UPPSALA UNIVERSITET

UPTEC F 23015

Examensarbete 30 hp

May 2023

Improved Furnace Control

System identification and model predicative control of Outokumpu’s reheating furnace

Oscar Holmqvist

Civilingenjörsprogrammet i teknisk fysik

UPPSALA
UNIVERSITET

Improved Furnace Control

Oscar Holmqvist

Abstract

This thesis investigates one option for improving the control of a reheating furnace used in heating steel slabs before hot rolling; an essential part of the steel manufacturing process. The furnace consumes a significant amount of energy, leading to high cost and high carbon dioxide emissions. The proposed solution is the implementation of a model predictive control (MPC) system to improve control and reduce fuel usage. The MPC system will be based on the use of system identification techniques to find a prediction model of the furnace, specifically using ARMAX models. An additional simulation model will be used to simulate the system, and to compare the performance of MPC and PID. The prediction model is found to have a normalized root mean squared error of over 91% for the first five minutes, suggesting that it has potential to be used for MPC. The simulation model has significant inaccuracies, due to the presence of unmeasured disturbances. The simulation results, although limited due to the inaccuracies of the simulation model, suggest that MPC is a viable option for improved control of the furnace. The use of MPC can potentially improve the repeatability of the heating process, resulting in improved steel quality and reduced defects. This thesis suggests that further investigation into the use of MPC for controlling reheating furnaces in the steel industry is worth pursuing, and could potentially bring significant benefits to both producers and the environment.

Teknisk-naturvetenskapliga fakulteten

Uppsala universitet, Utgivningsort Uppsala/Visby

Handledare: Per Åhman Ämnesgranskare: Mikael Sternad

Examinator: Tomas Nyberg

Populärvetenskaplig Sammanfattning

Stålindustrin står för en stor andel av Sveriges energiförbrukning. Ett steg i ståltillverkning är varmvalsning, vilket valsar rätblock av stål med hög temperatur till tunnare rullar. Denna process medför krav på stålets temperatur, vilket kontrolleras med en uppvärmningsugn. Temperaturen på ugnen kontrolleras genom att höja eller sänka gasflödet, och temperaturen på ugnen påverkar i sin tur temperaturen på stålet. I nuläget så används en uppsättning PID-regulatorer för att kontrollera temperaturen på ugnen. Mer specifikt så försöker regulatorn minimera skillnaden mellan den uppmätta ugnstemperaturen och den önskade temperaturen. Prestationen på regulatorn avgörs genom att bestämma ett antal parametrar, vilket är ett tidskrävande arbete utan ett entydigt svar. Denna uppgift görs desto svårare då det finns ett flertal begränsningar på ugnens beteende, till exempel den maximala temperaturen i ugnen eller hur snabbt temperaturen får öka.

Syftet med examensarbetet är att utveckla ett alternativ till den nuvarande kontrollmetoden. Mer specifikt kommer arbetet att vara centrerat kring en MPC-regulator. Denna metod bygger på en modell av ugnen, vilket kan användas för att approximera den framtida temperaturen givet hur mycket gasflöde man tillåter. MPC-regulatorn väljer ett gasflöde som ger det önskade beteendet på ugnen, vilket möjliggörs av systemmodellens förmåga att förutspå ugnens framtida temperatur. Denna metod skapar dessutom en möjlighet att inkludera en modell av stålets temperatur, så istället för att kontrollera temperaturen på ugnen så kan stålets temperatur direkt kontrolleras. MPC har även en förmåga att undvika att bryta begränsningar i systemets beteende.

Genom att använda mer effektiv kontroll så kan energiförbrukningen minska, och kvaliteten på stålet kan förbättras. Resultaten från detta examensarbete är begränsade till simulering, vilket begränsar jämförelsen till simuleringens noggrannhet. Resultatet pekar dock på att MPC har potential för att förbättra kontrollen av ugnen, och bör undersökas i större utsträckning.

Contents

1 Introduction **5**

2 System Identification **7**

2.1 Furnace model 7

2.2 Slab Model 14

3 Model Predictive Control **16**

4 Results **19**

4.1 Prediction 20

4.2 Simulation 24

4.3 Control 26

5 Limitations and Further Research **37**

6 Conclusion **38**

1 Introduction

The Swedish steel industry produced 4.7 million tons of steel in 2021, and uses 15 TWh of energy a year — about 15% of Sweden's total energy consumption [1] [2]. The industry produces approximately 6 million tons of carbon dioxide a year [3]. A key section in the steel manufacturing process is hot rolling, a process that takes steel slabs and rolls them into steel coils. Due to the thickness of the steel slabs it is necessary to soften the steel by heating it before rolling, which is done in a reheating furnace. The slab heating process takes a long time, and consumes a significant amount of fuel. The process is further complicated by the difference in temperature between the slabs and the furnace, which significantly decreases the temperature of the furnace. To heat the furnace multiple burners and sensors are used. Currently, the furnace is controlled by several individual proportional-integral-derivative (PID) controllers , one for each of the zones in the furnace. However, the interdependencies between different zones limits the effectiveness of this approach. By developing a control system that takes these interdependencies into account, the system can potentially be controlled more optimally. The desired effect of improving control is the reduction of fuel usage, which would benefit both the producers and the environment. Furthermore, better control and better repeatability of the process improves quality and leads to less defects in the steel. In some high alloy steels, reheating defects is a major reason for scraping.

This thesis will focus on Outukompu's reheating furnace in Avesta, Sweden. The system consists of eight zones, each with sensors and a burner. Slabs are inserted into the furnace, and then incrementally moved forward using walking beams. The furnace is divided into 3 sections, as shown in the image in Figure 1 such that the slabs enter into the section corresponding to zone 1 and 2, and exit in the section corresponding to zone 5, 6, 7, and 8. The interior of the furnace is geometrically simple; it has the shape of a rectangular box. This means that heat can transfer between the zones, and that adjacent zones should have some level of impact on each other. This can happen in several ways. A significant amount of the heat transfer within the furnace happens by radiation, as the walls radiate heat onto the slab.

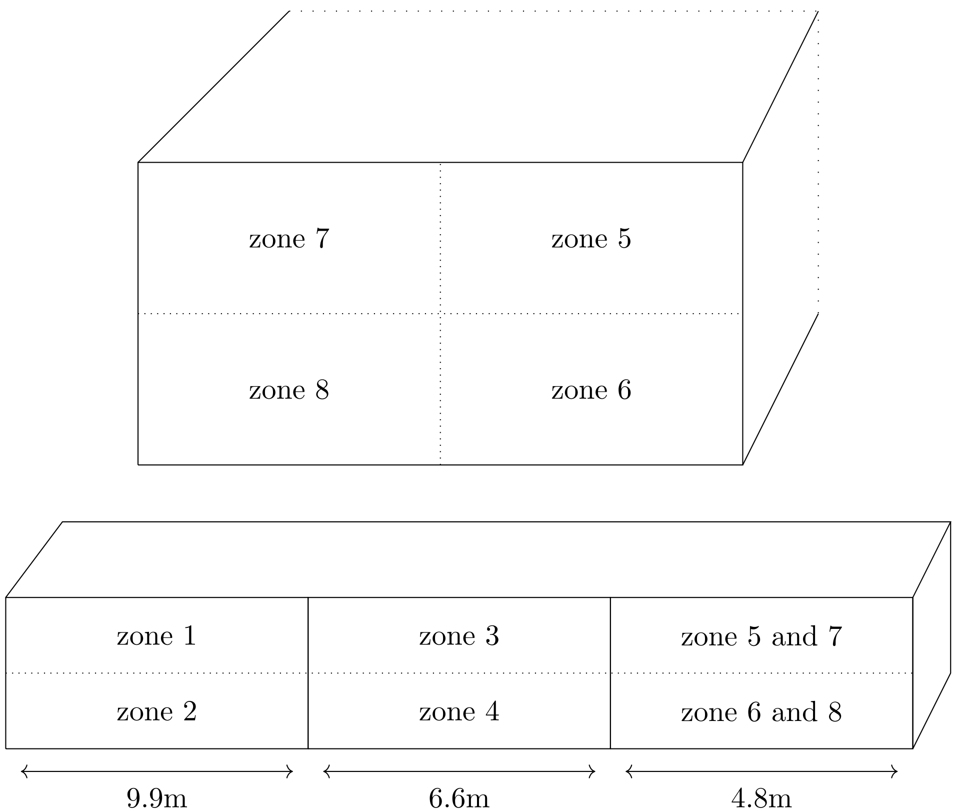

Figure 1: Back view (upper image) and side-view (lower image) of the furnace interior.

The aim of the heating process is to control the temperature of the steel slabs due for hot rolling. The slabs can be of varied thickness. Slabs of different thicknesses are naturally expected to have different heating rates. This is especially significant for the core temperature, which will heat more slowly for thicker slabs. As the temperature of the slabs is not measured after entering the furnace, a model is needed to approximate the slab temperature during the heating process. While there are sophisticated models used to calculate the heat distribution inside the slab, they are heavily integrated in the current system and not suitable for control optimization. Instead this thesis will limit the scope to implementing and using a simplified model that focuses on the average temperature of a slab, with a thickness of around 140 mm.

This thesis aims to use system data to identify a model that describes the furnace input-output characteristics, in order to use it for prediction and simulation. Furthermore, this thesis aims to use this model for model predictive control (MPC) of the furnace, with the goal of optimizing control and lowering fuel usage while maintaining the desired temperature of the slabs exiting the furnace. Finally, the accuracy of the model and the performance of the MPC will be discussed and compared to a PID-controller in a simulated environment.

