

LEVERAGING NEURAL NETWORKS IN A HYBRID MODEL PREDICTIVE CONTROL FRAMEWORK FOR DISTRICT HEATING NETWORKS

Dubon Rodrigue^{1,2,*}, Mohamed Tahar Mabrouk¹, Bastien Padeloup²,
Patrick Meyer², Bruno Lacarrière¹

¹IMT Atlantique, GEPEA, UMR CNRS 6144, Nantes, France

²IMT Atlantique, Lab-STICC, UMR CNRS 6285, Brest, France

*Corresponding Author: dubon.rodrigue@imt-atlantique.fr

ABSTRACT

District heating networks are urban-scale heat energy supply systems. These networks offer a flexible manner to supply heat energy by leveraging the thermal inertia of the network pipes and the presence of distributed heat sources. To exploit such flexibility and reach energy supply efficiency, optimal operational control strategies of the networks are essential. However, existing models developed for an optimal control strategy are hindered by their computational costs, rendering them impractical for large networks. The computational burden stems mostly from the iterative simulation of the networks required for each optimization resolution step, until optimal control parameters are attained. Indeed, each simulation of the network has computational costs proportional to the size of the networks. In this study, we propose a hybrid Model Predictive Control (MPC), a control framework using a hybrid model to simulate the network and optimize the mass flow rates and supply temperatures injected by the heat sources to fulfill substation energy demands with the lowest surplus and deficits of delivered heat. The proposed approach substitutes clusters of consumer substations with artificial Neural Networks (NN) models, trained to replicate real-time consumption patterns and leaving waters' temperatures evolution of the replaced clusters. This hybrid approach reduces computational costs while maintaining prediction errors below 0.52 %. Results demonstrate that replacing one-third of the network leads to a reduction of 9% of the computational time required by a physics-based MPC.

1 INTRODUCTION

District heating networks (DHNs) are an important component of urban energy infrastructure, offering a sustainable solution for heat energy distribution. As the emphasis on energy efficiency and greenhouse gas reductions grows, DHNs have seen increased interest in recent years. The recent trend involves integrating distributed heat sources with more renewable sources (Guelpa *et al.*, 2019), and lower generation temperatures (Buffa *et al.*, 2019). However, to achieve optimal energy efficiency in these complex systems, effective control on the heat sources is crucial (Wirtz *et al.*, 2021).

Model predictive control (MPC) has proven to be a very effective control method for DHNs due to its ability to combine system forecasting, anticipation and optimization (Jansen *et al.*, 2023). It operates by searching for the optimal control sequence over a predefined and finite horizon window. At the end of each optimization, only the first variables of the control sequence are applied, the horizon window is shifted forward, and the optimization process restarts over on the new horizon. This "rolling horizon" (Marquant *et al.*, 2015) approach allows the controller to adapt to changing conditions and uncertainties, making it well-suited for the dynamic nature of DHNs. MPC provides flexibility in defining its optimization functions, allowing for both linear (Zimmerman *et al.*, 2019) and non-linear (Sandou *et al.*, 2006) formulations depending on specific operational objectives. Notably, incorporating dynamic-based simulations of the network, as presented in Giraud *et al.* (2017) within the MPC framework

significantly enhances its accuracy. Dynamic simulations allow to capture subtle water temperatures changes and to leverage the thermal inertia of the pipes (Lund *et al.*, 2021). However, these dynamic models come at the cost of high computational burden, especially for large networks (Benonysson *et al.*, 1995). To address this challenge, Bavière and Vallée (2018) proposed a simplified approach that assumes that substation temperatures are weighted combinations of source unit temperatures. While it reduced computational burden, it relied on assumptions about mass flow rates that limit its applicability and compromise the accuracy compared to fully dynamic simulations (Jansen *et al.*, 2023).

On the other hand, data-driven models, such as machine learning (ML), have shown remarkable success in learning from data and replacing complex physical models in various applications. Integrating ML models into MPC to substitute computationally expensive physical models is a hybrid MPC approach (Zhang *et al.*, 2023) which reduces significantly computational costs. Recent related works include building indoor temperature optimal control by Zhang *et al.* (2023) and Ma *et al.* (2022). Ma *et al.* (2022) used a recurrent neural network with Long Short-Term Memory (LSTM) units (Hochreiter & Schmidhuber, 1997) to predict indoor temperatures based on forecasts and past states, combined with a particle-swarm optimization (PSO) algorithm (Kennedy & Eberhart, 1995) to control the frequency of a pump at the thermal inlet of a building to meet heating demands. Similarly, Zhang *et al.* (2023) employed LSTM-based networks for indoor temperature prediction in a building connected to a DHN substation. LSTM and Gated Recurrent Units (GRU) (Chung *et al.*, 2014) are designed to effectively capture long-term dependencies, making them ideal for systems with high inertia. While offering similar performance, GRU boast lower complexity and computational requirement compared to LSTM (Cahuantzi *et al.*, 2023).

