

# Field Performance of Commercial Building Load Flexibility Using Model Predictive Control

Ettore Zanetti<sup>1\*</sup>, David Blum<sup>1</sup>, Marco Pritoni<sup>1</sup>, Mary Ann Piette<sup>1</sup>,

<sup>1</sup> Building Technology & Urban Systems Division, Lawrence Berkeley National Laboratory,  
Berkeley, CA, USA  
ezanetti@lbl.gov

\* Corresponding Author

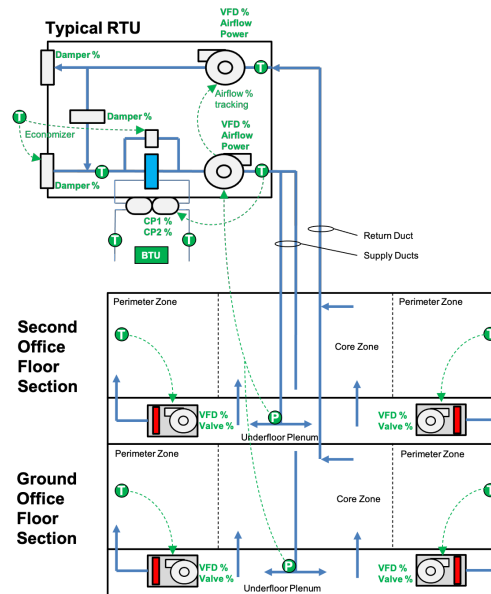
## ABSTRACT

Model Predictive Control (MPC) applied to buildings is starting to see some commercial adoption by companies. However, it is hard to estimate if relative energy cost savings are enough to justify the cost of MPC implementation with few reported demonstrations. In small commercial and residential buildings, a one-size-fits-all solution can help reduce implementation costs, while in very large buildings or districts the potential energy cost savings magnitude can cover more tailored solutions. This estimation becomes harder for medium to large commercial buildings, where a one-size-fits-all solution cannot be adopted and potential energy cost savings might not be sufficient to cover a tailored solution. Therefore, value propositions in addition to energy efficiency alone can make MPC technology more attractive through additional energy cost savings. One such value proposition is load shifting in response to dynamic electricity prices. On this aspect, MPC is a key technology to unlock building thermal mass for energy flexibility in response to electric grid conditions. This study shows the experimental results of MPC control of an office building in Berkeley, where different dynamic electricity price profiles were used in the MPC objective function to shift the building load and to calculate hypothetical electricity costs. Results show potential 50% cost savings with respect to the existing controller with the dynamic price scenario.

## 1. INTRODUCTION

Building electrification is key to decarbonizing the energy sector (U.S. DOE, 2024), and since buildings already account for more than 70% of the electricity use in the U.S., they also play an important role in providing demand flexibility to help mitigate the mismatch between supply and demand that results from increased electricity generation from renewable energy resources (Satchwell et al., 2021). Buildings can provide flexibility in a variety of ways (Tang, Wang, & Li, 2021), including shifting HVAC load in response to Demand Response (DR) events or variable highly dynamic electricity prices that discourage electricity use during peak hours. To ensure a comfortable environment while shifting the demand, the building controller needs to be aware of future building conditions and anticipate appropriate control actions. Model Predictive Control (MPC) can be an effective solution to this problem, since it can predict future building behavior based on a model with weather and price forecasts (Drgoňa et al., 2020). The steady rise in the number of simulation and experimental studies in the last two decades on this topic confirms the high potential of MPC for flexibility. However, MPC implementation still requires expert knowledge to setup and maintain, leading to uncertain economic benefits outside of research and development projects. This is particularly noticeable in medium to large commercial buildings, where more customized MPC solutions may be needed for each building. In addition, there remains a lack of real demonstrations showcasing price-responsive load shifting in medium to large commercial buildings through MPC, despite the increasing prevalence in utility programs.

The California Load Flexibility Research and Development Hub (CalFlexHub) initiative (CalFlexHub, 2024; Piette et al., 2022) is carrying out several experimental studies to demonstrate and analyze the capability of MPC for building flexibility at equipment (Mande, Ellis, et al., 2023), building (woo Ham et al., 2024) and district levels extending work from (Kim et al., 2022). This study, as part of CalFlexHub experimental efforts, focuses on medium to large commercial buildings. In particular, it extends the work carried out in (Blum et al., 2022), where an MPC aimed at minimizing energy consumption was deployed in a real office building at Lawrence Berkeley National Laboratory (LBNL). The work presented here aims the MPC at minimizing operating costs based on two dynamic price profiles (Gerke et al., 2024), one for Summer and one for Fall. Two sets of two-week tests were conducted in the two seasons, to show the



**Figure 1:** Design and control schematic in section view for one RTU of the UFAD HVAC system. Important sensor and control points are labeled, including temperatures (T) and static pressures (P).

MPC's ability to adjust the building demand in various price scenarios and weather conditions. This paper describes how the MPC was modified from the original formulation in (Blum et al., 2022) to take into account modifications in the underlying data and control infrastructure, as well as to fulfill the new optimization objectives. We present results from the field tests, and estimate demand flexibility metrics and cost savings for each test compared to baseline control in the building. The paper is structured in five sections. Section 2 summarizes the case study building. Section 3 summarizes the MPC implementation and updates from (Blum et al., 2022). Section 4 reports and discusses results. Finally, Section 5 provides conclusions.

