A Modeling and Closed-Loop Planning Strategy to Achieve Net-Zero Energy Use in Buildings

Michael J. Risbeck[a,*], Chenlu Zhang[a], Saman Cyrus[a], and Young M. Lee[a]

[a]AI Product Development Group, Johnson Controls International Plc., Milwaukee, WI 53202

Abstract

Given the challenge posed by climate change, there is increasing interest in operating "net-zero" energy buildings in which any energy consumed is balanced by corresponding renewable energy generation on site. Though there has been significant work on designing buildings to achieve this goal, less attention has been paid to the operational side to ensure that net-zero status is met and maintained. Therefore, we propose in this paper a modeling and closed-loop planning strategy to guide the operation of net-zero buildings. Planning is performed by iteratively solving an optimization problem to determine curtailment actions needed to achieve net-zero status by the end of the year. This optimization makes use of forecasted energy consumption and renewable energy generation, which can be provided by lightweight stochastic models fit to historical data. We illustrate the effectiveness of the proposed strategy through simulations, showing that closed-loop planning is an important feature to reliably meet net zero energy use.

Keywords

Net Zero Energy Buildings, Stochastic Modeling, Closed-Loop Planning

Introduction

In the United States, energy use in commercial and residential buildings accounts for 18% and 21% respectively of total energy consumption (Dunn, 2022). This energy is currently provided primarily by fossil fuels, and thus addressing climate change will require both reducing overall consumption and replacing it with renewable sources (Sartori et al., 2012). A primary goal is to achieve "net-zero" energy consumption so that the building is no longer a burden on overall global climate. Achieving this goal will require significant changes to both the design and operation of buildings (Lu et al., 2015).

Various definitions of "net-zero energy" exist (Sartori et al., 2012), differing in what quantities are balanced (e.g., energy import/export or load/generation) and over what time period the balance is calculated (typically a year, but monthly formulations are also possible). Additional weighting factors could be applied so that the balance is effectively on energy cost or carbon emissions (Torcellini and Crawley, 2006). In any case, achieving net-zero status generally requires careful planning and operation, which can be guided by mathematical modeling and optimization strategies.

Background and Scope

The topic of net-zero energy buildings has received much attention in the literature, for both the residential and commercial sectors. In residential buildings, the most efficient design depends on the available utilities, including the electricity grid and district heating or cooling (Wu and Skye, 2021). In commercial buildings, retrofit is likely required to reduce energy consumption to levels that could be met by renewable generation (Aksamija, 2016). To guide these design decisions, various optimization formulations have been developed (Longo et al., 2019), which can be further augmented with simulation tools (Attia et al., 2012). Design problems may require trading off multiple objectives, such as total cost and occupant comfort (Harkouss et al., 2018), which requires the use of multiobjective optimization to identify efficient solutions.

Once the required design elements have been put in place, it is necessary to operate buildings with high efficiency. Optimal efficiency often requires energy storage, which can be operated using model predictive control (MPC) or other advanced control strategies (Lu et al., 2015). These techniques can help reduce overall energy consumption without compromising service by temporally decoupling on-site utility production (e.g., chilled or hot water) from its actual consumption (Rawlings et al., 2018). Assuming proper system design, net-zero operation should be possible in theory. However, due to the inherent uncertainty in long-term energy production and generation, it is far from guaranteed that nominal operation will ultimately achieve the goal.

For our part, we are interested in generating target trajectories for energy consumption and renewable energy generation that can be used to guide operation. This process achieves much less attention in the literature, but it is nevertheless a critical component in practical net-zero buildings.

* E-mail: michael.james.risbeck@jci.com

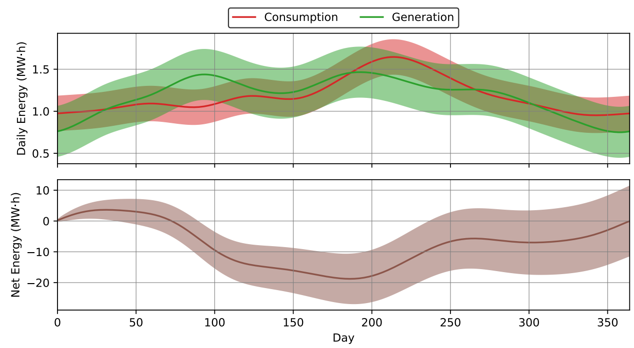

Figure 1: Example daily energy generation, consumption, and cumulative net for a net-zero building. Lines show medians while shaded regions show one standard deviation. See the examples section for details about the model.

