

Home Energy Management under Realistic and Uncertain Conditions: A Comparison of Heuristic, Deterministic, and Stochastic Control Methods

Michael Blonsky^{a,b}, Killian McKenna^a, Jeff Maguire^a, Tyrone Vincent^b

^aNational Renewable Energy Laboratory, 15013 Denver West Parkway, Golden, CO 80401, USA

^bColorado School of Mines, 1500 Illinois St, Golden, CO 80401, USA

Abstract

Home energy management systems (HEMS) have been shown to reduce energy bills and to provide grid services including peak demand reduction and demand flexibility. However, uncertainty in residential energy systems is a significant issue and can reduce the benefits of a HEMS to the homeowner or grid operator. Sources of uncertainty include weather forecasts, predictions of energy-related occupant activities (e.g., hot water draws), and parameter estimation for the building envelope and energy-consuming equipment.

This paper tackles the problem of uncertainty by developing a framework that simulates HEMS in uncertain conditions and evaluates the performance of multiple control strategies. A linear, reduced-order residential building model for model predictive control applications is derived and compared to a full-order model. Stochastic model predictive control is shown to perform better than deterministic and heuristic methods when considering realistic forecasts with uncertainty. The framework can evaluate the performance of HEMS in real-world applications, which can help de-risk HEMS deployment.

Keywords: home energy management systems, stochastic controls, model predictive control, residential energy modeling, distributed energy resources

Email address: Michael.Blonsky@nrel.gov (Michael Blonsky)

1. Introduction

The electric grid is experiencing an increase in both variable renewable generation resources and in electricity consumption due to the electrification of space heating and transportation [1]. These trends lead to large temporal changes in supply and demand and present new challenges in operating the power system, which increases the need for demand flexibility. The Brattle Group estimates approximately 200 GW of cost-effective demand flexibility potential in the United States by 2030, with much of the increase coming from residential loads, including smart thermostats, smart water heaters, electric vehicle (EV) chargers, and behind-the-meter batteries [2].

Each of these devices can provide demand flexibility as well as benefits to the homeowner when controlled through a home energy management system (HEMS). HEMS in the literature typically use model predictive control (MPC) to optimize energy cost and consumption while considering occupant comfort, device limitations, and other factors [3, 4, 5]; however, most studies use deterministic MPC (DMPC), which does not consider the uncertainty from forecasts, models, or sensor measurements. As shown in Section 2, these sources of uncertainty can all be significant and can lead to sub-optimal control performance in HEMS applications.

There are a few MPC methods that consider uncertainty in the control framework. Stochastic model predictive control (SMPC)—also referred to as chance-constrained optimization [3] and conditional value at risk [6, 7]—considers the probability distribution of uncertain variables within the optimization framework [8, 9]. Robust MPC considers extreme values for uncertain variables, which lead to improved performance in worst-case conditions but not necessarily in typical conditions [10]. Scenario-based MPC methods use an ensemble of model disturbances to quantify uncertainty [11, 12]. Stochastic dynamic optimization and other multistage optimization methods optimize over a set of discrete state transitions, which can greatly increase the computational complexity of the controller [3, 13].

Most papers in this research area often focus on novel control architectures and their performance benefits over existing strategies, but they pay little attention to quantifying uncertainty in their simulations or to assessing its impact on their findings. Forecast uncertainty has been considered in related fields including commercial buildings, EV fleets, and behind-the-meter battery systems [14, 15, 16, 17]. In particular, many papers present control strategies for commercial building heating, ventilating, and air-conditioning (HVAC) that

incorporate forecast and model uncertainty [10, 18, 19]. One study on commercial HVAC found that DMPC performed best at low levels of uncertainty, robust MPC performed best at intermediate levels, and heuristic controls performed best at high levels [20]. SMPC was used for a residential building with an EV and battery, but did not consider modeling uncertainty or any thermostatic loads like HVAC and water heating [21]. Another SMPC paper include these loads in the HEMS control but only considers uncertainty in weather forecasts [8]. Other HEMS control strategies consider occupancy forecasts but use simplified models with no model uncertainty [6, 11]. One paper quantifies uncertainty for residential HVAC energy for an aggregation of homes but does not consider uncertainty within a single home [22]. The authors are not aware of any studies on HEMS that consider both forecast and model uncertainty for multiple flexible loads within a home.

In this paper, multiple control strategies for HEMS are evaluated using a novel modeling and control framework that incorporates multiple sources of uncertainty. A heuristic control strategy that resembles standard residential building controls is used as a baseline. The framework incorporates MPC using an objective to minimize energy costs and occupant discomfort. It is run under deterministic (DMPC) conditions as well as stochastic (SMPC) conditions using the same objective and constraints. The control performance for each strategy is evaluated under various levels of uncertainty, accounting for the uncertainties addressed in Section 2. The key contributions of the paper include:

- Development of a HEMS framework that incorporates uncertainties caused by errors in the prediction forecast, control model, and sensor measurements, and that includes models for HVAC equipment, water heaters, EVs, and batteries
- Evaluation of heuristic, DMPC, and SMPC strategies using common forecasts methods and a high-resolution building model that can accurately simulate building dynamics for a typical building
- Model reduction and linearization of radiation and infiltration for a multi-node white-box envelope model for use in MPC
- A discussion on the sources of uncertainty in HEMS applications and their implications for HEMS control strategies

In Section 2, the sources of uncertainty in HEMS controls using MPC are described. In Section 3, the models used for evaluating the different controls strategies are

Nomenclature

Envelope Parameters

| | |
|--------------|--|
| H_i | Heat injected into node i (W) |
| H_s | Vector of heat gains from solar radiation on each exterior surface (W) |
| $H_{az,inf}$ | Heat due to infiltration from zone a to zone z (W) |
| H_{int} | Internal heat gains injected into the main indoor zone (W) |
| $H_{iz,rad}$ | Heat due to radiation for node i facing zone z (W) |
| $R_{az,inf}$ | Thermal resistance due to infiltration between zones a and z (K/W) |
| R_{ij} | Thermal resistance between nodes i and j (K/W) |
| $R_{iz,rad}$ | Thermal resistance due to radiation for node i facing zone z (K/W) |
| T_a | Ambient temperature ($^{\circ}$ C) |
| T_g | Ground temperature ($^{\circ}$ C) |
| T_m | Main indoor temperature ($^{\circ}$ C) |
| $T_{s,i}$ | Surface temperature of a boundary next to node i ($^{\circ}$ C) |

Equipment Parameters

| | |
|---------------|---|
| η_l | Efficiency of load l |
| θ_i | Water temperature of node i ($^{\circ}$ C) |
| θ_{wm} | Water mains temperature ($^{\circ}$ C) |
| E_l | Energy capacity of load l (kWh) |
| H_{wh} | Heat injected into the water tank (W) |
| P_l | Power of load l (kW) |
| P_{chg} | Battery charging power (kW) |
| P_{dis} | Battery discharging power (kW) |
| P_{drive} | EV power discharged due to driving (kW) |
| p_{drive} | Fraction of time that the EV is driving |
| P_{house} | Total house power (kW) |

| | |
|-------------|--|
| p_{leave} | Probability that the EV is leaving the parking space |
| P_{unc} | Uncontrollable load power (kW) |
| s_l | State of charge of load l |
| V_{draw} | Hot water flow rate (L/min) |

Stochastic Modeling Symbols

| | |
|---|--|
| $\sigma_{x_i}^2$ | Variance of random variable x_i |
| $\tilde{x} \sim \mathcal{N}(\hat{x}, \Sigma_x)$ | Gaussian random vector x with mean \hat{x} and variance Σ_x |
| A_i | State transition matrix for system i |
| B_i | Controllable input matrix for system i |
| C_i | Output matrix for system i |
| G_i | Uncontrollable input matrix for system i |
| K | Kalman gain |
| u_i | Controllable input vector for system i |
| x_i | State vector for system i |
| y_i | Output vector for system i |
| z_i | Uncontrollable input vector for system i |

Other Symbols

| | |
|--------------------------|---|
| β_{x_i} | z-score associated with a constraint on x_i |
| $\hat{\lambda}_{power}$ | Average cost of power over the forecast horizon (\$/kWh) |
| λ_i | Cost associated with cost term i |
| \bar{x}, \underline{x} | Maximum and minimum constraints on x |
| Φ^{-1} | Inverse cumulative distribution function of the normal distribution |
| h | Discrete time step for MPC horizon |
| I | Identity matrix |
| J_i | Cost function term i (\$) |
| k | Discrete time step |
| n_h | Number of steps in MPC horizon |
| t_s | Time resolution of simulation (s) |

described with a focus on the linear model used within the MPC framework. In Section 4, the heuristic control strategy and formulate the DMPC and SMPC problems are summarized. In Section 5, multiple forecast

techniques that are used in the MPC framework are defined. In Section 6, the performance of the different control strategies is compared. The paper concludes with a summary of findings and explorations for future work.

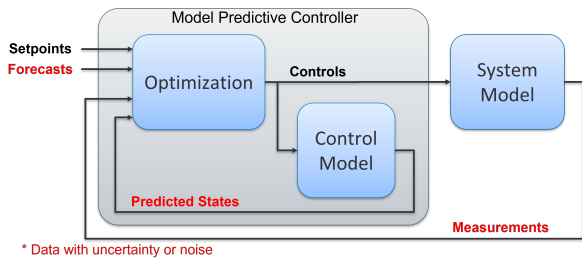


Figure 1: Schematic for a model predictive control algorithm

2. Uncertainty Considerations in HEMS

A literature review by Beaudin et al. finds that most studies on HEMS use an MPC framework, but most do not consider the impact of uncertainty on control performance [3]. When MPC methods are tested in a hypothetical scenario with exact forecasts, simplified system models, and little or no noise in communication signals, they are nearly always shown to be optimal strategies that perform better than heuristic or other methods; however, real systems are not so simple, and we argue that HEMS control strategies should be evaluated in conditions with realistic uncertainty levels to truly understand their value.