## 2. CASE STUDY

This section describes a summary of characteristics of the Building and its HVAC system. A more in depth description on these systems is reported in (Blum et al., 2022).

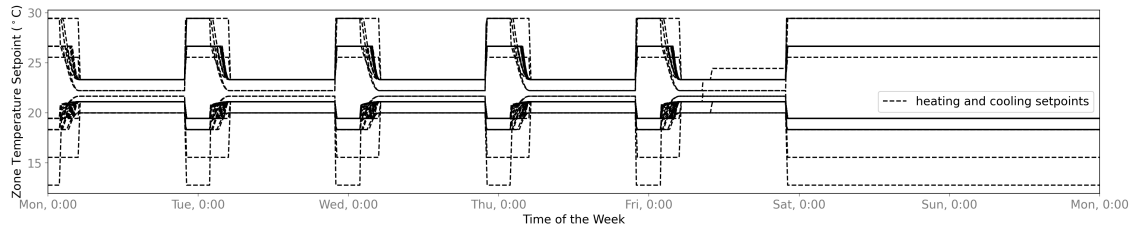
### 2.1 Building

Building 59 situated on the Lawrence Berkeley National Laboratory (LBNL) campus in Berkeley, CA, was built in 2015. Berkeley enjoys mild climate, falling under ASHRAE Climate Zone 3C. The MPC is applied to the HVAC system that serves the office area of the third and fourth floor, encompassing approximately 6,038 m<sup>2</sup> (65,000 ft<sup>2</sup>). The third floor predominantly comprises enclosed office areas, while the fourth floor is primarily open office space. The building structure is steel-framed with an exterior metal curtain wall system with integrated windows and foamed insulation core. In office areas, finished floors with carpeting are raised above structural concrete slabs, creating the plenum for the Underfloor Air Distribution (UFAD) system.

### 2.2 HVAC System

Heating, cooling, and ventilation are provided to the offices by a UFAD. The system uses four large rooftop units (RTUs) located on the roof with direct expansion (DX) coils to supply cool air to the common underfloor plenum on each level. Air diffuses directly to the core and perimeter zones through floor diffusers and can be pulled and heated to perimeter zones with fan-powered underfloor terminal units (UFT). Hot water for the UFT heating coils is produced by a 117 kW air-source heat pump (ASHP) located on the mechanical level, in the first floor of the building. The RTUs do not provide centralized heating to the spaces. This design is shown in cross-section for one RTU in Figure 1.

Each RTU serves the third and fourth floor offices between particular column lines, shown in Figure 1 as supply and return ducts, of the building, though the areas of service are not separated by internal wall partitions. The design air flow



**Figure 2:** Typical weekly schedule for all the setpoints in the different perimeter zones that have a thermostat. This weekly profile repeats for the whole year. Occupancy is considered from 5 a.m. to 10 p.m. on weekdays.

of each RTU is  $9.44 \text{ m}^3/\text{s}$ , with design minimum outdoor air flow rate of  $2.36 \text{ m}^3/\text{s}$  during occupied hours. For each RTU, the supply fan motors are 14.9 kW and the return fan motors are 5.6 kW, each equipped with variable frequency drives (VFD). The cooling capacity of each RTU is 104 kW with two 9.6 kW motor R410A scroll compressors with variable speed control of each compressor from 10% to 100%. There are 50 fan-powered underfloor terminal units (UFTs) with water heating coils to provide heating and reheat. The condenser water from the RTUs is cooled by cooling towers located next to the building on the mechanical level. These cooling towers are shared with the high performance computing cooling equipment, which dominates the load on the towers.

### 2.3 Baseline Control

The occupancy schedule for weekdays, from 5 a.m. to 10 p.m., determines the HVAC operation. To meet the minimum outside air requirement, during occupied times each RTU needs to provide a minimum airflow of  $2.36 \text{ m}^3/\text{s}$  or 5,000 CFM. The outside airflow set point is set to zero otherwise. The perimeter zone thermostat heating and cooling setpoints are also based on this occupancy schedule, as shown in Figure 2. The RTUs are turned off during unoccupied hours unless there is a call for heating or cooling as determined by the unoccupied set points.

