AI-Assisted Operation Planning and Strategy for a Hybrid Thermochemical System for Waste-Heat Upgrade

Munendra Pal Singh, and Ahmad Arabkoohsar

Technical University of Denmark, Copenhagen, Denmark, mpsingh.iitb@gmail.com

Abstract:

In an era of rapid digitalization, the resulting surge in energy demand is driving industries to run at higher loads for longer hours, thereby strengthening dependence on fossil fuels, showing the limits of intermittent renewables, and increasing the amount of waste heat (WH) that goes unused in the environment. Especially in Europe, a significant portion of the supplied process heat is discharged as medium- and high-temperature WH, which can be properly treated, reused, or integrated into other energy systems. Bridging this gap requires robust operational and strategic planning for industrial energy systems that is not only focused on WH cutting but also on recovery and upgrading, offering considerable scope to reduce primary fuel consumption, operational costs, and CO_2 emissions while enhancing industrial competitiveness.
Within this context, the EU-funded TechUPGRADE project investigates a hybrid heat-dispatch concept that couples a thermochemical heat transformer, an oil-based sensible storage tank, and an auxiliary electric heater. The integrated system upgrades waste heat to higher temperature levels. It redistributes it to multiple process users, thereby supplying high-temperature process heat and reducing the net heat load at industrial sites. This paper develops an AI-assisted optimization framework for planning and operating strategy such a system in two stages.
First, yearly data have been used to determine near-optimal sizes for the thermochemical unit, storage tank, and electric heater, along with an electricity price threshold that minimizes the levelized cost of heat (LCOH). Second, a rolling-horizon model predictive control (RH-MPC) layer operates the plant in real time. At each hour, machine-learning models forecast heat demand, waste-heat availability, and electricity prices over a 36-hour horizon. The MPC then solves a constrained cost-minimization problem that allocates waste heat, charges or discharges storage, and dispatches the electric heater while respecting equipment capacities and storage limits. Simulation results show that, compared with a static threshold-based strategy, the proposed RH-MPC achieves an operating cost reduction of approximately 18% and lowers electricity purchases by around 3–4%, while maintaining full process heat supply with zero unmet demand. A modest improvement in storage utilization (~5%) is also observed. These results highlight the role of coordinated dispatch across thermochemical upgrading, sensible storage, and electrified backup supply in improving system-level performance under dynamic conditions. Although demonstrated for the TechUPGRADE configuration, the methodology is objective-oriented and can be extended to other energy-intensive industrial sites integrating waste heat recovery and electrified heat generation within the EU decarbonization pathway.

Keywords: Thermo-chemical; Thermal Energy; Operation Planning, Neural Network, Optimization.

1. Introduction

The rapid digitalization of the industrial system is pushing high load requirements. It is difficult for renewable energy systems to satisfy this requirement because it is intermittent in nature. This continues to burden the fossil fuels, which leads to a large amount of thermal energy rejected as waste heat (WH). WH has high technical recovery potential, but is still especially seen in European industries, which are rejected without any treatment [1], [2]. In this direction, solutions like thermochemical heat transformers bridge a technology gap and lead to lower operating costs and reduced GHG emissions [3], [4]. Similarly, thermal energy-based storage and heating are a reliable solution during peak-hour operation [5]. However, successful usage of such systems depends on their operation over time in addition to the component efficiencies. The industry operated under a dynamic environment, the heat demand is fluctuating with time, and in turn, the WH availability depends on process operation. In addition, the electricity prices are driven by grid market volatility and the availability of renewable energy integration [6]. These factors introduce complex trade-offs between direct utilization, storage, and electrification [7]. This brings the importance of effective operational planning into picture, which ensures that there is an efficient balance between available WH, available storage systems, and favorable electricity price. In contrast, poorly designed operational strategies can introduce avoidable inefficiencies and unnecessarily increase the operational cost, affect system reliability, and lead to unmet process heat demands [8]. Such factors can affect the overall economic viability of the system and offset the expected benefit from the integration of advanced technologies. In terms of control strategies, conventional rule-based approaches have limitations wherein they are unable to take into account short-period fluctuations or predict future operating conditions. This paves the way for advanced control strategies, like model predictive control (MPC), which allows forecasts and system constraints to be incorporated into real-time decision-making. When combined with machine learning techniques for predicting demand, waste heat availability, and electricity prices, MPC offers a flexible framework for improving operational efficiency and system robustness.

Within this context, the EU-funded TechUPGRADE project investigates a heat-upgrading concept capable of converting low- or medium-temperature waste heat into process heat at medium-to-high temperature levels. In the present work, this technology is used as a case study to demonstrate AI-assisted operation planning and control. A hybrid system consisting of a thermochemical heat transformer (TCP), an oil-based sensible thermal storage tank (OST), and an auxiliary electric heater (EH) is considered. The proposed methodology explicitly links long-term system sizing with short-term operational planning, thereby addressing both design and real-time decision-making challenges. The approach is structured in two stages. In the first stage, annual profiles of waste heat availability, process heat demand, and electricity prices are used as inputs to a Genetic Algorithm (GA) that determines near-optimal component sizes and an electricity price threshold minimizing the levelized cost of heat (LCOH). In the second stage, robust real-time operation is achieved through a rolling-horizon Model Predictive Control (RH-MPC) framework. At each time step, the system optimizes its operation over a finite prediction horizon of 36 hours using forecasts of demand, waste heat, and electricity prices. While the optimization determines dispatch decisions for all components over the horizon, only the first control action is implemented before the horizon is shifted forward, and the process is repeated using updated system states and forecast information.

