The second baseline represents what might be achievable with batteries alone.

The alternative models were the most sophisticated, including a combination of batteries, TES, flexible building loads, and advanced supervisory controls. The alternative models can be used to compare different ZEB system designs and supervisory control concepts.

3.4 Zero Energy Building Subsystem Models

Many of the ZEB model inputs were based on specifications selected in the integration effort described in Section 2. The following sections describe the modifications made to various building systems to create the ZEB models. They also describe methods employed to model contributions to load profile variability at timescales as short as 1 minute.

3.4.1 PV and Battery

The team estimated the real VTIF building's annual PV generation to equal about 85% of its annual electricity use, but the building also makes heavy use of natural gas. When both energy sources are counted, the annual PV generation equals only about 23% of annual total energy use on a site energy basis or about 44% on a source energy basis.[5]

To create the ZEB baseline models, the team increased the modeled PV capacity to ensure that annual on-site renewable generation would equal or exceed the model's total annual energy use. The team also replaced natural gas end uses with electrical end uses to eliminate on-site use of fossil fuels and maximize the potential to match on-site generation with on-site loads.

In simulation studies, modelers have to make many decisions about which details to include for accuracy and which details to neglect for computational and analytical efficiency. For example, the real VTIF has a solar-tracking PV array, but modeling the benefit of solar tracking was not critical for this study. So when the team needed to increase PV size, the team modeled the new PV system as a fixed-angle horizontal array, as shown in Figure 5.

[5] Site energy is energy use within a facility's site boundary. Source energy includes "site energy plus the energy consumed in the extraction, processing, and transport of primary fuels such as coal, oil and natural gas; energy losses in thermal combustion in power generation plants; and energy losses in transmission and distribution to the building site" (DOE 2015).

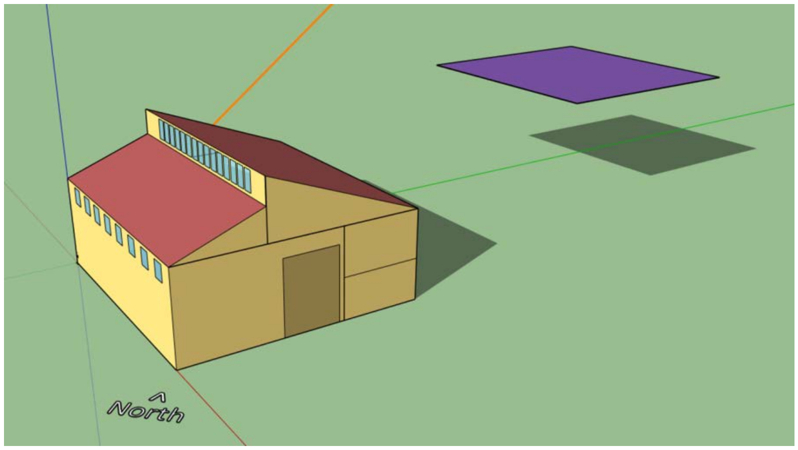

Figure 5. Relative position of building and PV array

On the other hand, it was important in the simulation study to model weather-induced subhourly variability in PV generation. To capture this kind of variability, the team provided the EnergyPlus model with 1-minute-interval insolation and weather data collected in 2014 from the Solar Radiation Resource Laboratory, which is located on the same campus as the VTIF.

After making these changes, the peak direct current PV output was 64 kW using a 1-minute time step, and the peak alternating current output from the inverter was 61 kW. The PV system size was chosen to be large enough that all models in this study, with or without energy storage, would be ZEBs. The baseline model without batteries uses the least energy overall, and it exports about 1.5 times more energy than it imports on an annual basis. Other models in this study have the same PV capacity but require more energy to offset storage losses, so they have import-to-export ratios less than 1.5.

Using a 1-minute simulation time step, a 15-minute interval for evaluating net load, and a 1-year simulation period, the baseline model without batteries had a peak demand of 32 kW and a peak export rate of 49 kW.

For cases with a battery, the team used the simpler of two battery options available in EnergyPlus to model a generic stationary battery, and the team selected performance characteristics achievable with lithium ion as a representative technology. Selected values included a charging efficiency of 90% and a discharging efficiency of 90%, which corresponds to a roundtrip efficiency of 81%. The team set the maximum charge and discharge rate at 30 kW based on an inverter presently installed at the VTIF.

3.4.2 HVAC and Thermal Energy Storage

During the integration effort described in Section 2, the team reviewed multiple alternatives to the existing evaporative cooling system and natural gas heating system. Key considerations for the new system were as follows:

- For both cooling and heating, high priorities included energy efficiency, controllability, and feasibility of integration with existing building systems.

- On the cooling side, the team needed a technology that could provide comparable space conditioning performance for baseline and alternative cases. The existing evaporative cooler could not be easily tied to TES, so it could not be used in both baseline and alternative models. This led the team to look at other cooling technologies that could be used for both cases.

- On the heating side, the team wanted to switch to an electricity-driven system to eliminate on-site use of fossil fuels and maximize the potential to use renewable generation on site.

Based on these criteria, heat pumps emerged as a strong candidate. Among these technologies, a ground-source heat pump system would have provided the greatest energy efficiency through the broadest range of weather, and they have been shown to be competitive on a life cycle cost basis for certain kinds of full-scale building construction projects. Ground-source heat pump capital costs, however, were beyond the budget for this comparatively small research project, so the team selected a split-system air-to-water heat pump (AWHP) technology for the integration effort described in Section 2.

The team used system specifications from the integration effort as a starting point for the simulation study. In the ZEB models, the team expanded the heat pump approach to cover cooling and heating requirements for the whole building. Figure 6 summarizes the resulting system topology used in the simulation study.

Each split-system AWHP consists of an outdoor unit and an indoor unit connected by a refrigerant loop. The outdoor unit exchanges heat with ambient air, serving as a condenser when the AWHP is in cooling mode and serving as an evaporator when the AWHP is in heating mode. The indoor unit contains a heat exchanger for transferring heat from a water loop that circulates chilled or hot water to other HVAC equipment in the building.

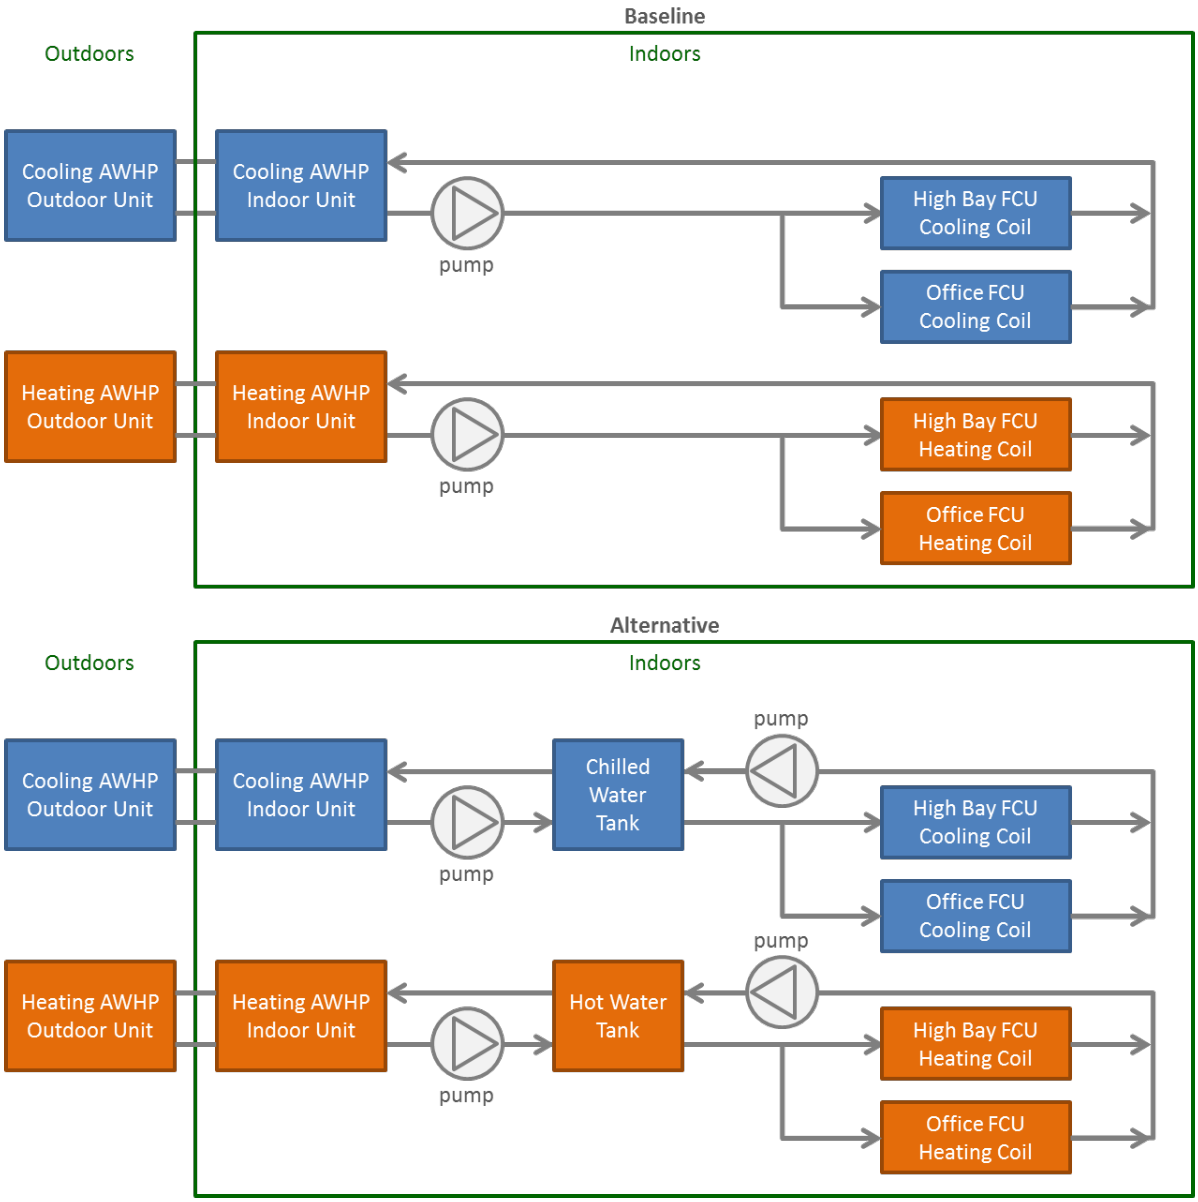

Figure 6. HVAC hydronic system conceptual plan for the baseline and alternative models

Each ZEB model includes two AWHPs: one for cooling and one for heating. Each thermal zone has one fan coil unit (FCU), which contains a chilled water coil and a hot water coil. The baseline and alternative models differ, however, in how water from the AWHPs is delivered to the FCUs:

- In the baseline models, chilled water is pumped from the cooling AWHP to the chilled water coil in each zone. Similarly, hot water is pumped from the heating AWHP to the hot water coil in each zone.
- In the alternative models, water is pumped from each AWHP to a tank that provides TES; when cooling or heating is needed, water is pumped from the appropriate tank to the FCUs.

To capture the effects of HVAC cycling on load profile variability, the team used the EnergyPlus EMS programming environment to define control deadbands for space temperature. This enabled the team to model HVAC control more realistically, capturing the effect of equipment turning on and off as space temperature enters and exits the control deadband. Building energy modelers may simplify these details for traditional design projects in which the focus is on calculating monthly and annual energy use—for example, they might estimate the hourly energy consumption for HVAC systems, assuming a certain amount of cycling per hour. In projects concerned with shaping a 15-minute-interval load profile, however, it is important to model HVAC cycling effects that can happen at scales between 1 and 15 minutes.

3.4.3 Lighting

EnergyPlus was selected for this project because of its abilities to model system interactions and supervisory control at appropriate timescales. As part of this effort, the team also modeled lighting systems with EnergyPlus and leveraged the EnergyPlus EMS feature to actuate high-resolution lighting controls. It should be noted that this program is not typically used for advanced modeling of high-resolution occupancy or daylighting control, but the level of accuracy achievable with EnergyPlus was sufficient for modeling lighting energy use and load profiles for the purposes of this study.

Each ZEB model included two primary lighting zones that align with the model's two thermal zones: the high-bay zone and the office zone. To enable fixture-based occupancy control, the team divided these zones further into additional spaces. These included seven high bay spaces, one mezzanine office space, one mezzanine laboratory space, and additional individual spaces for each enclosed closet and room in the building. For each space, the team defined individual lighting fixture power values.

3.4.4 Other Loads

The team used 1-minute-interval data collected in 2014 from an electricity meter at the VTIF to approximate electricity used by plug and process loads. This was an imperfect method because the meter did not isolate plug and process loads from other end uses. The meter mixed in lighting loads and a small fraction of HVAC loads, and it excluded some research equipment loads that would have been helpful to include. For this study, however, the metered dataset was an adequate proxy because it captured an important source of load profile variability at the 1-minute timescale. In this case, the team did not need the magnitude of plug and process loads to be highly accurate, but the variability of this end use needed to be realistic. Including such variability would help the model to demonstrate how a supervisory control would respond to a representative load profile.

In EnergyPlus, the team imported the 1-minute-interval dataset and divided the load between the two thermal zones. The team allocated 80% to the high bay zone and 20% to the office zone, basing the allocation roughly on floor area.

3.5 Supervisory Control Stages and Responses

The team used the EnergyPlus EMS programming environment to define supervisory controls based on the general strategy introduced in Section 2.2. The team wrote a combination of programs and subroutines to define device-level responses for each supervisory control stage, as well as conditions for transitioning between stages. Supervisory control stages and device-level responses for one approach tested in this study are summarized in Table 2. (An alternative approach is discussed in Section 3.6.4.)

Table 2. Summary of Supervisory Control Stages and Device-Level Responses

<table>
<tr><th>Stage</th><th>Description</th><th>Battery</th><th>Cold TES</th><th>Hot TES</th><th>Lighting</th></tr>
<tr><td>+4</td><td rowspan="4">Load-reducing actions needed</td><td rowspan="3">Discharge if conditions allow</td><td rowspan="4">Raise the tank temperature trigger for charging the tank</td><td rowspan="4">Lower the tank temperature trigger for charging the tank</td><td rowspan="4">Reduce levels while maintaining required minimum level of service</td></tr>
<tr><td>+3</td></tr>
<tr><td>+2</td></tr>
<tr><td>+1</td><td rowspan="3">Idle</td></tr>
<tr><td>0</td><td>Default control actions apply</td><td>Default action</td><td>Default action</td><td rowspan="5">Default action</td></tr>
<tr><td>-1</td><td rowspan="4">Load-increasing actions needed</td><td rowspan="4">Allow charging under more conditions</td><td rowspan="4">Allow charging under more conditions</td></tr>
<tr><td>-2</td><td rowspan="3">Charge if conditions allow</td></tr>
<tr><td>-3</td></tr>
<tr><td>-4</td></tr>
</table>

To define conditions for changing stages, the team defined multiple control parameters that can be adjusted to achieve different results. These parameters included:

- ***A low target net load and a high target net load.*** The low target net load and high target net load form a net load control deadband (e.g., -10 kW to 10 kW). When the facility's net load deviates from this range, a supervisory control stage change may occur, provided that other conditions are met.

- ***Triggers for jumping faster to Stage -4 or Stage +4.*** Smaller deviations from the net load control deadband will result in the stage moving only one level up or down. But if the net load falls below the trigger to jump to Stage -4 or exceeds the trigger to jump to Stage +4, a more dramatic system response is possible.

- ***A rolling average calculation timeframe for evaluating net load.*** This parameter mitigates the effects of high-frequency variability in net load. When deciding whether to change stages, the supervisory controller will evaluate the rolling average net load over a user-specified timeframe (e.g., 5 minutes). A shorter timeframe enables a faster response, while a longer timeframe improves controller stability.

- ***Minimum hold times for remaining in a particular stage before allowing different kinds of stage changes.*** This parameter provides device-level controllers with some time to respond to one stage change before another stage change occurs. Different subsystems have different response times. For example, batteries can generally respond more quickly than systems tied to building loads. Increasing hold times will allow more subsystems to respond to each given stage change, but if hold times are too long, the supervisory controller may not be able to dampen steep changes in net load.

An example of supervisory control parameter values and a resulting response is shown in Figure 7. Here, the low target net load is -10 kW, and the high target net load is 10 kW. The triggers for jumping to the lowest and highest stages are -12 kW and 12 kW, respectively. The rolling average calculation timeframe for evaluating net load is 5 minutes, denoted by the thin blue line. The thick green line marks the 15-minute average net load, which is commonly used by utilities for assessing demand charges.