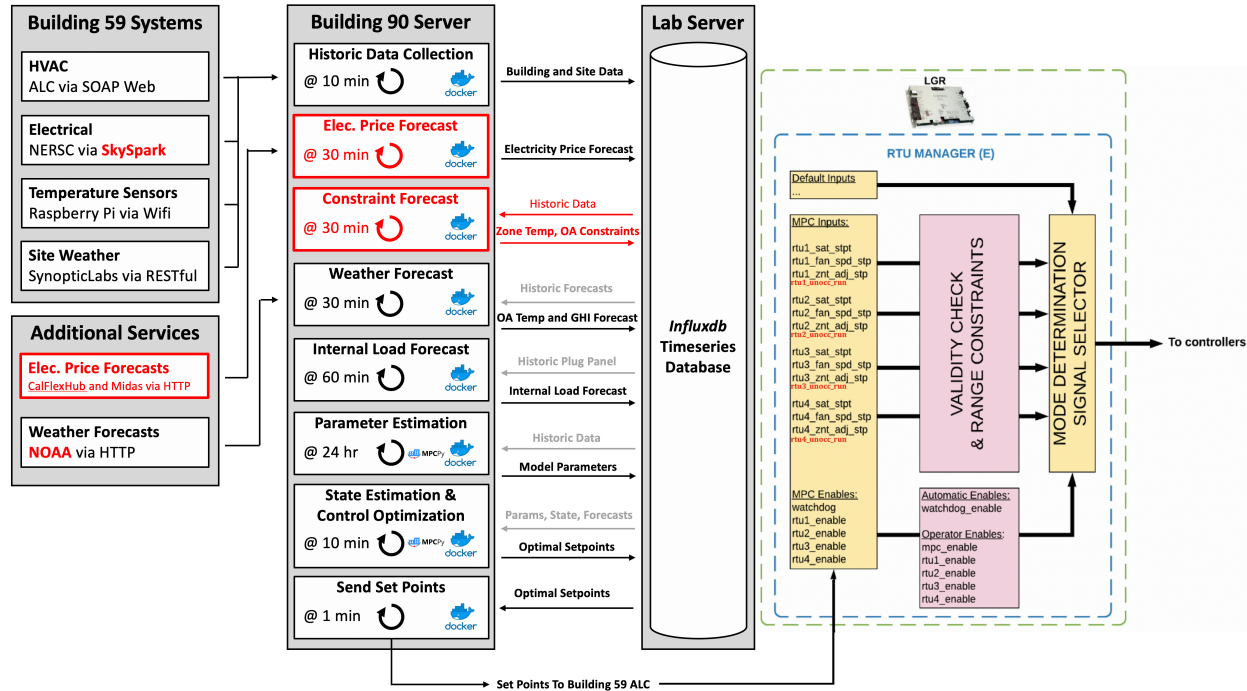
Perimeter zone thermostats control the UFTs fan speed for cooling and the hot water valve position for heating. However, to keep the noise of the UFT fans low for occupants and as a result of system commissioning, strict limits are placed on many UFT fan speeds typically ranging from a minimum of 20% and maximum of 50%. It is important to note that thermostats in the core zones of the building only trigger heating and cooling requests and do not control UFTs, as supply air diffuses directly from the underfloor plenums into core spaces. All four RTU supply fan speeds are controlled by a single PI controller taking a variable static pressure set point, between 3.75 and 12.5 Pa that depends on the number of zone cooling requests and the average underfloor plenum static pressure measurement of all units. The return fan speeds of each RTU are controlled to track a flow rate equal to 95% of the supply air flow rate minus an additional  $0.1 \text{ m}^3/\text{s}$ . The supply air temperature for each RTU is controlled to a variable set point, between 14.4 and 22.2 °C that depends on the number of zone heating and cooling requests, which is maintained by the DX coil and economizer when outside air dry bulb temperature conditions are appropriate. When not economizing, the outside air flow rate for each RTU is controlled to the design minimum outside air flow rate set point using outside air flow measurement stations in each RTU during occupied hours.

## 3. MPC IMPLEMENTATION

This section briefly summarizes the MPC implementation with a focus on updates from the study from (Blum et al., 2022). For a more in depth explanation on the MPC implementation please refer to (Blum et al., 2022).

### 3.1 Software Architecture, Data Exchange, and BMS Integration

The MPC software architecture is made up of eight main services that run in parallel, presented in Figure 3. Each service runs within its own Docker (*Docker*, 2024) container, though each container is based on the same Docker image. These services run on a dedicated server on campus and exchange data with various Building 59 and site data collection systems, web-based services, and the project database. This database is an instance of an Influxdb timeseries database (*InfluxDB*, 2024). The control optimization, state estimation, and parameter estimation services are implemented using the open-source software MPCPy (Blum & Wetter, 2017) commit ecb9833. This workflow utilizes a user-defined model written in Modelica along with user-defined data specified in a Python environment to automatically formulate and solve the appropriate optimization problems for each of the control optimization, state estimation, and parameter estimation services.



**Figure 3:** Schematic of the MPC software architecture, data integration, and building BMS communication, with updates from the previous study highlighted in red.

Figure 3 also shows the different systems and methods where data is collected. Separate Python modules were developed to collect data from each building system, since each building system requires use of a separate API or interface. The Historic Data Collection service runs on a single Python module that centralizes the use of each of the system-specific modules to collect all system data. The MPC solution is calculated by the Control Optimization service and stored in the database every 10 minutes. Every one minute, the most recent valid solution setpoints are sent to the building Building Management System (BMS) via the Send Set Points service. The integration logic between these setpoints and the Automated Logic Corporation (ALC) BMS is shown in Figure 3.

Highlighted in red in Figure 3 are the updates with respect to the previous study (Blum et al., 2022). In particular:

- The electrical plug load data previously obtained via an Elastic Search (ES) database Query are now obtained via a campus Skyspark API.
- A service to retrieve day-ahead electricity prices was implemented to obtain data from the CalFlexHub and MIDAS servers via HTTP. The electricity prices are used in the MPC objective function to estimate the cost.
- The weather forecast service was changed from using DarkSky (no longer free access) to the National Oceanic and Atmospheric Administration (NOAA) API.
- Since the study in (Blum et al., 2022) was carried out, the building facility managers implemented scheduling that reset zone temperature setpoints and outside air requirements based on the occupancy schedule. Therefore, an additional service was added to the MPC to retrieve historical temperature setpoints, outside air requirements, and occupancy schedule, and use them to define constraints in the MPC formulation.
- During the same control update, the facility managers implemented scheduling that turned off the HVAC system during unoccupied hours. This increased the MPC control space, allowing it to turn off the RTUs during unoccupied hours and turn them back on to pre-heat or pre-cool the building. For this reason, an additional setpoint, *unocc\_run*, was added to the list of points controlled by MPC. The facility staff implemented this change in ALC.

