HP-FLEX: Field demonstration of the semantics-driven configuration of a Model Predictive Control system to make heat pumps flexible

Lazlo Paul
Affiliate ASHRAE

Sang woo Ham, PhD
Affiliate ASHRAE

Anand Prakash
Student Member ASHRAE

Armando Casillas
Affiliate ASHRAE

Tao Yang, PhD
Affiliate ASHRAE

Marco Pritoni, PhD
Full Member ASHRAE

ABSTRACT

Model Predictive Control (MPC) has demonstrated significant potential for optimizing building operations and enabling demand flexibility. However, the widespread adoption of MPC is hindered by complex manual configuration and commissioning processes that must be conducted by control experts working alongside building operators. These challenges drive up costs and reduce scalability, particularly when technical human resources and building automation systems are limited, such as in small and medium commercial buildings (SMCBs). This paper demonstrates how semantic standards, specifically ASHRAE 223P, can accelerate the adoption of MPC applications for load flexibility in SMCBs. The authors present a replicable control framework, titled "HP-FLEX" that leverages a building's semantic model to bootstrap the required data configuration for an MPC controller developed for optimizing heat pump systems as flexible grid resources. The semantic model helps streamline the deployment workflow, particularly for site setup, data/control commissioning, and model setup. This integration enhances portability, transferability, and scalability of the HP-FLEX MPC, which has been developed to support MPC-based supervisory HVAC controllers in SMCBs. Additionally, the paper details the required building and thermostat metadata information to enable the HP-FLEX MPC based on field demonstrations. The new workflow was tested in a small commercial building located in California, U.S., and demonstrated a load shifting performance of 9% based on a dynamic pricing signal that varies by the hour. This work provides a practical pathway for transitioning sophisticated building applications from custom to standardized semantic representations, supporting the broader adoption of advanced control strategies like MPC. The framework also establishes a foundation for evolving metadata requirements as applications mature while maintaining compatibility with industry standards.

INTRODUCTION

Flexible electricity loads, particularly those within buildings, offer a promising avenue for mitigating grid stress and enhancing resilience (Neukomm et al. 2019). To shift electrical loads within buildings, advanced control approaches such as model predictive control (MPC) or reinforcement learning (RL) have been investigated (Touzani et al. 2021; Kim and Braun 2022). These approaches are typically applied for large commercial buildings with centralized building automation systems (Blum et al. 2022) rather than small and medium commercial buildings (SMCBs), which share 50% of the total floor area of all commercial buildings (EIA 2018). There are several barriers for deploying advanced controls in SMCBs. Most SMCBs use packaged HVAC systems that do not provide the data or control capabilities needed for these controls. Additionally, even though total floor area and energy use is similar to the totals for large commercial buildings, each SMCB has a lower budget for energy projects, meaning MPC solutions targeting SMCBs must be significantly less expensive to deploy.

Several recent studies have developed and demonstrated low-cost and practical MPC solutions that require only networked thermostats for each unit. Kim and Braun's (2018) MPC solution provided about 12% of energy and 18% of peak

Lazlo Paul, Sang woo Ham, Anand Prakash, Armando Casillas, Tao Yang, and Marco Pritoni are researchers at Lawrence Berkeley National Laboratory

demand through RTU coordination for several months of trials (Kim and Braun 2018). Kim and Braun integrated a load shifting feature into the MPC solution in a hierarchical MPC structure for ON/OFF staged packaged units in a laboratory building (Kim and Braun 2022). The solution showed 30% of demand cost savings and 40% of on-peak demand cost savings with less than 10% of total energy cost savings. The extended version of this MPC solution, called HP-FLEX MPC, now includes more systems applicable to SMCBs, such as dual fuel heating systems (i.e., those including heat pumps and gas furnaces) and variable refrigerant flow (VRF) systems (Ham et al. 2024).

While the HP-FLEX MPC is designed to be applicable for the vast majority of SMCBs, deployment of this technology can not be rapidly scaled due to two key challenges. First, the implementation of the MPC solution requires holistic understanding of the MPC technology and the deployment site for engineers, which is a key barrier in the deployment process for MPC (Ham et al. 2024). The second key challenge is the high cost of system integration. A recent study showed that approximately 20% of the total implementation time is spent on identifying and connecting the building data and control points to the MPC controller. Significant expertise and effort is expended in this process because of the heterogeneity of buildings and the lack of consistent documentation, which means experts in the MPC must piece together sparse information about the buildings and perform onboarding for each new site (Blum et al. 2019).

In the last decade, various semantic ontologies have emerged to address these challenges. These ontologies provide standard definitions of the concepts needed to model building systems to enable semantic interoperability (Pritoni et al. 2021). Semantic models made using these ontologies enable data integration from different sources while providing machine-readable understanding of the meaning of this data (Delgoshaei et al. 2022). This enables the development of more portable applications, which can be configured using semi-automated processes (Pritoni et al. 2024). Semantic ontologies such as Brick (Balaji et al. 2018) and the proposed ASHRAE Standard 223P (ASHRAE 2025) have been developed to support this automation in the deployment of building applications.