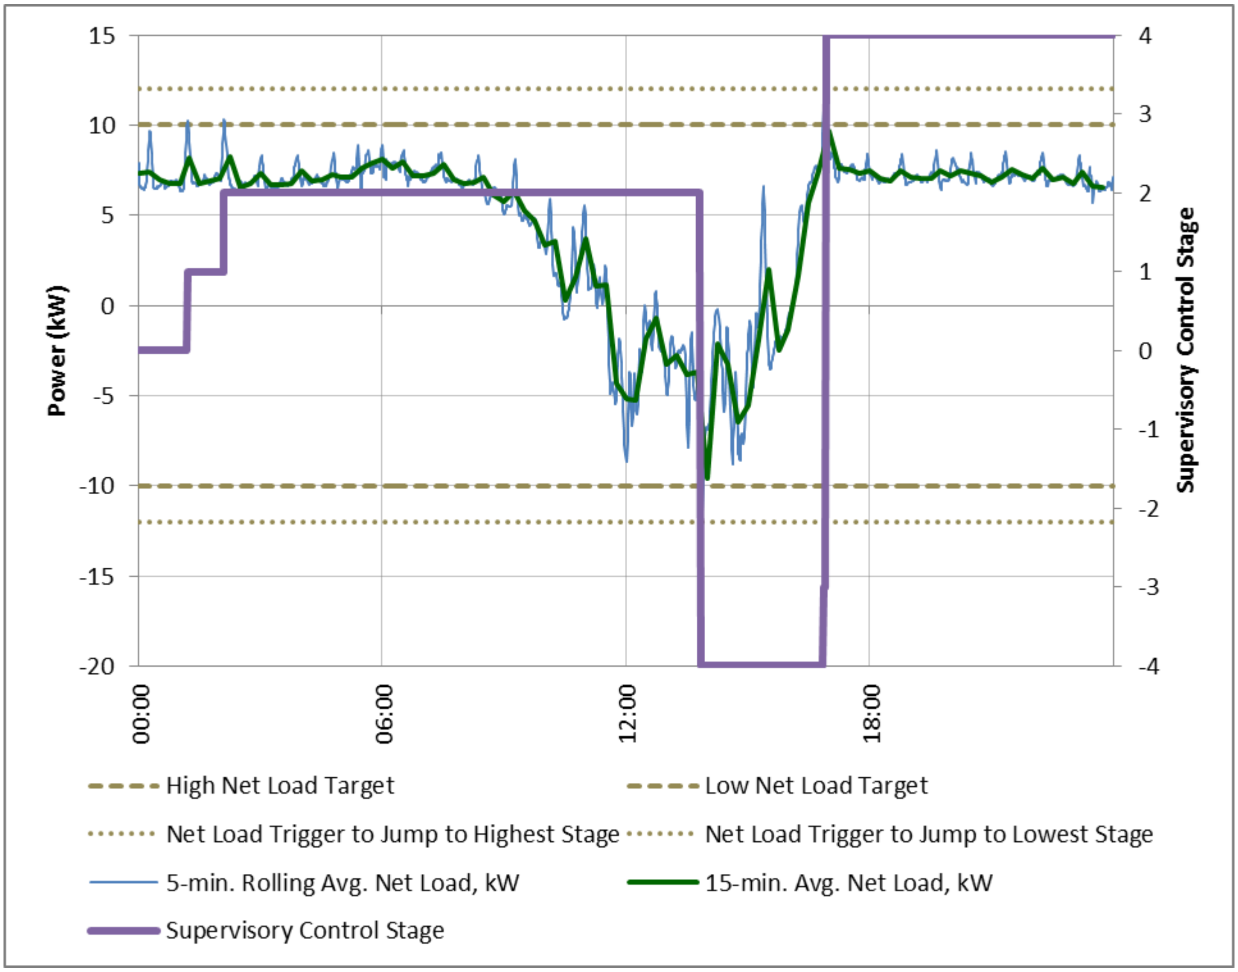

Figure 7. Example supervisory control parameters and system response

An example set of hold times is provided in Table 3. In this example, the system can step downward more quickly when the current stage is negative, and the system can step upward more quickly when the current stage is positive. The shortest hold times are reserved for the jump conditions on the right side of the table—this is based on the assumption that when the jump triggers are exceeded, a fast response is needed.

Table 3. Example Hold Times by Supervisory Control Stage

Stage	Description	Minimum Hold Time Before Moving Down 1 Stage	Minimum Hold Time Before Moving Up 1 Stage	Minimum Hold Time Before Jumping to Stage -4	Minimum Hold Time Before Jumping to Stage +4
+4	Load-reducing actions needed	10 minutes	N/A	1 minute	N/A
+3		10 minutes	5 minutes	1 minute	1 minute
+2		10 minutes	5 minutes	1 minute	1 minute
+1		10 minutes	5 minutes	1 minute	1 minute
0	Default control actions apply	10 minutes	10 minutes	1 minute	1 minute
-1	Load-increasing actions needed	5 minutes	10 minutes	1 minute	1 minute
-2		5 minutes	10 minutes	1 minute	1 minute
-3		5 minutes	10 minutes	1 minute	1 minute
-4		N/A	10 minutes	N/A	1 minute

In any given time step, the current supervisory control stage determines how device-level controllers will respond. Subsystem responses are described in the following sections.

3.5.1 Battery Response

The fastest device-level controller is the battery controller. In this study, battery response time was allowed to be as fast as the 1-minute simulation time step, though real-world response times can be even faster. HVAC and lighting systems can also turn on and off quickly, but the magnitude of their response is less consistent, and their availability for load management is constrained by their need to provide other building services.

If the supervisory controller stage allows the battery to charge or discharge, it will attempt to make up the difference between the current net load and the closest target net load. The only other constraint placed on battery charging and discharging is a SOC constraint. For example, the battery could be directed to maintain a SOC between 20% and 100%. Two detailed examples of alternative battery dispatch approaches are compared in Section 3.6.4.

3.5.2 HVAC and Thermal Energy Storage Response

Examples of HVAC and TES system responses are shown in Table 4 and Table 5. When the supervisory control stage is high (closer to Stage +4), the controller changes tank temperature set points to try to reduce demand. When the supervisory control stage is low (closer to Stage -4), the controller allows the chilled water tank to be "overcharged" to a lower set point, and it allows the hot water tank to be overcharged to a higher set point, provided that outside air conditions are favorable. The outside air temperature constraint prevents overcharging from occurring under inefficient operating conditions. This constraint is partly relaxed as stages approach Stage -4.

Table 4. Chilled Water System Responses by Supervisory Control Stage

<table>
<tr><th>Stage</th><th>Description</th><th>Chilled Water Tank High Temperature to Trigger Charging in Occupied and Preoccupancy Hours</th><th>Chilled Water Tank High Temperature to Trigger Charging in Other Hours</th><th>Outside Air Low Temperature Requirement for Overcharging of Chilled Water Tank</th></tr>
<tr><td>+4</td><td rowspan="4">Load-reducing actions needed</td><td>47°F</td><td rowspan="4">47°F</td><td rowspan="5">N/A</td></tr>
<tr><td>+3</td><td>46°F</td></tr>
<tr><td>+2</td><td rowspan="7">45°F</td></tr>
<tr><td>+1</td></tr>
<tr><td>0</td><td>Default control actions apply</td><td rowspan="5">45°F</td></tr>
<tr><td>-1</td><td rowspan="4">Load-increasing actions needed</td><td rowspan="2">65°F</td></tr>
<tr><td>-2</td></tr>
<tr><td>-3</td><td>70°F</td></tr>
<tr><td>-4</td><td>75°F</td></tr>
</table>

Table 5. Hot Water System Responses by Supervisory Control Stage

<table>
<tr><th>Stage</th><th>Description</th><th>Hot Water Tank Low Temperature to Trigger Charging in Occupied and Preoccupancy Hours</th><th>Hot Water Tank Low Temperature to Trigger Charging in Other Hours</th><th>Outside Air High Temperature Requirement for Overcharging of Hot Water Tank</th></tr>
<tr><td>+4</td><td rowspan="4">Load-reducing actions needed</td><td>118°F</td><td rowspan="4">118°F</td><td rowspan="5">N/A</td></tr>
<tr><td>+3</td><td>119°F</td></tr>
<tr><td>+2</td><td rowspan="7">120°F</td></tr>
<tr><td>+1</td></tr>
<tr><td>0</td><td>Default control actions apply</td><td rowspan="5">120°F</td></tr>
<tr><td>-1</td><td rowspan="4">Load-increasing actions needed</td><td rowspan="2">50°F</td></tr>
<tr><td>-2</td></tr>
<tr><td>-3</td><td>45°F</td></tr>
<tr><td>-4</td><td>40°F</td></tr>
</table>

3.5.3 Lighting Response

For lighting, the team considered three approaches, which are summarized in Table 6. The selected actions were informed by previous research on occupant detection and acceptance of various illuminance differences and change rates relative to a space's maximum illuminance (Akashi and Neches 2004). Other researchers have employed the findings of Akashi and Neches in subsequent demand response studies, one of which also informed the selection of strategies for this study (Newsham and Birt 2010).

The alternatives illustrate a range of actions with varying levels of impact. The included actions vary in terms of response time and detectability, but the team only included actions that research suggests will be acceptable to occupants. Resulting impacts on whole-building metrics are discussed further in Section 3.6.1.

In a real-world application, occupant comfort and acceptance of such lighting strategies should be evaluated with lighting-specific modeling tools or field observations to assess resulting illuminance, surface luminance, and transition times.

Table 6. Alternative Lighting System Responses by Supervisory Control Stage

<table>
<tr><th>Stage</th><th>Lighting Alternative 1</th><th>Lighting Alternative 2</th><th>Lighting Alternative 3</th></tr>
<tr><td>+4</td><td rowspan="4">Lighting power reduction that maintains the designed level of service:

Rapid dimming/response to an undetectable illuminance range, with a maximum power reduction of:
20% if the space is not daylighted
40% if the space is daylighted[b]</td><td rowspan="4">Lighting power reduction that maintains the designed level of service:

Slow dimming/response to an undetectable illuminance range, with a maximum power reduction of:
30% if the space is not daylighted
60% if the space is daylighted[b]</td><td>Maximum illuminance and lighting power reduction that impacts the designed level of service (detectable but acceptable to occupants):
40% if not daylighted
60% if daylighted
80% if daylighted and unoccupied[2]</td></tr>
<tr><td>+3</td><td>Alternative 2 as implemented:
30% if not daylighted
60% if daylighted[b]</td></tr>
<tr><td>+2</td><td>Alternative 2 with a maximum power reduction modification:
20% if not daylighted
40% if daylighted[b]</td></tr>
<tr><td>+1</td><td>Alternative 2 with a maximum power reduction modification:
10% if not daylighted
20% if daylighted[b]</td></tr>
<tr><td>0 to -4</td><td>Lighting baseline[a]</td><td>Lighting baseline[a]</td><td>Alternative 2 with a maximum power reduction modification:
20% if not daylighted
40% if daylighted[b]</td></tr>
</table>

[a] Includes automatic daylighting control only and low-frequency occupancy response (i.e., building-level occupancy schedules) comparable with a first-generation ZEB design.
[b] Power and illuminance reductions assumed to be equivalent for this study.

3.6 Illustrative Comparisons

To facilitate iteration and improvement of supervisory control designs, the team defined a wide range of variables in EnergyPlus EMS programs to represent important design parameters that could be changed later. To demonstrate how parametric analyses can assist engineers during the development of a ZEB supervisory controller, the following sections provide illustrative examples of controller design questions and how simulation tools like EnergyPlus can be used to help answer them.

Each example includes a discussion of one or more of the following metrics:

- ***Peak demand.*** This is the maximum net load, averaged over 15-minute intervals.
- ***Peak export rate.*** This is maximum export rate, averaged over 15-minute intervals. It is also equal to the absolute value of the minimum net load.
- ***Net load range.*** This is the difference between the maximum and minimum values for 15-minute-average net load. A ZEB supervisory controller could reduce a building's net load range by reducing its peak demand, peak export rate, or both.
- ***Total facility energy use.*** This is the total energy used by building subsystems. It excludes PV generation and battery losses.
- ***Total facility energy use plus battery losses.*** This is the sum of the total energy used by building subsystems and the energy losses associated with battery operation.
- ***Imported energy.*** This is the total energy delivered from the grid to the building.
- ***Battery roundtrip throughput.*** This is a simplified measure of battery cycling. For a given period of time, the battery roundtrip throughput is the sum of the battery's cumulative charge and cumulative discharge, divided by two. It neglects the impact of factors such as battery SOC and temperature. It can be replaced with more accurate metrics if a battery type or product has been specified.
- ***Ratio of exported electricity to imported electricity.*** This is a measure of how well a building is meeting or exceeding ZEB requirements. To be a ZEB, a building must have an export-to-import ratio of one or greater on annual basis.
- ***Ratio of PV generation used on site to PV generation exported.*** This is a measure of how well a building is able to leverage its own on-site generation and balance PV supply and demand within its site boundary.

Of these objectives, the first three are the most directly related utility concerns with ZEB load profile variability. Peak demand is of interest to both utilities and customers—the customer's interest in reducing peak demand is directly tied to a financial interest in reducing demand charges. Reductions in peak export rate and net load range are not typically incentivized by utility rate structures today, but this might change in the future as customer-sited renewable generation becomes more prevalent.

Reductions in battery cycling may extend battery life and would be of interest to the battery owner, which could be the building owner, the utility, or a third party. The remaining metrics in this list reflect building owner interests tied to operating costs or sustainability goals and policies.

3.6.1 How Do Lighting Control Choices Affect System Response Time?

Figure 8 illustrates how differences in lighting control algorithms can impact system response time. For a representative day in January, the figure shows simulated performance for three different models, labeled Alternative 1, Alternative 2, and Alternative 3. The models are similar except for how their lighting controls are defined:

- In all three cases, the supervisory controller applies wattage reductions based on whether or not daylight is available in the space at the time step. For Alternatives 2 and 3, daylight-based wattage reductions are only allowed in 1%–2% increments (assuming a 1-minute simulation time step).
- For Alternative 3, when the system is in supervisory control stage +4, the model uses a random number generator calculation to assign an occupancy state to each space in the high bay zone. It then reduces lighting power to the unoccupied spaces in this zone. The model included a 10-minute smoothing algorithm so that lighting conditions change at a reasonable frequency. This control strategy mimics a task-tuning lighting design approach, though occupant feedback about reduced illuminance was not modeled. In future model iterations, the implementation of an occupant feedback loop within the stages would allow use of task tuning as a demand management strategy while incorporating feedback about whether or not the intended level of service is maintained.

In Figure 8, between about 8:30 and 8:45 a.m., the supervisory control stage begins to drop, which means the building's net load is also dropping. In response, the lighting power is increasing. Compared to Alternative 1, Alternatives 2 and 3 have a slower lighting system response, taking more than 10 minutes longer to reach the peak load for the morning.

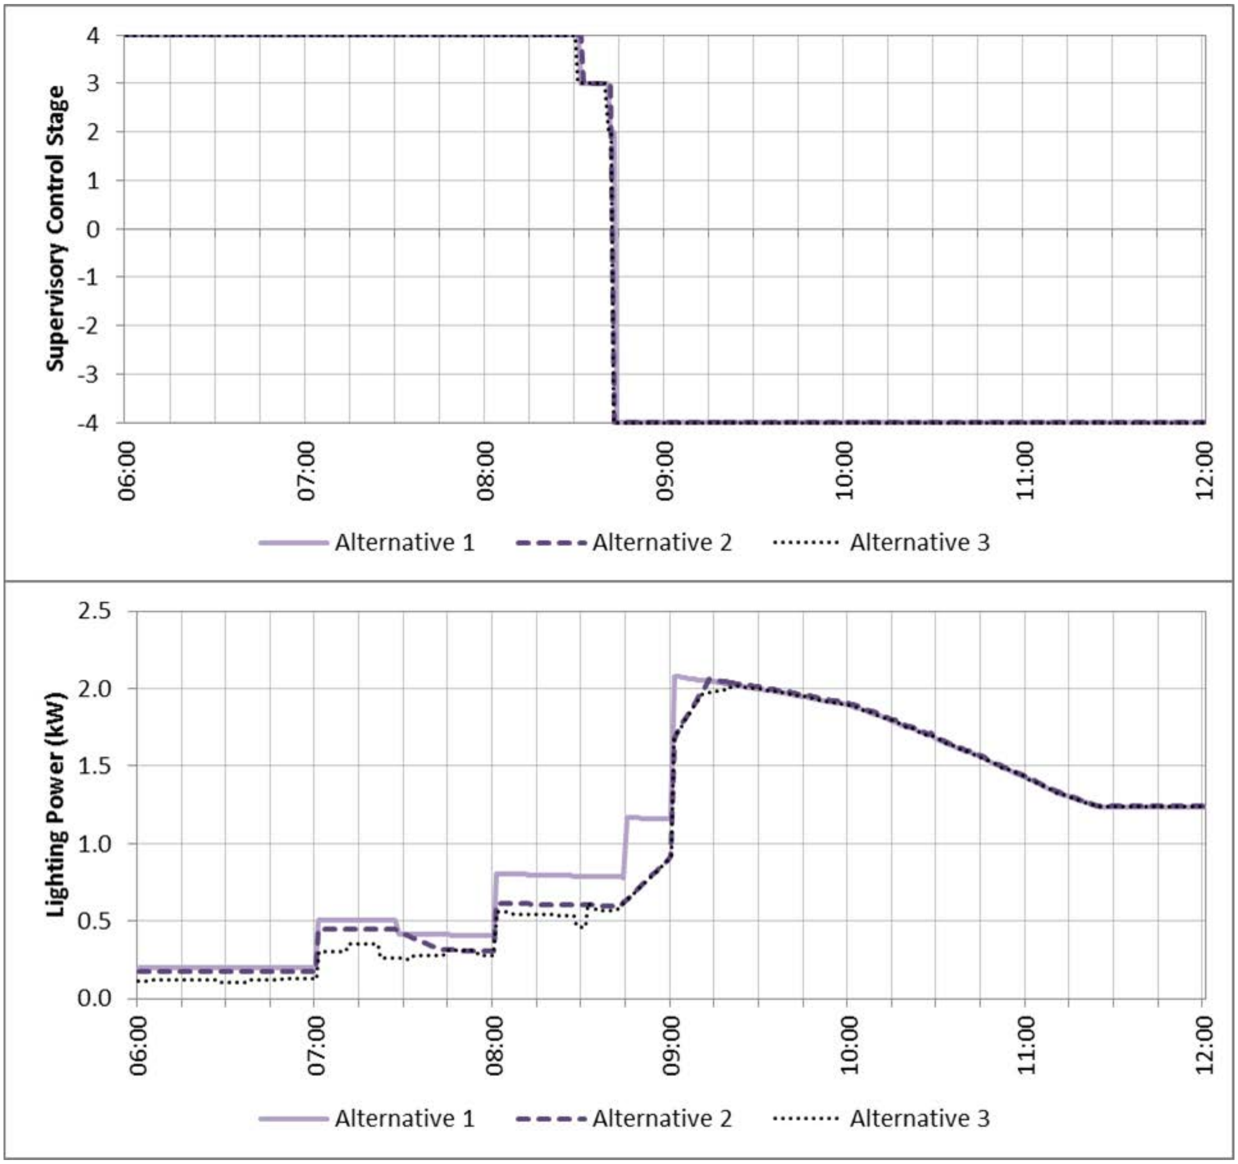

Figure 8. Lighting alternatives and their responses to supervisory control stages

When a difference in response time between alternative models is observed, further investigation may be required to determine whether the overall impact on building metrics is positive or negative. In this case, the three alternatives achieve similar whole-building results. Compared to Alternative 1, Alternative 2 used 0.3% less energy and lowered peak demand by 1.3% annually. Alternative 3 performed only a little better, using 0.5% less energy and lowering peak demand 2.5% compared to Alternative 1.

