

# Mixed-Integer Model Predictive Control of Variable-Speed Heat Pumps

Zachary Lee, Kartikay Gupta, Kevin J. Kircher, K. Max Zhang\*

*Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, New York 14853, United States*

---

## Abstract

Heating and cooling buildings cause a large and growing portion of the world's carbon emissions. These emissions could be reduced by replacing inefficient or fossil-fueled equipment with efficient electric heat pumps. Recent research has shown that variable-speed heat pumps (VSHPs) could also provide a variety of power system services, such as price- or carbon-based load shifting, peak demand reduction, emergency demand response, and frequency regulation. By providing some or all of these services, VSHPs could facilitate the integration of wind and solar power into the grid, accelerating decarbonization of other economic sectors. However, the provision of power system services poses new and challenging VSHP control problems. Model predictive control (MPC) is a promising approach to these problems. MPC combines weather and load predictions, online optimization, and a mathematical model of a VSHP's performance under varying ambient conditions. With few exceptions, existing VSHP models neglect two important features that arise in practice: (1) dependence of efficiency and capacity on compressor speed and indoor and outdoor temperatures, and (2) inability of VSHPs to operate at low compressor speeds. In this paper, we develop a new VSHP control model that considers both features. The model expresses the VSHP's power consumption and heat production as affine functions of the driving temperatures and compressor speed. The model also includes a binary variable that constrains the VSHP to be either turned off, or operating within a range of acceptable compressor speeds using mixed-integer programming (MIP). This VSHP model and optimization strategy is combined with neural network based heat load prediction in MPC simulations. These simulations suggest that this approach could reduce energy costs by 9 to 22% and carbon emissions by up to 22%, relative to MPC with existing VSHP models.

*Keywords:* Energy systems; built environment; thermal energy storage; electrification of the heating sector

---

## Highlights

- Mixed-integer programming (MIP) is used for variable speed heat pump (VSHP) control.

---

\*Corresponding author

*Email address:* kz33@cornell.edu (K. Max Zhang)

- MIP allows the use of a more complete VSHP model than the current state of the art.
  - The improved model reflects partial-load and temperature-dependent VSHP efficiency.
  - MIP reduces operating costs and energy use compared to the current state of the art.
- 

## 1. Introduction

With rapid growth in population and the floor area of conditioned buildings, global heating and cooling demand are projected to nearly double by 2050, which have enormous economic and environmental consequences [1, 2]. In the United States, for example, heating and cooling residences cost about \$88 billion in 2015 [3]. Heating and cooling costs are particularly burdensome for people living near the poverty line, often comprising 7–15% of their annual income [4, 5]. Space cooling directly contributes to critical peak load during high energy demand days, when the marginal electricity generation sources are most expensive and the atmosphere is most conducive to air pollution formation [6]. Moreover, 60% of United States homes use natural gas, oil, or propane as their main heating source, contributing to about 10% of the carbon emissions nationwide [3, 7]. Therefore, renewable heating and cooling (RHC), *i.e.*, meeting heating and cooling demand with clean, renewable options at costs competitive to fossil fuels, is a crucial component in the transition to sustainable energy systems.

Electric heating and cooling technologies such as air- and ground-source heat pumps have greatly advanced over the last decade. The leading air-source heat pump models have significantly improved their efficiencies at both extremely low (for heating) and high (cooling) ambient temperatures [8]. Powered by renewable electricity generation such as wind and solar, heat pumps can potentially become a viable RHC option. RHC viability increases when thermal energy storage (TES) is added to the heat pump system. TES decouples heat pump operation from heating or cooling demand, allowing the heat pump to operate when efficiency is high and when electricity is clean or inexpensive [9]. Optimal control of heat pumps with and without TES has become a large area of interest as integration of renewable energy sources requires more demand flexibility to match the fluctuations in electricity generation [10]. Several studies have already demonstrated the capability of controlling VSHPs without energy storage to provide grid frequency regulation [11, 12]. However, researchers have shown that coupling heat pumps with TES both increases the instantaneous power flexibility for frequency regulation [13] and enables demand response and load shifting services on longer time scales [14]. These ancillary services are valuable to grid operators and can provide new revenue to VSHP operators [15]. Therefore, we posit that effective heat pump system control mechanisms that have the ability to provide reliable grid services (e.g., frequency regulation, demand response, etc.) while maintaining thermal comfort are critical to realizing the potential of renewable-powered heat pumps as a viable RHC option.

A major challenge in VSHP control is balancing modeling accuracy with computational efficiency. VSHP power consumption and heat production are governed by coupled nonlinear differential equations, which are unsuitable for control and optimization purposes [16]. Thus, simplified but accurate models are needed for real-time control. Constant coefficient

45 of performance (COP) models were adopted in several heat pump control studies [17, 18, 14]. However, constant COP models neglect the dependency of heat pump performance on compressor speed, indoor temperature, and ambient air temperature, thus degrading controller performance [19, 20]. In [21], a partial-load and temperature-dependent COP model using a  
50 third-order polynomial fit was used to minimize the energy required to cool an office building. However, this model assumes operability over the entire range of partial-load ratios and does not consider the reliability constraints associated with running the compressor at low load ratios.

Recently, Kim *et al.* [11] demonstrated that in a given compressor speed operating range, steady-state VSHP performance behaves linearly with respect to ambient air temper-  
55 ature, compressor shaft speed, and indoor air temperature. They formulated a linear model to describe the steady-state heat pump thermal output and power consumption in their demonstration of frequency regulation applications. However, this linear model is unable to capture the true heat pump dynamics at low compressor speeds, as it can imply non-zero power consumption or heat output at zero compressor speed. More discussion of this aspect  
60 can be found in Section S1 of the Supplemental Information. It should be noted that this weakness was not an issue for Kim *et al.*, since their study focused on VSHP control in the operating range and did not consider low compressor speeds [11]. Nevertheless, problems can arise when these linear formulations are used in VSHP control over long time horizons and under varying boundary conditions. In such applications, VSHP models must be accurate  
65 over the entire range of possible compressor speeds and driving temperatures.