An important complication is that energy consumption and generation are generally not completely in-phase. Therefore, over the course of a year, net energy may become very positive or very negative before ultimately reaching zero at the end of the year. To illustrate, Figure 1 shows annual energy models based on data from a building in San Diego, CA (Silwal et al., 2021), adjusted so that mean values give net-zero annual operation. Despite the relatively constant weather conditions throughout the year, the building nevertheless goes through periods of slightly positive and significantly negative net energy. If this particular building happened to be at zero net energy in the middle of the year, then it is quite unlikely that net-zero status will be maintained through the end of the year. Thus, building operators are in need of continuous guidance to check whether they are on track for net-zero consumption and possibly take deliberate curtailment actions should energy consumption out-pace energy generation.

Paper Overview

Our goal in this paper is to present a modeling and closed-loop planning strategy to achieve net-zero energy use in buildings. We are primarily interested in commercial buildings, which are likely to have sufficient metering and higher operational flexibility, but in principle, the same techniques could be applied to residential buildings. We start by presenting an optimization problem that can be used to generate daily or weekly energy curtailment targets that are likely to deliver net-zero consumption by the end of the year. This optimization requires forecasts of baseline energy consumption and generation, and the objective is to minimize required curtailment from baseline consumption. Through regular re-optimization, these targets can be continually updated throughout the year as actual consumption and generation is realized. We next discuss the modeling needed to generate the forecasts used in the optimization problem. The proposed strategy is data-driven and is intended to be suitable in cases where only a modest amount of historical data is available. After presenting the overall framework, we illustrate its application using an open-source dataset. Finally, we conclude with a summary and future outlook.

Planning Optimization

A major component of our proposed strategy is the consumption planning optimization problem. The goal of this problem is to optimize energy targets to achieve net-zero status by minimizing required curtailment from baseline consumption levels. For our purposes, we adopt the load/generation definition of net-zero energy, as it is generally easier to measure and verify than alternatives such as import/export (Sartori et al., 2012). For brevity, we use the term "net-zero," but note that we are focused on net-zero *energy*. In addition, we use "generation" to refer specifically to renewable generation, which can include solar, wind, and geothermal sources. Note also that we refer to the system of interest as a "building," but this entity could also represent a campus or portfolio of multiple buildings that can operate in coordination (Peterson et al., 2015). We start with the formulation and later discuss how the optimization is applied in closed loop.

Problem Formulation

To model the problem, we define the ordered set of time periods $t \in \mathbb{T} := \{0, 1, \ldots, T-1\}$ for total horizon T. In general, we expect these time periods to be either days or weeks and the total horizon to span one calendar year, but of course other arrangements are possible. With this set, we define the following decision variables:

- X_t: cumulative net energy consumption at the beginning of time period $t \in \mathbb{T} \cup \{T\}$.
- C_t: fractional curtailment for consumption during period $t \in \mathbb{T}$.

The extra final variable X_T thus gives the cumulative net energy consumption at the end of the horizon, and the goal is to keep this value less than or equal to zero. Along with these variables, we define the following parameters:

- β_t: forecasted baseline energy consumption during time period $t \in \mathbb{T}$.
- γ_t: forecasted energy generation (or equivalent offsets) during time period $t \in \mathbb{T}$.
- $\phi_t(\cdot)$: cost function for curtailment during time $t \in \mathbb{T}$.
- $\bar{C}_t$: upper bound on curtailment for time $t \in \mathbb{T}$.