Recent literature has shown how semantic models can enable the development of portable applications that can be deployed across various buildings without significant manual reconfiguration. de Andrade Pereira et al. used semantic models to facilitate the development of portable rule-based demand flexibility applications (de Andrade Pereira et al. 2024) and Paul et al. proposed a semantics-driven middleware platform using ASHRAE 223P (Paul et al. 2025). Both of these works focus on rule-based controls simpler than MPC. Wan et al. designed a system identification for MPC using multiple semantic ontologies (Wan et al. 2025), which are also described in Chamari et al (Chamari et al. 2024). While Wan et al. presented a scalable data model focused on a part of MPC using real data, it uses multiple ontologies and a workflow relying on detailed information from building information models, which may not exist for many current buildings, particularly SMCBs and would be costly to put together. Prakash et al. reviewed the use of ontologies for MPC and noted that it is unclear if existing semantic ontologies represent the information necessary for MPC, posing a barrier to deployment (Prakash et al. 2024). Additionally, most existing building applications use proprietary or application specific representations of metadata. It is unclear how these applications may be updated to use recent standardized semantic metadata models, especially if these models are complex and utilize technology unfamiliar to the building industry, posing a barrier to the adoption of standards such as ASHRAE 223P.

The objective of this paper is to present a semantic modeling approach and software architecture used to upgrade an existing MPC controller (the HP-FLEX MPC) to use standardized semantic models. This approach is intended to show how semantic models can be used to reduce the expertise and labor required for the deployment of MPC, thus improving its scalability. Additionally, this approach demonstrates a method for upgrading an existing controller using an application-specific data modeling approach to using semantic models such as ASHRAE 223P, without requiring the user to have in-depth understanding of ASHRAE 223P or ontologies. This approach is also applicable to the design of new controllers that can natively use semantic models. This paper will introduce the HP-FLEX MPC technology and the workflow required for its deployment. Then it will present how semantics can be used to improve this workflow with a reference software architecture that has been tested on two different semantic ontologies, ASHRAE 223P and Brick. Finally, this paper will present a preliminary performance evaluation of the semantics-driven HP-FLEX MPC for a real site and a discussion on lessons learned.

HP-FLEX MPC

The HP-FLEX MPC is an advanced control strategy designed to optimize the operation of multiple heat pumps by strategically shifting energy consumption away from peak price time. Given the building gray-box model, weather forecast, and future price signal, the HP-FLEX MPC finds the optimal runtime fraction (RTF) trajectories of HPs to minimize the summation energy cost (i.e., HP power times electricity price), peak demand, and comfort violation term as a soft constraint during the prediction horizon. The HP power is modeled as constant, making the optimization problem linear. At each time step (30 minutes), the optimal RTF trajectories of HPs are found, and the predicted temperatures based on the optimal trajectories are used for the optimal setpoints. Figure 1 illustrates the conceptual difference between schedule-based and HP-FLEX MPC control during the cooling season. The schedule-based control widens the setpoint ranges during the unoccupied times to save energy, but it causes a peak in energy demand at the beginning of the occupancy period, especially when the multiple HPs simultaneously turn on to cool down the spaces. In contrast, utilizing weather forecasts and building thermal models, the HP-FLEX MPC pre-cools the building, reducing the morning peak and shifting energy outside of the high price period. This is achieved while maintaining thermal comfort.

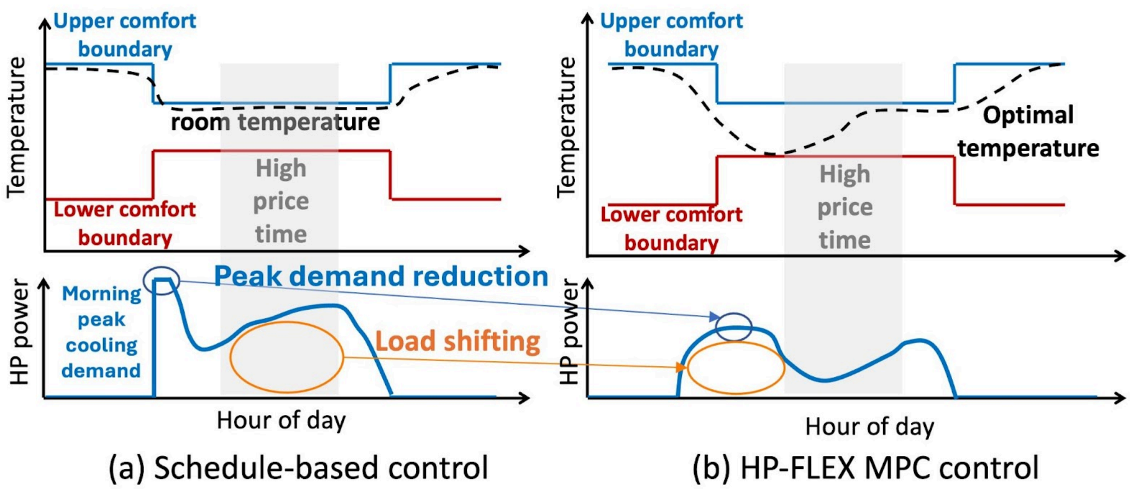

Figure 1: Conceptual diagram of HP-FLEX MPC for load shifting and peak reduction.

HP-FLEX MPC METADATA REQUIREMENTS

The predictive modeling within the MPC controller requires metadata spanning multiple domains including HVAC, controls, architecture, data storage, utility bills and occupant comfort. The first four domains - HVAC, controls, architecture, and data storage - are included in the semantic model because they represent the vast majority of data requirements for advanced building applications such as MPC. These domains describe the building topology and references to time series data from sources and actuators in the building and they provide the means to retrieve and contextualize data from the building, which is one of the most time consuming tasks in the workflow for instantiating MPC. The concepts from these domains are also well-represented by semantic ontologies such as ASHRAE 223P and Brick.

The HVAC metadata required by the MPC focuses on simple information regarding the type of HVAC system (e.g. Heat Pump Rooftop Units (HP RTUs) or mini-splits) and simple nameplate information including rated heating and cooling capacities and coefficients of performance under rated conditions. These simple elements of information allow the MPC to determine heat flows and electricity consumption using limited thermostat data.