Most hybrid MPC models aim to entirely replace physical modeling of energy systems, offering significant computational and operational benefits (Saloux *et al.*, 2023). The authors proposed black-box ML models to predict aggregated heating demand and boiler performance curves. However, this entire-system approach faces limitations due to data exhaustiveness. As highlighted by Stoffel *et al.* (2023), these "numerical twins" are only valid within the operational scenarios used for training, leading to unreliable predictions outside that domain. Substituting the full DHN dynamic simulation with an ML model would require vast, exhaustive data incorporating all possible scenarios, an impractical approach due to the exponential growth of possible scenarios.

Focusing on cluster-level substations presents a more efficient alternative with lower data requirements. Focusing the learning on the cluster's local physical dynamics for a wider range of temperatures and flow rates, regardless of network-wide scenarios and control strategies, this approach presents potentially better computational efficiency and lower data requirement. The utilization of clusters within DHNs has been explored for various applications, including optimal heat sources placement (Marquant *et al.*, 2018) and network topology reduction (Kane and Rolle, 2020). However, to our knowledge, the application of machine learning (ML) models to replace DHN substation clusters within a model predictive control (MPC) framework remains unexplored. This paper addresses this gap by proposing a hybrid MPC approach that incorporates trained ML models replicating the dynamic behavior and heat consumptions of individual clusters. The proposed framework aims to balance accuracy and computational efficiency. The rest of this paper is structured as follows: the methodology section introduces the proposed hybrid MPC approach, the results section evaluates its accuracy and computational cost reduction compared to a baseline MPC using a physical dynamic simulation model, and the conclusion section summarizes the key findings.

2 METHODOLOGY

2.1 Dynamic simulation model of the DHN

The dynamic simulation model of the District Heating Network (DHN) employed in this study leverages the node method introduced by Benonysson *et al.* (1995). This method represents the DHN as a directed graph, with nodes representing substations categorized as consumers (with heating demand powers) and sources (with heat generation powers). Heating powers injected by the sources are controlled by varying

the mass flow rates and temperatures of hot water injected into the network. Edges depict two identical transport pipes (supply and return) of the network, oriented arbitrarily in the supply flow direction. Overall, the modeling of the DHN is composed of both hydraulic and thermal models. Hydraulic model contains mass conservation at the substation nodes at every time step t (Equation 1) and pressure drops inside the pipes modeled by the Darcy-Weisbach equation (Equation 2) where ΔH_e indicates the head pressure loss through the pipes of the edge e , g_{rv} the gravity, u_e the water flows' velocity, d_e the diameter, l_e the length and f_e the friction factor of the pipes. This friction factor f_e is a function of the flows' Reynolds number and the pipes' effective roughness, further detailed in the work of Betancourt Schwarz *et al.* (2019). Equation 1 ensures mass balance at every substation node on both supply and return side. \dot{m}_s denotes the mass flow rates extracted by sources from the return pipe, heated and re-injected into the supply pipe at hot generation supply temperatures T_{ss} . Similarly, \dot{m}_c depicts the mass flow rates of supply water extracted by the consumers from the supply pipe to respond to heating demands (Equation 3) and re-injected into the mass flow rates at fixed and constant return temperature T_{rc} . The thermal model governs heat consumptions, enthalpy balances and heat transport (Equations 3-6). Equation 3 governs the consumption at consumer nodes and gives the effective heat power consumed to meet heating demands D^t at each time step. Deviations between heating demands and real heating power consumptions P_c may occur due to the delay time between the moment heat is generated at the sources and the moment that it is available to consumers, leading to consumptions surplus ($P_c > D$) or deficits ($P_c < D$). Thermal losses through the pipes may also contribute to this deviation. Enthalpy balance enforced at nodes for both return and supply pipes (Equations 4 and 5) translates to a weighted sum of incoming water temperatures. This modelling approach aligns with common approach in the literature, as detailed in Fang and Lahdelma (2015) work. Our model shares similarities with theirs, except for pipe modeling. Here, we employ the partial derivative heat transport equation (Equation 6) to model water temperature evolution through the pipes. This approach, based on the works of Betancourt Schwarz *et al.* (2019) and Giraud *et al.* (2015), utilizes the finite volume model and upwind schemes for spatial and temporal discretization, respectively.

$$\text{(mass balance)} \quad \forall(i, t), \dot{m}_{s_i}^t + \sum_{\{e \in E_i^+\}} \dot{m}_e^t = \dot{m}_{c_i} + \sum_{\{i \in E_i^-\}} \dot{m}_e^t \quad (1)$$

$$\text{(head loss)} \quad \forall e, \Delta H_e = \frac{4f_e l_e u_e^2}{2g_{rv} d_e} \quad (2)$$

$$\text{(consumption)} \quad \forall(i \in C), P_{c_i}^t = \dot{m}_{c_i}^t c_p (T_{s_i} - T_{rc_i}) \quad (3)$$

$$\text{(supply balance)} \quad \forall(i, t), \dot{m}_{s_i}^t T_{ss_i}^t + \sum_{\{e \in E_i^+\}} \dot{m}_e^t T_{s_i}^t = \left(\dot{m}_{c_i} + \sum_{\{e \in E_i^-\}} \dot{m}_e^t \right) T_{s_i}^t \quad (4)$$