The main contribution of this work is to demonstrate the importance of coordinated operation planning in hybrid industrial energy systems. In particular, the study establishes a consistent link between techno-economic system design and real-time operation through a forecast-driven RH-MPC framework, ensuring that sizing decisions remain aligned with actual operating conditions. The comparative results show that use of rolling-horizon MPC enables forecast-informed, feedback-based control that lowers operating costs and enhances grid utilization, while consistently meeting process heat demand more effectively than conventional static strategies. Taken together, these findings suggest that the proposed framework provides a practical and scalable solution for combining waste heat recovery with electrified heating, supporting wider industrial decarbonization efforts.

1.1. System Components and Operational Workflow

The proposed hybrid system consists of three major energy systems: a TCP, OST and EH to facilitate efficient waste heat recovery, upgrading, and utilization as shown in Figure 1. This hybrid system configuration has been designed in such way to prioritize waste heat recovery, utilize dependency on electricity grid while volatile variations in process heat demand, waste heat availability and electricity prices. The interaction between each system component is managed through a coordinated dispatch strategy, where available waste heat is upgraded directly to meet demand first, whereas excess energy is stored, and electrification is employed only when economically justified or highly required. This approach allows the system to respond flexibly to changing conditions, improves the overall utilization, and supports more cost-effective operation over time.

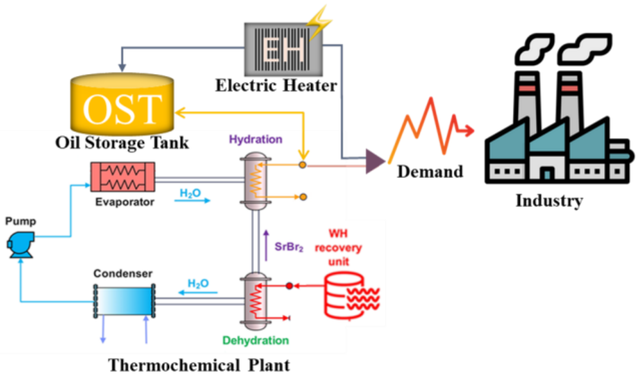

Figure 1: Hybrid system for heat upgradations.

Thermochemical Unit (TCP): The TCP operates with the highest priority, continuously utilizing the low-temperature unutilized waste heat as input energy. It consists of hydration and dehydration reactors, along with a condenser and evaporator, operating with the working pair $SrBr_2 \cdot H_2O$. In the dehydration reactor, heat input drives an endothermic reaction that converts hydrated $SrBr_2$ into anhydrous $SrBr_2$ and water vapor. The vapor is condensed and then re-evaporated to supply the hydration reactor at high temperature, where an exothermic reaction occurs as $SrBr_2$ absorbs water vapor, releasing high-temperature process heat. The solid material is cyclically transferred between the two reactors, enabling continuous heat upgrading, with an effective coefficient of performance (COP ≈ 0.605). A more detailed description of such a thermochemical reactor system can be found here [3], [4]. Any excess upgrade heat after fulfilling the process heat demand is stored in the buffer storage, ensuring availability during future peak demand periods while grid utilization is too expensive. During operation planning and execution, TCP output is constrained by either waste heat availability or system capacity.

Oil Storage Tank (OST): OST is a sensible thermal energy storage unit and a secondary thermal buffer in the hybrid system, which bridges the required heat generation from heat demand over time. It is charged using extra heat production through TCP and by the electric heater during periods of low electricity prices. However, when there is a peak demand periods or prices are not economically viable, the OST is discharged within its cap limit to supply process heat, by reducing dependence on external energy sources. This enhances the overall system flexibility and operating cost. Once the OST reaches its full nominal capacity, excess heat production or low-cost electricity can be exploited to increase the OST's effective thermal level, which further improves the system's ability to respond to future peak demands.

Electric Heater (EH): EH is an auxiliary energy system which enhances overall system flexibility by enabling grid utilization. It covers short-term process heat demand and charges OST during periods of low electricity prices, enabling exploitation of favorable market conditions and leading to low total operating costs. The electricity market prices drop significantly during high and sunny weather conditions, this is a good condition to run the EH on full load. When WH is unavailable or the TCP output is compromised, EH serves as a full backup to avoid a significant penalty and ensure a continuous process heat supply.

The present analysis assumes ideal operating conditions for all system components and therefore does not explicitly account for thermal or operational losses. This simplification is adopted to maintain a clear focus on overall system behavior and the effectiveness of the proposed control strategy. As a result, the reported performance metrics should be interpreted as optimistic estimates relative to those expected in a real industrial implementation. Incorporating loss mechanisms and non-ideal component behavior is planned as part of future work, with the aim of providing a more realistic assessment of system efficiency and operational performance.