Furthermore, VSHPs are designed to not operate below the minimum manufacturer-rated compressor speed. The range of compressor speeds below this value is sometimes referred to as the dead zone. Operating in the dead zone adversely affects the reliability of the compressor by reducing lubrication thickness and increasing frictional losses [22]. This can  
70 reduce the VSHP life and is generally not supported by manufacturers. To avoid this issue, if the desired heat load corresponds to a compressor speed in the dead zone, VSHPs will run at the minimum compressor speed for a shorter duration, increasing energy costs by cycling on and off [23].

The main objective of this paper is to present an effective method for VSHP optimiza-  
75 tion and control through mixed-integer programming (MIP). MIP is a method used to solve optimization problems when the decision variables are a mixture of continuous and integer variables. While MIP has been used in heat pump studies, the heat pumps are often considered as single or dual stage thermostatically controlled devices [24, 25]. To our knowl-  
80 edge, no studies have applied this method to VSHPs while considering both partial load and temperature-dependent COPs. In this paper, we use the linear VSHP model developed in [11] over the operating range of compressor speeds. This enables the use of online optimization, while capturing the dependence of COP on compressor speed, indoor temperature, and outdoor temperature. In addition, we introduce a binary variable into the optimization process to indicate whether the VSHP is turned on. This innovation allows the heat pump  
85 to turn off rather than operate below the minimum rated compressor speed, preventing operation in the dead zone and solving the issue of unrealistic model results at low compressor speeds. We demonstrate the effectiveness of the proposed MIP approach by applying model predictive control (MPC) to a realistic heat pump control problem. MPC has been widely employed in a range of heat pump and TES scenarios and has been shown to be an effective

90 way to optimize heat pump costs [26, 27, 28]. This simulation includes real-world electric-  
 ity rates, historical meteorological conditions, a high-fidelity building thermal model, and a  
 neural network based heat load prediction method to provide VSHP control performance for  
 a winter season.

This paper is organized as follows. Section 2 models the heat pump, TES, and building  
 95 thermal envelope. Section 3 formulates the control problem, describes the proposed mixed-  
 integer MPC algorithm, and discusses three simpler but more common controllers. Section  
 4 compares the performance of the four controllers. Section 5 concludes the paper.

## 2. Thermal modeling

We modeled a single corner apartment in a mid-rise apartment building. The unit is  
 100 equipped with a VSHP in series with TES. The VSHP pumps heat from the outside air into  
 the TES, which then serves the building heating demand. This design allows the TES to  
 function as a heat battery, giving the VSHP more operational flexibility. While we focused  
 on heating, our methods can also be used for cooling, as VSHPs are reversible. We model  
 the three physical components – the VSHP, TES and building envelope – in the remainder  
 105 of this section.

### 2.1. Variable speed heat pump

To capture the dynamics of a VSHP, we adopted the linear model formulated by Kim  
*et al.* [11] Shown in Equation 1, this model gives the steady-state power consumption  $P_t$   
 and heat production  $H_t$  at a time step  $t$  as affine functions of the compressor speed, indoor  
 110 temperature, and outdoor air temperature:

$$\begin{aligned} \text{Power Consumption:} \quad P_t &= \alpha_1 + \alpha_2\omega_t + \alpha_3T_{h,t} + \alpha_4T_{a,t} \\ \text{Heat Output:} \quad H_t &= \alpha_5 + \alpha_6\omega_t + \alpha_7T_{h,t} + \alpha_8T_{a,t}. \end{aligned} \tag{1}$$

Here  $\omega_t$  (rad/s) is the compressor speed,  $T_{h,t}$  ( $^{\circ}\text{C}$ ) is the TES temperature, and  $T_{a,t}$  ( $^{\circ}\text{C}$ ) is  
 the ambient temperature. The model coefficients  $\alpha_i$  can be fit to empirical performance data  
 via multiple linear regression. Experimental results in [29] show that the efficiency reduction  
 associated with startup and shutdown of a heat pump is 2% or less for a duty cycle of 15  
 115 minutes. Therefore, we assume this reduced performance during startup is negligible because  
 the MPC updates each hour. The VSHP efficiency, referred to as the COP, is defined as the  
 ratio of the heat output to the power consumption:

$$\text{COP} = \frac{P_t}{H_t}. \tag{2}$$

For this simulation, we used performance data from Carrier’s Infinity Series VSHP  
 (Model: 25VNA024A) [30] to determine the best-fit model coefficients. Table 1 shows the  
 120 best-fit coefficients. Figure 1 shows a subset of the manufacturer data and the resulting  
 linear equations for power consumption and heat output at an indoor air temperature of  
 21.1 $^{\circ}\text{C}$ . The R-squared values for power consumption and heat output are given as .996 and  
 .994, respectively. While publicly available data detailing partial load VSHP performance is  
 typically very limited, linear performance in the operating range was already experimentally

125 verified by Kim *et al.* [11]. Therefore more detailed partial load data is not required to verify  
linearity with respect to compressor speed.

Figure 1 also shows the implied deviation from the linear model at low shaft speeds. As  
compressor speed approaches zero, steady-state power consumption and heat output must  
also be zero for all driving temperatures. Therefore, there must be some nonlinear behavior  
130 as compressor speed approaches zero to rectify this offset. This implied nonlinear region also  
coincides with operation in the dead zone, the region below the manufacturer minimum  
rated compressor speed where operation significantly decreases the lifespan of the heat pump.