$$\text{(return balance)} \quad \forall(i, t), \dot{m}_{c_i}^t T_{rc_i}^t + \sum_{\{e \in E_i^-\}} \dot{m}_e^t T_{r_i}^t = \left(\dot{m}_{s_i} + \sum_{\{e \in E_i^+\}} \dot{m}_e^t \right) T_{r_i}^t \quad (5)$$

$$\text{(heat transport)} \quad \forall e, \frac{\partial T_e}{\partial t} = -\frac{4\dot{m}_e}{\pi \rho d_e^2} \frac{\partial T_e}{\partial x} - \frac{4h_e}{\rho c_p d_e} (T_e - T_{ground}) \quad (6)$$

2.2 Model Predictive Control of the DHN

Building on the DHN model described above, Model Predicted Control (MPC) aims to find the optimal sequence of supplied power by the sources. This sequence is determined by optimizing the injected supply mass flow rates \dot{m}_s^t and generation supply temperatures T_{ss}^t , across all time steps t . The objective is to minimize total heat production while satisfying consumer heating demands. We achieve this by minimizing the difference between the total consumed energy (integrated consumed heat powers P_c over the optimization horizon) and the total demanded heat energy. At each node, there may be either surpluses or deficits of heating power consumed compared to heating demands. In fact, both values reflect differences between the consumed and demanded heat by the consumers and differ only by their

sign. Consider δ_i^t the difference between heating consumption power and demanded power by the consumer i at time t . Integrating surplus ($\delta_i^t > 0$) and deficits ($\delta_i^t < 0$) of heating power consumptions over the optimization horizon gives total energy surplus and deficits of the network. We report these values in percentage of total heating demands of the network, denoted respectively as $\delta_{surplus}$ and deficits $\delta_{deficits}$ (Equation 7). In addition to minimizing total heat production, the MPC aims to keep these surplus and deficits percentage within an acceptable range (below 2% in our work). From the operational standpoint, minimizing the heat production by minimizing the heat supplied by the sources may potentially increase consumer deficits. Conversely, increasing heat supply can result in large total heat production and unnecessary surplus of energy. To capture this trade-off and achieve optimal DHN control, the MPC objective function (Equation 8) minimizes a function which balances these opposing factors.

$$\delta_{surplus} = 100 * \frac{\sum_{(i,t)} \max(\delta_i^t, 0)}{\sum_{(i,t)} D_i^t} \tag{7}$$

$$\delta_{deficits} = 100 * \frac{\sum_{(i,t)} \min(\delta_i^t, 0)}{\sum_{(i,t)} D_i^t}$$

$$J(\dot{m}_s, T_{ss}) = \sum_{t=1}^{N_t} \left[\sum_{\{i \in S\}} \dot{m}_{s_i}^t c_p (T_{ss_i}^t - T_{sr_i}) \right] + M \left(\max(0, (\delta_{surplus} - 2))^2 + \max(0, (\delta_{deficits} - 2))^2 \right) \tag{8}$$

$$s. t. \quad \forall (i \in S), \forall t, \dot{m}_{s_i}^t c_p (T_{ss_i}^t - T_{sr_i}) \leq P_{ss_i} \tag{9}$$

The MPC objective function J (Equation 8) consists of two main parts. The first term represents the total heat production across all sources. Each source has a maximum supply power capacity P_{ss} . T_{sr_i} are the time-dependent temperature of the return water at the source i . The second term penalizes total energy surpluses and deficits ratios, exceeding chosen threshold of 2%. Numbers M_1 and M_2 determine the weight of each penalty term (e.g., 1000). Both the objective function (Equation 8) and the constraint (Equation 9) are non-analytic functions due to their dependence on the dynamic simulation model (Equations 1-6). Evaluating the objective function requires a complete simulation of the DHN to calculate state variables (water temperatures, consumptions, mass flow rates) using the control variables as boundary states. This non-smoothness prevents the use of analytical optimization algorithms. Similar to Fang and Lahdelma (2015), we employ metaheuristic-based optimization algorithm which is, in our case, the Mesh Adaptive Direct Search or MADS (Audet and Dennis, 2006), influenced by the high computational cost of the dynamic simulation. MADS extends the Generalized Pattern Search approach by exploring denser set of directions, enabling fast convergence but not guaranteed to the global minimum. In our application, we focus on reducing the overall computational time, regardless the optimality of the found solution as long as $\delta_{surplus}$ and $\delta_{deficits}$ are below the 2% objective. In this work, generation supply temperatures are limited to 70°C - 80°C, while supply mass flow rates are determined by the mass balance constraint (Equation 1). The optimizer essentially allocates mass flow



Figure 1: Illustration of DHN simplification on an example cluster

rates and corresponding supply temperatures to minimize the total necessary heat production while respecting sources' capacities and consumers' heating demands. Due to the convergence difficulty towards the optimal global, the number of iterations serves as the stopping criteria for the optimization process.