From these definitions, we arrive at our planning optimization problem:

$$\min_{X_t, C_t} \quad \sum_{t \in \mathbb{T}} \phi_t(C_t) \tag{1a}$$

$$\text{s.t.} \quad X_{t+1} = X_t + \beta_t(1 - C_t) - \gamma_t \quad t \in \mathbb{T} \tag{1b}$$

$$0 \le C_t \le \bar{C}_t \quad t \in \mathbb{T} \tag{1c}$$

$$X_T \le 0 \tag{1d}$$

$$X_0 \text{ given} \tag{1e}$$

The overall goal is to minimize the total cost of curtailment actions such that the building achieves net-zero status (or better) at the end of the horizon. The resulting net-energy sequence X_t can serve as a target for the building to meet, while the optimized curtailment actions C_t indicate when and by how much energy consumption should be reduced.

We note that the intermediate variables X_t and the "dynamics" (1b) are included to illustrate the state-space structure of the system (with X_t serving as the states and C_t as the manipulated inputs). However, these variables and constraints could easily be pre-solved away, leaving just the bounds on C_t and a single constraint.

Cost Function and Solution Methods

The cost function (1a) consists of stage costs for taking curtailment actions for time period. Ideally, the functions $\phi_t(\cdot)$ would represent some actual tangible cost incurred by the building or its occupants for taking the corresponding curtailment action. For energy consumption associated with lighting or space heating/cooling, various visual and thermal comfort metrics exist that could be the basis of curtailment costs (Longo et al., 2019). However, these functions often require additional parameter tuning and can be challenging to balance. Thus, it is not likely that tangible cost functions will be available for most applications, and instead we opt for an empirical approach.

The overall design goal for the cost function is to appropriately balance required curtailment actions throughout time. One possibility is that curtailment actions should be spread evenly throughout each day. For this purpose, *quadratic* cost functions could be used, for example

$$\phi_t(C) := \omega_t C^2 \tag{2}$$

with the ω_t weights prioritizing particular time periods for curtailment (with a higher ω_t indicating that curtailment is less desirable for time t). This choice also reflects the fact that curtailment becomes increasingly difficult. For example, it may be easy to reduce a small fraction of HVAC energy use by delaying system startup in the morning when the building is sparsely occupied, but any additional reduction will begin to compromise occupant comfort and thus be undesirable.

An alternative possibility is that it is desired to make curtailment actions *sparse*, i.e., with $C_t = 0$ for most values of t. This situation would reflect the fact that building managers may prefer to operate normally on most days and then take intense curtailment actions on a small number of days. We note that this structure is similar to some grid-level demand-response programs in which only the most energy-intense periods are identified as "extreme days" in which consumers are specifically incentivized to reduce consumption (Albadi and El-Saadany, 2008). For these situations, a *nonconvex* or even *discontinuous* cost function could be adopted, e.g.,

$$\phi_t(C) := \omega_t \mathbf{1}_{>0}(C) \tag{3}$$

$$\approx \omega_t(1 - \exp(-\alpha C)) \tag{4}$$

with $\mathbf{1}_{>0}(\cdot)$ denoting the indicator function for the positive numbers. Of course, the discontinuity makes the optimization problem significantly more challenging, but it is likely to still be tractable for typical problem sizes.

To solve the optimization problem (1), we note that if the curtailment cost functions $\phi_t(\cdot)$ are convex, then the overall problem is convex. In the specific case of (convex) quadratic costs like (2), dedicated quadratic-programming solvers could be used, while for general nonlinear functions, a nonlinear-programming solver like IPOPT (Wächter and Biegler, 2006), can be applied. However, in cases where $\phi_t(\cdot)$ is nonconvex, some additional care is needed. If the function is at least smooth as in (4), then nonlinear-programming solvers could still be used, although it is possible that the solver converges to a locally optimal but globally suboptimal solution. For a discontinuous cost function like (3), the problem could be re-formulated using binary variables and solved via mixed-integer linear programming techniques. Alternatively, an effective strategy could be to discretize the problem (for appropriately chosen grid sizes on X_t and C_t) and apply dynamic programming. Such an approach would guarantee global optimality (up to discretization error). We note that because the problem (1) does not need to be solved very frequently (e.g., perhaps once per day or week), computational efficiency is not a major concern, as long as the problem does not become completely intractable.

