Building Technologies & Urban Systems Division
Energy Technologies Area
Lawrence Berkeley National Laboratory

Demystifying Thermal Energy Storage Integrated Heat Pump Systems: Development of Generalized Sizing and Control Algorithms for Demand Flexibility

Armando Casillas, Thomas Defauw, Sang woo Ham, Dre Helmns, Lazlo Paul, Anand Prakash, Weiping Huang, Peter Grant, Marco Pritoni

Lawrence Berkeley National Laboratory

Energy Technologies Area
August 2024

https://doi.org/10.20357/B7PW38

This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy,
Building Technologies Office, of the US Department of Energy
under Contract No. DE-AC02-05CH11231.

Demystifying Thermal Energy Storage Integrated Heat Pump Systems: Development of Generalized Sizing and Control Algorithms for Demand Flexibility

Armando Casillas, Thomas Defauw, Sang woo Ham, Dre Helmns, Lazlo Paul, Anand Prakash, Weiping Huang, Peter Grant, Marco Pritoni,
Lawrence Berkeley National Laboratory

ABSTRACT

As electrification and decarbonization goals become more commonplace across the country, the need for integrating thermal energy storage (TES) with HVAC to provide flexibility and load shifting is growing. Although there has been recent work related to the modeling and design of TES-integrated heat pump (HP) systems, investigation of generalized sizing and control methods for these systems remain limited. This paper details the development of generalized controls and sizing strategies applicable across different TES-integrated designs, two of which are discussed in this study. We demonstrate how model-based design enables an informed sizing and controls design process using the control-oriented Modelica language to generate high-fidelity models that accurately represent real-world behavior.

We detail our development and testing of both heuristic and model predictive control (MPC) algorithms to determine the optimal charge and discharge schedule with dynamic varying utility prices. Experimental results show MPC provides operating costs reduction of nearly 20% for a minimum TES sizing scenario. In addition, we provide a generalized and intuitive control algorithm with near-optimal performance to control HP + TES systems and test its performance in simulation. This generalized control algorithm is also used to drive the results of our cost analysis, providing insights for engineers designing new TES-integrated HVAC systems. The cost analysis demonstrates the tradeoff between higher initial hardware costs from larger equipment and the resulting operational cost benefits, and enables a cost-effective sizing method which is applicable to any system. The paper concludes with design recommendations for new integrated HP-TES control systems in buildings.

Introduction

Electrification of heating loads presents a significant decarbonization opportunity in all types of buildings. Although HP technology has become more commonplace as a means for electrifying these loads, there remain a number of barriers to scaled adoption and deployment. Among the barriers include the high cost of HP systems, compounded by potential additional costs related to electric panel upgrades and higher electricity costs (Rosenow et al., 2022). Typical air-to-water HP systems do not include any storage, preventing them from shifting electricity consumption to match times of renewable power production. Thermal energy storage systems bring the promise of higher flexibility for buildings while also serving as a remedy of the chronic oversizing seen in traditional HVAC design practices. Unlike traditional electrochemical battery systems, phase change material thermal energy storage systems are not subject to chemical degradation over time (they do not rely on a chemical reaction to store energy) and can integrate directly with an HVAC system (Gu et al., 2023). This allows the HP system to be designed to cover the majority of the conditioning hours throughout the year and can rely on the TES to provide additional capacity during peak demand hours when weather conditions are most challenging.

More importantly, these advantages allow for HPs to be downsized while maintaining the ability to adequately serve annual loads. Helmns et al's (2022) simulation study for instance, found TES to enable heat pump capacity reductions up to 60%. Recent market studies have revealed that first cost is a main driver for building owners looking to electrify their systems (Garcia et al., 2024), which currently include the procurement of the HP system (space heating and/or domestic water heating), and a possible costly panel upgrade, which may cost up to $5000 per panel (Walker, Casqueri-Moderego, and Less, 2023). In California, for example, about 30% of single-family homes have panel capacities less than 100 A thus needing panel upgrades to electrify (Pena et al. 2022), which serves as a conservative estimate for the widespread cost expected for electrifying residential buildings.

Despite the promise TES provides, it's ambiguous how to best size these systems in concert with HP systems. These gaps have been present for decades, with recommendations for standardized sizing tools dating back to the late 1980's (Dumortier et al. 1989) due to a tendency to oversize these systems, diminishing economic benefits of TES (DeForest et al. 2014). More recent studies have explored sizing methods such as Hao et al (2021) and Hirschey et al (2023) although these studies have not factored in the implications that sizing decisions have on control complexity and flexibility potential. Although multiple deployments of these systems have tried both rule-based and optimized control approaches (Behzadi 2022), controls for these systems also remain unstandardized. Standardizing these strategies allows for easier industry adoption and deployment of TES-integrated systems. Although control implementations have been demonstrated in the literature, some of the more advanced approaches can be difficult to understand or, and in the case of model predictive control, require time intensive model development, configuration and tuning to yield positive results (Cigler et al., 2013; Drgoňa et al. 2021).

The conventional process of designing and prototyping building technologies involves a series of sequential phases: system design, prototyping, limited laboratory testing, and extensive field tests (Naumann and Jenkins 1982) (Thomke 2003). Controls are typically developed independently, and the actual testing occurs in the field. This approach frequently results in suboptimal performance and the belated identification of control issues, necessitating costly measures to address these problems. As systems become more complex and involve integration of multiple technologies with high first costs, the risk of issues has more costly implications (Blum et al. 2021). A key method to de-risk these integrations is model-based design (MBD), which consists of an iterative modeling-design-simulation process, where systems and controls are developed together and validated in simulation before conducting expensive laboratory and field tests (Isermann 2014) (Wetter and Sulzer 2024). In this study we demonstrate how model-based design enables us to think about how we standardize sizing and controls of these systems and help us analyze how sizing and controls change based on how the TES is integrated.

