Helge Smedsrud

DYNAMIC MODELING AND CONTROL OF BROBEKK INCINERATION PLANT

MASTER THESIS
SPRING 2008

Norwegian University of Science and Technology
Faculty of Natural Sciences and Technology
Department of Chemical Engineering

ABSTRACT

Over the last year, Brobekk waste incineration plant in Oslo has suffered from transient cool-downs due to periodically increased loads in the district heating network it supplies. This situation is highly undesired, as it may lead to corrosion on the pipes in the furnaces. Over time, this may necessitate a run-down of the affected production line, and result in costly repairs and loss of income. A revision of the control structure related to the heat exchangers was planned to improve or eliminate this problem. For this purpose two mathematical models of the plant were made; one steady-state MATLAB/TOMLAB model for optimization purposes, and one dynamic Simulink model for control purposes. The optimization sought to reduce operational cost, while maintaining a steady water inlet temperature in the furnaces. This resulted in the plant operation space being divided into four regions distinguished by sets of active constraints on the manipulated variables. Four separate decentralized control structures were then devised. These were all able to successfully control the plant within their respective regions, and the maximum observed deviation in the furnace inlet temperature was only 30 % of the allowed deviation. Transitions between the regions were also examined, but this subject requires further studies.

I declare that this is an independent work according to the exam regulations of the Norwegian University of Science and Technology:

Trondheim, June 4, 2008

Helge Smedsrud

ACKNOWLEDGEMENTS

The author wishes to thank the following people for their contributions during the course of this project:

- *Sigurd Skogestad*, professor at the department of chemical engineering, for project supervision and discussions.

- *Johannes Jäschke*, Ph.D. student at the department of chemical engineering, for valuable comments, discussions, and aid.

- *Helge Mordt*, senior consultant at Prediktor AS, for providing background material for the project, as well as answering questions concerning the incineration plant and its control structure. Helge also arranged a greatly appreciated tour to the Brobekk plant.

- *Shridharakumar Narasimhan*, post-doctoral research fellow at the department of chemical engineering, for discussions concerning the control of the plant.

- *Magne Hillestad*, professor at the department of chemical engineering, for an enlightening discussion on pressure modeling.

CONTENTS

1. INTRODUCTION

With today's focus on climate and economic efficiency the cooperation between waste incineration plants (WIPs) and district heating networks (DHNs) has become increasingly interesting. The volume of the waste can be reduced by as much as 95-96 % through an incineration process (RenoSam and Rambøll, 2006), and large amounts of energy are released. The remaining residue, which contains non-flammable materials like glass and metal, is then either recycled or deposited at a designated location. Due to the residues' causticity it can successfully be used in the process of neutralizing acidic industrial waste.

The high temperatures and pressures in WIP-DHN systems, especially at the waste incineration plant, put forth many operational requirements; too high temperatures may lead to bending in the pipeline, whereas too low temperatures can result in condensation of acidic flue gas. Too high pressure can result in rupture of valves or bends, and too low pressure increases the risk of flashing. From a control point of view the apparent question is therefore: How can we develop a control structure that not only takes care of all these safety constraints, but also runs the plant as optimally as possible from an economic point of view? This is what will be studied in closer detail in this report. In Chapter 2 we discuss how one can model a WIP-DHN system mathematically, while Chapter 3 looks into the steady-state optimization of such a system, as well as the creation of a control structure for the incineration plant. Chapter 4 presents how the closed-loop dynamic model behaves for a wide range of disturbance combinations. Finally, Chapters 5 and 6 will discuss and conclude on the work.

2. BACKGROUND

2.1 BROBEKK INCINERATION PLANT

Brobekk incineration plant is located at Alnabru in Oslo, and was the first large-scale incineration plant in Norway when it was started up in 1967. The plant was upgraded in the eighties, and in 2002 a new, advanced treatment plant was built. During the summer of 2007 new heat exchangers were installed.

The Brobekk plant, as well as its sister plant at Klemetsrud, is operated by the waste recycling department (Energigjenvinningsetaten, EGE) of the City of Oslo. The plant sells its produced energy to Viken Fjernvarme AS, which operates a district heating network in Groruddalen and in Oslo city. The designed capacity of the plant is 100 000 metric tons of waste per year. The last years, however, the plant has been operating above this.

Figure 2.1 shows an overview of the WIP-DHN system in the Oslo area. The boxes under Viken Fjernvarme control represent local sub-stations where additional heat can be added to the DHN. The arrows represent flows of water.