### 3.2 Modeling and Parameter Estimation

The models described here are used by the MPC controller in real time to predict the performance of the building and optimize control. Important values to predict are building space temperatures, RTU air flow rates and temperatures, fan power consumption, DX compressor power consumption, and UFT heating demand. All the models used for the MPC can be considered grey-box, utilizing simplified, yet physically-based, models implemented in Modelica for use with MPCPy. In particular, an equivalent resistance capacitance circuit (R2C2) was used to model four thermal zones, one for each RTU. Field data was used to estimate the parameters of each model. Each RTU zone corresponds to a collection of several perimeter and core zones attributed to being served by the RTU. The overall temperature is estimated by weighting all the corresponding zones. The RTU model contains a mixing box, supply and return fans, and DX models. The mixing box mixes outside air and return air in a ratio specified by the outside air fraction control and supply air massflow rate signals. The supply fan follows a specified flow rate setpoint. As a model simplification, the return fan is assumed to move an equal amount of air. The power consumption of the fan is estimated as a cubic function of the air massflow rate. The DX is modeled as a simple heat exchanger removing sensible heat from the air stream. Its power consumption is estimated as a quadratic function of the heat flow rate.

The models parameters for the RC network, fan and DX power consumption curves, are considered time-invariant for each iteration of MPC optimization. Due to the simplification described above, real time-invariant parameters may not be able to properly characterize the behavior real system in all its operating points. To account for this mismatch, as well as other seasonal variations in performance, the model parameters for envelope and fan are updated daily. The new parameters are calculated with an optimization problem that minimizes the error between measurement and model prediction for the most recent data, split between training and validation periods, using up to a 20 days time period for the envelope, with minimum of one week depending on data quality, and 15 days for the fan, with a minimum of three. Because there are no direct measurement of the DX coil and the training process was more tailored, these parameters were determined only once with operational data from August 1, 2020 to October 18, 2020.

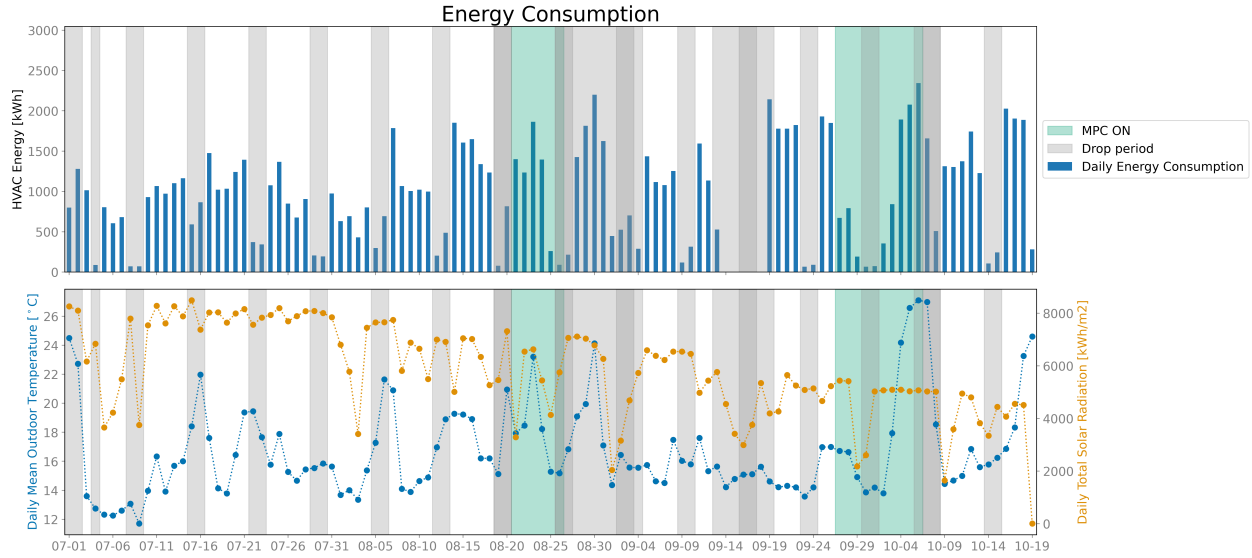
### 3.3 Forecasts

Weather forecasts are collected from NOAA (*National Oceanic and Atmospheric Administration*, 2024) and include drybulb temperature, humidity, dew point, pressure, wind speed, wind direction, and cloud cover. The global horizontal irradiation (GHI) is not provided by the station and it is instead estimated from cloud cover data, drybulb temperature, and clear-sky irradiation. This estimation is obtained using a data-driven model trained on historical data obtained from a weather station situated on LBNL's campus. The forecast data is provided with hourly time steps for a 24-hour period and updated approximately every 6 hours. The dynamic electricity price is collected from the CalFlexHub (*CalFlexHub*, 2024) and (*MIDAS*, 2024) servers. It consists of a 24-hour profile with a 15 min time step. There are four typical profiles, one for each season. These are synthetic profiles adapted from those defined by (Gerke et al., 2024) which use an approach similar to the CalFUSE program (Madduri, Gupta, Foudeh, & Phillips, 2022) to study demand response potential across California. The two profiles used in this study are shown in blue in Figure 5.