The paper is organized as follows: (1) Methods section, describing sizing strategies, our assessment of 3 control algorithms for 2 energy system designs, and our cost analysis, (2) Results section describe the products of our methodology, (3 & 4) Discussion and conclusions.

Methods

In this study we describe two energy system configurations for TES-integrated systems, analyze a number of controls and sizing methods and implement these algorithms to calculate first and operational costs of both system designs. Two TES-integrated systems are sized and controlled, the first is a series-integrated TES system and the second is a parallel-integrated TES system.

Energy System Configurations

The series-integrated TES system can be seen in Figure 1, in which a phase change material (PCM) serves as TES between an air-to-water HP and a load. This series configuration uses the TES as an intermediary ensuring a stable supply of thermal energy while simultaneously decoupling space loads from HP operation, and ensuring the HP receives water at a consistent, controlled temperature. The parallel integrated system (shown at the bottom of Figure 1) in contrast can serve the load directly with the heat pump or the TES. The heat pump is also responsible for charging the TES at appropriate times.

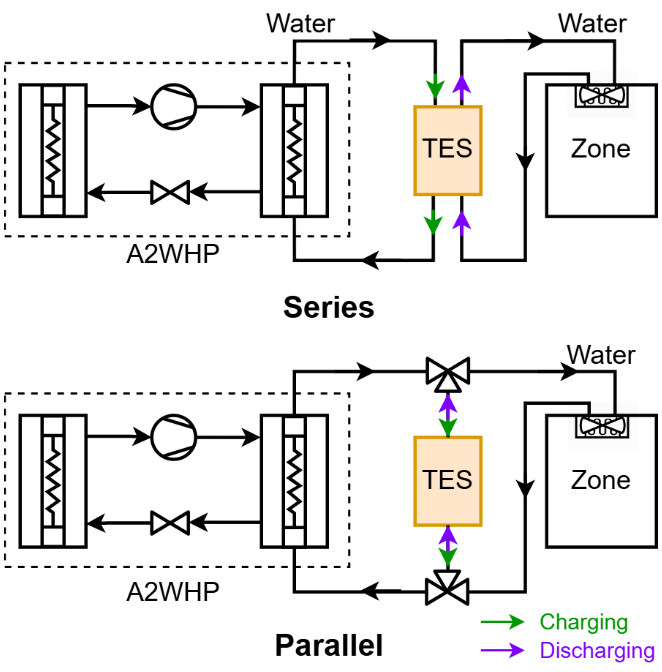

Figure 1. Conceptual diagrams of heat pump and PCM thermal energy storage system layout for both the series (top) and parallel (bottom) integrated systems

The parallel design was modeled in Modelica and tied into an experimental facility via a hardware-in-the-loop test design. This allows us to (1) test the performance of TES+HP sizing combinations using a realistic building thermal load profile, (2) determine the benefits of different control algorithms and their effect on flexibility potential. The series design was also modeled in Modelica and used to evaluate the performance of the proposed generalized control algorithm in simulation.

Operating Mode Definitions.

The series-integrated system (Figure 1 – Top) is designed similarly to a HP to buffer-tank design, where the HP's primary goal is to provide hot water to a middle storage device, while the storage device discharges to provide space conditioning to the zone. The following modes exist for this system are listed below and are summarized in Table 1:

Table 1. Summary of HP+TES control modes for both system designs

	Mode 1: Heat Pump Serves Zone Load	Mode 2: Heat Pump Serves PCM	Mode 3: PCM Serves Zone Load	Mode 4: Heat Pump Serves PCM while PCM Serves Zone Load
Series		x	x	x
Parallel	x	x	x	x

The parallel-integrated system is designed in such a manner in which the HP or the TES can provide space conditioning (Figure 1 – bottom). The TES is charged by the HP, when enabled, using a supply water temperature reset. There is another feasible mode in which the HP can provide both space conditioning and TES charging simultaneously, however in the case for PCM, the HP supply temperature would need to be exceptionally hot or cold and detrimental to system COP. An example of how the COP decreases at higher supply water temperatures can be found in publicly available manufacturer data (LG, 2020). The modes for both energy system configurations are explained below:

1. HP to Zone Mode
    a. HP provides space conditioning, no interaction with TES
2. TES Discharging Mode:
    a. TES provides space conditioning, when enabled, provided there is adequate state of charge (SOC).
    b. Capacity delivered to the zone is modulated by circulation pump controlled to a specified temperature difference
3. TES Charging Mode:
    a. Charging of TES is handled by the HP when the SOC is low and requires replenishment.
4. Simultaneous Charging and Discharging Mode
    a. Simultaneous charging and discharging is also possible in cases where the space calls for conditioning during times when the controller calls for HP charging

Control System Strategies

In addition to the two energy system configurations, three control system strategies were explored to determine the variation of performance relative to each system design. The three price-responsive algorithms, described below, determine the mode switching operation. The heuristic and MPC algorithms were conducted experimentally as part of a previous DOE-funded project, while the assessment of the generalized algorithm performance was conducted purely through simulation. We tested both control strategies in an experiment implementing the parallel design configuration. We also simulated the generalized control strategy and compared it to the heuristic baseline for the series design configuration. All control strategies are summarized in Table 2.