2 System Identification

This thesis will make use of system identification to find two models: one prediction model for the purpose of designing an MPC-controller, and one simulation model to compare the performance of the MPC-controller and the PID-controller. The latter is the currently used type of controller. This section is concerned with finding a prediction model and a simulation model that accurately describe the furnace temperature dynamics. Furthermore, a simple slab model will be proposed. Much of the knowledge and assumptions regarding the furnace is based on verbal communication with operators and Prevas, the company responsible for the current control system .

2.1 Furnace model

The goal of system identification is to find a model that accurately describes the input-output characteristics of a system. As described by Ljung [4], the system identification process has three key steps. First of all, some set of input-output data should be recorded. Designing experiments to be “maximally informative”, meaning that the data is sufficiently ”rich” in frequencies to distinguish between two models in a model set. However, designing experiments was not possible for this thesis, as the furnace is unavailable for experimentation. Therefore, the data used for identifying the system comes solely from normal production operation of the system, which is closed loop data. An example of the training data is shown in Figure 2. Using closed loop data can potentially negatively affect the model accuracy. The closed loop data is informative if and only if the reference value r is persistently exciting. Persistence of excitation can be understood as the number of harmonics in a signal, and determines the order a model can have while remaining distinguishable from other models in the model set. This is the case for the given dataset, which can be shown in Matlab with the *pexcit* function.

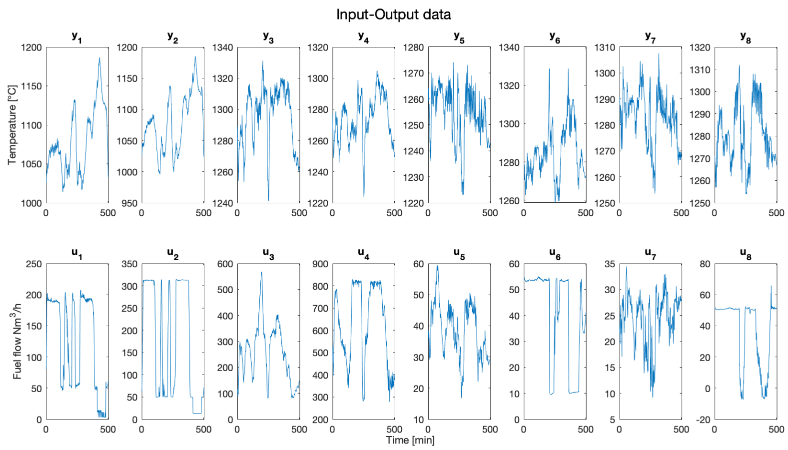

Figure 2: Example of furace input-output data. The top row shows the temperature signal for the different zones in thee furace, where signal y_i corresponds to the ith zone. Similarly, the bottom row shows the fuel flow for each section.

Second, model structure has to chosen. Here *a priori* knowledge of the system can be used to pick the model structure. With a very extensive understanding of the mechanisms of a system, a mathematical model can be designed. As the physical knowledge of this system is limited, the thesis will focus on black-box models. Black box, or parametric models work by optimizing a set of parameters to maximize fit between the model prediction and the input-output data. The set of all possible models with a given model structure constitutes the model set. The model parameters distinguish the different models in the model set, so the final step is to pick the best, or at least sufficient, model in the model set. This step involves finding hyper-parameters that give a good result, as well as optimizing the model. This means using the input-output data to estimate the model parameters.

Much of this system behavior is unknown, but the geometry and layout of the system is known. The furnace consists of eight zones, which can be divided up into three heating sections. Around the center of each zone there is a temperature sensor, and a heating flame. The measured temperature in each zone constitutes the output signal of the system, and the fuel flow, measured in units of normal cubic meters per hour (Nm^3/h), is the system input. The furnace is relatively narrow with regards to the slabs. The size of the furnace entails that multiple slabs are expected to be in each zone simultaneously. The flame heats the slabs directly by convection, as well as indirectly by radiation from the furnace walls. According to Prevas, the radiation from the walls account for 70-74% of the slab heating, and the atmosphere for another 20-25%. The convection from the flames gives less than 5% of the heat transferred to the slabs, i.e. the main purpose of the fuel burners is to heat the furnace, not the slabs. For this thesis, the slab temperature is assumed to be only dependent on the furnace temperature, not the amount of fuel used.

There are some behaviors that the system is expected to have. As heat flows from higher to lower temperature, the change in temperature for a zone should be increasing with the difference in temperature between the zone and a neighbouring zone. Furthermore, as heat is added by the flame, the temperature should increase. This knowledge can be incorporated in a physical model. In the process of searching for appropriate model structures, this was attempted — with no better results than other models.

There are likely several disturbances and characteristics of the system that are unmeasured and unknown. There are three known disturbances that are likely to have a significant effect on the system. The first one is the heat transfer from the furnace to the slabs. As the temperature of the slabs in the furnace are unmeasured, this is difficult to include in a system model. For a sufficiently accurate slab model the approximated value could be used as an input to the system, but this comes at the risk of an error feedback between the two models. It is simply not desirable to base the identified model on approximated data, which to some extent is inaccurate. The second one is the flow of gas inside the furnace, in combination with cold air from slabs entering or exiting the furnace. The flow of gas is not necessarily a disturbance, as it is inherent to the system. However, it might be too nonlinear for the model to accurately be able to describe it from the input-output data. The third disturbance arises in the cooling of the step-beams, which is done with water. This is necessary to prevent the beams from melting or being deformed, but it comes at the cost of heat being removed from the furnace. It affects both the slabs and the furnace, but to what extent is unknown. What all disturbances have in common is that they are unmeasured, and the characteristics of the disturbances are unknown.

The beginning of the furnace — the first half of zone 1 and 2 — is referred to as a dark zone, meaning that the temperature is neither directly measured nor controlled. The dark zone is by design heated by hot gas flowing backwards in the furnace. In addition to the heat supplied by hot air flowing backwards in the furnace, cold air enters when a new slab enters the furnace. The slab itself also has a significant effect on the temperature of the furnace, especially in the first section where it is usually much cooler than the furnace. More slabs also causes the inside of the furnace to have a higher heat capacity, essentially meaning that the rate of increase in temperature for a given amount of fuel will be lower. The presence of these disturbances in the first section makes the system model less accurate.

Given the limited understanding of the system itself and the large amount of input-output data available, a black box model is a suitable approach to modeling this system. Black-box models utilize input-output data in order to find the model parameters that minimizes the prediction error. The optimization of the parameters focuses on minimizing some loss function that describes the fit between the model prediction and the training data. A commonly used loss function is the mean-squared-error (MSE) of the prediction and the training data. Additionally the number of prediction steps on which to calculate the loss must be chosen. One approach is to minimize the one-step-ahead prediction. This approach aims to give accurate predictions for the next time-step. Another approach is the focus on simulation loss. This loss function instead uses the difference between the measured output and the model output over the entire training set. This puts less emphasis on accurately capturing short term behavior, but it can give the model stability and more accurate steady state values. As the name suggests, the latter approach is preferred for simulation of the system, while the former is preferred for model predictive control.

This thesis will focus mainly on the use of autoregressive moving averages with exogenous inputs (ARMAX). A linear ARMAX model uses a weighted linear combination of past input, outputs, and prediction residuals to predict future states of a system, as seen in Equation (1). The autoregressive part consists of a linear combination of past inputs from $t-1, t-2, ...t-n_a$. Similarly the exogenous inputs consist of a linear combination of past input from $t-n_k, ...t-n_t-n_b$, where n_k is the delay. These two parts constitute a linear ARX model, which can be written as a linear combination of the input output, by using the q operator. The q operator is defined as $q^{-n}y(t) = y(t-n)$, i.e. the value of the previous sample. What separates ARMAX from ARX is the use of residuals, i.e. past prediction errors, to make predictions about future states. The ARMAX includes moving averages, which is a linear combination of residuals from $t-1, ...t-n_c$. With the residuals, the predictions from the ARMAX cannot be written as a linear combination of input and outputs; the residual term has to be included. The residuals are not used for simulation, due to the absence of reference values.

$$y_i(t) = \sum_{j=1}^{n_y}\sum_{k=1}^{n_a} a_{jk}y_j(t-k) - \sum_{j=1}^{n_u}\sum_{k=n_k}^{n_b} b_{jk}u_j(t-k) - \sum_{j=1}^{n_y}\sum_{k=1}^{n_c} c_{jk}e_j(t-k) \tag{1}$$

Where n_y and n_u are the number of inputs and outputs, respectivly. Each y_i correspond to one of the output signals. The complete model can be written as a

$$\begin{aligned} A(q)y(t) &= B(q)u(t) + C(q)e(t) \\ A(q) &= I + q^{-1}A_1 + q^{-2}A_2 ... + q^{-n_a}A_{n_a} \\ B(q) &= q^{-nk}B_1 + q^{-n_k-1}B_2 ... + q^{-n_k-n_c}B_{n_b} \\ C(q) &= I + q^{-1}C_1 + q^{-2}C_2 ... + q^{-n_c}C_{n_c} \end{aligned} \tag{2}$$

Where $A_i \in \mathbb{R}^{ny \times ny}$, $B_i \in \mathbb{R}^{ny \times nu}$, and $C_i \in \mathbb{R}^{ny}$ are the weight matrices corresponding to q^{-i}, and T_s is the sample time. The ARMAX model structure outlined in Equation (1) model is further illustrated in Figure 3.

