Graphical Abstract

Model Predictive Control-Based Optimization of a Transcritical CO_2 Thermal Compressor Heat Pump Cycle Using a RNN-based Reduced Model

A. Salame, J. Peralez, M. Nadri, P. Dufour, V. Lemort

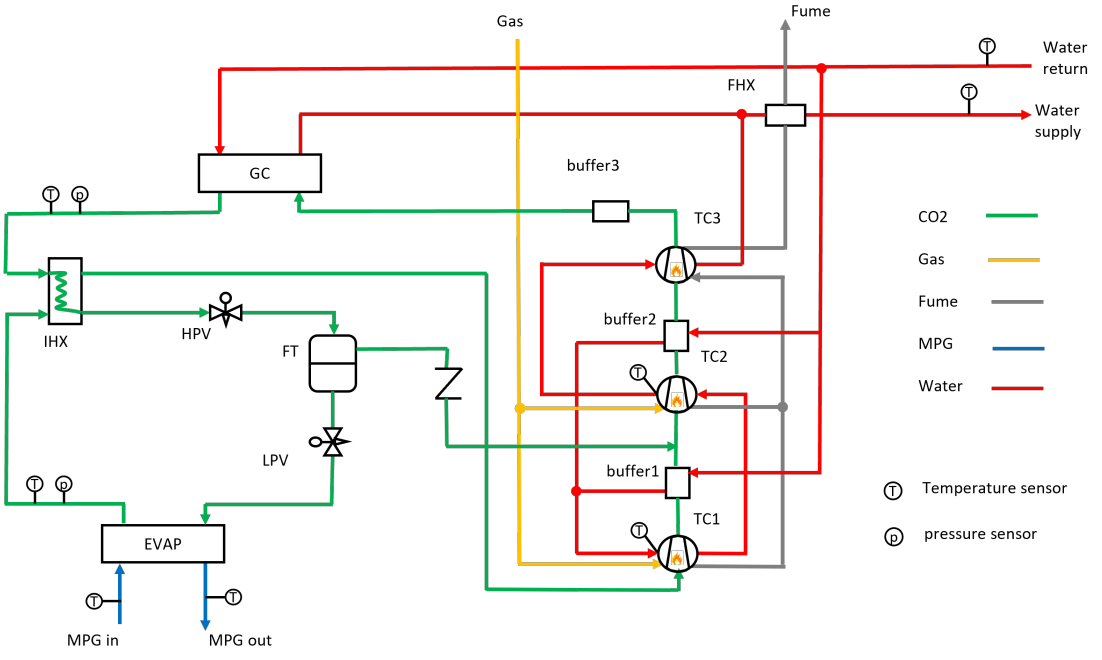

Highlights

Model Predictive Control-Based Optimization of a Transcritical CO_2 Thermal Compressor Heat Pump Cycle Using a RNN-based Reduced Model

A. Salame, J. Peralez, M. Nadri, P. Dufour, V. Lemort

- A physics-based dynamic model is developed using a finite-volume approach coupled with algebraic elements to capture the transient behavior of a multistage CO_2 thermal compressor heat pump (TCHP).
- The model is validated against experimental data for key outputs, including supply water temperature, high-side pressure, and thermal coefficient of performance (COP).
- Reduced-order RNN and LSTM models are trained using a combined dataset of experimental and simulated data to enable fast and accurate prediction.
- A model predictive control (MPC) strategy is proposed and implemented using both RNN and LSTM surrogate models for real-time residential heating control.
- MPC-RNN and MPC-LSTM strategies evaluated against reference PID baseline
- The MPC-LSTM strategy achieves a 20.56% improvement in thermal COP and a 58.69% reduction in temperature tracking error.

Model Predictive Control-Based Optimization of a Transcritical CO_2 Thermal Compressor Heat Pump Cycle Using a RNN-based Reduced Model

A. Salame[a,b,c], J. Peralez[a], M. Nadri[a], P. Dufour[a] and V. Lemort[b]

[a] *Univ Lyon, Université Claude Bernard Lyon 1, CNRS, LAGEPP UMR 5007, Villeurbanne, France*
[b] *University of Liège, Energy Systems Research Unit, , Liège, Belgium*
[c] *Boostheat Company, 69000, Lyon, France*

ARTICLE INFO

Keywords:
MPC, RNN, heat pump, CO2,

ABSTRACT

A Thermal Compressor Heat Pumps (TCHP) primarily utilizes thermal energy sources, such as natural gas or waste heat, to drive a compression cycle for efficient heating applications. TCHP that employ carbon dioxide (CO_2) as a refrigerant requires special consideration of its transcritical characteristics to maintain optimal performance. The unique properties of CO_2, particularly its high operating pressures and critical point behavior, require precise control strategies to ensure efficiency and system stability. To achieve these objectives, a model-based control framework is established here that automates main cycle actuators such as the high-pressure valve and the speed of the burner fan. A dynamic model based on finite volume approach is proposed to capture the transient processes in the TCHP cycle. The supply water temperature, thermal coefficient of performance (COP), and high-pressure outputs are validated with real experimental data. Two control strategies are tested on the dynamic validated model. The first method employs two independent proportional-integral-derivative (PID) controllers, each designed to achieve specific setpoints derived from offline optimal correlations. The second method utilizes model predictive control (MPC) to perform online optimization based on a defined objective function coupled to a vanilla recurrent neural network (RNN) and long short term memory (LSTM) models. The fast to compute vanilla RNN and LSTM models are deduced from real experimental data enriched with data generated from the validated dynamic model that is slow to compute. Both strategies are evaluated on short term, with variations of the supply water temperature setpoint.

1. Introduction

According to Eurostat, in 2024, the energy usage in households accounted for 25.8 % of the final energy consumption, in which 79 % are devoted to space cooling, space heating or water heating Eurostat (2024). In a world increasingly focused on energy efficiency and sustainable technologies, heat pumps have emerged as a pivotal solution for both residential and commercial heating and cooling need. One unconventional type of heat pump is the thermally driven heat pump (TDHP), which can be powered by waste heat, natural gas, or solar radiation. A specific TDHP application developed by Boostheat employs thermal compressors (TCs) instead of electric compressors, referred to as a thermal compressor heat pump (TCHP). Other than belonging to heat pumps family, a TCHP uses CO_2 as a refrigerant, encouraged today due to its environmentally friendly characteristics. However, a CO_2 heat pump has some specifications that needs to be followed in order to get its optimal performance. For example, the performance of CO_2 heat pump cycle is highly affected by the heat rejection pressure (high-pressure), as investigated by Nekså (2002) where he showed that higher coefficient of performance (COP) can be achieved when CO_2 is in its transcritical state. Several works were conducted to find optimal high-pressure correlation for different CO_2 heat pump applications Sarkar, Bhattacharyya and Gopal (2004) and Liao, Zhao and Jakobsen (2000), that eventually could not be generalized. Qi, He, Wang and Meng (2013) performed experimental studies on transcritical CO_2 heat pump systems and developed correlations for the optimal heat rejection pressure based on various parameters. Wang, Tuo, Cao and Xing (2013) examined the effect of heat rejection pressure on system performance in a transcritical CO_2 air-source heat pump (ASHP) water heater under a fixed return water temperature. For a TCHP application, Fallahsohi (2023) carried an energy analysis based on real data, and an optimal high-pressure correlation was derived as function of return water temperature Salame, Lemort, Dufour and Nadri (2024).

*Corresponding author: madiha.nadri-wolf@univ-lyon1.fr
ORCID(s):

Since an optimal high-pressure correlation could not be generalized, several works have proposed an online calculation of the optimal pressure. For example, Zhang and Zhang (2011) proposed a correlation-free online optimal control of heat rejection pressure by adding a dynamic correction term to the high-pressure setpoint. Since no system model was used, this method require online pressure measurements. Cecchinato, Corradi, Cosi, Minetto and Rampazzo (2012) proposed similar approach, but this time it was based on an online artificial neural network (NN) model. Aiming to find the best scenarios for running a transcritical CO_2 refrigeration cycle, Salazar and Méndez (2014) presented PID control schemes for thermal control. To avoid using a model, Hu, Li, Cao and Xing (2015) introduced an extremum-seeking control approach to optimize the COP of a transcritical CO_2 ASHP water heater system. Although extremum-seeking methods are model-free and can be reliable when used in real systems Rampazzo, Cervato, Corazzol, Mattiello, Beghi, Cecchinato and Virzi (2019), but they require more measurements than model based methods. Wang, Zhao, Zhou, Qiao and Cao (2021) proposed the use of a model predictive control coupled with a nonlinear data-driven state space derived from a physical dynamic model of an ASHP water heater. This approach directly maximizes the COP instead of explicitly defining an optimal high-pressure while ensuring that the supply water temperature follows the demand.

Due to the increasing complexity of processes with the increasing constraints and restrictions on new technological advances, advanced control strategies are becoming more necessary. Most of these strategies need reduced models that are both cost effective in terms of simulation and reliable in terms of accuracy. System identification is one of the fields that investigates obtaining a mathematical description of a dynamical system. By observing the input-output behavior of a system, a mathematical model can be derived Ljung (1999). Recently, with the rise of machine learning approaches, deep networks have been used to enhance system identification tools, and this interference is expected to increase Pillonetto, Aravkin, Gedon, Ljung, Ribeiro and Schön (2025). For example, a multi-layered NN was previously used for system identification by Jagannathan and Lewis (1996) proving reliability of system identification based on machine learning models. For time series applications, RNN emerges as a promising system identification tool, for its ability to maintain memory of previous inputs and states through hidden layers. Standard or Vanilla RNNs inspired from Elman (1990) often suffers from vanishing and exploding gradients due to unstable backpropagation and difficulty in learning long-range dependencies because past information fades over time. To overcome these issues, Hochreiter and Schmidhuber (1997) have proposed a long short-term memory (LSTM) NN as an improved RNN that handles gradient issues with gating mechanism, allowing the network to learn long dependencies more effectively. Wang (2017) proposed a new concept using LSTM NNs for dynamic system identification. Lanzetti, Lian, Cortinovis, Dominguez, Mercangöz and Jones (2019) proposed the coupling of a MPC with a RNN due its high scalability and the ability to learn multi-step ahead predictions of complex dynamic systems. RNN models are derived based on an off-line fitting of an input-output data from a validated physical model. Jung, da Costa Mendes, Önnheim and Gustavsson (2023) investigated practical aspects of implementing MPC with LSTM as predictors in a python framework by varying the hyperparameters and used solvers. Ławryńczuk and Zarzycki (2025) demonstrated that LSTM and GRU type RNNs are highly effective and increasingly vital for precise and computationally efficient MPC in nonlinear systems, particularly in industrial and embedded applications like HVAC and vehicles.

The capability of NNs to capture nonlinear behavior makes them ideal for integration with MPC, enhancing the accuracy and adaptiveness of real-time optimization in complex industrial processes Ren, Alhajeri, Luo, Chen, Abdullah, Wu and Christofides (2022). For instance, Turgut and Çoban (2020) developed an NN predictive control approach for a vapor compression refrigeration cycle, using the Whale Optimization Algorithm to determine control inputs at each time-step. The authors tested several objective functions by varying the consumption criterion (COP and exergy efficiency), while tracking a desired cooling load. No significant impact was observed. Most NN–MPC couplings have been applied at the high control level, particularly in building energy management systems Afram, Janabi-Sharifi, Fung and Raahemifar (2017). At this level, the typical control time step exceeds 15 minutes, whereas at the component (low) level, it is often below 1 second Evens and Arteconi (2025). As a result, transient dynamics at the low level are often neglected.

In contrast, this work proposes a RNN–MPC integration for both low- and high-level control, applied to a TCHP operating as a single-house heating unit, with a control time step of 1 second. The objective is to solve an optimization problem that ensures supply water temperature tracking (as a proxy for demand) while minimizing energy consumption. MPC is particularly suitable for this purpose due to its ability to manage multi-input multi-output systems and to anticipate future system dynamics through predictive modeling. Contributions of this work are:

1. Derivation of a physics-based dynamic model for a three-stage CO_2 heat pump cycle in a python framework in object-oriented forms. The transient responses are validated with real experimental data.

2. Integration of low- and high-level control using a MPC with a time-step of 1 second, instead of 1 second for low-level control and 1 hour for high-level control as is usually done.
3. Use of a vanilla RNN and LSTM as reduced-order models for prediction within the MPC framework, connecting supply water temperature (high-level) to internal inputs such as opening of expansion valve and burner fan speed (low-level).

The paper is organized as follows. In section 2, TCHP process is described and control problem is formulated. Section 3 derives a physics-based model for a TCHP and transient validation with real data is provided. Section 4 defines reduced models (vanilla RNN and LSTM), identified from real and physics-based (reference) model generated data, with the main variables satisfying the application purpose. Section 5 proposes a MPC design for the TCHP application and describes the previously applied PID-based strategy. Section 6 shows a validation of the control strategies named as PID, MPC-RNN, and MPC-LSTM applied on the reference model under different supply water temperature setpoints, providing a performance comparison based on temperature error and thermal COP. The work is then concluded in section 7, providing some insights about future possible work.