Table 2. Details of control strategies for testing a TES-integrated system

Strategy	Heuristic (Baseline)	MPC	Generalized
Description	Daily hour schedule based heuristic control to use PCM TES and avoid HP operation during peak cost time.	Model predictive control to minimize energy cost with respect to electricity cost signal.	Closed-loop algorithm that strives to turn on the heat pump from lowest to highest cost periods
Control Algorithm	Charges TES in low load and low-price windows, TES is discharged during high-price periods, HP is deactivated during high-price period, HP is operated to serve the load otherwise	Predicts thermal load and balances electricity price to minimize operational cost while taking full advantage of TES	Constructs an operation schedule for the HP that prioritizes the use of less expensive hours to charge the storage. Implements the first time step and repeats the process.
Price Signal	CalFlexHub Spring & Summer Highly Dynamic Price Signal (LBNL, 2024)	CalFlexHub Summer Highly Dynamic Price Signal (LBNL, 2024)	CalFlexHub Spring Highly Dynamic Price Signal (LBNL, 2024)
Evaluation Method	Simulation & Lab Experiment	Lab Experiment	Simulation

Heuristic and MPC Algorithms.

The heuristic (baseline) control strategy is schedule-based. It discharges the TES when the electricity cost is high and charges the TES overnight from 23:00 to 7:00 during the low-cost time. The cost signal used in this study has a high price during 16:00 to 22:00 and has the same profile for each day (Figure 6). Assuming the thermal load profile of an office building, the cooling load is highest in the afternoon. Therefore, the TES is set to be used for the late afternoon which has a high electricity price. Then, it is charged during the nighttime for the next day's use.

To better understand the benefits of optimized controls for a downsized TES-integrated system, a model predictive control (MPC) strategy was also tested. The MPC is designed to provide the optimal sequence of charge, discharge, and HP-only modes during the day based on the electricity cost signal and the predicted cooling load. After setting up the 3-zone gray-box (combination of data-driven and physics-based models) model, a black-box (data-driven model) optimizer is used to find the optimal sequence of mode profiles for a day.

The evaluation of the heuristic and MPC price-responsive algorithms was conducted in for the parallel system involved gathering experimental data from LBNL’s FLEXLAB facility (McNeil, Kohler, and Lee 2014). FLEXLAB is a highly instrumented and customizable testing environment, which allows researchers to implement a wide range of energy efficient technologies against an identical baseline cell to efficiently evaluate their added benefit to building performance. The experiments were set up in two identical building cells (X1A and X1B) in FLEXLAB with two identical HVAC systems. The HIL strategy in this FLEXLAB experiment uses physical HVAC distribution and delivery systems supplied by a physical water chiller, but the chiller and a downstream heating element are controlled by Modelica-based models running on a local server. The experimental tests used a 3 kW HP and a 14.5 kWH 8°C PCM TES in the virtual plant

Generalized Control Algorithm.

In the pursuit of standardization and ease of implementation, we also developed a simple generalized control algorithm for heat pump and thermal energy storage systems that leverages the best aspects of both the heuristic and MPC strategies. For a given set of electricity price and load forecasts, it finds a near-optimal operation schedule for the heat pump that guarantees thermal comfort while avoiding the use of electricity during expensive peak hours (it does not account for potential demand charges, as dynamic pricing is expected to replace demand charges in California [Matisoff 2020]). As with MPC, the algorithm is implemented in a closed loop, where the first step of the schedule is implemented at each time step, and the state of the system serves as a feedback.

The operation schedule prescribes the thermal power of the heat pump at each time step over the prediction horizon (time is typically discretized in hours and the horizon is generally set to 24 hours). The schedule is determined through an iterative, forward-moving process. Before the process begins, the heat pump’s power is initially set to 0 at each time step. The first hour in which the load cannot be met (i.e., insufficient SOC at that time) is identified. Then, the heat pump’s thermal power is increased during the lowest-cost hour(s) before that hour, until the load is satisfied at that hour. The power is increased only if it does not lead to exceeding the storage capacity. Additionally, if increasing the power results in reaching lower cost operating hours in the future, then the power is increased only enough to allow the storage to supply the load until that point. This process is repeated until the load is met at every hour in the prediction horizon.

The generalized control algorithm can be applied to both the series and parallel configurations, and is meant to be applicable to a wide range of configurations. For this reason, this control algorithm is used to determine the operating costs used in the generalized sizing method described in the next section. It assumes that the heating power of the heat pump can be adequately controlled, either directly or indirectly. In this study, we control it indirectly through the setpoint for the heat pump's supply water temperature, assuming that the temperature of the water returning from the PCM is roughly constant for a given state of charge. Direct control would involve modulating compressor speed, a control point which is only accessible to heat pump manufacturers.

The developed algorithm was implemented in Python and used to control a high-fidelity Modelica model of the system in series configuration. Its performance was compared against the heuristic algorithm, which turns off the heat pump during peak hours and evenly distributes the shifted load to the hours before the peak.

Functional Mockup Units (FMUs) for Testing:

The primary method of developing all of the aforementioned controls is co-simulation via Functional Mockup Units (FMUs). Akin to compressing a large collection of files into a .zip file, a functional mockup unit allows for data exchange between a detailed model and a control algorithm without having to compile and run a simulation in the native model compiler program, which often requires a software license and prior knowledge on how to operate the software. Interaction with the FMU is done through publicly available tools that can be leveraged for co-simulation, including PyFMI. For both the series and parallel systems presented earlier, we co-simulated our control strategies with the system model through a hardware-in-the-loop (HIL) approach. The model states such as SOC, water temperatures and system electric demand are fed back to the control algorithm which then decides any change in mode command and setpoints based on the feedback. Figure 2 illustrates this data exchange.