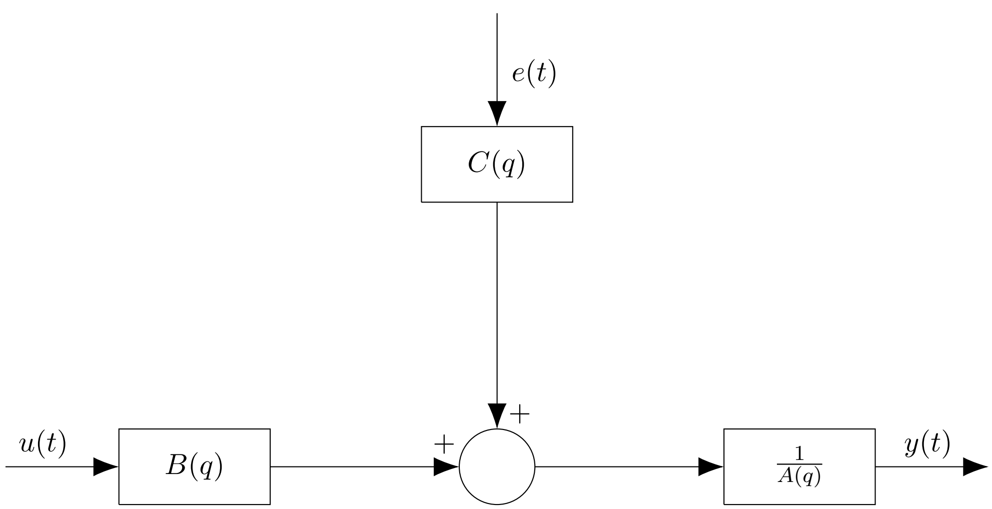

Figure 3: Block diagram of an ARMAX system.

The matrices $A(q)$, $B(q)$, and $C(q)$ are found by minimizing the prediction or simulation error with respect to the dataset. For a more thorough explanation, see chapter 7.2 in Ljung [4].

Training the ARMAX on unprocessed data might result in sub-optimal performance. The raw data should be preprocessed before it is used to train a model. There are several things that should be done before the data is ready for training. Faulty measurements due to sensor errors can give data that does not represent the behavior of the real system. This error can sometimes cause outliers in the dataset. Outliers that are likely to be caused by sensor error should be removed from the dataset. This process demands some *a priori* knowledge of the system behavior and what behaviour is reasonable. Here, large jumps in temperature are considered outliers.

The dataset contains missing data values. One option to deal with this is to simply disregard this sections with missing data. However, this comes at the cost of splitting the dataset into several smaller sets. Matlab's System Identification toolbox requires all experiments to be of the same size, meaning that this approach would lead to disregarding a lot of useful data points as the data subset size is determined by the smallest sequence without missing values. Another option is to estimate the missing data points using some algorithm. Here the missing data points have been estimated using a stencil operation, such that it is a weighted average of the adjacent data points. Sections with a large number of missing datapoints has been excluded, as the stencil estimate is unreliable for more than a few values.

The ARMAX model should be trained on data with an average value of 0, as pointed out by Ljung [4]. This means that the mean value of each sensor measurement should be subtracted from every datapoint in the dataset. The presence of feedback control and disturbances makes it difficult to find steady states of the system. However, if an input signal centered around some average causes the output signal to be centered around some output mean value, it is reasonable to infer that the value is a steady state for an input signal equal to the input average. Furthermore, the approximation of local linearity makes model prediction of values far from the operating state less accurate. It is of less interest that the model predicts a room temperature steady state with no input, and more important that it remains stable and accurate around the system's working state.

Models with higher complexity, e.g. nonlinear models, are increasingly able to capture more complex system behavior. A model with low complexity tends to be biased towards dominant features in the dataset. However, if the model is too complex it might overfit the dataset and be very sensitive to noise, as well as being unable to generalize well. In general, more complex models can be used for larger datasets. The model should balance these two sources of error.

Inaccuracies in the ARMAX model can arise in system disturbances, as well as nonlinearities that the model is unable to capture. The degree to which nonlinearities are present in the data affects the performance of the ARMAX model. If significant nonliniarities is present in the data, a nonlinear model such as the nonlinear ARX (NLARX) should be used instead. There are algorithms that can determine the level of nonlinearities in the data, more precisely the algorithm tests to what extent a NLARX would perform better than a linear ARX. The detection ratio indicates the amount of nonlinearities that are present in the system. This can indicate if an ARMAX model will be able to accurately capture the system behavior. A value under 0.5 indicates very limited presence of nonlinearities, a value over 2 indicates very significant nonlinearity, and a value close to 1 means that the test is unreliable and weak nonlinearities may be present. See Ljung [4] for further detail. In this case, the presence of nonlinearities in the system is limited, with values close to or smaller than 0.5, according to Matlab's function *isnlarx*. However, there are non-linearities present in zone 6 and 8, the two diagonally opposite zones in the last section. While the detection ratio does not indicate a very significant amount of nonlinearities in these zones, it is still expected to be a source of error for the ARMAX model. However, the amount of nonlinearities present does not strongly suggest that a nonlinear ARX should be used instead of the ARMAX.

Regularization is a technique that adds a penalty to high values in the model, for an ARMAX model values in A, B, or C. Using quadratic regularization, or L_2-regularization, increases the penalty quadratically. The amount of regularization is determined by the parameter λ. Regularization forces the model to not be over reliant on certain features, and can thus prevent non-realistic behavior that still gives good accuracy.

The above mentioned hyperparameters must be tuned to ensure good model performance. Much of the time for this thesis has been spent on finding appropriate hyperparameters, especially for the ARMAX model structure. The structure that best combines prediction or simulation accuracy, and expected behavior is an ARMAX with input and output signals being communicated between every zone, and $n_a = n_b = n_c = 4$ and input delay $n_k = 6$. For regularization the parameter $\lambda = 20$ is used.

2.2 Slab Model

The loss function of the MPC is based on the predicted slab temperature. Since the temperature is not measured after the slab enters the furnace, a model is needed to predict it. Developing an accurate model is outside the scope of this thesis. The model used in this thesis is a simple first order model. Equation (3) shows the equation used for approximating the average slab temperature, where α is a constant estimated to maximize the fit between the slab model and the training data. Note that the data used to estimate the heating coefficient is itself an approximation from the 2DSTEELTEMP model, although this model is significantly more sophisticated [5]. The reason for not using the 2DSTEELTEMP model is twofold. First of all, it is not linear. This means that the prediction can not be written as a set of linear equality constraints. Furthermore, the current implementation limits the usage to simulation, and can not be used as a predictor. Improving the slab heating model deserves more attention and is an area of interest for further research. While it is beneficial to have a simplified prediction model, a more sophisticated model should be used in parallel to provide feedback after every time-step. However, the performance of the slab model is sufficient for simulation and prediction. Figure 4 shows the performance of the slab model compared to the reference value. With an NRMSE of 80%, it is accurate enough to give some insight into how well MPC and PID works.

$$y_{t+1}^{slab} = y_t^{slab} + \alpha \left(y_t^{furnace} - y_t^{slab} \right) \tag{3}$$

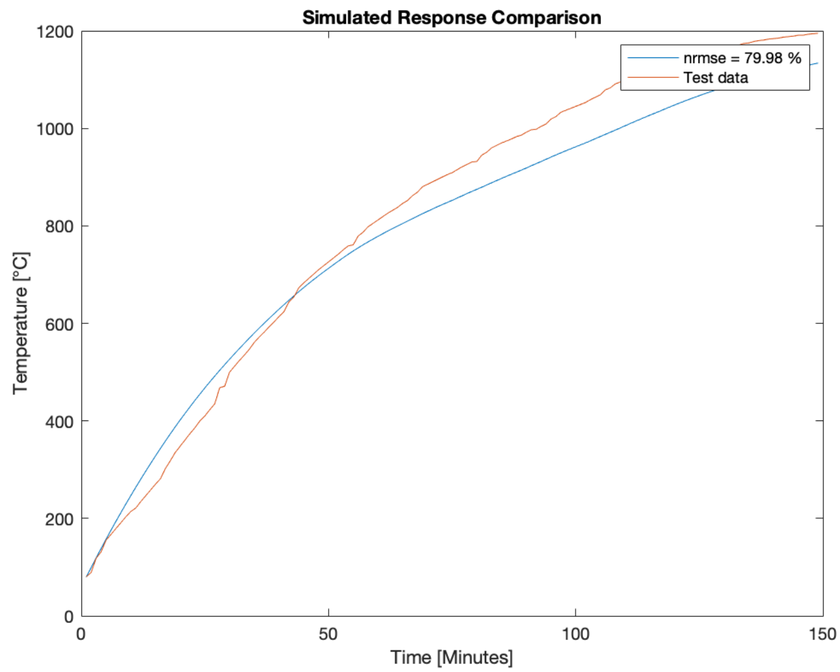

Figure 4: Comparison of slab heating model and reference data.