Attribute	Coefficient	Value	Units
Power	$\alpha_1$	-.5922	kW
	$\alpha_2$	.0042	kW s rad <sup>-1</sup>
	$\alpha_3$	.0321	kW °C <sup>-1</sup>
	$\alpha_4$	-.0023	kW s °C <sup>-1</sup>
Heat Output	$\alpha_5$	-.5091	kW
	$\alpha_6$	.0203	kW s rad <sup>-1</sup>
	$\alpha_7$	-.0258	kW °C <sup>-1</sup>
	$\alpha_8$	.1592	kW °C <sup>-1</sup>

Table 1: Coefficients of Fit for Heat Pump Power Consumption and Heat Output




Figure 1: VSHP performance curves at  $T_{indoor} = 21.1^{\circ}C$  for (a) power consumption and (b) heat output. The shaded dead zone in (b) shows where the model breaks down and gives unrealistic heat output at zero compressor speed. (Print in color)

## 2.2. Thermal energy storage

Heat is typically stored in TES by filling an insulated tank with a storage material  
and then changing the material's temperature or phase. Common storage materials include  
135 water, rocks, bricks, paraffins, fatty acids and salt hydrates. In this paper, we adopted a

general model that can represent a variety of temperature- and phase-change thermal storage technologies. The basic model is a first-order continuous-time linear system:

$$\begin{aligned}\dot{E}(t) &= \beta E(t) + \dot{Q}(t) \\ 0 &\leq E(t) \leq E_{\max}.\end{aligned}\tag{3}$$

Here  $t$  (h) denotes time,  $E$  (kWh) is the stored heat, and  $\dot{Q}$  (kW) is the net heat flow into storage. For a given storage technology, the parameter  $\beta$  (1/h) and the storage capacity  $E_{\max}$  (kWh) can be defined in terms of the storage dimensions and thermophysical properties. For control purposes, we considered a continuous time span  $t \in [0, K\Delta t]$ , divided into  $K$  discrete steps of duration  $\Delta t$  (h), indexed by  $k = 0, \dots, K$ . Assuming that the net heat flow  $\dot{Q}$  is constant over each time step,  $\dot{Q}(t) = \dot{Q}(k\Delta t)$  for all  $t \in [k\Delta t, (k+1)\Delta t)$ . With  $E$  and  $\dot{Q}$  at a time step  $k$  denoted by  $E_k$  and  $\dot{Q}_k$ , respectively, the discrete-time analog of the continuous-time storage model (3) is

$$\begin{aligned}E_{k+1} &= aE_k + b\dot{Q}_k, \quad k = 0, \dots, K-1 \\ 0 &\leq E_k \leq E_{\max}, \quad k = 1, \dots, K.\end{aligned}\tag{4}$$

The discrete-time parameters are defined in terms of  $\beta$  and  $\Delta t$ :

$$a = e^{\beta\Delta t}, \quad b = \frac{a-1}{\beta}.$$

These definitions follow from the solution of the governing differential equation (3).

### 2.3. Building envelope

To determine space heating load, we simulated a mid-rise corner apartment facing southeast in Binghamton, NY, using EnergyPlus [31]. EnergyPlus is an open-source whole-building simulation software package that considers indoor and outdoor air temperatures, solar radiation, and other weather values, as well as internal heat loads and building occupancy. The simulated building is the mid-rise commercial reference building developed in [32], which uses the most common construction materials, energy usage, and occupancy schedules. The indoor thermostat setpoint is 21.7 °C. The outside weather conditions came from three weather files for Binghamton, NY: a Typical Meteorological Year (TMY3) and real data from the years 2001 and 2014 [31, 33]. The TMY3 file for Binghamton is derived from 24 years of historical weather data and provides a weather profile that represents a typical year [33]. The TMY3 and 2001 weather files provide a range of conditions that we used to train our model. We then simulated the controller performance under the 2014 weather data.

### 2.4. State-space system model

By combining the linear model in Equation 1 with the discrete time energy storage formulation in Equation 4, a state space model is developed:

$$x_{k+1} = ax_k + \alpha_6 bu_k + w_k.\tag{5}$$

Here the state  $x_k$  (kWh) is the stored TES energy  $E_k$ , the control input  $u_k$  (rad/s) is the compressor speed  $\omega_k$ , and the disturbance  $w_k$  is a function of the outdoor air temperature ( $^{\circ}\text{C}$ ) and building heat demand (kW).

Specific disturbances can be defined for different TES types and system configurations. In our simulation, we modeled a stratified water tank, the most common form of residential TES. Stratified water tanks can be closely approximated to have two layers: a hot layer and cold layer with time-invariant temperatures  $T_h$  and  $T_c$ , respectively [34]. Depending on the amount of thermal energy stored, the thermocline separating these layers will move up and down in the tank, where the maximum and minimum energy storage limits correspond to the minimum and maximum operating height of the thermocline. For a stratified water TES in series with a VSHP, the system’s disturbance is defined as

$$w_k = b [\alpha_5 + \alpha_7 T_h - T_c/R + (\alpha_8 + 1/R) T_{a,k} - Q_{d,k}].$$

Here  $R$  ( $^{\circ}\text{C}/\text{kW}$ ) is the thermal resistance of the thermal storage tank wall and  $Q_{d,k}$  (kW) is the space heating demand.

### 3. Control and prediction

In this section, we formulate the problem of optimally operating the combined heat pump and thermal storage system using MPC. MPC is a control method that predicts system behavior to develop an optimal control schedule for a receding time horizon. MPC is constructed of three main components: (1) a system control model, (2) an optimization framework, and (3) an accurate prediction model for future system inputs.

#### 3.1. Control policies

MPC requires a detailed system control model to determine optimal control schedules. Increasing the physical accuracy of the control model should generally increase MPC performance. However, increasing complexity can cause models to switch from convex to non-convex. While convex optimization algorithms can quickly provide absolute optimal control, non-convex optimization is generally more computationally expensive and relies on heuristic solution methods that can only provide near optimal control. Therefore, the challenge is to design a physically realistic control model that does not significantly hurt optimization time and performance. In our study, we analyzed four VSHP control policies, each with decreasing complexity. The proposed MIP control is used in Policy 1, while Policies 2, 3, and 4 are used as a comparison and represent simpler, more commonly used VSHP control models. The minimum compressor speed constraint in Policy 1 is embedded into the optimization through MIP, which allows the VSHP to turn off when below the minimum compressor speed. For Policies 2, 3, and 4, this constraint is applied by post-processing the control schedule after the optimization process. If the optimization calls for a compressor speed below the minimum speed, the speed is set at the minimum and run for a percentage of the time step to provide an equivalent amount of heat. The true dynamics of the thermodynamic system are updated using Equation 1, which was experimentally verified in [11] to be applicable across the range of operating compressor speeds. Table 2 summarizes the details of each of the control policies.