2. Problem statement

2.1. TCHP Process Overview

The TCHP, developed by Boostheat, serves as a heating application for residential use. TCHP cycle represented in Figure 1 consists of three TCs, three buffers, one gas cooler (GC), one evaporator (EVAP), an internal heat exchanger (IHX), flash tank (FT), two electronic expansion valves (EEVs), and a fume heat exchanger (FHX). At evaporation pressure, CO_2 enters the first-stage compressor and exits at a higher pressure. It is then cooled inside buffer 1 with a water flow and mixed with the vapor exiting the FT through a one-way valve before entering the second TC. The CO_2 then exits at a higher intermediary pressure and is cooled again inside the second buffer before entering the third TC. The buffers after each TC also act as dampers for the pressure waves resulting from compression. At its highest pressure level and high temperature, CO_2 rejects heat into the water entering the GC. It is then cooled again while passing inside one side of the IHX by the CO_2 exiting the EVAP, which passes through the other side. After exiting the IHX, the CO_2 refrigerant is expanded in a high-pressure valve (HPV), then enters a FT, where the CO_2 in its gaseous state is re-injected between the first and second TCs through a one-way valve, and the remaining liquid exits and expands in a low-pressure valve (LPV) before finally entering the EVAP.

For the water side (heating circuit): a smaller portion of the water flow goes to the buffers and the TCs, where it is subdivided again between the first two buffers and collected before consequently entering the coolers of the TCs. The larger part of the entering water goes to the gas cooler, where the two parts mix again before entering the FHX, where the resulting fumes from the combustion process are cooled.

On the TCs side, the top parts of the first two TCs, i.e. heaters are connected to a burner fan, which is used to deliver the methane/air mixture into the top surfaces, where the gas mixture is burned to heat up. The resulting fume from both chambers is then directed to the top surface of the third TC, before exiting into the FHX. The EVAP is heated by a mixture of Monopropylene Glycol-water (MPG), which is cooled by a chiller to a certain temperature value to mimic cold weather temperatures. MPG is used because of its low freezing point, thus preventing freezing.

While heat exchangers allow heat transfer between the refrigerants and external fluids, the TCs as well as the expansion valves dictate the refrigerant mass flow rate and pressure differences between each two consecutive stages. The constant change of thermophysical and heat transfer properties of the refrigerant causes high nonlinearities and complex couplings inside. This forms a challenging task for dynamic modeling and optimal control attempts. In the next section, we will formulate the control problem of the TCHP process and give the control structure connecting the TCHP cycle to the water circuit.

2.2. Control Problem Formulation

The TCHP follows a certain control logic that guarantees its proper behavior. This control is usually hierarchical, where a goal (the user asks for a certain room temperature) is specified, then a high-level control (i.e. supervisor) guarantees it by providing setpoints for the low level cycle controllers that ensure that these setpoints are followed. The major actuators guiding the TCHP operation and corresponding measurements are presented in Figure 2. The connections between them are further explained:

- The motor speeds of the three TCs ($\omega_{m1}, \omega_{m2},$ and ω_{m3}) where the first one affects the thermodynamic condition at the outlet of the EVAP and the second two affect the thermodynamic conditions at the inlet of the GC.

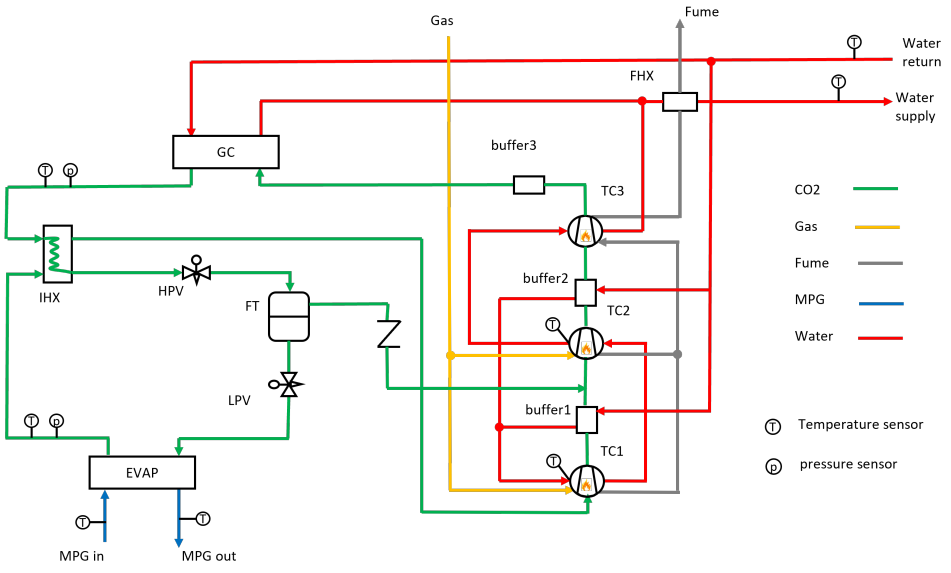

Figure 1: Test bench layout of the TCHP cycle.

- The burner fan speed ω_{bf} regulates the amount of gas being burned, thus controlling the heaters temperatures (T_{heater1} and T_{heater2}) and consequently the mass flow rate induced by the TCs.
- The mass flow rate leaving the GC is directly affected by the opening of the HPV (φ_{HPV}), which is commonly used to control the high-pressure (p_{gc}) at the GC.
- The mass flow rate entering the EVAP is directly affected by the opening of the LPV (φ_{LPV}), which is usually used to control the superheat (defined as the difference between the EVAP outlet temperature ($T_{evap,out}$) and the saturation temperature).
- The water and MPG pumps on the GC and EVAP sides regulate the mass flow rates $\dot{m}_{\text{w}}$ and $\dot{m}_{mpg}$ and control temperatures differences on the secondary fluids sides.
- The return water temperature $T_{\text{w, return}}$ is influenced by room and outdoor temperatures, while the inlet MPG temperature $T_{mpg,in}$ is only influenced by the latter. They are both considered as floating inputs or otherwise disturbances.

The supply water temperature setpoint $T^{\text{sp}}_{\text{w, supply}}$ is defined according to a an offline-correlation as function of outdoor temperature and a desired room temperature that is assumed fixed at 20 °C in this work. The return water recovers heat from several heat exchangers (GC, buffers, TCs coolers, and FHX) before being supplied to the user. Except for the FHX, other heat exchangers are primarily influenced by the heater temperatures T_{heater1} and T_{heater2}, which exhibit slower dynamics compared to the burner fan speed ω_{bf}. In contrast, ω_{bf} directly affects the heat exchange rate in the FHX. Therefore, we focus on the inputs that have the most significant impact on the supply water temperature $T_{\text{w, supply}}$, namely: φ_{HPV}, ω_{m2}, ω_{m3}, ω_{bf}, $\dot{m}_{\text{w}}$, and $T_{\text{w, return}}$.

This study focuses on ω_{bf} and φ_{HPV} as primary control inputs due to their significant impact on performance; the former which is the major responsible for running the TCs (inducing mass flow rates) and more directly the heat recovered by FHX. The latter directly influences the high-pressure which have already been proven to be highly influencing to the overall performance. $\dot{m}_{\text{w}}$ and $T_{\text{w, return}}$ are treated as disturbances, while ω_{m1} and ω_{m2} are fixed at a constant value (150) rpm due to their low influence. Keeping ω_{m1}, ω_{m2}, and $\dot{m}_{\text{w}}$ constant ensures consistent electric power consumption for fair performance comparisons. On the EVAP side, superheat is regulated via a closed-loop PID controller adjusting φ_{LPV}.

The nonlinear and dynamic nature of the system, especially the fast responses of ω_{bf} and φ_{HPV}, complicates control design, requiring a robust strategy. A MPC framework is developed to optimize these inputs while achieving the

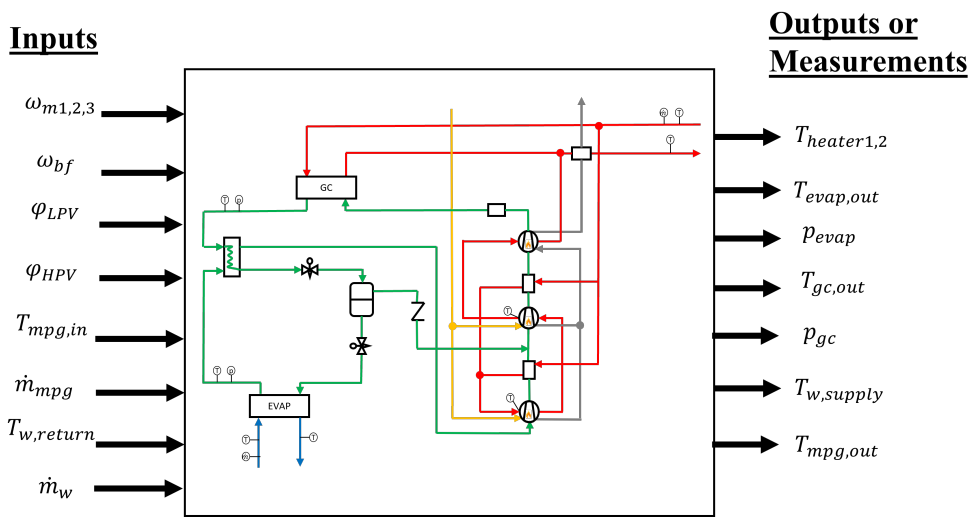

Figure 2: TCHP inputs and outputs (measurements).

supply water temperature setpoint and minimizing energy consumption. Energy use is mainly tied to gas consumption, controlled by ω_{bf}, unlike conventional electric heat pumps where compressor speed is the primary variable. The MPC objective minimizes the supply water temperature deviation from setpoint and ω_{bf}, a proxy for energy use. The MPC leverages a reduced data-driven RNN model, derived from real data and validated against a physical model, as detailed in the next section.

3. Dynamic Model for the TCHP

To develop a model for the TCHP, individual components are first modelled. CoolProp by Bell, Wronski, Quoilin and Lemort (2014) is used to compute the thermophysical properties of CO_2 in a Python framework. In conventional electric-driven VCC systems, the expansion valve and compressor are assumed to operate in a steady-state, as their dynamics occur more rapidly compared to those of the heat exchangers. The use of a TC instead of a conventional compressor changes this rule. In fact, the slowest dynamics in a TCHP cycle correspond to the TCs due to the high inertia of their heaters, which also causes the slow changes in the refrigerant characteristics (mass flow rate, temperature, and pressure). Since we can not represent TC dynamically, as it would largely exhaust the computation cost, we consider the heater temperature as an input rather than burner fan speed. This allows us to safely represent the TC with steady-state regression models. The dynamic variations of the EEVs are significantly faster than those of the heat exchangers. Consequently, the EEVs and TCs are modeled statically using algebraic equations, while the dynamic behavior of the heat exchangers and flash tank is captured through differential equations. To ensure a simpler model run, the following assumptions are adopted: i) No pressure drop is considered on either side of the fluid, ii) Mass variation is only accounted for in the last control volume of components: GC, EVAP, buffer1, and FT, while no variation is considered in the IHX and buffer2. iii) Heat conduction fluxes are neglected and iv) The secondary fluids are assumed to be incompressible

3.1. Heat Exchangers

Beyond the primary heat exchangers (GC and EVAP), TCHP cycle incorporates several additional heat exchangers, including IHX, FHX, TC coolers, and two buffers. The GC and EVAP, constructed as brazed plate heat exchangers made of copper, are specifically designed to withstand the high-pressure conditions associated with CO_2 (see Table 1). These heat exchangers are spatially discretized into several control volumes (see Figure 3). The components such as FT and buffers are treated as single control volumes. For the CO_2 working fluid, the continuity principle yields the following expression for the time derivative of density in the i-th control volume:

$$a_i \frac{\partial p_i}{\partial t} + b_i \frac{\partial h_i}{\partial t} = \dot{m}_{f,j-1} - \dot{m}_{f,j}, \tag{1}$$

Table 1
Design parameters of the GC and EVAP.