The temperature of a slab is a function of the furnace temperature. Likewise, the furnace temperature is dependent on the temperature of the slab. When the slab enters the furnace, it increases the heat capacity of the system, but since the slab usually has a much lower temperature it adds relatively little heat to the system. The heating of the slab means an equal amount of heat is removed from the gas and the walls of the furnace, thus lowering the temperature. It would be beneficial to the prediction and simulation models to have constant measurements of the slab temperature. Unfortunately, this is not available. The estimated slab temperature could be used, but using an estimated input for the furnace in addition to the estimated input to the slab could cause an error feedback loop. The only measurement of the slab is right before it enters the furnace. This should mostly only affect the first section of the furnace, but it is also there that the slab should have the largest cooling effect on the furnace. The rate of decrease in temperature should be proportional to the mass of the slab, and to the difference in temperature between the furnace and the slab. Therefore including this difference as an input signal has the potential to supply the model with further relevant information about the system behavior.

3 Model Predictive Control

A system model can be used for model predictive control. MPC uses a model to estimate future system outputs based on future inputs and past system data. Furthermore, MPC finds an input that minimizes a given loss function. The loss function is based on the predicted system responding to the input signal in the time-steps within the prediction horizon. The number of time-steps the model predicts forward is referred to as the prediction horizon, whereas the number of time-steps for which the input is optimized is called the control horizon. At each time-step t the future inputs are optimized up until $t + \tau$, then the input for the time $u(t + 1)$ is used as the next input signal to the system. Here τ_c refers to the control horizon, and τ_p refers to the prediction horizon. In this thesis $\tau_c = \tau_p$ is used, and will be referred to simply as τ

In order to highlight eventual improvements of the model, the old implementation must be considered. Right now the system is controlled by eight PID-controllers. Tuning the controllers for a complex MIMO system is not straightforward. In the current implementation, the values of the PID-controller are set in such a way that the constraints should not be broken. In addition to the practical difficulties in setting the values to achieve this for several controllers in a very interconnected system; avoiding any constraint violation comes at the cost of sub-optimal fuel usage.

MPC has the potential to improve on many of the shortfalls of PID-control. While a PID-controller has to be manually tuned to avoid violation of the constraints, they can be explicitly included in the model predictive controller. This is of increasing importance as the set of constraints increase in size. Additionally, having a model that accurately describes the system enables the chosen input-signal to be chosen by optimizing some loss function. The loss function can also include a term set to penalize fuel usage, with the aim of decreasing the total amount of fuel used.

MPC can efficiently control (MIMO) systems, for which an accurate proportional-integral-derivative (PID) controller might be hard to implement, especially with multiple constraints that needs to be taken into account. It is also able to optimize the input with respect to a user-designed loss function. In the case of the reheating furnace, the difference between the target and the measured output signal does not need to constitute the loss function. Additionally, MPC has the benefit of allowing for constraints in the input and output values of the system. Since the input is optimized for every time-step, giving the desired constraints to the optimiser ensures that e.g. the maximum furnace temperature is never exceeded. MPC is an optimization problem. The problem includes the model, the constraints, and the loss function. It is beneficial to have linear constraints and a linear model, as it allows for the utilization of more efficient optimization algorithms. The MPC optimization problems with a linear model and constraints can be written as (4).

$$\begin{aligned} \min_{x} \quad & (r-x)^T H (r-x) + \mu x^T F x \\ \text{s.t.} \quad & A_{eq} x = b_{eq} \\ & Ax \leq b \\ & lb \leq x \leq ub \end{aligned} \tag{4}$$

In Equation (4) the reference signal r is a vector, where the elements corresponding to y_{slab} are used as a reference signal that sets the target for the controller. This process is illustrated in figure 5. For this thesis, r_k refers to the slab target temperature. The MPC calculates the input signal u for all time steps $t+1.., t+\tau$, but only actually uses $u(t+1)$. It then receives feedback, for which the optimization of $u(t+1)$ is used.

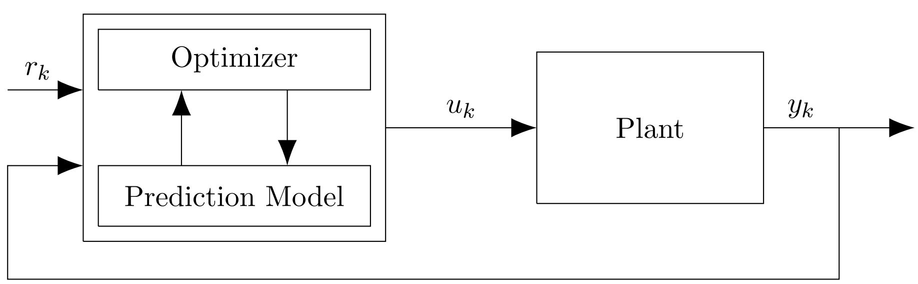

Figure 5: Block diagram of an MPC-controller.

In Equation (4), x is a vector including the furnace temperature, the fuel used, and the slab temperature, $x = [y_{furnace}(t+1), ...y_{furnace}(t+\tau), u(t+1), ...u(t+\tau), y_{slab}(t+1), ...y_{slab}(t+\tau)$. Using this notation allows the model predictions to be written as a set of linear constraints, where A_{eq} and b_{eq} determine the linear constraints imposed by the model predictions for each time step, such that $Ay(t)) = Bu(t) + Ce(t)$. There are no residuals for simulation, i.e. $e(t) = 0$. The two terms only leave the future input signal as free variables to optimize. The loss function consists of two matrices, H and F. H is determined by the difference between the desired temperature of the slabs from $t+1$ to $t+\tau$. Contrasting a PID-controller, the MPC allows for every slab to be taken into consideration when determining the future input signal. F penalizes the amplitude of the input signal, in order to decrease fuel consumption. The penalty coefficient μ should be set such that it lowers fuel consumption without being large enough to significantly impact the final temperature of the slabs. The level of acceptable impact must be determined by the user. A comparison of different values for μ is included in Section 4.3. A clear advantage of MPC on this form is that the loss function is quadratic, which allows for more efficient optimization.

There are additional demands on the performance of the controller. If the furnace gets too hot the slabs can melt or become deformed, so there is a maximum allowed temperature. In Equation (4), the upper and lower bounds for all variables are set by ub and lb, respectively. The burners have different effects, so the maximum possible fuel flow varies. The minimum is naturally set at zero. There is also a maximum allowed rate of temperature increase for the slabs. This rate of increase is set to not exceed 4 °C a minute, in order to avoid material tensions from the heat gradient in the slab. All of these additional constraints can be expressed as the inequality constraints A and b, making them easy to incorporate into existing optimizers. There are no nonlinear constraints present in the controller.

The choice of optimization algorithm is important, not only for the computational efficiency, but also for the result. Using an interior point algorithm adds error from the barrier method when the solution is close to any of the constraints. This is not desirable, as the MPC will likely find a solution close to the constraints, especially the constraint on maximum increase in temperature per timstep. The error from using an interior point algorithm negatively impacts the accuracy model prediction, and should not be used. Sequential quadratic programming offers a better solution, due to its higher accuracy for solutions close to the constraints.

Since the sample and control time of the system is one minute, there are no significant requirements on the speed of the control. One minute is more than enough time for the optimization problem to be solved. This of course depends on the model size and the control horizon. It might also be preferable to use a shorter τ, due to the drop in accuracy over time, shown in

Figure 8. Having higher accuracy signifies that the predicted behavior is closer to the real behavior, and that the control signal will have the desired effect. However, the prediction horizon has to be long enough for the next input to actually affect the system output. In other words, the prediction horizon should be greater than n_k.

The choice of loss function or criterion is essential for achieving the desired performance. In contrast to a PID-controller, MPC uses optimization based controls. The loss function can contain terms explicitly dependent on the output signal, such as the difference between a reference signal and the measured output signal. It can also have terms indirectly dependent on the output signal. In this case, the slab temperature constitutes such a term. The loss function used in this paper is solely dependent on the slab temperature and the amount of fuel used. Controlling the slab temperature rather than the furnace temperature will hopefully add additional flexibility to the controller, which can enable it to decrease the fuel usage. However, this comes at the cost of potential errors in the slab model accumulating. Additionally, an input penalty term can be added to the loss function. The impact of this term is decided by changing the parameter μ. Decreased fuel usage can have adverse effects on the controllers ability to control the slab temperature. Therefore the value of μ should be set by the user such that it does not impact the slab-temperature beyond some accepted tolerance. Given acceptable controller performance it is desirable to penalize the fuel usage to the greatest extent possible, both for environmental and economic reasons.

4 Results

The results are divided up into two sections: the results from the system identification process and the result from MPC control of the furnace. The former is focused on presenting the accuracy and behavior of the ARMAX model, as well as its limitations. The latter revolves around the extent to which MPC can control the system in the simulation environment and the difference between the simulation and the real system. This section will also include a comparison between an MPC-controller and a PID-controller. As PID-control is the current method for controlling the system, this will serve as a point of reference to the MPC.

4.1 Prediction

This thesis uses an ARMAX model, however, other models were explored. In a comparison between multiple models, specifically ARX, ARMAX, state space model, and different types of nonlinear ARX (see Ljung [4]), the ARMAX model had the best prediction accuracy for the dataset. The model is trained on a dataset of 30000 samples, and tested on a test dataset of 1200 samples.

$$\text{NRMSE} = \left(1 - \frac{\sum_{i=1}^{n} ||y_i - \tilde{y}_i||}{\sum_{i=1}^{n} ||y_i - \bar{y}||}\right) \tag{5}$$

Training a model with focus on prediction 1-step-ahead prediction yielded good results for MPC. Figure 6 and 7 show the results for 1-step-ahead and 15-step-ahead prediction on the test dataset. The prediction accuracy is calculated as normalized-root-squared-mean-error (NRMSE), seen in Equation (5). Here, y_i is the ith measured output, $\tilde{y}_i$ the ith prediction, and $\bar{y}$ the mean of all samples. A value of one is a perfect fit, whereas a value of zero means that the prediction is as accurate as the average value of the output signal. There is no lower range limit to NRMSE.