2.3 Hybrid MPC

Standard MPC suffers from high computational cost due to repeated calls of the dynamic simulation during optimization. For large DHN, this cost becomes impractical. Our hybrid MPC approach proposes a solution by combining dynamic simulation with machine learning models. Predetermined *Clusters* of substation nodes (*i.e.*, sets of connected consumer nodes), selected according to expertise analysis, are replaced with trained Neural Network (NN) models which mimic their physical behavior. These NN models eliminate the need to simulate the replaced clusters, significantly reducing the computational burden to simulate the network. Figure 1 exemplifies this approach with cluster of substation nodes (1, 2, 3). Here, the physical cluster is replaced by a single "numerical twin node v_{NN} " modeled by a trained NN model. This substitution reduces the network's topological complexity (see reduced DHN in Figure 1). However, the objective function (Equation 8) must remain consistent with the original and reduced DHN. This implies maintaining identical return temperatures at the sources and total deficits and surpluses of consumptions. To ensure this consistency, the "numerical twin nodes" preserves heat losses and transport delays of the replaced clusters. Additionally, the twin nodes maintain identical

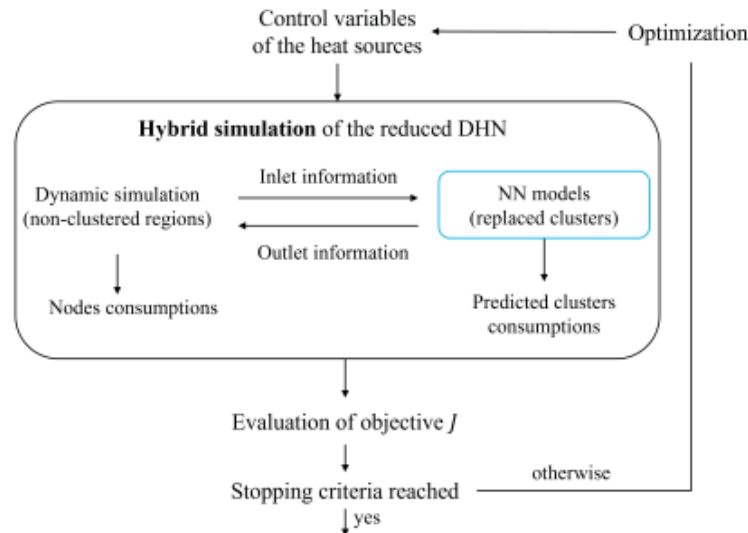


Figure 2: Hybrid MPC flow chart

consumptions as the original clusters, preserving $\delta_{surplus}$ and $\delta_{deficits}$. Figure 1 illustrates this concept. The twin node preserves outlet information (water temperatures) upstream of node 0 and downstream of nodes 4 and 5. In essence, the NN models predict both outlet information and real consumption of the replaced cluster. They take inlet information (water temperatures and mass flow rates) from the surrounding network and heating demands $\{D_1, D_2, D_3\}$ of the original cluster composing nodes as input. This creates a hybrid simulation framework where the dynamic simulation on the original network is replaced with a combination of NN models and a reduced network simulation. In this work, the considered clusters include only consumer substations. The focus of this paper lies on the potential of the hybrid simulation within the MPC framework, deferring a detailed explanation of the hybrid simulation for future studies. Figure 2 outlines the core process of the proposed "hybrid MPC" approach. The core idea consists in leveraging the use of hybrid simulation on the reduced DHN with lower computational load and preserving the physical equivalency with the original DHN.



Figure 3: NN models architecture

2.4 Neural Network models

The considered input and output features are highly time-dependent. To effectively capture the thermal inertia of the replaced substations clusters, we rely on many-to-one recurrent neural networks architecture. Such type of NN considers sequences of the input features with fixed length, called prediction length, and predicts the output feature at given time step. We consider a prediction length of 60 corresponding to the number of dynamic steps per time-steps used to simulate the DHN. More precisely, to predict output features at time t , input features between time steps $(t - 60)$ and t are fed into the neural networks. The chosen NN architecture utilizes two layers of gated recurrent units (GRUs) followed by a hidden dense layer and the output dense layer leading to the output features (see Figure 3). Dense (hidden and outputs) and GRUs layers use rectified linear unit and hyperbolic tangent respectively as activation functions. We have performed a random search on pre-defined values of hyperparameters including the number of units at each GRU layer, the number of activated neurons at the hidden dense layer, the batch size and the rate of dropout (Srivastava *et al.*, 2014). The hyperparameters leading to the best learning performances use 20 units for each GRU per layer, 80 units at the hidden dense layer, batch size of 64 and dropout rate of 0. Relying on this NN architecture, we compare two configurations of the NN models. The first consists in training a single NN model per cluster which predicts both outlet information (*i.e.*, temperatures of outgoing waters) and cluster real consumption (*i.e.*, total of composing nodes consumptions). While the second consists in training two NN models per cluster using the same presented architecture which respectively predict the outlet information and the cluster consumption (*i.e.*, total of composing nodes consumptions). These two versions of the NN models are respectively referred as “single model” and “two models” configurations. Indeed, the intuition is that the first configuration focuses on prediction speed while the second approach should emphasize precisions in the predicted values.