The controls metadata includes information about the available control and datapoints (e.g. temperature setpoints), type of control used (e.g. On/Off, multi-stage), and other parameters (e.g. thermostat hysteresis, setpoint resolution). This data is required to convert data into and out-of the specific formats needed by the MPC. This ensures that input data can be utilized by the HP-FLEX MPC and that the MPC's output can be written back to the thermostat in a supported format. The MPC requires a minimum of three data points for control: thermostat temperature (the measured temperature at the thermostat), cooling status (the active stage of cooling), and the heating status (the active stage of heating). Additional data, such as outdoor weather, is used for system identification, and other points are used for monitoring the building and MPC.

Architecture metadata is classified into four different kinds: site, space, window, and HVAC zone, which is associated with the HVAC system and controls metadata. The site metadata includes the site name, timezone, location, and

National Oceanic and Atmospheric Administration (NOAA) station code for weather forecasting and estimation of outdoor air temperature if no sensors are present. The space metadata includes the area of each space, and the window metadata includes the area, tilt, and azimuth of the window, used for system identification of the building thermal model (Ham et al. 2024). The HVAC zone information links the architectural metadata to the HVAC and controls metadata, indicating the spaces and windows in a zone, and which controller and HVAC unit serve each zone. Zones are generally conditioned by a packaged HVAC unit and some controller, commonly a thermostat.

References to the time-series data for each data point and the endpoints for actuator commands are included in the database metadata. This metadata serves to map the database ids to the semantic model so that historical data can be queried and used by the MPC. ASHRAE 223P has a concept for external references that is extended by the ref-schema to include more kinds of references (Ref Schema 2025). Comfort boundary and price signal information consist of time-series data for comfort temperature ranges and electricity price information. Currently, we use existing CSV- or JSON-formatted data that was originally developed for the HP-FLEX MPC before the semantic integration. Metadata for these concepts is not currently represented by building ontologies because it is used in advanced demand-flexible controls but is uncommon in other applications (Prakash et al. 2024).

MPC TRADITIONAL DEPLOYMENT PROCESS

The process for deploying the MPC without using semantic models is shown in Figure 2. This workflow generally involves three different roles – the facility manager, the building engineer, and the MPC expert. The facility manager handles long-term operation of the building. The building engineer has understanding of control networks and mechanical systems and supports setup of a control network and controls platform enabling the MPC, which may be a Building Automation System, Energy Management and Information System, or other middleware. The MPC expert has a complete understanding of the MPC and what it needs to operate. Different business models may evolve supporting the deployment of the MPC involving different types of organization, which is especially likely because of the range of ownership structures in SMCB, but these roles serve as a likely starting point based on the experience of the research team.

The first step consists of building metadata collection and a site inspection that generally involves all the roles. During this phase, the HVAC and architecture metadata is collected. The MPC expert must review the installed equipment and determine if the MPC can be applied, or if there must be changes in the MPC formulation to fit the building. Collected metadata is stored in MPC specific configuration files, which will later be used to set up the MPC, but are difficult to use for setting up other applications, such as data dashboards used for monitoring during MPC operation.

Next, the control devices and controls platform are set up. This task generally involves the facility manager and building engineer. This step may involve setting up networking for BACnet devices or smart thermostats, using point lists and device manuals to set up data collection, and setting up a control platform for data storage and control. The building engineer and facility manager try to produce and collect as detailed information as possible on the data available from the control devices, including point names and descriptions, units, whether points are writable, and data resolutions. Information based on device documentation can often be incorrect. For example, points indicated as writable may not actually be writable, or the resolution for writing and reading points may differ from the documentation. For this reason, there is generally also a data commissioning step.

Data commissioning consists of checking available data for errors or unique features. Errors may include incorrect configuration of BACnet networks or device drivers, which may be able to be corrected by the building engineer, such as misconfigured read and write capabilities. Other issues may lead to changes in the MPC configuration. Mechanical or networking issues like network instability may require other methods to be addressed. This step is critical for determining how the MPC can write data to the network. Because of the variation in building control methods, the MPC expert must then review the collected metadata and determine which data points map to the variables used by the MPC, and understand how to map the MPC outputs to the available control points. This process is time consuming and requires an in-depth understanding of MPC and building systems. The knowledge for identifying and mitigating these issues is gained with experience, but may be compiled into fault detection rules that can ameliorate each subsequent deployment of the MPC.

After data commissioning, the configuration files defined in the previous stages can be used to configure the MPC. During this process, several steps are required including system identification for the building model, validation of the building model for performance and reliability, and finally deployment of the MPC. This process is well defined in previous work (Ham et al. 2024). It is theoretically possible to automate these steps, but system identification and configuration of the MPC can be complicated due to the heterogeneity of buildings and the relationship between thermal dynamics and confounding factors in measurement like sensor placement and accuracy. This work is done most effectively by an MPC expert.

Finally, during MPC operation monitoring is required. Monitoring generally relies on alerts and data dashboards. Alerts determine fault conditions affecting the software, such as data outages, or faulty conditions of the HVAC system, such as insufficient heating or cooling service. It is rare that SMCBs have dedicated facility staff, so alerts provide a method of addressing faults before occupant comfort is significantly affected. Data dashboards organize and present building data in an easy-to-understand and interactive way, allowing MPC performance to be continually assessed. These dashboards also enable fault detection and diagnosis of the HVAC system. Alerts and data dashboards are useful features for improving building operation, and are often deployed as standalone applications using EMIS platforms (Lin et al. 2022). Setting up these dashboards and alerts is a time consuming process on its own, which relies on additional data collection and data mapping beyond what is required for the MPC application.