205 *3.1.1. Policy 1: Mixed-integer MPC*

MIP is used to solve optimization problems when a portion of the decision variables are constrained to be integer values. Constraining some variables to be integer values causes the problem to become non-convex, and a global optimal solution becomes more difficult to find. Common algorithms for solving MIP problems include branch and bound as well as  
 210 heuristic methods such as tabu search and simulated annealing [35].

Despite the increased difficulty of finding an optimal solution, MIP allows the use of a more physically realistic control model that considers COP dependence on compressor speed, outdoor air temperature, and TES temperature. It solves the issue of unrealistic model results at low compressor speeds by using a binary variable in the optimization to  
 215 allow the heat pump to turn off when below the minimum compressor speed. When the heat pump is on, it operates between the minimum and maximum compressor speeds and calculates power consumption and heat output using Equation 1. When the heat pump is off, power consumption and heat output are zero, satisfying the boundary condition at zero compressor speed. This model is shown in Equation 6 with  $\omega_{\min}$  and  $\omega_{\max}$  as the minimum  
 220 and maximum rated compressor speeds, respectively,

$$P_{t,MIP}, H_{t,MIP} = \begin{cases} P_t, H_t & \omega_{\min} \leq \omega \leq \omega_{\max} \\ 0 & 0 \leq \omega < \omega_{\min} \end{cases} \quad (6)$$

*3.1.2. Policy 2: MPC with linear COP*

Policy 2 does not consider COP dependence on compressor speed and only considers COP as a linear function of outdoor air temperature and TES temperature. This model is therefore convex and can provide an absolute optimal solution. To determine the linear  
 225 function for COP, Equation 1 was evaluated at the maximum compressor speed over the operating range of winter temperatures to create a linear regression model with the resulting linear equation:

$$COP = .0446T_{amb} + 2.77. \quad (7)$$

*3.1.3. Policy 3: MPC with Constant COP*

Policy 3 assumes a constant COP model regardless and neglects COP dependence on  
 230 compressor speed and driving temperatures. While seasonal average VSHP COPs are often given by the manufacturer in the form of the heating seasonal performance factor (HSPF), it is important that the constant COP used in the optimization is sized much lower to account for times of low ambient air temperature. Since VSHP COP and capacity are small when ambient temperatures are low, using an inflated COP can cause the control scheme to be  
 235 unable to meet the required heat load at those times. For this reason, the constant COP for Policy 2 was conservatively set at 2.0, which corresponds to the COP at the low range of operating temperatures.

*3.1.4. Policy 4: No TES or optimization*

Policy 4 uses only a VSHP without optimization or thermal energy storage to provide  
 240 a baseline comparison. In this policy, the VSHP exactly provides the heat load from the residence, which is the most common way for VSHPs to operate. Any control optimization



worth implementing should first be able to outperform Policy 4. We assumed that the indoor air temperature remains constant at the same setpoint used in the EnergyPlus simulation at 21.7 °C.

Policy	Control Policy	Convexity	Thermal Energy Storage
1	Mixed-integer	Non-convex	Yes
2	Linear COP	Convex	Yes
3	Constant COP	Convex	Yes
4	None	N/A	No

Table 2: Summary of Optimization Policy Details

### 245 3.2. Optimization framework

The optimization framework must ensure that the VSHP electricity cost is minimized while maintaining thermal comfort of the occupants. The discretized VSHP electricity cost  $g_{e,k}$  at a time step  $k$  can be written as:

$$g_{e,k}(u_k, w_k) = \psi_{e,k} \Delta t P_k y_k. \quad (8)$$

Here  $\psi_{e,k}$  (\$/kWh) is the utility’s hourly residential time-of-use rate given deterministically. In this simulation, we use New York State Electric and Gas Corporation’s (NYSEG) time-of-use electricity supply rate [36]. We only considered the electricity supply rate since taxes and delivery charges contributing to the total electricity cost are time invariant and vary geographically. The time step is given by  $\Delta t$  (h). The heat pump power consumption at a time step  $k$  is given by  $P_k$  (W) and is calculated using Equation 1. The binary variable used to implement MIP is given by  $y_k$  and can have values of zero or one. When  $y_k$  is zero, the heat pump is turned off and the electricity cost is zero. When  $y_k$  is one, the compressor must operate between the minimum and maximum speeds. For Policies 2, 3, and 4,  $y_k$  was set constant at one and the minimum compressor speed was set at zero, since the control post-processing described in Section 3.1 satisfies the compressor speed constraint.

In addition, the TES should not be over or undercharged, as this could be unsafe, decrease efficiency, and decrease occupant comfort. Therefore, an upper limit and lower limit to the amount of energy stored in the thermal storage was imposed, given by  $E_U$  and  $E_L$ , respectively. Moreover, having a certain amount of energy storage stored at all times is preferable, as prediction and model uncertainties can cause the scheduled heat supply to be insufficient to meet heat load. Therefore, a tunable penalty parameter  $\psi_L$  was included to penalize when thermal energy is below a desired percentage of TES capacity  $\gamma$ ,

$$g_{d,k}(x_k) = \psi_L (\gamma E_U - x_k)_+, \quad (9)$$

Together, these constraints and cost functions construct the optimization framework and are shown in Equation 10, where  $N$  is the control horizon. To minimize this function we

used the Gurobi solver [37] in the convex optimization software CVX [38], which are capable  
of efficiently solving mixed-integer linear programs.

$$\begin{aligned}
& \underset{u}{\text{minimize}} && \sum_{k=1}^N (g_{e,k} + g_{d,k}), \\
& \text{subject to} && E_L \leq x_{k+1} \leq E_U, \quad k = 1, \dots, N, \\
& && y_k \omega_{\min} \leq u_k \leq y_k \omega_{\max} \quad k = 1, \dots, N.
\end{aligned} \tag{10}$$

### 3.3. Data-Driven Heat Load Predictions

To efficiently plan the future control scheme, MPC requires predictions of future heat  
load requirements and outdoor air temperature. We assumed a perfect temperature forecast  
since publicly available day-ahead temperature forecasts are often within 1°C of the actual  
temperature values [39]. To predict heat load, we used a Long Short Term Memory (LSTM)  
recurrent neural network. LSTM networks excel in learning patterns in sequence data by  
encoding important information into memory cells that can be used many time steps in the  
future. This memory cell mitigates the problem of exploding and vanishing gradients when  
using a large number of previous time steps as an input to the network.