Figure 1 shows a schematic of an MPC algorithm and highlights the data that are likely to be stochastic. MPC requires forecasts for model inputs in the horizon window that often cannot be perfectly estimated at the controller run time. The MPC control model is not a perfect representation of the system model, causing the predicted states and outputs to deviate from the actual model states. Finally, sensor measurements from the system model can be noisy or biased. Next, we elaborate on each of these sources with a focus on HEMS applications.

2.1. Forecast Uncertainty

Forecasts for weather and occupant behavior likely have the largest amount of uncertainty in residential applications due to their stochastic nature. Residential energy consumption is highly dependent on both occupant activity and weather, making these forecasts critical pieces of information when optimizing for future control decisions. For example, preconditioning air to shift HVAC consumption away from a peak demand period requires knowledge of weather variables that are used to calculate HVAC loads [18, 23]. Other examples include preheating a water heater before a large water draw event or charging an EV such that it is fully charged before the occupant leaves [24].

Short-term weather forecasts are often accurate for ambient temperature, but solar irradiance is more difficult to predict [25]. Occupant behavior is highly unpredictable, and good models often require data for a particular person or household [26]. Most weather variables (e.g., temperature, irradiance) are more continuous and smoother than many occupancy variables, for example, hot water draws or discrete activity changes such as plugging in an EV or turning on an oven [27]. HEMS often use a single forecast for uncontrollable loads rather than estimating consumption from each type of activity [5].

In many papers, forecast uncertainty is quantified using heuristics [8] or is not defined at all. Auto-regressive moving average models have been used to estimate uncertainty for weather variables [28, 29, 30] as well as for occupancy variables [31]. Other forecasting methods usually do not quantify uncertainty, for example, sophisticated weather forecasts [13, 15] and Markov chain-based stochastic occupancy models [22, 29].

2.2. Model Uncertainty

An MPC framework requires an underlying model that represents the system, but it is often an imperfect approximation of the system. Many MPC frameworks use a linear model, which is required to use fast and convex optimization methods, including linear programming, quadratic programming, and mixed-integer linear programming methods [3, 17, 32]. Linear models must either simplify or ignore nonlinearities such as device efficiency, device degradation, and radiative thermal processes. Models for specific homes often lack detailed data, including envelope and equipment parameters as well as training data for black box models, which can lead to inaccurate model parameters [10, 22]. Because every house is different, it can be very time consuming and costly to create an accurate model representation [33].

Another key source of uncertainty arises from the time resolution of the model and the MPC controller. Time resolutions may be constrained by the time resolution or update frequency of the forecast, by measurement or actuator update rates, or by the computational complexity of the MPC algorithm. Time resolution may also be purposefully reduced to eliminate the need to consider fast system dynamics, for example, equipment cycling behavior or high-resolution occupant activities [34, 35]. Although low-resolution controllers might work well for some objectives—for example, reducing energy costs with an hourly electricity price—they are likely to reduce the performance of any objective with

real-time impacts on the occupant, for example, thermal comfort or convenience associated with EV charging.

2.3. Measurement Uncertainty

Uncertainty can also arise from sensor measurements [4, 36]. Sensor noise is not typically an issue, but sensors can have low time resolution or low precision, which leads to reduced control performance [37]. Sensors can also be biased, especially if the sensor is not directly capturing the variable of interest. For example, a thermostat located near a vent or a sunny window might report a temperature that is not accurately representing the house temperature [26]. Some sensors might also have limited communication abilities, for example, an EV charger that can only measure the EV state of charge when the vehicle is parked and plugged in.

3. Model Description

To simulate model uncertainty in the control framework, two models are defined for each end use: a system model and a control model for the MPC (see Figure 3). This section describes the linearization and other techniques used to derive the MPC models.

The system models are taken from OCHRE™, a residential energy model with controllable device models for each load examined in this paper. OCHRE is a white-box model that uses weather and occupancy schedules to calculate heat generation and power from each device. More details on OCHRE models can be found in [35] and [24]. The MPC models described in this section are derived from the OCHRE models.

3.1. Controllable Load Model

In general, a linear model is defined for each controllable load $l \in \mathcal{L}$ in state space form:

$$\begin{aligned} x_l(k+1) &= A_l x_l(k) + B_l u_l(k) + G_l z_l(k) \\ y_l(k) &= C_l x_l(k) \end{aligned} \quad (1)$$

where x_l is the state vector, u_l is the controllable input vector, z_l is the uncontrollable input vector, y_l is the output vector, and A_l , B_l , G_l , and C_l are state space system matrices.

In this paper, the set of controllable loads is $\mathcal{L} = [hvac, wh, ev, batt]$ corresponding to an air conditioner, water heater, electric vehicle, and battery, respectively. Each load model contains at least one state and at least one controllable input that corresponds to the load power. Most models include at least one uncontrollable input that corresponds to an uncertain variable that must be forecasted. If the output vector y_l is not defined, it is assumed that $y_l = x_l$, and C_l is the identity matrix.

3.2. HVAC and Envelope

The system model for the HVAC equipment and the building envelope is a nonlinear, white-box model created using OCHRE. It includes a linear, multi-node equivalent circuit model for conduction and convection between air zones and the boundaries between them, as well as nonlinear elements for infiltration, ventilation, and long-wave radiation. The OCHRE model also includes a thermostat controller with a deadband which sets the power based on a temperature set point.

To develop a linear envelope model for the MPC, infiltration, interior long-wave radiation, and exterior long-wave radiation are linearized separately. For each heat transfer pathway, resistors are added to OCHRE's equivalent circuit model (see [35], figures 1 and 2), as shown in Figure 2. In the figure, T_a is the ambient temperature, T_z is the temperature of an interior zone, $T_{s,i}$ is the surface temperature of boundary i , and subscripts 1, 2, and 3 correspond to boundary materials connected to either the ambient or an interior zone.

For infiltration (and mechanical ventilation), a resistance $R_{az,inf}$ is added between an interior air zone and the ambient air zone. The resistance value is calculated by linearizing the heat gain at an operating point:

$$R_{az,inf} = \frac{\Delta T_{op}}{H_{az,inf}} \quad (2)$$

where ΔT_{op} is the temperature difference between zone z and ambient at the operating point, and $H_{az,inf}$ is the infiltration and ventilation heat gains to zone i at the operating point. $H_{az,inf}$ also depends on wind speed, which must also be defined in the operating point.

We use a similar approach to linearize the interior and exterior radiation. The radiation equation for the surface of boundary material i is:

$$H_{iz,rad} = \epsilon_i \sigma a_i (T_{z,rad}^4 - T_{s,i}^4) \quad (3)$$

where ϵ_i is the emissivity of material i , a_i is the area of the material i , and σ is the Stefan-Boltzmann constant. For the ambient zone, we assume $T_{z,rad} = T_a$, ignoring effects of the sky radiation temperature. For interior zones, $T_{z,rad}$ depends on the surface temperatures of all boundaries facing the zone. Linearizing with the operating point $T_{z,rad} = T_{s,i} = T_{op,rad}$ gives:

$$\frac{1}{R_{iz,rad}} = \frac{dH_{iz,rad}}{d(T_{z,rad} - T_{s,i})} = 4\epsilon_i \sigma a_i T_{op,rad}^3 \quad (4)$$

where $R_{iz,rad}$ is a resistance value that accounts for radiation effects. For exterior radiation, the resistor is applied in parallel with an exterior film resistance that models the effects of convection (see Figure 2, $R_{a1,filn}$).




Figure 2: A linear envelope model with infiltration and interior and exterior radiation elements

For interior radiation, the resistor is applied between the boundary surface and a fictional node $T_{z,rad}$ representing all of the radiation within an interior zone. Using Kirchoff's Current Law, this network topology ensures that all radiative heat within a zone sums to zero. The fictional node and all boundary surface nodes are eliminated by using the star-mesh transform to generate resistance values between each boundary surface within an interior zone.

The updated envelope model is converted into a linear, discrete-time state space model with states for each interior node temperature and inputs for each exterior node temperature and each interior node heat gain (see [35] Eq. 2–3). OCHRE's envelope model includes many interior nodes (30–45 for a typical home), which can lead to slow computational speed when used within an MPC framework. Balanced truncated model reduction is used to reduce the number of states of the envelope model [38].

Before the model reduction, unused inputs and outputs are eliminated and inputs are normalized based on typical variability. The main indoor temperature is the only output, and the HVAC power is the only controllable input, i.e., $y_{hvac} = T_m$ and $u_{hvac} = P_{hvac}$. The uncontrollable inputs are $z_{hvac} = [T_a, T_g, H_{int}, H_s^T]^T$, where T_g is the ground temperature, H_{int} is the internal heat gains in the main zone, and H_s is a vector of the heat gains from solar radiation on each exterior surface (e.g., walls, roof, attic walls), which depends on the solar irradiance and the area and absorptivity of each surface. The state vector x_{hvac} is an output of the model reduction

algorithm and represents the temperatures of the interior zones and boundaries.

Internal heat gains include heat from occupants, other equipment, and solar radiation from windows. The total heat gain to the main zone H_m is a combination of the internal heat gains and the heat from the HVAC equipment. When an air conditioner is used:

$$H_m(k) = -\eta_{hvac} P_{hvac}(k) + H_{int}(k) \quad (5)$$

where η_{hvac} is the HVAC coefficient of performance and is assumed to be constant.