Model	Number of Plates	Total Volume [L]	Max Temp [°C]	Max Pressure [bar]	Total Area [m^2]
GC (C042)	56	1.65	200	140/30	1.0422
EVAP (C096)	28	1.917	200	100/30	1.235

where the coefficients a_i and b_i are defined as:

$$a_i = V_i \left(\frac{\partial \rho_i}{\partial p_i} \right)_{h_i}, \tag{2}$$

$$b_i = V_i \left(\frac{\partial \rho_i}{\partial h_i} \right)_{p_i}. \tag{3}$$

For the energy conservation equation on the refrigerant side, the expression involving the heat transfer rate $\dot{Q}$ is:

$$c_i \frac{\partial p_i}{\partial t} + d_i \frac{\partial h_i}{dt} = \dot{m}_{f,j-1} h_{j-1} - \dot{m}_{f,j} h_j - \dot{Q}_{f,i}, \tag{4}$$

where the coefficients c_i and d_i are given by:

$$c_i = V_i \left[h_i \left(\frac{\partial \rho_i}{\partial p_i} \right)_{h_i} - 1 \right], \tag{5}$$

$$d_i = V_i \left[h_i \left(\frac{\partial \rho_i}{\partial h_i} \right)_{p_i} + \rho_i \right]. \tag{6}$$

Here, the thermodynamic derivatives $\left(\frac{\partial \rho}{\partial h} \right)_p$ and $\left(\frac{\partial \rho}{\partial p} \right)_h$ are evaluated at constant pressure and constant enthalpy, respectively. In the IHX, Equation (4) is applied to both CO_2 streams, while no mass variation is assumed to take place. For heat exchangers with water or MPG as the secondary fluids, the temperature variation is described by:

$$\left(mc \frac{\partial T}{\partial t} \right)_{sf,i} = \dot{Q}_{sf,i} + \dot{m}_{sf,j} h_{sf,j} - \dot{m}_{sf,j-1} h_{sf,j-1}. \tag{7}$$

The temperature variation of the heat exchanger wall is modeled as:

$$\left(mc \frac{\partial T}{\partial t} \right)_{wall,i} = -(\dot{Q}_{f,i} + \dot{Q}_{sf,i}). \tag{8}$$

The heat transfer rates for the working fluid (CO_2) and secondary fluid are expressed as:

$$\dot{Q}_{f,i} = U_{f,i} A_i (T_{wall,i} - T_{f,i}), \tag{9}$$

$$\dot{Q}_{sf,i} = U_{sf,i} A_i (T_{wall,i} - T_{sf,i}). \tag{10}$$

Accurate determination of the heat transfer coefficient (U_f) is critical for modeling heat exchange between CO_2 and the secondary fluid. This coefficient varies with the phase of CO_2 (supercritical, sub-critical, or two-phase) and the direction of heat transfer (cooling or heating). Although extensive research has been conducted to define suitable correlations for CO_2, no universally reliable generalization has emerged. Consequently, we developed custom correlations for the heat exchangers by tuning the Dittus-Boelter correlation with experimental data from Fallahsohi (2023). The Nusselt number (Nu) correlations, based on CO_2 phase, are as follows:

In the supercritical region ($p > 73.8$ bar):

$$Nu = 0.005 Re^{0.9} Pr^{0.18}. \tag{11}$$

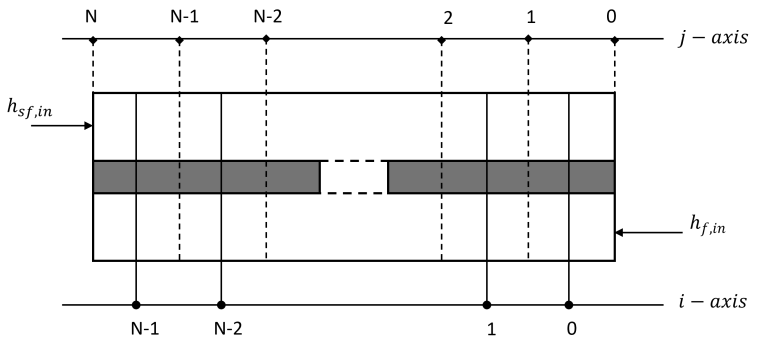

Figure 3: Spatial discretization of a heat exchanger into N control volumes. Each CV on the $i-axis$ carries uniform pressure and enthalpy. While $j-axis$ carries uniform mass flow rates information.

Similarly used by Yoon, Kim, Hwang, Kim, Min and Kim (2003) but with different constant parameters. In the subcritical region ($p < 73.8$ bar) and two-phase ($0 < q < 1$):

$$Nu = 0.1 Re^{0.8} Pr^{0.6} q^{0.2}. \tag{12}$$

In single-phase conditions, the standard Dittus-Boelter correlation is applied:

$$Nu = 0.023 Re^{0.8} Pr^{0.4} \tag{13}$$

The single phase Dittus-Boelter is used for the secondary fluid heat transfer coefficient U_{sf}. For simulation purposes, the GC, EVAP, and IHX are discretized into 8, 12, and 4 control volumes, respectively, which offer a good trade-off between model accuracy and computational efficiency. The buffers and FT are modeled with a single control volume each since heat exchange rate is not critical in these components. In fact, these components are mainly needed for creating pressure differences between stages. For the FT, the model by Qiao, Aute and Radermacher (2015) is adopted. The coolers, key heat recovery components in the TCHP, are modeled using regression techniques, to be elaborated in Thermal compressors section.

3.2. Electronic Expansion Valves

The TCHP cycle employs two electronic expansion valves: the low-pressure valve (LPV) and the high-pressure valve (HPV). The LPV regulates the mass flow rate entering EVAP by adjusting its opening in percentage (φ_{LPV}), while the HPV regulates the mass flow rate leaving GC via its opening in percentage (φ_{HPV}).

Assuming the valves to be adiabatic and the flow to be isenthalpic, the mass flow rates through the valves are determined by the flow coefficient, the flow area, the inlet density, and the pressure difference across the valve:

$$\dot{m}_{\text{LPV}} = (C_v.A)_{\text{LPV}} \sqrt{\rho_{\text{ft,l}}(p_{\text{ft}} - p_{\text{evap}})}, \tag{14}$$

$$\dot{m}_{\text{HPV}} = (C_v.A)_{\text{HPV}} \sqrt{\rho_{\text{gc}}(p_{\text{gc}} - p_{ft})}. \tag{15}$$

The combined value of the product $(C_v.A)$ can be obtained by a regression relation with percentage valve opening φ fitted with experimental data:

$$C_v.A = c_0 + c_1.\varphi + c_2.\varphi^2. \tag{16}$$

Where the c_0, c_1, c_2 parameters are tuned using steady-state data from Fallahsohi (2023). In reality, the valve opening does not change instantaneously and their behavior is expected to follow the following first order forms:

$$\dot{\varphi}_{\text{HPV}} = \frac{1}{\tau_{\text{HPV}}}(\varphi_{\text{HPV, target}} - \varphi_{\text{HPV}}), \tag{17}$$

$$\dot{\varphi}_{\text{LPV}} = \frac{1}{\tau_{\text{LPV}}}(\varphi_{\text{LPV, target}} - \varphi_{\text{LPV}}). \tag{18}$$

Time constants of $\tau_{\text{HPV}} = 20s$ and $\tau_{\text{LPV}} = 10s$ are calibrated to replicate the real transient responses.

Table 2
Prediction accuracy for each output: MAPE (%) for flow and power variables, MAE K for temperatures, and R^2 for all outputs.

Output Variable	MAPE (%) / MAE (K)	R^2
$\dot{m}_{TC1}$	13.1 %	0.92
$T_{TC1,dis}$	0.9 K	0.987
$\dot{Q}_{\text{cooler1}}$	2.3 %	0.984
$\dot{m}_{TC23}$	9.6 %	0.94
$T_{TC23,dis}$	1.1 K	0.98
$\dot{Q}_{\text{cooler23}}$	2.8 %	0.96

3.3. Thermal Compressors

The TC exhibits the slowest dynamics in the system due to high thermal inertia on the heater. As such, heater temperatures are treated as independent variables in regression models, so the inertia or slow dynamic is embedded inside the derived models. Step tests on the third motor speed (ω_{m3}) showed negligible impact on system variables and are thus excluded. The third TC heater temperature, lacking direct measurements, is combined with the second TC to avoid model stiffness from coupling algebraic mass flow and differential pressure equations. Regression models for the second and third TCs (TC23) define the mass flow rate entering the GC ($\dot{m}_{TC23}$), discharge temperature ($T_{TC23,dis}$), and heat recovery ($\dot{Q}_{\text{cooler23}}$):

$$LR_{23,i}(p_{r23}, T_{r2}, \omega_{\text{m2}}), \quad i \in \{1, 2, 3\}, \tag{19}$$

where p_r and T_r being the pressure and temperature ratios. Temperature ratio is the ratio of heater to water temperatures (both in Kelvin). In this way the inertia of the heater is embedded inside the correlation. For the first TC (TC1), regression models define the mass flow rate leaving the EVAP ($\dot{m}_{TC1}$), discharge temperature ($T_{TC1,out}$), and heat recovery ($\dot{Q}_{\text{cooler1}}$):

$$LR_{1,i}(p_{r1}, T_{r1}, \omega_{\text{m1}}), \quad i \in \{1, 2, 3\}. \tag{20}$$

The accuracies of these models (mean absolute error (MAE), mean absolute percentage error (MAPE), and R^2) are summarized in Table 2. Linear regression (LR_i) was selected to avoid overfitting given the limited sample size. The primary limitation of these models is their exclusivity to a certain data range. The full representation of these linear regression models can be found in appendix A.

3.4. Fume Heat Exchanger

After the combustion process takes place, part of the heat is directly absorbed by the first two TCs, then it goes to the third and then the remaining enters the FHX to recover it with circulating water. From previous experimental work, the fume temperature entering the FHX was estimated as function of burner fan speed:

$$T_{fume,in} = 0.0625\omega_{\text{bf}} + 193.26. \tag{21}$$

The temperature of the water entering the FHX is estimated as:

$$T_{fhx,w,in} = (\dot{Q}_{\text{gc}} + \dot{Q}_{\text{cooler1}} + \dot{Q}_{\text{cooler23}} + \dot{Q}_{buffer1,2})/(\dot{m}_{\text{w}} cp_w) + T_{\text{w, return}}. \tag{22}$$

The dynamic change of the water temperature entering FHX is respected, as heater temperature is embedded inside the heat transfer rates expressions in (22). The heat exchange rate in the FHX is considered faster than other components as it is directly influenced by the burner fan speed rather by the heater temperature. So the dynamics of the heat exchanger is dominant. In other words, when changing the burner fan speed, the FHX has faster influence on varying water temperature compared to other heat exchangers mainly depending on heater temperature. The heat recovered by the FHX can be estimated using logarithmic method:

$$\dot{Q}_{\text{fhx}} = AU_{\text{fhx}} \cdot \Delta T_m, \tag{23}$$

$$\Delta T_m = \frac{(T_{fhx,fume,in} - T_{fhx,w,out}) - (T_{fhx,fume,out} - T_{fhx,w,in})}{\ln\left(\frac{T_{fhx,fume,in} - T_{fhx,w,out}}{T_{fhx,fume,out} - T_{fhx,w,in}}\right)}, \tag{24}$$

where AU_{fhx} is a tuned correlation of dependent variables $T_{fhx,w,in}$ and ω_{bf} (appendix C). The fitting resulted in a MAPE of 15.8 % (Figure 4). After we distinguished between the influence of the burner fan speed and heater temperatures, a connection is needed and necessary to be made. This is carried out in the next subsection where we attempt to model dynamic change of heater temperature as function of burner fan speed.

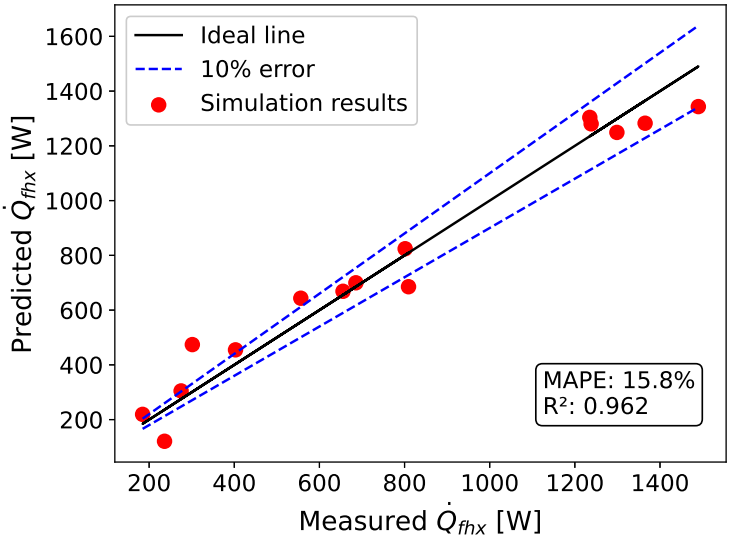

Figure 4: Parity plot comparing the predicted and measured heat exchange rates in the fume heat exchanger.

3.5. Thermal Compressor Heater

We intend in this section to model the dynamic change of the heaters temperatures. The heater temperature is heated by gas on one side (externally) and cooled by CO_2 from another (internally). The study of the combustion chamber is out of this work scope, and the internal modeling of the TC is also omitted. One solution is to treat it as a black box. The corresponding inputs of the external and internal heat transfer are burner fan and motor speeds respectively, so a system identification based on real experimental step tests is carried to find the good models. The identification is done with 3 different data-sets, first one is a pseudo random binary sequence (PRBS) on burner fan speed, while second contains steps tests, and third is step tests on motor speed. The input ranges are: $80 \leq \omega_{\text{m1}}, \omega_{\text{m2}} \leq 240$ rpm, and $1950 \leq \omega_{\text{bf}} \leq 9500$ rpm. A PRBS is needed here to cover the entire range of burner fan speed.

The defined state vector is $[T_{\text{heater1}},\ T_{\text{heater2}}]^T$, the input is ω_{bf}, and the disturbance vector is $[\omega_{\text{m1}},\ \omega_{\text{m2}}]^T$. The output vector is equivalent to the state vector, implying full-state measurement.

The system identification model is nonlinear and can be expressed as:

$$\begin{bmatrix} T_{\text{heater1}} \\ T_{\text{heater2}} \end{bmatrix}(k+1) = \beta_0 + \beta_1 \begin{bmatrix} T_{\text{heater1}} \\ T_{\text{heater2}} \end{bmatrix}(k) + \beta_2 \begin{bmatrix} \omega_{\text{bf}} \\ \omega_{\text{bf}}^2 \end{bmatrix}(k) + \beta_3 \begin{bmatrix} \omega_{\text{m1}} \\ \omega_{\text{m2}} \end{bmatrix}(k), \tag{25}$$

where $\beta = [\beta_0, \beta_1, \beta_2, \beta_3]$ is retrieved from the minimization of the ordinary least squares error between the measured and predicted T_{heater1} and T_{heater2}. The results in Figures 5 and 6 show relatively high accuracy, admitting MAEs of 15.86 K and 16.77 K and R^2 scores of 0.96 and 0.94 for heaters temperatures 1 and 2, respectively (Appendix B). The identified heaters temperatures models admit an acceptable fit towards real data, while accuracies can be further improved. However, since it is also out of this work scope, these models are deemed reliable and are used as part of the TCHP model.