The main neural network architecture consists of an LSTM layer followed by three fully  
connected layers and an output layer. The previous 22 hours are used as an input into the  
LSTM layer, which has an auxiliary output layer to facilitate the LSTM layer training. This  
auxilliary output ensures that the LSTM layer can learn the patterns of the time-series data  
without them being distorted by the fully connected layers. The 24-hour ahead temperature  
predictions are then concatenated with the LSTM output as input into the fully connected  
layers. The main output layer then outputs the 24-hour ahead heat load predictions. The  
network is trained using the mean squared error loss function. Figure 2 shows the full neural  
network architecture. Optimal hyperparameters were determined through a randomized grid  
search and are shown in Table 3.

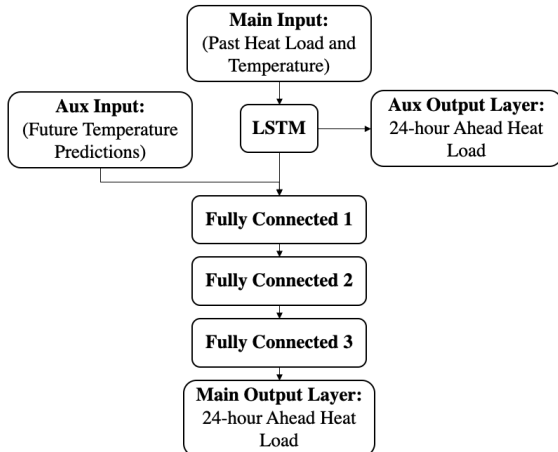


Figure 2: Neural network architecture

Layer	Neurons	Activation
LSTM	399	tanh
Fully Connected 1	139	sigmoid
Fully Connected 2	84	sigmoid
Fully Connected 3	323	relu

Hyperparameters	Value
Optimizer	Adam
Learning Rate	.003

Input	Value
Previous Time Steps	22

Table 3: Neural network hyperparameters obtained through randomized grid search

Hourly training data came from the EnergyPlus simulation for an average mid-rise corner  
apartment in Binghamton, NY, over the the course of the TMY3 and 2001 winter seasons

(December 1 to February 28). The model was trained on 3,592 hourly samples with 634 hourly samples used for model validation and hyperparameter tuning. The model was then used to predict the hourly heat load for the same simulated building using real weather data from the specific year 2014 in a continuous learning model. Each week, the previous week's hourly samples were added to the training dataset, and the neural network was retrained on the slightly larger dataset. This most closely models a real world scenario where prediction models can continuously improve from the constant stream of new data. Figure 3 shows the February 2014 heat load predictions for 1 hour ahead, 12 hours ahead, and 24 hours ahead. These figures show that apart from extremely high or low values, the prediction error is low and therefore does not compound to significantly effect heat load predictions many hours in the future.

#### 4. Results and discussion


Simulations provided optimal control schedules for each of the control policies for the winter of 2014. The simulations were run continuously from December 1st, 2013 until February 28th, 2014 with an MPC horizon of 24 hours updating each hour. Figure 4 shows a subset of the optimal control schedules and the corresponding temperature and electricity price for a five day range in January. Table 4 gives the total energy consumption and electricity cost for the entire winter. It is important to note that this table shows the electricity supply cost, which is only about 30% to 50% of the total electricity cost depending on residence location and taxes.

##### 4.1. Control policy performance


Policy 1 (MIP) provided both the lowest energy consumption and the lowest electricity cost because it can optimize using all relevant knowledge: dynamic electricity price, compressor frequency, and ambient air temperature. Figure 4b shows Policy 1's preference for running less often but with higher power consumption, since Equation 1 results in a higher COP at higher compressor speeds. Policy 2 behaved similarly to Policy 1, concentrating power consumption toward high efficiency, low cost hours. However, the linear COP model of Policy 2 does not capture the COP dependence on compressor speed and therefore operated at low compressor speeds for some hours, reducing the efficiency. Policy 3 concentrated operation toward times of low electricity cost during the night. However, this is often the

Policy	Model	Total Energy Consumption (kWh)	Electricity Supply Cost (\$)
1	MIP	846	32.28
2	Linear COP	921	35.43
3	Constant COP	1086	41.66
4	No TES or optimization	1086	48.77


Table 4: Total simulation results for winter of 2014: Using MIP in Policy 1 significantly reduces electricity usage compared to simpler, more commonly used VSHP control policies.



(a) 1 Hour Ahead Heat Load Predictions



(b) 12 Hours Ahead Heat Load Predictions



(c) 24 Hours Ahead Heat Load Predictions

Figure 3: Heat load predictions for (a) 1 hour ahead and (b) 12 hours ahead, and (c) 24 hours ahead. Apart from extreme values, the prediction error is low and does not accumulate for predictions many hours ahead.

least efficient time to operate due to low outdoor air temperatures, resulting in no decrease in energy consumption compared to the baseline Policy 4. The non-convexity of Policy 1 also incurred only a marginal increase in average computational time of 19% when compared to the convex models in Policies 2 and 3.