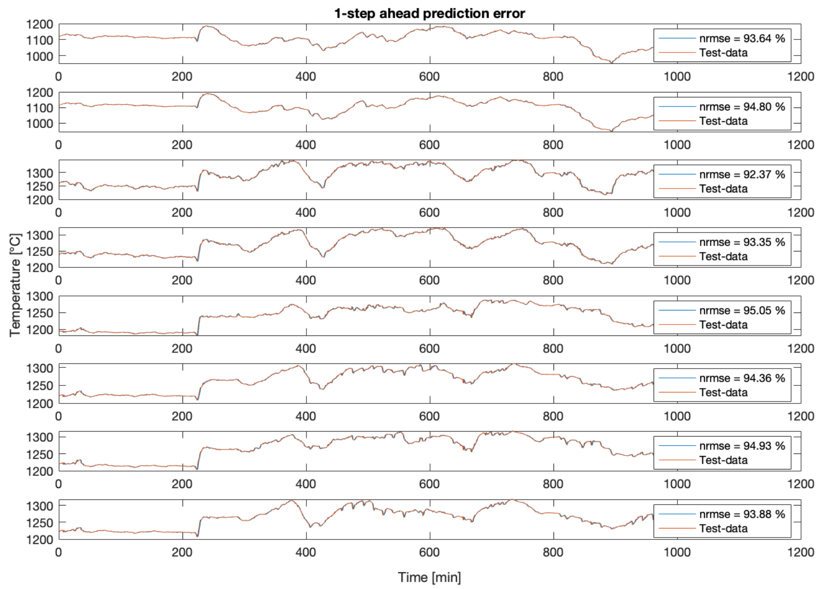

Figure 6: ARMAX 1-step-ahead prediction. The temperatures of the different zones are shown in increasing order, from top to bottom.

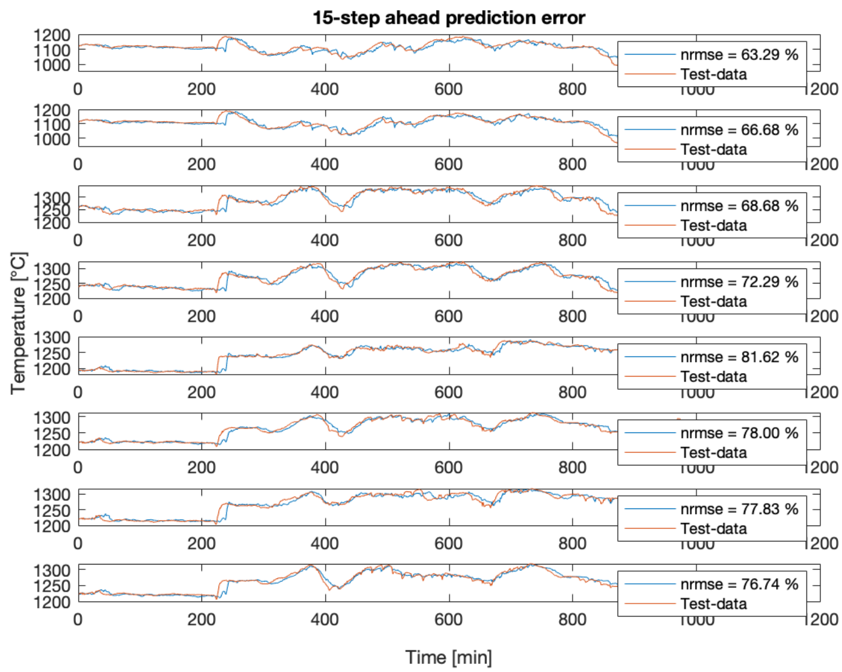

Figure 7: ARMAX 15-step-ahead prediction. The temperatures of the different zones are shown in increasing order, from top to bottom.

The first section of the furnace has a faster decrease in prediction accuracy with increasing τ, especially for zone 1 and 2, as demonstrated in Figure 8. This is not unexpected, since a large part of this section is a "dark zone", the cold slabs enter here, and there are disturbances in the form of gas backflow and air inflow. The errors caused by these disturbances accumulate over time. However, the accuracy of this section of the furnace has less importance for MPC. The difference between slabs and the furnace temperature is usually large, meaning that the heating rate of the slabs is not significantly impacted. In contrast, the last section is of high importance, since it is used to set the final temperature of the slabs before they exit the furnace. The error is further eliminated by the presence of feedback in the MPC-controller.

The prediction accuracy decreases as the prediction horizon increases. For the first couple of step-ahead predictions the accuracy is very good, with NRMSE over 90%. This naturally decreases over time, as prediction errors due to unmeasured disturbances and nonlinearities accumulate. Figure 8 shows how the NRMSE of the prediction accuracy decreases as the length of the prediction horizon increases. The accuracy is initially high and gradually decreases. The rate of decrease in accuracy differs between the zones, once again highlighting that zone 1 and 2 are more affected by disturbances and other causes of error. The rate of decrease in accuracy is of importance when setting the prediction horizon for MPC. While it is beneficial to take a longer period of time into account when optimizing the input signal, decreasing accuracy can cause the control to have a different effect than desired. Therefore, the prediction horizon should be limited to a range which has relatively good accuracy. However, being able to accurately predict what effect the input signal will have on the next couple of time-steps is the most important factor, in order to avoid unexpected behavior that will not be eliminated by feedback. Simply put, if the predicted behavior for the next few steps is very inaccurate, there is a significant problem. In this regard the model performs quite well, having an average NRMSE of 91% over the first 5 step-ahead predictions.

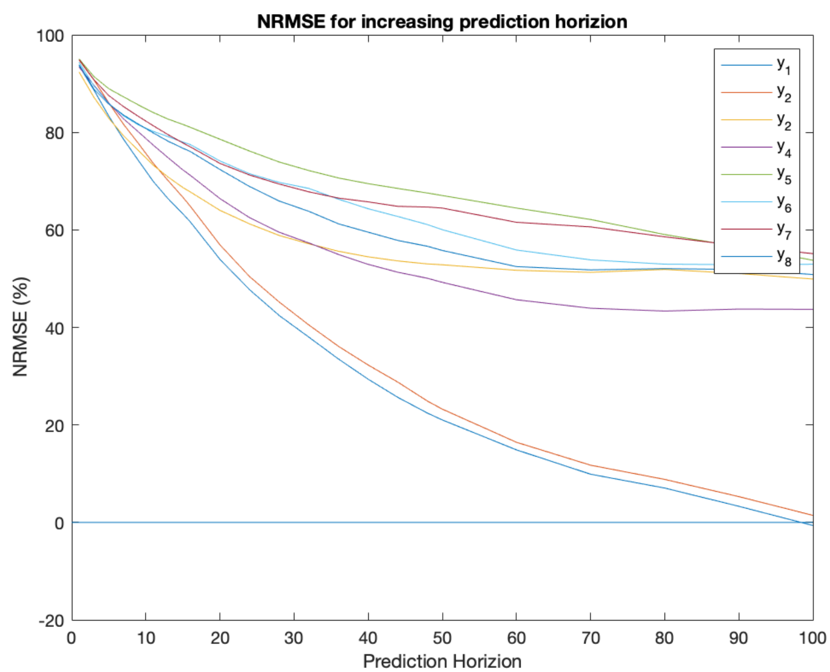

Figure 8: Prediction error over increasing τ.

The graph shows the difference in prediction accuracy for the different zones. It is clear that zone 1 and 2 perform worse than the other zones, which is likely due to the cool slabs entering, as well as the cold air that comes with the slabs. It is also likely a result of the length of the first zone, most of which is not directly controlled. Furthermore, the later zones usually have heated slabs present, and thus more heat. Therefore, disturbances such as cool air should not cause the same decrease in temperature.

4.2 Simulation

For the purpose of simulation, the long term stability of the model is of higher importance. While training the model with a prediction focus yields good accuracy, it comes at the cost of neglecting the long-term accuracy of the model. A simulation model is trained with emphasis on long-term accuracy. As such, it performs better over a long simulation period, but performs poorly for short-term prediction. For the purpose of MPC, the prediction model is the preferred choice over the simulation model. Since MPC works under feedback at every time-step, low frequency error does not accumulate over time. However, since the physical model is unavailable for experiments, the thesis will proceed with the simulation model for the MPC simulation.

The accuracy of the simulation model is limited. The effect of unmeasured disturbances, and to some extent nonlinearities, accumulate over time, causing the simulation to increasingly deviate from the behavior of the physical system. Figure 9 shows the comparison between the simulation model and the testing data. The accuracy here is significantly lower than the short term accuracy of the prediction model. While this is not relevant for implementing the MPC for the physical system, it limits the conclusions that can be drawn from the following results. Furthermore, since the accuracy of the simulation model degrades over time, it is best suited for relatively short simulations.