2. Operation planning and Strategy

2.1. Optimization Problem Formulation

The operation and planning of the proposed hybrid system are formulated as a constrained cost minimization problem. The objective is to ensure a reliable supply of process heat while minimizing the overall operating cost through coordinated use of waste heat, thermal storage, and grid electricity. To guarantee system reliability, unmet heat demand is penalized within the optimization framework.

The objective function is defined as:

$$\min \sum_{t=1}^{T} \left[price(t) \cdot \frac{EH_{dem}(t) + EH_{ch}(t)}{\eta_{EH}} + C_{WH} \cdot WH_{used}(t) + \lambda \cdot Unmet(t) \right] \tag{1}$$

where $price(t)$is the electricity price, C_{WH} is the cost of WH, $EH_{dem}(t)$ is the EH supply to demand, $EH_{ch}(t)$ is the EH charging to storage, $\eta_{EH}$is the heater efficiency, and $Unmet(t)$ represents unmet heat demand penalized by a large factor λ. The objective function accounts for the cost of electricity consumption, the utilization cost associated with WH, and a penalty for unmet process heat demand. Electricity costs arise from both the direct process heat supply and the charging of storage through the EH. In addition, to ensure reliable system operation, unmet demand is heavily penalized within the optimization

At each time step t, the heat balance constraint is given by:

$$TCP_{dem}(t) + EH_{dem}(t) + Dis(t) = heat_{demand(t)} - Unmet(t) \tag{2}$$

where $TCP_{dem}(t)$is the thermochemical heat supplied to demand and $Dis(t)$is the storage discharge.

The storage energy balance follows:

$$SOC(t+1) = SOC(t) + TCP_{ch}(t) + EH_{ch}(t) - Dis(t) \tag{3}$$

subject to:

$$0 \leq SOC(t) \leq OST_{cap} \tag{4}$$

where $SOC(t)$is the state of charge of the storage system and $OST_{cap}$is the storage capacity. The thermochemical unit is constrained by both installed capacity and waste heat availability:

$$TCP_{dem}(t) + TCP_{ch}(t) \leq \min\left(TCP_{cap}, COP \cdot waste_{heat(t)}\right) \tag{5}$$

The electric heater operation is limited by:

$$EH_{dem}(t) + EH_{ch}(t) \leq EH_{cap} \tag{6}$$

In addition, charging and discharging rates are bounded:

$$Dis(t) \leq Dis_{max}, TCP_{ch}(t) + EH_{ch}(t) \leq Ch_{max} \tag{7}$$

A key feature of the formulation is the use of an electricity price threshold P_{th}, which governs the cheap and expensive grid hours. When $price(t) \leq P_{th}$, the electricity is cheap, leading to extensive use of EH to allow charging of OST and, if needed, meet the remaining demand as well. On the other hand, while electricity is too expensive, storage discharge is prioritized before using the electric heater, that minimizing reliance on grid electricity.

The optimization is implemented in two stages. In the first stage, annual data are used to determine near-optimal system sizes ($TCP_{cap}, OST_{cap}, P_{th}$) by minimizing the LCOH. In the second stage, the RH-MPC strategy is applied. At each time step, forecasts of heat demand, waste heat, and electricity price are used over a 36 h horizon, while only the first control action is executed, and the process is repeated in a receding-horizon manner. Since stage one and stage two share the same cost model, the sizing decisions carry through to the real-time controller without any reformulation, which keeps the system efficient across varying conditions and reliable in its heat supply.

2.2 Predictive Control and Operational Strategy

Operation planning is becoming progressively more essential for such hybrid energy systems in industry, where multiple systems are integrated, and process heat demand, waste heat availability, and electricity prices are subject to inherent variability. In such systems, decisions made at each time step directly influence future system states, particularly through the state of charge of thermal storage. Therefore, operational planning must be conducted in multiple stages to be both forward-looking and adaptive, ensuring a reliable heat supply while minimizing operational costs.

To address this challenge, a rolling-horizon control framework based on Model Predictive Control (MPC) is employed. In this formulation, system operation is represented as a sequence of short-term intervals and overlapping optimization windows, with feedback controls. At each time step, an optimization problem is solved over a finite prediction horizon, while only the first control action is implemented. The horizon is then advanced, and the optimization is repeated using updated system states and newly available information. This receding-horizon structure enables control decisions to continuously adapt to evolving operating conditions. Within this framework, MPC offers a systematic means of integrating forecasted variables together with operational constraints, allowing informed and forward-looking decision-making under dynamic conditions.

In the present study, machine-learning-based forecasting models are employed to generate short-term predictions of heat demand, waste heat availability, and electricity prices over a 36-hour horizon. We use a gradient-boosting regression approach trained on historical time series (annual datasets) with lagged inputs. The data are split into training and validation sets as usual, and model performance is evaluated using standard error metrics. The resulting forecasts reach a mean absolute percentage error (MAPE) of roughly 2 to 4 percent, which is adequate for the short-term dispatch decisions we need. These forecasts feed the MPC at each step, which then fixes the dispatch of the TCP, OST, and EH under the relevant capacity and storage constraints.