Without any optimization or TES, Policy 4 performed significantly worse than the models using TES and MPC, showing that MPC can have a large positive impact on reducing operating costs. However, this reduction in operating costs is limited by the accuracy of the prediction model. For example, a perfect prediction model of future heat load and outdoor temperature can further reduce the operating cost by up to 6% when compared to the pro-

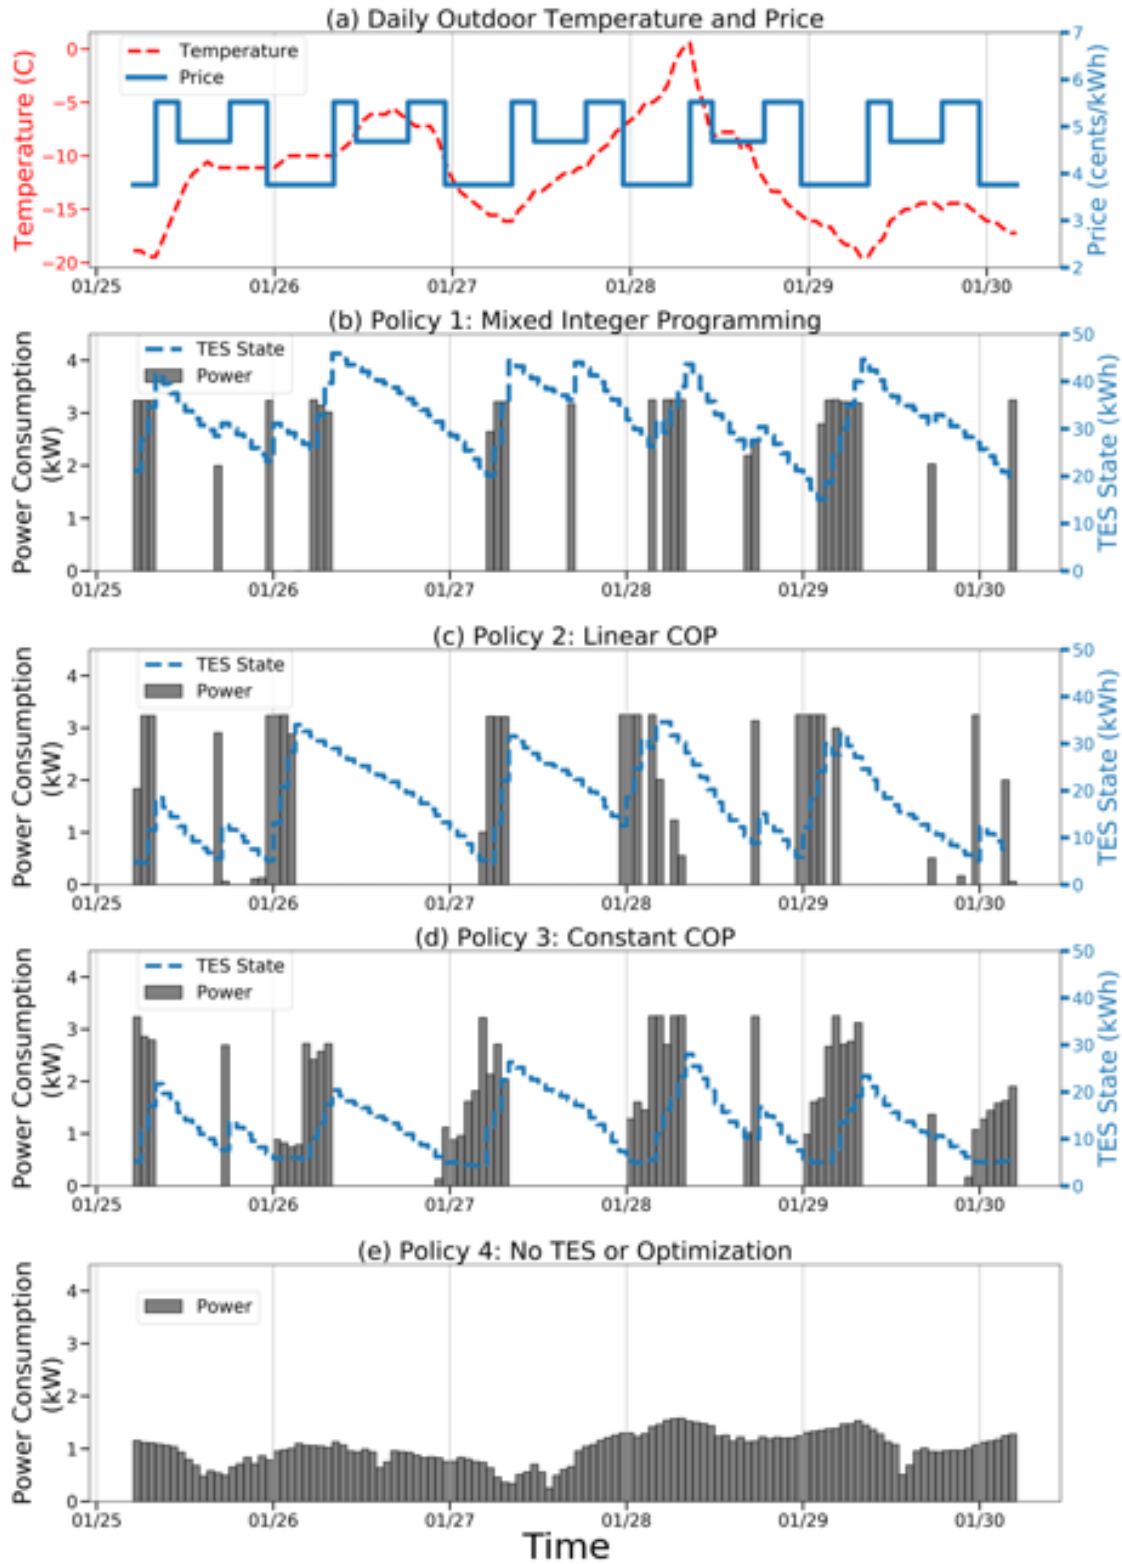


Figure 4: A subset of the control plots for a five day period in January for each control policy. Policy 1 (MIP) runs less frequently at a higher average speed compared to each of the comparison studies, giving a higher average COP and less total energy consumption and cost.

posed prediction model. Therefore, an accurate prediction model is required to capture the full potential benefit of MPC. While the proposed mixed-integer MPC significantly reduces operating costs, the total economic viability of implementing TES and MPC into a home heated by a VSHP requires a more detailed economic analysis on capital costs specific to the residence.