3.3. Water Heater

We use OCHRE's heat pump water heater and 2-node water tank models for the system model. The 2-node tank model includes nonlinearities due to heat transfer from water draws. Similar to the HVAC model, OCHRE's water heater model uses a thermostat with deadband control that converts a temperature setpoint into an on/off power signal. The temperature setpoint can be controlled externally, but the thermostat ensures that the deadband remains within the thermal comfort region.

We define a 1-node linear model for the control model that accounts for convection and conduction similarly to the envelope model described in Section 3.2. Heat injections into the tank come from water heater power and water draws:

$$H_{wh}(k) = \eta_{wh} P_{wh}(k) + V_{draw}(k) \rho c_p (\theta_{wm} - \theta_1(k)) \quad (6)$$

where η_{wh} is the coefficient of performance of the heat pump water heater, θ_1 is the tank temperature, θ_{wm} is the water mains temperature, V_{draw} is the water draw flow rate, ρ is the density of water, and c_p is the heat capacity of water. To linearize the biquadratic term for water draws, we assume constant temperature differences. We note that this formulation does not include a backup heating element for the heat pump water heater, which is included in OCHRE's model.

The linear state space equation for the water heater has one state for the water tank temperature, $x_{wh} = \theta_1$, and one controllable input, $u_{wh} = P_{wh}$. The uncontrollable inputs are $z_{wh} = [V_{draw}, \theta_a]^T$, and we assume that the water mains temperature is known.

3.4. Electric Vehicle

OCHRE's EV model tracks EV charging but does not model battery discharging during driving [24]. To account for discharging, a term is added that decreases the EV state of charge (SOC) when the EV is driving:

$$\frac{ds_{ev}(k)}{dt} = \frac{1}{E_{ev}}(\eta_{ev}P_{ev}(k) - P_{drive}p_{drive}(k)) \quad (7)$$

where s_{ev} is the EV SOC, E_{ev} is the EV energy capacity, η_{ev} is the charge efficiency, P_{ev} is the EV charging power, P_{drive} is the power lost while driving, and p_{drive} is the percentage of time that the EV is driving.

The linear state space equation for the EV has one state for the EV SOC, $x_{ev} = s_{ev}$. The only uncontrollable input is $z_{wh} = p_{drive}$.

3.5. Battery

We use OCHRE's linear battery model to track the battery SOC. Due to separate charging and discharging efficiencies, the controllable input is split into two variables, i.e., $u_{batt} = [P_{chg}, P_{dis}]^T$, and we define $P_{batt} = P_{chg} - P_{dis}$. The battery model has one state for the battery SOC, $x_{batt} = s_{batt}$. There are no uncontrollable inputs for the battery model, i.e., z_{batt} is an empty vector. We note that $A_{batt} = 1$ because there is no self-discharge modeled.

3.6. Combined House Model

The four controllable equipment models are combined in parallel to create a linear model for the house:

$$\begin{aligned} x(k+1) &= Ax(k) + Bu(k) + Gz(k) \\ y(k) &= Cx(k) \end{aligned} \quad (8)$$

where:

$$\begin{aligned} x &= [x_{hvac}^T, \theta_1, s_{ev}, s_{batt}]^T \\ u &= [P_{hvac}, P_{wh}, P_{ev}, P_{chg}, P_{dis}]^T \\ z &= [T_a, T_g, H_{int}, H_s^T, V_{draw}, \theta_a, p_{drive}]^T \\ y &= [T_m, \theta_1, s_{ev}, s_{batt}]^T \\ A &= \text{diag}([A_{hvac}, A_{wh}, A_{ev}, A_{batt}]) \end{aligned}$$

and the other matrices are defined similarly to A . The whole house power is defined as:

$$P_{house}(k) = \sum_{l \in \mathcal{L}} P_l(k) + P_{unc}(k) \quad (9)$$

where P_{unc} is the power from all uncontrollable house loads. The uncontrollable load power is considered as an uncertain input, similar to the inputs in z .

3.7. Stochastic House Model

The control model in (8) can be used directly in a deterministic control framework. For SMPC, the uncontrollable inputs become random variables, and noise \tilde{v} is added to the outputs. The stochastic model becomes:

$$\begin{aligned} \tilde{x}(k+1) &= A\tilde{x}(k) + Bu(k) + G\tilde{z}(k) \\ \tilde{y}(k) &= C\tilde{x}(k) + \tilde{v} \end{aligned} \quad (10)$$

where:

$$\begin{aligned} \tilde{x}(k) &\sim \mathcal{N}(\hat{x}(k), \Sigma_x(k)) \\ \tilde{z}(k) &\sim \mathcal{N}(\hat{z}(k), \Sigma_z(k)) \\ \tilde{y}(k) &\sim \mathcal{N}(\hat{y}(k), \Sigma_y(k)) \\ \tilde{v} &\sim \mathcal{N}(0, \Sigma_v) \end{aligned}$$

are all Gaussian random vectors. The mean and variance of all the uncontrollable inputs distributions $\tilde{z}(k)$ are based on forecast estimates (see Section 5), and we assume that there is no covariance between the inputs. The diagonals of the covariance matrices correspond to the variance of individual input or output variables, which we denote as $\sigma_{x_i}^2$, for example, $\text{diag}(\Sigma_y) = [\sigma_{T_m}^2, \sigma_{\theta_1}^2, \sigma_{s_{ev}}^2, \sigma_{s_{batt}}^2]$.

The probability distributions of the state and output matrices are calculated using a Kalman filter [39]:

$$\begin{aligned} \hat{x}(k) &= \hat{x}^{(-)}(k) + K(y^{(-)}(k) - C\hat{x}^{(-)}(k)) \\ \Sigma_x(k) &= (I - KC)\Sigma_x^{(-)}(k) \\ \hat{y}(k) &= C\hat{x}(k) \\ \Sigma_y(k) &= C\Sigma_x(k)C^T \end{aligned} \quad (11)$$

where $y^{(-)}$ is the most recent house status received from OCHRE, and:

$$\begin{aligned} K &= \Sigma_x^{(-)}(k)C^T(C\Sigma_x^{(-)}(k)C^T + \Sigma_v)^{-1} \\ \hat{x}^{(-)}(k+1) &= A\hat{x}(k) + Bu(k) + G\hat{z}(k) \\ \Sigma_x^{(-)}(k+1) &= A\Sigma_x(k)A^T + G\Sigma_z(k)G^T \end{aligned} \quad (12)$$

The equations above can be used directly when k is the current time. When calculating the state and output distributions in the future, $y^{(-)}(k)$ is unknown and we set $\hat{x}(k) = \hat{x}^{(-)}(k)$. We also assume a closed-loop control strategy that accounts for past disturbances and reduces the expected uncertainty in future states. This is achieved by ignoring the $\Sigma_x(k)$ term when calculating $\Sigma_x^{(-)}(k+1)$.

3.8. Model Implementation

The implementation of the HEMS controller and OCHRE is described in Figure 3. Before the simulation begins, forecasts are generated using time-series data from weather and occupancy schedules. Every 30 minutes, the HEMS controller receives the most recent house status $y^{(-)}(k)$ from OCHRE. It receives forecasts from a forecast generator of the form $Z(k) = \{z(k, h) : h \in [k, k + n_h - 1]\}$, where $z(k, h)$ is the forecast at time h evaluated at time k , and n_h is the number of steps in the MPC horizon. The HEMS controller uses an MPC optimization to solve for $u(h)$, and the control model predicts the future states $x(h+1)$ throughout the forecast horizon. The controller sends the optimal set points $y(k+1)$ to OCHRE, which implements the set point at 1-minute resolution. This process repeats every 30 minutes throughout the simulation.

Note that the HEMS controller sends the temperature and SOC set points $y(k+1)$, not the power set points $u(k)$, to OCHRE. This allows OCHRE's device controllers to determine the device power at 1-minute resolution, which enables equipment cycling for the air conditioner and the water heater. Providing OCHRE with these set points reduces the effects of model and forecast uncertainty in achieving the desired set points, but it could cause the device power to deviate from the power expected by the MPC. This control signal works best for the MPC formulation in this paper; however, power set points might be more appropriate if a demand charge or a demand response event is involved and precise power controls are desired.

We also note that the house status $y^{(-)}(k)$ includes all state variables in the deterministic case. This gives the DMPC perfect knowledge of the initial system state. In the SMPC case, the house status only includes output variables and the Kalman filter is used to estimate the states.

4. Control Formulation

We evaluate three types of HEMS controls: heuristic controls, DMPC, and SMPC. In this section, these three

control strategies are outlined for each controllable device. All controllers operate at a 30 minute time resolution.

For multiple equipment types, a soft constraint is used to limit discomfort or inconvenience costs. Soft constraints were chosen to more easily compare the different control strategies and to define how to quantify the costs when constraints are violated [40].

4.1. Heuristic Controls

The heuristic controls for HVAC and water heating use a thermostat control with a deadband and a precooling strategy designed around a time-of-use rate. The top of the HVAC deadband defines the comfort temperature \overline{T}_m that is used to evaluate all of the HVAC control strategies, with a comfort cost incurred when that threshold is exceeded. The water heater control uses a minimum comfort temperature $\underline{\theta}_1$, which is also used to evaluate all control strategies.

The heuristic controller charges the EV at its maximum power immediately upon arrival. This control strategy ensures that the EV SOC will be at its maximum possible value when the EV leaves, but it disregards the variable cost of charging at different times of day.

The heuristic battery controller performs energy arbitrage based on a time-of-use rate. The battery begins charging at 10 AM and begins discharging at the start of the peak period.

