

Data-driven Demand Response Modeling and Control of Buildings with Gaussian Processes

(Full paper)

Truong X. Nghiem and Colin N. Jones

Abstract—This paper presents an approach to provide demand response services with buildings. Each building receives a normalized signal that tells it to increase or decrease its power demand, and the building is free to implement any suitable strategy to follow the command, most likely by changing some of its setpoints. Due to this freedom, the proposed approach lowers the barrier for any buildings equipped with a reasonably functional building management system to participate in the scheme. The response of the buildings to the control signal is modeled by a Gaussian Process, which can predict the power demand of the buildings and also provide a measure of its confidence in the prediction. A battery is included in the system to compensate for this uncertainty and improve the demand response performance of the system. A model predictive controller is developed to optimally control the buildings and the battery, while ensuring their operational constraints with high probability. Our approach is validated by realistic co-simulations between Matlab and the building energy simulator EnergyPlus.

I. INTRODUCTION

Electric grids are undergoing a substantial change in scale and complexity. The growing addition of renewable energy sources and distributed storage systems, and the mass adoption of electric vehicles will all contribute to the significant increase in complexity and volatility of future grids. Demand response (DR), or demand side management (DSM) in general, represents a significant but still underutilized resource for supplying reliability services to the grid. A wide range of DR programs and approaches have been offered by utilities or investigated by researchers [1].

We are interested in a demand tracking control problem in which an energy customer adjusts its power demand in real time to track a reference demand signal, often for a limited window of time. Let r_t denote the reference demand at time step $t \geq 0$. Here, time is discretized in equal intervals of length $T > 0$, and the power demand is averaged over each interval. The sampling time T may vary between application and customer types, from seconds in fast Ancillary Services to minutes in slow DR with commercial buildings [2].

Consider an aggregation of commercial buildings as an energy customer to track a reference demand signal. This is generally achieved by controlling the heating, ventilating,

and air-conditioning (HVAC) systems of the buildings as the HVAC system is often the largest electricity load in a commercial building. However, other building systems such as the lighting system can also be manipulated to adjust the whole building power demand. A common approach adopted by many research works is to identify a detailed mathematical model of the building's HVAC system and the building's thermal model, then use an optimization-based method to directly control certain field devices, such as dampers, pumps, or fans [3]. While these methods can achieve the best performance, they are expensive for several reasons: (a) they require a major retrofit of the building to install sensors, actuators, and other equipment; (b) they require upgrading the building energy management system (BEMS) to collect data and implement new control logics, which will incur significant cost; (c) they require detailed modeling of the building's thermal dynamics and HVAC system, which is time consuming and costly [4]. For these reasons, their benefits may not be attractive enough considering the cost and effort of deployment, especially for existing commercial buildings. Consequently their penetration into practical, real-world applications is still at a low level.

A much simpler, hence more practical, approach to DR with buildings is by global setpoint change rules, i.e., certain global setpoints in the HVAC system are adjusted according to predefined rules to change the building's power demand. Such rules can include setbacks for thermostat setpoints, the supply air temperature setpoint and the chilled water supply temperature setpoint, and also fixed dimming levels of lights. Specifically, in the "Global Temperature Adjustment" (GTA) strategy, the thermostats of the building's zones are governed by one or a few global setpoints, which are manipulated to receive the change in power demand of the HVAC system [5]. The simplicity of this approach makes it very cost-effective and practical, hence it is widely adopted in practice. This approach can be enhanced by learning a model from the setpoint change to the building's power demand, then using this model to devise good setpoint values. This requires either an upgrade of the BEMS for learning and optimization software, or services from an external provider. In the latter case, the fact that an external party has access to internal data of the BEMS, for model identification and operational purposes, and can impose direct setpoint changes in the system might cause reluctance of the building's operator.

We propose a simpler approach to overcome these practicality issues. An artificial DR control signal $u_t \in [-1, 1]$ is

The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 307608: BuildNet, and the Swiss National Science Foundation under the GEMS project (Green Energy Management of Structures, grant n. 200021 137985).

The authors are with the Automatic Control Laboratory, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland.




Fig. 1. Architecture of the demand tracking system. In blue are control signals, while in black are physical or measurement signals. Each building implements its own DR rule, hidden from the controller, to respond to the DR control signal u .

sent to the BEMS at each time step t , which requests it to relatively increase (when $u_t > 0$) or decrease (when $u_t < 0$) its power demand, or to keep the nominal operational setting (when $u_t = 0$). The maximum upward and downward power demand adjustments are expected to be attained at $u_t = 1$ and $u_t = -1$ respectively. Each building is free to implement any suitable rule to respond to the request, e.g., by adjusting its global temperature setpoint proportionally to u_t (GTA strategy). *These rules are assumed to be designed to satisfy all the operational constraints of the building, e.g., maintaining the occupant comfort, while reacting appropriately to u_t .* An aggregator learns a *data-driven* Gaussian Process (GP) model of how the buildings respond to the DR control signal. A GP model is a non-parametric statistical model, which can predict the building power demand in response to u_t and also provide a measure of its confidence in the prediction. This model is used to calculate the DR control signal sent to the buildings in real time, in order to track the reference demand signal. *We note that the DR rules implemented at the building to respond to the DR control signal are completely unknown to the aggregator.* To improve the tracking performance, an auxiliary battery is employed and controlled by the aggregator to compensate for prediction uncertainty of the model. The overall system architecture is depicted in Figure 1.

Advantages: The proposed scheme will be simple to deploy since it is based on the rule-based DR and does not require specialized software in the BEMS except for a communication channel to receive the DR control signal. Furthermore, the aggregator does not need access to any internal data of the BEMS, nor direct change of any setpoints. This ensures data confidentiality and operational safety of the building system, which are much preferred by the building's operator. Finally, since each building's energy system and requirements are unique, the abstract interface between the buildings and the aggregator by a simple DR control signal removes the need for customizing the aggregator's design to each and every building. Therefore, the aggregation, learning, and control algorithms can be applied to a wide range of building types without modifications.