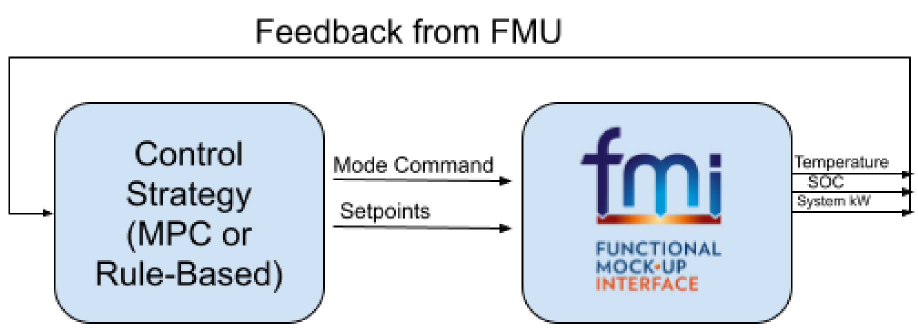

Figure 2. Data exchange between control strategy and the modeled HP+TES plant

Equipment Sizing & Cost Analysis

Recommended Equipment Sizing.

Selecting the right storage and heat pump capacity is critical in order to bring down the costs of such systems. While increasing the capacity of the equipment generally yields lower operating costs, it also comes with higher investments for capital cost. (Zheng, Ma and Wang 2015) In the series integrated system, the HP is insulated from the hourly changes in zone load and only serves to charge the TES. However, the TES needs to be able to serve the load for the entirety of the high price periods or demand response events to provide demand flexibility. The HP must also be large enough to charge the TES in an adequate time period in order to avoid the

TES falling below a minimum SOC. The parallel integrated system can also only serve the load either through the heat pump or TES, and therefore has an equivalent sizing strategy.

The sizing strategy for these systems is based on a method established by Hao et al. (2021) and was conducted using time series data from the peak cooling day within the observation period (August to October 2021). Using the peak day ensures that the system will be capable of meeting the highest cooling demands we can reasonably expect for the site. In order to create the sizing map, we assumed an 8-hour charging window while the discharge period is assumed to cover 4 hours during high price periods. The resulting sizing map for this case study is shown in Figure 3. This case study represents system sizing at the intersection of the blue line and orange line, which denotes the minimum heat pump capacity needed to serve the peak thermal load at the plant sizing determined. This is equivalent to point A as seen in Figure 3. There are several assumptions made in this case study. First is that the charge and discharge efficiencies of the PCM are treated as 100% and that the PCM is not charge rate limited. This means that it can charge using whatever load the HP can provide, and discharge at the rate required by the site. Specified off-peak times were used as the charging window and on-peak times as the discharging window. The length of time the PCM is allowed to charge or required to discharge affects the maximum recommended capacity, shown in red and orange on Figure 3. These assumptions were useful in this case study, but there is room for future improvement. The charge and discharge windows could be optimized, taking the control strategies into account, and the efficiency and power constraints of the PCM could be considered.

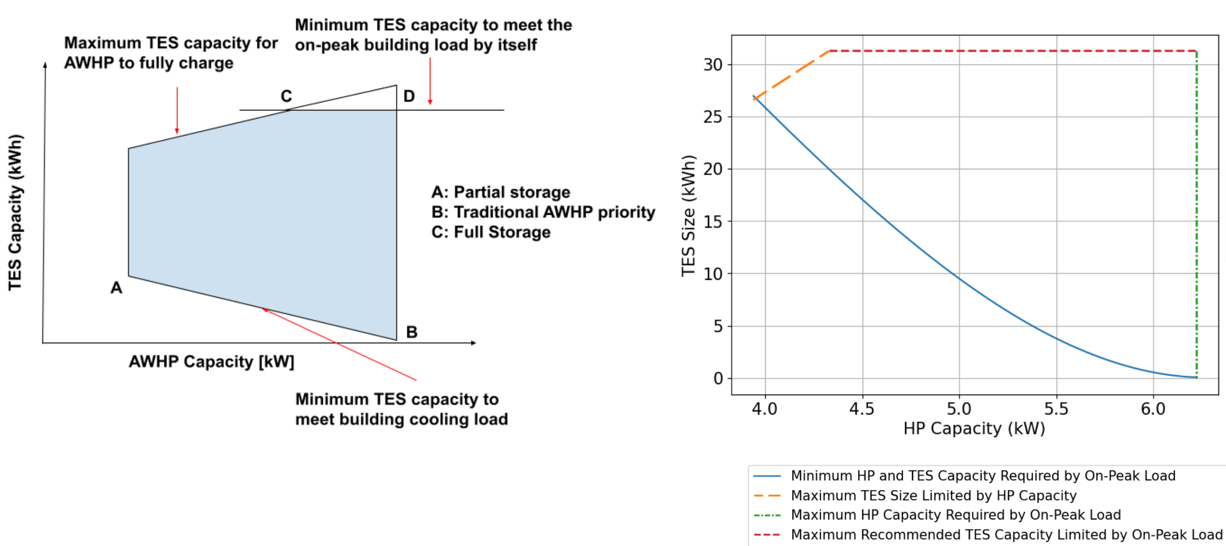

Figure 3. Sizing map that identifies the relationship between TES and AWHP capacity for (A) partial storage, (B) traditional AWHP priority, and (C) full storage scenarios, reproduced from (Hao et al. 2022) [left]. Sizing diagram for generated using field data [right]

Cost Analysis for Sizing.

By estimating both the annualized capital expenditures (CAPEX) and the operating expenditures (OPEX) of HP + TES systems, we define a generalized cost-based sizing approach which simply consists in selecting the equipment sizes (HP thermal power and TES capacity) that yield the lowest overall costs over the lifetime of the equipment.