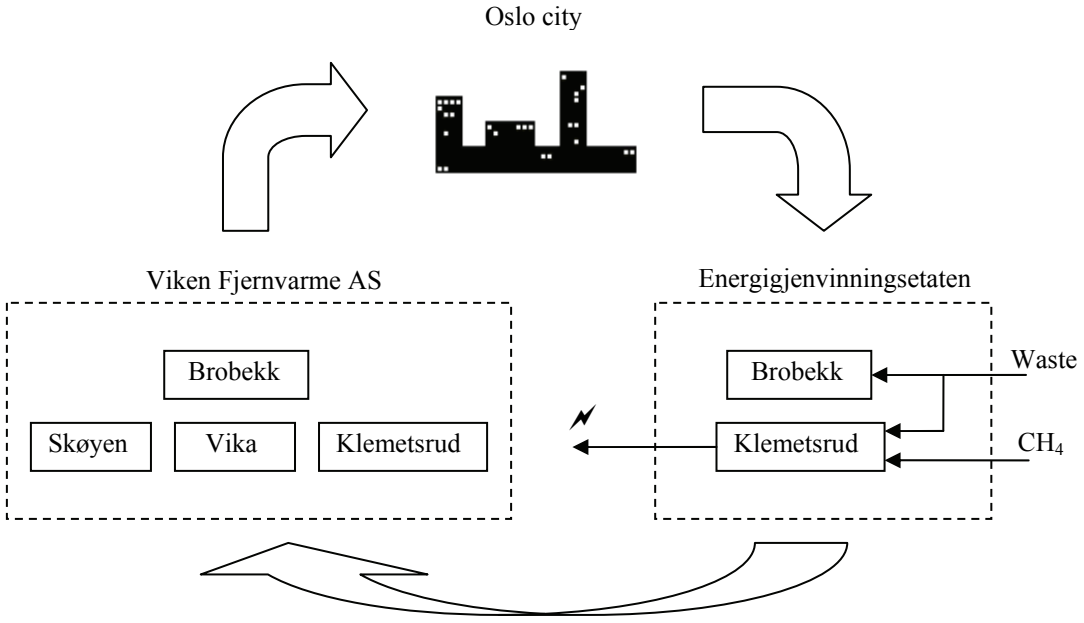

Figure 2.1: Waste recycling and district heating in Oslo city.

With the installation of new heat exchangers during the summer of 2007, Prediktor AS – a Fredrikstad based IT company – was contracted for revising the control structure. Soon after it turned out that the implemented control structure was not able to run the plant as desired. Early in the morning and late in the afternoon, when the DHN users take out much heat, the plant may get temporarily cooled down. This can happen because the present control structure allows Viken Fjernvarme (henceforth called Viken) to take out more energy from the system than the 32 MW produced at the Brobekk plant (henceforth called Brobekk).

Figure 2.2 shows a process diagram of one of the two identical heat exchanger lines at Brobekk, as well as Viken's side of the plant. The variables are explained in Tables 2.1 to 2.3.

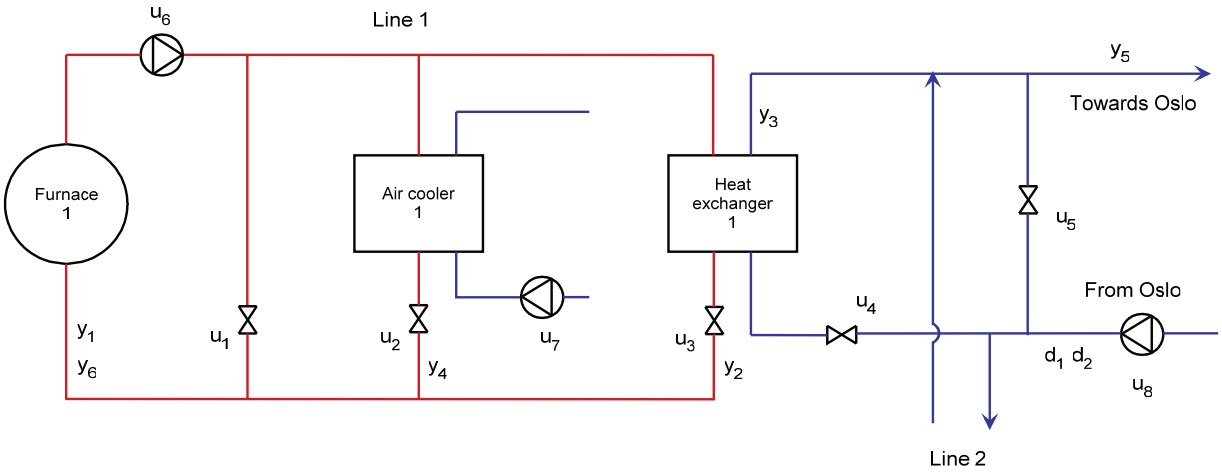

Figure 2.2: Schematic overview of the Brobekk plant.

Table 2.1: Manipulated variables (MVs) in the Brobekk plant.

Shorthand notation	Description	Quantity
u_1	Brobekk side bypass valve	2
u_2	Brobekk side AC[1] valve	2
u_3	Brobekk side HE[2] valve	2
u_4	Viken side HE valve	2
u_5	Viken side bypass valve	1
u_6	Brobekk side pump duty	2
u_7	Brobekk side fan duty	2
u_8	Viken side pump duty	1

Table 2.2: Controlled variables (CVs) and measurements in the Brobekk plant.

Shorthand notation	Description	Quantity
y_1	Water temperature at furnace inlet	2
y_2	Water temperature at HE primary side outlet	2
y_3	Water temperature at HE secondary side outlet	2
y_4	Water temperature at AC primary side outlet	2
y_5	Water temperature towards Viken/Oslo	1
y_6	Total flow rate of water, Brobekk side	2

Table 2.3: Disturbances in the Brobekk plant.

Shorthand notation	Description	Quantity
d_1	Water temperature from Oslo	1
d_2	Total flow rate of water, Viken side	1

[1] Air cooler.
[2] (Primary) heat exchanger.

2.2 MODELING

In order to reproduce the characteristics of a system with flowing fluids, pumps, fans, and valves were included in the model. The following simplifications and assumptions were made:

- Since the system contained only pressurized water and non-compressed air, thermodynamical and material properties like heat capacities and densities were assumed constant. Average values for the respective temperature intervals were used.
- The pressure drop across many modern heat exchangers is very small, so this was neglected in the models.
- Isothermal flow was assumed through the pumps, fans, and valves. This was done due to the low pressure differences in the system.