Closed-Loop Implementation

The overall purpose served by the planning problem (1) is to provide guidance to building managers in how to operate their buildings so as to achieve net-zero energy consumption by the end of the year. A key requirement is that this guidance is updated regularly based on actual performance so far and likely consumption/generation in the future. Thus, we propose to solve the planning problem repeatedly in sequence with a shrinking horizon similar to batch MPC. At the beginning of each time period, mean forecasts (conditioned on values observed so far) for β_t and γ_t to the end of the year are generated from the stochastic models, and the planning optimization (1) is solved using those values. The optimized value of curtailment C_t^* at the current time t is then implemented. In a real building, implementing the suggested curtailment is a nontrivial process, and would likely be performed by adjusting setpoints or schedules for the HVAC system, lighting, etc. After operating with these actions, actual values of the *curtailed* consumption $(1 - C_t)\beta_t$ and generation γ_t are realized. The hypothetical *baseline* consumption β_t and current net energy X_t can then be calculated, after which the process is repeated at the next time period, with the horizon now shrinking by one period.

Extensions

Before moving on, we briefly discuss some possible extensions to the planning problem (1). As formulated, we assume that energy generation is given as a parameter and that the only actions to reduce net energy consumption are to curtail consumption. However, these assumptions may not hold in more general settings. One possibility is that the use of generation equipment is discretionary. For example, wind turbines could be shut down when not needed to reduce equipment wear. In such cases, additional decision variables could be added to choose how much generation capacity to employ. Of course, an appropriate cost term would need to be added to the objective function, which may be hard to tune against the existing curtailment cost, but in principle, it would be possible. Alternatively, it may be possible to purchase energy offsets that get credited as reduced energy consumption. Thus, decision variables to purchase offsets could be added to the formulation, either constrained to a maximum budget or added as a separate term in the cost function. This extension becomes particularly interesting in the stochastic case where prices of offsets may fluctuate throughout the year, and thus it is beneficial to make purchases when prices are relatively low but only if it will eventually become necessary to achieve net-zero status. In either case, the primary change to the formulation is that the net-energy calculation in (1b) is augmented with the additional terms necessary to model the particular effect.

In some buildings, energy consumption may be measurable across separate categories (e.g., HVAC, lighting, or plug loads) that could be curtailed separately. In addition, the system of interest could a campus or portfolio of multiple buildings with independent consumption measurements and curtailment. To model such situations, extra indices $k \in \mathbb{K}$ could be added, with the curtailment actions extended to C_{kt} to consider per-category and/or per-building curtailment. Of course, extension to categorial forecasts β_{kt} and cost functions $\phi_{kt}(\cdot)$ would also be required, which could potentially complicate tuning and modeling. However, the additional specificity in curtailment recommendations would likely be valuable to building managers. Finally, it may be desirable to achieve net-zero status over multiple overlapping windows (e.g., over an annual period at the beginning of each month). For this situation, new indices $w \in \mathbb{W}$ could be added, giving a separate state X_{tw} for each window with the same dynamics as in (1b) and a window-dependent ending time T_w for the net-zero constraint (1d). With this modification, the horizon of each open-loop problem would extend to cover the end of all windows containing the current day, and would thus remain roughly the same length in contrast to the shrinking horizon for the single-window case.

Stochastic Modeling

A key requirement of the planning problem discussed in the previous section is the (forecasted) baseline energy consumption β_{kt} and energy generation γ_t. To solve the deterministic problem, we require mean values of these parameters for each time point remaining in the horizon, while in the stochastic formulation, we require multiple possible realizations. Although there are many sophisticated machine-learning models that could be built for this purpose, we note that many buildings are unlikely to have significant amounts of historical data. Thus, the approach we describe in this section reflects a balance between meeting model requirements while being amenable to small and possibly dirty training datasets.

Model Training

To fit our requirements, we fit each model in two separate pieces: (1) a deterministic but possibly nonlinear model for the time-varying mean; and (2) a stochastic linear autoregressive (AR) model for deviation from the mean. Mathematically, we use x_t to denote our quantity of interest, μ_t to denote the value of the mean model, and e_t to denote the deviation from the mean. As before, $t \in \mathbb{T}$ is the time index, which is discrete and finite. In the interest of brevity, we will model each unique quantity separately so that x_t is a scalar. However, because the different variables are likely to be correlated (e.g., higher energy generation from solar panels is likely to be correlated with higher energy consumption for space cooling), it may be useful to model them together. The approach proposed here does generalize to the vector case, although additional training data may be required to avoid overfitting to spurious correlations.