4.2. Deterministic MPC

The DMPC formulation uses the linear house model described in Section 3.6 to minimize energy costs and costs associated with occupant comfort and convenience. All terms in the objective function include a "cost coefficient" λ that normalizes the term into units of dollars. The first term defines the energy costs:

$$J_{power}(k) = \lambda_{power}(k)t_s P_{house}(k) \quad (13)$$

where λ_{power} is a time-varying price of electricity, and t_s is the control time resolution.

The objective includes occupant comfort and convenience costs for HVAC, water heating, and EV charging, which depend on air temperature, hot water temperature, and EV SOC, respectively. The EV control is designed to minimize any inconvenience associated with a low SOC at the time the EV leaves [24]. The costs are defined as:

$$\begin{aligned} J_{hvac}(k) &= \lambda_{hvac} |T_m(k) - \overline{T}_m|_+ \\ J_{wh}(k) &= \lambda_{wh} V_{draw} |\overline{\theta}_1 - \theta_1(k)|_+ \\ J_{ev}(k) &= \lambda_{ev} (1 - s_{ev}(k)) p_{leave}(k) \end{aligned} \quad (14)$$

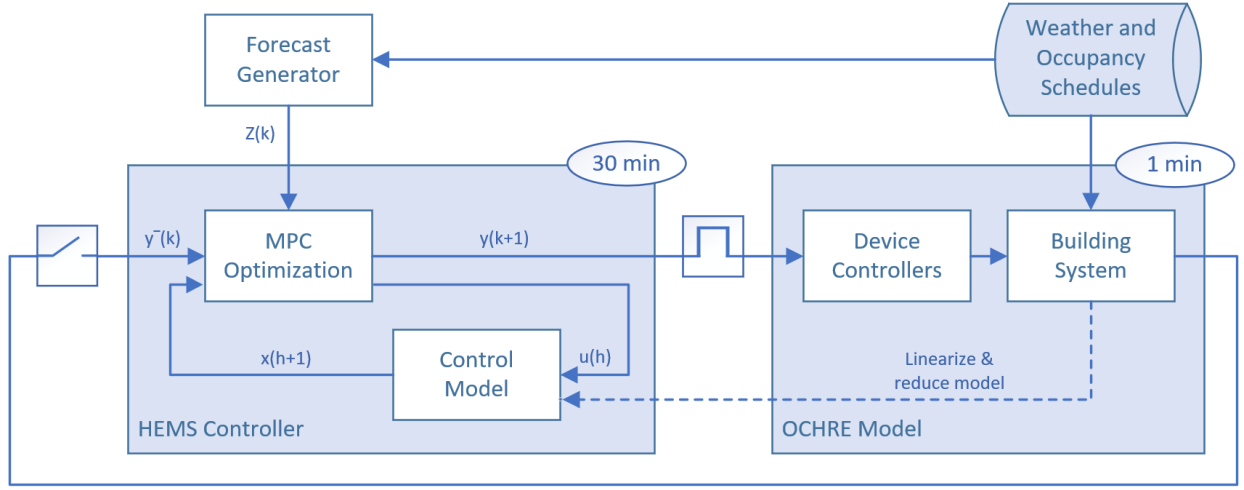


Figure 3: Flow diagram showing the inputs and outputs of the HEMS controller and OCHRE

where λ_l is a comfort cost associated with each device, $p_{leave}(k) = p_{drive}(k) - p_{drive}(k-1)$ is the probability that the EV is leaving during time step k , and $|\cdot|_+ = \max(\cdot, 0)$. J_{hvac} only accounts for discomfort from hot temperatures, but a similar cost term could be included for cold temperatures.

For the battery, a terminal cost term is included, which is designed to make the battery control indifferent to discharging at the end of the horizon:

$$J_{batt}(k+n_h) = -\hat{\lambda}_{power}(k)E_{batt}\eta_{dis}s_{batt}(k+n_h) \quad (15)$$

where $\hat{\lambda}_{power}$ is the average cost of electricity over the forecast horizon, E_{batt} is the battery energy capacity, and η_{dis} is the battery discharge efficiency. Note that this objective only considers the SOC at the end of the horizon, i.e., at time $k+n_h$.

We also include a term to smooth the power profile of each controllable device. This term can help reduce device degradation and reduce the set of optimal solutions, which improves the solver's consistency across multiple scenarios. The peak cost term is:

$$J_{peak}(k) = \lambda_{peak} \sum_{l \in \mathcal{L}} \max_{h \in [k, k+n_h-1]} |P_l(h)| \quad (16)$$

where λ_{peak} is a small positive number.

The DMPC objective function at current time k with n_h steps in the horizon is:

$$J(k) = \sum_{h=k}^{k+n_h-1} (J_{power}(h) + J_{hvac}(h+1) + J_{wh}(h+1) + J_{ev}(h+1)) + J_{batt}(k+n_h) + J_{peak}(k) \quad (17)$$

The deterministic optimization problem is defined as:

$$(\mathcal{P}_1) \min_{u(h)} : J(k) \quad (18)$$

$$\text{s.t.} \quad x(h+1) = Ax(h) + Bu(h) + Gz(h)$$

$$y(h) = Cx(h)$$

$$0 \leq u_l(h) \leq \bar{u}_l \quad \forall l \in \mathcal{L}$$

$$\underline{s}_{batt} \leq s_{batt}(h+1) \leq \bar{s}_{batt}$$

$$s_{ev}(h+1) \leq 1$$

$$0 \leq P_{ev}(k) \leq \bar{P}_{ev}(1 - p_{drive}(k))$$

$$\forall h \in [k, k+n_h-1]$$

where, in general, $\bar{\cdot}$ and $\underline{\cdot}$ are the minimum and maximum values of a given variable. Note that there is no minimum SOC constraint for the EV because of the uncertainty due to driving loss. Because p_{drive} is either 0 or 1, the maximum EV power is constrained to zero when the EV is driving and to the maximum charging power when the EV is parked.

All constraints are linear, and the objective function can be transformed into a linear function through the addition of slack variables. This makes \mathcal{P}_1 a linear optimization problem that can be solved very quickly and guarantees a globally optimum solution.

4.3. Stochastic MPC

Under stochastic conditions, all uncontrollable inputs and all outputs are assumed to be Gaussian random variables. The stochastic optimization incorporates uncertainty in these variables by using chance constraints based on a probability ϵ_{x_i} of exceeding a constraint [9]:

$$\Pr(\tilde{x}_i \leq \bar{x}_i) \geq 1 - \epsilon_{x_i} \rightarrow \beta_{x_i} = \Phi^{-1}(1 - \epsilon_{x_i}) \quad (19)$$

where Φ^{-1} is the inverse cumulative distribution function of the normal distribution, and β_{x_i} is a z-score associated with the variable x_i .

For soft constraints, a back-off magnitude of $\beta_{x_i}\sigma_{x_i}$ incorporates the chance constraint in the objective function. The updated comfort costs from Equation (14) for the SMPC formulation are:

$$\begin{aligned} J_{hvac}(k) &= \lambda_{hvac} |\hat{T}_m(k) - \bar{T}_m + \beta_{T_m} \sigma_{T_m}(k)|_+ \\ J_{wh}(k) &= \lambda_{wh} \hat{V}_{draw} |\bar{\theta}_1 - \hat{\theta}_1(k) - \beta_{\theta_1} \sigma_{\theta_1}(k)|_+ \\ J_{ev}(k) &= \lambda_{ev} (1 - \hat{s}_{ev}(k)) (\hat{p}_{leave}(k) + \beta_{p_{leave}} \sigma_{p_{leave}}(k)) \end{aligned} \quad (20)$$

where $\sigma_{p_{leave}}$ is the standard deviation of \tilde{p}_{leave} .

The SMPC optimization problem is very similar to Problem (\mathcal{P}_1). The only changes are updating the cost terms from Equation (14) to Equation (20) and using the mean values of each random variable throughout the formulation, for example, the state equation becomes:

$$\hat{x}(h+1) = A\hat{x}(h) + Bu(h) + G\hat{z}(h)$$

We note that p_{drive} is no longer a Boolean input. No back-off magnitude is used in the maximum EV SOC constraint because EV charging is fully deterministic. The battery model has no uncertain inputs, so the stochastic controls are identical to the deterministic controls for this device.

5. Schedule and Forecast Descriptions

Forecasts were developed for uncertain inputs z , including weather data and occupancy schedules. Multiple forecasts were derived from the same underlying data, but with varying levels of complexity and accuracy. Common forecast methods from practice and from research literature were chosen, although we note that more sophisticated forecast methods can lead to forecasts with lower uncertainty.

5.1. Weather Schedule and Forecast

All simulations were performed using measured weather data from Denver, CO. Weather data were collected in real time in July 2021 using an online application programming interface (API) [41]. The same weather API was used to get weather forecast data. The weather data and forecast were collected every 1 hour, and the forecast had a duration of 48 hours and a time resolution of 1 hour.

Weather data used in the simulation include ambient temperature, pressure, relative humidity, cloud cover, and wind speed. Cloud cover data were converted

to irradiance values using a clear-sky scaling method [42, 43].