The next chapters describe how the Brobekk plant was modeled for steady-state (Section 2.2.1) and dynamical (Section 2.2.2) purposes. All symbols are described in Appendix B, while the values used in the models are shown in Appendix D.

2.2.1 STEADY-STATE MODEL

Steady-state models are commonly used for optimization purposes or for sizing process equipment. Since these models contain no dynamics, a simplification concerning the heat exchangers can be done, as is shown below.

Heat exchangers
The equations below describe a heat exchanger using the number of transfer units or NTU method (Hertzberg, 2008; Mathisen et al., 1993). First, a heat capacity flow rate ratio β is calculated:

$$\beta = \frac{w^c c_p^c}{w^h c_p^h} \tag{2.1}$$

Then the total number of transfer units is calculated:

$$\alpha = \frac{UA}{w^h c_p^h} + \varepsilon \frac{UA}{w^c c_p^c} \tag{2.2}$$

with ε being -1 for counter-current heat exchangers, and +1 for co-current heat exchangers. After this follows the calculation of a variable arising from the integration of the heat exchanger's energy balance:

$$\gamma = \exp(-\alpha) \tag{2.3}$$

Equations 2.1 to 2.3 are then combined into a dimensionless transfer function matrix describing the temperature changes in the heat exchanger:

$$\mathbf{D} = \begin{bmatrix} \dfrac{\gamma(1-\beta)}{\gamma-\beta} & \dfrac{\beta(\gamma-1)}{\gamma-\beta} \\ \dfrac{\gamma-1}{\gamma-\beta} & \dfrac{1-\beta}{\gamma-\beta} \end{bmatrix} \tag{2.4}$$

The outlet temperatures themselves, as well as the conducted heat, can then be calculated from the following expressions:

$$\mathbf{T}_o = \mathbf{D}\mathbf{T}_i \tag{2.5}$$

$$q = w^h c_p^h \left(T_i^h - T_o^h\right) \tag{2.6}$$

As can be seen from Equation 2.4, if β becomes unity then **D** will become singular, and the simulation will fail. This can be prevented by using a Taylor expansion whenever β lies between $1 \pm \delta$, where δ is an arbitrary small value[3]. First, the following expression is calculated:

$$\eta = \frac{UA}{w^c c_p^c} \tag{2.7}$$

Then follows a series expansion of the exponential term in Equation 2.3:

$$S = \sum_{i=1}^{\infty} \frac{(-\eta\delta)^i}{(i+1)!} \tag{2.8}$$

In this work S was truncated after $i = 5$. The new γ and **D** take the following forms:

$$\gamma = -\eta\delta(1+S)+1 \tag{2.9}$$

$$\mathbf{D} = \begin{bmatrix} \dfrac{\gamma}{1+\eta(1+S)} & \dfrac{\eta\beta(1+S)}{1+\eta(1+S)} \\ \dfrac{\eta(1+S)}{1+\eta(1+S)} & \dfrac{1}{1+\eta(1+S)} \end{bmatrix} \tag{2.10}$$

After this, Equations 2.5 and 2.6 can be used as when β was outside the $1 \pm \delta$ region.

[3] In the work presented in this report a value of 1×10^{-5} was used.

Fans
By assuming a low pressure difference across the fans, as well as a horizontal position, the Bernoulli equation can be used to calculate the fan duty (Roald, 2001):

$$P = \frac{1}{\eta}\left[\frac{w^3}{2\rho^2}\left(\frac{1}{A_2^2}-\frac{1}{A_1^2}\right)+\frac{\Delta p w}{\rho}\right] \tag{2.11}$$

Mixers
The mixing units were modeled using the following simple expressions, under the assumption of instant and homogeneous mixing:

$$w_{tot} = \sum_{i=1}^{n} w_i \tag{2.12}$$

$$T_{tot} = \frac{\sum_{i=1}^{n} w_i T_i}{w_{tot}} \tag{2.13}$$

Pumps
By neglecting any temperature rise in the water during the travel through the pumps, they too could be described using the Bernoulli equation. The elevation difference was set to zero, and the pipe diameter was assumed equal before and after the pump.

$$p_2 = p_1 + \frac{\rho P \eta}{w} \tag{2.14}$$

Valves
The valves were modeled using a standard valve equation on the form

$$Q = K_V \sqrt{\frac{\rho_\circ}{\rho}\Delta p} \tag{2.15}$$

with the flow factor K_V either being constant (Brobekk side valves) or being calculated from model fitted manufacturer data (Viken side valves);

$$K_V = -605.97u^3 + 2986u^2 + 1081.2u \tag{2.16}$$

where u represents normalized valve position. Since the flow rate was unidirectional in the model, a signum function was not included in Equation 2.15.