The main advantage of the rolling-horizon MPC is that it can act on expected future conditions rather than the current hour alone. The controller can choose to charge the storage during a low-price window or when waste heat is available in surplus and draw on that stored energy later when demand or price is high. A static or rule-based controller cannot reproduce this behavior because the forward picture of demand and price is precisely the information it does not have. Continuously updating forecasts and system states at each time step allows the operation to remain responsive and robust under realistic uncertainty. At the same time, the proposed framework is inherently scalable, as it can be applied to larger systems without requiring changes to its underlying structure. Taken together, this approach provides a practical and flexible foundation for real-time operational planning, delivering improved economic performance without compromising the reliability of process heat supply.

3. Results and Discussion

The performance of the proposed strategy is divided into two parts: first, to optimize the system sizes; and second, to examine these systems under volatile industrial data. Industrial waste heat, process heat requirements, and dynamic electricity prices are considered three input parameters; these inputs exhibit substantial temporal variability in the Danish context. The waste heat and process heat demand profile depends entirely on production schedules, whereas, especially in the Nordic countries, electricity prices are strongly shaped by a high share of wind generation. Sometimes, with clear intra-day and seasonal fluctuations tied to renewable output and to market coupling effects (e.g., Nord Pool market dynamics), electricity prices can occasionally approach zero or become negative. This motivates shareholders to fully utilize these periods.

The present robust operation planning focuses on how the system performs under such dynamics under coordinated operation planning and on the impact of this coordination on system performance. The result and discussion are structured in three parts: (i) optimal system sizing, (ii) robust operation planning under RH-MPC strategy, and finally (iii) quantitative performance comparison between RH-MPC and static operation.

3.1. System Sizing Outcomes

In order to determine the optimal system sizing, a Genetic Algorithm (GA) is used to determine suitable values for TCP capacity, storage capacity (OST), along with the electricity price threshold $\mathrm{P_{th}}$. In GA, a techno-economic approach is followed, where the levelized cost of heat (LCOH) serves as the objective function. The GA iteratively minimizes the LCOH by considering system performance over an entire year, so that variability in heat demand, waste heat availability, and electricity prices are also captured. Whereas EH is not included as a decision variable in the optimization, as it primarily serves as an auxiliary backup energy system. Therefore, its capacity is fixed to the maximum observed heat demand. However, its operation is still constrained through charging and discharging limits at each time step. The electricity price threshold $\mathrm{P_{th}}$ is optimized within bounds derived directly from the dataset, rather than assuming an arbitrary range. In this work, the lower and upper bounds are defined as 20% and 95% of the maximum electricity price, respectively.

The GA shows stable convergence, terminating when the change in fitness value becomes very small, which means the solution is not improving further. The resulting optimal configuration includes a TCP capacity of **4.0 MW**, an OST capacity of **5.0 MWh**, and an EH capacity fixed at **6.0 MW**. The optimal electricity price threshold is found to be **1613.53 DKK/MWh**. This value reflects realistic operating conditions and avoids bias toward unrealistically low thresholds. From a system perspective, the storage unit sizes seem very small. Still, it can play an important role in managing temporal mismatches between supply and process heat demand by storing excess heat during low-demand periods and releasing it during peak demand. At the same time, the optimized price threshold ensures the electric heater is used only when electricity is economically viable, reducing unnecessary operating costs. The final objective function value is 723.19. Notably, this value remains unchanged after rounding the design variables to practical sizes, indicating that the solution is robust and not sensitive to discretization.

Table 1. *Optimized system design variables, corresponding search bounds, units, and final values obtained from the genetic algorithm.*

Variable	Symbol	Lower & Upper Bound	Unit	Optimized Value
Thermochemical capacity	TCP	$0 - 10$	MW_{th}	4.0
Storage capacity	OST	$0 - 20$	MWh_{th}	5.0
Electric Heater	EH	$-$	MW_{th}	6.0 (max Demand)
Electricity price threshold	P_{th}	$0.2 \times P_{\max} - 0.95 \times P_{\max}$	DKK/MWh	1613.53

3.2. Operation planning and Strategy

After obtaining the optimal system sizes, the system capacities are imported into the RH-MPC framework, where robust operation planning is performed, a 36-hour short-term window has been chosen. In the RH-MPC workflow, a machine learning regression model (Least Squares Boosting) has been used for short-term prediction. In order to train the model, previously used data for sizing was used to generate 36-hour-ahead forecasts. During planning and execution, at each time step, forecasts of heat demand, waste heat availability, and electricity prices are used to determine an optimal dispatch plan; however, only the first hour of operation is implemented. The system states are then updated using real-time data, and the optimization is repeated for the next horizon. This receding-horizon approach allows the control strategy to continuously adjust to forecast deviations and adapt to changing operating conditions.