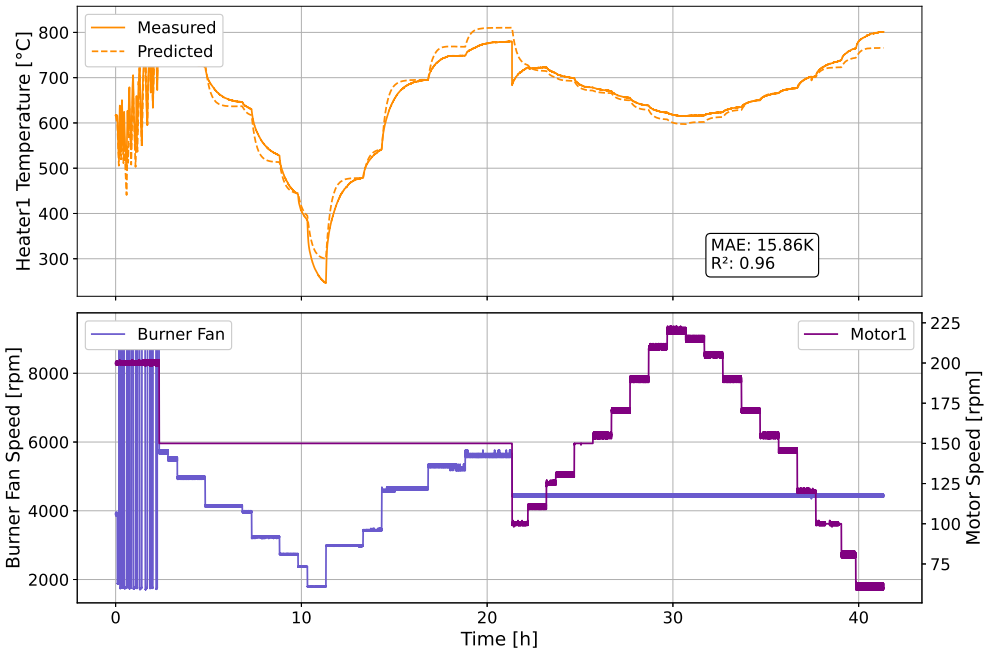

Figure 5: Comparison of observed and predicted T_{heater1} temperature over time. MAE is 15.86 K and R^2 is 0.96, indicating a good fit. The bottom plot shows the burner fan speed and motor speed over time, illustrating their influence on T_{heater1} dynamics.

3.6. TCHP Performance Characterization

In an electrically driven heat pump, the COP is defined as the ratio of the total recovered heat to the electric power consumed by the cycle. In this work, heat is the primary source of energy. With a small contribution of electricity compared to heat and for simplicity, TCHP is evaluated with a thermal COP, defined as the ratio of total heat recovered $\dot{Q}_{\text{heat, tot}}$ to total heat consumed $\dot{Q}_{\text{comb}}$ (combustion power):

$$\text{COP}_{\text{thermal}} = \frac{\dot{Q}_{\text{heat, tot}}}{\dot{Q}_{\text{comb}}}. \tag{26}$$

The total heat recovered is distributed as follows:

$$\dot{Q}_{\text{heat, tot}} = \dot{Q}_{\text{gc}} + \dot{Q}_{\text{cooler1}} + \dot{Q}_{\text{cooler23}} + \dot{Q}_{buffer1,2} + \dot{Q}_{\text{fhx}}. \tag{27}$$

And the combustion power is:

$$\dot{Q}_{\text{comb}} = \dot{m}_{\text{CH}_4} \cdot \text{LHV}, \tag{28}$$

where $\dot{m}_{\text{CH}_4}$ is the methane gas mass flow rate which is proportional to the burner fan speed and LHV is its lower heating value. The resulting state vector of a TCHP

$$\begin{aligned} X = [&p_{\text{gc}}, \bar{h}_{\text{gc}}, \bar{T}_{\text{gc,wall}}, p_{\text{evap}}, \bar{h}_{\text{evap}}, \bar{T}_{\text{evap,wall}}, p_{\text{buff1}}, h_{\text{buff1}}, T_{\text{buff1,wall}}, h_{\text{buff2}}, \\ &T_{\text{buff2,wall}}, \bar{h}_{\text{ihx1}}, \bar{h}_{\text{ihx2}}, \bar{T}_{\text{ihx,wall}}, p_{\text{ft}}, h_{\text{ft}}, T_{\text{heater1}}, T_{\text{heater2}}]^\top. \end{aligned} \tag{29}$$

The fluid enthalpies $\bar{h}$ and wall temperatures $\bar{T}_{wall}$ are represented as vectors for the EVAP, GC, and IHX, where their lengths are 12, 8, and 4 respectively. The input and output vectors are respectively (Figure 2):

$$U = [\varphi_{\text{HPV}}, \varphi_{\text{LPV}}, \omega_{\text{bf}}, \omega_{\text{m1}}, \omega_{\text{m2}}, T_{\text{w, return}}, \dot{m}_{\text{w}}, T_{\text{mpg, in}}, \dot{m}_{\text{mpg}}]^\top, \tag{30}$$

$$Y = [T_{\text{heater1}}, T_{\text{heater2}}, p_{\text{gc}}, p_{\text{evap}}, T_{\text{gc,out}}, T_{\text{evap, out}}, T_{\text{w, supply}}, T_{\text{mpg, out}}]^\top. \tag{31}$$

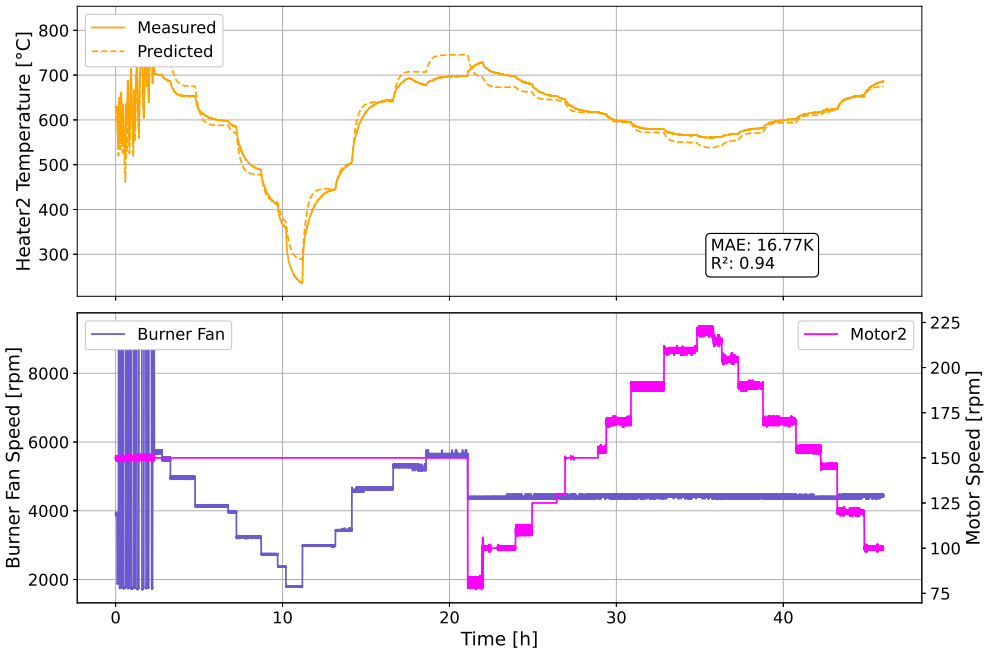

Figure 6: Comparison of observed and predicted T_{heater2} temperature over time. MAE is 16.77 K and R^2 is 0.94, indicating a good fit. The bottom plot shows the burner fan speed and motor speed over time, illustrating their influence on T_{heater2} dynamics.

Tests	MAPE [%] or MAE [K]		R^2 [-]	
	HPV	BF	HPV	BF
p_{gc}	1.8	2.88	0.92	0.52
$T_{\text{w, supply}}$	0.33 K	0.18 K	0.6	0.8
$\text{COP}_{\text{thermal}}$	2.74	4.26	0.48	0.62

Table 3
Performance evaluation of model predictions compared to real data under different tests.

From Eqs. (1)–(8) and (25), the resulting model of a TCHP is a set of ordinary differential equations in the following general form:

$$\dot{X} = F(X, U). \tag{32}$$

The resulting model is complex—high-order, non-linear, and stiff—making it impractical for efficient MPC predictions and optimizations. Consequently, a simpler, reduced model is required to act as an effective prediction model for MPC, where accuracy is critical for success. This need can be satisfied by nonlinear identification models, such as NNs like RNNs and LSTMs. The main output representing the objective of the TCHP application is the supply water temperature $T_{\text{w, supply}}$, so the reduced output vector is defined as $y = T_{\text{w, supply}}$, while the dominant reduced input vector is $u = [\varphi_{\text{HPV}}, \omega_{\text{bf}}]^\top$, and the disturbance vector $d = [\dot{m}_{\text{w}}, T_{\text{w, return}}]^\top$ (justified in section 2.2).

The real TCHP machine was excited with step changes on ω_{bf} (BF test) and φ_{HPV} (HPV test). Figures 7 and 8 show the model validation with real data when varying burner fan speed and high-pressure valve opening respectively. The model well captures the transient behavior of the main performance factors. The resulting MAPE or MAE and R^2 of p_{gc}, $T_{\text{w, supply}}$, and $\text{COP}_{\text{thermal}}$ of the resulting two tests are summarized in Table 3. With these results, the model can be considered validated and will be referred to as the 'reference model' throughout this paper.

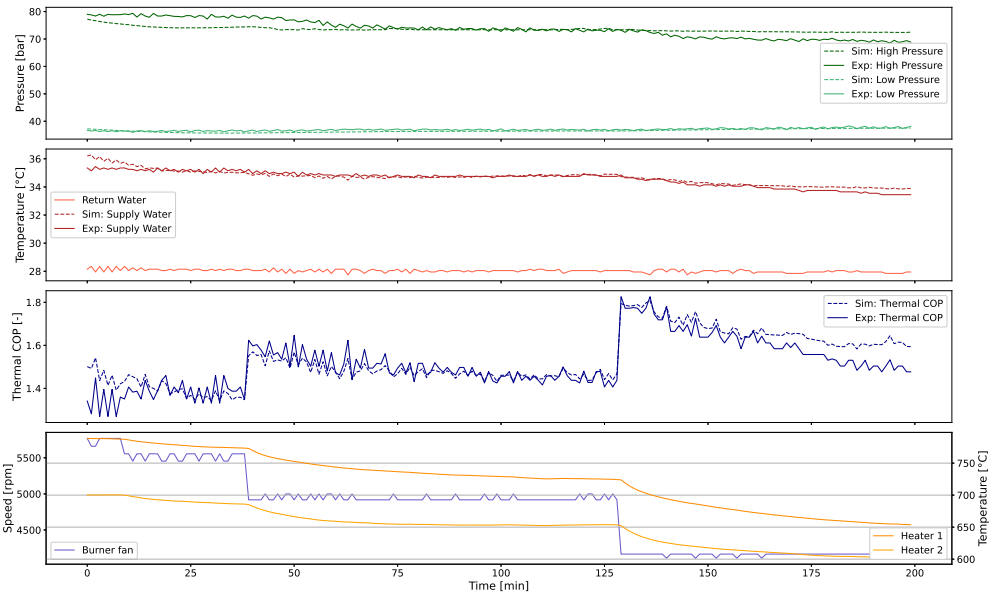

Figure 7: Dynamic model validation with real experimental data varying the burner fan speed (BF test).

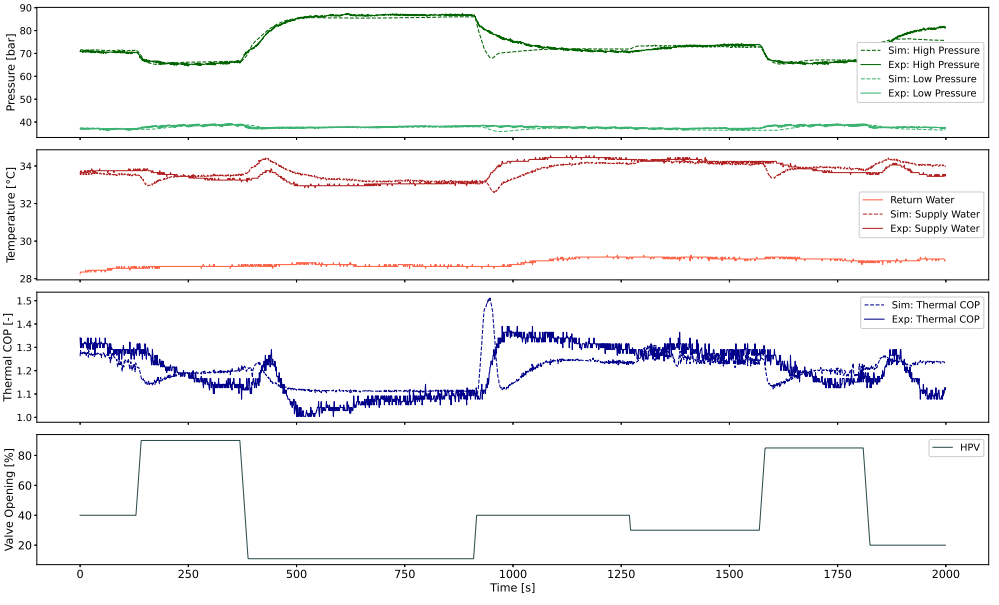

Figure 8: Dynamic model validation with real experimental data varying the high-pressure valve opening (HPV test).