Comparisons with the Simulink model presented in the next chapter revealed that the heat transfer coefficients – which determine U in the heat exchangers – needed to be modeled with a dependency on the cold water (or air) flow rate in order for the heat exchangers to give

the same behavior. The following empirical correlations were used, with h^h and h^c being assumed identical:

$$\text{HE:} \quad h = 0.0004\left(w^c\right)^3 - 0.1503\left(w^c\right)^2 + 21.917w^c + 7615.8 \tag{2.17}$$

$$\text{AC:} \quad h = 0.4219\left(w^{air}\right)^3 - 20.438\left(w^{air}\right)^2 + 432.3w^{air} + 666.09 \tag{2.18}$$

2.2.2 DYNAMIC MODEL

The dynamic Simulink model was developed for testing the control structure and simulation purposes. In order to get the proper dynamic behavior from the heat exchangers an approach using a cell model with ordinary differential equations was chosen (Mathisen et al., 1993). This means that the heat exchanger was divided into perfectly and instantly mixed tanks, each featuring a hot side, a wall side, and a cold side element (Figure 2.3). The idea is that this will approximate the logarithmic mean temperature difference or ΔT_{LM} of the heat exchanger as the number of cells increases. In our model ten cells – all having identical properties – were used. The furnaces were for simplicity modeled using the NTU method, as the flow rate and inlet temperature of the flue gas was assumed constant.

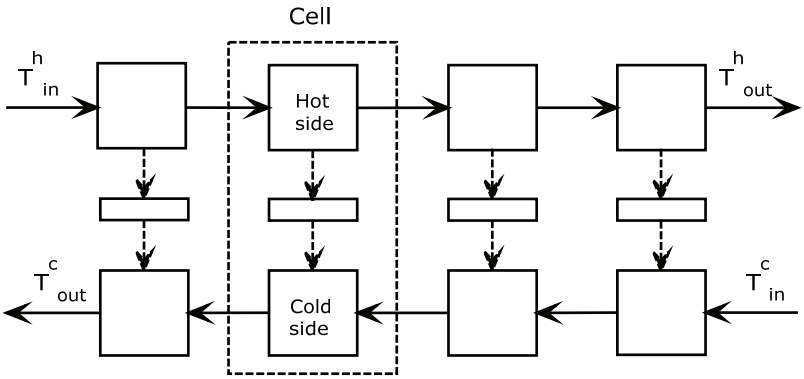

Figure 2.3: Cell model of a heat exchanger with no back-mixing.
The middle elements represent the wall separating the shell and tube sides.

The fans, the Brobekk side pumps, and all valves were modeled as in the previous chapter, with the exception that first order transfers functions with $\tau = 1.5$ seconds were added in order to include some dynamics. The equation for the Viken side pump was rewritten such that the pump duty was calculated, instead of being set by a controller. This was required because the cold water flow rate needed to be given to the Simulink model as a disturbance.

Heat exchangers

It was assumed that each cell was perfectly homogenous, and that no back-mixing occurred. Also, the mixing was instantaneous. The equations below describe a shell-and-tube heat exchanger, for simplicity, but the heat exchangers at Brobekk are of the plate-and-shell type. The parameters for the units were taken from manufacturer data, but were adjusted so that the heat exchangers would give correct steady-state values.

Hot side: $$\frac{\mathrm{d}T^h(i)}{\mathrm{d}t} = \left(T^h(i-1) - T^h(i) - \frac{h^h A}{w^h c_p^h N} \Delta T^h(i) \right) \frac{w^h N}{\rho^h V^h} \quad (2.19)$$

Wall side: $$\frac{\mathrm{d}T^w(j)}{\mathrm{d}t} = \left(h^h \Delta T^{wh}(j) - h^c \Delta T^{wc}(j) \right) \frac{A}{\rho^w c_p^w V^w} \quad (2.20)$$

Cold side: $$\frac{\mathrm{d}T^c(j)}{\mathrm{d}t} = \left(T^c(j-1) - T^c(j) + \frac{h^c A}{w^c c_p^c N} \Delta T^c(j) \right) \frac{w^c N}{\rho^c V^c} \quad (2.21)$$

Due to the stiffness in the dynamic model, with fast responses in the pressure, and slower responses in the heat transfer, the choice of solver fell on `ode15s`. However, when heavy data storing was implemented, this solver tended to give some annoying and unexplainable noise. Another stiff-problem solver, `ode23s`, was in these situations found to eliminate this problem, although it was solving the model somewhat slower.

2.3 STATEFLOW

Stateflow is a design environment for creating state machines and flow charts in Simulink. It features a clean and easy-to-understand graphical user interface, and uses blocks as states, and arrows as transitions between the states (Figure 2.4).

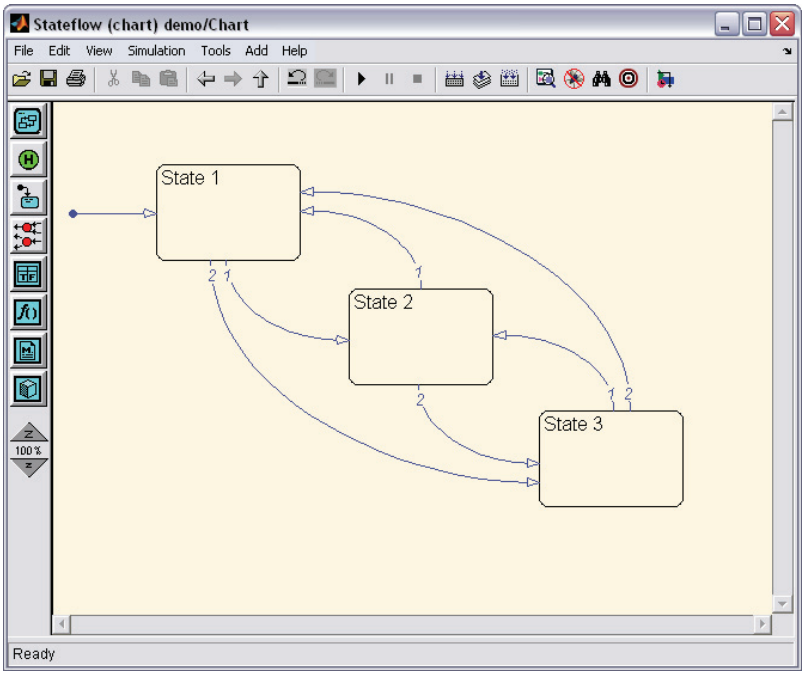

Figure 2.4: The Stateflow chart window, showing three state blocks and six state transitions. The straight arrow indicates to Stateflow what state is to be treated as default.