3 EXPERIMENTS AND CASE STUDY DHN

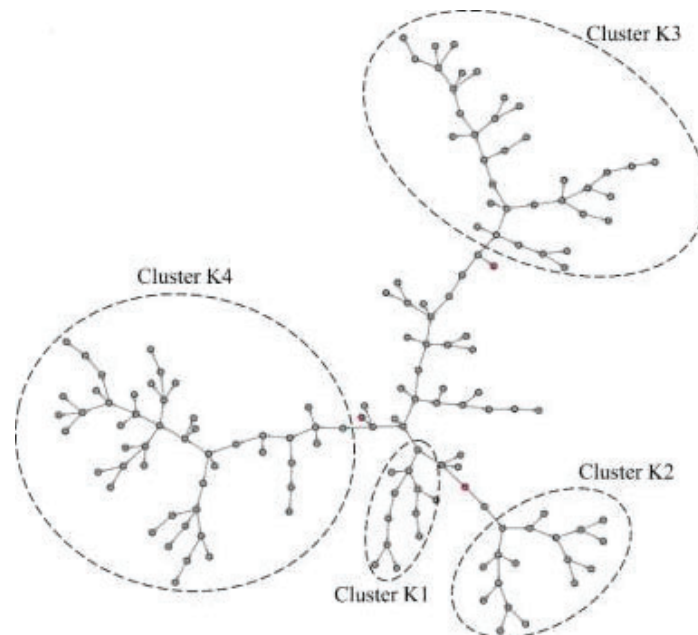


Figure 4: Synthetic District Heating Network case study, with considered clusters delimited in dashed lines

To validate the proposed hybrid MPC, we employ a synthetic DHN comprising 150 substation nodes. This case study network includes 3 distributed heat sources and 147 heat consumers, connected by 149 edges in total (149 supply pipes and 149 return pipes). Four clusters, denoted K1, K2, K3 and K4 (see Figure 4), are selected for substitution with NN models. The selection of the clusters, including their number and size (*i.e.*, number of composing nodes) are supported by operational considerations which are not detailed in the current paper. Listed clusters (*i.e.*, K1, K2, K3 and K4) are order by their sizes with cluster K1 the smallest (13 nodes) and cluster K4 the largest (50 nodes).

3.1 Data generation

To train and effectively test the NN models, we simulated the case study DHN using realistic heating demand forecasts from four distinct months (January, April, August, November) across 2880 hours (see Figure 5). Heating demands shown in this figure include both space heating and domestic hot water demands. Each hour is discretized by the simulation into 1 minute time steps, resulting in a total of 172,680 simulation steps, representing the data points. To ensure the NN models can handle the entire range of generation supply temperature (70°C - 80°C) which will be explored during the optimization processes, we have generated our simulation data with sinusoidal supply temperatures varying between these limits over a 3-day period (72-hour steps). This approach avoids abrupt temperature changes while covering the entire relevant domain. Since the optimization process determines the distribution of mass flow rates among sources, we have generated random dispatch values for all distributed sources. Furthermore, the values of dispatch which guide the supply mass flow rates are constrained by the sources' maximum capacities and must sum to one, reflecting the optimization's dispatch decisions.

3.2 Data preprocessing and split

The previously described simulation setup leads to a dataset composed by time varying variables (*i.e.*, waters temperatures, heating consumptions and mass flow rates) with length of 172,680. This dataset is split into training and test sets. Splitting this temporal data for training, validation, and testing purposes is crucial. To ensure the NN models' ability to generalize effectively, the test set must be independent of the training and validation sets. Therefore, we adopted a seasonal split strategy. Data from the first three months is used for training and validation, while data from the remaining month (November) serves as the test set for evaluating the models' performance. For illustration purpose of the data preprocessing, consider $X \in R^{(n_t \times n_i)}$ and $Y \in R^{(n_t \times n_o)}$, respectively the time-varying input and output features for a given replaced cluster (details in section 2.3). n_t is the total number of time steps in the data, 129,600 for the training data and 43,200 for the test. n_i and n_o are respectively the number the input features required and the output features to predict. X is divided into sequences $X_t \in R^{(60 \times n_i)}$. Each X_t is associated with the corresponding $Y_t \in R^{(1 \times n_o)}$. The couple (X_t, Y_t) constitutes a set. The training sets are then shuffled and split into training and validation sets. Validation sets include 20% of the overall training sets.




Figure 5: Case study DHN aggregated heating demands for four distinct months

3.3 Training procedure

The training process aims to minimize the mean squared error (Equation 10) between predicted output features \hat{y}_i and ground truth values from the dynamic simulation y_i for each set inside the training sets within a batch of size 64. This minimization is performed using the Adam optimizer (Kingma and Ba, 2014) for a maximum of 40 epochs. An early stopping mechanism halted the training process if the mean squared error measured on the validation sets after each optimization step does not improve over 10 consecutive epochs in order to preventing overfitting. Additionally, we decrease the learning rate by a factor of 0.1 every 10 epochs. Our experiments have shown better learning performances using such step-wide decreases.

$$\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2 \quad (10)$$

4 RESULTS

Learning performances of the NN models over the considered clusters are presented in this section along with the hybrid MPC results.