Internal loads consist of various internal heat gain sources present in the building such as plug load, lighting, and occupant heat gains. Plug load was used as a proxy of total internal loads and a simple data-driven model is used to predict the plug load based on the hour of the day and day of week. The model uses the average value of the same hour of day and the same day of week of the previous 90 days. Temperature setpoint forecasts, used to define operational constraints in the MPC control problem formulation, are based on the weekly schedule shown in Figure 2 established by the building operators. Temperature setpoints correspond to cooling and heating setpoints, creating a comfort band. Each zone with a thermostat is assigned a weekly schedule and from there a weighted average is obtained for each of the aggregated zones used by the MPC. The other occupancy-based set point is the minimum outside airflow requirement.

### 3.4 Control optimization and state estimation

The optimal control problem is defined over a 24-hour horizon. The objective is minimizing HVAC operating cost, calculated as the sum of the electrical consumption of supply and return fans, DX compressor, and UFT reheat power multiplied by the dynamic electricity price, while keeping the zone measured temperature between cooling and heating setpoints. These set points are implemented as soft constraints with slack variables in the objective function. Additional soft constraints include regularization terms on all control variables to avoid oscillatory behavior and the lower limit of the supply air temperature set point, 17 °C, to ensure problem feasibility in the case the outside air temperature is too low to maintain the minimum supply air temperature and minimum outside air flow requirement. Finally, there are hard constraints on the upper bound of the air supply temperature, 21.1 °C, and the minimum outside airflow requirement,



**Figure 4:** HVAC daily energy consumption (top) and daily average external dry bulb temperature and daily horizontal radiation (bottom). Highlighted in green is the MPC testing period and in grey the days that were removed from the test.

2.36  $m^3/s$  during occupied hours and zero otherwise.

The state estimation consists of a simple moving horizon estimator. The estimator solves an optimization problem from the start of historic time horizon (previous 24h) to the current time, where the objective is minimizing the error between measured and estimated zone temperature and the optimization variables are the state temperatures at the initial time. Then, the state temperatures at current time are used to initialize the optimization at the current time step.

## 4. RESULTS

### 4.1 Flexibility Showcase

This section showcases the results obtained during the tests performed in the Summer and Fall 2023. Figure 4 shows the overall period under analysis (7/1-10/19) and the MPC testing periods (8/21-8/26 and 9/27-10/6) highlighted in green. The grey shaded areas are periods not considered in the results analysis. This is to avoid weekends where the HVAC was mostly off and periods where HVAC operation was abnormal due to faults or data failures that inhibited the ability for the MPC to operate continuously. Figures 5a and 5b show that the MPC successfully shifted demand from periods with high electricity prices, to periods with lower prices. This behavior clearly differs from the baseline. In fact, calculating the Energy Shed,  $E_{shed}$ , defined as

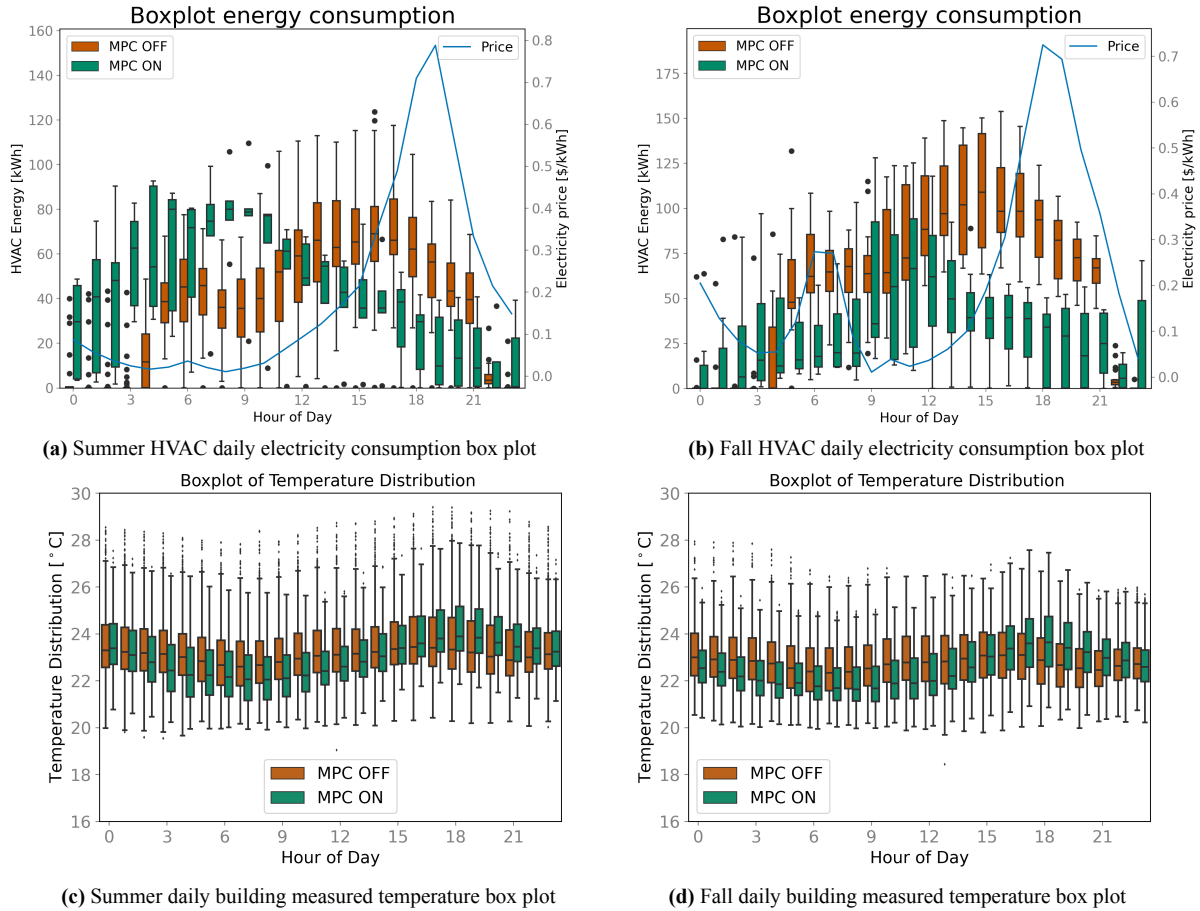
$$E_{shed} = \frac{Energy_{MPCON_{peakhours}}}{Energy_{MPCOFF_{peakhours}}} \quad (1)$$