Related Work: A comprehensive review of building modeling and control methods can be found in [3] and the references therein. A broad range of data-driven modeling, assessment, and control methods for DR with buildings have

been investigated in the literature. Regression trees were used in [6] to model and compute setpoint schedules of buildings for DR. Regression and GP models were investigated for forecasting short- and long-term building energy consumption in [7]. Simulation studies with EnergyPlus and regression models were used in [8] to quantify the flexibility of buildings for DR using setpoint change rules. The authors of [9], after comparing four data-driven inverse modeling methods for building energy predictions, concluded that the Gaussian approaches were accurate and highly flexible, and the uncertainty measures could be helpful for certain applications involving risks.

The rest of this paper is organized as follows. In Section II, a data-driven modeling method for predicting building power demand based on Gaussian Processes will be presented. Section III develops a model predictive control formulation for the demand tracking problem using the Gaussian Process model. The proposed scheme is illustrated and validated by a detailed simulation case study in Section IV.

II. GAUSSIAN PROCESSES FOR DEMAND PREDICTION

A. Introduction to Gaussian Processes

This section briefly introduces Gaussian Process (GP) modeling with applications in control. For more details, we refer the reader to [10] for general GP discussions and [11] for applications of GPs in dynamic systems.

A GP can be viewed as an extension of the multivariate Gaussian distribution to functions. It is a collection of random variables, any finite number of which have a joint Gaussian distribution. Specifically, consider noisy observations y of an underlying function $f : \mathbb{R}^n \mapsto \mathbb{R}$ through a Gaussian noise model: $y = f(x) + \mathcal{N}(0, \sigma_n^2)$, $x \in \mathbb{R}^n$. Given the regression vectors $X = [x_1, \dots, x_N]$ and the corresponding observed outputs $Y = [y_1, \dots, y_N]^T$, we wish to find the distribution of the output y_* corresponding to a new input vector x_* . Assuming a GP structure $y \sim \mathcal{GP}(\mu, k; \theta)$, y is fully specified by its mean function $\mu(x)$ and covariance function $k(x, x')$,

$$\mu(x) = \mathbb{E}[f(x)]$$

$$k(x, x') = \mathbb{E}[(f(x) - \mu(x))(f(x') - \mu(x')))] + \sigma_n^2 \delta(x, x')$$

where $\delta(x, x')$ is the Kronecker delta function. The hyperparameter vector θ parameterizes the mean and covariance functions. The covariance matrix K is defined to have elements $K_{ij} = k(x_i, x_j)$.

The covariance function $k(x, x')$ indicates how correlated the outputs are at x and x' . Intuitively speaking, $k(\cdot, \cdot)$ relates nearby inputs (in a sense defined by the covariance function) so that they will give similar predictions. This is an advantage of GP modeling as it specifies the structure of the covariance matrix of the input variables rather than a fixed structural input-output relationship. It is therefore highly flexible and can capture complex behavior with fewer parameters. There exists a wide range of covariance functions and combinations to choose from [10].

Given the training data, the hyperparameters can be estimated by maximizing the likelihood: $\arg \max_{\theta} \Pr(Y|X, \theta)$.

Once the hyperparameters are determined, the GP can be used to predict the distribution of the output y_* corresponding to a new input vector x_* : $y_* \sim \mathcal{N}(\bar{y}_*, \sigma_*^2)$, where

$$\bar{y}_* = g_m(x_*) := \mu(x_*) + K_* K^{-1} (Y - \mu(X)) \quad (1a)$$

$$\sigma_*^2 = g_v(x_*) := K_{**} - K_* K^{-1} K_*^T \quad (1b)$$

$$K_* = [k(x_*, x_1), \dots, k(x_*, x_N)], K_{**} = k(x_*, x_*). \quad (1c)$$

A key advantage of GP regression is that it provides the predictive posterior distribution of the output instead of a point estimate. This posterior distribution carries the full information of the prediction, including for example its confidence level. While the predicted mean is often used as the best guess of the output, the full distribution can be used in a meaningful way. In our method, the variance of the predictive distribution is incorporated into the optimization to control the system with a quantifiable confidence guarantee.

1) *Gaussian Processes for Dynamical Systems*: As GPs can be used to model nonlinear functions, they are suitable for modeling dynamical systems. This is achieved by feeding delayed input and output signals back to the model as regressors [11]. In such cases, the model is said to be autoregressive, whose current output depends on its past inputs and outputs. Specifically, in control systems, it is common to use as the regressors of a dynamical GP

$$x_t = [y_{t-l_y}, \dots, y_{t-1}, u_{t-l_u}, \dots, u_t, w_{t-l_w}, \dots, w_{t-1}, w_t]$$

where t denotes the time step, u the control input, w the exogenous disturbance input, y the (past) output. Here, l_y , l_u , and l_w are respectively the lags for autoregressive outputs, control inputs, and disturbances. Note that u_t and w_t are the current control and disturbance inputs. The vector of all autoregressive inputs can be thought of as the current state of the model. A dynamical GP can then be trained from data in the same way as any other GPs.

When a GP is used for control or optimization, it is usually necessary to simulate the model over a finite number of future steps and predict its multistep-ahead behavior. Because the output of a GP is a distribution rather than a point estimate, the autoregressive outputs fed to the model beyond the first step are random variables, resulting in more and more complex output distributions as we go further. Therefore, a multistep simulation of a GP involves the propagation of uncertainty through the model. There exist several methods for uncertainty propagation in GPs [12], [11]. We mention here two simulation methods for autoregressive GPs.