4.1 NN models learning performances

The performance of the NN models was evaluated using the Mean Absolute Percentage Error (MAPE) between predicted and simulated values (from dynamic simulation) over the test set. Table 1 reports the performance of both considered NN models' configurations to preserve the temperatures flowing out the replaced clusters and their thermal power and energy consumption. According to presented values, the "two models" configuration offers no significant improvement in learning and generalization performances compared to the "single model" configuration, while increasing computational complexity. Indeed, "two models" per cluster imply twice the training processes and prediction times. Therefore, we focused on the "single model" configuration for further analysis. Notably, the size of the clusters does not display strong correlation with the NN model performance. However, the number of output features, linked to the number of outgoing waters' temperatures to learn, directly impacts the performances of the model. More output features mean indeed more outgoing waters' temperatures to learn and predict and, thus, higher physical complexity to apprehend. Crucially, the NN models successfully conserved the thermal energy consumption of the replaced clusters (Table 1). In the reported values, the NN models show satisfying results in predicting the consumption power and energy of the replaced clusters. Results highlight higher error in power predictions than energy. The power errors likely stem from overestimation and underestimation at each time step, leading high range variations with mean values presented in Table 1. These errors in powers cancel off due the integration over the test sets, leading to low errors in consumption energy predictions. From our application, high accuracy on the consumption energy infers high accuracy on the $\delta_{surplus}$ and $\delta_{deficits}$ over the optimization horizon. Our experiments demonstrated a mean error of 0.52 % or less, implying that clusters' consumptions energy was accurately predicted over the test sets.

Table 1: NN models learning performances results

	Cluster K1	Cluster K2	Cluster K3	Cluster K4
Number of clustered nodes	13	18	36	50
Number of output features	3	2	2	2
Temperature error				
▪ Single model	0.58 %	0.06 %	0.22 %	0.23 %
▪ Two models	0.59 %	0.27 %	0.22 %	0.22 %
Consumption power error				
▪ Single model	2.31 %	0.92 %	2.65 %	1.79 %
▪ Two models	1.47 %	0.76 %	2.66 %	1.56 %
Consumption energy error				

▪ Single model	0.07 %	0.52 %	0.43 %	0.24 %
▪ Two models	0.27 %	0.51 %	0.42 %	0.26 %

4.2 Hybrid MPC results

The primary aim of the proposed hybrid MPC is to generate optimal control variables of the DHN while reducing the computational time. Therefore, achieving significant time reduction is crucial alongside accuracy. Our experiments with results presented in Table 2 used an optimization horizon of 10 hours, with a 3-hours sliding window advancing by one hour. We tested two decision time steps referred as dt : 1 hour and 0.2 hour (*i.e.*, 12 minutes). To assess the accuracy of hybrid MPC, we simulated the original network using the control variables obtained from the hybrid MPC on the reduced networks. Results in Table 2 use the “single model” configuration. Deficits and surpluses obtained on the original network stay within our chosen threshold (*i.e.*, below 2 %). This suggests that the hybrid MPC on the reduced networks have successfully mimicked the physical MPC on the original network by successfully preserving the values of the objective function J with acceptable errors. For further validation, we compare the hybrid MPC performances with a rule-based control strategy (Jansen *et al.*, 2023) that maintains a constant generation supply temperature of 80°C. Dispatch of the supply mass flow rates is guided by a preference order on the sources based on production cost and capacity. Results in Table 2 show that the rule-based controller overproduces due to the high fixed generation temperature, for both dt values. With 1-hour dt , the rule-based exhibits high consumption surplus due to high temperature incoming flows and large decision time. However, with 12-minute dt , it achieves comparable consumption deficits and surpluses to the hybrid MPC. Indeed, with lower decision time, the consumers can adapt more dynamically their consumption rates in response to the delivered heat. Nevertheless, the rule-based controller tends to overproduce which leads to high heat energy produced by the sources for the entire DHN operation. Additionally, as shown in Table 1, the low performance of the neural network (NN) models in predicting the clusters’ leaving temperatures affects the hybrid MPC’s ability to accurately evaluate the power generated by the sources. This might lead to underestimating or overestimating the heat that needs to be produced which we observe for clusters K1 and K4 (see Table 2).