In this study, we assume that the cost of installing the equipment is not significantly affected by equipment size, and that no panel upgrades are required as a result of installing the heat pump. Therefore, the capital expenditures are the costs from purchasing the equipment only. The capital expenditures are expressed as $/year by dividing these costs by the lifetime of the equipment (typically assumed to be 20 years). Prices for different HP and TES sizes are obtained by scaling from reference prices using a polynomial fit, as shown in Figure 4. The operating costs are estimated from a yearly simulation (with past weather, load, and electricity price data) in which the system is controlled using the generalized control algorithm described in the previous section. When available in simulation, we recommend using the control algorithm that will be used in the physical system to improve the cost estimation.

Finally, the total expenditures (TOTEX) are obtained in $/year by summing the capital expenditures and the operating costs. With this metric, it is possible to compare the costs associated with different combinations of equipment sizes and select the most cost-effective solution, via a parametric sweep of simulations. It is important to note that the results obtained with this method are case specific, as they strongly depend on the available electricity prices and typical heating loads expected in the building.

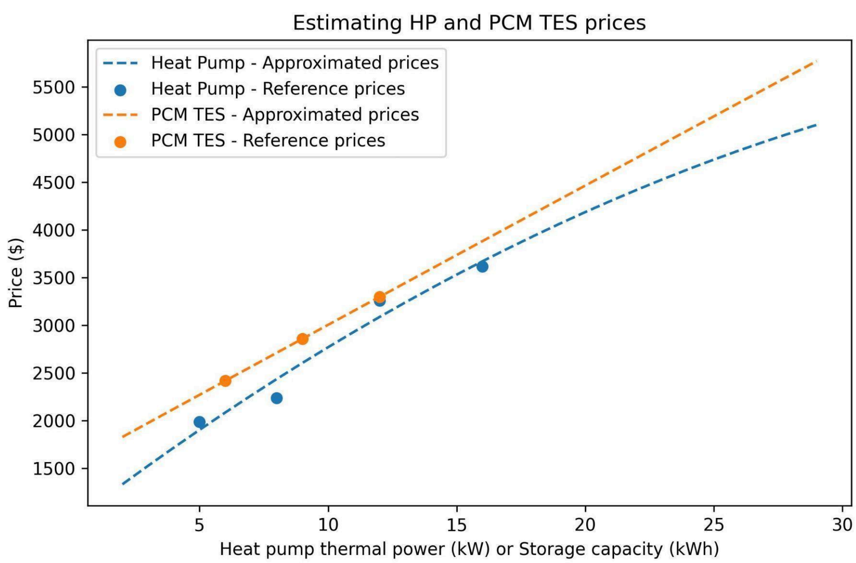

Figure 4. Estimating CAPEX costs for different sizes of HP and PCM TES, based on reference prices

Note: Reference prices were obtained from online sources (Midsummer, 2024) for the Samsung Gen 6 R32 Monobloc Heat Pump and from SunaAmp thermal batteries.

Results

Control Strategies Testing

Parallel system performance with heuristic vs MPC control algorithms.
Figure 5 shows the power profile for the heuristic control and MPC strategies. The MPC algorithm demonstrated the capability to predict thermal loads and strategically discharge the

TES, thereby shifting the building load from the morning high-price period to the lower-price period at night. The lower TES SOCat the end of the day meant a more substantial overnight thermal demand as the HP charged the TES in preparation for the next day. Despite the high evening peak, the overall demand is lower for the MPC. The MPC algorithm demonstrated its ability to anticipate thermal loads, unlike the heuristic controller, which initiated TES discharge well before the high-priced period, per the set schedules that could not capture the highly dynamic price signal.

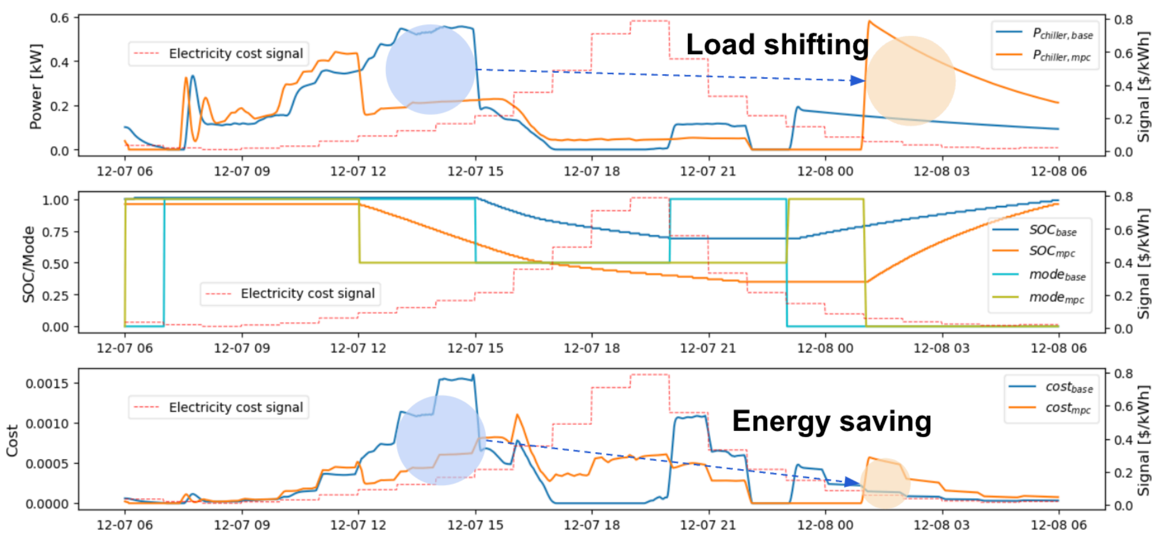

Figure 5. Power profiles, averaged across testing days, for the baseline scenario and the MPC scenario for FLEXLAB HIL test.