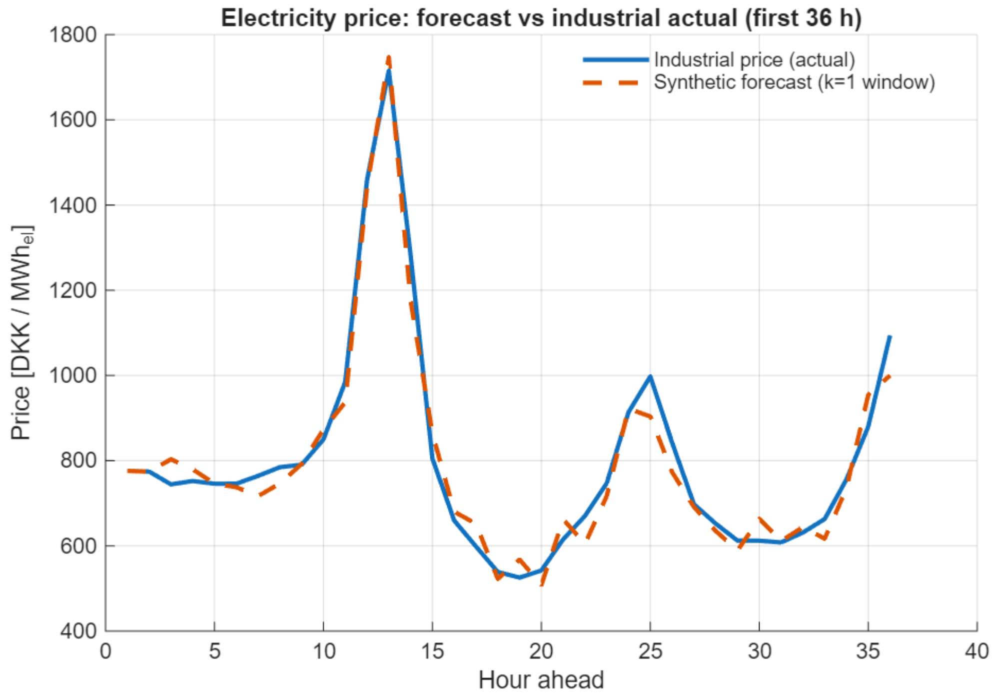

Figure 2. *Comparison of forecasted and actual electricity prices over the 36-hour horizon, illustrating the accuracy of short-term price prediction used for operational planning.*

As can see in the Figure 2, comparison between forecasted and actual electricity prices over the 36 h horizon. The results show that the forecasting model captures the overall trend and major fluctuations and peaks in electricity prices, despite minor deviations. The prediction accuracy, with a mean absolute percentage error (MAPE) of approximately 2–4%, is sufficient to support reliable short-term decision-making within the MPC framework. Accurate price forecasting is critical, as it directly influences the operation of the electric heater and storage charging strategy.