In situations such as these in which whole-building results are similar across alternatives, the designer should evaluate the alternative strategies' impacts on occupant comfort metrics to identify the preferred option.

3.6.2 How Do Thermal Energy Storage Temperature Choices Affect Whole-Building Performance?

Like batteries, TES systems have a tendency to increase overall system energy use, though there are some exceptions to this general rule (Kung, Deru, and Bonnema 2013, 9). If TES is integrated with a heat pump, the TES system can be used to shift the timing of the heat pump's electricity use and potentially increase its efficiency. Whether this efficiency gain outweighs TES energy losses will depend on project-specific design decisions and site-specific weather conditions.

In this study, HVAC and TES performance is also affected by the design of the supervisory controller and other building systems. In particular:

- Space conditioning needs will impact the availability of TES for load shaping.
- Supervisory control design will affect the timing and magnitude of TES charging and discharging, which in turn will affect TES losses and HVAC energy use.
- Lighting system efficiency and control changes will affect thermal loads, which in turn will affect size requirements for HVAC and TES components.

These interactive effects can be explored through parametric analyses in which one or more system characteristics are varied and one or more performance metrics are evaluated. To illustrate this with an example comparison, the team considered the following question: "How does the hot water tank's maximum temperature set point affect whole-building performance metrics?"

To answer this question, it is necessary to select metrics by which alternative models can be compared. Suppose for this example that the following metrics are of interest to project stakeholders: peak demand, peak export rate, and total facility energy use.

Ideally, it would be desirable to reduce all three of these values, but these priorities may conflict. Between peak demand reduction and peak export reduction, the former application is more familiar to commercial building owners, particularly those with demand charges on their utility bills. For these stakeholders, peak demand reduction will generally take precedence, provided that their buildings are in locations that have no penalties for exporting power.

Nonetheless, as ZEBs and distributed generation become more prevalent, utility stakeholders may assign the greatest value to buildings that can reduce both peak demand *and* peak exports. This is because a building's net load range is a measure of one aspect of ZEB load profile variability. A ZEB supervisory controller could reduce a building's net load range by reducing its peak demand, its peak export rate, or both.

Regardless of how these metrics are prioritized, it can be difficult to predict how various design parameters will affect these outcomes. It is therefore helpful to model a range of cases, using the EnergyPlus simulation engine to account for physical system interactions and subtler effects arising from the difference in inputs.

For the following example, the team varied the maximum allowable temperature of the hot water tank from 120°F, which is relatively low for hydronic heating, to 180°F, which is the upper limit of an insulated water tank product that the team had previously considered installing in the VTIF. All other design parameters were held constant.

Figure 9 shows that whole-building energy use increased with maximum hot water tank temperature, indicating that thermal losses from the tank dominated other effects. The primary driver here is that radiative losses are higher when stored water is farther from room temperature. Additionally, closer inspection of daily data showed that the modeled system has insufficient storage capacity to avoid running the AWHP when the weather is very cold. Other design alternatives may be able to overcome this obstacle by increasing storage capacity or switching to a different type of system (e.g., switching to a ground-source heat pump).

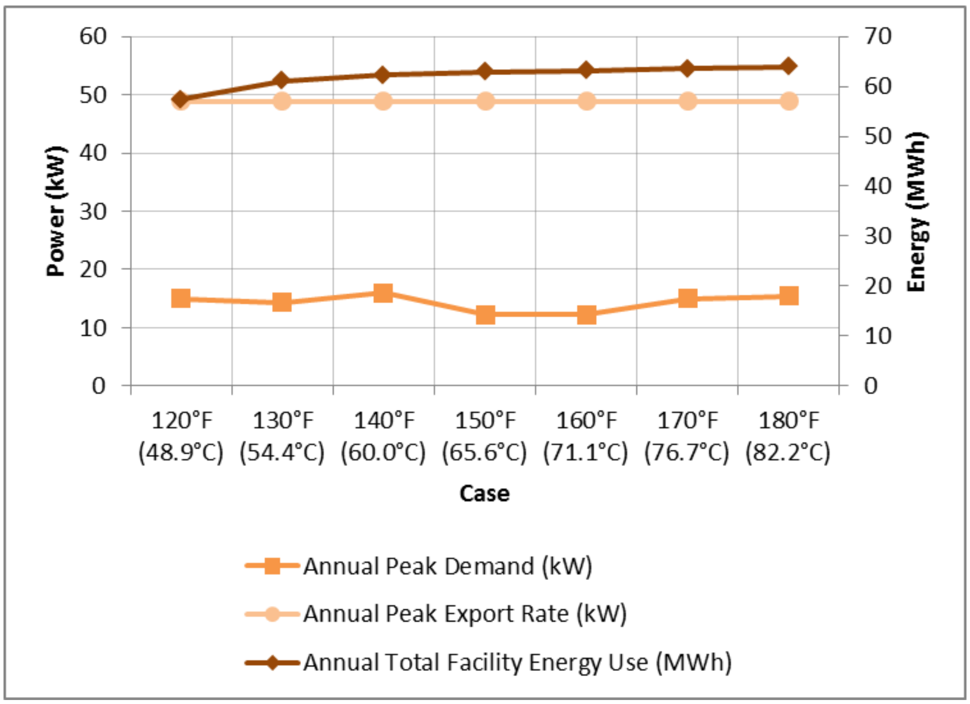

Figure 9. Maximum hot water tank temperature versus selected annual metrics

Figure 9 also shows that maximum hot water tank temperature has a negligible effect on the annual peak export rate. For these models, it appears that other design parameters have a stronger effect on this metric.

On the other hand, the maximum hot water tank temperature appears to have a variable effect on annual peak demand. To examine this relationship more closely, peak demand can be evaluated at smaller timescales.

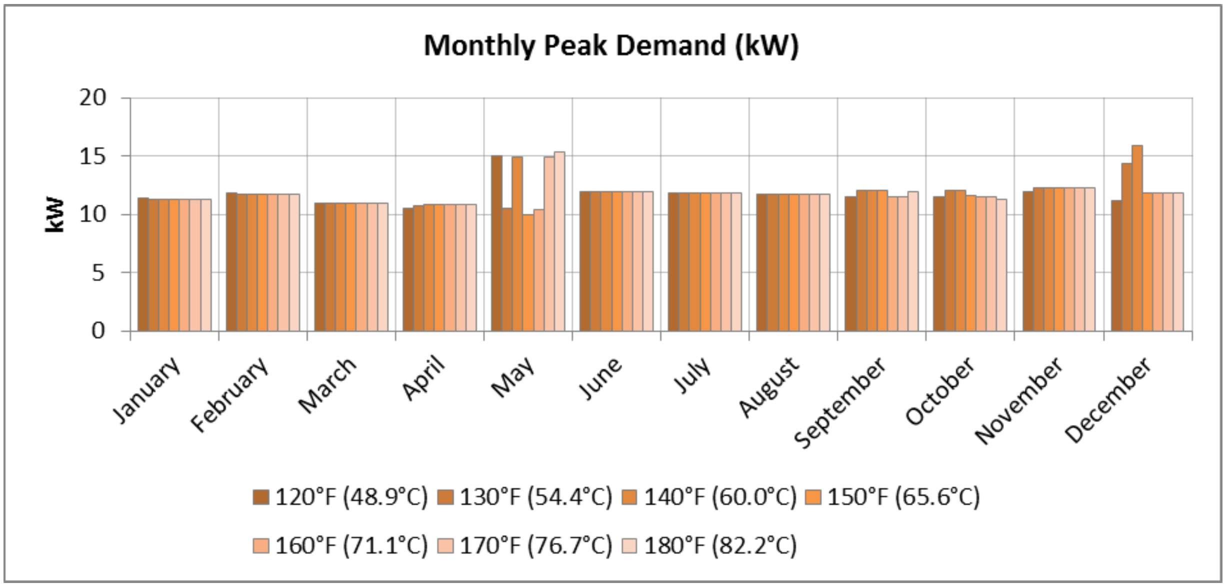

Figure 10. Maximum hot water tank temperature versus monthly peak demand

Figure 10 shows that monthly peak demand spikes in May and December for a subset of the modeled cases. The relationship between these incidents and tank temperature is not immediately obvious, but more clues may be found with subhourly data. A review of higher-resolution outputs reveals that there was a spike in plug and process load electricity on the morning of May 9, and the models differed in their ability to dampen this spike. Two models with different responses to this event are shown in Figure 11.

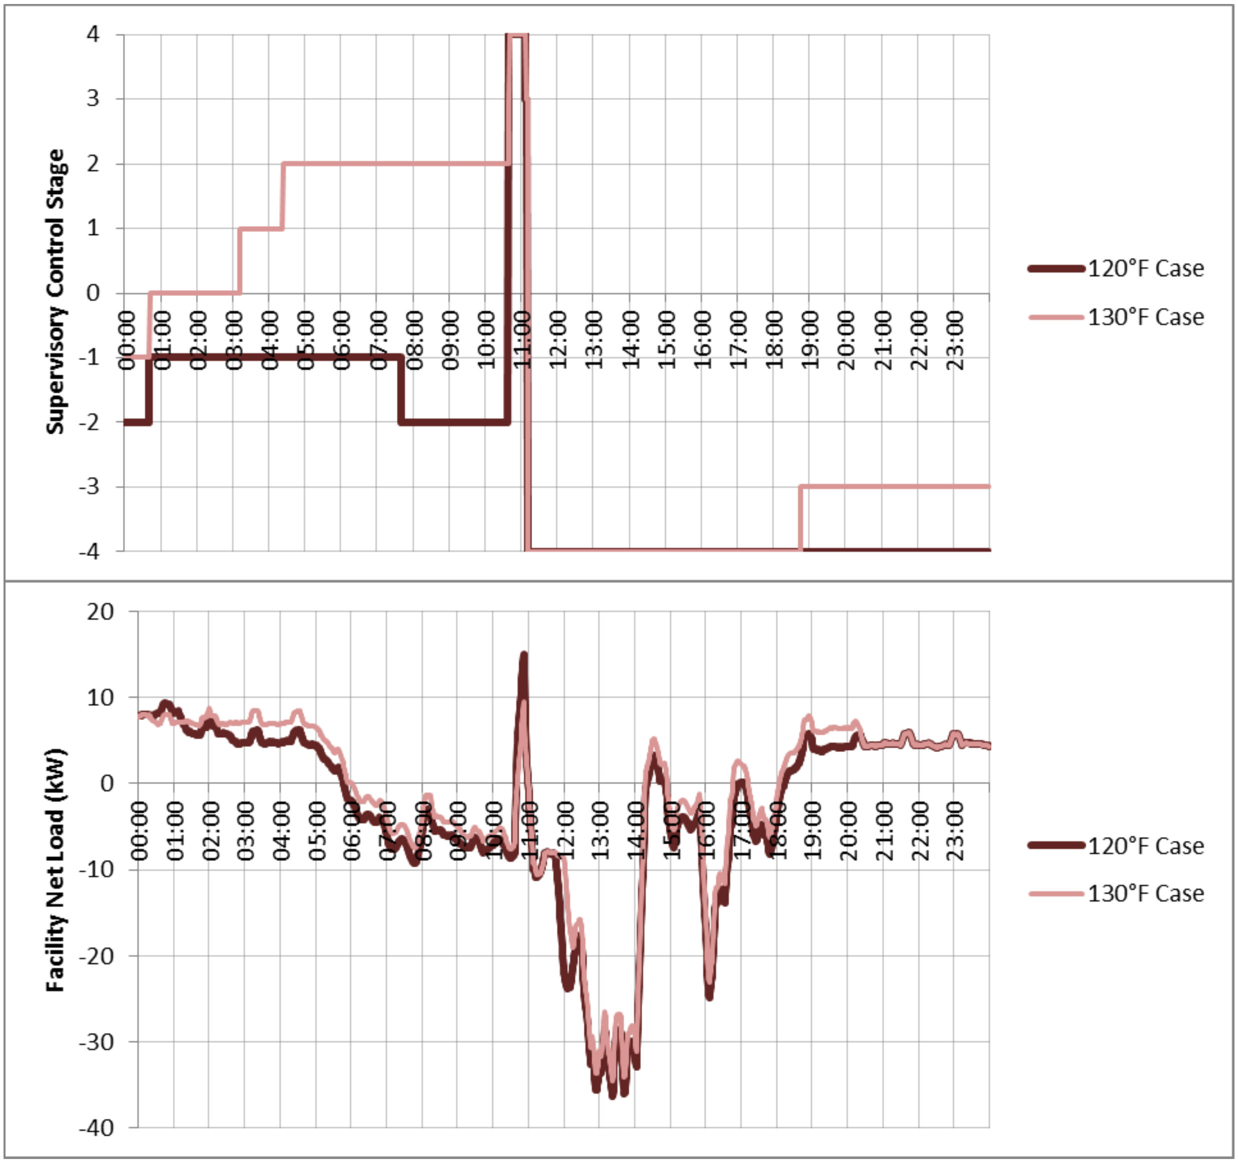

Figure 11. Differences in whole-building response to a load spike on May 9

The models have the same supervisory control logic and battery capacity, so it might seem reasonable to expect that they would respond equally well to this spike in load. The reason the models have different outcomes is that differences in the maximum hot water tank temperature cause each building to have a different load profile history leading up to the May 9 load spike.

As a result, some of the models are in a different supervisory control stage at the start of the load spike. Additionally, the team had selected supervisory control logic that only allows the battery to discharge in Stages +2 and higher. Thus, the models that can respond most quickly are the ones that are already in Stage +2 or higher when the load spike begins. Other models take a few extra minutes to switch to a stage where the battery can discharge—in this case, that extra time proves too long to suppress a spike in facility peak demand.

The December spike was a little different in that it was triggered by a drop in PV production, but the factors at work were similar. Once again, the tank temperature did not directly cause the spike, but it did contribute to differences in model conditions right before the spike. In this case, some models had empty batteries at the time of the drop in PV production—for these models, no electrical storage was available to discharge, and other resources are unable to respond. The result for these models was a spike in peak demand. Other models started the event with adequate storage and were thus able to discharge electricity to dampen the net load.

In this illustrative example, the maximum hot water tank temperature indirectly affected peak demand, but it was not a good predictor of whole-building performance. This suggests that other design parameters might have a stronger influence and could be tuned to improve the performance of the supervisory controller. The next section examines one of these parameters more closely.

3.6.3 How Do Net Load Control Targets Affect Whole-Building Performance?

Net load control targets are one example of the many adjustable inputs in the supervisory control algorithm. As discussed in Section 3.5, the team defined the net load control targets to include two values: a high target and a low target.

Together, the two target values form a net load control deadband. The supervisory controller will attempt to keep the facility net load within this range. If the net load moves outside of the deadband for a long enough time, the supervisory control stage will change, triggering additional responses from device-level controllers.

In the following illustrative example, the team compared multiple scenarios with different net load control targets. For simplicity, the net load control deadbands were defined to be symmetrical about 0 kW—that is, if the top of the deadband is x, then the bottom of the deadband is $-x$. In practice, however, the two targets could be decoupled.

To further simplify this example, the team also defined the following parameters to be a function of the net load control targets. In practice, these parameters could also be decoupled:

- The net load trigger for jumping to the highest stage is equal to $x + 2$ kW, where x is the high net load target.
- The net load trigger for jumping to the lowest stage is equal to $-x - 2$ kW, where $-x$ is the low net load target.
- If the net load is low, and if the battery is allowed to charge, the battery controller will attempt to bring the net load up to a value of $-x + y$, where $-x$ is the low net load target and y is 10% of the deadband range.
- If the net load is high, and if the battery is allowed to discharge, the battery controller will attempt to bring the net load down to a value of $x - y$, where x is the high net load target and y is 10% of the deadband range.

For this example, the battery is allowed to discharge down to a SOC of 20%. The team fixed the battery capacity at 420 kilowatt-hours (kWh). This is roughly 10 times the capacity of a research battery that was previously installed at the VTIF. (For a separate example that explores the effect of varying battery capacity, see Section 3.6.4.)

Keeping all other design parameters constant, the team modeled 20 alternative cases with net load targets ranging from ±1 kW to ±20 kW. The team evaluated the modeled results in terms of the following performance objectives:

- Reduce peak demand
- Reduce peak export rate
- Reduce net load range
- Reduce total facility energy use
- Reduce battery losses
- Reduce imported energy
- Reduce battery cycling as measured by battery roundtrip throughput
- Ensure the ratio of exported energy to imported energy is greater than one
- Increase the ratio of PV generation used on site to PV generation exported.

A comprehensive multiparameter optimization of the supervisory controller design was outside the scope of this study, but such an analysis could be conducted with EnergyPlus and OpenStudio tools. (For further discussion, see Section 5.2.) The following figures demonstrate the use of simulation results in an analogous but more simplified set of comparisons.