For each of the states in the figure above, three main situations may be defined; `entry`, `during`, and `exit`. All actions defined under these labels will only happen if the state block is in the proper situation. During a simulation the currently active state or transition is colored boldface blue, in order to easily see where the simulation is. Transitions between different states are defined by naming the transition arrow with the following syntax:

```
event[condition]{condition_action}/transition_action
```

Stateflow uses the same logical operators as MATLAB and Simulink, and the Stateflow chart behaves like other Simulink blocks with input and output ports. For more information concerning Stateflow in general the reader is referred to the homepage of the manufacturer (The MathWorks, 2008).

3. CONTROL STRUCTURE DESIGN

3.1 SELF-OPTIMIZING CONTROL

The issue of determining what outputs to control in a plant is of great importance – not only for economical reasons, but also for safety and environmental reasons. In many situations this is directly given from the process itself, but if the plant gets more complex – and as the number of outputs increases – this task may no longer be equally straightforward. A concept which is useful in this context is self-optimizing control (Skogestad, 2000). According to Skogestad, the idea behind self-optimizing control is «when we can achieve an acceptable loss with constant setpoint values for the controlled variables without the need to reoptimize when disturbances occur.» This means that we want to find and control outputs whose optimal values are non-varying when disturbances occur, or outputs whose optimal values are varying so little that we can use constant set points and still have an acceptable loss. In his paper, Skogestad (2000) presents four requirements which a controlled output should meet:

1. Its optimal value should be insensitive to disturbances.
2. It should be easy to measure and control accurately.
3. Its value should be sensitive to changes in the manipulated variables.
4. For cases with two or more controlled outputs, the selected outputs should not be closely correlated.

Self-optimizing control is sought after the optimally constrained outputs have been controlled at their constraints (also known as active constraint control), and is thus considering the remaining unconstrained degrees of freedom.

3.2 DEGREES OF FREEDOM ANALYSIS

Skogestad (2000) presents the following rule for calculating the number of degrees of freedom (DOFs) available for optimization of a plant:

$$N_{opt} = N_m - N_0 \tag{3.1}$$

where N_m is the number of control degrees of freedom in the plant; and N_0 is the number of degrees of freedom with no steady-state effect on the cost function (e.g. most liquid levels). As N_m is fourteen – based on Table 2.1 – and with N_0 being zero (since all valves and duties in the model will affect the cost) we are left with fourteen degrees of freedom for optimization. However, we must subtract all variables – inputs and outputs – which are constrained in order to find the number of free degrees of freedom for optimization:

$$N_{opt,free} = N_{opt} - N_{active} \tag{3.2}$$

3.3 OPTIMIZATION

The previous chapter revealed that there are fourteen manipulated variables or degrees of freedom for optimization in the Brobekk plant. The task is now to find their optimal values for different disturbance combinations. For this purpose the steady-state model developed in Section 2.2.1 will be used. The software being used is MATLAB and an additional optimization toolbox called TOMLAB. The objective is to minimize the operational cost subject to operational constraints (for a list of these, see Appendix E). This gives the following cost function J:

$$J = u_6 + u_7 + u_8 \tag{3.3}$$

3.3.1 PROBLEM DEFINITION

As shown in Equation 3.3, the optimization of the Brobekk model was done with respect to the sum of the pump and fan duties. For simplicity it was assumed that the two lines at Brobekk were symmetric and that both lines received that same amount of cold water (i.e. the two u_4 valves had identical positions). The natural choice of disturbances where the temperature and flow rate of the cold water from Oslo. A disturbance grid was set up with ΔT = 0.1 °C, running from 65 to 90 °C, and with Δw = 1.6 t/h, running from 500 t/h to 900 t/h. This resulted in 63001 grid points, and the solver[4] used just under one hour to solve the problem. For each point y_1 was constrained to 126 °C, while y_6 was constrained to 250 t/h. Outputs y_3 and y_5 were allowed to vary.

3.3.2 RESULTS

The results from the optimization are shown in Figures 3.1 and 3.2. The blank line which is observed in the plots is due to the NTU method. The reason is that the solver did not converge the model in the region where the heat capacity flow rate ratio β equaled $1 \pm \delta$, due to the non-smoothness of the model. However, values from inside this field may be obtained by interpolating from nearby points. Outputs y_1 and y_4 are not presented; the first is constant at 126 °C, while the latter only exists in the model whenever the air coolers are active, and then it too is constant at 126 °C.