Finally, reducing energy consumption can have a direct impact on reducing carbon emissions. A recent report from the New York Independent System Operator (NYISO) stated that the marginal energy source in Binghamton and Central New York is almost always natural gas [40]. In this case, assuming that a decrease in energy consumption corresponds to a proportional reduction in carbon emissions, Policy 1 (MIP) reduces carbon emissions by 8.1%, 22.1%, and 22.1% relative to Policies 2, 3, and 4, respectively.

#### *4.2. Infeasibility of the linear model without MIP*

Additional simulations were run using the linear program relaxation of the MIP formulation of Policy 1. This control policy used the linear model of Equation 1 without the integer variable, allowing the compressor speed to range from zero to the maximum. If the optimal compressor speed fell below the minimum, the compressor speed was post-processed in the same format of Policies 2 through 4. However, this formulation proved to be unrealistic and caused feasibility errors during the optimization process. This is because the linear model gives non-zero heat output at zero compressor speed, disallowing the heat pump to turn off. Under certain boundary conditions, this causes the TES to charge indefinitely, violating the TES limit constraints. Therefore, this type of control policy does not provide a valid comparison to Policies 1 through 4.

#### *4.3. Effect of TES on control performance*

The potential benefit of TES on MPC performance is highly dependent on the variance of both outdoor air temperature and electricity time-of-use rates. All else equal, with a higher difference between the maximum and minimum of these values, the optimization can take advantage of times of even higher efficiency or lower electricity costs. When compared to the baseline Policy 4, the addition of cost-based optimization in Policy 3 provided a large decrease in energy cost of 14.6%. However, the addition of the complete efficiency-based optimization in Policy 1 resulted in a further cost reduction of 22.5% when compared to Policy 3. This difference implies that the efficiency-based optimization is more effective than the cost-based optimization for this scenario. Binghamton, NY, with an above average number of cloudy days, experiences relatively small temperature swings, limiting the potential benefit of efficiency-based optimization. Therefore, simulating the climate in Binghamton serves as a baseline to show that including the complete cost- and efficiency-based optimization of Policy 1 should outperform other control policies even more in most other climates.

The size of the TES had relatively little effect on the efficiency of the control schemes as long as the TES has the capacity to provide heat load for the control horizon. The energy stored in the TES is optimized to be low to reduce energy loss to the ambient, as well as to minimize costs across the control horizon. Increasing the control horizon could take better advantage of increased TES capacity; however, this comes with a reduction in forecasting accuracy and increased computational time. Other effects that involve the

unsteady thermodynamics of the TES could have an effect on the optimization but were assumed to be negligible due to the relatively long hourly time scale of the optimization.

## 375 5. Conclusion

This paper presents a computationally efficient but realistic variable speed heat pump control optimization model using MIP that solves the unphysical characteristics of linear VSHP models at low compressor speeds. In addition, it prevents the operation of the heat pump in the compressor dead zone, which reduces the life span of the heat pump. By using  
380 machine learning methods to predict future heat load and ambient air temperatures, this method can provide optimal VSHP control schedules for a given time horizon using MPC. While existing VSHP control models often do not consider all variables affecting VSHP COP, using MIP allows the use of a more complete VSHP model that captures COP based on outdoor air temperature, indoor temperature, and compressor speed. The proposed method  
385 outperforms existing VSHP control policies by reducing electricity costs by 9 and 22% and reducing carbon emissions by up to 22%. Since VSHPs and TES can provide ancillary services to the grid in the form of load shifting or frequency regulation, future work can explore expanding this control scheme to include grid signals on faster time scales. Overall, this control method can go toward reducing electricity costs for residences with VSHPs and  
390 TES, diminishing barriers to more efficient and renewable heating and cooling options.

## Acknowledgments

The authors acknowledge the support from the National Science Foundation (NSF) under grant 1711546 and the valuable discussions with Prof. H. Ezzat Khalifa at Syracuse University and Prof. Youngjin Kim at Pohang University of Science and Technology, Korea.

## 395 Declaration of Interest

Declarations of interest: none

## References

- [1] B. Dean, J. Dulac, K. Petrichenko, P. Graham, Towards zero-emission efficient and resilient buildings: Global Status Report, Global Alliance for Buildings and Construction (GABC), 2016.  
400
- [2] D. Ürge-Vorsatz, L. F. Cabeza, S. Serrano, C. Barreneche, K. Petrichenko, Heating and cooling energy trends and drivers in buildings, *Renewable and Sustainable Energy Reviews* 41 (2015) 85–98.
- [3] US Energy Information Administration, Residential Energy Consumption Survey, <https://www.eia.gov/consumption/residential/data/2015/index.php/>, 2018.  
405
- [4] A. Drehobl, L. Ross, Lifting the high energy burden in Americas largest cities, American Council for an Energy-Efficient Economy (ACEEE), 2016.
- [5] NYSERDA Low- to moderate-income market characterization report, New York State Energy Research and Development Authority (NYSERDA), 2017.
- [6] X. Zhang, K. M. Zhang, Demand response, behind-the-meter generation and air quality, *Environmental Science & Technology* 49 (2015) 1260–1267.  
410
- [7] S. Billimoria, M. Hennen, L. Guccione, L. Louis-Prescott, The Economics of Electrifying Buildings: How Electric Space and Water Heating Supports Decarbonization of Residential Buildings, Rocky Mountain Institute, 2018.
- [8] R. Aldrich, J. Grab, D. Lis, Northeast/Mid-Atlantic Air-Source Heat Pump Market Strategies Report 2016 Update, Northeast Energy Efficiency Partnerships, 2017.  
415
- [9] S. N. Palacio, K. F. Valentine, M. Wong, K. M. Zhang, Reducing power system costs with thermal energy storage, *Applied Energy* 129 (2014) 228–237.
- [10] P. Sorensen, N. A. Cutululis, A. Viguera-Rodríguez, L. E. Jensen, J. Hjerrild, M. H. Donovan, H. Madsen, Power fluctuations from large wind farms, *IEEE Transactions on Power Systems* 22 (2007) 958–965.  
420
- [11] Y. J. Kim, L. K. Norford, J. L. Kirtley, Modeling and Analysis of a Variable Speed Heat Pump for Frequency Regulation Through Direct Load Control, *IEEE Transactions on Power Systems* 30 (2015) 397–408.