Where  $n_{peakhours}$  are the 6 hours when the electricity price is the highest during the day, and  $Energy_{MPCON_{peakhours}}$  and  $Energy_{MPCOFF_{peakhours}}$  are the HVAC energy consumption during the peak hours for MPC and baseline. The MPC was able to reduce HVAC energy use during these peak hours by 41% and 32% in the Summer and Fall respectively. Furthermore The HVAC  $E_{shed}$  can be expressed as an average demand decrease *Intensity*, defined as

$$Intensity = \frac{Energy_{MPCOFF_{peakhours}} - Energy_{MPCON_{peakhours}}}{n_{peakhours} Area_{building}} \quad (2)$$

and it can be interpreted as the average power shed during peak hours per square meter, with values of 5.5 and 9.6  $W/m^2$  for Summer and Fall. The effect of the demand shift is reflected in the average internal temperature distribution of the building shown in Figure 5c and 5d.

The MPC tends to pre-cool the building, especially in the Summer during nighttime unoccupied hours, so it can reduce HVAC consumption during the peak hours while maintaining similar temperature levels with respect to the baseline



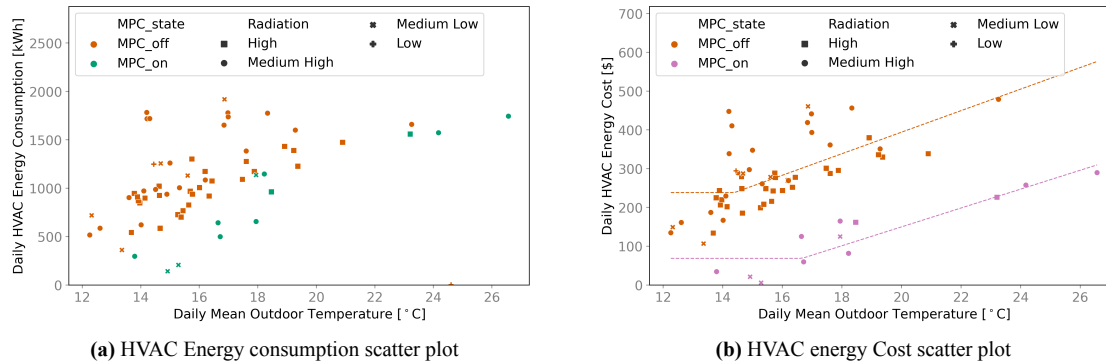
**Figure 5:** Figures a) and b) show the hourly distribution of HVAC energy consumption for each day of the analysis period and figures c) and d) show the building measured temperature hourly distribution, including all measured zones, for the Summer test period (7/1-8/27) and Fall test period (9/4-10/19), with excluded days shown in Figure 4. In addition, on the right y-axis of Figures a) and b) is reported the dynamic electricity prices used for Summer and Fall.

controller. Furthermore, looking at the box plots' whiskers, the MPC seems to keep a tighter temperature boundary with respect to the baseline controller. The only exception is late afternoon in the Fall as shown in Figure 5 plot d), where the edge of the whiskers are above the baseline controller. The rationale behind this can be observed in Figure 4 specifically in the lower plot. Towards the end of the MPC test (10/4-10/6), the outdoor air temperature reached unusually high levels, causing the HVAC system to face challenges in maintaining temperatures within acceptable limits.

Figure 6 shows daily energy and cost as a function of daily mean outside air temperature under baseline and MPC control. Figure 6b shows the daily electricity cost of the HVAC system for the testing period (7/1-10/19) for the MPC and baseline controllers. Under MPC control, the average daily cost is half of that under baseline control. This is calculated by taking the ratio between the MPCOFF and MPCON average daily costs for the testing period as shown in Equation 3.

$$CostSavings = \frac{\sum_{k=1}^{n_{daysMPCON}} DailyCost_{MPCON}}{\sum_{k=1}^{n_{daysMPCOFF}} DailyCost_{MPCOFF}} \quad (3)$$

The decrease in HVAC electricity cost does not come at the expense of increased energy consumption. In fact, as shown in Figure 6a, the average daily MPC energy consumption is approximately 10% lower than the baseline controller. The calculation for this percentage was carried out in a similar fashion as the HVAC cost.