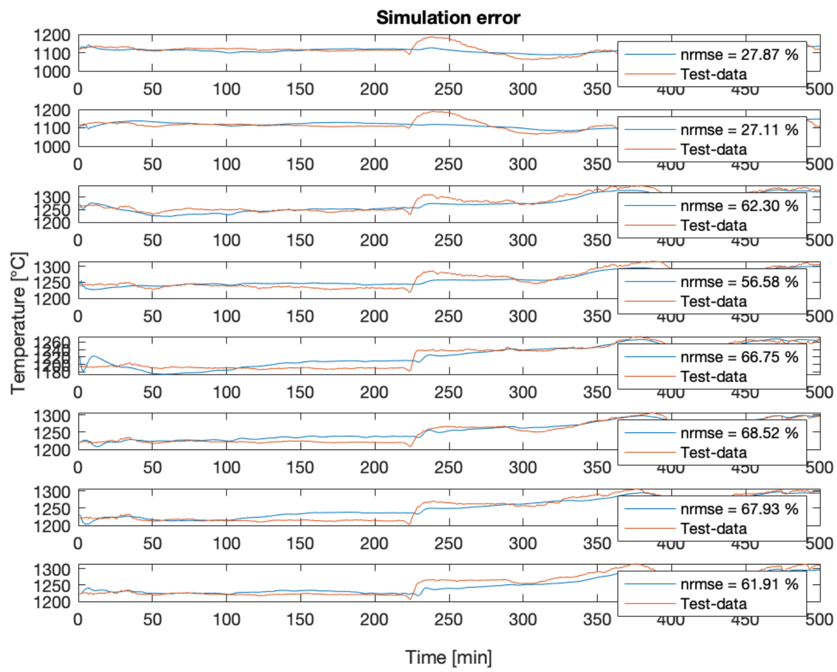

Figure 9: Simulation accuracy for ARMAX. The temperatures of the different zones are shown in increasing order, from top to bottom.

The model has similar levels of accuracy for the training data. This indicates that it is biased, meaning that it is unable to accurately capture the behavior of the system. It is clear from the training-dataset simulation accuracy that there are significant inaccuracies in the simulation model, especially in the first two zones. This issue is not unique to a linear model. The NARX also produces low accuracy prediction simulation results. This implies that the inaccuracy does not arise in the ARMAX mainly due to the ability to capture nonlinearities, but more so unmeasured disturbances that are not captured in the input-output data. This is further supported by the nonlinear detection ratio, which shows that no nonlinear are detected for a time period of 5000, more than long enough to not be a significant problem for neither the simulation model nor the prediction model. The lack of experiments and physical insight makes it difficult to validate certain characteristics of the model. One of the few things known is that the model should be stable. However, a model trained on the input-output data is not guaranteed to be stable. In fact, not using regularization yielded an unstable model. However, tuning the hyperparameters, especially regularization with the regularization parameter $\lambda = 20$, gave a stable model. While the model is stable, the steady states are not necessarily correct. There are no experiments available that can be used to verify the steady state value for a step function, nor any experiments that show the cooldown rate in the absence of disturbances and slabs. The lack of physical insight makes it difficult to validate the model on characteristics other than the prediction accuracy.

An additional problem with simulation is the absence of the weighted sum of the residuals for predicting the next time-step. The moving average term constantly corrects any drift from the real values when predicting. When simulating, no reference values are available. The absence of the residuals makes the simulation result from ARMAX less reliable.

4.3 Control

Using MPC to control the simulated system yields acceptable slab temperatures, although this depends on the MPC (hyper)parameters. However, for MPC to be an option to PID that is worth further consideration it also needs to outperform PID-control. Due to the inaccuracies in the simulation model and lack of disturbances, it is hard to conclude anything about the performance of MPC and PID for the real system. Instead the comparison will be limited to the simulation model, and the extent to which the two methods can be used to control it. Therefore, the results mainly serve to highlight differences and limitation of the two methods.

The desired characteristics of a controller for this system is to achieve a final slab temperature of 1250°C, while using as little fuel as possible. The maximum amount of fuel used by the furnace is 2162 Nm^3 per hour. Over the simulation period of 338 minutes, this gives a maximum of 730756 Nm^3. From the system input data, it is clear that the furnace typically uses around 65% of this, which usually give satisfiable results in regards to the final slab temperature. This data is from a period during which 15 slabs entered the furnace, the same as the simulation. This is not unique to this subset of the data; similar amounts of fuel is typically used.

The goal of the furnace is to have a homogeneous slab temperature of 1250 °C, or as close as possible. The simple slab model used in this thesis offers no insight to heterogeneity: the mean temperature is the measurement on which the evaluation will rely. This is a weakness that could be further improved upon. For the evaluation of the control system, ten slabs enter with 10 minute intervals, and stay in the furnace for 178 minutes: 83 minutes in zone 1, 55 minutes in zone 2, and 40 minutes in zone three. This is similar to the real system in operation. The difference is that the real system is paused by the operator if slabs are not sufficiently hot, and they stay in the same position until they are. However the purpose of this simulation is to highlight the performance of MPC. Distribution of final slab temperature is a good metric of this, along with how much fuel is used.

The PID-controller serves as a comparison for the MPC. It also highlights some of the problems of the simulation model. The PID-controller referenced here is a set of eight SISO controllers, one for each zone. The parameters of the PID-controllers are set using Ziegler–Nichols' method, as described by Glad and Ljung [6]. Figure 10 and 12 shows the input and output signal of the PID-controlled system. The PID does manage to control the slab temperature adequately.

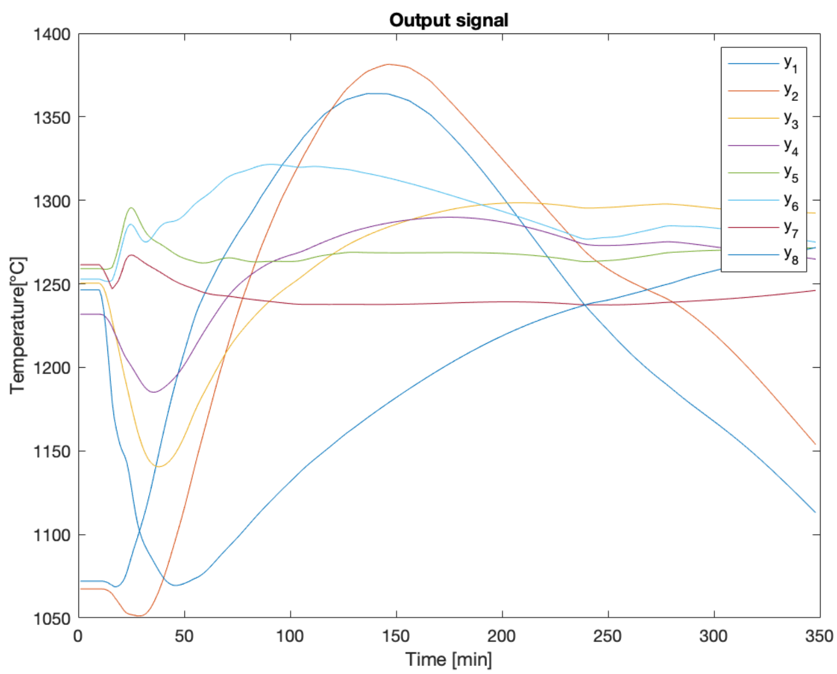

Figure 10: System output with PID-controller.

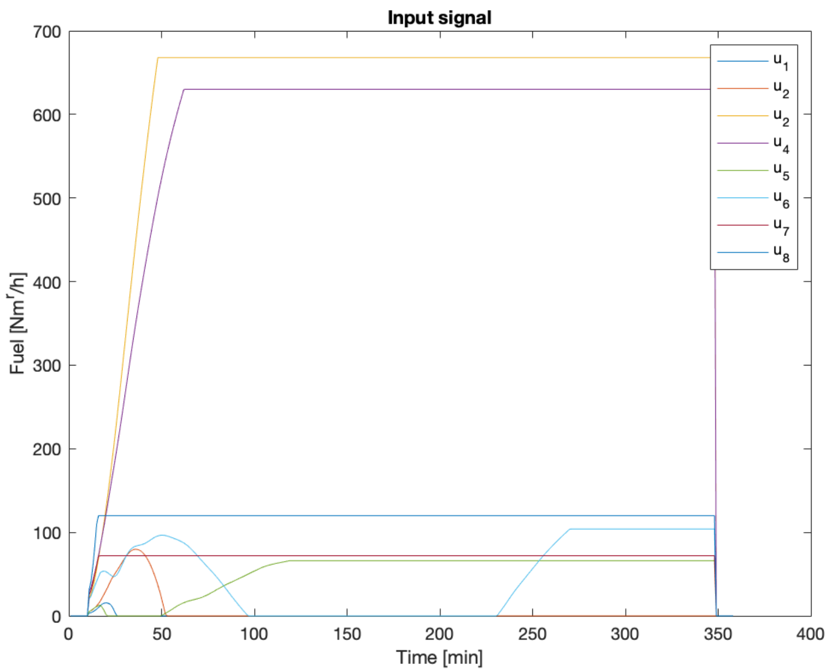

Figure 11: System input with PID-controller.