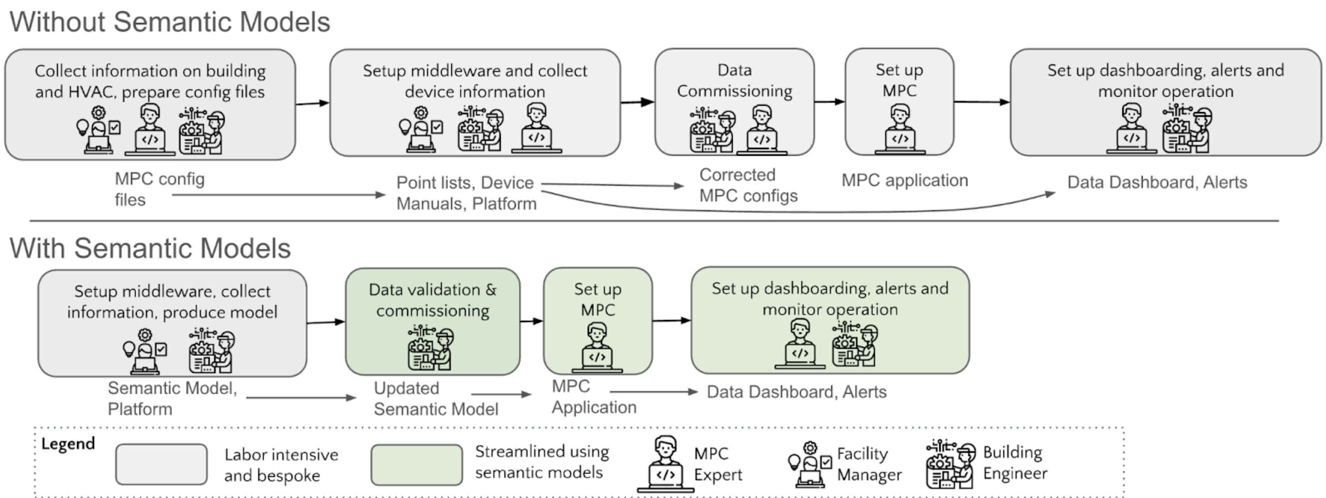

Figure 2: Workflow for implementing the MPC in a building with and without the semantic model-driven approach developed in this paper

HP-FLEX MPC SEMANTICS-DRIVEN DEPLOYMENT PROCESS

In this paper we present a new workflow leveraging semantic models to improve the scalability of this MPC by digitizing the process and enabling a streamlined and efficient designation of responsibilities for the described roles. This workflow uses templates to bootstrap semantic model creation. Templates abstract the details of the semantic ontologies and allow us to guide the user through the process of generating the semantic models without requiring understanding of the underlying semantic technologies. This is enabled using the BuildingMOTIF package that supports the creation of semantic models based on concepts described by Fierro et al. (2022). This significantly reduces the amount of effort and the opportunity for syntactic mistakes in the creation of the semantic models. The software architecture used to enable the semantic model-driven approach is shown in Figure 3.

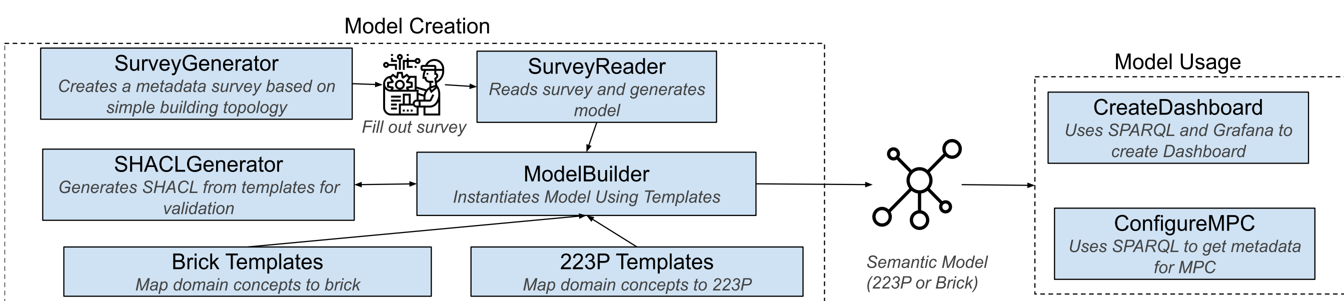

Figure 3: Software architecture for semantic modeling. Brick and 223P Templates abstract details of particular ontologies from higher level software. The ModelBuilder and SHACLGenerator have easy to use functions to generate semantic models using the templates. The survey generator prepares a simple survey requesting nameplate data. The SurveyReader uses this to create the model. CreateDashboard and ConfigureMPC configure the target applications

Using the template-based workflow, metadata collection is guided using a series of csv forms that are simple to fill out. The vast majority of SMCBs use packaged HVAC units feeding one zone each, allowing the structure of the building as pertains to MPC to be described using a simple array indicating the amount of zones and amount of windows and spaces per zone, as shown in Figure 4 in Metadata Collection. This topology of the building may optionally be described in more detail if desired or if a more complex system is modeled, such as a VRF system. This information is used to generate a series of CSV forms that are prefilled with information identifying each hvac unit, zone, space, and window that only require simple nameplate information, like the area of each zone. This process also uses a CSV point list identifying a semantic template for each point. The process of classifying points using classes, tags, or in this case semantic templates is still labor intensive, and under investigation for further automation (Koh et al. 2018; Mishra et al. 2020).