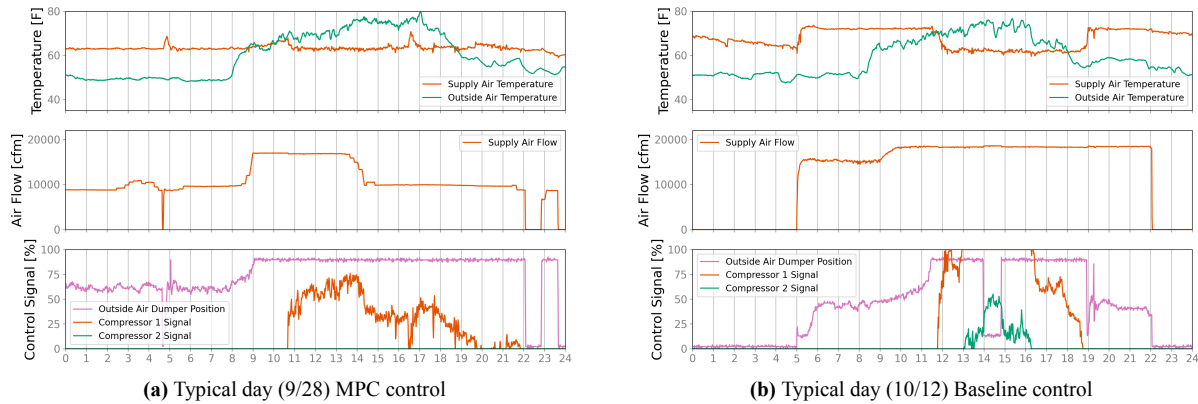


**Figure 6:** Figure a) and b) show respectively the daily HVAC energy consumption and cost plotted against the daily mean outdoor temperature. The excluded days and plotted days for MPC on and off are shown in Figure 4. The days are further divided between different intensity of daily solar radiation by using different shapes. On plot b) are also shown piece-wise linear regression lines for MPC ON and MPC OFF (dashed lines).

Looking at typical day RTU operation in the MPC and baseline controller cases in Figure 7 helps to better understand how the MPC achieves such results. Figure 7a shows the typical day (9/28) operation for Summer and Fall for the MPC and Figure 7b for the baseline controller (10/12) for RTU1. These two days were chosen because they have a comparable outdoor temperature profile and radiation. Also, while all RTU operate in a similar fashion, RTU1 was chosen as representative. Looking at the top plots of Figure 7, the first noticeable difference is the air supply temperature under MPC control being lower than it is under baseline control in the morning hours. Here, the MPC takes advantage of the lower external temperature during the night to pre-cool the building without turning on the compressor and while electricity prices are cheap. It does so even overnight by turning on the RTUs with modest fan speed. During occupied hours and peak price hours, the pre-cooling allows the MPC to satisfy the building cooling load at reduced supply and return fan speeds, which greatly contribute to the HVAC overall electricity consumption. The combined effect of pre-cooling and lower supply air flow during the central hours of the days also allows the MPC to make less use of the direct expansion coil, as shown in the bottom plot of Figure 7.

#### 4.2 Discussion

While the new MPC implementation was effective in shifting energy and reducing cost, updating the software from (Blum et al., 2022) required significant effort. This was quantified in around 4 person months of full time work, with 50% dedicated to integrating new data streams and replacing those that were no longer available (details in Section 3.1). The remaining 50% was spent in a variety of tasks, including 1) monitoring the data quality for the streams necessary to properly operate the MPC, especially during testing periods, 2) checking the validity of the MPC results before field tests, 3) updating and testing the MPC-ALC interface, 4) addressing unforeseen circumstance, such as modifying the code to incorporate a new mode of operation for wildfires. In this mode, adjustments were made to the MPC constraints regarding the minimum outside air damper position. Specifically, the minimum was reduced to zero, and an upper limit was implemented to restrict outside air intake during wildfires. Furthermore, 5) the BMS control logic and configuration had to be updated to allow MPC to turn RTUs on during unoccupied hours, without repercussions on other parts of the control logic. This was only possible through close collaboration with facility managers. A lesson learned from the experiment was that as the number of underlying data sources and streams increases, so do the chances of failures and data interruptions. These consideration raise a question of economic feasibility of MPC, especially for large commercial buildings. In such buildings, that are typically heterogeneous and complex, it is harder to reduce labour time required to setup the MPC and continuously updating it. To tackle this matter and streamline commercial MPC, the research community should establish target or acceptable implementation cost ranges based on the cost savings achievable in various scenarios. Yet, there is a lack of comprehensive experimental studies in literature that estimate the yearly cost savings for large commercial buildings and under dynamic electricity prices. For this reason, the authors have been carrying out additional tests in Winter and Spring seasons at the same site and analyzing past literature to start defining values for this target cost for different types of buildings and HVAC systems.



**Figure 7:** RTU1 example day plot. Supply (orange) and outside (green) air temperatures (top), supply (orange) air flow rate (second), compressor 1 (orange), compressor 2 (green) and outside air dumper (0 closed, 100 open) control signals (third) under MPC ON (a) and MPC OFF (b) control. MPC ON is on 9/28 and MPC OFF is on 10/12.