The weather forecast used in this study did not provide an estimate of the uncertainty in the forecast. The temperature distribution was calculated using the difference between the actual temperature and the forecast, aggregated by the time difference between the current time k and the prediction time h , denoted $\Delta k = h - k$:

$$\begin{aligned} T_a^{(b)}(\Delta k) &= \frac{1}{n_k} \sum_{k=1}^{n_k} (T_a^{(f)}(k, h) - T_a(h)) \\ \sigma_{T_a}^2(\Delta k) &= \frac{1}{n_k - 1} \sum_{k=1}^{n_k} (T_a^{(f)}(k, h) - T_a(h) - T_a^{(b)}(\Delta k))^2 \end{aligned} \quad (21)$$

where $T_a^{(f)}(k, h)$ is the forecasted temperature at time h evaluated at time k , $T_a^{(b)}(\Delta k)$ is the bias in the forecast, and n_k is the number of time steps with valid data. There was valid weather and forecast data for almost every hour for the month of July. The ambient temperature distribution is:

$$\tilde{T}_a(k, h) \sim \mathcal{N}(T_a^{(f)}(k, h) - T_a^{(b)}(h - k), \sigma_{T_a}^2(h - k)) \quad (22)$$

Figure 4 shows the temperature forecast at noon on the first simulation day, with the shaded region showing the 95% confidence interval. The bias of the temperature forecast with a time difference of less than 4 hours was within 1°C, and the standard deviation with a time difference of 1 hour was 0.6°C. The uncertainty increased for the first 3 hours of the forecast and then remained stable near 2°C. We note that this analysis covers only temporal uncertainty; there could also be spatial uncertainty because of weather differences between the weather station and the house.

The solar heat gain distribution was calculated using a slightly different method to keep uncertainty low when solar irradiance is low, for example, at night. The standard deviation of global horizontal irradiance (GHI) was calculated in the same way as (21) and was used to calculate the solar heat gain standard deviation:

$$\sigma_{H_s}(k, h) = H_s^{(f)}(k, h) \frac{\sigma_{GHI}(h - k)}{\mu_{GHI}(h - k)} \quad (23)$$

where $H_s^{(f)}(k, h)$ is the forecasted solar gain vector, and $\mu_{GHI}(\Delta k)$ is the average forecasted GHI. The uncertainty follows a similar pattern to temperature; the standard deviation increases during the first few hours of the forecast and then remains stable near 100 W/m².

5.2. Occupant Schedule and Forecast

A single-occupant schedule was used for all simulations. The schedule was generated using the Res-

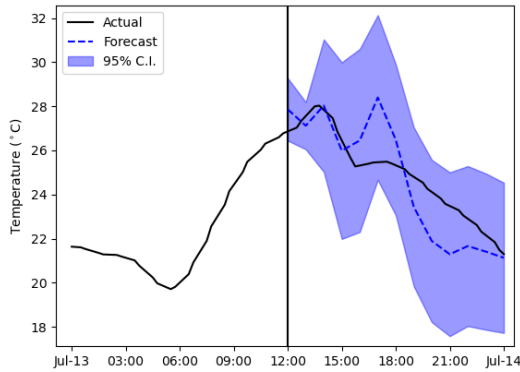


Figure 4: Temperature forecast and 95% confidence interval at noon on one of the simulation days

StockTM stochastic occupancy generator for a typical single-family home with a single occupant in Denver, CO [44]. The schedule has 1-minute resolution profiles for occupancy, power usage for uncontrollable loads, and hot water draws for appliances and fixtures. The EV schedule was generated from OCHRE using parking event data from EVI-Pro [45]. Note that because the EV schedule came from a different data source, the occupancy and EV schedules do not directly match.

The schedule was run using OCHRE for a full year with typical meteorological year weather data to get an annual profile of key variables to be used for forecasts. Key variables include driving status (for p_{drive} and p_{leave}), total hot water draws, total uncontrollable load power, and internal heat gains. Uncontrollable power includes appliances, lighting, and other electrical equipment that are not modeled in the MPC framework. Internal heat gains incorporate heat gains from occupants, equipment, radiation from windows, and infiltration.

Three types of forecasts were generated from the annual simulation data. All forecasts use a time resolution of 30 minutes and define the uncertainty of the random variable by assuming a Gaussian distribution at each time step. An exact forecast directly uses the results of the simulation as the mean and sets the standard deviation of the variable to zero.

A second forecast uses a backward-looking 30-day moving horizon to estimate random variables based on the time of day. For example, the water draw at noon on July 1 is estimated using the water draw values at noon from the month of June. The sample mean and sample variance are used to define the random variable distribution at that time.

A third forecast uses an autoregressive integrated

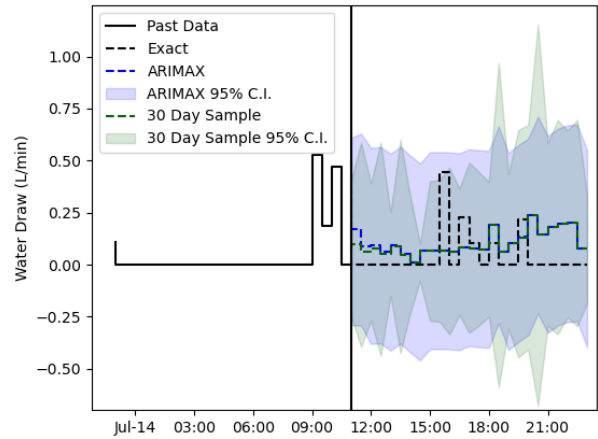


Figure 5: Water draw forecasts and confidence intervals at 11AM on one of the simulation days

moving average exogenous (ARIMAX) filter to estimate random variable distributions. The 30-day horizon forecast was used as the exogenous component. The ARIMAX filter uses an order of (2, 1, 1) for autoregressive, difference, and moving average parameters [46].

It is very difficult to measure internal gains in a real environment, and it would be impossible to use an ARIMAX forecast without data measurements. For this reason, an ARIMAX forecast for internal gains is not calculated; instead, the 30-day horizon forecast is used for the ARIMAX scenarios.

Figures 5 and 6 show the three forecast options at a single simulation time step for the water draw and EV driving variables, respectively. The figures show the mean and the 95% confidence interval for the 30-day horizon and ARIMAX forecasts. The 30-day moving horizon forecasts have negligible bias. The average standard deviations are 0.2 L/min for the water draw forecast and 0.35 for the EV state, which are both high relative to the average actual values. In the near term, the ARIMAX forecast does a better job of estimating the EV driving state and reducing its uncertainty by using current and past data in its prediction; however, the two long-term estimates perform about the same, with the ARIMAX predicting larger uncertainty. The ARIMAX method also does not significantly impact the water draw forecast, which tends to be more stochastic.

6. Simulation Results and Discussion

Simulations were run with multiple combinations of the three control types described in Section 4 and the three forecast types described in Section 5. First, the impacts of the model reduction and linearization on the

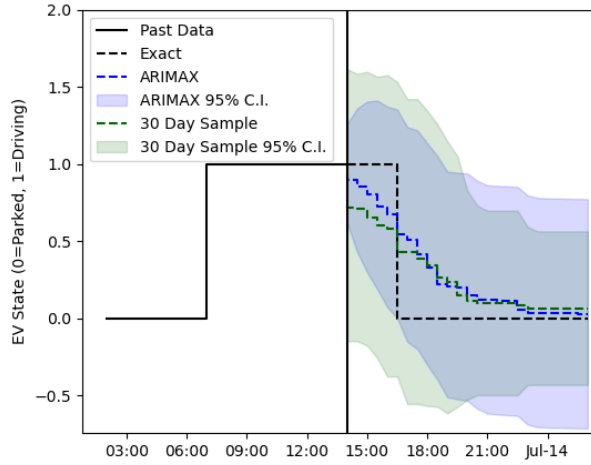


Figure 6: EV driving forecasts and confidence intervals at 2PM on one of the simulation days

Table 1: Key modeling and control parameter values

| Parameter | Value | Parameter | Value |
|---------------------------------|----------|------------------------|-----------------|
| ΔT_{op} | 10°C | n_h | 48 |
| $T_{op,rad}$ | 20°C | λ_{power} | \$0.10–0.27/kWh |
| η_{hvac} | 3.8 | λ_{hvac} | \$0.125/°C-hr |
| $(\theta_1 - \theta_{wm})_{op}$ | 33.1°C | λ_{wh} | \$0.01/°C-L |
| η_{wh} | 4.5 | λ_{ev} | \$1/kWh |
| P_{drive} | 0.65 kW | λ_{peak} | 10^{-5} /kW |
| E_{ev} | 32.5 kWh | T_m | 22.7°C |
| E_{batt} | 12 kWh | θ_1 | 46.1°C |
| | | ϵ_{x_i} (all) | 20% |

envelope model are shown. Next, the heuristic control and DMPC strategies are compared. The DMPC and SMPC are then compared under different levels of forecast uncertainty, and optimal control considerations in HEMS applications are discussed.

All simulations were run with the same house model, equipment characteristics, weather data, and occupancy schedules. Scenarios without MPC were run for 14 days in July, and scenarios with MPC were run for 13 days. Key simulation parameters are shown in Table 1. The electricity price λ_{power} was taken from Xcel Energy Colorado’s summer time-of-use rate [47], which includes an off-peak rate of \$0.10/kWh, a shoulder period from 1–3 PM at \$0.19/kWh, and a peak period from 3–7 PM at \$0.27/kWh.

The values of the cost coefficients associated with occupant comfort were estimated using assumptions for a person’s willingness to pay for the desired comfort metric. For example, the water heater cost coefficient assumes a \$5 cost for a 50 L shower at 10°C colder than

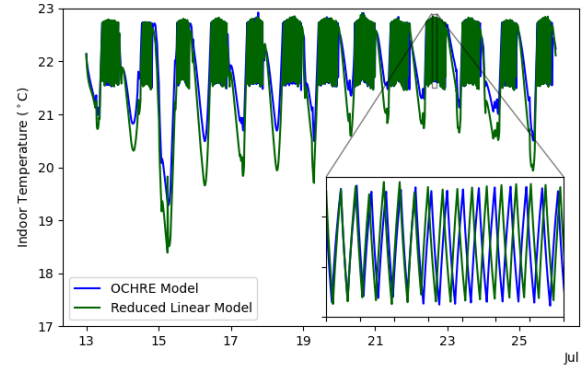


Figure 7: Simulated indoor temperature with the full and reduced-order linearized models

the minimum desired temperature.