For the mean model, we can use any appropriate function of time. A relatively simple approach is to use a finite basis of spline functions that covers the time horizon. Letting ψ_{it} give the value of the ith basis function for time t, we can fit the model via linear least squares as

$$\min_{\alpha_i} \quad \sum_t |x_t - \mu_t|^2 \text{ s.t. } \mu_t = \sum_i \alpha_i \psi_{it}$$

in which the x_t are known values from the training data set, and the α_i are optimized model coefficients. By using spline functions, we ensure that the mean model is smoothly varying (as expected for the quantities being modeled), and we can still fit parameters in cases where isolated x_t samples are corrupted or missing, simply by removing those terms from the objective function. If other independent variables are known (e.g., the building occupancy schedule) or can be reasonably forecasted (e.g., the weather), then can be included as additional model inputs with associated trainable parameters.

Once the mean model μ_t is fit, we can move on to the deviation model for $e_t := x_t - \mu_t$. As mentioned above, we assume a linear autoregressive model such that

$$e_t = \sum_{n=0}^{N-1} a_n e_{t-n} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma^2) \tag{5}$$

for a suitably chosen model order N. The coefficients a_n in the AR model (5) are fit via constrained least-squares regression, i.e.,

$$\min_{a_n} \quad \sum_{t=N}^{T-1} \left(e_t - \sum_{n=0}^{N-1} a_n e_{t-n} \right)^2 \text{ s.t. } \sum_{n=0}^{N-1} |a_n| \leq 1 - \delta \tag{6}$$

in which the values of e_t are computed from x_t and μ_t. The constraint is added to ensure that the resulting AR model is stable with tolerance $\delta \in [0, 1)$, which ensures that resulting samples will remain bounded. The noise variance σ^2 is then taken as the optimal value of the objective function in (6), which corresponds to the variance in the model fit. Note that the appropriate model order can be chosen from standard methods, e.g., by examining the autocorrelation plot for e_t.

Model Sampling

In order to iterate the model (5), we need to know values of the previous error terms $\mathbf{e}_0 := (e_{-N}, \ldots, e_{-1})$. When starting at least N steps into the horizon, these values can be calculated from the actual realizations of x_t and mean model μ_t. However, when starting from the beginning, one should start from the stationary distribution $\mathbf{e}_0 \sim \mathcal{N}(0, \boldsymbol{\Sigma}_0)$. Once the $\mathbf{e}_0$ is known, the full sequence $\mathbf{e} := (e_0, \ldots, e_{T-1})$ can be obtained by iterating (5), randomly sampling ε_t each step or keeping $\varepsilon_t \equiv 0$ to obtain the mean value. Alternatively, we note that the full sequence is jointly normal $\mathbf{e} \sim \mathcal{N}(0, \mathbf{\Sigma})$, which can be useful for calculating quantiles or confidence intervals for the resulting trajectories. In both cases, the covariance $\mathbf{\Sigma}_0$ or $\mathbf{\Sigma}$ can be obtained by applying the Yule-Walker equations to the AR coefficients a_n and noise variance σ^2.

We note that the *linear* definition of the noise $e_t = x_t - \mu_t$ means that the realized values of x_t could take on any real value. Thus, quantities like energy consumption that must be nonnegative could take on invalid values. In practice, this is not a significant concern if the μ_t are sufficiently far from zero relative to the random noise, and individual invalid samples could be clipped. However, for sign-restricted quantities with structural zeros (e.g., space heating loads, which are typically zero during summer months), it may be useful to use a *multiplicative* definition of noise $e_t = \log(x_t/\mu_t)$. Some special handling is needed when x_t or μ_t are zero, but overall, this approach will likely give more realistic samples.