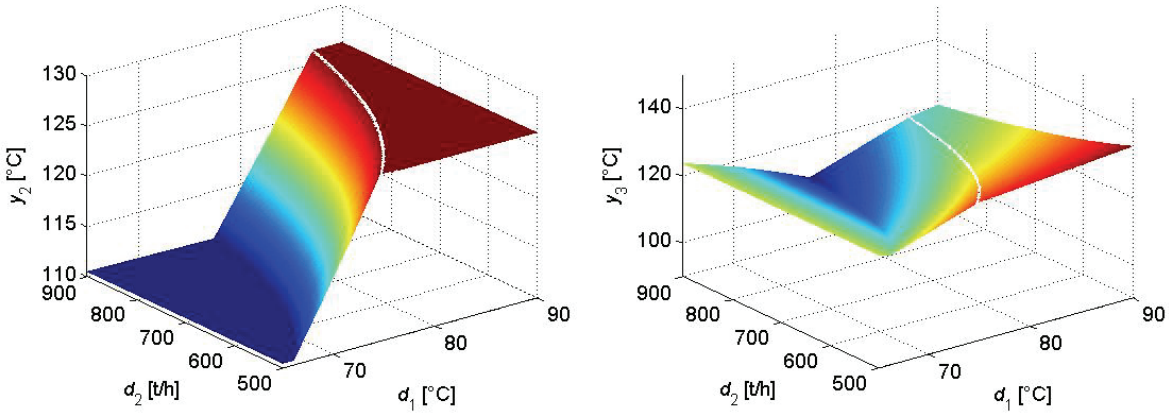

[4] Minos; a large-scale, sparse, general, nonlinear solver (Tomlab, 2008).

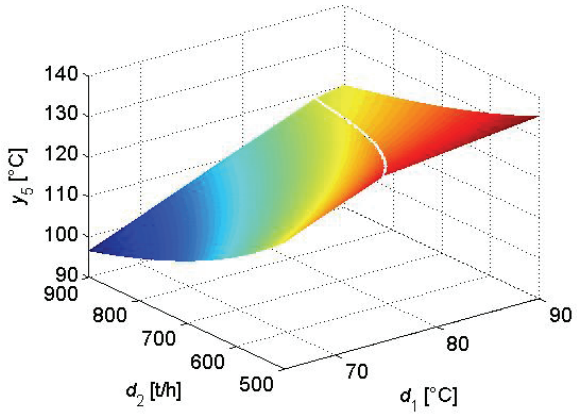

Figure 3.1: Optimal values of selected outputs as a function of d_1 and d_2.