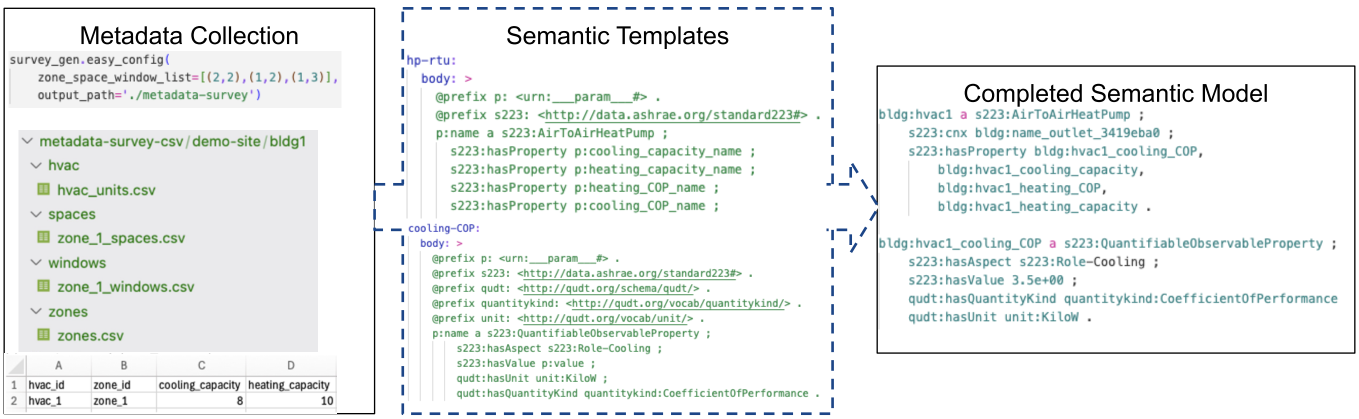

Figure 4: Images of software during metadata collection and model creation. Shows use of 'easy_config' function and simplified building topology to create a survey that will be filled out by the building engineer. Semantic Templates have snippets of ASHRAE 223P used to bootstrap the process of creating a complete semantic model, shown in TTL format[1]

Through this workflow, all of the middleware collection is shortened into a single process accomplished by the building engineer with the help of the facility manager. This can streamline configuration of downstream applications such as data commissioning, MPC, and dashboarding for monitoring because the identity of each point can be queried from the semantic model, and timeseries data can be queried using a reference to the database. In the traditional workflow, configuration files would have to be prepared for each application, which are only useful in conjunction with semantic metadata that is already hard coded within the applications. The process of creating a semantic model brings the hardcoded semantic information that would have been within the MPC out to a semantic model, so that it can be shared between all the applications. In this process, the semantic information that would have been within the MPC application is encoded in the semantic templates. This makes the semantic modeling process easier and defines concepts used by the MPC developer according to a specific ontology, like 223P.

Applications can be configured using semantic queries as shown in Figure 5. Querying can use only the concepts defined in an ontology, such as 223P, as shown in the lower path of Figure 5 in Semantic (SPARQL) Query. However, using our tools, an application developer can also utilize templates as a shortcut to make query authorship simpler and less verbose. As described above, the templates created for this paper define the concepts that the MPC application developer uses for a building in terms of the lower level concepts the ontologies use. Inference can use these templates to label entities in the semantic model with the higher level concepts used by the MPC application developer. Figure 5 shows how inference can be used to label bldg:rtu_1_cc as a hpf:cooling-capacity, which is one of the templates defined for this workflow, based on the ASHRAE 223P representation of bldg:rtu_1_cc. This label is then used for querying. This inference is enabled by SHACL

[1] https://www.w3.org/TR/turtle/

(Knublauch 2017) and the SHACLGenerator software module, which generates SHACL shapes based on the templates.

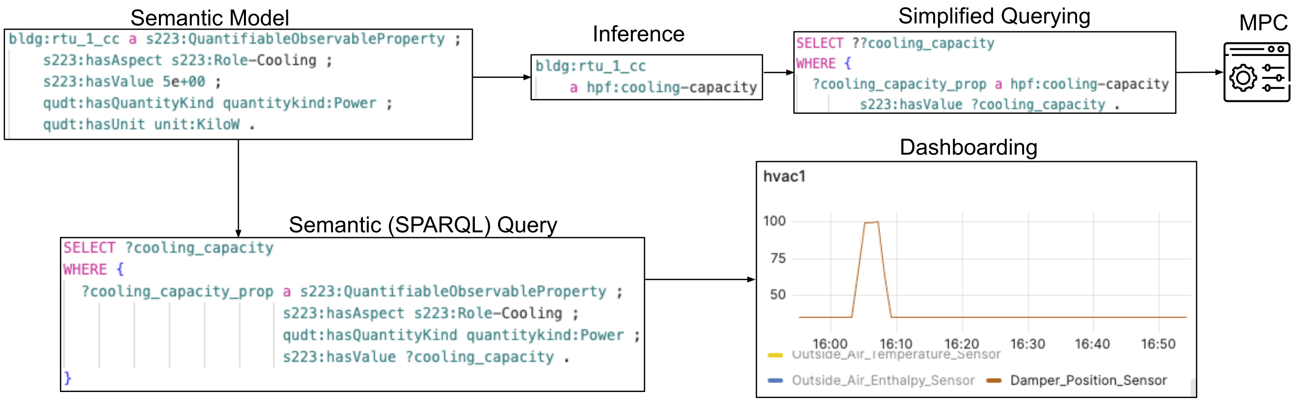

Figure 5: Querying of the semantic model to set up the MPC, supported by inference, and the dashboard using exclusively concepts defined in the ontology, ASHRAE 223P

SHACL can be used to specify the expected structure of semantic data and can be used for inference and validation of semantic models. Inference can add information to semantic models based on what is present in the model. Validation can check that the semantic model fits the expected structure, and if not, it can provide feedback identifying the missing or incorrect information in the model. This is used to check that a model complies with an ontology or that it has the information required to support a specific application. The SHACLGenerator module is used to generate SHACL that can perform inference, labeling models where subgraphs within the model match the templates, and to generate SHACL that can be used for validation, ensuring that information required by the applications is present. How templates may be used for this purpose is further described in Fierro et al. (2022).