Illustrative Examples

To illustrate the proposed framework, we make use of an open-source dataset for the Trade Street building on the UC-San Diego campus (Silwal et al., 2021). The data file provides daily total energy consumption and on-site generation from photovoltaic (PV) panels, spanning roughly four years of operation but with some missing points. We will start by building stochastic models for generation and consumption, and then use those models for net-zero planning.

Stochastic Modeling

For the building of interest, we require an annual model for daily energy generation and consumption. We start by fitting mean models for these values via linear regression against a basis of cubic splines (13 knots with a periodicity constraint) covering the year. Figure 2 shows these mean models and the underlying datasets. From these results, we see that although the daily data is very noisy, the proposed models capture the overall behavior. We note that for planning purposes, the *cumulative* values are the most important, and fortunately the cumulative models are even more accurate due to noise cancelation from day to day. We see also that the uncertainty region (to be discussed later) captures the spread of the data, except perhaps for extreme samples in the generation data (likely due to extremely cloudy days that prevent PV generation). Thus, these models are suitable for the proposed use case.

Before moving on to the noise models, we note that there is a strong 7-day periodic component in the consumption data due to the weekly occupancy cycle of the building. Of course, the simple mean model shown before does not capture that behavior. Thus, to improve model predictions, we fit an alternative consumption model (denoted "Consumption*") by adding seven binary features for the day of the week with corresponding regression coefficients. In the interest of brevity, we do not show this model in Figure 2, but we will use that model in the following discussion.

With the mean models in hand, we can now fit autoregressive noise models for each. A key decision in this process is what model order to use. To help make this determination, Figure 3 plots the autocorrelation function of the mean-model deviations in the dataset. The most striking feature in this plot is the periodic behavior for the "Consumption" model, which derives from the occupancy pattern as just discussed. Fortunately, when adding the extra day-of-week features to the "Consumption*" model, this behavior goes away. Although there is not clear cutoff in autocorrelation, we choose 7 days as the order for each AR model and fit coefficients via linear regression. The resulting CVRMSE of the models is 14.9% for Generation, 18.9% for Consumption, and 13.0% for Generation*. Thus, we see that the inclusion of the extra features in the "Consumption*" model does lead to significant reduction in error. Though predictions could likely be made more accurate by applying more advanced modeling strategies, the models obtained here are sufficiently accurate for our purposes.

Closed-Loop Planning

With the stochastic models fit in the previous section, we can apply the proposed planning framework. Unfortunately, the baseline model fits indicate that the building annually consumes 13% more energy than it generates. Thus achieving net-zero status via curtailment would likely cause too large of a disruption to building occupants. To make the simulation more realistic, we assume that additional PV generation capacity is added so that mean total consumption is only 5% above generation. (We accomplish this by simply rescaling the predictions of the generation model.) This gap is narrow enough that it could be closed through operation but still wide enough to illustrate the planning problem.

To simulate performance of the closed-loop planning procedure, we draw random samples from each model and use those values as the actual consumption and (baseline) generation that gets revealed over the course of the simulation. For simplicity, we assume that the suggested curtailment fraction C_t can be implemented exactly, and thus the resulting curtailed consumption is set equal to the (random) realized baseline consumption value reduced by the optimized curtailment. We refer to the resulting system trajectory as "Closed-Loop". For comparison, we compute a "Perfect" solution, obtained by solving (1) using the actual realized values for β_t and γ_t (i.e., assuming we have access to a perfect forecast). In addition, we simulate a "Naive" solution that optimizes (1) once using mean values of parameters and uses the resulting curtailments C_t regardless of the actual parameter realizations.

We begin by illustrating a closed-loop planning trajectory. Figure 4 shows open-loop trajectories for curtailment and net energy for a single simulation. Note that "Open-Loop" indicates the *planned* trajectories from the open-loop optimization problems, while "Mean" shows the *planned* net-energy trajectory for the "Naive" solution. We see from this plot that, although the open-loop plans are made using forecasted values, the resulting closed-loop trajectory stays reasonably close to the perfect solution. The only real defect is the higher variance in curtailment at the end of the year, as aggressive action is needed to make up for daily prediction error. Never-