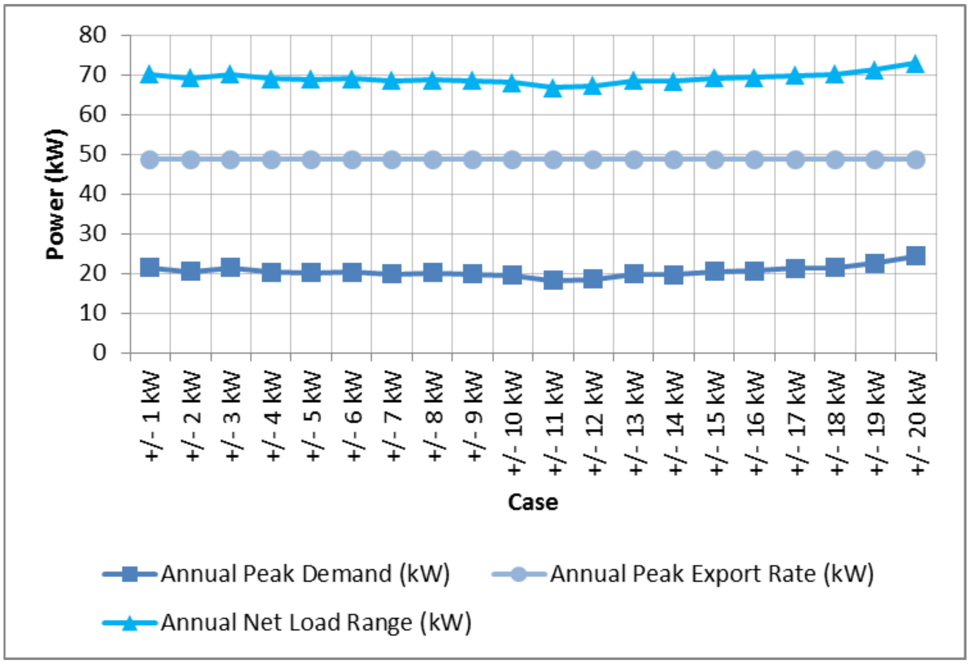

Figure 12. Annual peak demand, peak export rate, and net load range for the alternative models

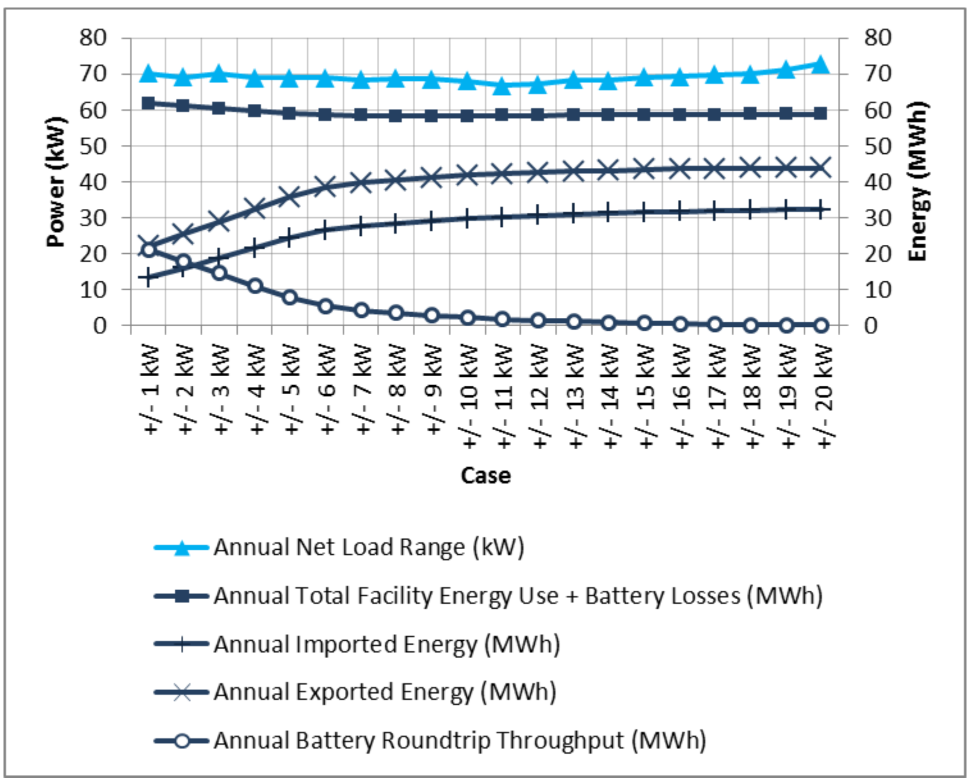

Figure 13. Annual net load and energy metrics for the alternative models

Figure 12 illustrates how the selected annual metrics vary as the net load control deadband is widened from a low of ±1 kW to a high of ±20 kW. This plot indicates that, within this range of cases, peak demand can be lowered most effectively by selecting net load targets of -11 kW and 11 kW. Peak export rate was relatively consistent, indicating that other design parameters should be adjusted if this metric is a priority for reduction. The net load range is a function of both peak demand and peak export rate, so it also has a minimum value when the net load targets are ±11 kW.

Figure 13 shows that there is a partial tradeoff between reducing total energy use (total facility energy use plus battery losses) and reducing imported energy. Total energy use decreases and imported energy increases as the net load deadband is relaxed from ±1 kW to about ±8 kW.

The figure also illustrates a tradeoff between battery roundtrip throughput and imported energy. Throughput was selected as a simplified proxy for battery cycling because the battery was modeled in a generic fashion without specifying a particular type of battery product. (This approach neglects depth of discharge, temperature, and other effects that could be modeled if more battery characteristics are specified.) If two models achieve the same whole-building performance but differ in battery cycling, the model with less cycling would be preferable in that it would potentially allow the battery to last longer. However, the figure shows that imported energy increases as battery roundtrip throughput decreases. Imported energy is informative as a proxy for greenhouse gas emissions.

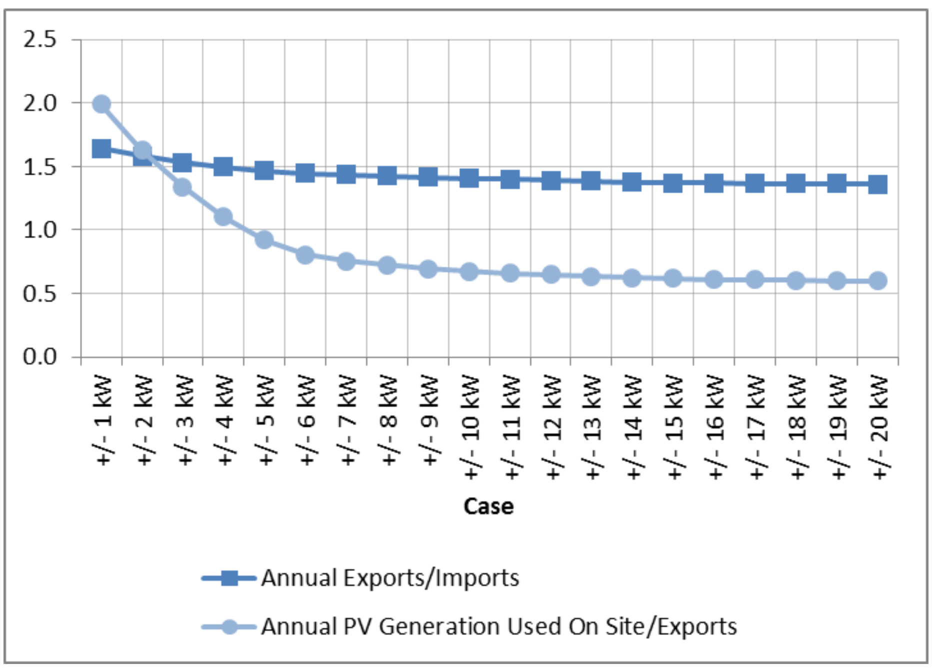

Figure 14. Annual export ratio metrics for the alternative models

Figure 14 shows two additional sustainability metrics that vary with the size of the net load deadband. One implication is that, while experimenting with different control strategies, it is important to check that the selected strategy does not compromise the ability of the building to meet the ZEB definition. (In this figure, each case does meet the ZEB requirement, which is that the ratio of exported energy to imported energy must be greater than one.)

The second metric in this figure is the ratio of PV generation used on site to the amount of PV generation that is exported. A higher ratio would be considered more favorable to a building owner, but this value also decreases as the net load deadband is relaxed.

3.6.4 How Do Battery Control Choices and Battery Capacity Affect Whole-Building Performance?

Another way to vary the supervisory control algorithm is to change the way that the battery is dispatched. In this example, the team considered four battery control cases, including two baseline models and two alternative models:

- ***Baseline A:*** A first-generation ZEB with no energy storage.
- ***Baseline B:*** A ZEB with batteries but no other load-shaping resources or supervisory control.
- ***Alternative A:*** A next-generation ZEB with a moderately constrained battery.
    - When evaluating whether to change the supervisory control stage at the end of each time step, the supervisory controller calculates the 5-minute rolling average of the net load and compares it to the net load deadband.
    - The battery will attempt to make up any difference between the instantaneous (1-minute) net load and the net load deadband. However, the battery is only allowed to charge if the supervisory control stage is -2 or lower, and it is only allowed to discharge if the stage is +2 or higher.
- ***Alternative B***: A next-generation ZEB with a minimally constrained battery.
    - When evaluating whether to change the supervisory control stage at the end of each time step, the supervisory controller calculates the 5-minute rolling average of the "non-battery net load" and compares it to the net load deadband. The "non-battery net load" equals the building load minus on-site PV generation. It excludes battery actions.
    - The battery will attempt to make up any difference between the instantaneous (1-minute) net load and the net load deadband. It can charge or discharge in any supervisory control stage other than Stage 0.

In both alternative models, the supervisory controller attempts to leverage lighting and TES responses before leveraging the battery. Alternative A does this by limiting battery actions to fewer stages—this is the approach originally described in Table 2 of Section 3.5. Alternative B disregards battery contributions when evaluating the deviation of the current net load from the net load deadband—this approach is summarized in Table 7.

Table 7. Summary of Supervisory Control Stages and Device-Level Responses for Alternative B

<table>
<tr><th>Stage</th><th>Description</th><th>Battery</th><th>Cold TES</th><th>Hot TES</th><th>Lighting</th></tr>
<tr><td>+4</td><td rowspan="4">Load-reducing actions needed</td><td rowspan="4">After other devices have acted, charge or discharge as needed to maintain net load within the deadband</td><td rowspan="4">Raise the tank temperature trigger for charging the tank</td><td rowspan="4">Lower the tank temperature trigger for charging the tank</td><td rowspan="4">Reduce levels while maintaining required minimum level of service</td></tr>
<tr><td>+3</td></tr>
<tr><td>+2</td></tr>
<tr><td>+1</td></tr>
<tr><td>0</td><td>Default control actions apply</td><td>Idle</td><td>Default action</td><td>Default action</td><td rowspan="5">Default action</td></tr>
<tr><td>-1</td><td rowspan="4">Load-increasing actions needed</td><td rowspan="4">After other devices have acted, charge or discharge as needed to maintain net load within the deadband</td><td rowspan="4">Allow charging under more conditions</td><td rowspan="4">Allow charging under more conditions</td></tr>
<tr><td>-2</td></tr>
<tr><td>-3</td></tr>
<tr><td>-4</td></tr>
</table>

The main advantage of Alternative B is its faster response time. With Alternative B, the battery may briefly be directed to charge in positive stages (Stage +1 through +4) or discharge in negative stages (Stage -1 through -4) if there is a momentary difference in sign between the instantaneous 1-minute net load and the 5-minute rolling average net load. This situation can occur if the building experiences a spike in net load like the one depicted in Figure 11.

Using the same battery capacity and net load deadbands from example in Section 3.6.3, the team compared the four battery control cases using various metrics. The results are summarized in Figure 15 through Figure 22.

Figure 15 shows that Alternative B outperformed the other models in terms of peak demand reduction. The figure also shows that the optimal net load deadband was different for each model—in the case of Alternative B, the optimal net load deadband was ±13 kW.

The peak export rate was relatively constant for the included cases. As a result, the net load range in Figure 16 (which is a function of peak demand and peak export rate) follows the same general trend as Figure 15.

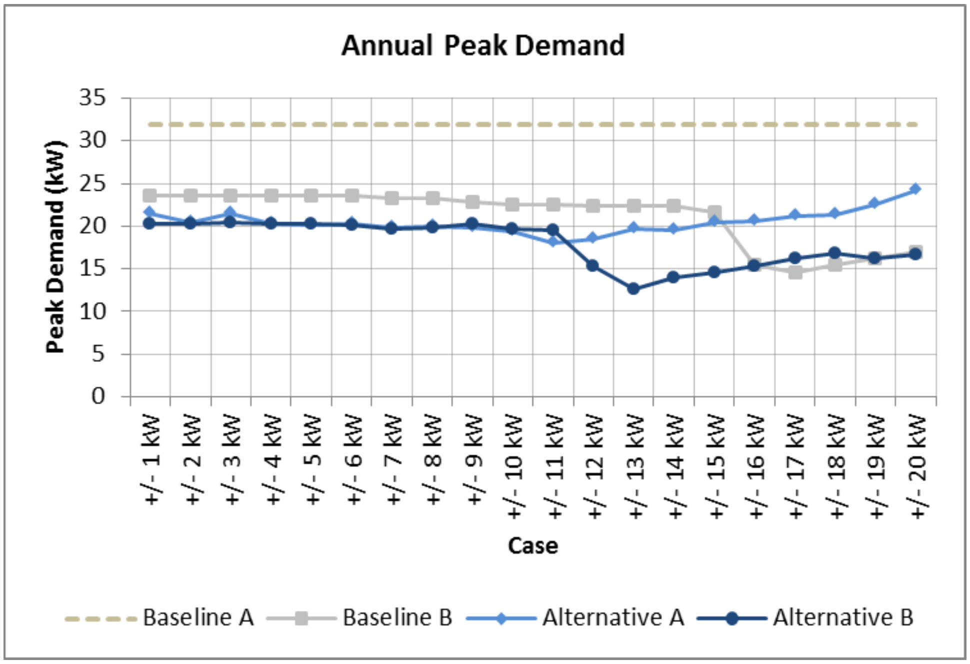

Figure 15. Peak demand for baseline and alternative models

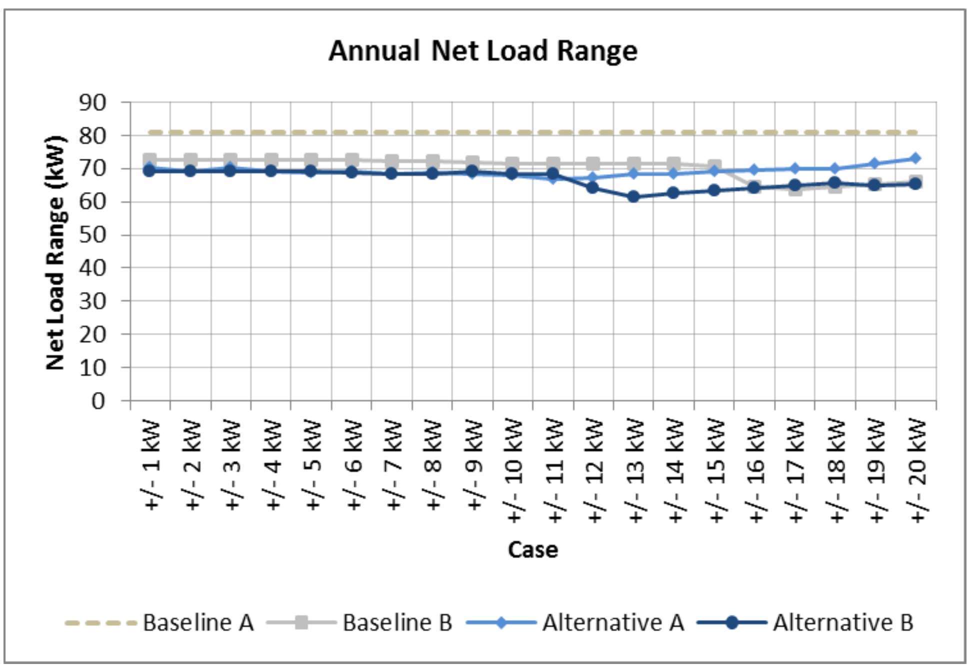

Figure 16. Net load range for baseline and alternative models

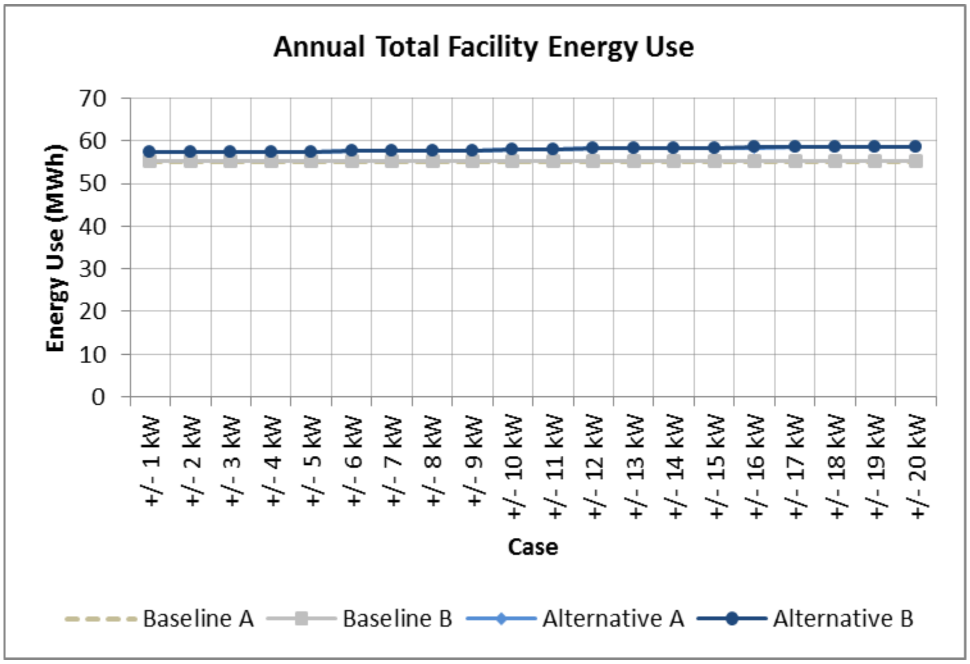

Figure 17. Total facility energy use for baseline and alternative models

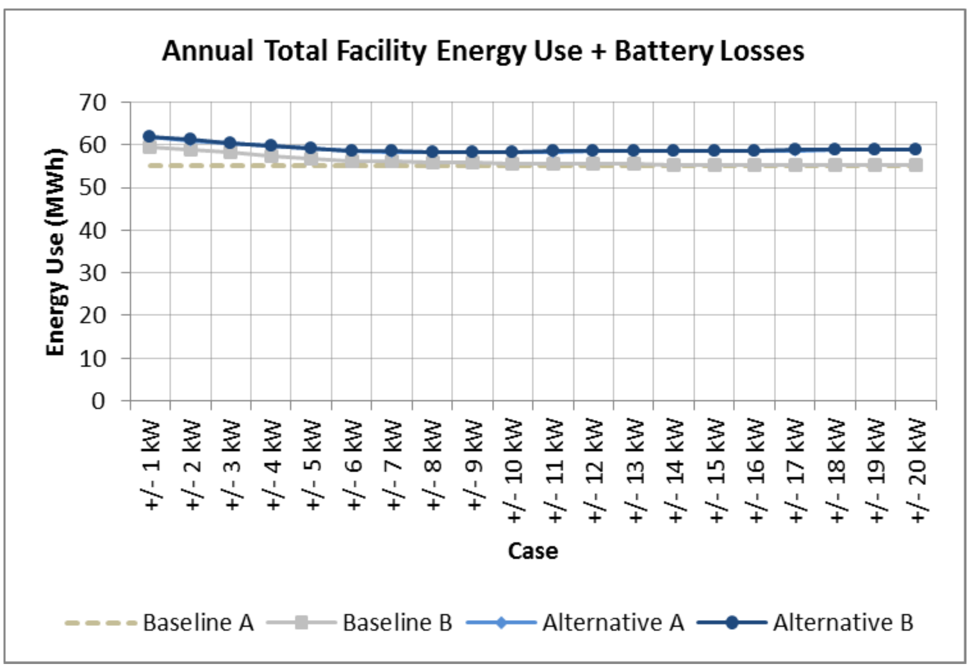

Figure 18. Total facility energy use plus battery losses for baseline and alternative models

Figure 17 and Figure 18 show that both alternative models use more energy than the baseline models. This indicates that other design parameters would have to be adjusted to further improve the alternative models with respect to this efficiency metric. Figure 18 also shows that there is an energy penalty associated with narrow net load deadbands.