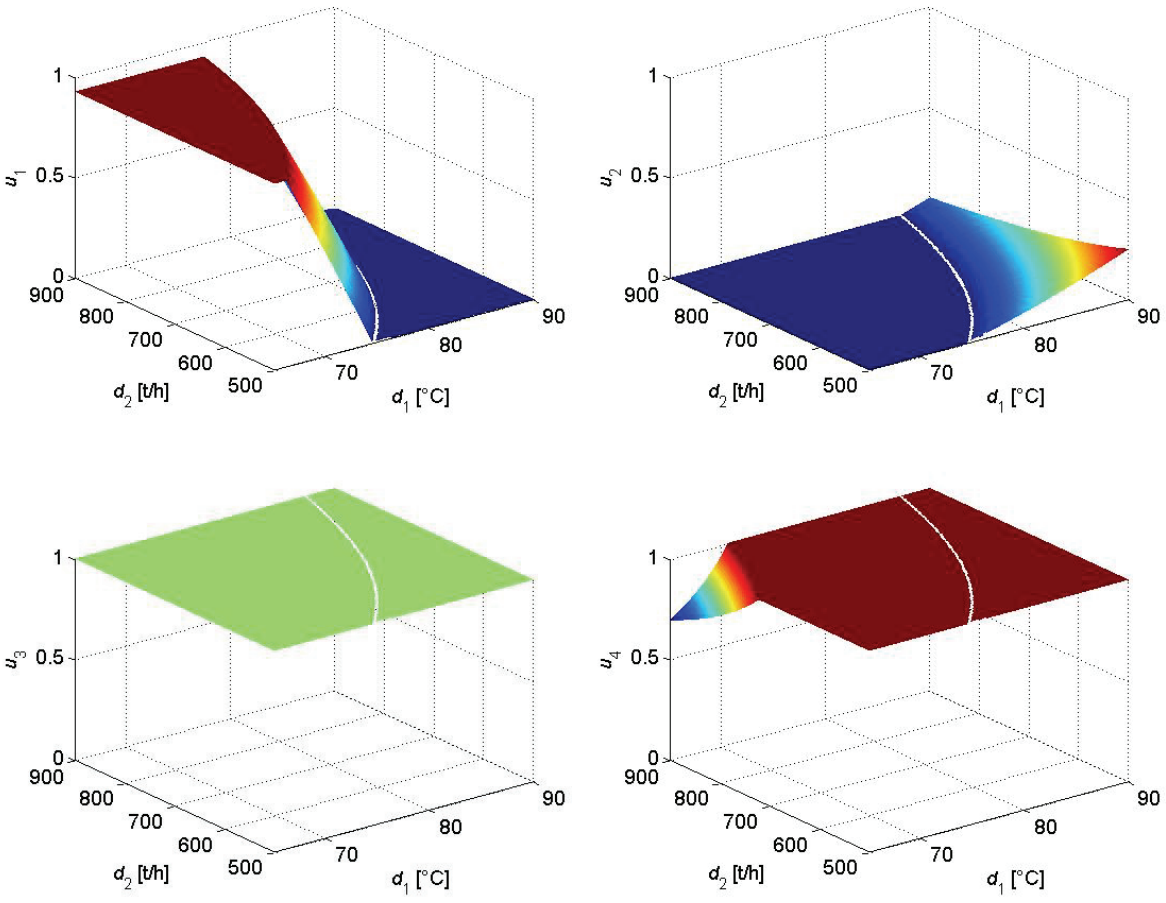

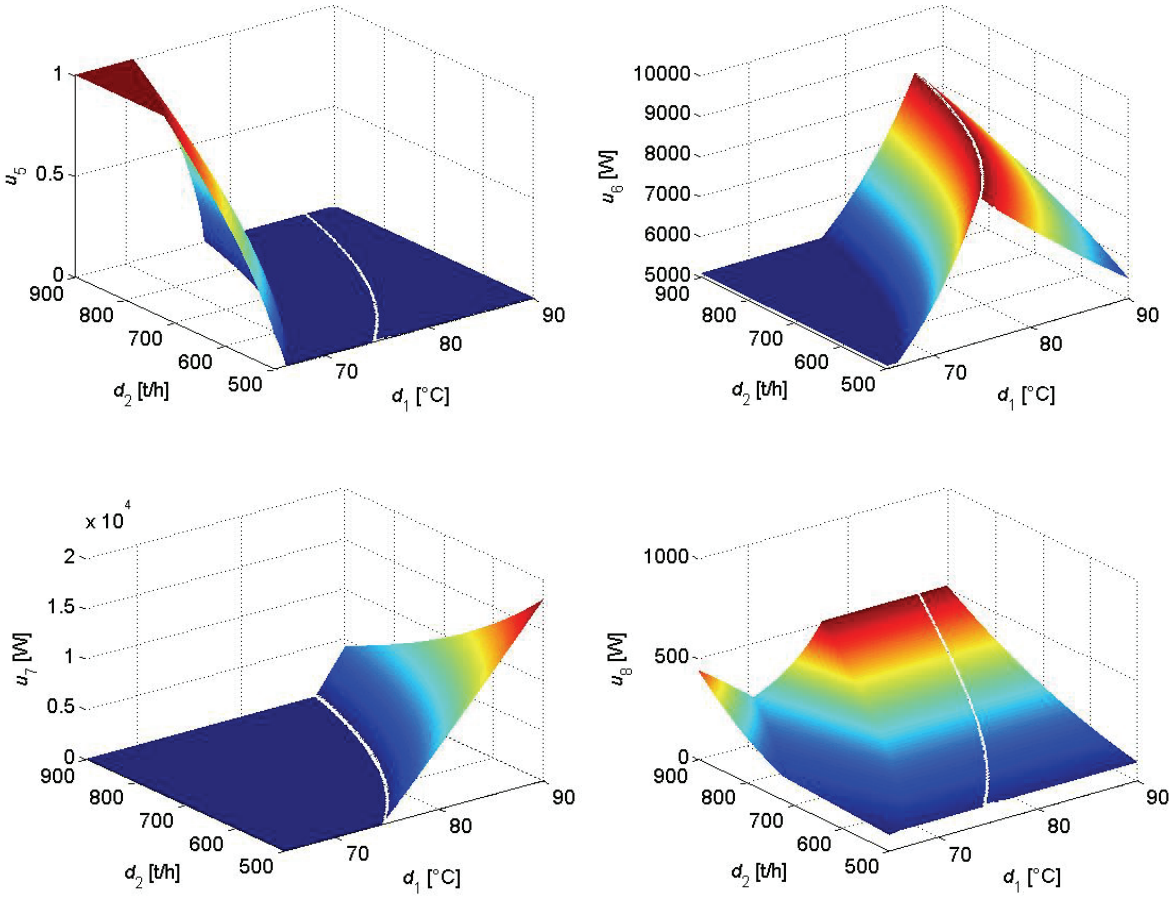

Figure 3.2: Optimal values of inputs as a function of d_1 and d_2.
The valve positions, u_1 to u_5, have been normalized.

Examination of where the manipulated variables were at constraints revealed that the plant operation could be divided into four distinct regions, named α, β, γ, and δ (Figure 3.3). The optimal values of the inputs and outputs in these regions are given in Tables 3.1 and 3.2. A missing entry means that the variable is varying.