- The *Monte-Carlo method* obtains samples of the output distribution under input uncertainty, which can be seen as a Gaussian mixture. This Gaussian mixture becomes more complex in later steps of the simulation, therefore efficient numerical algorithms must be implemented. This method can achieve good prediction accuracy at the expense of high computational load. It is also general, i.e., it can be used with any covariance functions.
- The *zero-variance method* does not propagate uncertainty. At each step, the autoregressive outputs are replaced by their corresponding expected values. Obviously, this method will underestimate the variances

of the output distributions. However, its computational simplicity is attractive, especially in optimization applications where the GP must be simulated for many times. In such cases, if the prediction error caused by not propagating uncertainty is insignificant, the zero-variance method can and should be used. For more detailed discussions on this topic, see [11], [12].

B. Data-driven Modeling of Building Power Demand with Gaussian Processes

As introduced in Section I, in our proposed scheme, each building implements its own rule to respond to the DR control signal u . Because the specific rules are unknown to the aggregator, we employ a data-driven GP regression approach to learn from experiment data a black-box model of the building power demand in response to the DR signal. The data-driven GP modeling approach consists of the following steps.

1) *Experiment and data collection*: In this step, data are collected from experiments with the actual buildings. A test DR signal, which is generated randomly or follows a predefined pattern, is sent to the buildings. Their power demands in response to the test signal are recorded, together with other measurements such as the climate condition (e.g., the ambient air temperature and humidity). These additional inputs represent the exogenous disturbances w to the model, while u will be the DR signal input and y the power demand output. Data points are associated with the time-stamps at which they are collected, which include the time of day, the day of week, and the full date. The time-stamps usually reveal the usage or load patterns of the buildings and therefore are important features of the model.

The experiment typically lasts several weeks to collect enough data. In general, more experiment data and a broader spectrum of the inputs, especially of the disturbances, can increase the prediction accuracy and confidence of the GP. *Because our control scheme does not directly change the building setpoints nor control its equipment but lets the building controller apply its own appropriate rules, the experiment will not adversely affect the building's operation or the comfort of its occupants.* The experiment may be repeated seasonally to adjust the model.

2) *Feature selection*: Because buildings are dynamical systems, an autoregressive GP is usually needed to model the response from the DR control signal to the power demand. Often we do not know a priori what the lag values are and which of the above inputs are relevant to predicting the output. *Feature selection* concerns with selecting the relevant inputs for the model. GPs provide an automatic tool for determining the relevance of each input, using *Automatic Relevance Determination* (ARD) covariance functions. An ARD covariance function has a length scale hyperparameter for each input, which indicates how much the input value must vary to significantly influence the output value. An input with a very large length scale is irrelevant to the model. Therefore, we can train a GP with all the inputs up to maximum lag values and with an ARD covariance

function, then remove the irrelevant inputs. This feature selection method will be demonstrated in the case study.

3) *Model selection, training and validation*: The mean function $\mu(x)$ is often chosen to be zero, though other options can be useful in certain cases [10]. Selecting the covariance function of the GP often requires domain insights and trial-and-error. Common choices are the Squared Exponential (SE) covariance function with ARD

$$k(x, x') = \sigma_f^2 \exp \left[-\frac{1}{2} \sum_{d=1}^D \left(\frac{x_d - x'_d}{\lambda_d} \right)^2 \right] \quad (2)$$

and the Rational Quadratic (RQ) covariance function

$$k(x, x') = \sigma_f^2 \left[1 + \frac{1}{2\alpha} \sum_{d=1}^D \left(\frac{x_d - x'_d}{\lambda_d} \right)^2 \right]^{-\alpha} \quad (3)$$

Here, α , σ_f and λ_d (length scales) are hyperparameters, and D is the input dimension. Periodic behaviors can be modeled with periodic covariance functions [10].

The experiment data collected in the first step are partitioned into two separate sets: a training data set and a testing data set. The GP's hyperparameters are tuned on the training data. The resulting GP is then validated on the testing data. Several covariance functions can be compared based on their validation results to select the best model. If the results are unsatisfactory, we may need to go back to an earlier step, for example to change the covariance function structure or to select a different set of inputs.

III. DEMAND TRACKING CONTROL

This section formulates a model predictive control (MPC) approach for the demand tracking problem. The control system consists of three main components shown in Figure 1:

- A collection of buildings which respond to a DR control signal by unknown but fixed rules, resulting in demand variations. Their responses are modeled by a GP.
- A battery whose state of charge (SoC) can be measured and whose charge/discharge power can be controlled. The battery helps improve the tracking quality by absorbing the output uncertainty of the GP.
- A controller which receives the demand tracking request from the utility and computes the DR signal to the buildings and the power of the battery to optimally track the reference demand signal.

We first present the models and notations used in this section, followed by the formulation of the controller. Illustrative simulation results will be discussed in Section IV.

A. Building's Demand Response Model

In Section II-B we discussed how a GP could be trained to learn the buildings' aggregate demand response to the DR signal from experiment data. The GP can predict the expected demand as well as provides probabilistic information about its prediction in the form of the variance. As we will show later, this additional information is invaluable in the tracking control method, particularly for controlling the battery.