Successful MPC prediction allowed for the effective shift of thermal loads to periods with lower prices, resulting in a 10% reduction in HVAC energy use and a 24% reduction in total electricity costs compared to baseline operations with heuristic control. Results are summarized in Table 2.

Table 2. Summary of results for FLEXLAB

Test	Heuristic (Baseline) Total	MPC Building Total (% Reduction)	Heuristic (Baseline) HVAC Total	MPC HVAC Total (% Reduction)
Average daily energy consumption (kWh)	12.4 kWh	11.9 kWh (-4%)	3.23 kWh	2.91 kWh (-10%)
Average daily cost ($)	$2.60	$2.00 (-24%)	$0.39	$0.30 (-23%)

Series system performance with heuristic vs generalized control algorithm.

FMUs were designed to assess the performance of the generalized control algorithm on the parallel and series system configurations. As a reminder, the generalized control algorithm outputs a schedule of thermal power at which the heat pump should operate for every hour in the prediction horizon.

In the series system, the heat pump can only serve the PCM, and the temperature returning from the PCM is assumed to be roughly constant. Therefore, the thermal power of the heat pump can be controlled using the setpoint for the supply water temperature (the heat pump's mass flow rate is assumed to be a function of the supply water temperature, as is typical of commercially available heat pumps).

In the parallel system, the heat pump can either serve the PCM or the Zone, or both simultaneously. In this implementation, the hour is split into two sections: the heat pump first serves the zone directly until the zone's heating load is satisfied, and then switches to charging the PCM for the remainder of the hour. In both modes, the heat pump operates at the same thermal power, modulated by the supply water temperature setpoint. When the HP is serving the zone, the setpoint is dynamically adjusted as a function of the measured return water temperature. When the HP is serving the PCM, we use the same assumption as with the series system in which we consider the return water temperature to be a function of the state of charge.

Figure 6 shows an example of a 24-hour simulation with the FMU of the series system controlled with the generalized control algorithm (solid blue line) and the baseline algorithm (dashed blue line). The heat pump power was indirectly controlled by dynamically adjusting the supply water temperature setpoint and assuming that the water returning from the PCM can be calculated from its state of charge. When the blue line is above the red line, the heat pump is supplying more heat than required by the zone, and this excess energy is charging the storage. The performance is satisfactory, and peak hours (indicated by the higher electricity prices) are successfully avoided by both algorithms. Given the small storage size used in this simulation (10.3 kWh) and large heating loads, both algorithms are only capable of discharging for 2 consecutive hours at a time. The main difference between the two stands on how the charging is distributed across the hours. The generic algorithm makes the most out of the least expensive hours of the day (12 to 14), whereas the baseline charges equally in all hours before a price peak, regardless of the electricity price.

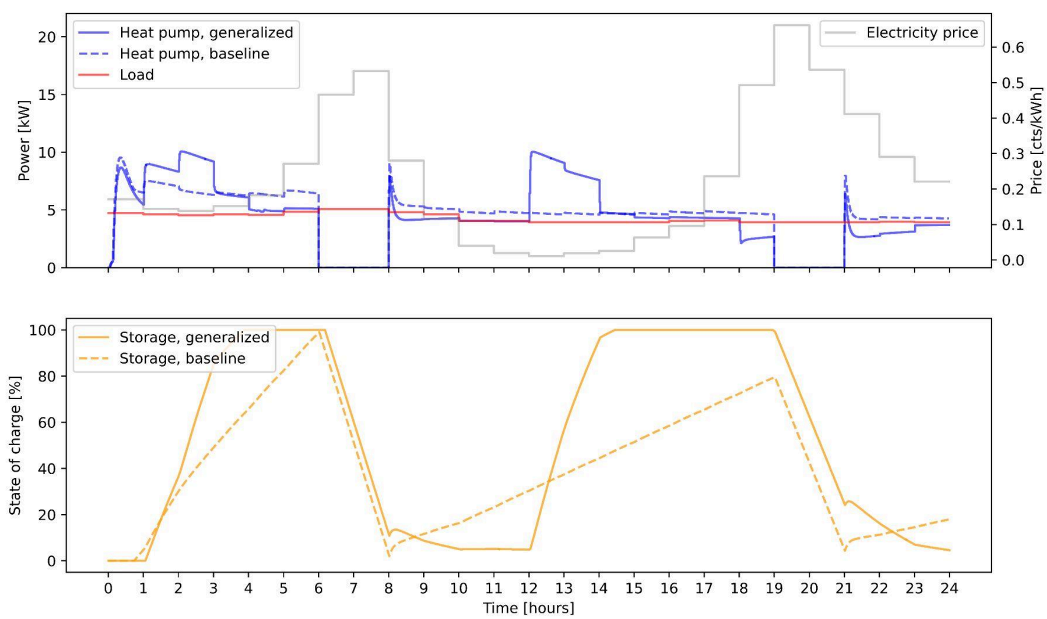

Figure 6. 24-hour simulation of the series system on the corresponding FMU, where the thermal power of the HP is indirectly controlled through the supply water temperature setpoint

Similarly to the MPC developed for the parallel system configuration, the performance of the generalized algorithm was compared to a baseline algorithm. The baseline would turn off the heat pump during pre-designated peak hours (2 hours in the morning peak, 2 in the evening peak) if the storage could provide the load during those hours, and evenly distribute the shifted load to the hours before the peaks. Table 3 summarizes the results obtained with both algorithms when simulating the series system configuration used to heat a house from January to April (included), using past weather and load data. For these simulations, the system was sized according to the strategy described in the 'Recommended Equipment Sizing' subsection, which recommended a 10.35 kW heat pump with 10.3 kWh thermal energy storage.