So far, we have described the workflow for building the semantic model from scratch. Using validation to ensure the model has the needed information and inference to add information to support querying is not interesting if the model has been made entirely using this workflow, as we can guarantee that the needed semantic information is present and labeled as desired. However, there may be future scenarios in which a semantic model already exists for a building. If a semantic model for the building were to already exist, this validation process could be used to check if the semantic model would support the MPC application. If not, this process would identify the exact information that should be added to the model so that it supports the MPC application, as shown in Figure 6. Inference, as shown in Figure 5 could then make that model easier to query over. Enabling this functionality can eliminate much of the tedious metadata collection required to enable the MPC.

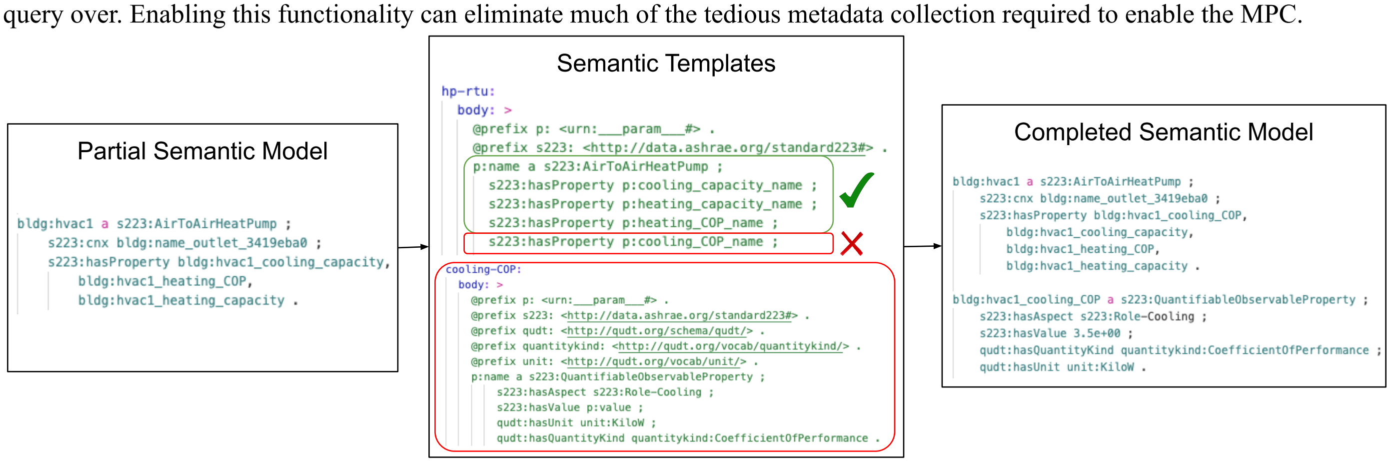

Figure 6: Visualization of the validation of a partial semantic model against the requirements defined within a semantic template. The missing information is identified to help produce a completed semantic model.

IMPLEMENTATION

The semantic-driven HP-FLEX MPC has been effectively implemented in a college building in California, utilizing a workflow based on a semantic model, with ongoing experimental assessments. The building consists of two thermal zones, each conditioned by a HP rooftop unit. The operation schedule for this building runs from 5:00 AM to 8:00 PM, maintaining a temperature range between 20°C (68°F) and 22.2°C (72°F). The HP-FLEX MPC was deployed during the winter season, leveraging dynamic electricity pricing signals as outlined by (Piette et al. 2022). Operations using HP-FLEX MPC and Baseline control were conducted in alternating weeks, with approximately one month of data collected for each case to facilitate performance evaluation (Figure 7).

In the Baseline scenario, a noticeable spike in power demand for heating occurs between 5:00 and 7:00 AM, coinciding with slight overlaps in the peak price brackets from 6:00 to 7:00 AM. Once the desired temperature is reached, the HPs cease operation, leading to a sharp decrease in power usage, followed by a slight rebound around 8:00 AM. In contrast, the HP-FLEX MPC proactively engages heating early in the morning, ensuring that room temperatures remain consistently close to the setpoints. This results in a more uniform load profile during the initial operational period, starting at 5:00 AM. Unlike the Baseline approach, the HP-FLEX MPC minimizes abrupt shut-offs and rebounds between 7:00 and 9:00 AM, leading to a steadier load profile. Consequently, the HP-FLEX MPC shifts approximately 9% of the load from the high-price morning period of 5:00–11:00 AM to earlier morning hours. It is important to note that the peak pricing period does not overlap with occupied hours, limiting the amount of shifted load on this site. Furthermore, in the Baseline scenario, the heat pumps start operating at 5:00 AM, and the room temperature reaches its setpoint around 6:00 AM. In contrast, the HP-FLEX MPC proactively maintains the room temperature close to the setpoint, achieving the desired temperature by 5:00 AM. This can be viewed as an enhancement in occupant comfort.