6.1. Comparison of Linear Models

The impacts of the model reduction and linearization of radiation and infiltration are shown in figures 7 and 8. Figure 7 shows the difference in the modeled indoor air temperature between the full OCHRE model and the linearized model described in Section 3. In both models, the air conditioner turns on around midday and cycles according to the deadband control with a constant set point. At night, the air conditioner turns off, and the temperature floats below the set point. During a 10- to 15-hour window each night, the maximum temperature deviation between the simulations is less than 1°C.

Figure 8 compares the air-conditioning energy consumption during the peak period under different modeling conditions to assess the impacts of linearization and model reduction. Three models—OCHRE, a full-order linear model, and the reduced-order linear model—were run with a constant HVAC set point and the heuristic precooling control. The linear models overestimate the HVAC consumption compared to the OCHRE model, and they underestimate the energy savings due to precooling controls. The reduced linear model underestimates the difference by 14%, indicating that the MPC might undervalue precooling controls.

6.2. Deterministic MPC

We next compare the control performance between the baseline heuristic controls and the DMPC with an exact forecast. Although this DMPC case does not include forecast or measurement uncertainty, there is still some model uncertainty due to the model linearization and the difference in time resolution between OCHRE and the control model.




Figure 8: Cooling energy consumption during on-peak hours using a constant thermostat set point and a pre-cooling control. The difference in on-peak energy consumption between the two controls is shown at the right.

Figure 9 shows the state of each controllable equipment and the total house power for these two cases for a subset of the simulation time. The DMPC reduces total house power during the shoulder and peak periods. Both controllers adjust the HVAC set point before the peak period, but the air temperature is very similar for the rest of the day. The baseline heuristic control maintains a constant hot water set point, whereas the DMPC keeps the set point very low except when large water draws are expected during the peak period. The DMPC EV controller delays charging until after the peak period and slowly reaches full SOC right before the EV leaves. It also charges the battery more slowly than the scheduled heuristic controller.

Note that the heat pump water heater temperature stays very close to the set point temperature in the DMPC control because the deadband size reduces as the set point approaches the minimum comfort temperature. Although this helps achieve low energy costs and low water discomfort, it leads to a lot of equipment cycling, which can reduce the lifetime of the equipment.

Imperfect forecasts reduce the DMPC performance, particularly for the water heater controls. Figure 10 shows the difference in water heater results from the DMPC cases with an exact forecast and the ARIMAX forecast. The ARIMAX case tends to have a lower set point temperature because the forecast expects more frequent but smaller water draws. As shown in Table 2, the forecast error leads to more discomfort due to low water temperatures, but it also slightly reduces the water heater energy cost. Both controls cycle the water heater frequently when no draws are expected.

The DMPC forecasts also impact the EV controls as shown in Figure 11. The baseline controller always charges the EV as fast a possible, and the DMPC with an exact forecast always charges the EV as slow as pos-

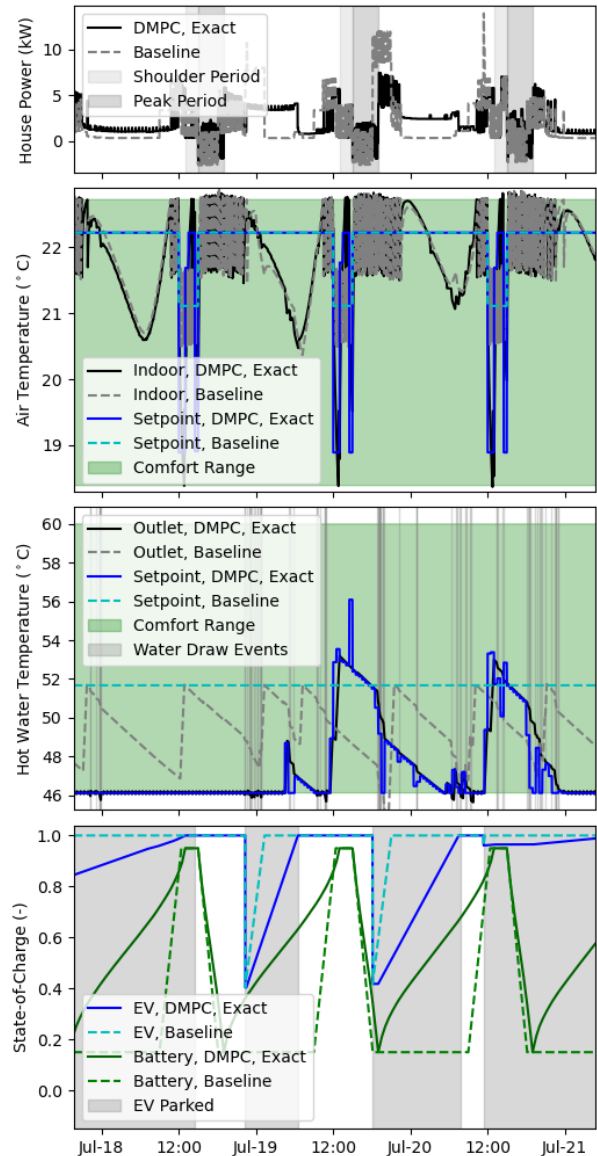


Figure 9: Simulation results for the baseline heuristic control case and the DMPC case with exact forecasts

sible while achieving full charge before the EV leaves. The DMPC with the ARIMAX forecast often charges in between these two cases, delaying charging during the peak period but charging more quickly after that. We note that the energy cost due to EV charging was identical for all MPC cases.

6.3. Stochastic MPC

Compared to the DMPC cases with imperfect forecasts, the SMPC is able to reduce hot water comfort costs as well as water heater cycling. Figure 12 shows

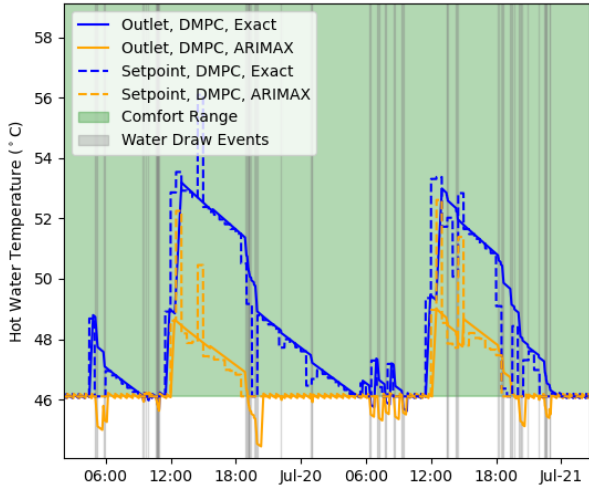


Figure 10: Water heater simulation results for the DMPC cases with exact and ARIMAX forecasts

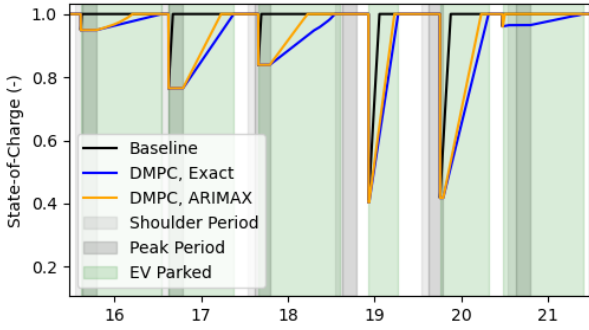


Figure 11: EV SOC for the baseline heuristic case and the DMPC cases with exact and ARIMAX forecasts

the water temperature for the DMPC and SMPC cases with the ARIMAX forecast. The SMPC case maintains a higher set point most of the time, leading to less water discomfort, less water heater cycling, and slightly higher energy costs. The SMPC water heater set point depends on the distribution of water draw volume for the next MPC time step. The set point temperature increases when either the mean value increases or when the variance increases. The set point also increases significantly before the peak demand periods.

The SMPC led to small differences in the HVAC controls as shown in Figure 13. Although both the DMPC and SMPC controls precool the space before the peak period, the SMPC case reduces the cooling set point by approximately 0.1°C for the rest of the day. This has little impact on the nighttime air temperature, but it does reduce the temperature while the air conditioner is on. This leads to a small increase in energy costs and a very

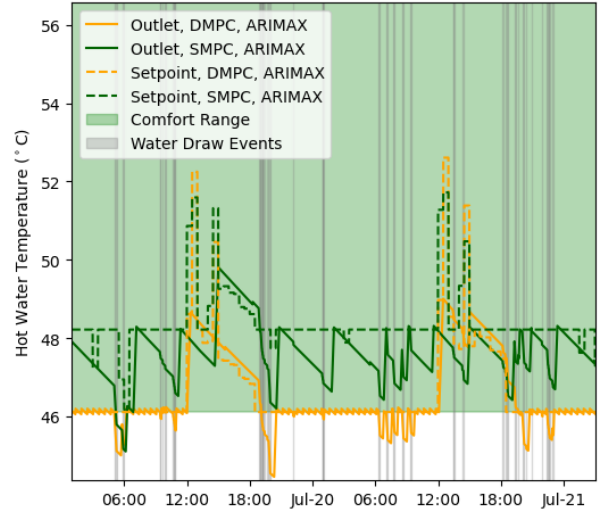


Figure 12: Water heater simulation results for the DMPC and SMPC cases with ARIMAX forecasts

Table 2: Control performance results for each case, including energy costs (J_{power}), water heater comfort costs (J_{wh}), and water heater cycles

| Control Type | Forecast Type | J_{power} (\$/mo) | J_{wh} (\$/mo) | WH Cycles (1/day) |
|--------------|---------------|---------------------|------------------|-------------------|
| Heuristic | N/A | 119.4 | 1.1 | 2.0 |
| DMPC | Exact | 94.6 | 2.4 | 27.9 |
| DMPC | 30-day | 94.2 | 9.2 | 24.5 |
| DMPC | ARIMAX | 94.2 | 7.6 | 22.3 |
| SMPC | 30-day | 96.9 | 1.4 | 4.0 |
| SMPC | ARIMAX | 97.0 | 1.0 | 4.2 |

small decrease in air comfort as quantified by the cost J_{hvac} .