We will use the subscript notation $\bullet_{\tau|t}$ to denote the value of a variable at time step τ given the current information available at t . It is a measured value if $\tau < t$, e.g.,

$y_{t-1|t}$ is the measured power demand in the previous time step. When $\tau \geq t$, the value is predicted, e.g., $y_{t|t}$ is the predicted power demand in the current time step. At τ , the input vector $x_{\tau|t}$ of the GP consists of the measured and predicted outputs $\mathbf{y}_{\tau|t} = [y_{\tau-l_y|t}, \dots, y_{\tau-1|t}]$, control inputs $\mathbf{u}_{\tau|t} = [u_{\tau-l_u|t}, \dots, u_{\tau|t}]$, and disturbances $\mathbf{w}_{\tau|t} = [w_{\tau-l_w|t}, \dots, w_{\tau|t}]$, as well as the time value τ :

$$x_{\tau|t} = [\mathbf{y}_{\tau|t}, \mathbf{u}_{\tau|t}, \mathbf{w}_{\tau|t}, \tau].$$

The GP regression (1) gives us the distribution of the power demand $y_{\tau|t} \sim \mathcal{N}(\bar{y}_{\tau|t}, \sigma_{y,\tau|t}^2)$ where

$$\bar{y}_{\tau|t} = g_m(x_{\tau|t}), \quad \sigma_{y,\tau|t}^2 = g_v(x_{\tau|t}). \quad (4)$$

Note that $x_{\tau|t}$ may contain predicted values. In particular:

- Control inputs $u_{k|t}$, for $k \geq t$, are decision variables.
- We assume that predicted disturbances $w_{k|t}$, for $k \geq t$, are available, e.g., from short-term forecasts.
- Predicted outputs $y_{k|t}$, for $k \geq t$, are random variables at previous steps. Here, we adopt the zero-variance simulation method for simplicity, i.e., $y_{k|t}$ take the predicted means $\bar{y}_{k|t}$, which are deterministic values. Uncertainty propagation can be incorporated for more accurate predictions as discussed in Section II-A, however the following derivations will be more complex.

As it is assumed that the DR rule of each building is designed to be compliant with its own operational requirements, there are no further constraints on the GP except that u is in the range $[-1, 1]$.

B. Battery's Model and Constraints

For simplicity, we assume an ideal lossless battery model

$$s_{t+1} = s_t + T b_t \quad (5)$$

where b_t is the battery's power during the time step t and s is the battery's SoC. Here, b is positive if the battery is charging and negative if discharging. The battery is subject to two operational constraints: its power must be bounded by $b_{\min} \leq b \leq b_{\max}$, and its SoC must stay in a safe range $[s_{\min}, s_{\max}]$ where s_{\max} is the fully-charged level and s_{\min} is the lowest safe discharged level.

C. Tracking Constraint

The link between the buildings and the battery is the tracking constraint, which states that their total power $p = y + b$ should track the reference power demand r . In this way, the battery serves two purposes:

- Because its power can be controlled precisely, it helps reject the uncertainty of the GP, i.e., it compensates for the discrepancy between the predicted power by the GP and the actual power.
- It acts as an energy buffer to increase the tracking capability of the controller.

Although ideally p_t should track r_t exactly at any time t , this strict constraint may be infeasible in certain circumstances, for example when r is outside the DR capability of the system. Therefore, we introduce a slack variable δ_t , which is the tracking error: $\delta_t = r_t - p_t$. The controller tries to keep $\delta = 0$, however when exact tracking is

TABLE I
CATEGORIZATION OF VARIABLES; $t \leq \tau \leq t + H - 1$; $k < t$.

Decision variables	Deterministic: $u_{\tau t}, \delta_{\tau t}$
Exogenous variables	Deterministic: $r_{\tau t}, w_{\tau t}$
Current states	Deterministic: y_k, u_k, w_k, s_t
Derived variables	Random: $y_{\tau t}, b_{\tau t}, s_{\tau+1 t}$

impossible, it will maintain the operational safety of the system while keeping the tracking errors as small as possible. This objective will be reflected in the cost function. In this formulation, we consider δ as a decision variable and write the battery power in terms of the reference r , the building power y , and the slack δ :

$$b_t = r_t - \delta_t - y_t. \quad (6)$$

This equation is made possible by the fact that the battery can be controlled precisely in real-time.

D. Model Predictive Tracking Control

We adopt the MPC approach [13] to solve the demand tracking problem. At the core of the MPC are the models of the system, which are used to predict future system states given the current state and current and forecast future inputs. Suppose the current time step is t and the MPC horizon is $H > 0$. To formulate the MPC, the equations predicting the values of the system's variables over the horizon must be derived, based on the system models. We first categorize all variables into: *decision variables*, which we aim to compute; *exogenous variables*, such as the disturbances; *current states*, i.e., available measurements at the current time t ; and *derived variables*, which are computed from the other variables. We also differentiate between deterministic and random variables. Table I summarizes this categorization. Note that for simplicity the exogenous variables, namely the future reference signal r and disturbances w , are assumed to be deterministic, meaning that we know their future values exactly, e.g., from accurate short-term forecasts. Uncertainties in these variables can be incorporated into the formulation as follows: uncertainty in r will increase the variance of b , and uncertainties in w will propagate to y through the GP.

Formulating the MPC then amounts to deriving the distributions of the derived random variables conditioned on the other variables. For example, we wish to derive the probability distribution of each $y_{\tau|t}$ given the current states and all decision and exogenous variables. For notational brevity, we will drop the conditioning part of the notations and simply write $\Pr(y_{\tau|t})$, with the implicit interpretation that it is a conditional probability.