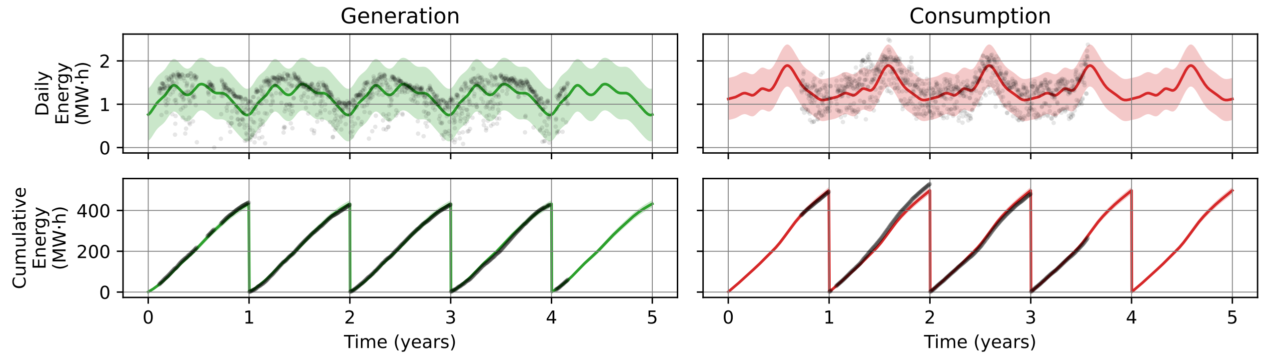

Figure 2: Mean model fits to example dataset for generation and consumption. Solid lines show mean values, while shaded regions show two standard deviations for the predictions (as derived from the AR noise models). Points show samples from the datset, with missing values in the cumulative plot assumed equal to the mean model.

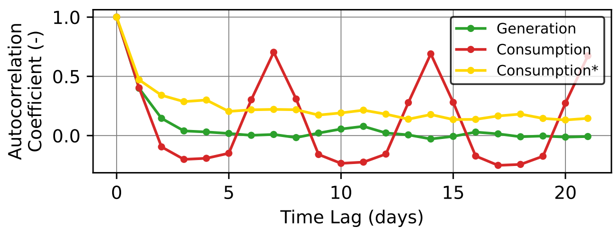

Figure 3: Autocorrelation plots for deviations from the mean models. Due to missing values, each lag is normalized by the number of valid pairs in the dataset.

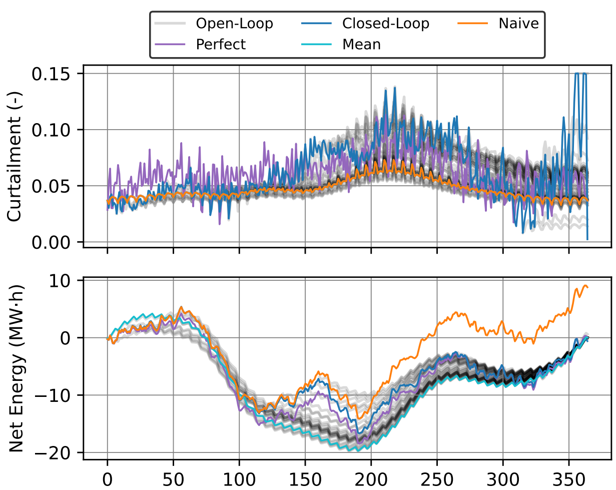

Figure 4: Example annual planning trajectories. Note that only every 7th open-loop solution is plotted.

theless, the system ultimately does achieve net-zero status for the year. By contrast, the Naive solution does not correctly revise its plan and thus finishes the year with significant net-positive energy consumption.

To assess overall performance, we repeat the previous simulations for 1,000 random samples from the generation and consumption models. For each simulation, we compute the total curtailment cost by summing $\phi_t(C_t)$ and check the final net-energy value X_t. Ideally, the costs are low, and net-energy is at or below zero. Figure 5 shows distributions of these values for the Perfect, Closed-Loop, and Naive solutions. From these results, we see that the Closed-Loop solution is performing reasonably close to the Perfect solution and significantly better than the Naive solution. We note that the median Closed-Loop curtailment cost is 33% higher than Perfect, which is reasonable considering that the cost function is quadratic. (Note that the Naive cost is constant, as it uses the same mean C_t sequence for all realizations.) In addition, final net energy for Closed-Loop has significantly less variance compared to the Naive solution, indicating that it consistently achieves the net-energy goal without unnecessary curtailment. The overall conclusion is that the Closed-Loop implementation is critical to achieving the net-zero goal and thus fills an important need in the net-zero buildings space.