4. Neural Network System Identification

The performance of MPC critically depends on the accuracy of the underlying prediction model. To construct this model, a system identification approach based on NNs is employed. While conventional feedforward NNs offer one-way connections between neurons, RNNs introduce feedback connections that enable information to persist across time steps, making them well-suited for time-series modeling (Figure 9).

In this work, we implement two reduced-order models for system identification: a vanilla RNN, originally proposed by Elman (1990), and a LSTM, introduced in Hochreiter and Schmidhuber (1997). Both models are designed to capture the dynamic evolution of supply water temperature, using data collected from exciting real system and reference model.

Table 4
Operating Ranges for Input Parameters and Outputs

Parameter	Description	Range
φ_{HPV}	High-pressure valve opening	[11, 100] %
ω_{bf}	Burner fan speed	[2000, 9500] rpm
$T_{\text{w, return}}$	return water temperature	[20, 45] °C
$\dot{m}_{\text{w}}$	Water mass flow rate	[0.13, 0.45] kg/s
$T_{\text{w, supply}}$	supply water temperature	[25, 55] °C

To ensure consistency, both models adopt an encoder-decoder architecture: the encoder processes a sequence of past measurements and control inputs, while the decoder recursively predicts future outputs over a specified prediction horizon, conditioned on known future control inputs. The training process for both models is described below.

The vanilla RNN and LSTM input vector $z(k) = [y(k-1), u(k), d(k)]^\top$ is constructed by concatenating measured output $y = T_{\text{w, supply}}$, reduced control input vector $u = [\omega_{\text{bf}}, \varphi_{\text{HPV}}]^\top$, and disturbance vector $d = [\dot{m}_{\text{w}}, T_{\text{w, return}}]^\top$. All input features are normalized using a standard scaler to ensure numerical stability during training. To enrich the available real data, the validated model is excited using PRBS and step changes with varying amplitudes of u and d, ensuring sufficient variability in both inputs and outputs while remaining within the operational bounds defined in Table 4. The dataset comprises 60,000 samples collected at one-second intervals, where around 30 % are from real experimental data and the rest are from reference model predictions. Each of the real and simulated datasets is chronologically split into training (80%), validation (10%), and testing (10%) sets to preserve temporal continuity. To enable multi-step prediction, the data is transformed into overlapping sequences using a sliding-window approach. Each training sequence consists of 10 past time steps (e.g. sequence length $l_{seq} = 10$ as input to the encoder, and the decoder predicts the next 10 time steps (e.g. prediction length $l_{pred} = 10$).

Both the vanilla RNN and LSTM models are implemented with one recurrent layer of 16 hidden units, which has been found sufficient, followed by a fully connected output layer. A dropout rate of 10% is applied between layers to reduce overfitting. Models are trained using the Adam optimizer with a learning rate of 0.001, and the mean squared error (MSE) is used as the loss function. Training was performed for a maximum of 200 epochs using mini-batches of size 32, with early stopping based on the validation loss. The best-performing model on the validation set was evaluated on the test set using both the MSE and R^2 metrics.

The vanilla RNN training took approximately 21 minutes, during which MSE validation loss decreased from 5.22 K^2 to 0.0871 K^2. Similarly, the LSTM training lasted 40 minutes, with MSE validation loss decreasing from 3.7 K^2 to 0.06 K^2, respectively. This shows an impressive ability of RNN type models to learn nonlinear complex dynamic models. Next, the different specifications of vanilla RNN and LSTM are explained.

4.1. Vanilla RNN System Identification

Since a RNN represents a discrete-time dynamical system, it is particularly well-suited for modeling sequential data where temporal dependencies play a critical role. In contrast to classical state-space models, where the state vector $x(k)$ has a direct physical interpretation, RNN maintains a latent internal state $h(k)$ that evolves recursively and captures the system memory. This hidden state implicitly encodes temporal correlations through a first-order Markov structure. A vanilla RNN used for output prediction is described by the following equations:

$$h(k) = \sigma(W_z z(k) + W_h h(k-1) + b_h), \tag{33}$$

$$\hat{y}(k) = W_y h(k) + b_y, \tag{34}$$

where $z(k)$ includes past control inputs and outputs, σ is a nonlinear activation function (e.g., tanh), and W_z, W_h, W_y, b_h, and b_y are matrices and vectors learned during training by minimizing the error between predicted outputs $\hat{y}(k)$ and measured output $y(k)$, as described earlier in this section.

As a prediction model for MPC, the vanilla RNN should be evaluated over multiple future steps. Figure 10 illustrates the model's performance on a test sequence of length 5050. The vanilla RNN is initialized five times using 10 previous samples, and then recursively predicts the next 500 time steps.

The supply water temperature $T_{\text{w, supply}}$ is predicted with a mean squared error (MSE) of 0.84 K^2 and an R^2 score of 0.975.

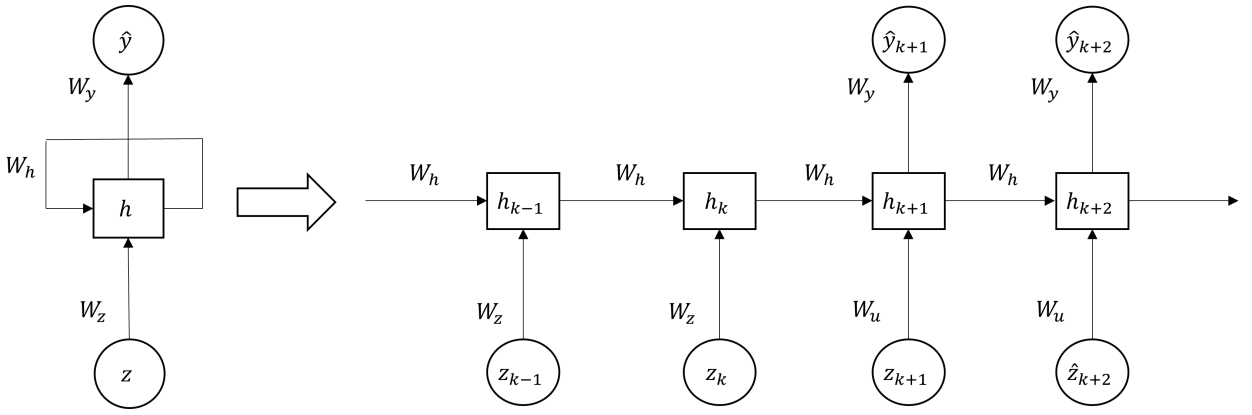

Figure 9: Structure of a vanilla RNN unrolled over time steps. At each time step k, the model receives an input vector $z(k) = [y(k-1), u(k), d(k)]^\top$ and updates its hidden state $h(k)$ via $h(k) = \sigma(W_z z(k) + W_h h(k-1) + b_h)$, where σ is a nonlinear activation function (e.g., *tanh*), and W_z, W_h, and b_h are learned parameters. The predicted outputs is computed as $\hat{y}(k) = W_y h(k) + b_y$. For clarity, biases are omitted from the illustration but are included in the model.

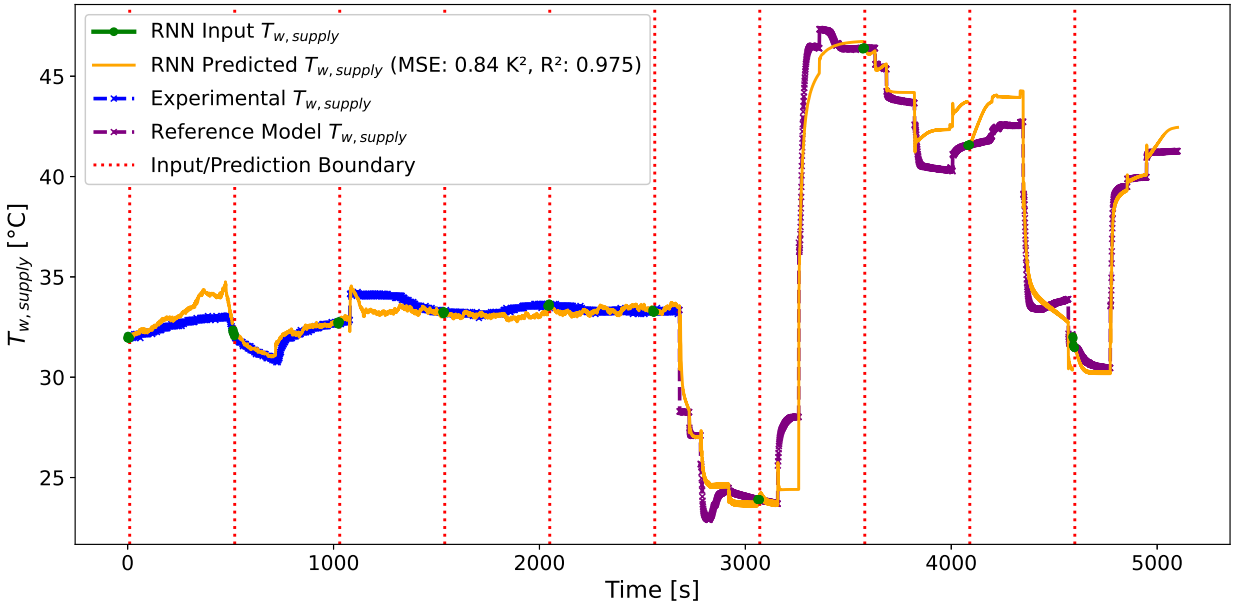

Figure 10: Validation of the vanilla RNN model on a sample from the test set. The model is initialized using 10 previous samples, followed by prediction of the next 500 samples. The transition between initialization and prediction is indicated by the red dashed line. The plot shows the predicted and true values of the supply water temperature $T_{\text{w, supply}}$. The RNN achieves a mean squared error of 0.84 K^2 and an R^2 score of 0.975, indicating high predictive accuracy over the forecasting horizon.

4.2. LSTM System Identification

LSTM is a type of RNN that incorporates internal gating mechanisms to address the vanishing gradient problem. Potentially, this makes LSTM more suitable for long input sequences and, consequently, for system identification tasks. LSTM networks extend vanilla RNNs by introducing a memory cell $c(k)$ that enables long-term information retention. The hidden state $h(k)$ serves as short-term memory, while a gating mechanism regulates the flow of information. Each LSTM cell includes an input gate $i(k)$, forget gate $r(k)$, candidate gate $\tilde{c}(k)$, and output gate $o(k)$. The update equations

are defined as follows:

$$
\begin{aligned}
i(k) &= \sigma(W_i z(k) + R_i h(k-1) + b_i),\\
r(k) &= \sigma(W_f z(k) + R_f h(k-1) + b_f),\\
\tilde{c}(k) &= \tanh(W_g z(k) + R_g h(k-1) + b_g),\\
o(k) &= \sigma(W_o z(k) + R_o h(k-1) + b_o),\\
c(k) &= f(k) \odot c(k-1) + i(k) \odot g(k),\\
h(k) &= o(k) \odot \tanh(c(k)),\\
\hat{y}(k) &= W_y h(k) + b_y.
\end{aligned}
\tag{35}
$$

Here, σ denotes the sigmoid activation function, tanh the hyperbolic tangent, and $\odot$ the Hadamard (element-wise) product. The LSTM weights are grouped as $\{W_i, W_f, W_g, W_o\}$ for the network input $z(k)$ and $\{R_i, R_f, R_g, R_o\}$ for the hidden state $h(k)$. The corresponding bias vectors are $\{b_i, b_f, b_g, b_o\}$. This formulation follows the standard LSTM structure. The implementation procedure follows the same setup described earlier in this section.

Figure 11 shows the validation of the LSTM prediction over the next 500 steps, using the same setup as for the vanilla RNN. The supply water temperature $T_{\text{w, supply}}$ is predicted with a mean squared error of 0.76 K^2 and an R^2 score of 0.977.