We note that there is no change in the DMPC and SMPC controls for the EV and battery. In theory, the SMPC EV controls should be more conservative and might charge the EV sooner to ensure a low EV inconvenience cost J_{ev} ; however, this behavior was not shown in the simulation results.

6.4. Control Performance

The control performance for all the simulations is shown in Table 2. The energy cost J_{power} and water heater comfort cost J_{wh} were the only costs from the MPC objective that significantly contributed to the overall cost. All other comfort costs were exactly zero except for J_{hvac} , which was no larger than \$0.04/mo for some scenarios due to small temperature deviations that exceeded the comfort range. Costs were calculated using the power and states from the system model in

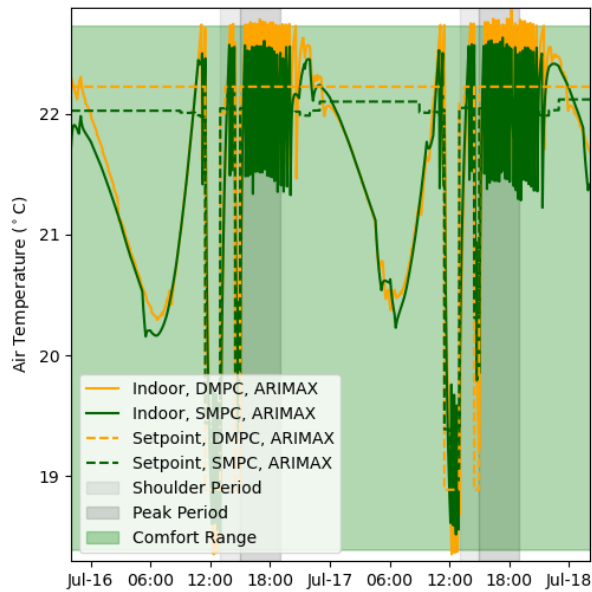


Figure 13: HVAC simulation results for the DMPC and SMPC cases with ARIMAX forecasts

OCHRE. The costs were aggregated over the full simulation time of 13 days and then scaled to show monthly costs. The average number of water heating (WH) heat pump cycles per day are also reported from OCHRE. We note that the electric resistance element was not turned on in any of the cases.

As expected, the DMPC with the exact forecast reduces the total cost the most when considering only the objective function costs and not the water heater cycles. All MPC cases have a smaller energy cost than the heuristic case, primarily because of the delayed EV charging after the peak period. The 30-day horizon forecasts perform very similarly to the ARIMAX forecasts and had very similar control strategies.

The SMPC cases outperform all other cases excluding the case with an exact forecast. Compared to the heuristic control, they significantly reduce the energy cost and slightly increase the number of water heater cycles. Compared to the DMPC cases, they have slightly higher energy costs and much lower water discomfort costs. The SMPC also had much less water heater cycling, which can reduce the equipment lifetime; while water heater cycling is not considered in the control objective, the results show this additional benefit of the SMPC.

The SMPC control had the largest impact on the device with the highest uncertainty in the control variable: the water heater. Due to difficulty in forecasting water draws, the water heater model inputs and outputs had

significant uncertainty, which caused a large difference between the DMPC and SMPC set points. The improvements in the SMPC performance highlights the importance of stochastic controls, especially in systems with high levels of uncertainty.

There is an interesting parallel between the differences in set points between the DMPC and SMPC cases and the deadband of the heuristic control. The SMPC back-off magnitudes effectively act as a deadband that buffers the control variable from exceeding its constraints. In this sense, the SMPC algorithm combines the best features of the heuristic control and the DMPC control. In addition, the SMPC improves upon the heuristic control by varying the size of the “deadband” based on the uncertainty in the control variable. We also note that the back-off magnitudes are easily adjustable by updating the risk tolerance of exceeding a constraint ϵ_{x_i} .

7. Conclusions

In this paper, a framework is developed for simulating HEMS controllers that incorporates model and forecast uncertainty. The framework includes a linear, reduced-order building model that can be derived from white-box equipment and envelope parameters. Heuristic, DMPC, and SMPC strategies are evaluated, and we find that the SMPC algorithm performs best when using common forecasts methods that incorporate uncertainty. The framework can evaluate the performance of HEMS in real-world applications, which can help de-risk HEMS deployment.

The simulation results are limited to a single test case and might not be generalizable for cases with different equipment parameters, weather and occupant schedules, or objective functions; however, SMPC is shown to incorporate the features of both MPC and heuristic controls and is likely to provide benefits for systems with high levels of uncertainty.

There are many extensions to the modeling and control methods developed in this paper. More sophisticated forecasts are likely to reduce uncertainty and improve control performance. Many advanced control features could be incorporated, for example, quadratic or mixed-integer programming, parameter-varying models, or asynchronous controls. We also suggest further study into a combination of SMPC and heuristics to optimally set the deadband size for devices with equipment cycling.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was authored by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by the U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Solar Energy Technologies Office. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