Table 2: Performances of the Hybrid MPC by substituting the clusters, the physical MPC on the original network and the rule-based control for 10-time steps optimization horizon using two decision time steps ‘dt’ of 1 hour and (0.2) hour

Clusters	Physical MPC	Hybrid MPC				Rule-based
	-	K1	K2	K3	K4	-
Nb. nodes	0	13	18	36	50	0
dt (hour)	1 (0.2)	1 (0.2)	1 (0.2)	1 (0.2)	1 (0.2)	1 (0.2)
Produced energy (MWh)	202.8 (218.3)	206.6 (203.5)	202.7 (218.4)	203.2 (213.6)	203.6 (209.7)	235.7 (232.4)
$\delta_{deficits}$	1.46 (0.71)	1.37 (0.20)	1.44 (0.68)	1.38 (0.57)	1.40 (0.57)	1.32 (0.35)
$\delta_{surplus}$	0.41 (0.62)	0.68 (0.51)	0.39 (0.63)	0.40 (0.45)	0.44 (0.36)	5.94 (0.95)

In addition to the accuracy, we compare the computational time taken by the hybrid MPC on the reduced networks and by the physical MPC on the original network (Figure 6). Our MADS algorithm performs at most 20 iterations per time step over the total optimization horizon. It means that the dynamic simulations or the hybrid simulations are run 200 times in our experimental horizon (*i.e.*, 10-time steps). Computational time reduction by the hybrid simulation cumulates over the iterations and the optimization leading to noticeable computation gain. Each black-reverse triangle in Figure 6 represents a reduced network by substituting the associated cluster. The largest cluster considered in this study (*i.e.*, cluster K4) contains 50 nodes representing (1/3) of the original DHN number of nodes reduces the

required computational time by a factor of 9% compared to physical MPC on the original DHN. In fact, the results in Figure 6 suggests a trend of computational time reduction with the size of the considered clusters. Notably, the substitution of the smallest cluster (*i.e.*, cluster K1) has not given any reduction in computational time. In fact, the reduction in physical pipes and nodes, and thus in physical equations to solve, has not compensate the NN models inference time. Therefore, by extrapolation, larger clusters may in fact lead to significant computational time gain in controlling the large-scale case study DHN.




Figure 6: Computational time comparison between physical MPC on original DHN and hybrid MPC on reduced networks

Table 1 results have indeed shown no direct correlation between the NN performances and the size of the clusters. Although, some clusters may present higher physical complexity, leading to more difficulty of the NN models to learn them. Therefore, a compromise between NN models' performances and potential computational time reduction must be evaluated before clusters substitutions.

5 CONCLUSIONS

This paper introduces a novel hybrid MPC approach that reduces the computational cost of controlling large-scale District Heating Networks. The approach integrates dynamic simulation with machine learning by using neural networks models to learn the physical behavior of substation clusters and replace them within the MPC optimization process. Our results show that the NN models effectively predict both outlet water temperatures and aggregated heating consumption across different control scenarios during a month-long test, with measured absolute percentage errors remaining below 0.5%. The proposed hybrid MPC has achieved a 9% decrease in computational time required to derive optimal control variables for the studied 150-node DHN. Interestingly, the results suggest a correlation between computational time reduction and the size of the replaced clusters which enforces the potential the proposed approach on larger clusters. However, results have also highlight learning performances disparities among the considered clusters which are not directly correlated with the size of the substituted clusters. This finding suggests that cluster size might not be the main factor influencing the NN models' ability to learn cluster physical dynamics. In contrast, the number of temperatures to learn and predict seem to have more impact the NN models' performances. Future work will focus on investigating the key physical factors or clusters internal topology metrics that most significantly affect the NN models' performances. In addition, to further reduce computational time, future research will focus on simplifying the NN models.

ACKNOWLEDGMENT

We would like to thank the CNRS for granting this work access to the HPC/AI resources of the “Institut du développement et des ressources en informatique scientifique” under the allocation 20XX-AD011014376, made by GENCI.

NOMENCLATURE

Abbreviations

DHN	District Heating Network
ML	Machine Learning
NN	Neural Networks
GRU	Gated Recurrent Unit

Topology

S	List of source nodes inside the DHN
C	List of consumer nodes inside the DHN
E_i^+	List of edges pointing towards the node i (incoming supply pipes)
E_i^-	List of edges pointing away from the node i (incoming return pipes)

Physical values (not defined in the text)

T_e	Temperatures of waters traversing the pipes of the edge e	(K or °C)
T_{r_i}	Temperatures of return waters at the substation node i	(K or °C)
T_{s_i}	Temperatures of supply waters at the substation node i	(K or °C)
\dot{m}_e	Mass flow rates of waters traversing the pipes of the edge e	(kg/s)
c_p	Water specific heat capacity	(J/kg/K)
ρ	Water density	(kg/m ³)
g_{vr}	Gravity intensity	(m/s ²)
d_e	Diameter of pipes composing the edge e	(m)
ΔH_e	Head pressure losses through pipes of the edge e	(m)
h_e	Convective loss coefficient of pipes composing the edge e	(W/K/m ²)
l_e	Length of pipes composing the edge e	(m)
f_e	Friction factor of pipes composing the edge e	(-)
T_{ground}	Temperature of ground outside the pipes	(K or °C)

Subscript

x	Flow axis
i	Substation node (source or consumer) i
e	edge e (containing a supply and a return pipe)