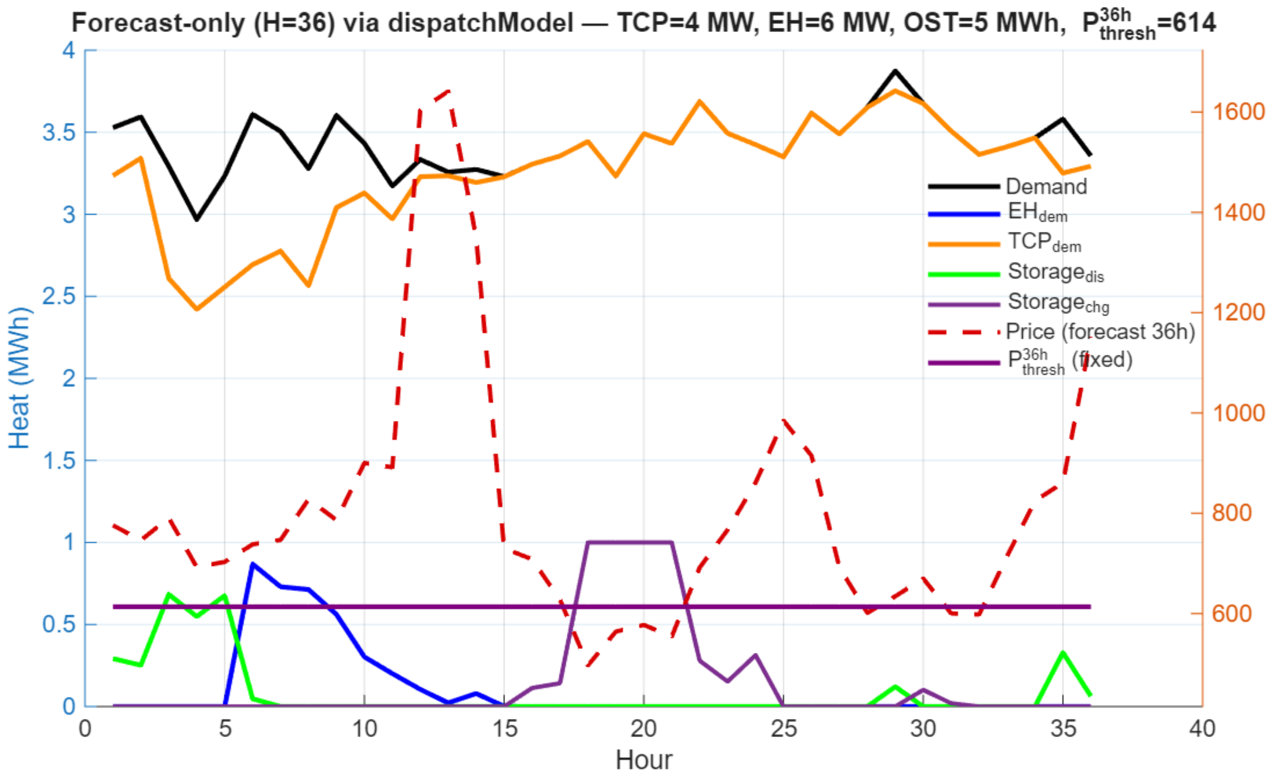

Figure 3. *Forecast-only (open-loop) dispatch results.*

Based on the forecast data, plan dispatch without any feedback correction. As shown in Figure 3, based on the forecast data, a new optimal threshold ($P_{th} = 614\ DKK/MWh$) has been calculated to distinguish between cheap and expensive grid hours. All systems are working together to meet the process heat demand. The general operational pattern is normal as expected; however, the lack of real-time updates leads to suboptimal decisions that can significantly disrupt future execution dynamics. This highlights the limitations of open-loop operation under uncertain conditions.

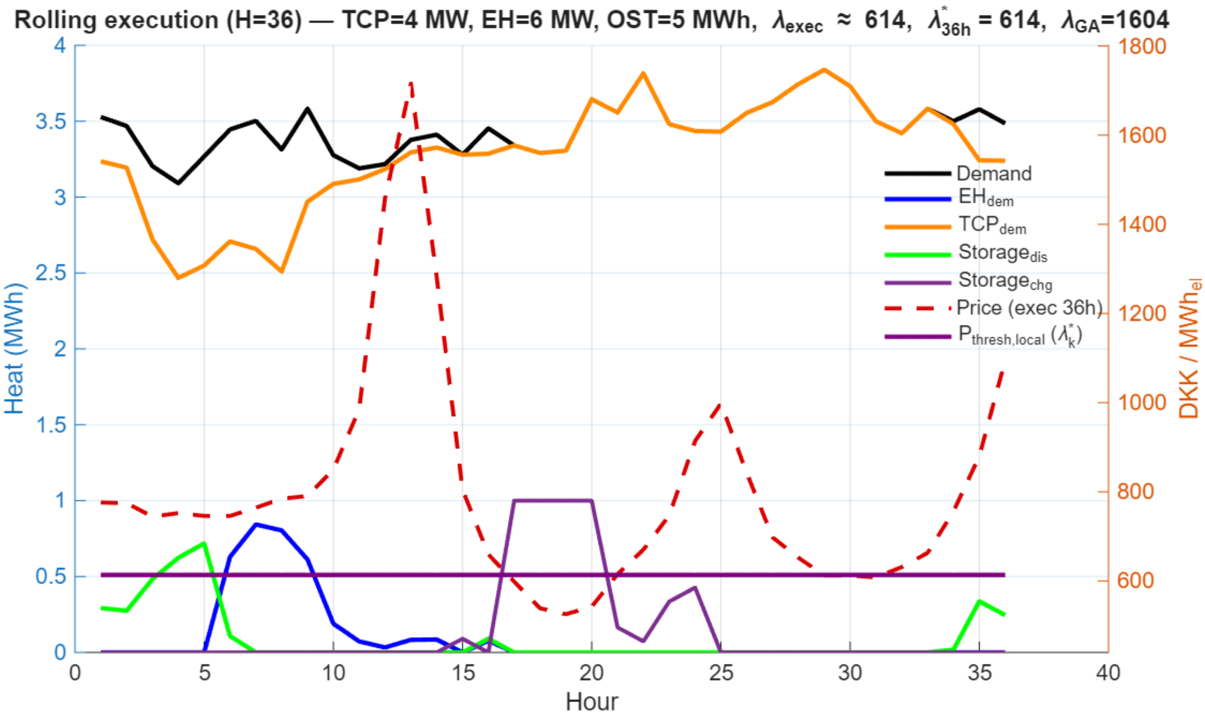

Figure 4. *Rolling-horizon MPC execution results showing system dispatch, including TCP output, EH operation, OST charge/discharge, and electricity price threshold over the 36-hour horizon.*

The system operation under the RH-MPC framework is illustrated in Figure 4. The results show that TCP serves as the primary process heat supplier, meeting the majority of process heat demand. This confirms the maximum utilization and upgradation of WH availability. It is also clear to observe that the EH and OST systems provide additional flexibility to bridge the short-term mismatches between supply and process heat demand. The RH-MPC dynamically adjusts system operation through feedback based on forecasted operating conditions. During periods of low electricity prices, the controller allows the electric heater not only to support demand but also to take advantage of cheap grid power to charge and raise the storage's thermal temperature. On the other hand, during high-price periods, the system prioritizes storage discharge and TCP operation, thereby reducing usage of expensive electricity to minimize the LCOH objective. This behavior exhibits the adaptive and anticipatory nature of MPC.