Figure 19 shows that both alternative models import less energy than Baseline B when the net load deadband is tight, but the alternative models import more energy when the net load deadband is greater than roughly ±6 kW.

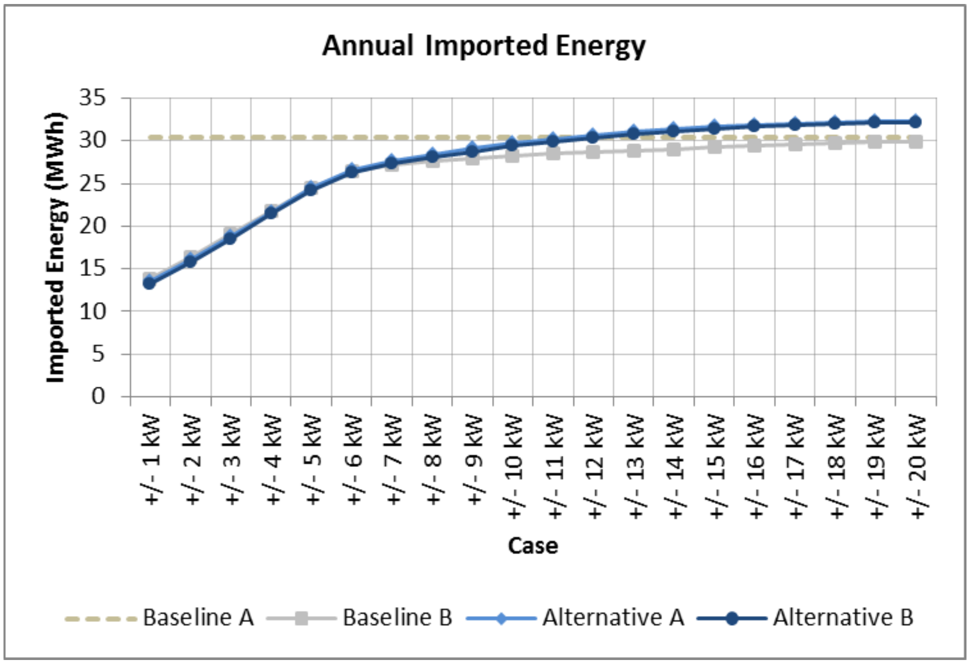

Figure 19. Imported energy for baseline and alternative models

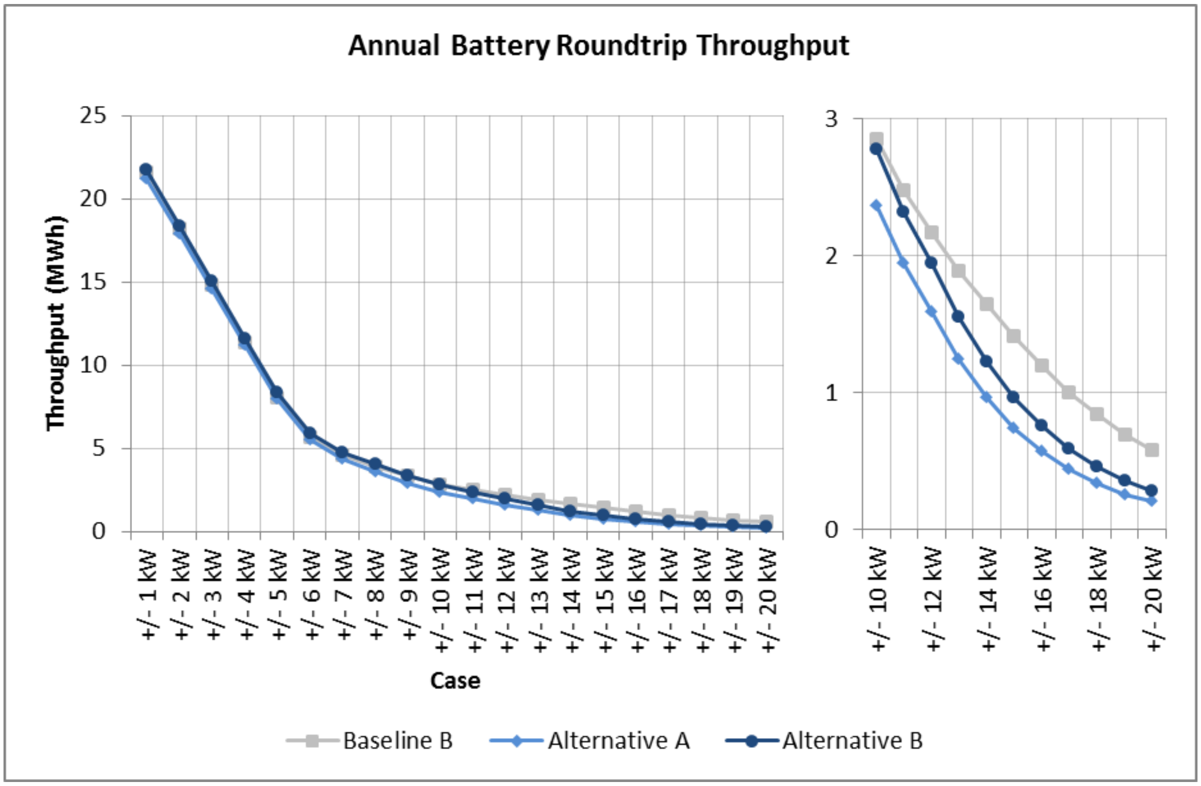

Figure 20. Battery roundtrip throughput for baseline and alternative models

Figure 20 indicates that the four battery control approaches are similar in terms of battery roundtrip throughput when the net load deadband is narrow. Both alternatives outperform Baseline B, however, when the net load deadband is ±10 kW or greater, as shown on the right side of the figure.

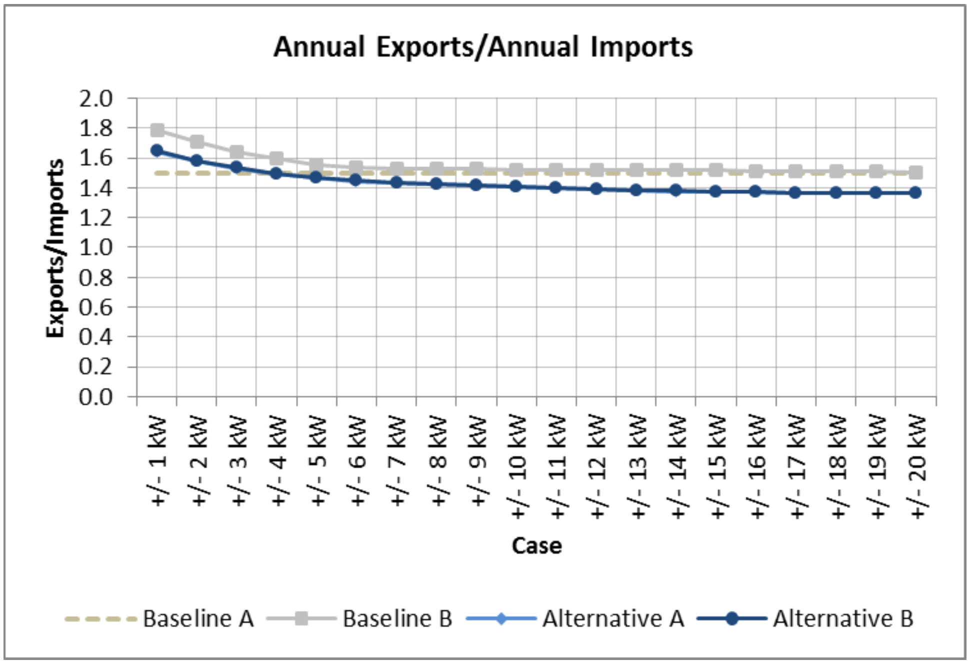

Figure 21. Ratio of exported energy to imported energy for baseline and alternative models

Figure 21 shows that all four battery control approaches were able to preserve ZEB status by maintaining a ratio of exported energy to imported energy of one or more. Baseline B outperformed the other models by this metric. The two alternative models outperformed Baseline A when the net load was tight, but their export-to-import ratios dropped as the net load deadband increased.

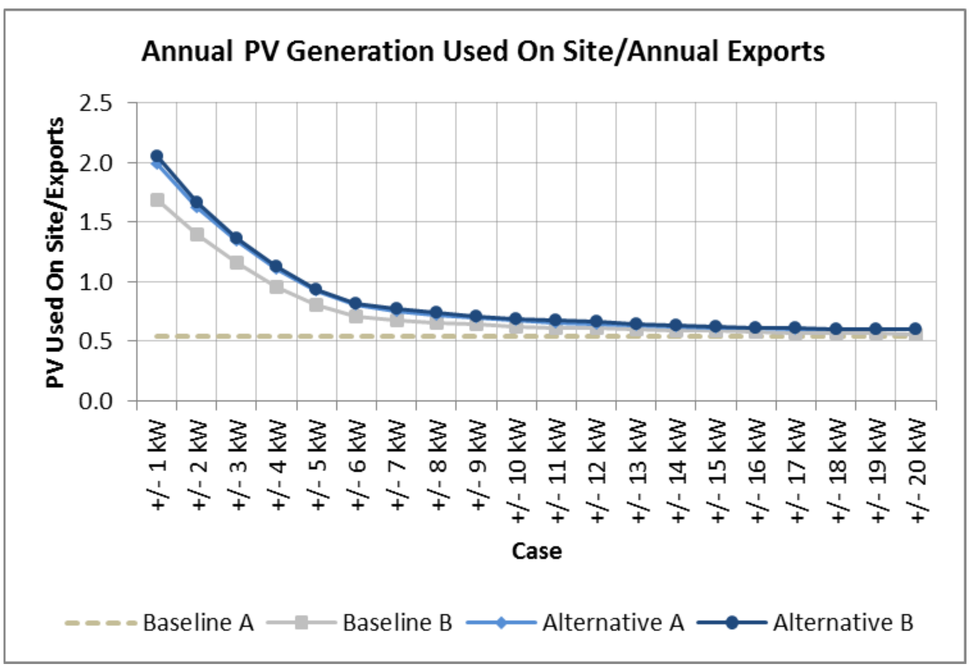

Figure 22. Ratio of PV generation used on site to PV exports for baseline and alternative models

Figure 22 shows that the alternative models had similar ratios of PV generation used on site to PV exported. Both alternative models outperformed the baseline models by this metric.

Overall, no single model outperformed the others by all metrics. In real-world applications, site-specific priorities will determine which metrics and performance objectives are most important. In an iterative design process, these priorities will inform which control concepts are the best candidates for further exploration and variation.

To illustrate an extension of this example, suppose that reducing peak demand is the top priority of the objectives reviewed thus far. Referring back to Figure 15, the optimal net load deadbands for Baseline B, Alternative A, and Alternative B appeared to be, respectively, ±17 kW, ±11 kW, and ±13 kW. If these deadbands are acceptable, it may be informative to examine how sensitive each model's peak demand reductions are to changes in other design parameters, such as battery capacity. A battery capacity comparison is illustrated in Figure 23.

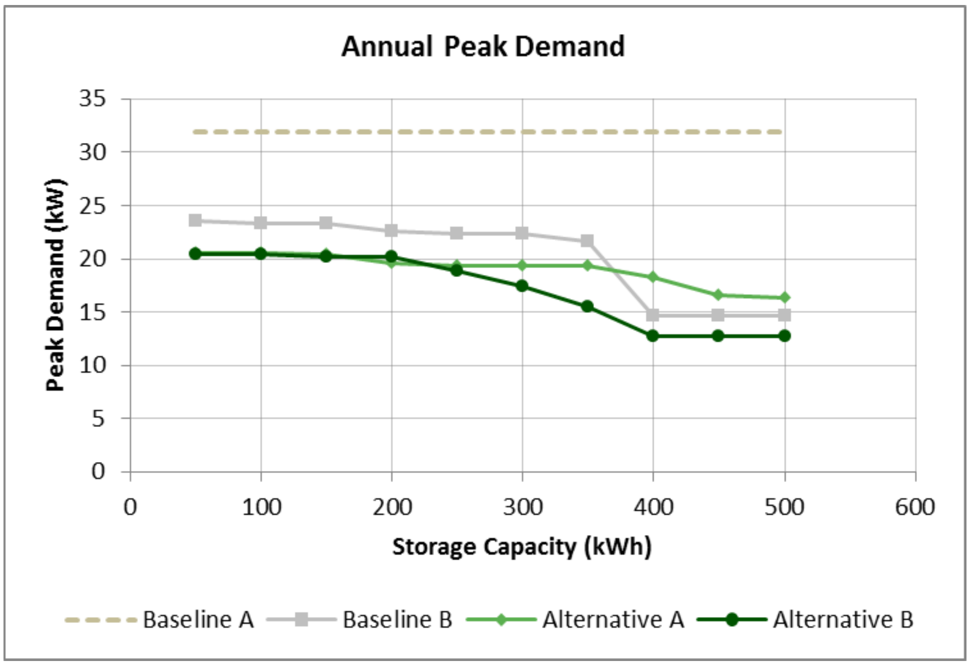

Figure 23. Impact of battery storage capacity on peak demand for select models

Figure 23 shows that within the included range of battery capacities, Alternative B is capable of the greatest peak demand reduction relative to Baseline A. It also suggests that increasing storage capacity between 400 and 500 kWh will not benefit peak demand reduction, though it could be considered for other metrics.

For each of the metrics discussed in this section, performance may be improved by adjusting additional design parameters. This can be done manually for a limited set of scenarios, but if the designer is willing to consider more variations to combinations of design parameters, a multiparameter simulation-based analysis facilitated by a tool such as OpenStudio can potentially provide a more efficient, comprehensive, and effective workflow.

4 Characterizing and Synthesizing Load Variability at the Minute Scale

To manage and shape the load profile of a ZEB or other building with on-site renewable generation, a supervisory controller must be able to account for variability in loads, storage, and generation systems at the 1-minute timescale. Therefore, obtaining a realistic representation of short-term load variability is a critical consideration for model-based design and evaluation of supervisory control strategies.

This section presents a preliminary exploration of alternative methods for characterizing and synthesizing load variability in buildings for which 1-minute-interval data are unavailable. With further development, such methods could help researchers and technology developers to address load profile information gaps for more buildings. This would allow them to extend the supervisory controller modeling strategies in Section 3 to more ZEB projects and other sites that integrate on-site generation, energy storage, and flexible building loads.

4.1 Approaches to Consider

For the VTIF case study presented in Section 3, the team was able to model 1-minute-scale variability in HVAC, lighting, and PV systems with a combination of whole-building energy simulation tools and high-resolution weather data. Additionally, the team was able to approximate remaining loads with 1-minute-interval data from a high-resolution VTIF meter and incorporate the data as an input into the models. Unfortunately, access to high-quality, high-resolution data is the exception, not the norm.

Few resources exist to assist a typical energy modeler with accurate representations of minute-scale load variability in an energy model. Published data sets, such as those available on OpenEI, typically supply demand data at an hourly or 15-minute resolution (OpenEI 2016). Utility interval data available to customers—e.g., from smart meters—also typically use a 15-minute reporting interval to match the utility's contractual peak demand billing interval.

To reduce peak demand charges, supervisory control algorithms must respond *within* the 15-minute demand period. Therefore, hourly and 15-minute-interval load data are of limited use for the purpose of modeling, developing, and testing supervisory control algorithms for load shaping.

In rare cases, an energy modeler may have access to high-resolution, 1-minute-scale demand data from an existing building. If so, then the energy modeler may consider leveraging the existing building's observed load data and incorporating similar variability into the energy model of the new building. The modeler can consider two approaches:

1. The modeler may *extract* a time-series representation of short-term variability from the observed data using filtering techniques, rescale it, and subsequently apply it to the energy model as a time series input.

2. Using the characteristics of the variability in observed data, the modeler may *synthesize* a similar time series of randomized artificial variability to use in the energy model.

Both approaches are viable but limited by the quality and applicability of the observed data set. The following sections present a preliminary exploration of these concepts. Future work could include a broad, systematic characterization of short-term load variability in real buildings that can provide high-quality input data for either approach.

