LEVERAGING DECENTRALIZED BUFFERS FOR PEAK SHAVING AND LOAD SHAPING IN 100% RENEWABLE DISTRICT HEATING

Summary

Flexible charging strategies for decentralized buffer storages can be used to shave thermal load peaks or shape loads to support the transition to only renewable heat sources and eliminate peak load gas boilers. A neural network model is trained to forecast the resulting substation thermal power and secondary return temperature for operational planning. Training data are created using physics-based. Polysun simulations for different space heating (SH) and domestic hot water (DHW) profiles and two hydraulics (buffer tank plus DHW tank or freshwater station). Peak loads are reduced by 31.9% compared to direct DHW consumption for the configuration with DHW tank and by 16.3% with central freshwater station. Forecast errors are 2.6% (power) and 0.7% (return temperature) with DHW tank; 2.9% and 1.7% with freshwater stations. Generalization capabilities are tested for wide SH/DHW ratios. With awareness of some limitations, the approach enables proactive scheduling and aligning the substation loads with renewable availability in the DH heat generation.

Keywords: district heating, decentralized buffer storage, thermal load peak shaving, neural network

Field of application for the findings: district heating operation

Addressed audience: researchers

1. Introduction

District heating (DH) networks can significantly contribute to the transition toward a fossil-free building heat supply, particularly in densely populated areas, as they facilitate the integration of renewable energy sources and enable coordinated management of heat generation and demand. Many DH networks still rely on gas boilers to cover peak thermal loads. To transition to 100% renewable DH systems, system flexibilities must be exploited and more sophisticated control strategies developed that can adapt to variable and only partially dispatchable renewable heat sources. Unlike fossil heat sources, renewable heat supply exhibits limited dispatchability and schedulability, making flexibility options essential for reliable, efficient operation. Flexibility exists on both, the network and consumer level. On the network side, central buffer storages [1] and the thermal inertia of the network itself [2,3] can be leveraged for peak shaving and temporal shifting of demand. On the consumer side, existing approaches utilize the thermal mass of buildings [4] and user behavior (Knudsen et al. 2025) to modulate demand. However, decentralized buffer storages which exist in most multi-residential buildings remain underutilized as a flexibility resource. These storages are often charged using simple temperature threshold rules, which tend to synchronize charging events and amplify rather than alleviate peak loads. To address these challenges, this work proposes control strategies for decentralized buffers that can be used for responding to current network needs - either heat surplus or scarcity - and a fast forecasting model that predicts the resulting thermal load profiles from customers under such control strategies. Space heating (SH) and domestic hot water (DHW) demand profiles used as an input into these models remain unchanged; flexibility comes solely from the use of the buffer storage. This work includes two goals: (i) develop and evaluate decentralized buffer charging strategies capable of producing either damped thermal load profiles (peak shaving) or intentionally shaped profiles (artificial peaks at favorable times), and (ii) forecast the resulting substation load profiles from given SH and DHW inputs to support efficient real-time control of DH systems. For this task, a neural network approach is chosen, where a Convolutional Neural Network (CNN)-Long Short-Term Memory (LSTM) model is trained for each control strategy. The existence of many different hydraulic connection configurations on the consumer side, makes it also necessary to train different models for both control strategies for two main categories of hydraulic configurations (DHW tank or freshwater stations). By shifting the focus from purely local temperature thresholds to network-needresponsive logic, the proposed approach aims to reduce reliance on fossil peak coverage and enable alignment of consumer-side loads with availability of renewable heat. The forecasting component provides operators with fast, integrable predictions of consumer load under different buffer control modes, facilitating proactive scheduling and source dispatching of DH systems.

2. Methods

2.1 Development of control strategies and data generation

Control strategies are developed using the software Polysun [5]. This software is suitable for the simulation of energy systems and comes with components such as buffer storage, heat exchanger and the possibility to define consumption profiles for SH and DHW. Different hydraulic configurations can be set up with these components. Figure 1 shows an exemplary hydraulic setup for a parallel supply of SH and DHW preparation from a buffer tank with a DHW tank. Polysun allows flexible control of the pumps by custom-scripting of control algorithms. Using this functionality, different buffer charging strategies can be developed and yearly simulations are conducted for each, using diverse SH and DHW load profiles. These simulation results serve as training and test data for the neural network model. Established forecasting tools for SH and DHW are assumed to provide the exogenous inputs, which remain unaltered by this control. For each hydraulic configuration, SH and DHW demands are varied within system-feasible bounds to ensure reliable demand coverage, and the DHW-to-SH energy ratio is systematically adjusted to represent different building insulation qualities. The outputs of the model are the thermal power at the substation and secondary return temperature.