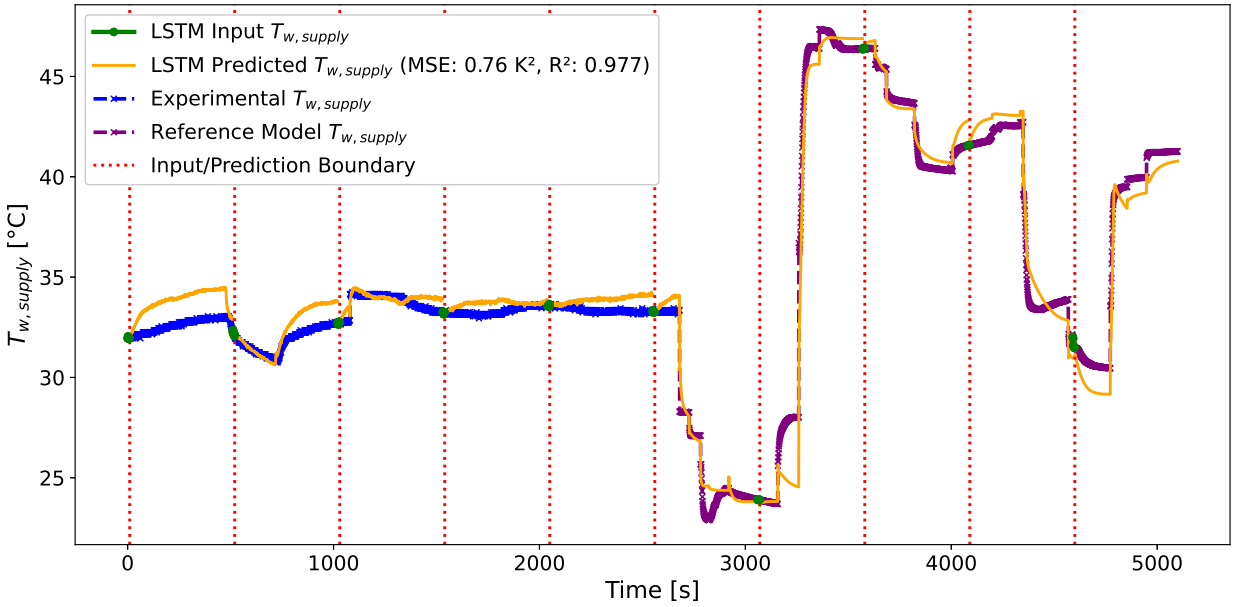

Figure 11: Validation of the LSTM model on a sample from the test set. The model is initialized using 10 previous samples, followed by prediction of the next 500 samples. The red dashed line separates the initialization phase from the prediction phase. The plot shows the predicted and true values of the supply water temperature $T_{\text{w, supply}}$. The LSTM achieves a mean squared error of 0.76 K^2 and an R^2 score of 0.977, indicating high prediction accuracy across time.

Both the vanilla RNN and LSTM models achieve satisfactory prediction accuracy and can be used as forecasting models in a MPC with a prediction horizon of 10 steps, for which the reduced-order models are expected to perform even better than over the 500-step horizon. LSTM performs slightly better, likely due to its enhanced internal structure.

5. MPC Design

The control strategies compared in this study are illustrated schematically in Figures 12 and 13. The baseline PID strategy relies on predefined correlations: g_1 for determining the heater temperature setpoint $T^{\text{sp}}_{\text{heater2}}$ as a function of

the supply water temperature setpoint $T^{\text{sp}}_{\text{w, supply}}$ and another function g_2 for setting the high-pressure target $p^{\text{sp}}_{\text{gc}}$ based on the return water temperature to maximize efficiency.

Model Predictive Control (MPC) is a model-based control strategy that has gained widespread adoption in industrial applications requiring optimal, constraint-aware decision-making. The performance of MPC depends primarily on four components: (1) an accurate prediction model (in this case, the previously derived vanilla RNN and LSTM models), (2) the formulation of the cost function, (3) constraint handling, and (4) the choice of optimization algorithm or solver. Since prediction model accuracy cannot be guaranteed over long horizons, incorporating feedback from the real system can help correct deviations and improve robustness Rawlings (2000).

In the MPC strategy, the desired setpoint temperature $\hat{y}^{\text{sp}} = T^{\text{sp}}_{\text{w, supply}}$ is directly fed to the MPC controller. The controller computes the optimal burner fan speed $u_1 = \omega_{\text{bf}}$ and the optimal high-pressure valve opening $u_2 = \varphi_{\text{HPV}}$ by minimizing the following objective function:

$$\begin{aligned}\min_{u_1,u_2} J = \alpha_1 \sum_{j=1}^{N_p} (\hat{y}(k+j|k) - y^{\text{sp}}(k+j|k))^2 + \sum_{j=0}^{N_c-1} \left(\alpha_2 (u_1(k+j|k))^2\right) \\ + \sum_{j=0}^{N_c-1} \left(\alpha_3 (\Delta u_1(k+j|k))^2 + \alpha_4 (\Delta u_2(k+j|k))^2\right) + \alpha_5 \sum_{j=1}^{N_p} L_{\text{constraint}}(T_{\text{heater2}}(k+j|k)),\end{aligned} \tag{36}$$

subject to:

$$\hat{y}(k+j|k) = f_{\text{RNN}}(y(k-l_{seq}), \dots, y(k), u(k-l_{seq}), \dots, u(k), d(k-l_{seq}), \dots, d(k)), \qquad \forall j = 1, \dots, N_p, \tag{37}$$

$$u_{\min} \le u(k+j|k) \le u_{\max}, \qquad \forall j = 0, \dots, N_c - 1. \tag{38}$$

In this formulation, the prediction model f_{RNN} corresponds to either the trained vanilla RNN or LSTM model derived in the previous section. The soft constraint term $L_{\text{constraint}}(T_{\text{heater2}}) = \max\left(0,\, T_{\text{heater2}} - T_{heater2,\max}\right)$ is introduced to penalize violations of the maximum allowable heater temperature, ensuring that safety constraints are respected.

The weight vector $\alpha = [\alpha_1, \alpha_2, \alpha_3, \alpha_4]$ adjusts the relative importance of each objective term:

- $\alpha_1 > 0$ penalizes the tracking error between the predicted supply water temperature and its setpoint.
- $\alpha_2 > 0$ discourages excessive use of burner fan speed ω_{bf}.
- $\alpha_3 > 0$ penalizes variations in burner fan speed $\Delta\omega_{\text{bf}}$.
- $\alpha_4 > 0$ penalizes large changes in high-pressure valve opening $\Delta\varphi_{\text{HPV}}$.
- $\alpha_5 > 0$ discourages exceeding $T_{heater2,max}$.

N_p and N_c denote the prediction and control horizons, respectively. The notation $k + j|k$ refers to the predicted value at time $k + j$ based on information available at time k. Equation (38) enforces hard constraints on the control inputs $u(k + j|k)$, while additional penalties on $\Delta u(k + j|k)$ help limit abrupt changes in control actions that could destabilize the system.

The MPC controller operates at each time step k as follows: let $u^{opt}(k+1) = \{u^{\text{opt}}_{k+1|k}, \dots, u^{\text{opt}}_{k+N_c|k}\}$ denote the optimal control sequence obtained by solving the objective function. Only the first element of this sequence, $u^{opt}_{k+1|k}$, is applied to the TCHP system. At the next time step $k + 1$, the system state is updated, and the optimization problem is solved again. This receding horizon approach enables continuous real-time optimization of control inputs.

The prediction horizon is set to match the capacity of the reduced-order models, i.e., $N_p = l_{\text{pred}} = 10$, while the control horizon is fixed at $N_c = 2$. The weighting factors in the objective function are selected as: $\alpha_1 = 0.95$, $\alpha_2 = 0.002$, $\alpha_3 = 0.02$, $\alpha_4 = 0.002$, and $\alpha_5 = 0.001$. MPC operates with a sampling time of 1 second. The method used for solving the objective function is 'Nelder-Mead' built in python 'minimize' function. In the next section, we validate the defined control strategies on the TCHP reference model.

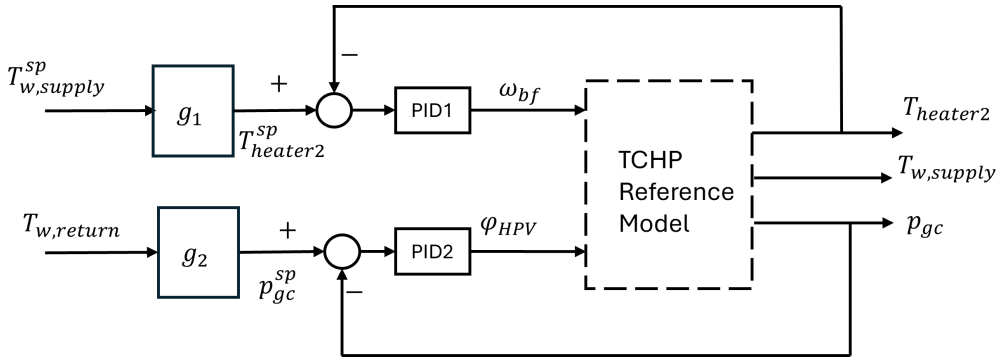

Figure 12: PID-based control strategy: The TCHP reference model is controlled using two parallel PID loops. The first loop adjusts the burner fan speed ω_{bf} to track $T^{\text{sp}}_{\text{heater2}}$ which is correlated to $T^{\text{sp}}_{\text{w, supply}}$ with a g_1 function. The second loop controls the high-pressure valve opening φ_{HPV} to regulate the system pressure p_{gc}, where the setpoint $p^{\text{sp}}_{\text{gc}}$ is determined via a mapping function g_2, based on the return water temperature $T_{\text{w, return}}$.

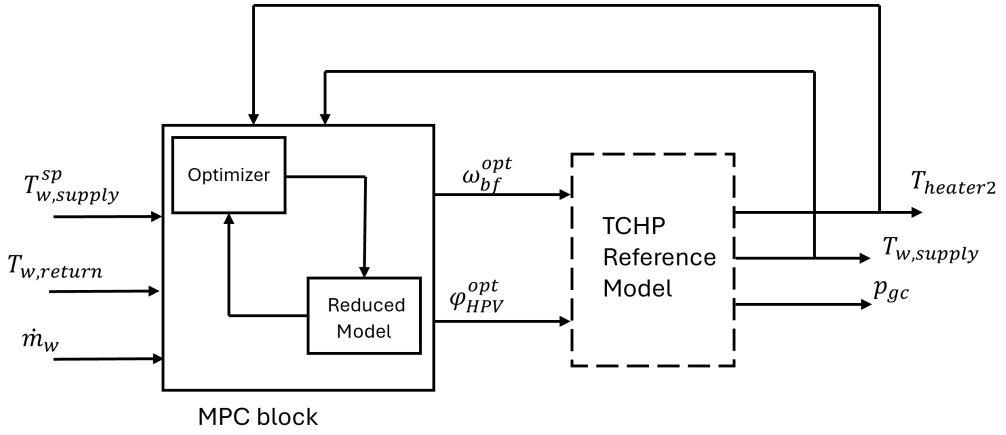

Figure 13: MPC-based strategy: This model predictive control framework relies on a learned model of the system based on vanilla RNN or LSTM as reduced models. The optimizer computes the optimal control actions $\varphi^{\text{opt}}_{\text{HPV}}$ and $\omega^{\text{opt}}_{\text{bf}}$ over a prediction horizon to track the $T^{\text{sp}}_{\text{w, supply}}$, given the current return water temperature $T_{\text{w, return}}$ and water mass flow rate $\dot{m}_{\text{w}}$. The optimizer uses predictions from the reduced model, which captures the system dynamics. The optimized inputs are applied to the TCHP reference model, and the resulting outputs $T_{\text{w, supply}}$ and T_{heater2} are fed back to the controller.

6. Results and Discussions

The main control objective is to achieve the desired supply water temperature setpoint $T^{\text{sp}}_{\text{w, supply}}$ with minimal gas consumption. In this section, the performance of the control strategies is validated on the TCHP reference model by simulating 5000 seconds, where $T^{\text{sp}}_{\text{w, supply}}$ is changed each 1000 second. As mentioned in section 2.2, the step changes in $T^{\text{sp}}_{\text{w, supply}}$ are induced by varying the outdoor temperature between 2 °C and 7 °C. Each increase in setpoint corresponds to an increase in heating energy demand and vice versa. The transient performance of the control strategies is monitored and a comparison of their rise time, settling time, overshoots, average temperature error, and average thermal COP is conducted. A key point is that disturbances are present, and the robustness of the control strategies is also being assessed. The inlet glycol mixture temperature is assumed to have a pinch of 5 K relative to the outdoor temperature. φ_{LPV} is regulated with a PID to maintain a superheat of 5 K.

The transient responses of the TCHP system under the control strategies are shown in Figures 14 and 15, corresponding to the outputs (p_{gc}, $T_{\text{w, supply}}$, and T_{heater2}) and inputs (φ_{HPV} and ω_{bf}), respectively. Both MPC-RNN and MPC-LSTM are able to follow the setpoint, with MPC-LSTM achieving a smaller error, likely due to its slightly better predictive accuracy. The PID strategy was also able to stay near the desired $T^{\text{sp}}_{\text{w, supply}}$, but exhibited significant overshoots, especially during increases in setpoint. This is mainly because ω_{bf} regulates T_{heater2} and not $T_{\text{w, supply}}$

Labels	Average COP [-]	Average Error [K]	Rise Time [s]	Settling Time [s]	Computation Time [s]
MPC-LSTM	1.82	0.09	5	30	7111
MPC-RNN	1.72	0.094	5	30	6580
PID	1.5	0.2	5	200	4082

Table 5
Performance metrics for PID, MPC-RNN, and MPC-LSTM strategies.

directly, meaning that good tracking of $T^{\text{sp}}_{\text{heater2}}$ does not necessarily translate to good tracking of $T^{\text{sp}}_{\text{w, supply}}$. This also results in a long settling time of 200 s. The limitation with the PID strategy lies not in the controller itself -as the corresponding setpoint are well followed- but in the control structure. In contrast, the MPC strategy directly connects ω_{bf} to $T_{\text{w, supply}}$ using a reduced model (RNN or LSTM), avoiding this issue. This reflects MPC's ability to act as both a high- and low-level controller. The TCHP system combines slow and fast dynamics—T_{heater2} responds more slowly than $T_{\text{w, supply}}$, primarily due to the FHX, which is directly influenced by ω_{bf}. This highlights the need for coordination between high- and low-level control, which this work directly addresses. Unlike other studies, we fully adopt a 1-second time step for both control levels. This ensures harmony between the two and avoids neglecting the internal dynamics.