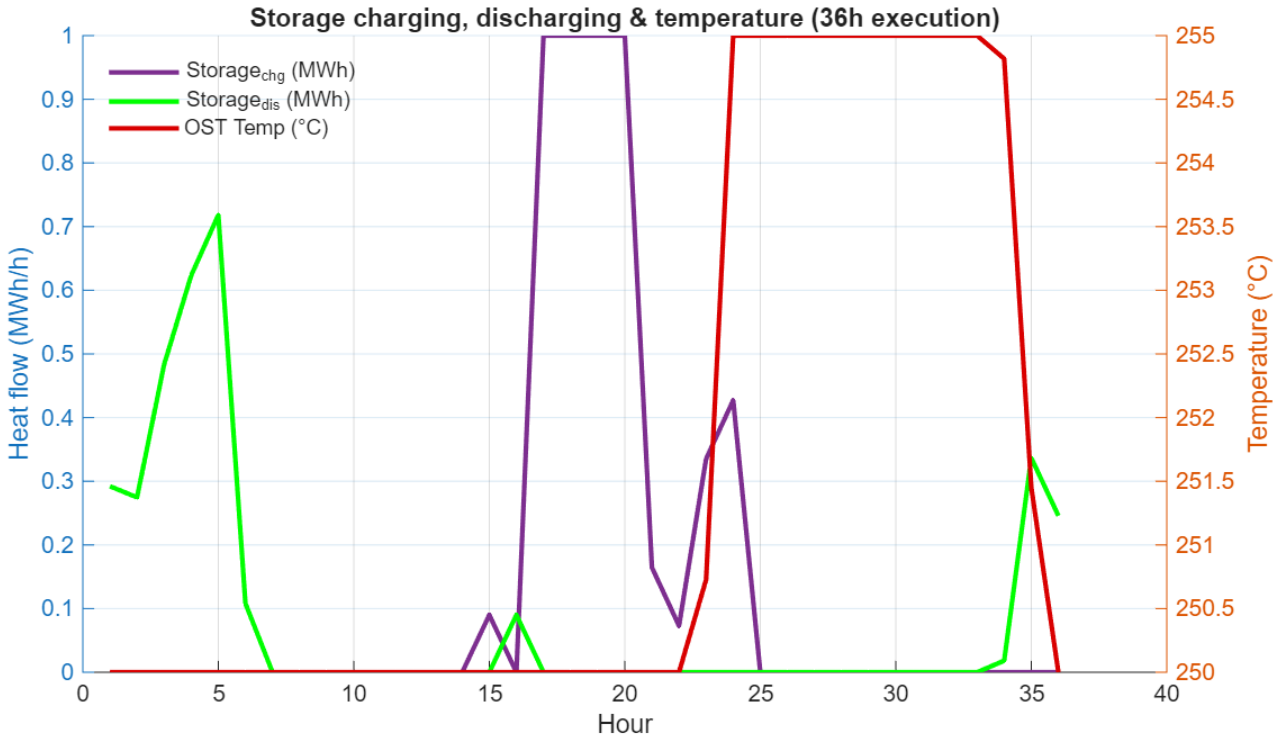

Figure 5. Thermal storage behavior shows charging, discharging, and temperature profiles, demonstrating energy shifting between low- and high-price periods.

Figure 5 illustrates the evolution of the storage state of charge (SOC), with the initial level set at 50% of the total capacity. Under the RH-MPC strategy, the SOC trend looks more consistent and stable. The charging and discharging patterns are clearer, and decisions are better aligned with system conditions. In particular, charging tends to occur when excess thermochemical output is available or when electricity prices are relatively low, while stored energy is later used to support demand and limit grid use during more expensive periods.

For the static case, the behavior is different. Since decisions are based solely on current conditions, the system is less effective at coordinating storage operations over time. This leads to a more irregular SOC profile and reduced ability to shift energy between low- and high-cost periods. Despite these improvements, the overall impact on storage utilization through RH-MPC remains modest. This is largely due to the relatively small storage capacity, which limits the system's ability to shift large amounts of energy. In practical applications, increasing storage size could enhance the benefits of such control strategies.

These observations highlight the advantage of the rolling-horizon MPC approach in adapting to real-time conditions and improving overall system operation. To further quantify these differences, a comparative assessment with the static strategy is presented in the following section.

3.3. Comparative Analysis of RH-MPC and Static Threshold

In this section, the adopted operating strategy, rolling-horizon model predictive control RH_MPC), is compared with a conventional static threshold-based dispatch. In the static case, decisions are made using only information from the current time step, without considering how demand, waste heat availability, or electricity prices may evolve. Whereas, the RH-MPC strategy incorporates a 36-hour forecast of these variables at each step. An optimal dispatch plan is generated, and dispatch over the full horizon, but only the first control action is applied. Then the system is updated with real data, and the process is repeated. This receding-horizon formulation enables the controller to continuously adapt and step-by-step adjust as conditions change.

The performance comparison is summarized in Table 2. The RH-MPC approach reduces the total operating cost from 54,228 DKK to 44,423 DKK, a 18.08% decrease. The same trend is seen in the cost per unit of supplied heat. A reduction in electricity consumption is also observed, suggesting that the MPC framework better coordinates WH use and EH operation more effectively. OST operation also shows a smaller but

noticeable improvement: increased discharge leads to a 5.06% rise in storage utilization, suggesting more effective use of the available thermal storage. In both strategies, the process heat demand is fully satisfied, as no unmet demand is observed.
Overall, the main benefit of the RH-MPC approach lies in its ability to lower operating costs through forward-looking and feedback-based control. The impact on storage utilization is present but limited, which is mainly due to the relatively small storage capacity selected during the sizing stage.