Table 3. Summary of results for simulating the series system using past load and weather data

Test	Heuristic HVAC Total (Baseline)	Generalized HVAC Total (% Reduction)
Average daily energy consumption (kWh)	47.3 kWh	47.3 kWh
Average daily cost ($)	$9.02	$7.61 (-16%)

These results show that, in this specific scenario, the generalized algorithm was able to save about 34% on electricity costs through load shifting, which is roughly 16% more than what the baseline algorithm was capable of achieving. Both these algorithms use the storage to displace heating loads over time, which explains why the amount of energy remains the same.

Equipment Sizing Simulations

While the results above demonstrate the benefits of different control strategies for operating costs, we would also like to explore the effects of varying HP + TES sizes on total costs. Here, we apply the Cost Analysis for Sizing to determine the most cost-effective combination of heat pump thermal power and thermal storage capacity. The electricity prices used in the simulation are the highly dynamic CalFlexHub (LBNL, 2024) seasonal electricity prices. The operating costs are estimated using the generalized control algorithm, which is independent of the energy system configuration (parallel or series).

Figure 7 provides sample results for a single-family residence in a cold climate, where the total annual costs are estimated for a range of equipment sizes through a parametric sweep. From the plots, we can see that, at first, the additional expenses for purchasing larger equipment are compensated by the reduced operating costs resulting from the ability to shift heating loads to hours in which electricity prices are low. However, above a certain size, the total costs increase as these additional expenses are no longer outweighed by the operational cost benefits.

The red dot on Figure 7 shows the point at which the total costs are minimized, which is a compromise between large equipment for low operating costs (OPEX) and small equipment for low capital costs (CAPEX). To reiterate, the optimal sizing of the equipment will depend on the control algorithm, weather, load and electricity prices associated with the area and building in which the system is installed. In this case, the cold climate, high heating demand, and very dynamic prices used in the simulation recommend large equipment (22 kW HP and 48 kWh TES). In many use cases, this will lead to additional costs such as panel upgrades which could change the recommended equipment size. Given the currently available thermal battery sizes, some options might also not be feasible as of today.

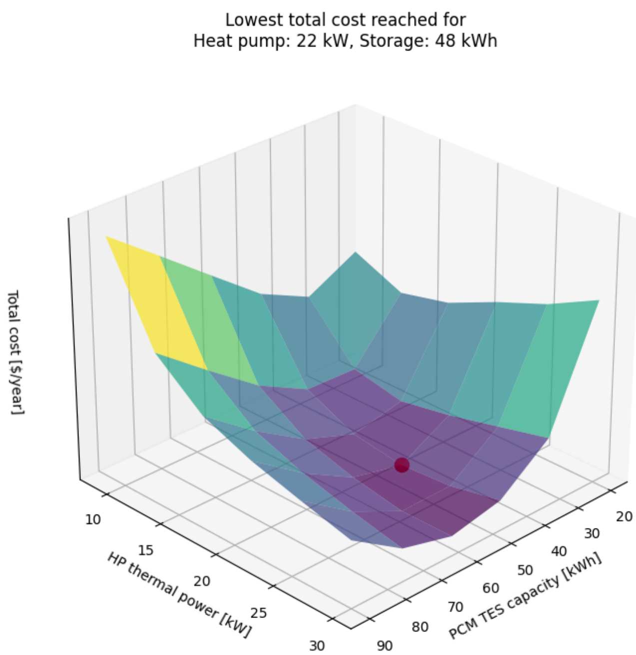

Figure 7. Plot showing the total yearly price (TOTEX = CAPEX + OPEX) of a HP and PCM thermal energy storage system for different equipment sizes. The red dot represents the optimal combination of equipment capacities for this particular example.

Note: The operating cost analysis used to generate Figure 7 considered a yearly weather and heating load data from a 2,200 square foot house in Western Massachusetts with a maximum hourly cumulative heating requirement of 8.4 kWh. With similar load data for any other building, an alternate TOTEX surface map could be generated to identify optimal sizes for the HP + TES.

Discussion

This paper presented three different control approaches for integrated HP + TES systems. The presented MPC for the parallel system reduced building costs for the designed case by 24% but is customized to a specific situation and requires tuning prior to deployment. The simple price-responsive algorithm is intuitive and performs close to the equivalent optimization problem at 16% reduction, but could be challenging to set up depending on the system configuration and available control inputs (controlling the HP's thermal power is generally not directly possible). The two approaches have different, viable paths to market. MPC could be viable in products that are designed/tuned once and sold repeatedly, akin to the smart controls in a Nest thermostat, e.g. a prefabricated, plug-and-play HP + TES system with pre-installed controls. The simple price-responsive algorithm is designed to avoid the tuning required by MPC, which makes it more viable for installations designed in the field or for startup companies that want a control algorithm to leverage with low development budgets.

While the technical potential of integrated HP + TES systems is tremendous, market adoption has been slow due in large part to the challenges of 1) each site requiring customized design without supporting tools, and 2) limited applicable control algorithms. The sizing tools presented here can create a foundation for design tools that support industrial practitioners identifying a high-performing, minimal cost system, while leveraging model-based design strategies enables rapid integration of data describing the specific sites where the system will be installed. This enables designs of systems which are right-sized for specific buildings, reducing both the CAPEX and OPEX of the installed system.

The simple price-responsive algorithm presented here was designed for AWHP + TES systems, but fundamentally operates by identifying the lowest cost times to operate equipment. It could be expanded to other situations where a) building needs, b) operating costs, and c) equipment capabilities are known. Ongoing conversations include collaboration with industry to determine the optimal times to charge/discharge PCM, leveraging the algorithm to identify the optimal times to send load shifting control signals to end-uses. The sizing methods presented here should work just as well for other TES-integrated systems such as in-duct PCM and would likely need only minimal modification to also include other end uses such as heat pump water heaters.