## 5. CONCLUSIONS

This study demonstrated the ability of the developed MPC to shift HVAC energy consumption in a commercial building in response to two different dynamic electricity price profiles. The MPC reduced the daily HVAC energy cost by an average of 50% with respect to a conventional control strategy. Thermal comfort was maintained during MPC operation, the temperature range in zones was tighter than conventional control, and building operators indicated there were no complaints by occupants during the MPC tests. Maintaining the MPC platform running and ready to be operational required significant continuous effort due to issues that included data stream interruptions, server restarts, and software updates, making data management a significant issue to address. Lastly, even though the results of this work are encouraging to demonstrate the ability for MPC to provide demand flexibility, they still do not depict a full picture of building operation under MPC control for the whole year, which is necessary to calculate the viability of MPC, and it is the subject of future work.

## NOMENCLATURE

ALC	Automated Logic Corporation
API	Application Programming Interface
ASHP	Air Source Heat Pump
BMS	Building Management System
DX	Direct Expansion
DR	Demand Response
ES	Elastic Search
GHI	Global Horizontal Irradiation
HVAC	Heating, Ventilation, and Air Conditioning
LBNL	Lawrence Berkeley National Laboratory
MPC	Model Predictive Control
NERSC	National Energy Research Scientific Computing Center
NOAA	National Oceanic and Atmospheric Administration
RC	Resistance-Capacitance
RTU	Roof-Top Unit
UFAD	Underfloor Air Distribution
UFT	Underfloor Terminal Unit
VAV	Variable Air Volume
VFD	Variable Frequency Drive

## REFERENCES

- Blum, D., Wang, Z., Weyandt, C., Kim, D., Wetter, M., Hong, T., & Piette, M. A. (2022). Field demonstration and implementation analysis of model predictive control in an office hvac system. *Applied Energy*, 318, 119104.
- Blum, D., & Wetter, M. (2017). MPCPy: An open-source software platform for model predictive control in buildings. *Proceedings of the 15th IBPSA Conference*, 1381-1390.
- Calflexhub. (2024). Retrieved from <https://calflexhub.lbl.gov/>
- Docker. (2024). Retrieved from <https://www.docker.com>
- Drgoňa, J., Arroyo, J., Figueroa, I. C., Blum, D., Arendt, K., Kim, D., ... others (2020). All you need to know about model predictive control for buildings. *Annual Reviews in Control*, 50, 190–232.
- Gerke, B. F., Stuebs, M., Murthy, S., Khandekar, A., Cappers, P., Brown, R. E., & Piette, M. A. (2024). Potential bill impacts of dynamic electricity pricing on california utility customers. doi: <https://doi.org/10.20357/B7WK65>
- Influxdb. (2024). Retrieved from <https://www.influxdata.com/time-series-platform/>
- Kim, D., Wang, Z., Brugger, J., Blum, D., Wetter, M., Hong, T., & Piette, M. A. (2022). Site demonstration and performance evaluation of mpc for a large chiller plant with tes for renewable energy integration and grid decarbonization. *Applied Energy*, 321, 119343. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0306261922006894> doi: <https://doi.org/10.1016/j.apenergy.2022.119343>
- Madduri, A., Gupta, A., Foudeh, M., & Phillips, P. (2022). *White paper: Advanced strategies for demand flexibility management and customer der compensation* (Tech. Rep.). California Public Utilities Commission, Energy Division. Retrieved from <https://www.cpuc.ca.gov/-/media/cpuc-website/divisions/energy-division/documents/demand-response/demand-response-workshops/advanced-der---demand-flexibility-management/ed-white-paper---advanced-strategies-for-demand-flexibility-management.pdf>
- Mande, C., Ellis, M. J., et al. (2023). Supervisory multi-objective economic model predictive control for heat pump water heaters for cost and carbon optimization. *ASHRAE Transactions*, 129.
- Midas. (2024). Retrieved from <https://midasapi.energy.ca.gov/>
- National oceanic and atmospheric administration. (2024). Retrieved from <https://www.noaa.gov/>
- Piette, M., Liu, J., Nordman, B., Smith, S., Brown, R., & Pritoni, M. (2022). Accelerating decarbonization with the california load flexibility research and deployment hub.
- Satchwell, A., Piette, M. A., Khandekar, A., Granderson, J., Frick, N. M., Hledik, R., ... others (2021). *A national roadmap for grid-interactive efficient buildings*. Retrieved from <https://gebroadmap.lbl.gov/A%20National%20Roadmap%20for%20GEBs%20-%20Final.pdf>
- Tang, H., Wang, S., & Li, H. (2021). Flexibility categorization, sources, capabilities and technologies for energy-flexible and grid-responsive buildings: State-of-the-art and future perspective. *Energy*, 219, 119598.
- U.S. DOE. (2024). *Decarbonizing the u.s. economy by 2050: A national blueprint for the buildings sector*. Retrieved from <https://www.energy.gov/eere/articles/decarbonizing-us-economy-2050>
- woo Ham, S., Paul, L., Kim, D., Pritoni, M., Brown, R., & Feng, J. D. (2024). Decarbonization of heat pump dual fuel systems using a practical model predictive control: Field demonstration in a small commercial building. *Applied Energy*, 361, 122935.

## ACKNOWLEDGMENT

This research was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technologies of the U.S. Department of Energy, under Contract No. DE-AC02-05CH11231, and by the California Energy Commission and the California Load Flexibility Research and Deployment Hub. The authors would like to thank Christopher Weyandt, Norm Bourassa, John Elliot, and Raphael Vitti from Lawrence Berkeley National Laboratory, USA very much for their support of this study, providing access to facilities and data, and help in the overall implementation and testing of the MPC controller in Building 59.