- 425 [12] E. Vrettos and G. Andersson, Scheduling and Provision of Secondary Frequency Reserves by Aggregations of Commercial Buildings, *IEEE Transactions on Sustainable Energy* 7 (2016) 850–864.
- [13] C. Finck, R. Li, R. Kramer, W. Zeiler, Quantifying demand flexibility of power-to-heat and thermal energy storage in the control of building heating systems, *Applied Energy* 430 209 (2018) 409–425.
- [14] B. Alimohammadisagvand, J. Jokisalo, S. Kilpeläinen, M. Ali, K. Sirén, Cost-optimal thermal energy storage system for a residential building with heat pump heating and demand response control, *Applied Energy* 174 (2016) 275–287.
- [15] B. Kirby, Ancillary Services: Technical and Commercial Insights, Wartsila North America Inc, 2007.  
435
- [16] X.-D. He, Dynamic modeling and multivariable control of vapor compression cycles in air conditioning systems, Ph.D. thesis, Department of Mechanical Engineering, Massachusetts Institute of Technology, 1996.
- [17] M. Awadelrahman, Y. Zong, H. Li, C. Agert, Economic Model Predictive Control for Hot Water Based Heating Systems in Smart Buildings, *Energy and Power Engineering* 440 9 (2017) 112–119.
- [18] R. Halvgaard, N. K. Poulsen, H. Madsen, J. B. Jorgensen, Economic model predictive control for building climate control in a smart grid, 2012 IEEE PES Innovative Smart Grid Technologies (ISGT) (2012) 1–6.
- 445 [19] C. Verhelst, F. Logist, J. V. Impe, L. Helsen, Study of the optimal control problem formulation for modulating air-to-water heat pumps connected to a residential floor heating system, *Energy and Buildings* 45 (2012) 43–53.
- [20] A. Bloess, W.-P. Schill, A. Zerrahn, Power-to-heat for renewable energy integration: A review of technologies, modeling approaches, and flexibility potentials, *Applied Energy* 450 212 (2018) 1611–1626.
- [21] T. Zakula, P. R. Armstrong, L. Norford, Modeling environment for model predictive control of buildings, *Energy & Buildings* 85 (2014) 549–559.
- [22] S. Shao, W. Shi, X. Li, H. Chen, Performance representation of variable-speed compressor for inverter air conditioners based on experimental data, *International Journal of Refrigeration* 455 27 (2004) 805 – 815.
- [23] N. Saraf, Predictive Control for Residential Capacity Controlled Heat Pumps in a Smart Grid Scenario, Master’s thesis, Delft University of Technology, 2015.
- 460 [24] M. C. Bozchalui, S. A. Hashmi, H. Hassen, C. A. Canizares, K. Bhattacharya, Optimal operation of residential energy hubs in smart grids, *IEEE Transactions on Smart Grid* 3 (2012) 1755–1766.

- [25] G. Bianchini, M. Casini, A. Vicino, D. Zarrilli, Demand-response in building heating systems: A model predictive control approach, *Applied Energy* 168 (2016) 159 – 170.
- [26] D. Fischer, T. R. Toral, K. B. Lindberg, B. Wille-Haussmann, H. Madani, Investigation of thermal storage operation strategies with heat pumps in German multi-family houses, *Energy Procedia* 58 (2014) 137–144.
- [27] T. Q. Pean, J. Salom, R. Costa-Castello, Review of control strategies for improving the energy flexibility provided by heat pump systems in buildings, *Journal of Process Control* (2018).
- [28] Candanedo, Jose and Dehkordi, Vahid R., Simulation of Model-based Predictive Control Applied to a Solar-assisted Cold Climate Heat Pump System, *International High Performance Buildings Conference* (2014).
- [29] M. Uhlmann, S. S. Bertsch, Theoretical and experimental investigation of startup and shutdown behavior of residential heat pumps, *International Journal of Refrigeration* 35 (2012) 2138–2149.
- [30] Carrier, 25VNA Infinity Variable Speed Heat Pump with Greenspeed Intelligence 2 to 5 Nominal Tons, <https://www.carrier.com/residential/en/us/products/heat-pumps/25vna0>, 2011.
- [31] D. B. Crawley, C. O. Pedersen, L. K. Lawrie, F. C. Winkelmann, Energyplus: Energy simulation program, *ASHRAE Journal* 42 (2000) 49–56.
- [32] M. Deru, K. Field, D. Studer, K. Benne, B. Griffith, P. Torcellini, B. Liu, M. Halverson, D. Winiarski, M. Rosenberg, et al., U.S. Department of Energy Commercial Reference Building Models of the National Building Stock, *National Renewable Energy Laboratory*, 2011.
- [33] S. Wilcox, W. Marion, Users Manual for TMY3 Data Sets (Revised), *National Renewable Energy laboratory*, 2008. doi:10.2172/928611.
- [34] J. Waluyo, M. Abd Majid, Temperature profile and thermocline thickness evaluation of a stratified thermal energy storage tank, *International Journal of Mechanical and Mechatronics Engineering* 10 (2010).
- [35] A. Richards, J. How, Mixed-integer programming for control, *Proceedings of the 2005, American Control Conference*. (2005) 2676–2683 vol. 4.
- [36] New York State Electric and Gas Corporation, Electricity service rate, 2018. Service class 12, Rate No. 115-12-00.
- [37] Gurobi Optimization, Gurobi Optimizer Version 8.1, <http://www.gurobi.com/>, 2018.
- [38] M. Grant and S. Boyd, CVX: Matlab Software for Disciplined Convex Programming, Version 2.1, Build 1123, <http://cvxr.com/cvx/>, 2018.

- [39] B. Rose, E. Floehr, Analysis of High Temperature Forecast Accuracy of Consumer Weather Forecasts from 2005-2016, Intellovations, LLC, 2017.
- [40] D. B. Patton, P. LeeVanSchaick, J. Chen, R. P. Naga, 2017 State of the Market Report for the New York ISO markets, New York Independent System Operator, Inc., 2017.