The average thermal COP and temperature error of the three strategies during simulation are demonstrated with bar plots in Figure 16. Left figure shows a highest thermal COP by MPC-LSTM, directly followed by MPC-RNN, and a significantly lower value by PID. Right figure shows that the smallest average temperature error is achieved by MPC-LSTM, directly followed by MPC-RNN, and a significantly higher by PID.

Table 5 summarizes the performance metrics of the applied control strategies. Compared to the PID strategy ($\text{COP}_{\text{thermal}} = 1.57$), both MPC-RNN and MPC-LSTM show clear improvements, achieving higher average thermal COP values of 1.71 (+14.26 %) and 1.82 (+20.56%), respectively. Similarly, the average temperature error decreases from 0.20 °C under PID to 0.092 °C (-56%) with MPC-RNN and 0.09 °C (-58%) with MPC-LSTM. The average temperature errors from both MPC strategies are relatively close, with even a similar behavior on ω_b and T_{Theater2} as seen in Figures 15 and 14. However, the higher thermal COP when using MPC-LSTM shows that it is closer to the optimal high-pressure compared to MPC-RNN. For the PID strategy that defines fixed correlations (g_1 and g_2) leads to poor performance because (1) it is difficult to consider all influencing inputs within the correlation, (2) such correlations are derived under steady-state assumptions, neglecting transient behavior, and (3) system behavior may change over time. Additionally, the difference in the dynamic response between burner fan speed and high-pressure valve opening reveals the need for a controller capable of handling both variables simultaneously and in real time. The strength of MPC lies in its ability to regulate multiple inputs while satisfying system constraints and continuously optimizing performance.

The results demonstrate the effectiveness of the MPC-based approach in balancing energy efficiency and system demand. MPC's ability to anticipate future behavior enables earlier reactions to disturbances. The use of a prediction length equal to the prediction horizon avoided iterations in the optimization process, significantly reducing computation time. Nevertheless, computation time still increased from 4082 s (PID) to 6580 s (MPC-RNN) and 7111 s (MPC-LSTM) over the carried simulation. This cost could be reduced further by using approximate MPC, for instance, by generating optimal control inputs offline across multiple operating points and learning a control policy via a neural network. Another alternative is to linearize the trajectory online, enabling faster optimization using simpler solvers. For long-term evaluation, a building model is required. Although a 1-second time step is coherent for internal control, it can become a computational burden in longer simulations.

Despite their effectiveness, a limitation of both MPC-RNN and MPC-LSTM is the relatively high number of hyperparameters that must be tuned to ensure stability and good performance. Finally, the proposed control strategy is also applicable to electrically driven heat pumps by replacing the burner fan speed with electric motor speed as the manipulated variable.

7. Conclusion

A novel advanced control strategy based on model predictive control (MPC) was proposed for a CO_2 multi-stage thermal compressor heat pump (TCHP) application for residential heating. At first, the process physics and control structure are explicitly described. A physics-based dynamic model was derived with coupling of energy and mass balances equations on a finite volume discrete space combined with algebraic representations of thermal compressors

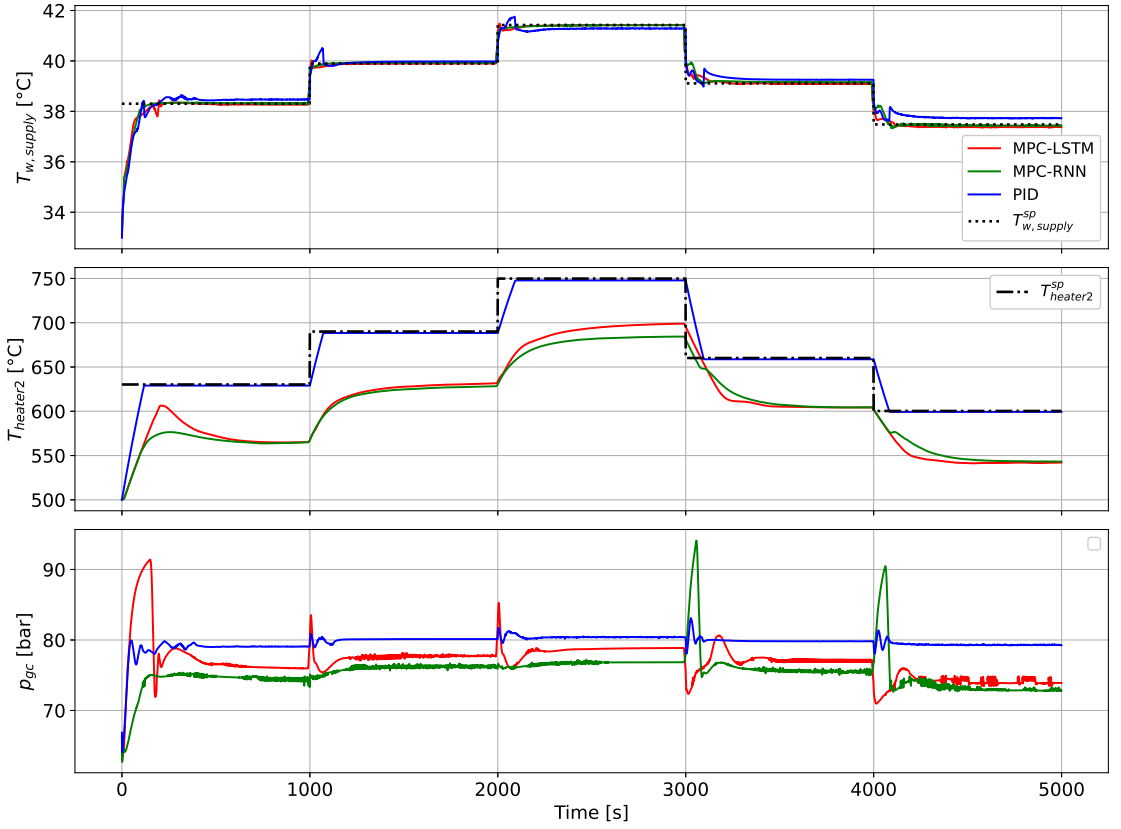

Figure 14: The resulting supply water temperature $T_{\text{w, supply}}$, second heater temperature T_{heater2}, high-pressure p_{gc} values, and their corresponding setpoints when applying control strategies PID, MPC-RNN, MPC-LSTM.

and expansion valves. The transient behavior of the dynamic model is validated with real data at most significant inputs and outputs. The model becomes a reference model throughout the paper. Reference model is excited to generate data and is used with the real experimental data to derive reduced models of vanilla recurrent neural network (RNN) and long short term memory (LSTM) nature, which are able to mimic the dynamic response of the TCHP with high validity. These reduced models are then coupled with a MPC controller to serve as predictions models resulting in MPC-RNN and MPC-LSTM strategies that perform online-optimizations. These strategies are tested on the reference model and are compared with the previous one based on PID that perform on offline correlations. Compared to PID strategy, MPC-LSTM strategy increases the thermal COP by 20.56 % and reduces the error by 58.69 %. MPC-RNN strategy increases the thermal COP by 14.26 % and reduces the error by 56.59 %. This work is also a call for the necessity of coupling high- and low- levels controls with a common time step of 1s. A MPC with an accurate nonlinear reduced model makes this possible.

MPC demonstrates high effectiveness in multi-input control of heat pump cycles, which are known to be nonlinear and strongly coupled. It allows better setpoint tracking and dynamic performance compared to PID by optimizing control inputs at each step and handling system constraints. However, a disadvantage is the high number of hyperparameters introduced when coupling RNNs with MPC, which requires careful tuning and can affect robustness. Although the computation cost was not significantly higher than that of PID, further reduction is still possible to enable more efficient deployment in real-time applications. For example, an approximate MPC from learning the optimal control on the whole range of operation can potentially reduce computational costs. Further work could also include low pressure valve opening to the reduced models to increase their degrees of freedom. A lot can further be done, but what is sure is that more intelligent operations are more and more needed in industries today as they can potentially increase performance and respect added restrictions.

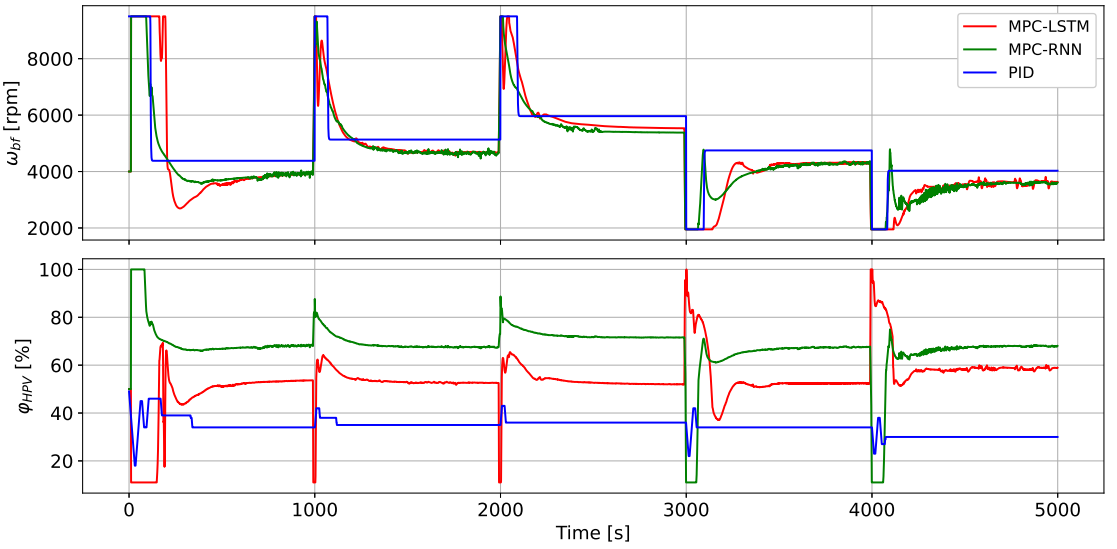

Figure 15: The resulting control inputs of burner fan speed ω_{bf} and high-pressure valve opening in percentage φ_{HPV} from control strategies PID, MPC-RNN, and MPC-LSTM.

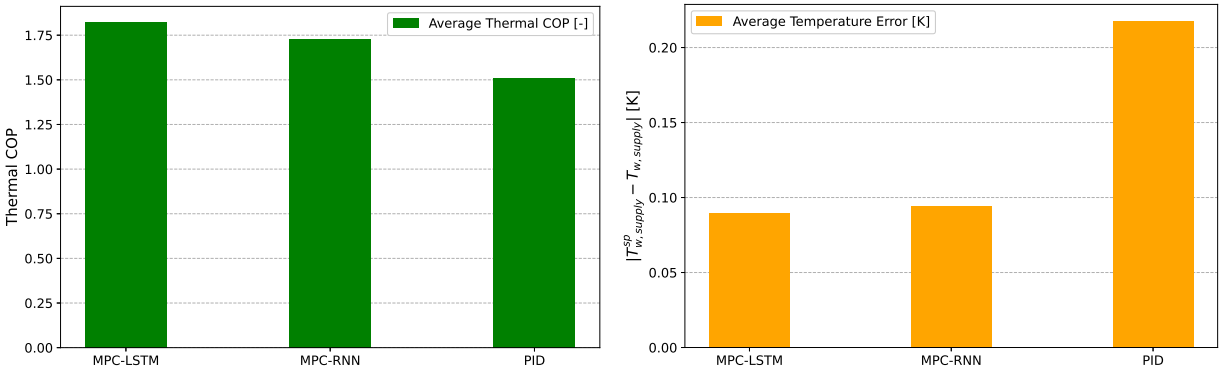

Figure 16: The average thermal COP and temperature error of the three control strategies during simulation are demonstrated with bar plots in left and right figures respectively.

References

Afram, A., Janabi-Sharifi, F., Fung, A.S., Raahemifar, K., 2017. Artificial neural network (ann) based model predictive control (mpc) and optimization of hvac systems: A state of the art review and case study of a residential hvac system. Energy and Buildings 141, 96–113. URL: https://www.sciencedirect.com/science/article/pii/S0378778816310799, doi:https://doi.org/10.1016/j.enbuild.2017.02.012.

Bell, I., Wronski, J., Quoilin, S., Lemort, V., 2014. Pure and pseudo-pure fluid thermophysical property evaluation and the open-source thermophysical property library coolprop. Industrial & Engineering Chemistry Research 53, 2498–2508.

Cecchinato, L., Corradi, M., Cosi, G., Minetto, S., Rampazzo, M., 2012. A real-time algorithm for the determination of r744 systems optimal high pressure. International Journal of Refrigeration 35, 817–826. URL: https://www.sciencedirect.com/science/article/pii/S0140700712000060, doi:https://doi.org/10.1016/j.ijrefrig.2012.01.005.

Elman, J.L., 1990. Finding structure in time. Cognitive Science 14, 179–211. URL: https://www.sciencedirect.com/science/article/pii/036402139090002E, doi:https://doi.org/10.1016/0364-0213(90)90002-E.