1) *Predicted Power Demands of Buildings*: At each time step τ in the horizon, the predicted output is given by Eq. (4). Recall that only the predicted output means are propagated. Therefore $y_{\tau|t}$ is independent of the realization of the random variables $\mathbf{y}_{\tau|t}$ in previous steps. It follows that the random variables $y_{\tau|t}$, for all $t \leq \tau \leq t + H - 1$, are conditionally independent given the non-derived variables. Consequently, the sum of any subset $\mathcal{I} \subseteq \{t, \dots, t + H - 1\}$ of them is also a Gaussian random variable

$$\sum_{\tau \in \mathcal{I}} y_{\tau|t} \sim \mathcal{N}(\sum_{\tau \in \mathcal{I}} \bar{y}_{\tau|t}, \sum_{\tau \in \mathcal{I}} \sigma_{y,\tau|t}^2). \quad (7)$$

2) *Predicted Battery Power and SoC*: From Eq. (6), the predicted battery power at time $\tau \geq t$ has the Gaussian distribution $b_{\tau|t} \sim \mathcal{N}(\bar{b}_{\tau|t}, \sigma_{b,\tau|t}^2)$, where

$$\bar{b}_{\tau|t} = r_{\tau} - \delta_{\tau|t} - \bar{y}_{\tau|t}, \quad \sigma_{b,\tau|t}^2 = \sigma_{y,\tau|t}^2. \quad (8)$$

The battery's dynamics in Eq. (5) result in the future battery's SoC as, for $\tau \geq t$, $s_{\tau+1|t} = s_t + T \sum_{k=t}^{\tau} (r_k - \delta_{k|t}) - T \sum_{k=t}^{\tau} y_{k|t}$. It follows that $s_{\tau+1|t} \sim \mathcal{N}(\bar{s}_{\tau+1|t}, \sigma_{s,\tau+1|t}^2)$ where

$$\bar{s}_{\tau+1|t} = s_t + T \sum_{k=t}^{\tau} \bar{b}_{k|t}, \quad \sigma_{s,\tau+1|t}^2 = T^2 \sum_{k=t}^{\tau} \sigma_{y,k|t}^2 \quad (9)$$

in which Eq. (7) is used to derive the variance.

Recall that the battery's power and SoC are bounded (see Section III-B). In the probabilistic setting of the MPC, these constraints lead to corresponding chance constraints. We wish to guarantee that at each time step, the power and SoC constraints are satisfied with probability at least $(1 - \epsilon_p)$ and at least $(1 - \epsilon_s)$, respectively, where $0 < \epsilon_p, \epsilon_s \leq \frac{1}{2}$ are given constants. Specifically, for each τ in the horizon,

$$\Pr(b_{\min} \leq b_{\tau|t} \leq b_{\max}) \geq 1 - \epsilon_p \quad (10a)$$

$$\Pr(s_{\min} \leq s_{\tau|t} \leq s_{\max}) \geq 1 - \epsilon_s \quad (10b)$$

where $b_{\tau|t}$ and $s_{\tau|t}$ are Gaussian random variables whose means and variances are given in Eqs. (8) and (9).

3) *Objective Function*: The tracking controller aims to minimize the tracking error $\delta_{\tau|t}$, therefore we select the quadratic objective function to minimize $J = \sum_{\tau=t}^{t+H-1} \delta_{\tau|t}^2$. We may also consider additional objectives, such as reducing the battery's charge-discharge frequency to improve its lifetime. These objectives can be encoded as extra weighted terms to the above tracking objective function. In this paper, we will only consider the tracking objective.

4) *MPC Formulation*: Putting everything together, we obtain the following stochastic MPC formulation:

$$\begin{aligned} & \underset{\mathbf{u}(\cdot), \delta(\cdot)}{\text{minimize}} && J = \sum_{\tau=t}^{t+H-1} \delta_{\tau|t}^2 \\ & \text{subject to} && \text{Equations (4), (8), (9)} \end{aligned}$$

Chance constraints (10)

$$u_{\tau|t} \in [-1, 1], \quad \forall t \leq \tau \leq t + H - 1$$

At each t , the above optimization is solved to find optimal DR signal $u_{\tau|t}^*$ and optimal slacks $\delta_{\tau|t}^*$, $t \leq \tau \leq t + H - 1$; of which the first optimal decisions $u_{t|t}^*$ and $\delta_{t|t}^*$ are applied. At the next time step $t + 1$, the optimization is repeated with newly available information (i.e., new measurements and updated forecasts) and the shifted horizon $[t + 1, t + H]$.

The above MPC formulation is difficult to solve due to the two-sided chance constraints (10). In a recent work [14], it has been shown that any two-sided linear chance constraint of the form $\Pr(a \leq x^T \xi \leq b) \geq 1 - \epsilon$ is convex in a, b and x given that $\epsilon \leq \frac{1}{2}$. Here, $a \in \mathbb{R}, b \in \mathbb{R}, x \in \mathbb{R}^n$, and ξ is a known jointly Gaussian vector. Furthermore, there exists a computationally tractable second-order cone (SOC) approximation of this chance constraint, stated in Lemma 1.

Lemma 1 (Adapted from [14, Lemma 16]): Let $\xi \sim \mathcal{N}(\mu, \Sigma)$ be a jointly distributed Gaussian random vector with known mean μ and covariance matrix Σ , and $0 < \epsilon \leq \frac{1}{2}$. Let $LL^T = \Sigma$ be the Cholesky decomposition of

Σ , and $\Phi^{-1}(\cdot)$ the inverse cumulative distribution function of the standard Gaussian distribution. The following constraints, with the auxiliary variable γ and $\epsilon^* = \epsilon/1.25$,

$$\begin{aligned} \gamma &\geq \|L^T x\|_2, & a - b &\leq 2\Phi^{-1}(\epsilon^*/2)\gamma \\ a - \mu^T x &\leq \Phi^{-1}(\epsilon^*)\gamma, & \mu^T x - b &\leq \Phi^{-1}(\epsilon^*)\gamma \end{aligned}$$