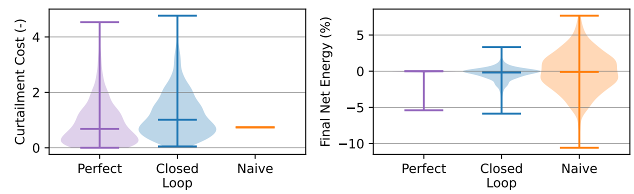

Figure 5: Performance results across 1,000 random simulations. Ranges show min, median, and max values, while shaded regions give a kernel density estimate for the distribution.

Conclusions

In this paper, we have presented a lightweight modeling and closed-loop planning strategy to achieve net-zero energy use in buildings. To forecast baseline energy consumption and generation, stochastic models can be built by combining a deterministic mean model with a linear AR noise model. These forecasts can then be used to regularly solve a planning optimization problem that suggests required curtailment actions in order to meed the net zero constraint by the end of the year. We have illustrated through examples that the proposed models give good accuracy while still capturing the inherent randomness in the modeled quantities, and that the closed-loop planning process is successful at guiding net-zero energy operation, with only slightly higher curtailment intensity compared to a perfect solution. Future work will focus on extending the planning problem to include additional decision variables and integrating the curtailment recommendations with building control systems to actually achieve the desired energy consumption.

References

Aksamija, A. (2016). Regenerative design and adaptive reuse of existing commercial buildings for net-zero energy use. *Sustainable Cities and Society 27*, 185–195.

Albadi, M. and E. El-Saadany (2008). A summary of demand response in electricity markets. *Electric Power Systems Research 78*, 1989–1996.

Attia, S., E. Gratia, A. De Herde, and J. L. Hensen (2012). Simulation-based decision support tool for early stages of zero-energy building design. *Energy and Buildings 49*, 2–15.

Dunn, D. R. (2022). Monthly Energy Review, April 2022, Section 2. Technical Report DOE/EIA-0035(2022/4), US Energy Information Administration.

Harkouss, F., F. Fardoun, and P. H. Biwole (2018). Multi-objective optimization methodology for net zero energy buildings. *Journal of Building Engineering 16*, 57–71.

Longo, S., F. Montana, and E. R. Sanseverino (2019). A review on optimization and cost-optimal methodologies in low-energy buildings design and environmental considerations. *Sustainable Cities and Society 45*, 87–104.

Lu, Y., S. Wang, and K. Shan (2015). Design optimization and optimal control of grid-connected and standalone nearly/net zero energy buildings. *Applied Energy 155*, 463–477.

Peterson, K., P. Torcellini, and R. Grant (2015). A common definition for zero energy buildings. Technical Report DOE/EE-1247, US Department of Energy.

Rawlings, J. B., N. R. Patel, M. J. Risbeck, C. T. Maravelias, M. J. Wenzel, and R. D. Turney (2018). Economic MPC and real-time decision making with application to large-scale HVAC energy systems. *Computers & Chemical Engineering 114*, 89–98.

Sartori, I., A. Napolitano, and K. Voss (2012). Net zero energy buildings: A consistent definition framework. *Energy and Buildings 48*, 220–232.

Silwal, S., C. Mullican, Y.-A. Chen, A. Ghosh, J. Dilliott, and J. Kleissl (2021). Open-source multi-year power generation, consumption, and storage data in a microgrid. *Journal of Renewable and Sustainable Energy 13*(2), 025301.

Torcellini, P. A. and D. B. Crawley (2006). Understanding zero-energy buildings. *ASHRAE Journal 48*(9), 62–69.

Wächter, A. and L. T. Biegler (2006). On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming 106*(1), 25–57.

Wu, W. and H. M. Skye (2021). Residential net-zero energy buildings: Review and perspective. *Renewable and Sustainable Energy Reviews 142*, 110859.