Eurostat, 2024. Energy use in eu households in 2022 lowest since 2016. URL: https://ec.europa.eu/eurostat/en/web/products-eurostat-news/w/ddn-20240605-2. accessed: Jun. 05, 2024.

Evens, M., Arteconi, A., 2025. Heat pump digital twin: An accurate neural network model for heat pump behaviour prediction. Applied Energy 378, 124816. URL: https://www.sciencedirect.com/science/article/pii/S0306261924021998, doi:https://doi.org/10.1016/j.apenergy.2024.124816.

Fallahsohi, H., 2023. Energy analysis of the heat pump operating with supercritical co_2 driven by thermal compression. American Journal of Engineering Research (AJER) 12, 61–74. doi:https://www.ajer.org/papers/Vol-12-issue-11/12116174.pdf.

Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. Neural Comput. 9, 1735–1780. URL: https://doi.org/10.1162/neco.1997.9.8.1735, doi:10.1162/neco.1997.9.8.1735.

Hu, B., Li, Y., Cao, F., Xing, Z., 2015. Extremum seeking control of cop optimization for air-source transcritical co2 heat pump water heater system. Applied Energy 147, 361–372. URL: https://www.sciencedirect.com/science/article/pii/S0306261915002949, doi:https://doi.org/10.1016/j.apenergy.2015.03.010.

Jagannathan, S., Lewis, F., 1996. Identification of nonlinear dynamical systems using multilayered neural networks. Automatica 32, 1707–1712. URL: https://www.sciencedirect.com/science/article/pii/S0005109896800070, doi:https://doi.org/10.1016/S0005-1098(96)80007-0.

Jung, M., da Costa Mendes, P.R., Önnheim, M., Gustavsson, E., 2023. Model predictive control when utilizing lstm as dynamic models. Engineering Applications of Artificial Intelligence 123, 106226. URL: https://www.sciencedirect.com/science/article/pii/S0952197623004104, doi:https://doi.org/10.1016/j.engappai.2023.106226.

Lanzetti, N., Lian, Y.Z., Cortinovis, A., Dominguez, L., Mercangöz, M., Jones, C., 2019. Recurrent neural network based mpc for process industries, in: 2019 18th European Control Conference (ECC), pp. 1005–1010. doi:10.23919/ECC.2019.8795809.

Liao, S., Zhao, T., Jakobsen, A., 2000. A correlation of optimal heat rejection pressures in transcritical carbon dioxide cycles. Applied Thermal Engineering 20, 831–841. URL: https://www.sciencedirect.com/science/article/pii/S1359431199000708, doi:https://doi.org/10.1016/S1359-4311(99)00070-8.

Ljung, L., 1999. System identification (2nd ed.): theory for the user. Prentice Hall PTR, USA.

Nekså, P., 2002. Co2 heat pump systems. International Journal of Refrigeration 25, 421–427. URL: https://www.sciencedirect.com/science/article/pii/S0140700701000330, doi:https://doi.org/10.1016/S0140-7007(01)00033-0.

Pillonetto, G., Aravkin, A., Gedon, D., Ljung, L., Ribeiro, A.H., Schön, T.B., 2025. Deep networks for system identification: A survey. Automatica 171, 111907. URL: https://www.sciencedirect.com/science/article/pii/S0005109824004011, doi:https://doi.org/10.1016/j.automatica.2024.111907.

Qi, P.C., He, Y.L., Wang, X.L., Meng, X.Z., 2013. Experimental investigation of the optimal heat rejection pressure for a transcritical co2 heat pump water heater. Applied Thermal Engineering 56, 120–125. URL: https://www.sciencedirect.com/science/article/pii/S1359431113002238, doi:https://doi.org/10.1016/j.applthermaleng.2013.03.045.

Qiao, H., Aute, V., Radermacher, R., 2015. Transient modeling of a flash tank vapor injection heat pump system – part i: Model development. International Journal of Refrigeration 49, 169–182. URL: https://www.sciencedirect.com/science/article/pii/S0140700714001820, doi:https://doi.org/10.1016/j.ijrefrig.2014.06.019.

Rampazzo, M., Cervato, A., Corazzol, C., Mattiello, L., Beghi, A., Cecchinato, L., Virzi, A., 2019. Energy-efficient operation of transcritical and subcritical co2 inverse cycles via extremum seeking control. Journal of Process Control 81, 87–97. URL: https://www.sciencedirect.com/science/article/pii/S0959152419304226, doi:https://doi.org/10.1016/j.jprocont.2019.06.008.

Rawlings, J., 2000. Tutorial overview of model predictive control. IEEE Control Systems Magazine 20, 38–52. doi:10.1109/37.845037.

Ren, Y.M., Alhajeri, M.S., Luo, J., Chen, S., Abdullah, F., Wu, Z., Christofides, P.D., 2022. A tutorial review of neural network modeling approaches for model predictive control. Computers & Chemical Engineering 165, 107956. URL: https://www.sciencedirect.com/science/article/pii/S0098135422002927, doi:https://doi.org/10.1016/j.compchemeng.2022.107956.

Salame, A., Lemort, V., Dufour, P., Nadri, M., 2024. Improving performance of a co_2 thermally driven heat pump based on a new optimal high-pressure correlation, in: Proceedings of Carnot24 Liège Symposium.

Salazar, M., Méndez, F., 2014. Pid control for a single-stage transcritical co2 refrigeration cycle. Applied Thermal Engineering 67, 429–438. URL: https://www.sciencedirect.com/science/article/pii/S1359431114002312, doi:https://doi.org/10.1016/j.applthermaleng.2014.03.052.

Sarkar, J., Bhattacharyya, S., Gopal, M.R., 2004. Optimization of a transcritical co_2 heat pump cycle for simultaneous cooling and heating applications. International Journal of Refrigeration 27, 830–838. URL: https://www.sciencedirect.com/science/article/pii/S0140700704000520, doi:10.1016/j.ijrefrig.2004.03.006.

Turgut, M.S., Çoban, M.T., 2020. Neural network predictive control of a vapor compression cycle. Arabian Journal for Science and Engineering 45, 779–796. URL: https://doi.org/10.1007/s13369-019-04149-2, doi:10.1007/s13369-019-04149-2.

Wang, S., Tuo, H., Cao, F., Xing, Z., 2013. Experimental investigation on air-source transcritical co2 heat pump water heater system at a fixed water inlet temperature. International Journal of Refrigeration 36, 701–716. URL: https://www.sciencedirect.com/science/article/pii/S014070071200271X, doi:https://doi.org/10.1016/j.ijrefrig.2012.10.011.

Wang, W., Zhao, Z., Zhou, Q., Qiao, Y., Cao, F., 2021. Model predictive control for the operation of a transcritical co2 air source heat pump water heater. Applied Energy 300, 117339. URL: https://www.sciencedirect.com/science/article/pii/S0306261921007479, doi:https://doi.org/10.1016/j.apenergy.2021.117339.

Wang, Y., 2017. A new concept using lstm neural networks for dynamic system identification, in: 2017 American Control Conference (ACC), IEEE, Seattle, WA, USA. pp. 5324–5329. doi:10.23919/ACC.2017.7963782.

Yoon, S.H., Kim, J.H., Hwang, Y.W., Kim, M.S., Min, K., Kim, Y., 2003. Heat transfer and pressure drop characteristics during the in-tube cooling process of carbon dioxide in the supercritical region. International Journal of Refrigeration 26, 857–864. URL: https://www.sciencedirect.com/science/article/pii/S0140700703000963, doi:https://doi.org/10.1016/S0140-7007(03)00096-3.

Zhang, W.J., Zhang, C.L., 2011. A correlation-free on-line optimal control method of heat rejection pressures in co2 transcritical systems. International Journal of Refrigeration 34, 844–850. URL: https://www.sciencedirect.com/science/article/pii/S0140700711000272, doi:https://doi.org/10.1016/j.ijrefrig.2011.01.014.

Ławryńczuk, M., Zarzycki, K., 2025. Lstm and gru type recurrent neural networks in model predictive control: A review. Neurocomputing 632, 129712. URL: https://www.sciencedirect.com/science/article/pii/S0925231225003844, doi:https://doi.org/10.1016/j.

Table 6
Symbols

Symbol	Description	Unit
p	Pressure	bar
h	Specific enthalpy	J/kg
T	Temperature	°C
ρ	Density	kg/m^3
m	Mass	kg
A	Area	m^2
V	Volume	m^3
$\dot{m}$	Mass flow rate	kg/s
ω	Speed	rpm
U	Heat transfer coefficient	W/(m^2.K)
$\dot{Q}$	Heat exchange rate	W
$\left(\frac{\partial \rho}{\partial h}\right)_p$	Density derivative w.r.t. enthalpy at constant pressure	kg^2/(J.m^3)
$\left(\frac{\partial \rho}{\partial p}\right)_h$	Density derivative w.r.t. pressure at constant enthalpy	s^2/m^2
c_p	Specific heat capacity	J/kg·K
C_v	Valve flow coefficient	-
φ	Valve opening in percentage	%
N_p	Prediction horizon for MPC	-
N_c	Control horizon of MPC	-
l_{seq}	Sequence length for RNN model	-
l_{pred}	Prediction length for RNN model	-
Nu	Nusselt number	-
Pr	Prandtl number	-
q	Quality vapor	-

neucom.2025.129712.

A. Regression Models

The derived linear regression (LR) models for the mass flow rate, discharge temperature and cooler heat exchange rate on the first TC are respectively:

$$\dot{m}_{TC1}[\text{kg/s}] = -0.0382p_{r1} + 0.0056T_{r1} + 0.00014\omega_{\text{m1}}[\text{rpm}] + 0.0135, \tag{39}$$
$$T_{TC1,dis}[^\circ\text{C}] = 75.1p_{r1} + 15.52T_{r1} - 0.27\omega_{\text{m1}}[\text{rpm}] - 40.4, \tag{40}$$
$$\dot{Q}_{\text{cooler1}}[W] = 2270p_{r1} + 96.1T_{r1} - 0.35\omega_{\text{m1}}[\text{rpm}] - 2335. \tag{41}$$

While on the second two TCs are:

$$\dot{m}_{TC23}[\text{kg/s}] = -0.025p_{r1} + 0.0162T_{r1} + 0.0000097\omega_{\text{m1}}[\text{rpm}] + 0.00094, \tag{42}$$
$$T_{TC23,dis}[^\circ\text{C}] = 20p_{r1} + 11.45T_{r1} - 0.058\omega_{\text{m1}}[\text{rpm}] - 5.75, \tag{43}$$
$$\dot{Q}_{\text{cooler23}}[W] = 1101.1p_{r1} + 331T_{r1} + 7.42\omega_{\text{m1}}[\text{rpm}] - 2128.1. \tag{44}$$

B. Thermal Compressor Heater

Model accuracy is assessed using the following metrics:

- **Mean Absolute Percentage Error (MAPE):**

$$\text{MAPE} = \frac{100}{N}\sum_{k=1}^{N}\left|\frac{X(k+1) - \hat{X}(k+1)}{X(k+1)}\right|$$

.

Table 7
Abbreviations

Abbreviation	Description	Unit
CO_2	Carbon dioxide	-
CH_4	Methane	-
MPG	Mono-propylene glycol	-
VCC	Vapor Compression Cycle	-
ASHP	Air-Source Heat Pump	-
TCHP	Thermal Compressor Heat Pump	-
TDHP	Thermally Driven Heat Pump	-
TC	Thermal Compressor	-
GC	Gas Cooler	-
EVAP	Evaporator	-
IHX	Internal Heat Exchanger	-
FT	Flash Tank	-
EEV	Electronic Expansion Valve	-
HPV	High-Pressure Valve	-
LPV	Low-Pressure Valve	-
FHX	Fume Heat Exchanger	-
PID	Proportional-Integral-Derivative Controller	-
MPC	Model Predictive Control	-
NN	Neural Network	-
RNN	Recurrent Neural Network	-
LSTM	Long Short Term Memory	-
LHV	Lower heating value of methane	J/kg
COP	Coefficient of performance	-

Table 8
Subscripts and Ratios

Symbol	Description	Unit
p_r	Pressure ratio	-
T_r	Temperature ratio	-
w	Water	-
f	CO_2 fluid	-
m	Motor	-
sf	Secondary fluid	-
b	Burner fan	-
α	Weighting factor in MPC objective function	-
β	Parameters of T_{heater} model	-
Δ	Difference	-
τ	Time constant	s

- **Coefficient of Determination (R^2):**

$$R^2 = 1 - \frac{\sum_{k=1}^{N} \left(X(k+1) - \hat{X}(k+1)\right)^2}{\sum_{k=1}^{N} \left(X(k+1) - \bar{X}\right)^2},$$

where:

- $X(k+1)$ is the true value at time step $k+1$,
- $\hat{X}(k+1)$ is the predicted value,
- $\bar{X}$ is the mean of the true values over all time steps.

C. Fume Heat Exchanger

The correlation defined for the heat transfer coefficient on the FHX as function of water mass flow rate, methane gas flow rate and entering water temperature to the FHX is as follows:

$$AU_{\text{fhx}}[W/K] = 5.98 \times 10^{-4}\omega_{\text{bf}}[\text{rpm}] + 0.687T_{fhx,w,in}[^{\circ}\text{C}] - 16.34 \tag{45}$$