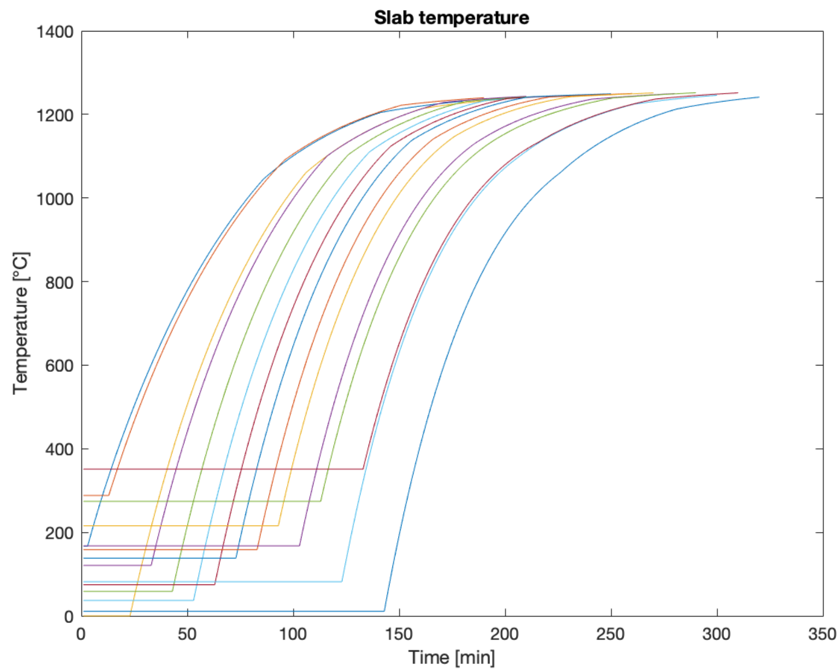

Figure 12: Slab temperature controlled by PID-controller. Each line represents the temperature of a different slab, with a different initial temperature.

The simulated PID-controlled system highlights some potentially unrealistic behavior. While experimental verification of the systems step response is not available, the initial overshoot in zone 1 and 2 is not found in the training or testing data. This overshoot reaches temperatures higher than the maximum value found in the training data, and exceeds the maximum allowed temperature in the furnace. This mainly highlights a problem with the simulation model. The overshoot cannot be dampened by the PID, as it does not seem to be the naturally step response of the system. As the controller can only add heat to the system, it can not dampen the overshoot. Moreover, the effect of disturbances seen in the system data is significantly larger than the disturbances included in the simulation model. The only disturbance present in the simulation is the initial effect of the cold slabs entering, for the first $nk + nb$ timesteps for each slab. While the initial temperature decrease is included as a measured disturbance, the simulation model does not seem to be affected by this to the same extent as the real system is to a cold slab entering the furnace. This limits the scope of the comparison to what extent the two methods can be used to control the (almost) undisturbed system.

During the simulation, the PID-controller does manage to heat the slabs to an acceptable temperature. If the system is allowed to reach a steady state value, the PID-controller has excellent performance. However, the presence of disturbances makes a steady state value unrealistic, and the results from doing so are even more idealized then the simulation model already is. Figure 13, 14, and 15 show the typical performance for MPC-control of the simulated system. The experimental conditions are the same as for the PID, with the exception of the starting point which is not at steady state. While some input signals are at maximum capacity most of the duration of the simulation, some other input signals display behavior very different from PID-control. While some input signals are not smooth, the output signal remains fairly constant for most of the simulation. This is a relatively narrow window for the slabs to heat up to the desired temperature. When slabs exit a zone, the input signal is decreased in order to minimize the amount of fuel used. Similarly to the PID-controller, the MPC chooses to keep the furnace temperature fairly constant throughout the simulation. While this is not surprising, it would be more difficult with the addition of disturbances.

The overshoot found with the PID-controller is not found for the MPC-controller. This suggests that overshoot is likely a result of some correlation between the output signal for zone 1 and 2, and the input or output signal for the other zone. This highlights the fact that the furnace is a MIMO system, and connecting several SISO controllers does not give full control. It is likely that the simulation model is biased, as the behavior in Figure 10 is not found in the training data. Despite this, the differences in behavior for the two controllers show an advantage of MPC for MIMO systems. Within the context of the simulation, it is clear that MPC provides better control.

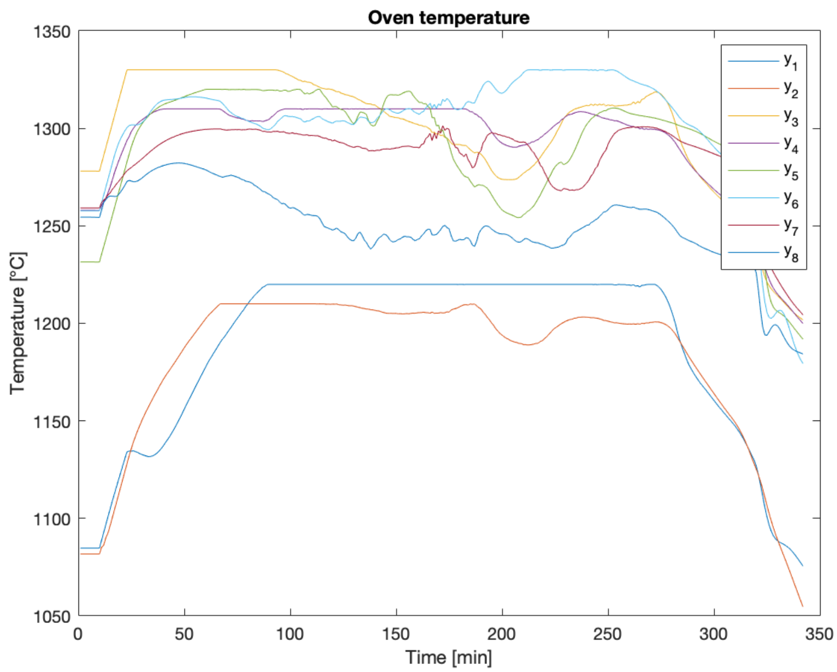

Figure 13: Furnace fuel flow controlled by MPC-controller with $\mu = 0.01$.

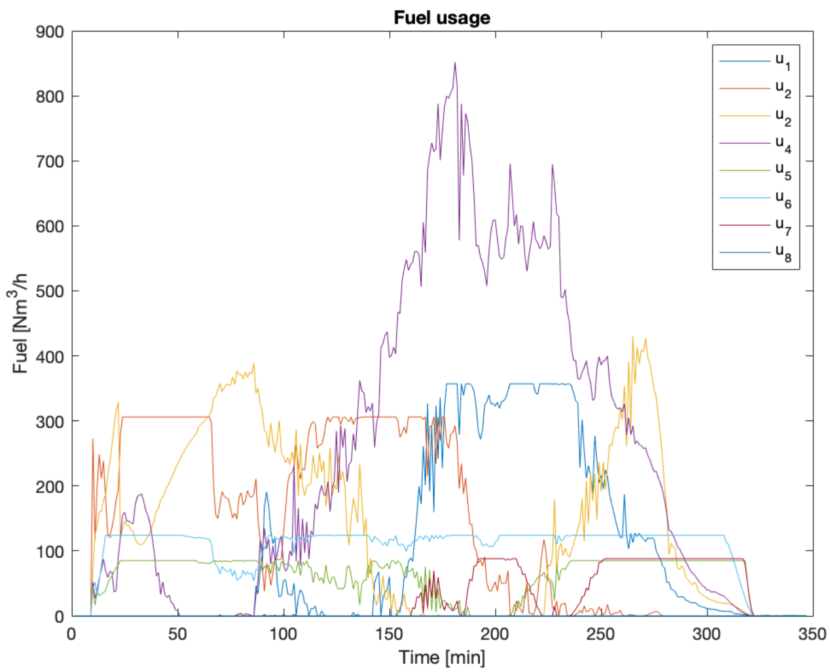

Figure 14: Furnace temperature controlled by MPC-controller with $\mu = 0.01$.

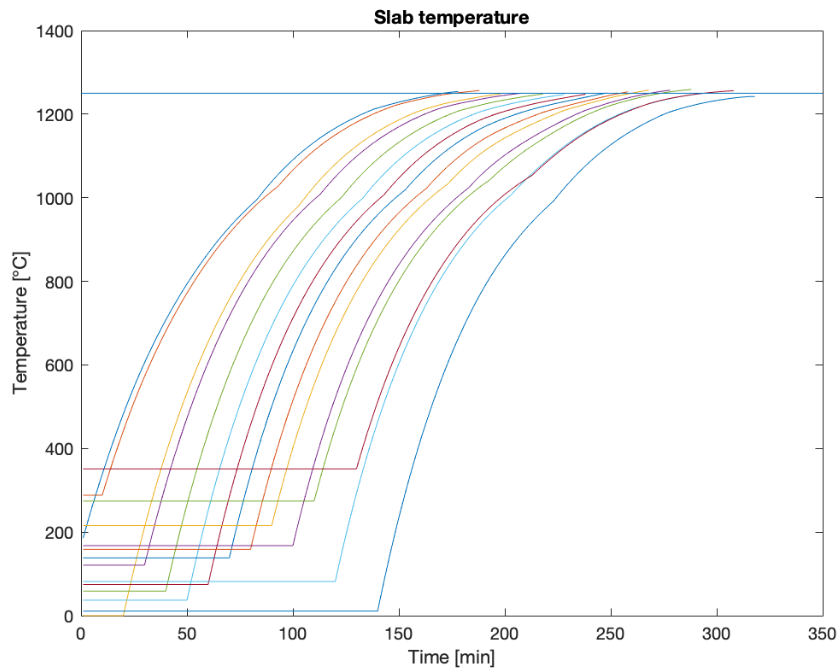

Figure 15: Slab temperature controlled by MPC-controller with $\mu = 0.01$. Each line represents the temperature of a different slab, with a different initial temperature.

The simulated comparison of the PID and MPC-controllers provides some insight to the behavior of the furnace. Although the comparison is limited by accuracy of the simulation, the behavior can be used as an indicator of the relative performance of the two methods. Moreover, the results show that the implementation of MPC works, and could be used (with some modification) to control the physical furnace.