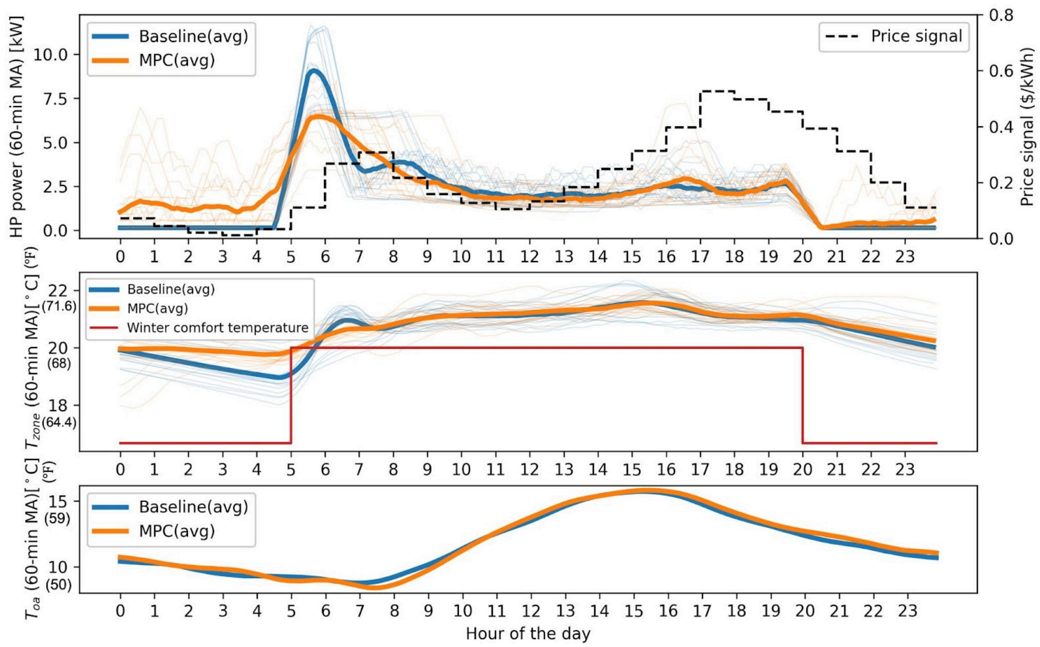

Figure 7: Preliminary performance comparison of daily average HP power profiles (Top), average thermostat temperature profiles (Middle), and average outdoor air temperature (Bottom) between Baseline and HP-FLEX MPC for heating season with dynamic electricity price signal; all data is 60-min moving-averaged which is typical peak demand billing window.

DISCUSSION AND CONCLUSION

In this paper, we have defined a simple workflow and tool for creating a semantic model to support advanced building applications, particularly the HPFLEX-MPC. For the SMCBs studied, this tool was easily used by individuals with little experience with semantic ontologies. The simplicity of this tool was enabled by two things: 1) the lower complexity of SMCBs compared to large buildings, which have many more possible system configurations, increasing the complexity of semantic modeling; and 2) the ability of these ontologies to model systems at different levels of specificity. This allows us to model only the information needed by particular applications. We have tested this approach with 223P and Brick with SHACL, but expect that it should be applicable for other ontologies, such as Haystack with XETO (Xeto 2025). 223P was a particularly good fit for this approach because it describes components at a more granular level, allowing users to compose them in a more flexible manner. Future work should investigate the possibility of translating concepts between different ontologies, such as Brick, ASHRAE 223P, and Haystack, by examining more in depth the best practices used to create these templates.

While we achieved ease-of-use in creating semantic models for SMCBs in the example presented , a significant amount of more effort would be required to develop a solution that is both easily usable and generalizable to multiple buildings, applications, and various vendor technologies. As the applications and system-types grow, the challenge of effectively organizing and composing templates for modeling will become increasingly difficult. We expect control vendors will likely need to develop similar tools and workflows to scale advanced control applications across the large building stocks.

ACKNOWLEDGMENTS

This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technologies Office, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231, by the New York State Energy Research & Development Authority (NYSERDA) through the NextGen HVAC Innovation Challenge program and by California Energy Commission through grant EPC-19-013.

REFERENCES

Andrade Pereira, Flavia de, Lazlo Paul, Armando Casillas, Anand Prakash, Weiping Huang, Marco Pritoni, Conor Shaw, Susana Martín-Toral, Donal Finn, and James O' Donnell. 2024. "Enabling Portable Demand Flexibility Control Applications in Virtual and Real Buildings." *Journal of Building Engineering*, January. https://doi.org/10.1016/j.jobe.2024.108645.

ASHRAE. 2025. "ASHRAE Standard 223 User Documentation." 2025. https://docs.open223.info/intro.html.

Balaji, B., A. Bhattacharya, G. Fierro, J. Gao, and J. Gluck. 2018. "Brick: Metadata Schema for Portable Smart Building Applications." *Applied Energy*.

Blum, D. H., K. Arendt, L. Rivalin, M. A. Piette, M. Wetter, and C. T. Veje. 2019. "Practical Factors of Envelope Model Setup and Their Effects on the Performance of Model Predictive Control for Building Heating, Ventilating, and Air Conditioning Systems." *Applied Energy* 236 (February):410–25.

Blum, D., Z. Wang, C. Weyandt, D. Kim, M. Wetter, T. Hong, and M. A. Piette. 2022. "Field Demonstration and Implementation Analysis of Model Predictive Control in an Office HVAC System." *Applied Energy* 318:119104.

Chamari, Lasitha, Shalika Walker, Lu Wan, Ekaterina Petrova, and Pieter Pauwels. 2024. "Portable Model Predictive Controller System Design for Demand Side Management Using Semantic Web Technologies." In *CIBW782024*.

Delgoshaei, Parastoo, Mohammad Heidarinejad, and Mark A. Austin. 2022. "A Semantic Approach for Building System Operations: Knowledge Representation and Reasoning." *Sustainability* 14 (10): 5810.

"Energy Information Administration (EIA)- about the Commercial Buildings Energy Consumption Survey (CBECS)." n.d. Accessed March 29, 2023. https://www.eia.gov/consumption/commercial/data/2018/bc/html/b1.php.