4.2 Characteristics of Short-Term Load Variability in Buildings

Broadly speaking, short-term variability in building load may be divided into variability from HVAC sources and non-HVAC sources. Variability related to HVAC systems includes cycling of fans and compressors in response to heating, cooling, and ventilation demand. Sudden changes in HVAC load can also be triggered by step changes in HVAC system set points, such as when transitioning from an unoccupied schedule to an occupied schedule. Because HVAC energy consumption is demand-driven, HVAC-related variability is best captured in the output of a whole-building energy model. If appropriate HVAC component models are used, these phenomena are readily captured by an energy model running at a sufficiently small time resolution, as discussed in Section 3.4.2.

Non-HVAC load variability is attributable to random or semi-random variation in a building's lighting and miscellaneous equipment loads due to occupant behavior, independent equipment operation, or the response of building systems to varying environmental conditions. If a particular source of variability can be captured as a function of a known physical input, it is often preferable to explicitly model the relationship. For example, if lighting demand in a daylighted building is a quantifiable function of solar irradiance, then high-resolution irradiance data may be used as an input to generate a variable lighting load. More often, however, the driving mechanisms for non-HVAC loads are unknown, and empirical methods must be used to represent them.

Even given random inputs, such as occupant behavior, aggregate building loads exhibit clear patterns. As individual loads turn on and off, the aggregate load rises and falls in a sequence of clearly identifiable pulses of varying magnitude and duration. Typical features include a clearly identifiable base load trend, large load spikes, steep transitions, and short-duration periods during which demand remains relatively level. These features are often obscured when a lower-resolution reporting interval is used. Figure 24, which plots the VTIF electrical demand for June 19, 2014, illustrates the loss of detail that occurs when the data are displayed at a 15-minute time resolution instead of a 1-minute resolution.

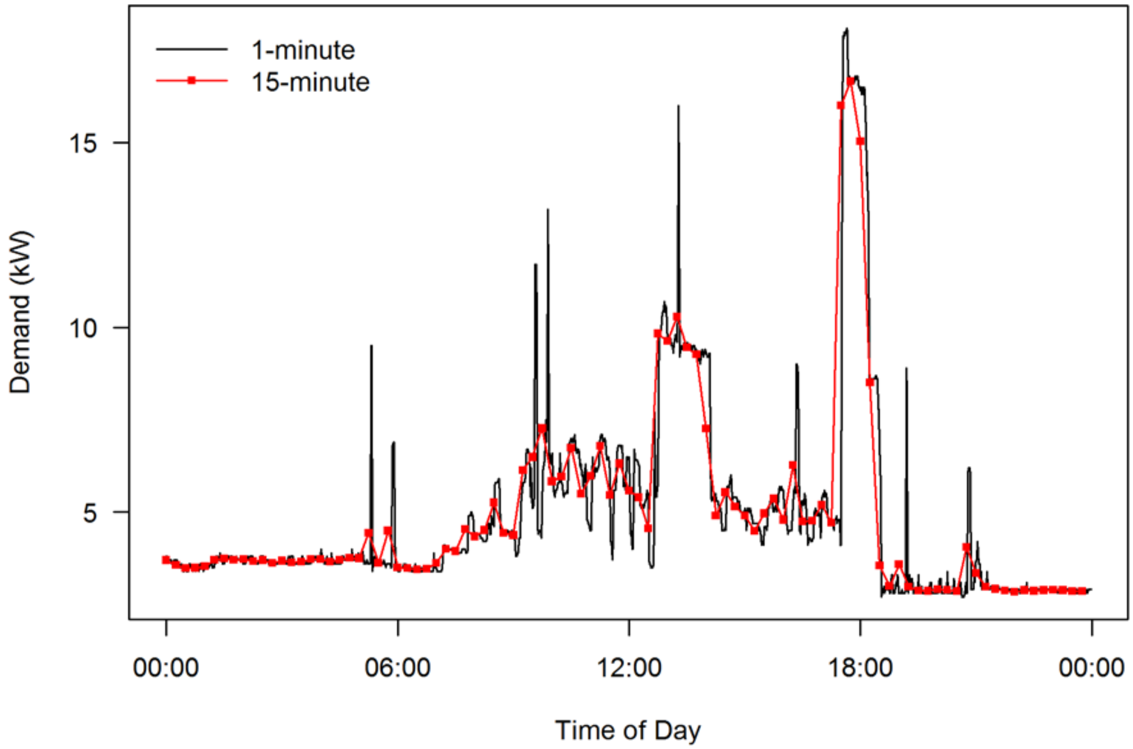

Figure 24. Vehicle Testing and Integration Facility electrical demand for June 19, 2014

The signatures of individual loads may be regularly or randomly timed, and different loads cycle at different intervals. Examining the power spectrum of the load in the frequency domain can identify key characteristics of the load shape. For instance, Figure 25 shows the power spectrum of the VTIF load for June 2014.

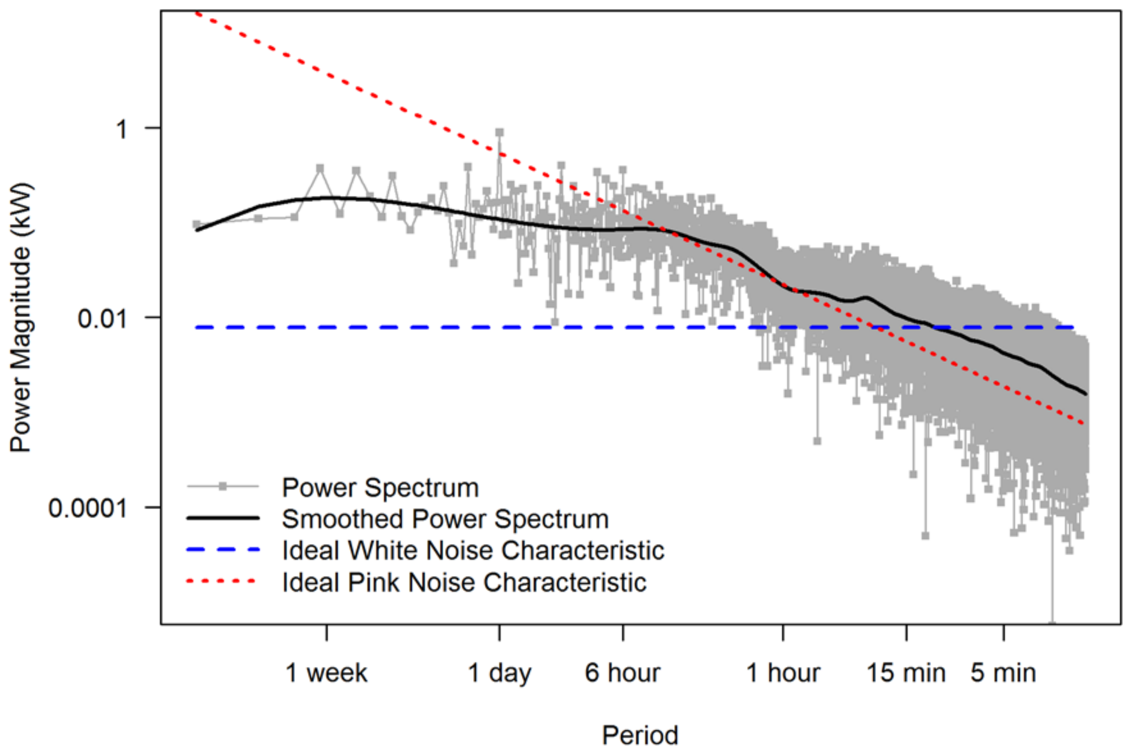

Figure 25. Power spectrum of Vehicle Testing and Integration Facility electrical demand for June 2014

At low frequencies, the VTIF load variability has a relatively flat power spectrum reminiscent of white noise. At high frequencies, the power content is inversely proportional to frequency—similar to pink noise but with a slightly slower decay characteristic.

Frequency-domain analysis can also expose regular timing patterns, as illustrated by the power spectrum of RSF I demand data for January 2014 (Figure 26). The sharp spike at 1.67 millihertz (mHz) and its harmonics indicates a cyclic load with a regular pattern of exactly 10 minutes. The smaller but wider hump at approximately 0.4 mHz indicates a significant cyclic load with a variable period of approximately 40 minutes. Examination of time series data verifies the presence of these two cyclic load patterns (Figure 27).

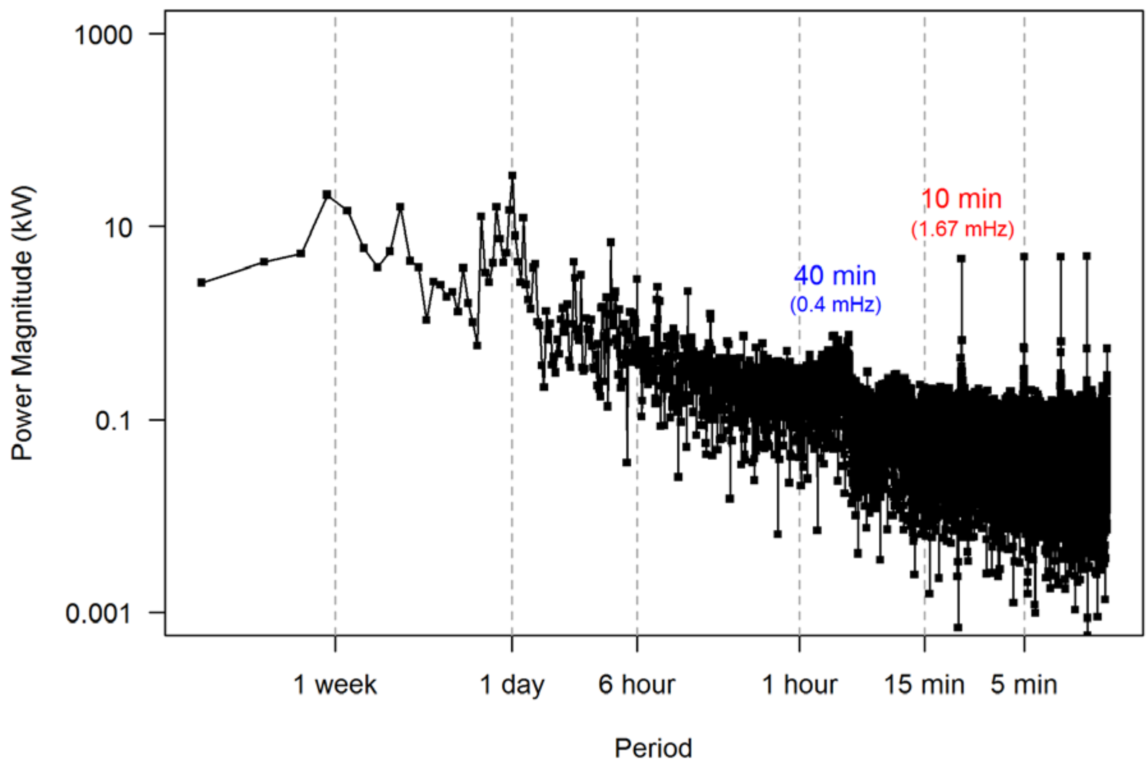

Figure 26. Power spectrum of Research Support Facility I electrical demand for January 2014

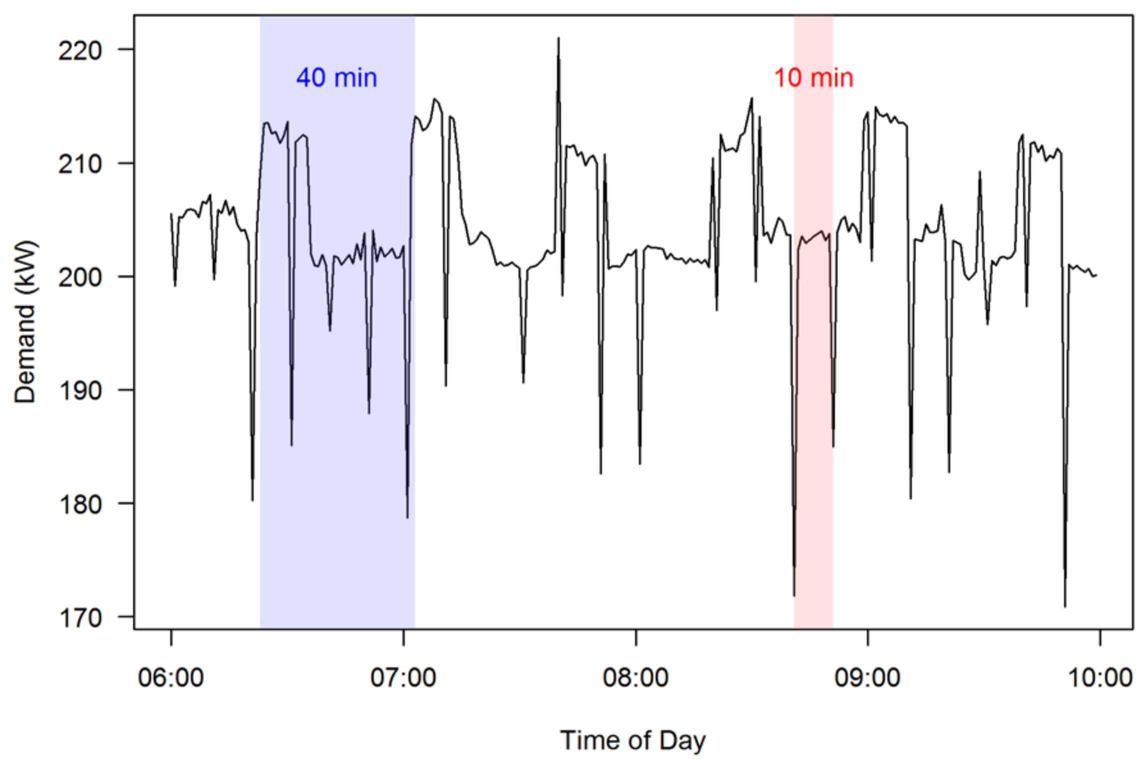

Figure 27. Research Support Facility I demand data for January 11, 2014, illustrating periodic load patterns revealed by frequency-domain variability analysis

4.3 Variability Extraction

One approach to modeling short-term load variability is to extract a time series of observed load variability, or variability signature, from a measured data set. The extracted variability signature may then be superimposed on the output of an energy model to create a realistic load profile. The variability signature should not alter the modeled building's mean load or daily demand profile. This goal is met if the variability signature:

- Has zero mean
- Has a power spectrum with negligible low-frequency content.

Passing a measured load through a high-pass filter achieves these goals. The cutoff frequency for the filter should correspond to the smallest time resolution of the underlying energy model. For instance, if the underlying model uses an hourly schedule, then the variability signature should contain only subhourly features. An ideal filter will accurately preserve the shape and magnitude of short-term variability while completely eliminating long-term variation. Practical filters, however, always have some pass-band distortion and stop-band leakage.

The team encountered a tradeoff between the quality of the filtered signal and the level of low-frequency rejection obtained by filtering. The type of digital filter applied (one-sided or two-sided, linear or nonlinear) also strongly influences the shape of the filtered signal. For example, Figure 28 compares four variability signatures obtained by filtering the VTIF electrical demand data for June 19, 2014.

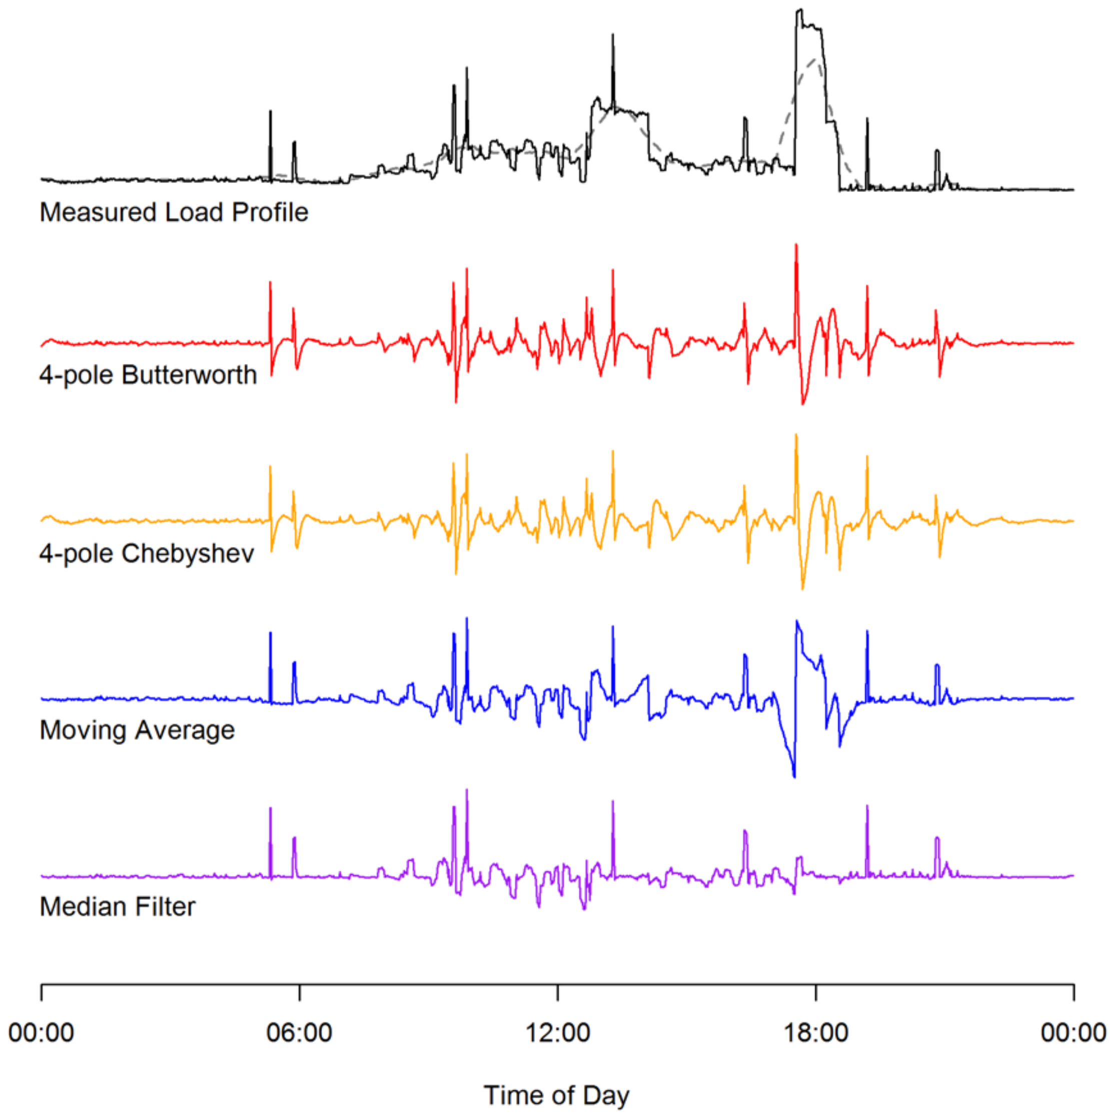

Figure 28. Comparison of variability signatures obtained using several high-pass filtering techniques

The Chebyshev and Butterworth variability signatures were obtained by passing the measured data through one-sided, high-pass, autoregressive moving average filters of type Chebyshev and Butterworth, respectively. (These are two common filter topologies used in digital signal processing.) The moving average and median filter variability signatures were obtained by subtracting a smoothed time series (using a centered 1-hour moving average and a centered 1-hour median filter, respectively) from the original time series.