The MPC simulation is not without problems. Apart from the differences between the simulation and the real system, the MPC-controller is not ideal. Since it starts with and empty furnace, the initial focus is solely to increase the temperature in section 1. For the simulation model, this initial overshoot brings the system away from the quasi-equilibrium position, and it results in all zones dropping in temperature. This is most likely due to the simulation model having high bias towards the training data, that fails to accurately describe the behavior of the furnace due to an absence of disturbances. The overshoot does not occur in the real furnace. However, it still highlights that the MPC has flaws apart from the model prediction errors. Increasing the prediction horizon would remove this, but this comes at the cost of nonoperational increase in computational time, more than the sampling time allows for. While this problem arises in the simulation setup, it might occur naturally in production. Therefore a solution to this problem should be included in any real system implementation of MPC.

The distribution of slab temperature for different μ is shown in Figure 16. Here, 15 slabs with initial temperatures of between 0 and 400 °C are used. Each slab was simulated to be in the furnace for 178 minutes. The amount of fuel used in the simulations is shown in Figure 17. Both figures make it clear that MPC with lower μ outperforms the PID-controller in the simulation environment, both in terms of final slab distribution and fuel usage.

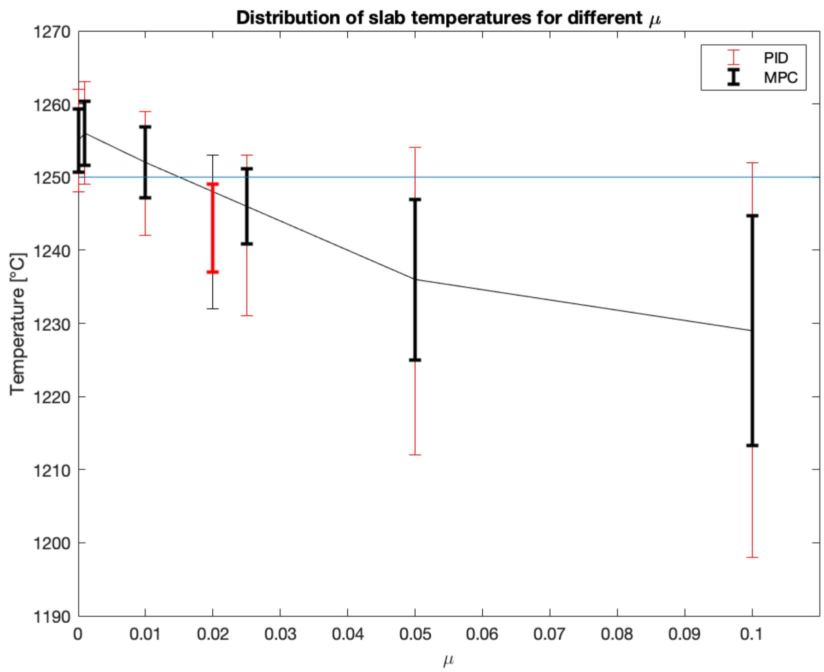

Figure 16: Distribution of slab temperatures when entering the furnace for different μ. The thin bars show min and max values, and the black bars show the length of 1 standard deviation from the average value. The red bar shows the comparative performance of the PID-controller, and the blue reference line is the target value.

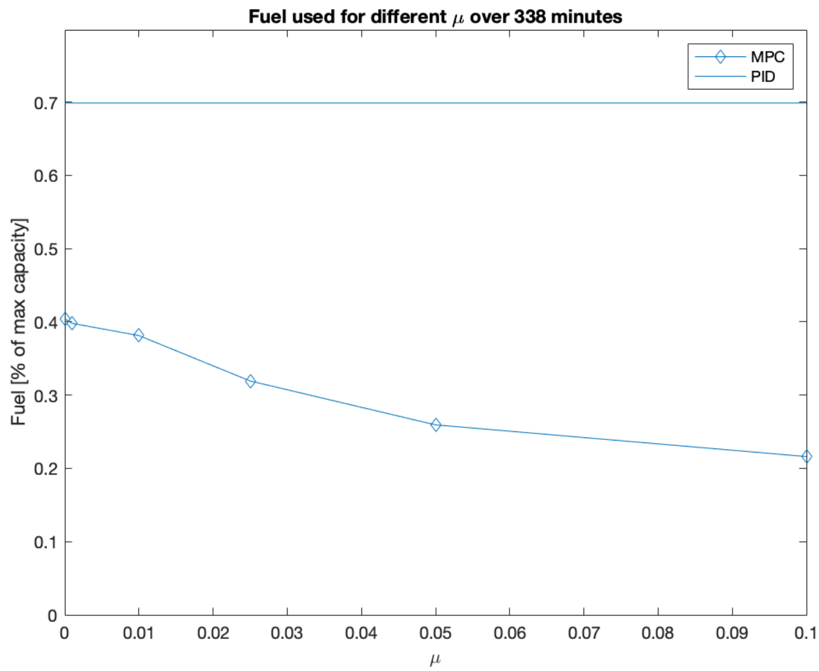

Figure 17: Percentage of maximum fuel capacity for different μ.

Figure 16 and 17 showw the performance for different μ. It is clear that using a high μ greatly reduces the fuel usage, but it comes with slab temperatures that are in an unacceptable range. Smaller values of μ perform well, while still managing to lower the fuel usage. While the decrease in fuel usage is relatively small, it still has a rather large impact on the fuel consumption over time due to the large amounts of fuel used. The PID has a somewhat worse performance than MPC for all the input penalty terms.

This thesis makes no claim that this PID-controller has the best set parameters, it could certainly be improved. The result does hint at MPC having an advantage over PID-control. This conclusion is only valid for the simulated system and further experimentation on the real system is needed so that this is the case for the physical system. However, it still gives an indication of the behavior. While any conclusion is not guaranteed to transfer to the physical furnace, it still suggests that MPC might have some benefits over PID-control for this system.

5 Limitations and Further Research

There are several improvements that could be made to further investigate the result and develop a stronger controller. The first one is data acquisition. The current data available is from the closed loop system. Having an open loop system with different input signals could make the process of system identification easier and more accurate. It could also help highlight a wider range of values. Additionally, the disturbances could to a larger extent be accounted for. The best case scenario would be to measure the slab temperature inside the furnace, which would allow for more direct control. Additionally this would supply information about the cooling effect that the slabs have on the furnace. However, this is an expensive option. A different option would be to use a more accurate slab heating model. If the error in the model is very low, it could potentially be useful to determine the cooling effects on the slab.

Another clear improvement would be the model validation. While the simulation model allows for some comparison of the two control methods, it is limited. The above mentioned improvements mainly concern the simulation methods, and with access to the real system for testing this model would be obsolete. It is reasonable to suspect that the prediction model could perform well on the real system, but without access to it is difficult to determine with any certainty. While the PID-controller's steady state is better in terms of slab temperature than the MPC, the real furnace does not allow for a steady state due to the disturbances. The main question that this thesis leaves open is what method is more suitable for controlling the real system with disturbances present.

When using a nonzero input signal penalty μ, the extent to which the fuel flow is penalized is dependent on the number of slabs present in the furnace. Since every slab contributes to the overall cost function, more slabs gives the H matrix a larger weight in comparison to the F matrix (see Equation (4)). Therefore, when very few slabs are present in the furnace, the MPC might mainly aim to optimize the fuel usage, rather than the slab temperature. A solution to this problem would be to divide μ with the number of slabs in the furnace, which would give a similar fuel penalty regardless of the number of slabs in the system. On the other hand, the cost of lower quality increases with the number of slabs, so not scaling μ might have some benefits.

6 Conclusion

While the result is not conclusive, it does suggest that MPC could be a useful substitute to PID-control for Outukompu's furnace in Avesta. This thesis can be seen as an indicator that this subject is worth further investigation, and that the result has the potential to yield results that are beneficial for the amount of fuel used and the ability to effectively control the temperature of the slabs. This result is solely based on the simulation model, and does not necessarily transfer to the real system. Much further development and experimentation is needed before arriving at any such result, but the limited evidence in this thesis suggests MPC should be considered as the furnace control system, due to good prediction model accuracy and lowered simulation fuel usage.

As for the simulation model, it can not be considered good enough for reliable testing of the system. In order to achieve this, the disturbances should be measured or modeled. As mentioned, measuring or accurately modeling the slab temperature can potentially greatly increase the accuracy of the simulation. For now, the usage is limited to getting an overview of control performance for an undisturbed system of similar characteristics as the furnace.

References

[1] Jernkontoret. *Roadmaps for fossil free competitiveness.* URL: https://fossilfrittsverige.se/en/roadmap/the-steel-industry/. Accessed: 11.01.2023.

[2] Jernkontoret. *The steel industry.* URL: https://www.jernkontoret.se/en/the-steel-industry/industry-facts-and-statistics/energy-use/. Accessed: 11.01.2023.

[3] Statistical Database. *Total emissions and removals of greenhouse gases by greenhouse gas and sector.* URL: https://www.statistikdatabasen.scb.se/pxweb/en/ssd/START__MI__MI0107/TotaltUtslappN/. Accessed: 11.01.2023.

[4] Lennart Ljung. *System Identification, Theory for the user.* Prentice-Hall, Inc., 1999.

[5] Einar E. Madsen et al. *STEELTEMP 2D.* SWEREA, Institute for Energy Technology, 2012.

[6] Torkel Glad and Lennart Ljung. *Reglerteknik, Grundläggande teori.* Studentlitteratur AB, 2006.