Fierro, Gabe, Avijit Saha, Tobias Shapinsky, Matthew Steen, and Hannah Eslinger. 2022. "Application-Driven Creation of Building Metadata Models with Semantic Sufficiency." In *Proceedings of the 9th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*. Vol. 1. New York, NY, USA: ACM. https://doi.org/10.1145/3563357.3564083.

Ham, Sang Woo, Donghun Kim, and Lazlo Paul. 2024. "Design and Experimental Performance of Practical MPC for Multi-Zone VRF System for Small and Medium Commercial Buildings." In *International High Performance Buildings Conference*.

Ham, Sang Woo, Lazlo Paul, Armando Casillas, Anand Prakash, Donghun Kim, Richard Brown, Marco Pritoni, and Peter Grant. 2024. "Practical Challenges of Model Predictive Control (MPC) for Grid-Interactive Small and Medium Commercial Buildings." In *\textitProceedings of the 2024 ACEEE Summer Study on Energy Efficiency in Buildings*. Washington, D.C.: American Council for an Energy-Efficient Economy.

Ham, Sang Woo, Lazlo Paul, Donghun Kim, Marco Pritoni, Richard Brown, and Jingjuan(dove) Feng. 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.

Holger Knublauch, Dimitris Kontokostas. 2017. “Shapes Constraint Language (SHACL).” W3C. July 20, 2017.

Kim, D., and J. E. Braun. 2018. “Development, Implementation and Performance of a Model Predictive Controller for Packaged Air Conditioners in Small and Medium-Sized Commercial Building Applications.” *Energy and Buildings* 178:49–60.

Kim, Donghun, and James E. Braun. 2022. “MPC Solution for Optimal Load Shifting for Buildings with ON/OFF Staged Packaged Units: Experimental Demonstration, and Lessons Learned.” *Energy and Buildings* 266 (July):112118.

Koh, Jason, Dezhi Hong, Rajesh Gupta, Kamin Whitehouse, Hongning Wang, and Yuvraj Agarwal. 2018. “Plaster: An Integration, Benchmark, and Development Framework for Metadata Normalization Methods.” In *Proceedings of the 5th Conference on Systems for Built Environments*. New York, NY, USA: ACM. https://doi.org/10.1145/3276774.3276794.

Lin, Guanjing, Hannah Kramer, Valerie Nibler, Eliot Crowe, and Jessica Granderson. 2022. “Building Analytics Tool Deployment at Scale: Benefits, Costs, and Deployment Practices.” *Energies* 15 (13): 4858.

Mishra, Sakshi, Andrew Glaws, Dylan Cutler, Stephen Frank, Muhammad Azam, Farzam Mohammadi, and Jean-Simon Venne. 2020. “Unified Architecture for Data-Driven Metadata Tagging of Building Automation Systems.” *Automation in Construction* 120 (103411): 103411.

Neukomm, M., Valerie Nubbe, and R. Fares. 2019. “Grid-Interactive Efficient Buildings Technical Report Series: Overview of Research Challenges and Gaps.” Office of Energy Efficiency & Renewable Energy, U.S. Department of Energy. https://doi.org/10.2172/1577966.

Paul, Lazlo, Flavia De Andrade Pereira, Anand Krishnan Prakash, Sang Woo Ham, Jingjuan Dove Feng, Rich Brown, and Marco Pritoni. 2025. “Open Building Operating System: A Grid-Responsive Semantics-Driven Control Platform for Buildings.” *Science and Technology for the Built Environment*, January, 1–18.

Piette, Mary Ann, Jingjing Liu, Bruce Nordman, Sarah Smith, Brown Richard, and Marco Pritoni. 2022. “Accelerating Decarbonization with the California Load Flexibility Research and Deployment Hub.” In . LBNL. https://doi.org/10.20357/B79S3T.

Prakash, Anand Krishnan, Flavia De Andrade Pereira, Mario Bergés, Marco Pritoni, and Burcu Akinci. 2024. “Ontologies at Work: Analyzing Information Requirements for Model Predictive Control in Buildings.” In *Proceedings of the 11th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, 214–18. New York, NY, USA: ACM.

Pritoni, Marco, Drew Paine, Gabriel Fierro, Cory Mosiman, Michael Poplawski, Avijit Saha, Joel Bender, and Jessica Granderson. 2021. “Metadata Schemas and Ontologies for Building Energy Applications: A Critical Review and Use Case Analysis.” *Energies* 14 (7): 2024.

Pritoni, Marco, Michael Wetter, Lazlo Paul, Anand Prakash, Weiping Huang, Steven Bushby, Parastoo Delgoshaei, et al. 2024. “Digital and Interoperable: The Future of Building Automation Is on the Horizon. What’s in It for Me?,” August.

“Ref Schema.” 2025. Github. 2025. https://github.com/gtfierro/ref-schema/tree/main.

“Xeto.” 2025. Github. 2025. https://github.com/Project-Haystack/xeto

Touzani, Samir, Anand Krishnan Prakash, Zhe Wang, Shreya Agarwal, Marco Pritoni, Mariam Kiran, Richard Brown, and Jessica Granderson. 2021. “Controlling Distributed Energy Resources via Deep Reinforcement Learning for Load Flexibility and Energy Efficiency.” *Applied Energy* 304 (December):117733.

Wan, Lu, Ferdinand Rossa, Torsten Welfonder, Ekaterina Petrova, and Pieter Pauwels. 2025. “Enabling Scalable Model Predictive Control Design for Building HVAC Systems Using Semantic Data Modelling.” *Automation in Construction* 170 (105929): 105929.