The Chebyshev and Butterworth filters capture the leading edge of load shifts effectively but yield undesirable overshoot and ringing at the trailing edges of events, producing considerable distortion in load shape compared to the original signal. The moving average filter yields less distortion, but still produces undesirable ramps instead of sharp transitions at either edge of longer events. Qualitatively, the variability signature produced by the median filter is the most visually faithful to the original observed variability.[6] In contrast, the Chebyshev and Butterworth filters provide the best low-frequency rejection and the median filter provides the worst (Figure 29). The variability signature from the median filter achieves less than 10:1 low-frequency signal attenuation.

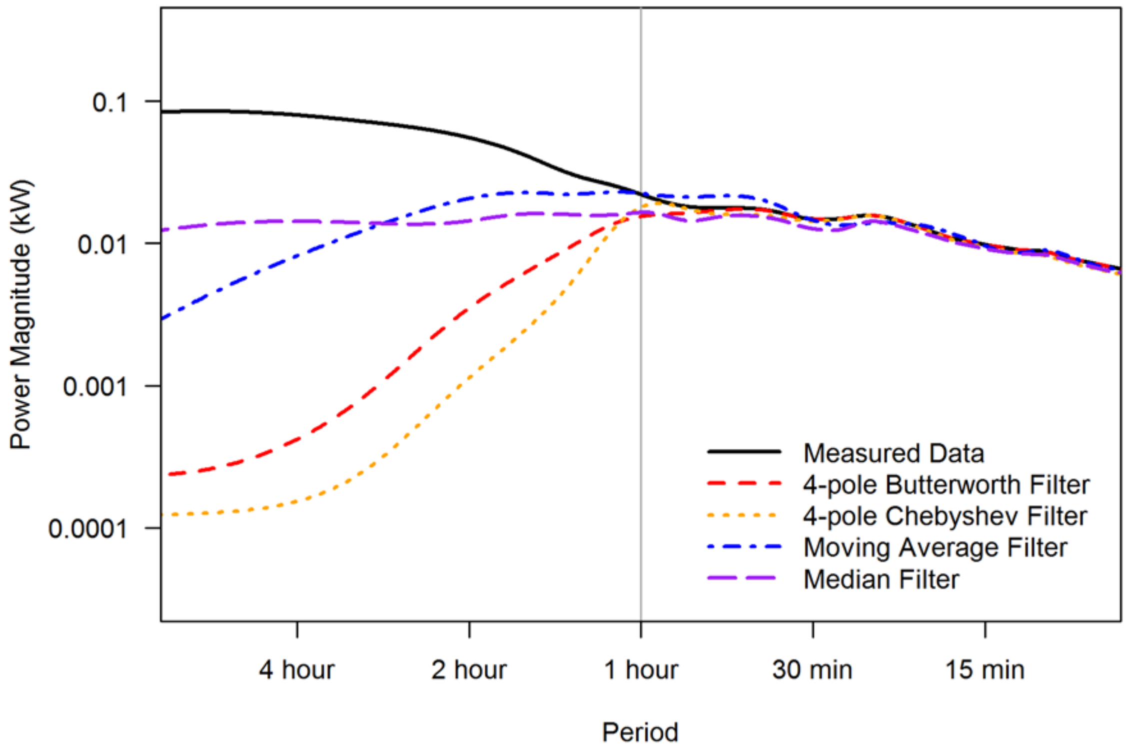

Figure 29. Comparison of power spectra and low-frequency cutoff of variability signatures obtained using several high-pass filtering techniques

Regardless of the filtering method selected, the variability signature obtained from variability extraction must be rescaled prior to applying it to a model of a different building. One potential strategy is to multiply the signature by the ratio of the total energy consumption of the modeled building to the total energy consumption of the building that provided the variability signature. The rescaled signature may then be added to the internal electrical load of the modeled building. Future research is needed to determine how similar two buildings would need to be for this strategy to be effective.

4.4 Variability Synthesis

The goal of variability synthesis is to create artificial short-term load variability with realistic features and a power spectrum similar to that which occurs in real buildings. The main advantage of variability synthesis over extraction is the ability to rapidly generate many variability signatures for stochastic simulation. A key disadvantage is that creating a synthetic variability

---

[6] The exception is during periods of monotonically changing load, during which small individual step changes may be lost due to the nonlinear nature of the median filter. An example of this occurs at approximately 18:30 in Figure 23.

signature with known periodic loads is more difficult than simply using an observed signature that contains the periodic loads of interest.

Imposition of independently sampled variation ("white" noise) does not achieve the goal of creating a variability signature with a realistic power spectrum, because white noise has a level power spectrum. "Red" noise (the integration of "white" noise) produces the desirable high-frequency roll-off, but the resulting time series is not stationary, contains significant low-frequency content, and lacks the sharp transitions and level periods typical of real load variability. Therefore, more sophisticated variability synthesis techniques are required.

The team explored two promising techniques for variability synthesis:

1. Frequency-domain variability synthesis using a reverse discrete Fourier transform (DFT)

2. Time-domain variability synthesis using recursive conditional sampling.

With either technique, the modeler must have a well-characterized load profile or variability signature to use as a template. Thus, variability synthesis requires measured variability to use as a template and good variability analysis or extraction techniques.

4.4.1 Frequency Domain Synthesis

Frequency-domain variability synthesis requires a frequency-domain model of the template variability signature—that is, a model for the high-frequency portion of the signal's power spectrum. To synthesize variability, the team added randomly sampled noise to the fitted model to create a randomized power spectrum that closely resembles the power spectrum of the original signal. Typically, the magnitude of the power spectrum can be modeled as a function of frequency (potentially with a logarithmic transformation), while the phase angle is uniformly distributed within the range $(-\pi, \pi]$.

Because of the symmetry inherent in the DFT, generating N samples in the time domain requires a synthetic power spectrum vector containing only $N/2 - 1$ samples. A possible procedure is as follows:

1. Given a desired sampling frequency f_s, generate a uniformly spaced frequency vector with $N/2$ elements from f_s/N to $f_s/2$. (N must be an even number.) Drop the last element of this vector, which corresponds the Nyquist frequency, $f_s/2$. The resulting vector f should have exactly $N/2 - 1$ elements.

2. Let f_c denote the desired cutoff frequency for the variability signature. (For instance, $f_c = 1/3600$ if synthetic subhourly variability is desired.) For each frequency f in the resulting vector, create a randomly sampled magnitude and angle using the following heuristic:

   a. Sample the magnitude x at each frequency f according to:

$$x(f) = \begin{cases} 0, & f < f_c \\ \overline{x}(f) + X(\omega), & f \geq f_c \end{cases} \tag{1}$$

in which $\overline{x}(f)$ denotes the best-fit model of the (one-sided) power spectrum of the variability at high frequency and $X(\omega)$ denotes a random sample drawn from the probability distribution of the residuals of the model. The probability distribution used may be normal (Gaussian) or some other distribution, such as a non-parametric empirical distribution.

b. Sample the angle ϕ at each frequency f according to:

$$\phi(f) = \begin{cases} 0, & f < f_c \\ \Phi(\omega) \sim U(-\pi, \pi), & f \geq f_c \end{cases} \tag{2}$$

in which $\Phi(\omega) \sim U(-\pi, \pi)$ denotes a sample drawn from the uniform distribution from $-\pi$ to π.

3. From the sampled magnitude and angle vectors, construct the complex vector $z(f) = x(f)\angle\phi(f)$.

4. Compute z^*, the complex conjugate of z.

5. Construct the complex vector $\zeta = [0, z/2\ ,0, z'/2]$, in which z' contains the elements of z^* in reverse order. Vector ζ contains the correct symmetry required to achieve a real-valued inverse DFT. The division of z and z' by two preserves amplitudes within the inverse DFT input.

6. Take the inverse DFT of vector ζ.

The result of this process is a synthetic time-domain variability signature with exactly N elements and sampling frequency f_s. For the amplitude of the variability signature to match the amplitude of variability in the original waveform, N must be equal to the number of elements in the time-domain sample used to generate the DFT for the frequency-domain model. If not, then the synthetic waveform may be scaled as desired until the amplitudes match, or until the amplitude is appropriate for the modeled building to which the signature will be applied.

Variability signatures generated using this procedure contain no spurious low-frequency content, but may not exhibit the shelves and sharp transitions typical of real building loads. Neither will this process preserve any distinctly periodic patterns present in the original signal.

4.4.2 Recursive Conditional Sampling

Variability synthesis using recursive conditional sampling does not use the frequency-domain power spectrum but rather requires a filtered time-domain variability signature to emulate. The method develops an empirical probability distribution for the power level conditional on the value of the previous sample. Generating a synthetic variability signature requires recursively sampling from this conditional probability distribution, feeding each newly generated sample forward as input to generate the next sample.

A simple implementation of this procedure that uses binning and kernel density estimation is as follows:

1. Given a known variability signature v, create K nonoverlapping bins that fully span the values in v. The bins may be uniform or nonuniform; a possible rationale for using nonuniform bins is to divide the data into several distinct operating regions observed visually in the original data.

2. For each $k \in \{1, \ldots, K\}$, select all samples in v for which the *previous* sample is in bin k, then use the selected samples to generate empirical probability distribution D_k using kernel density estimation. D_k represents the probability distribution for the variability given that the previous sample was in bin k.

3. To generate a variability signature x with N elements, set sample $x_1 = 0$ or use a randomly generated value for x_1 that falls within the range of v. Set sample index $i = 1$.

4. Increment i.

5. Determine k such that x_{i-1} falls within bin k.

6. Set sample x_i equal to a random sample drawn from conditional distribution D_k.

7. Repeat steps 4–6 until $i = N$.

This process will yield a randomized variability signature with approximately the same range and probability distribution as the original data. In addition, the recursion tends to generate realistic shelves and sharp transitions that mimic those in the original variability signature. However, the process will not reproduce periodic patterns, nor is it guaranteed to yield a signature without a bias or other spurious low-frequency content.

4.4.3 Case Study

As a case study, the team applied both methods of synthetic variability generation to the VTIF electrical demand data for June 2014. For templates, the frequency-domain synthesis method used the Fourier transform of the original load data and the recursive conditional sampling method used a variability signature extracted using a 1-hour median filter.

To develop a model for frequency-domain synthesis, the team modeled the logarithm of the power magnitude at high frequency as a function of the logarithm of the frequency. The resulting best fit linear model (Figure 30) has the form

$$log(\text{Magnitude}) = -0.8575\, log(\text{Frequency}) - 10.43 \quad (3)$$

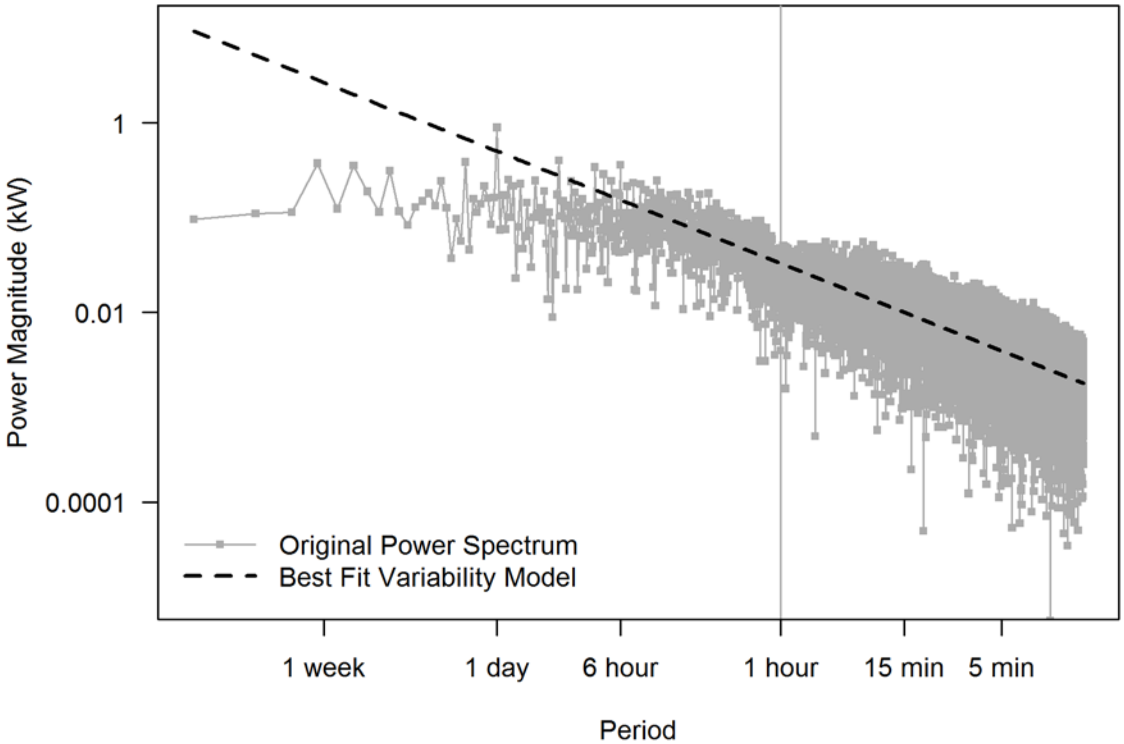

Figure 30. Best fit frequency-domain model for subhourly variability in Vehicle Testing and Integration Facility load for June 2014

The standard deviation of the residuals for this model is 0.6481. The residuals follow a slightly skewed distribution with a long lower tail (Figure 31).

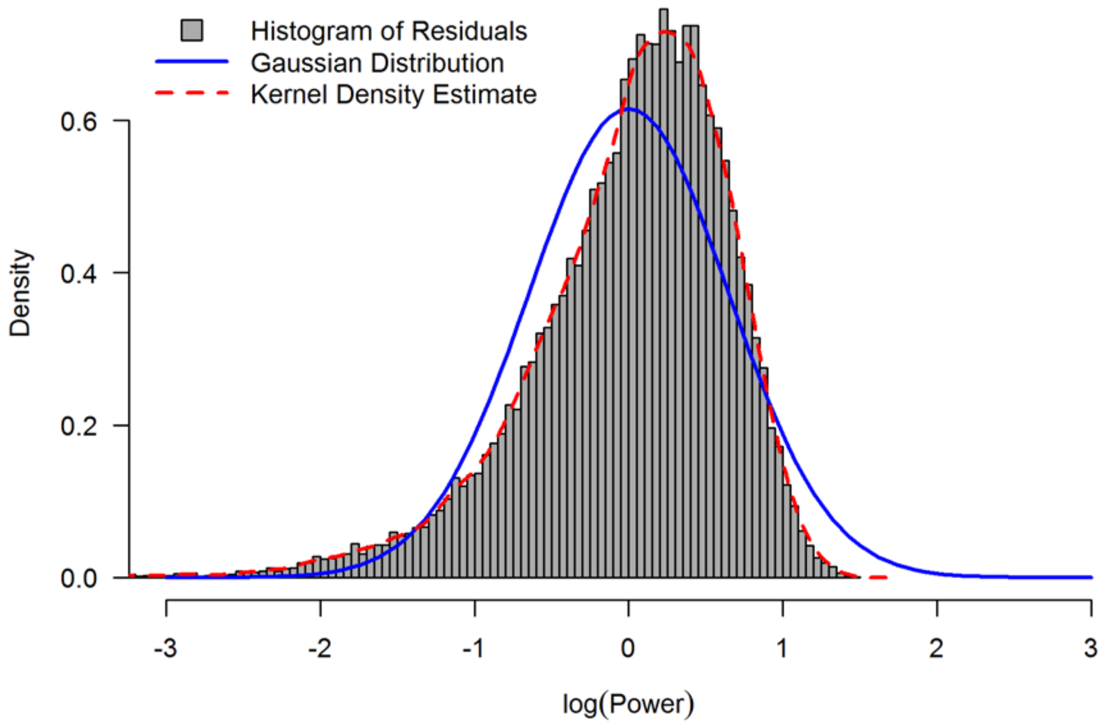

Figure 31. Histogram of residual errors for the best-fit model of subhourly variability shown in Figure 30

Following the procedure outlined in Section 4.4.1, the team then synthesized a similar variability spectrum by adding a randomly sampled level of noise to the modeled power magnitude at each frequency, with the sample drawn from a normal distribution with mean zero and standard deviation 0.6481, to match the observed distribution of the residuals.[7] (In a slight variation from (1), the team sampled and added this noise prior to reversing the logarithmic transformation.) Similarly, the team randomly sampled phase angles from the uniform spectrum $(-\pi, \pi]$ per (2). After constructing the randomized frequency-domain spectrum, the team applied an inverse DFT to obtain a synthetic time-domain variability signature.