Up-script

t	Current time step t
-----	-----------------------

REFERENCES

- Audet, C., and Dennis, J. E., 2006, Mesh Adaptive Direct Search Algorithms for constrained optimization, *SIAM Journal on Optimization*, vol 17, no 1, p. 188-217
- Bavière, R., Vallée, M., 2018, Optimal temperature control of large scale district heating networks, *International Symposium on District Heating and Cooling*, Energy Procedia: p. 69-78
- Benonysson A., Bøhm B., Ravn F. H., 1994, Operational optimization in a district heating system, *Energy convers manage*, vol 36, no 5, p. 297-314
- Betancourt S., M., Mabrouk, M. T., Silva, C. S., Harrant, P., Lacarrière, B., 2019, Modified finite volumes method for the simulation of dynamic district heating networks, *Energy*, vol 182, p. 954-964

- Buffa, S., Cozzini, M., D'Antoni, M., Baratieri, M., Fedrizzi, R., 5th generation district heating and cooling systems: A review of existing cases in Europe, *Renew sust energy rev*, vol 104, p. 504-522.
- Cahuantzi, R., Chen, X., Güttel, S., 2023, A comparison of LSTM and GRU networks for learning symbolic sequences, *Intelligent Computing*, vol 2, p. 771-785
- Chung, J., Gulcehre, C., Cho, K., Bengio, Y., 2014, Empirical evaluation of gated recurrent neural networks on sequence modelling, *ArXiv:1412.3555*
- Fand, T., and Lahdelma, C., 2015, Genetic optimization of multi-plant heat production in district heating networks, *appl energy*, vol 159, p. 610-619
- Giraud, L., Bavière, R., Paulus, C., Vallée, M., Robin, J. F., 2015, Dynamic modelling, experimental validation and simulation of a virtual district heating network, *ECOS'2015*
- Giraud, L., Merabet, M., Baviere, R., Vallée, M., 2017, Optimal control of district heating systems using dynamic simulation and mixed integer linear programming, *International Modelica Conference*, Linköping Electronic Conference Proceedings, p. 141-150
- Guelpa, E., Marincioni, L., Verda, V., 2019, Towards 4th generation district heating: Prediction of building thermal load for optimal management. *Energy*, vol 171, p. 510–522
- Hochreiter, S., Schmidhuber, J., 1997, Long short-term memory. *Neural Comput*, vol 9, p. 1735–80
- Jansen, J., Jorissen, F., Helsen, L., 2023, Optimal control of a fourth generation district heating network using an integrated non-linear model predictive controller, *Appl therm eng*, vol 223, p. 120030
- Kane, M., and Rolle, J., 2020, “quantum networks”: a new approach for representing a network and evaluating hydraulic and thermal losses in district heating/cooling systems, *ECOS'2020*
- Kennedy, J., Eberhar, R., 1995, Particle swarm optimization. *International Conference on Neural Networks*, vol 4, p. 1942–1948
- Kingma, D. P. and Ba, J., 2015, Adam: A method for stochastic optimization, *International Conference for Learning Representations*, vol 3
- Lund, H., Østergaard, P. A., Nielsen, T. B., Werner, S., Thorsen, J. E., Gudmundsson, O., Arabkoohsar, A., Mathiesen, B. V., 2021, Perspectives on fourth and fifth generation district heating, *Energy*, vol 227, p. 120520
- Ma, L., Huang, Y., Zhang, J., Zhao, T., 2022, A model predictive control for heat supply at building thermal inlet based on data-driven model, *Buildings*, vol 12, no 11, p. 1879
- Marquant J. F., Bollinger, L. A., Evins, R., Carmeliet, J., 2018, A new combined clustering method to analyse the potential of district heating networks at large-scale, *Energy*, vol 156, p. 73–83
- Marquant J. F., Evins, R., Carmeliet, J., 2015, Reducing computation time with a rolling horizon approach applied to a MILP formulation of multiple urban energy hub system, *International Conference on Computer Science*, vol 51, p. 2137–2146
- Saloux, E., Runge, J., Zhang, K., 2023, Operation optimization of multi-boiler district heating systems using artificial intelligence-based model predictive control: Field demonstrations, *Energy*, vol 285, p. 12952
- Sandou, G., Font, S., Tebbani, S., Hiret, A., Mondon, C., 2006, Predictive Control of a Complex District Heating Network, *IEEE conference on Decision and Control, and the European Control conference*, vol 2005
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R., 2014, Dropout: A Simple Way to Prevent Neural Networks from Overfitting, *J. Mach. Learn. Res.*, vol 15, no 56, p 1929-1958
- Stoffel, P., Henkel, P., Rätz, M., Kümpel, A., Müller, D., 2023, Safe operation of online learning data driven model predictive control of building energy systems, *Energy and AI*, vol 14, p. 100296
- Wirtz, M., Neumaier, L., Remmen, P., Müller, D., 2021, Temperature control in 5th generation district heating and cooling networks: An MILP-based operation optimization, *appl energ*, vol 288, p. 116608
- Zhang, Z., Zhou, X., Du, H., Cui, P., A new model predictive control approach integrating physical and data-driven modelling for improved energy performance of district heating substations, *energ buildings*, vol 301, p. 11368
- Zimmerman, N., Kyprianidis, K., Lindberg, C.F., 2019, Achieving lower district heating network temperatures using feed-forward MPC, *materials*, vol 12, no 15, p. 2465