References

- [1] M. Blonsky, A. Nagarajan, S. Ghosh, K. McKenna, S. Veda, B. Kroposki, Potential Impacts of Transportation and Building Electrification on the Grid: A Review of Electrification Projections and Their Effects on Grid Infrastructure, Operation, and Planning, *Current Sustainable/Renewable Energy Reports* 6 (4) (2019) 169–176. doi:10.1007/s40518-019-00140-5.
- [2] R. Hledik, A. Faruqui, T. Lee, J. Higham, The National Potential for Load Flexibility, Tech. rep., The Brattle Group (2019).
- [3] M. Beaudin, H. Zareipour, Home energy management systems: A review of modelling and complexity, *Renewable and Sustainable Energy Reviews* 45 (2015) 318–335. doi:10.1016/j.rser.2015.01.046.
- [4] A. Roth, Grid-interactive Efficient Buildings Technical Report Series: Whole-Building Controls, Sensors, Modeling, and Analytics, Tech. rep., US Department of Energy, Building Technologies Office (2019).
- [5] X. Jin, K. Baker, D. Christensen, S. Isley, Foresee: A user-centric home energy management system for energy efficiency and demand response, *Applied Energy* 205 (2017) 1583–1595. doi:10.1016/j.apenergy.2017.08.166.
- [6] Z. Wu, S. Zhou, J. Li, X. P. Zhang, Real-time scheduling of residential appliances via conditional risk-at-value, *IEEE Transactions on Smart Grid* 5 (3) (2014) 1282–1291. doi:10.1109/TSG.2014.2304961.
- [7] S. Paul, N. P. Padhy, Resilient Scheduling Portfolio of Residential Devices and Plug-In Electric Vehicle by Minimizing Conditional Value at Risk, *IEEE Transactions on Industrial Informatics* 15 (3) (2019) 1566–1578. doi:10.1109/TII.2018.2847742.
- [8] K. Garifi, K. Baker, D. Christensen, B. Touri, Stochastic Home Energy Management Systems with Varying Controllable Resources, *IEEE Power and Energy Society General Meeting 2019-Augus*. doi:10.1109/PESGM40551.2019.8973708.
- [9] T. A. N. Heirung, J. A. Paulson, J. O’Leary, A. Mesbah, Stochastic model predictive control — how does it work?, *Computers and Chemical Engineering* 114 (2018) 158–170. doi:10.1016/j.compchemeng.2017.10.026.
- [10] M. Maasoumy, M. Razmara, M. Shahbakhti, A. S. Vincentelli, Handling model uncertainty in model predictive control for energy efficient buildings, *Energy and Buildings* 77 (2014) 377–392. doi:10.1016/j.enbuild.2014.03.057.
- [11] H. Wu, A. Pratt, S. Chakraborty, Stochastic optimal scheduling of residential appliances with renewable energy sources, *IEEE Power and Energy Society General Meeting 2015-Sept*. doi:10.1109/PESGM.2015.7286584.
- [12] T. H. Pedersen, S. Petersen, Investigating the performance of scenario-based model predictive control of space heating in residential buildings, *Journal of Building Performance Simulation* 11 (4) (2018) 485–498. doi:10.1080/19401493.2017.1397196.
- [13] M. Yousefi, N. Kianpoor, A. Hajizadeh, M. Soltani, Smart Energy Management System for Residential Homes Regarding Uncertainties of Photovoltaic Array and Plug-in Electric Vehicle, in: *IEEE International Symposium on Industrial Electronics*, Vol. 2019-June, Institute of Electrical and Electronics Engineers Inc., 2019, pp. 2201–2206. doi:10.1109/ISIE.2019.8781471.
- [14] D. Thomas, O. Deblecker, C. S. Ioakimidis, Optimal operation of an energy management system for a grid-connected smart building considering photovoltaics’ uncertainty and stochastic electric vehicles’ driving schedule, *Applied Energy* 210 (2018) 1188–1206. doi:10.1016/j.apenergy.2017.07.035.
- [15] F. Oldewurtel, A. Parisio, C. N. Jones, D. Gyalistras, M. Gwender, V. Stauch, B. Lehmann, M. Morari, Use of model predictive control and weather forecasts for energy efficient building climate control, *Energy and Buildings* 45 (2012) 15–27. doi:10.1016/j.enbuild.2011.09.022.
- [16] D. Azuatalam, K. Paridari, Y. Ma, M. Förstl, A. C. Chapman, G. Verbič, Energy management of small-scale PV-battery systems: A systematic review considering practical implementation, computational requirements, quality of input data and battery degradation, *Renewable and Sustainable Energy Reviews* 112 (2019) 555–570. doi:10.1016/j.rser.2019.06.007.
- [17] Y. Iwafune, T. Ikegami, J. G. D. S. Fonseca, T. Oozeki, K. Ogimoto, Cooperative home energy management using batteries for a photovoltaic system considering the diversity of households, *Energy Conversion and Management* 96 (2015) 322–329. doi:10.1016/j.enconman.2015.02.083.
- [18] Y. Ma, J. Matusko, F. Borrelli, Stochastic model predictive control for building HVAC systems: Complexity and conservatism, *IEEE Transactions on Control Systems Technology* 23 (1) (2015) 101–116. doi:10.1109/TCST.2014.2313736.
- [19] X. Zhang, S. Grammatico, G. Schildbach, P. Goulart, J. Lygeros, On the sample size of randomized MPC for chance-constrained systems with application to building climate control, in: *2014 European Control Conference, ECC 2014*, Institute of Electrical and Electronics Engineers Inc., 2014, pp. 478–483. doi:10.1109/ECC.2014.6862498.
- [20] M. Maasoumy, M. Razmara, M. Shahbakhti, A. Sangiovanni Vincentelli, Selecting building predictive control based on model uncertainty, in: *Proceedings of the American Control Conference*, Institute of Electrical and Electronics Engineers Inc., 2014, pp. 404–411. doi:10.1109/ACC.2014.6858875.
- [21] M. Rahmani-Andebili, H. Shen, Energy scheduling for a smart home applying stochastic model predictive control, in: *2016 25th International Conference on Computer Communications and Networks, ICCCN 2016*, Institute of Electrical and Electronics Engineers Inc., 2016. doi:10.1109/ICCCN.2016.7568516.
- [22] M. Hu, F. Xiao, Quantifying uncertainty in the aggregate energy flexibility of high-rise residential building clusters considering stochastic occupancy and occupant behavior, *Energy* 194 (2020) 116838. doi:10.1016/j.energy.2019.116838.
- [23] K. Garifi, K. Baker, B. Touri, D. Christensen, Stochastic Model Predictive Control for Demand Response in a Home Energy Management System, *IEEE Power and Energy Society General Meeting 2018-Augus*. doi:10.1109/PESGM.2018.8586485.
- [24] M. Blonsky, P. Munankarmi, S. P. Balamurugan, Incorporating residential smart electric vehicle charging in home energy management systems, *IEEE Green Technologies Conference 2021-April (2021)* 187–194. doi:10.1109/GREENTECH48523.2021.00039.
- [25] K. Doubleday, V. Van Scyoc Hernandez, B. M. Hodge, Benchmark probabilistic solar forecasts: Characteristics and recommendations, *Solar Energy* 206 (2020) 52–67. doi:10.1016/j.solener.2020.05.051.
- [26] W. Jung, F. Jazizadeh, Human-in-the-loop HVAC operations: A quantitative review on occupancy, comfort, and energy-efficiency dimensions (apr 2019). doi:10.1016/j.apenergy.2019.01.070.
- [27] L. Serafini, E. Principi, S. Spinsante, S. Squartini, Multi-Household Energy Management in a Smart Neighborhood in the Presence of Uncertainties and Electric Vehicles, *Electronics* 2021, Vol. 10, Page 3186 10 (24) (2021) 3186. doi:10.3390/ELECTRONICS10243186.
- [28] A. Afram, F. Janabi-Sharifi, Theory and applications of HVAC control systems - A review of model predictive control (MPC) (feb 2014). doi:10.1016/j.buildenv.2013.11.016.
- [29] B. Dong, K. P. Lam, A real-time model predictive control for building heating and cooling systems based on the occupancy

- behavior pattern detection and local weather forecasting, *Building Simulation* 7 (1) (2014) 89–106. doi:10.1007/s12273-013-0142-7.
- [30] S. R. Cominesi, M. Farina, L. Giulioni, B. Picasso, R. Scattolini, A Two-Layer Stochastic Model Predictive Control Scheme for Microgrids, *IEEE Transactions on Control Systems Technology* 26 (1) (2018) 1–13. doi:10.1109/TCST.2017.2657606.
- [31] M. Blonsky, K. McKenna, T. Vincent, A. Nagarajan, Time-of-use and Demand Charge Battery Controller using Stochastic Model Predictive Control, *IEEE SmartGridComm* (November). doi:10.1109/SmartGridComm47815.2020.9302943.
- [32] P. Munankarmi, J. Maguire, S. P. Balamurugan, M. Blonsky, D. Roberts, X. Jin, Community-scale interaction of energy efficiency and demand flexibility in residential buildings, *Applied Energy* 298 (2021) 117149. doi:10.1016/j.apenergy.2021.117149.
- [33] A. Roth, J. Reyna, Innovations in Building Energy Modeling Research and Development Opportunities Report for Emerging Technologies Disclaimer, Tech. rep., US Department of Energy, Building Technologies Office (2016).
- [34] E. McKenna, M. Thomson, High-resolution stochastic integrated thermal-electrical domestic demand model, *Applied Energy* 165 (2016) 445–461. doi:10.1016/j.apenergy.2015.12.089.
- [35] M. Blonsky, J. Maguire, K. McKenna, D. Cutler, S. P. Balamurugan, X. Jin, OCHRE: The Object-oriented, Controllable, High-resolution Residential Energy Model for Dynamic Integration Studies, *Applied Energy* 290 (2021) 116732. doi:10.1016/j.apenergy.2021.116732.
- [36] M. Maasoumy, B. Moridian, M. Razmara, M. Shahbakhti, A. Sangiovanni-Vincentelli, Online simultaneous state estimation and parameter adaptation for building predictive control, in: *ASME 2013 Dynamic Systems and Control Conference, DSCC 2013, Vol. 2, American Society of Mechanical Engineers (ASME)*, 2013. doi:10.1115/DSCC2013-4064.
- [37] M. Sofos, J. Langevin, M. Deru, E. Gupta, K. S. Benne, D. Blum, T. Bohn, R. Fares, N. Fernandez, G. Fink, S. Frank, J. Gerbi, J. Granderson, D. Hoffmeyer, T. Hong, A. Jiron, S. Johnson, S. Katipamula, T. Kuruganti, J. Langevin, W. C. Livingood, R. Muehleisen, M. Neukomm, V. Nubbe, P. Phelan, M. Piette, J. Reyna, A. Roth, A. Satre-Meloy, M. Specian, D. Vrabie, M. Wetter, S. Widergren, *Innovations in Sensors and Controls for Building Energy Management: Research and Development Opportunities Report for Emerging Technologies*, Tech. rep., U.S. DOE EERE (2020). doi:10.2172/1601591. URL <http://www.osti.gov/servlets/purl/1601591/>
- [38] A. J. Laub, M. T. Heath, C. C. Paige, R. C. Ward, Computation of System Balancing Transformations and Other Applications of Simultaneous Diagonalization Algorithms, *IEEE Transactions on Automatic Control* 32 (2) (1987) 115–122. doi:10.1109/TAC.1987.1104549.
- [39] R. E. Kalman, A new approach to linear filtering and prediction problems, *Journal of Fluids Engineering, Transactions of the ASME* 82 (1) (1960) 35–45. doi:10.1115/1.3662552.
- [40] A. Mirakhorli, B. Dong, Model predictive control for building loads connected with a residential distribution grid, *Applied Energy* 230 (2018) 627–642. doi:10.1016/j.apenergy.2018.08.051.
- [41] OpenWeatherMap One Call API. URL <https://openweathermap.org/api/one-call-api>
- [42] PVlib Documentation. URL https://pvlib-python.readthedocs.io/en/stable/generated/pvlib.forecast.ForecastModel.cloud_cover_to_irradiance.html
- [43] D. P. Larson, L. Nonnenmacher, C. F. Coimbra, Day-ahead forecasting of solar power output from photovoltaic plants in the American Southwest, *Renewable Energy* 91 (2016) 11–20. doi:10.1016/J.RENENE.2016.01.039.
- [44] J. Chen, R. Adhikari, E. Wilson, J. Robertson, A. Fontanini, B. Polly, O. Olawale, Stochastic simulation of residential building occupant-driven energy use in a bottom-up model of the U.S. housing stock, Arxiv. URL <https://arxiv.org/abs/2111.01881v2>
- [45] Alternative Fuels Data Center: Electric Vehicle Infrastructure Projection Tool (EVI-Pro) Lite Assumptions and Methodology (2020). URL <https://afdc.energy.gov/evi-pro-lite/load-profile/assumptions>
- [46] S. Seabold, J. Perktold, Statsmodels: Econometric and Statistical Modeling with Python, in: *Proc. of the 9th Python in Science Conf, no. Scipy, 2010*, pp. 92–96. URL <http://statsmodels.sourceforge.net/>
- [47] Time of Use — Xcel Energy (2022). URL <https://co.my.xcelenergy.com/s/billing-payment/residential-rates/time-of-use-pricing>