is an SOC approximation that guarantees

$$\Pr(a \leq x^T \xi \leq b) \geq 1 - \epsilon.$$

This result allows us to conservatively approximate the chance constraints (10). We first rewrite, for each τ , $b_{\tau|t} = \bar{b}_{\tau|t} + \sigma_{y,\tau|t}\xi$ using Eq. (8), where $\xi \sim \mathcal{N}(0, 1)$. Then (10a) is equivalent to $\Pr(b_{\min} - \bar{b}_{\tau|t} \leq \sigma_{y,\tau|t}\xi \leq b_{\max} - \bar{b}_{\tau|t}) \geq 1 - \epsilon_p$, which is approximated by Lemma 1 as, $\forall \tau \in [t, t+H-1]$,

$$b_{\min} - \bar{b}_{\tau|t} \leq \Phi^{-1}(\epsilon_p^*)\sigma_{y,\tau|t}, \quad (11a)$$

$$\bar{b}_{\tau|t} - b_{\max} \leq \Phi^{-1}(\epsilon_p^*)\sigma_{y,\tau|t}, \quad (11b)$$

$$b_{\min} - b_{\max} \leq 2\Phi^{-1}(\epsilon_p^*/2)\sigma_{y,\tau|t}. \quad (11c)$$

Here $\epsilon_p^* = \epsilon_p/1.25$, and the auxiliary variable γ can be dropped for $\gamma = \|\sigma_{y,\tau|t}\|_2 = \sigma_{y,\tau|t}$.

Similarly, (10b) is approximated by first rewriting $s_{\tau|t} = \bar{s}_{\tau|t} + \sigma_{\tau|t}\xi$, where $\sigma_{\tau|t}$ is the vector $[\sigma_{y,t|t}, \dots, \sigma_{y,\tau|t}]^T$, and $\xi \sim \mathcal{N}(0, \mathbf{I})$ is a vector of $(\tau - t + 1)$ independent standard Gaussian random variables. Applying Lemma 1, we obtain the following constraints for all $t \leq \tau \leq t+H-1$, with the auxiliary variable γ_τ and $\epsilon_s^* = \epsilon_s/1.25$,

$$\gamma_\tau \geq \|\sigma_{\tau|t}\|_2, \quad s_{\min} - s_{\max} \leq 2\Phi^{-1}(\epsilon_s^*/2)\gamma_\tau \quad (12a)$$

$$s_{\min} - \bar{s}_{\tau|t} \leq \Phi^{-1}(\epsilon_s^*)\gamma_\tau, \quad \bar{s}_{\tau|t} - s_{\max} \leq \Phi^{-1}(\epsilon_s^*)\gamma_\tau \quad (12b)$$

The final MPC formulation with chance constraint approximations is given below:

$$\underset{u(\cdot), \delta(\cdot)}{\text{minimize}} \quad J = \sum_{\tau=t}^{t+H-1} \delta_{\tau|t}^2 \quad (13)$$

subject to Equations (4), (8), (9)

Constraints (11), (12)

$$u_{\tau|t} \in [-1, 1], \quad \forall t \leq \tau \leq t+H-1$$

Due to the GP model (4), the optimization problem (13) is a non-convex nonlinear program, which can be solved by a nonlinear optimization solver such as IPOPT [17].

IV. SIMULATION CASE STUDY

A. The Case Study

In this case study, we deployed a group of three buildings: a medium office building, a large hotel, and a supermarket. These buildings are modeled in EnergyPlus, taken from the U.S. Department of Energy's Commercial Reference Building Models (<http://energy.gov/eere/buildings/commercial-reference-buildings>). EnergyPlus is a high-fidelity whole building energy simulation software popularly used in research and industry. The library of reference building models represent roughly 70% of commercial buildings in the U.S. and are highly realistic. Therefore, we may consider these models as real physical buildings. Each building model was modified to receive global temperature setpoints from a rule-based DR controller. An aggregator, developed in Matlab, receives whole building power demand measurements from

the EnergyPlus models and computes and sends a single DR signal to all three local DR controllers. This co-simulation of multiple EnergyPlus and Matlab instances were developed and deployed with the OpenBuildNet framework [15] on two networked computers.

Each local DR controller implements a simple GTA strategy, which adjusts the global thermostat setpoints proportionally to the DR control signal by $-2u$ ($^\circ\text{C}$). Specifically, the setpoint is raised by up to 2°C to reduce cooling power demand when $u < 0$, or lowered by up to -2°C to increase demand when $u > 0$, compared to the nominal setpoint ($u = 0$). Here we assume that the $\pm 2^\circ\text{C}$ GTA is within the operational range and acceptable comfort range of each building. The sampling time T is 15 minutes. The weather profile was of Chicago (U.S.), which is available from the standard EnergyPlus' weather dataset. All analyses and simulations were performed for the summer weather, however similar results could be obtained for other seasons following the same approach. The usage patterns of the buildings, represented by various schedules (occupancy, equipment operation, etc.), were included in the EnergyPlus models. On a regular summer day, the nominal aggregate demand of the three buildings peaks at between 700 and 800 kW.

B. Gaussian Process Model Development

This section reports the development of a GP model for the collective demand response behavior of the buildings.

1) *Experiment* : An experiment was carried out by simulating the building models and their DR rules with i.i.d. uniformly random DR signal $u_t \in [-1, 1]$ during the first four weeks of June. The aggregator was not active during the experiment. Besides the excitation signal u , only three measurements were collected from EnergyPlus at each time step: aggregate power of the buildings, outside air temperature and humidity. The experiment data was partitioned into two sets: 75% of the data points, or three weeks of data, for training and the rest for validation. Because the measurements have vastly different scales (e.g., $u \in [-1, 1]$ while y is in the hundreds of kW), they were normalized into the range $[-1, 1]$ for model training purpose, and the validation results were de-normalized for reporting purpose.