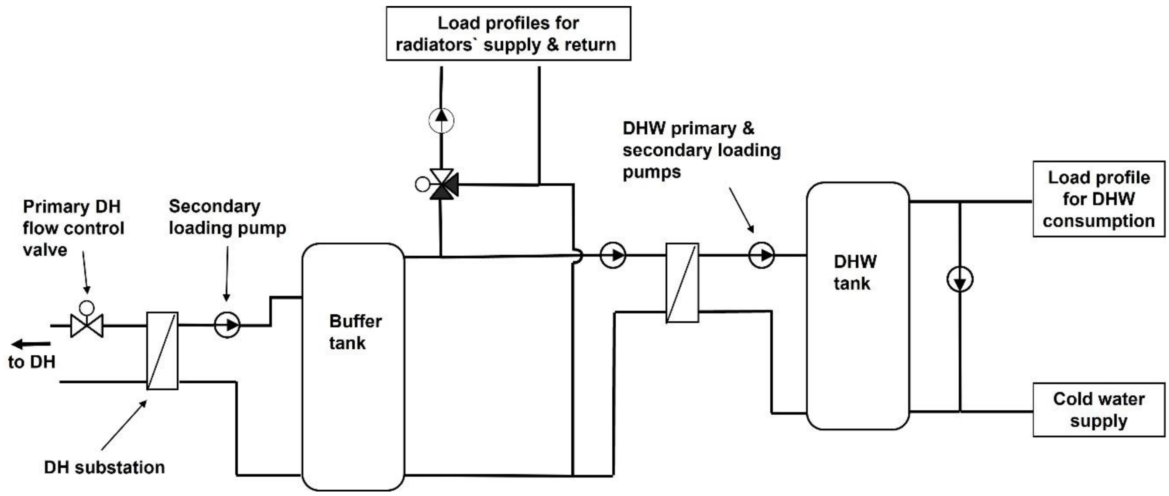

Fig.1: Schematic setup of a hydraulic configuration with DHW storage and parallel SH and DHW supply from a buffer storage

2.2 Forecasting model (CNN-LSTM dual backbone)

The forecasting models are set up using Keras [6] with the TensorFlow [7] backend. Each consists of two separate, architecturally identical backbones: one predicts the thermal power at the substation, the other the secondary return temperature. Each backbone consists of a CNN part to extract local features from the input sequences and an LSTM part to learn temporal dependencies. The training of models is performed on the HPC system of the University of Innsbruck, on an a30 GPU of the leo5 Compute Cluster.

2.3 Evaluation strategy

Models are trained and tested on full simulated years for each control strategy and for each hydraulic configuration. Performance is evaluated against Polysun outputs for both, power and secondary return temperature, using the mean absolute error (MAE) metric. To assess model robustness, inputs cover a broad range of SH and DHW energy demands and profiles, and the annual DHW-to-SH energy ratio is varied within configuration-feasible bounds to represent different building insulation qualities (better insulation lowers annual SH, increasing the annual DHW/SH ratio).

3. Results

Comparing the temperature treshold-based, step-wise buffer charging control to the pure sum of SH and DHW power consumption, which corresponds to the thermal load without buffer storage, it exhibits significant substation peaks. The peak-shaving buffer charging control strategy, on the other hand, can lower the direct consumption peaks. The yearly average reduction of the thermal peak loads is around 53.3% for the peak-shaving control compared to the step-wise control and 31.9% compared to the pure consumption. This percentual reduction refers to all load peaks whose prominence was above the 70th percentile of the prominence distribution of all peaks and that have a minimum distance of 12 h on the time axis. For a hydraulic configuration with central freshwater station instead of a DHW storage, the thermal peak loads are around 47.8% less for the peak-shaving control compared to the step-wise control and 16.3% less compared to the pure consumption. A CNN-LSTM model was trained for the peak-shaving control concept for the hydraulic configuration shown in Fig. 1 and tested for a variant with 218 MWh yearly SH demand and 88 MWh yearly DHW demand. Comparing the results to the Polysun outputs revealed a yearly mean percentual error of 2.6% of the neural network model for the power predictions and 0.7% for the return temperature predictions. The generalizability of the model for yearly SH demands up to 350 MWh and for annual SH to DHW demand ratios from 0.8 up to 8.2 was shown to be good. A model with the same architecture trained for the hydraulic configuration with central freshwater station instead of a DHW storage, has for the same SH and DHW input profiles a yearly MAE of 2.9% for the power predictions compared to the Polysun output and 1.7% for the secondary return temperature.

4. Discussion

The peak-shaving buffer charging strategy substantially reduces load peaks, and the CNN–LSTM model provides fast forecasts of substation thermal power and secondary return temperature, enabling operational planning. The runtime for a full year forecast on a single Graphical Processing Unit (GPU) with 24 GB high-bandwidth memory (CUDA 12.6 and driver version 560.35) is around 40 s, facilitating real-time use and integration into supervisory control or model predictive control in a digital twin as a surrogate model. Nevertheless, the model`s performance is inherently data-dependent. Polysun simulations are physics-based and the software is validated against measurement data, yet experimental calibration of the CNN-LSTM model remains open. Circulation and other losses are not explicitly modeled, they enter only via training data and may differ across buildings, introducing systematic bias if uncorrected. Forecasts seem to remain usable, but peak-power errors may increase for substantially different tank sizes. Less common hydraulic setups, which cannot be represented by one of the two presented configurations require separate training of new models. Keeping these limitations in mind, the method supports alignment of consumer-side loads with renewable heat availability in a DH network and reduces reliance on fossil peak coverage.

5. References

[1] R. Egging-Bratseth, H. Kauko, B. R. Knudsen et al., “Seasonal storage and demand side management in district heating systems with demand uncertainty,” *Applied Energy*, vol. 285, p. 116392, 2021.

[2] J. L. Svensen, “Peak Shaving in District Heating Utilizing Adaptive Predictive Control,” *Energies*, vol. 15, no. 22, p. 8555, 2022.

[3] T. van Oevelen, L. Scapino, J. A. Koussa et al., “A case study on using district heating network flexibility for thermal load shifting,” *Energy Reports*, vol. 7, pp. 1–8, 2021.

[4] J. Lind, E. Möllerström, H. Averfalk et al., “Energy flexibility using the thermal mass of residential buildings,” *Energy and Buildings*, vol. 301, p. 113698, 2023.

[5] Vela Solaris AG, *POLYSUN*, Vela Solaris AG.

[6] F. e. a. Chollet, *Keras*, Google LLC, 2015.

[7] TensorFlow Developers, *TensorFlow*, Google LLC, 2015.