Table 2. *MPC vs. static dispatch performance comparison over a 36h horizon.*

Metric	Static_OpenLoop	RH_MPC	Improvement
Total cost (DKK)	54,228	44,423	**−18.08%**
Cost per demand MWh (DKK/MWh)	439.02	356.90	**−18.71%**
Electricity use (MWh)	6.2356	6.0069	**−3.67%**
EH heat supplied (MWh	3.5837	3.4283	**−4.34%**
Storage discharge (MWh	3.0138	3.1907	**+5.87%**
Storage utilization (% demand)	2.4399	2.5634	**+5.06%**
Unmet demand	0	0	**0%**

4. Conclusion

This study presents an AI-Asset-integrated framework for jointly optimizing the system sizing and operation of a hybrid thermochemical heat upgrade system comprising a thermochemical unit (TCP), an oil storage tank (OST), and an electric heater (EH) under the rolling horizon approach. The system sizing was performed using a Genetic Algorithm (GA), followed by an operational planning strategy based on a rolling-horizon Model Predictive Control (RH-MPC) approach. The GA-based sizing leads to a configuration with moderate thermochemical capacity and comparatively small storage. This suggests that, under the current techno-economic assumptions, the system tends to favor direct use of upgraded waste heat rather than relying heavily on temporal shifting through storage. When this design is operated under the RH-MPC framework, the system responds more flexibly to changing conditions, as dispatch decisions are continuously adjusted using short-term forecasts and updated system states. Compared with the static threshold strategy, the RH-MPC approach leads to a clear reduction in operating cost. The total cost decreases by 18.08%, while electricity consumption is reduced by 3.67%. At the same time, the full heat demand is met throughout the operation. The impact on storage utilization is smaller. Storage use increases by 5.06%, which indicates some improvement, but the effect remains limited. This is mainly due to the relatively small storage capacity used in the system. With a small storage size, the ability to shift energy over longer periods is restricted. Even so, the results show better coordination between storage charging and discharging. This suggests that the control strategy uses the available storage more effectively than the static approach. Overall, the main strength of the RH-MPC strategy lies in improved economic performance. This is achieved through more responsive and coordinated operation of the system components. The results also highlight the importance of storage sizing. Larger storage capacities could further increase the benefits of advanced control methods. Despite this limitation, the proposed framework remains both applicable and scalable. It can be extended to other industrial systems where waste heat recovery and electrified heating are combined. Future work will focus on including forecast uncertainty, evaluating larger storage systems, and analyzing long-term performance under a broader range of operating conditions.

Acknowledgments

This work has been carried out within the framework of the European Union's Horizon Europe program under Grant Agreement No. 101103966 (Thermochemical Heat Recovery and Upgrade for Industrial Processes: TechUPGRADE). The authors gratefully acknowledge the support of the Thermal Energy Section at the Technical University of Denmark (DTU) for providing the necessary infrastructure and research facilities that enabled this study.

References

[1] Z. Su *et al.*, “Opportunities and strategies for multigrade waste heat utilization in various industries: A recent review,” *Energy Convers. Manag.*, vol. 229, Feb. 2021, doi: 10.1016/j.enconman.2020.113769.

[2] R. Agathokleous *et al.*, “Waste heat recovery in the EU industry and proposed new technologies,” *Energy Procedia*, vol. 161, pp. 489–496, 2019, doi: 10.1016/j.egypro.2019.02.064.

[3] H. R. Rahbari and A. Arabkoohsar, “Hydration heat transformer: A groundbreaking technology for sustainable process heating,” *Energy*, vol. 337, p. 138757, Nov. 2025, doi: 10.1016/J.ENERGY.2025.138757.

[4] H. R. Rahbari, B. Elmegaard, E. Bellos, C. Tzivanidis, and A. Arabkoohsar, “Thermochemical technologies for industrial waste heat recovery: A comprehensive review,” *Renewable and Sustainable Energy Reviews*, vol. 215, p. 115598, Jun. 2025, doi: 10.1016/J.RSER.2025.115598.

[5] H. Jouhara, N. Khordehgah, S. Almahmoud, B. Delpech, A. Chauhan, and S. A. Tassou, “Waste heat recovery technologies and applications,” *Thermal Science and Engineering Progress*, vol. 6, pp. 268–289, Jun. 2018, doi: 10.1016/J.TSEP.2018.04.017.

[6] “Nord Pool.” Accessed: Apr. 19, 2026. [Online]. Available: https://www.nordpoolgroup.com

[7] J. J. Aguilera, R. Padullés, W. Meesenburg, W. B. Markussen, B. Zühlsdorf, and B. Elmegaard, “Operation optimization in large-scale heat pump systems: A scheduling framework integrating digital twin modelling, demand forecasting, and MILP,” *Appl. Energy*, vol. 376, p. 124259, Dec. 2024, doi: 10.1016/J.APENERGY.2024.124259.

[8] M. Sadi and A. Arabkoohsar, “Advanced AI technique for optimizing the operation strategy of an industrial heat pump coupled with thermal energy storage,” *J. Energy Storage*, vol. 146, p. 119877, Feb. 2026, doi: 10.1016/J.EST.2025.119877.