2) *Feature Selection*: We observed that for each building, its operation schedules were similar between weekdays but were reduced significantly on Saturdays and switched off on Sundays. Hence, its response to the DR signal was similar between weekdays but very different between a weekday and a non-weekday. We could train a single model which takes into account the day of week, or train a model for weekdays and another model for non-weekdays. The second option was adopted since the accuracy of the specialized models would be higher than that of a universal model. In fact, we would only perform the study for weekdays because on weekends, the DR capability of the buildings was severely limited. The day of week input was thus excluded as it was unnecessary.

We chose an autoregressive GP with candidate inputs

$$x_t = [y_{t-l_y}, \dots, y_{t-1}, u_{t-l_u}, \dots, u_t, w_{t-l_w}, \dots, w_t, \hat{h}_t]$$

where y is the power demand, u the DR control signal, w the disturbances, and $\tilde{h}_t \in [0, 24]$ the current hour-of-day. Here, y and u are scalars, while $w = [\theta, \eta]$ where θ and η are respectively the ambient air temperature and humidity.

The lags l_y, l_u, l_w are unknown in advance. To select the relevant features, we used the SE covariance function with ARD (2) to train models with different orders, compared them on the validation dataset and examined their input length scales. Three initial GP models \mathcal{M}_n of order $n = 2, 3, 4$, with $l_y = n, l_u = l_w = n - 1$, were considered. The log likelihood of the identified model ℓ_1 was maximized during the training. For validation, three performance measures on the mean predictions were calculated: the root mean squared error (RMSE), the standardized mean-squared error (SMSE), and the mean standardized log loss (MSLL). Details on these measures can be found in [11, p. 78]. Essentially, the smaller they are, the better the predictions are.


The feature selection results are reported in Table II, in which extremely large length scales are replaced by ∞ and some of them are excluded due to lack of space. Observe that past weather conditions have infinitesimal influence on the current power demand. By examining the length scales, we determined the lag orders are: $l_y = 3, l_u = 1, l_w = 0$.

3) *Covariance Selection* : As mentioned earlier, the usage patterns of the buildings during weekdays are periodic. With this insight, we selected the covariance function as the product $k(x, x') = k_{nt}(x_{nt}, x'_{nt})k_h(\tilde{h}, \tilde{h}')$, where k_{nt} is a covariance function for non-temporal inputs x_{nt} and k_h is a periodic covariance function for the temporal input \tilde{h} ; $x = [x_{nt}, \tilde{h}]$. Covariance selection was performed by trying different choices for k_{nt} and k_h and selecting the one with the best performance on validation data. Finally, we selected the SE covariance function for k_{nt} and the periodic RQ covariance function for k_h .


4) *Model Validation*: The final GP model was validated on the testing dataset, resulting in good performance scores: RMSE = 14.579 (kW), SMSE = 0.0055, MSLL = -2.4358. To test the model for multistep-ahead predictions, two 24-hour-ahead prediction simulations were performed. The first simulation used the Monte-Carlo method with 400 samples per step, and the second simulation used the zero-variance method. The results are reported in Figure 2, which plot the true values y_t , the predicted mean \bar{y}_t and 95% confidence intervals (2σ), and the absolute residuals $|e_t| = |y_t - \bar{y}_t|$. In both cases, the prediction accuracy was good. Figure 2c plots the difference between the output standard deviations σ_{MC} and σ_{ZV} of the two methods. The error was small, averaged at 1.04 (kW), hence the variance estimates obtained by the zero-variance method were sufficiently accurate. Therefore, this method could be used for predicting future outputs in control and optimization.

C. Demand Tracking Control Simulation Results


In this simulation study, we considered a DR event lasting for four hours between 13:00 and 17:00 on July 7th. This day is outside the training period for the GP model and has an unusually high outside air temperature profile, which



(a) Monte-Carlo simulation results.



(b) Zero-variance simulation results.



(c) Standard deviation difference between methods.

Fig. 2. 24-hour ahead simulation results.

prompts the DR event. During the event, the system was requested to track a reference demand signal which was Δ_P (kW) below the baseline demand. The baseline was obtained by simulating the buildings in their nominal modes on the same day; alternatively, it could be calculated by simulating the GP with a zero DR control signal. Given Δ_P , the reference signal is $r_t = p_t^* - \Delta_P$ when t is between 13:00 and 17:00, where p_t^* is the baseline demand. The chance constraint thresholds were both 90%, i.e., $\epsilon_p = \epsilon_s = 0.1$.


The aggregator's MPC formulation (13) was implemented in Matlab with CasADi [16] and IPOPT [17] for solving the nonlinear program. Using the OpenBuildNet framework, we performed many simulations for different Δ_P values, different battery sizes, and three different horizons $H = 4, 8, 12$ (respectively 1, 2, 3 hours). The MPC problem size depends on the horizon H and is largest when $H = 12$. On an Apple Mac Pro 2013 with Matlab 2015b and utilizing only one CPU core, the largest MPC problem took less than a minute to solve at each time step, which was fast enough given the sampling time of 15 minutes.

Figure 3 illustrates the simulation results for $\Delta_P = 69$ kW, $H = 8$, $b_{\max} = 30$ kW, battery size $s_{\max} = 50$ kWh, and the initial battery SoC at half capacity. The top figure plots the baseline p^* (dashed line), the reference r (solid line), and the buildings' aggregate power y (dash-dotted line). The

TABLE II

FEATURE SELECTION RESULTS: $\lambda_{x,i}$ IS THE LENGTH SCALE OF INPUT $x(t-i)$; RELEVANT FEATURES ARE HIGHLIGHTED.