Conclusion

In this study we explored more effective ways of sizing and controlling heat pump systems when integrated with TES. Our cost analysis calculates the total cost of all sizing combinations to determine optimal sizing based on minimized cost. These methods will be needed to scale the deployment of these systems, further equipping manufacturers, engineers, and contractors with tools to consider the interactions between heat pump and TES sizing, controls and total costs. Ultimately, the future to keep in mind is one where adopters of this technology can easily size, purchase and reap the benefits of these innovative system designs.
In future work, a sizing tool that takes into account regional differences in utility rate structures and thermal load could further provide value to the aforementioned stakeholders.

Acknowledgements

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.

References

Behzadi, A., Holmberg, S., Duwig, C., Haghighat, F., Ooka, R., & Sadrizadeh, S. (2022). Smart design and control of thermal energy storage in low-temperature heating and high-temperature cooling systems : A comprehensive review. *Renewable and Sustainable Energy Reviews*, *166*(February), 112625. https://doi.org/10.1016/j.rser.2022.112625

Blum, D., Arroyo, J., Huang, S., Drgoňa, J., Jorissen, F., Walnum, H. T., ... & Helsen, L. (2021). Building optimization testing framework (BOPTEST) for simulation-based benchmarking of control strategies in buildings. Journal of Building Performance Simulation, 14(5), 586-610

Cigler, J., Gyalistras, D., Sˇiroký, J., Tiet, V., & Ferkl, L. (2013a). Beyond theory: The challenge of implementing model predictive control in buildings. Proceedings of 11th rehva world congress, Clima, Prague, Czech Republic.

Deforest, N., Mendes, G., Feng, W., Lai, J., Marnay, C., & Berkeley, L. (2014). Optimal Deployment of Thermal Energy Storage under Diverse Economic and Climate Conditions.

Drgoňa, J., Arroyo, J., Cupeiro, I., Blum, D., Arendt, K., Kim, D., Helsen, L. (2020). *Annual Reviews in Control All you need to know about model predictive control for buildings*. *50*(May), 190–232. https://doi.org/10.1016/j.arcontrol.2020.09.001

Garcia, J., Feng D., Medal L., Grant P., Woolley J., (2024), How many engineers does it take to sell a novel heat pump? Listening to the market and addressing key barriers in cold climate multifamily buildings so the heat pump sells itself

Gu, H., Chen, Y., Yao, X., Huang, L., & Zou, D. (2023). Review on heat pump (HP) coupled with phase change material (PCM) for thermal energy storage. Chemical Engineering Journal, 455(September 2022), 140701. https://doi.org/10.1016/j.cej.2022.140701

Hao, K., Kim, D., & Braun, J. E. (2022). Comparing the economic performance of ice storage and batteries for buildings with on-site PV through model predictive control and optimal sizing. Journal of Building Performance Simulation. https://doi.org/10.1080/19401493.2022.2084161

Helmns, D.; Casillas, A.; Prakash, A.; Blum, D.; Pritoni, M.; Dutton, S., et al. (2022). Heat Pumps with Phase Change Thermal Storage: Flexible, Efficient, and Electrification Friendly. Lawrence Berkeley National Laboratory. http://dx.doi.org/10.20357/B7R015 Retrieved from https://escholarship.org/uc/item/017427ct

Hirschey, J., Gluesenkamp, K. R., Shen, B., Li, Z., & Graham, S. (2023). *A proposed methodology to reduce heat pump size with integrated thermal energy storage 1*. (May), 1–13.

Isermann, R. (2014). *Model-Based Design Workflow for Engine Control Systems*. 52–60.

LBNL, (2024). calflexhub.lbl.gov/calflexhub-affiliate-program/

LG, “Performance Data, Heating Operation,” Therma V R32 Monoblok datasheet, 2020.

Matisoff, Daniel C., et al. "A review of barriers in implementing dynamic electricity pricing to achieve cost-causality." Environmental Research Letters 15.9 (2020): 093006. Available at https://iopscience.iop.org/article/10.1088/1748-9326/ab9a69/meta

Mcneil, A., Kohler, C., & Lee, E. S. (2014). *High Performance Building Mockup in FLEXLAB*.)

Midsummer, (2024), www.midsummerwholesale.co.uk

Naumann, J. D., & Jenkins, A. M. (1982). Prototyping: the new paradigm for systems development. Mis Quarterly, 29-44

R. Kammerud, B. Birdsall, D. Dumortier, B. Andersson, J. Eto, W.L. Carroll and F. Winkelmann, (1989) "Thermal Energy Storage System Sizing" LBL 27049. Applied Science Division, LBL, Berkeley, California 94720, 1989.

Rosenow, J., Gibb, D., Nowak, T., & Lowes, R. (2022). *Heating up the global heat pump market*. *7*(October), 10–13. https://doi.org/10.1038/s41560-022-01104-8

Thomke, S. H. (2003). Experimentation matters: unlocking the potential of new technologies for innovation. Harvard Business Press.

Walker, I. S., Casquero-modrego, N., Less, B. D., Brown, R., & Wei, M. (2023). *Challenges and Opportunities for Home Decarbonization*. (March).

Wetter, M., & Sulzer, M. (2024). A call to action for building energy system modelling in the age of decarbonization. Journal of Building Performance Simulation, 17(3), 383-393.

Zheng, W., Ma, K., & Wang, X. (2015). TE-Shave: Reducing data center capital and operating expenses with thermal energy storage. IEEE Transactions on Computers, 64(11), 3278-3292.