To develop the conditional probability models for the recursive conditional sampling method, the team divided the samples in the template variability signature among 22 bins from -7 kW to +15 kW, with width 1 kW, spaced. the team then dropped each bin that contained less than 20 samples. For each of the remaining 11 bins, the team developed a corresponding empirical probability distribution function for the next sample as a function of the current sample using kernel density estimation. As an example, Figure 32 shows the empirical probability distribution function obtained for the $(0,1]$ kW bin.

[7] In this case, the team sampled from the normal (Gaussian) distribution, but the process is similar for other distributions, such as the kernel density estimate shown in Figure 26.

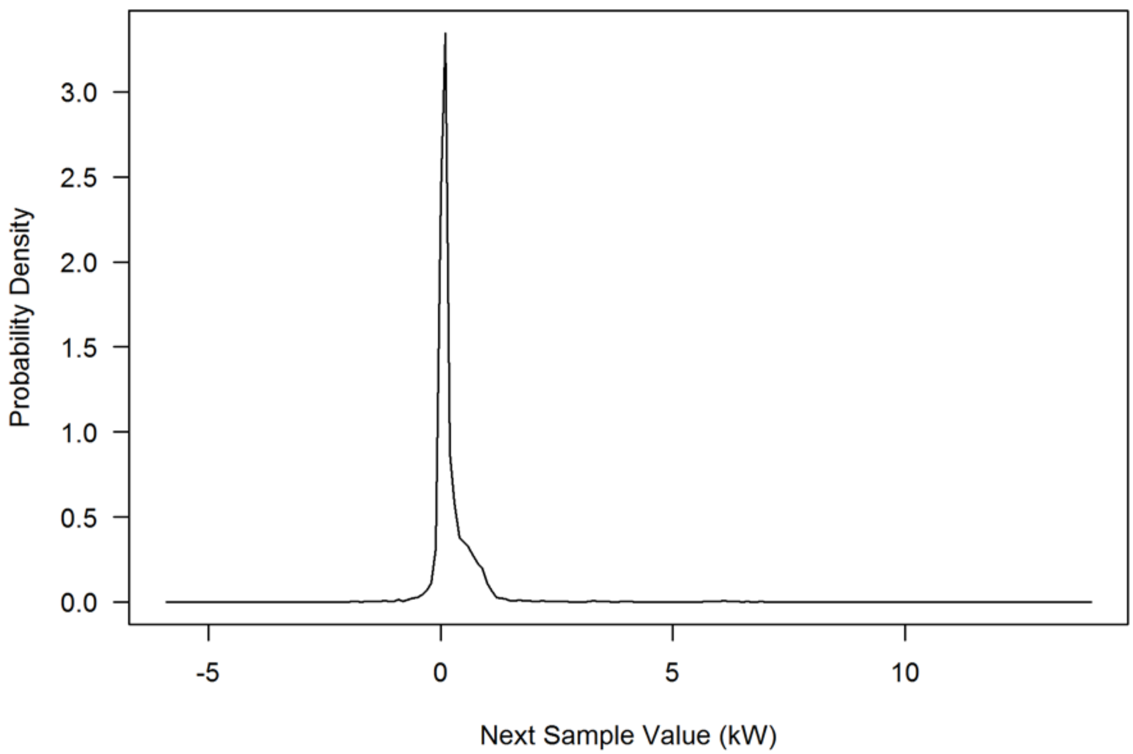

Figure 32. Empirical probability density function of the next sample of June 2014 Vehicle Testing and Integration Facility electrical demand given that the current sample falls in the (0,1] kW bin

Figure 32 illustrates how building load is most likely to remain unchanged, with only a small probability of changing by a large amount. Recursive sampling from these probability distributions therefore yields the prolonged shelves and rare but sharp transitions typical of real building loads.

The team then generated a random variability signature from this family of probability distributions using the recursive conditional sampling algorithm outlined in Section 4.4.2. Figure 33 compares the original variability signature to the two synthetic signatures for a sample weekday.

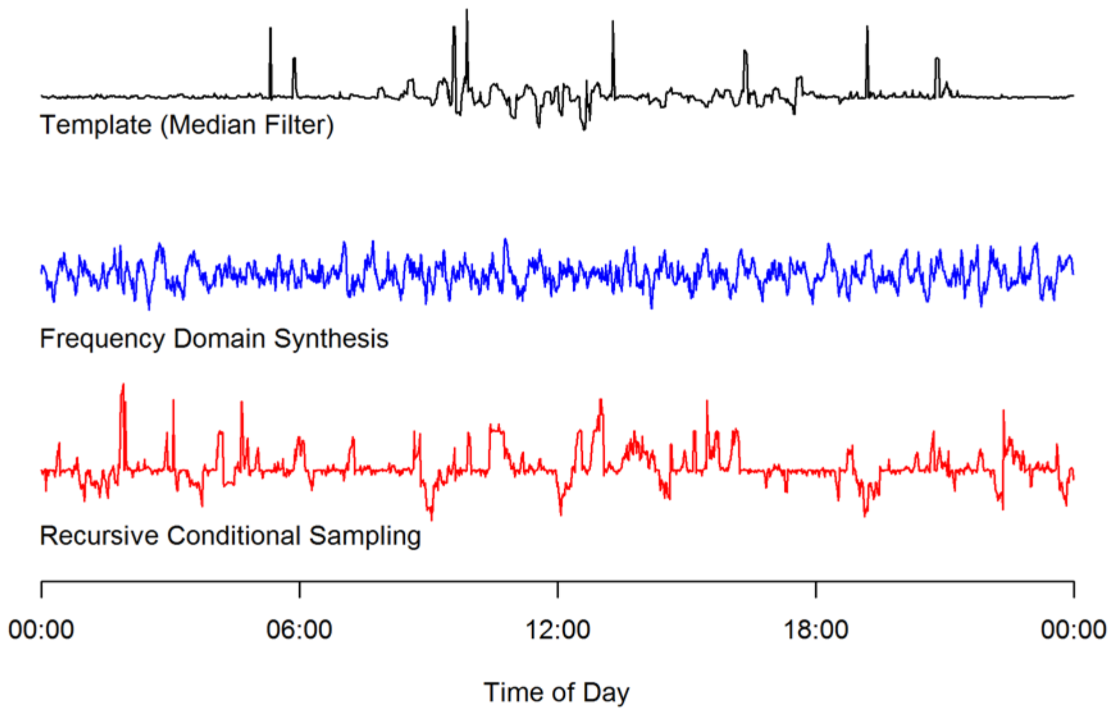

Figure 33. Comparison of synthetic variability signatures generated by two methods: frequency-domain synthesis and recursive conditional sampling

The recursive conditional sampling method generates a variability signature with the shelves, spikes, and ramps typical of real building loads. Therefore, the signature obtained from recursive conditional sampling is much more visually similar to the template signature than the output of the frequency-domain synthesis method. While this is a significant advantage of the recursive conditional sampling method, it is unclear whether visual similarity to real buildings is required for a variability signature to be sufficient for testing supervisory control algorithms.

The frequency domain synthesis method does maintain two advantages over recursive conditional sampling. First, it requires fewer input data: three values (two coefficients of a linear frequency-domain model and the standard deviation for the model residuals) are sufficient. Contrast this with the recursive conditional sampling method, which requires a family of empirical conditional probability distributions developed from a detailed template signature. Second, unlike recursive conditional sampling, frequency domain synthesis provides a variability signature with no low-frequency content (Figure 34). Therefore, a signature obtained from frequency domain synthesis will not unintentionally distort the daily load profile when overlaid on a lower resolution time series.

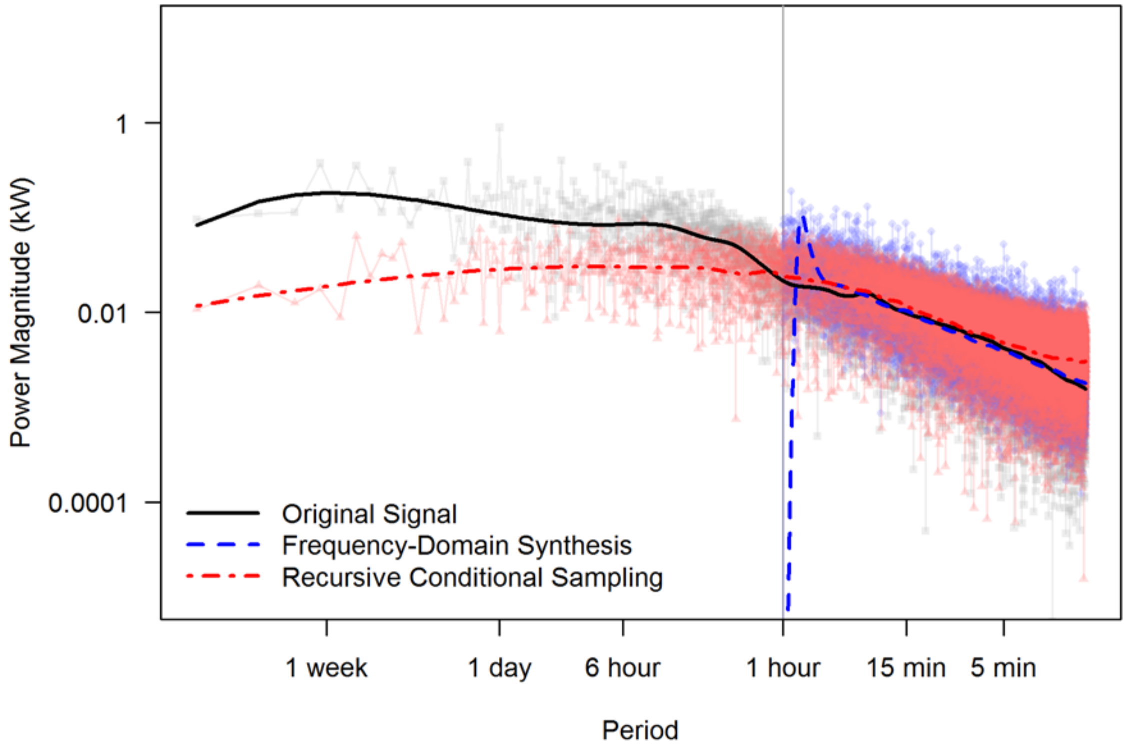

Figure 34. Comparison of power spectra of synthetic variability signatures generated by frequency-domain synthesis and recursive conditional sampling with the power spectrum of the original load data

Other points of comparison raise further research questions:

- The observed variability is higher during the daytime and tapers off at night. Do synthetic variability signatures need to duplicate this behavior?
- Are separate variability signatures needed for weekdays and weekends?
- Is it critical to synthesize a variability signature with peaks of the same magnitude as the observed signature?

These and other similar questions are suggested for future research in Section 5.3.

5 Conclusions

ZEB supervisory controller performance is sensitive to multiple design parameters. These can range from control choices (e.g., net load targets, battery constraints) to system characteristics (e.g., battery capacity). Furthermore, different control objectives can work in opposition to each other, such as when peak demand is reduced at the expense of an increase in imported energy and battery cycling. Developers of supervisory controls for next-generation ZEBs will therefore require tools that can navigate complex parameter spaces.

This project demonstrated that a simulation-based approach can help modelers draft, test, compare, and improve ZEB supervisory control strategies based on a variety of metrics. Through illustrative examples, the team also showed that carefully designed systems combining batteries, TES, flexible loads, and novel supervisory controls can sometimes outperform systems that rely solely on batteries for ZEB load profile management.

To conduct such investigations, it was important to use a whole-building energy modeling tool capable of capturing interactions between building loads, generation, and energy storage systems with 1-minute simulation time steps. The team also identified potential methods for characterizing and synthesizing load data at the 1-minute timescale for buildings lacking metered data at that resolution.

The following sections highlight additional findings from the real-world building integration effort, the demonstration of novel capabilities for simulation-based ZEB controller design, and the exploration of methods for characterizing and synthesizing minute-scale variability in building loads.

5.1 Integration Considerations

When selecting equipment and control strategies for next-generation ZEB systems, many decisions will be driven by project-specific priorities. Key considerations include:

- Will the control approach be purely heuristic, or will it include real-time optimization?
- Will there be a learning component to the control algorithm?
- Will there be a predictive component to the control system?
- Will the supervisory controller be part of or separate from the BAS?
- Will the supervisory controller send many specific commands to many low-level actuators? Or will it send fewer commands and allow device-level controllers to be more autonomous?
- Will the control approach involve communication with utility control and signaling systems, and if so, will the communication be one-way or two-way?
- How many inputs and outputs will be required?

- How will the number of inputs and outputs affect installation costs?
- What communication protocols will be required to complete the system?
- Will communication across firewalls be required?

Options for energy storage and flexible loads will also vary in terms of availability and responsiveness. In general, batteries can respond more quickly than lighting and HVAC-integrated TES systems. This suggests potential tradeoffs between dispatch order and different metrics. For example, if controlling net load within a tight deadband is paramount for a particular application, then the ideal control algorithm might rely heavily on dedicated batteries. On the other hand, if a wider deadband is tolerable, and if dispatching other resources can potentially reduce battery size or extend battery life, then the ideal control algorithm might leverage multiple resources. Tradeoffs such as these can be evaluated with simulation tools that capture interactions between flexible loads, generation, and energy storage systems.

5.2 Modeling Strategies

During the design of advanced load-shaping controllers and compatible building systems, whole-building energy simulation may be necessary if one or more of the following conditions are true for a particular project:

- Building loads will change as a result of supervisory controller action.
- The behavior of energy storage systems and building loads is interrelated.
- There is a need to model interactions between building subsystems.

A wide range of ZEB system specifications and supervisory controller design parameters can be evaluated, refined, and optimized with the assistance of whole-building energy simulation. When selecting a modeling tool, it is important to consider what physical effects are necessary to capture and what time resolution is necessary to model. For this study, the team used EnergyPlus and OpenStudio tools, including the EnergyPlus EMS programming feature, to model integrated systems with a 1-minute simulation time step.

A comprehensive multiparameter optimization of the supervisory controller design was outside the scope of this study, but such an analysis could be conducted with OpenStudio tools. In such an analysis, multiple metrics could be prioritized and used to guide an iterative controller design process. Potential metrics include (but are not limited to):

- Peak demand
- Peak export rate
- Net load range
- Net load ramp rate

- Total facility energy use
- Total facility energy use plus battery losses
- Battery cycling metrics, such as roundtrip throughput
- Imported energy
- Net imported energy (imported energy minus exported energy)
- Ratio of exported energy to imported energy
- Ratio of renewable generation used on site to renewable generation exported
- Ratio of renewable generation used on site to total facility energy use
- Fossil fuel use
- Greenhouse gas emissions
- Life cycle cost.

5.3 Variability Analysis

For HVAC, lighting, and PV systems, the team identified methods for modeling subsystem loads and generation at 1-minute timescales through a combination of whole-building energy simulation tools and high-resolution weather data.

To model other loads, the team developed and explored a preliminary set of techniques for the analysis, extraction, and synthesis of short-term load variability signatures for commercial buildings. Further work is needed, however, to determine the best ways to apply such techniques to test building controls and equipment. To advance industry knowledge and practices in this area, future research areas of interest include:

- **Investigation of the impact of short-term load variability on building equipment and controls.** It is not clear what features of variability drive the behavior (or misbehavior) of building equipment and controls. Therefore, research is needed to determine which characteristics of variability must be modeled faithfully and which may be approximated. Such research could lead to standard short-term variability templates that may be used to check proper functionality in building equipment and control systems.

- **Investigation of typical short-term load variability.** Very few high-resolution load data sets are available for short-term variability analysis. Future work could include collecting and analyzing high-resolution data sets from a variety of commercial building types in order to establish typical variability patterns and characteristics. Establishing variability signatures that are representative of the commercial building stock at large will provide modelers with higher quality inputs.

- **Refine variability extraction and synthesis methods.** Future research should refine variability extraction and synthesis methods to best represent the variability features deemed most critical to represent accurately. Such features are as of yet unknown but might include a lack of low-frequency content, steep ramps, or large spikes. Refined synthesis methods might also replicate differences in variability between occupied and unoccupied hours, between weekdays and weekends, and between cold and hot seasons.

References

Akashi, Yukio, and Jason Neches. 2004. "Detectability and Acceptability of Illuminance Reduction for Load Shedding." *Journal of the Illuminating Engineering Society* Winter 2004:3–13.

DOE (U.S. Department of Energy). 2012. *The Design-Build Process for the Research Support Facility.* DOE/GO-102012-3293. Accessed from http://www.nrel.gov/docs/fy12osti/51387.pdf.

DOE (U.S. Department of Energy). 2015. *A Common Definition for Zero Energy Buildings.* DOE/EE-1247. Accessed from http://energy.gov/sites/prod/files/2015/09/f26/bto_common_definition_zero_energy_buildings_093015.pdf.

DOE (U.S. Department of Energy). 2016a. "EnergyPlus." Accessed from https://energyplus.net/.

DOE (U.S. Department of Energy). 2016b. "OpenStudio." Accessed from https://www.openstudio.net/.

DOE (U.S. Department of Energy). 2016c. "Application Guide for EMS." Accessed from https://energyplus.net/documentation.

Kung, Feitau, Michael Deru, and Eric Bonnema. 2013. *Evaluation Framework and Analyses for Thermal Energy Storage Integrated with Packaged Air Conditioning.* Golden, CO: National Renewable Energy Laboratory (NREL). NREL Report No. TP-5500-60415. http://www.nrel.gov/docs/fy14osti/60415.pdf.

Newsham, Guy, and Benjamin Birt. 2010. "Demand-Responsive Lighting: A Field Study." *Leukos* 6:3, 203–226. http://ies.tandfonline.com/doi/abs/10.1582/LEUKOS.2010.06.03002.

OpenEI. 2016. "Open Energy Information." Accessed from http://en.openei.org/wiki/Main_Page.

Office of Federal Sustainability. 2015. "Implementing Instructions for Executive Order 13693: Planning for Federal Sustainability in the Next Decade." Accessed from https://www.whitehouse.gov/sites/default/files/docs/eo_13693_implementing_instructions_june_10_2015.pdf.

State of California. 2016. "Building Energy Efficiency Program." Accessed from http://www.energy.ca.gov/title24/.