Model	ℓ_1	RMSE	SMSE	MSLL	$\lambda_{y,4}$	$\lambda_{y,3}$	$\lambda_{y,2}$	$\lambda_{y,1}$	$\lambda_{u,3}$	$\lambda_{u,2}$	$\lambda_{u,1}$	$\lambda_{u,0}$	$\lambda_{\theta,1}$	$\lambda_{\eta,1}$	$\lambda_{\theta,0}$	$\lambda_{\eta,0}$
\mathcal{M}_2	4151	21.31	0.0119	-2.0456			28.14	6.04			26.28	7.95	∞	∞	3.70	12.75
\mathcal{M}_3	4188	21.15	0.0118	-2.0923		63.32	24.07	4.71		∞	35.32	6.79	∞	∞	3.51	11.20
\mathcal{M}_4	4224	20.80	0.0114	-2.1655	470.34	72.11	26.63	4.78	∞	∞	49.59	6.92	∞	∞	3.97	11.73


Fig. 3. Simulation results for $\Delta_P = 69$ kW: power (top), DR signal (second), battery power (third), battery SoC (bottom).

DR event period is shaded. Observe that the buildings alone could not track the reference exactly: $y_t > r_t$, the dash-dotted line is above the solid line during the DR event. This slack between the buildings' power and the reference was fully compensated by the battery as shown in the third figure, and **the system successfully tracked the reference demand**. The battery SoC is plotted in the bottom figure.

The flexibility of the system, i.e., the maximum reduction Δ_P it is able to track exactly, is determined by the MPC horizon H and the battery size. By performing many simulations, we were able to find the system's flexibility for various combinations of H and battery size, summarized in Figure 4. Obviously, as the battery size or the MPC horizon increases, the flexibility also increases. However, we found that increasing the battery size, hence its cost, significantly (50kWh to 75kWh to 100kWh) did not result in significant increases in the DR flexibility. Increasing the MPC horizon, on the other hand, could achieve similar increases in the DR flexibility while being much easier and cheaper to implement.

V. CONCLUSIONS

This paper presented a data-driven approach to provide demand response services, in particular reference demand tracking, with buildings and energy storage. The key advantage of the proposed scheme is its simplicity in implementation. It requires minimal changes in the building's equipment and energy management system. Furthermore, it achieves a separation of concerns between the building's level and

Fig. 4. Flexibility for different MPC horizons H and battery sizes s_{\max} .

the aggregator's level through an abstract DR control signal and through the use of Gaussian Process models. We showed that a GP model trained on experiment data could achieve accurate predictions of a building's power demand with a quantifiable confidence measure. We developed an MPC formulation for the aggregator's algorithm and showed the effectiveness of the proposed approach in a realistic simulation case study.

Several extensions are currently being pursued. We are developing an analytical method to quantify the system's flexibility without resorting to brute-force simulations. Uncertainty propagation will be incorporated into the scheme for more accurate predictions with the GP. Large-scale distributed optimization algorithms for controlling a large number of buildings and batteries to provide DR services are being investigated. Finally, we will study the effectiveness of the proposed approach for other demand response and ancillary service applications.

REFERENCES

- [1] P. Siano, "Demand response and smart grids—a survey," *Renewable and Sustainable Energy Reviews*, 2014.
- [2] F. A. Qureshi, T. T. Gorecki, and C. N. Jones, "Model predictive control for market-based demand response participation," in *IFAC Proceedings*, vol. 47, no. 3. Elsevier, 2014, pp. 11 153–11 158.
- [3] M. Maasoumy, A. Sangiovanni-Vincentelli, *et al.*, "Smart connected buildings design automation: Foundations and trends," *Foundations and Trends in Electronic Design Automation*, vol. 10, pp. 1–143, 2016.
- [4] D. Sturzenegger, D. Gyalistras, M. Morari, and R. S. Smith, "Model predictive climate control of a swiss office building: Implementation, results, and cost–benefit analysis," *IEEE Trans. Control Syst. Technol.*, vol. 24, no. 1, pp. 1–12, 2016.
- [5] D. S. Watson, S. Kiliccote, N. Motegi, and M. A. Piette, "Strategies for demand response in commercial buildings," in *2006 ACEEE Summer Study on Energy Efficiency in Buildings*, 2006.
- [6] M. Behl, T. X. Nghiem, and R. Mangharam, "DR-Advisor: A data driven demand response recommender system," in *Proceedings of CISBAT*, 2015.
- [7] H. Y. Noh and R. Rajagopal, "Data-driven forecasting algorithms for building energy consumption," in *Proc. of SPIE*, 2013.
- [8] R. Yin, E. C. Kara, Y. Li, *et al.*, "Quantifying flexibility of commercial and residential loads for demand response using setpoint changes," *Applied Energy*, 2016.

- [9] Y. Zhang, Z. O'Neill, B. Dong, and G. Augenbroe, "Comparisons of inverse modeling approaches for predicting building energy performance," *Building and Environment*, vol. 86, pp. 177–190, 2015.
- [10] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- [11] J. Kocijan, *Modelling and Control of Dynamic Systems Using Gaussian Process Models*. Springer, 2016.
- [12] A. Girard, "Approximate methods for propagation of uncertainty with gaussian process models," PhD thesis, University of Glasgow, 2004.
- [13] J. M. Maciejowski, *Predictive control: with constraints*. Pearson, 2002.
- [14] M. Lubin, D. Bienstock, and J. P. Vielma, "Two-sided linear chance constraints and extensions," *ArXiv e-prints*, Feb. 2016, arXiv:1507.01995v2 [math.OC].
- [15] T. X. Nghiem, A. Bitlislioglu, T. Gorecki, F. A. Qureshi, and C. N. Jones, "OpenBuildNet framework for distributed co-simulation of smart energy systems," in *ICARCV'16*, 2016.
- [16] J. Andersson, "A General-Purpose Software Framework for Dynamic Optimization," PhD thesis, KU Leuven, 2013.
- [17] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical programming*, vol. 106, no. 1, pp. 25–57, 2006.