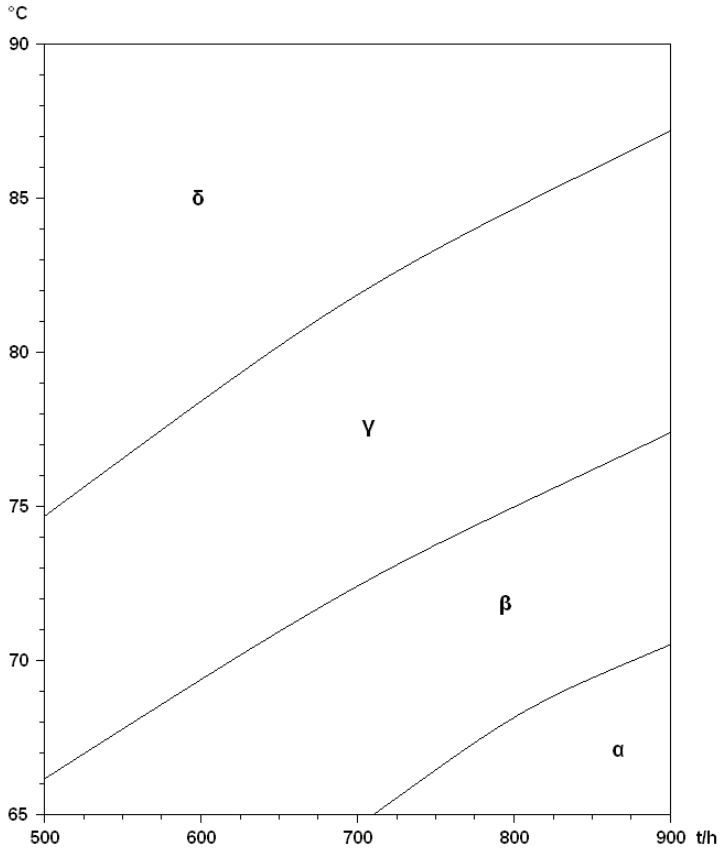

Figure 3.3: Representation of operational regions at the Brobekk plant as a function of d_1 (y axis) and d_2 (x axis).

Table 3.1: Optimal values for inputs in the Brobekk model. The numbers represent percent values of the maximum value.

Region	u_1 [%]	u_2 [%]	u_3 [%]	u_4 [%]	u_5 [%]	u_6 [%]	u_7 [%]	u_8 [%]
α	92.6	0	100		100	6.78	0	
β	92.6	0	100	100		6.78	0	
γ		0	100	100	0		0	
δ	0		100	100	0			

Table 3.2: Optimal values for the outputs in the Brobekk model.

Region	y_1 [°C]	y_2 [°C]	y_3 [°C]	y_4 [°C]	y_5 [°C]	y_6 [t/h]
α	126			-		250
β	126			-		250
γ	126			-		250
δ	126	126		126		250

Using the information in the tables above, we can explain the four regions as follows:

- *Region α*
  The conditions (d_1 and d_2) at Viken's side allow them to take out all the produced energy. Thus, the air coolers are not active (i.e. u_2 and u_7 are zero). In order to minimize pressure drop, the remaining valves on the Brobekk side are either fully open or close to fully open. The same reasoning is true for Viken's side. An additional effect on their side is the large flow rates which are encountered in this region: In order not to cool down the plant, the bypass u_5 must be fully open.

  Based on Equation 3.2, the number of degrees of freedom available for optimization is:

$$\begin{aligned} N_{opt,free} &= 14 - 2\times u_2 - 2\times u_3 - 1\times u_5 - 2\times u_7 - 1\times u_8 - 2\times y_1 - 2\times y_6 \\ &= 2 \end{aligned}$$

- *Region β*
  In this region two Viken side phenomena that may heat up the plant are encountered: In one scenario the flow rate decreases to 500 t/h, while in the other scenario the temperature goes to above 75 °C. Both of these two situations generally imply a decreased energy need. The answer is to start closing bypass valve u_5 while opening u_4 completely, resulting in the heat transfer in the primary heat exchangers being maintained. The DOF analysis gives

$$\begin{aligned} N_{opt,free} &= 14 - 2\times u_2 - 2\times u_3 - 2\times u_4 - 2\times u_7 - 1\times u_8 - 2\times y_1 - 2\times y_6 \\ &= 1 \end{aligned}$$

- *Region γ*
  The risk of overheating is now increasing. The bypass valve u_1 thus starts to close, while bypass valve u_5 is now completely closed. Valve u_4 is fully open in order to maintain the heat transfer. The DOF analysis shows that there are 0 degrees of freedom left for optimization:

$$\begin{aligned} N_{opt,free} &= 14 - 2\times u_2 - 2\times u_3 - 2\times u_4 - 1\times u_5 - 2\times u_7 - 1\times u_8 - 2\times y_1 - 2\times y_6 \\ &= 0 \end{aligned}$$

- *Region δ*
  With a surplus of energy in the system the air coolers become active. All bypass valves are now completely closed, and u_2 and u_7 come into action. This region also has 0 degrees of freedom left for optimization:

$$\begin{aligned} N_{opt,free} &= 14 - 2\times u_1 - 2\times u_3 - 2\times u_4 - 1\times u_5 - 1\times u_8 - 2\times y_1 - 2\times y_2 - 2\times y_6 \\ &= 0 \end{aligned}$$

From the four regions one point in each region was chosen as that region's nominal operating point (Table 3.3). These points were chosen such that they would lie close to, or in, the center of the region.

Table 3.3: Center points for the four operational regions.

Region	d_1 [°C]	d_2 [t/h]
α	67.0	840.8
β	70.0	700.0
γ	77.5	700.0
δ	86.0	740.0

3.4 CONTROL STRUCTURE DESIGN

With optimal input and output values obtained from the optimization, and with the DOF analysis in Chapter 3.2 in mind, one can finally start designing a control structure for the plant. Because the plant operation turned out to be divided into regions, separate control structures were devised for each region. For all regions u_6 controls y_6, while u_8 defines the Viken side flow rate, d_2.

3.4.1 REGION α

Region α is characterized by low cold water temperatures, and medium to high cold water flow rates. The DOF analysis showed that we have two degrees of freedom available for optimization, and from Table 3.1 we have two manipulated variables available; u_1 and u_4. Among the outputs we need to control y_1. One alternative is to keep u_1 constant – as Table 3.1 recommends – and then use u_4 for control of y_1. This control structure is shown in Figure 3.4. The steady-state gain k for the pairing u_4-y_1 is -39.35, which seems promising. However, because of the relatively large time constant (Figure C.5), the control will be somewhat slow.

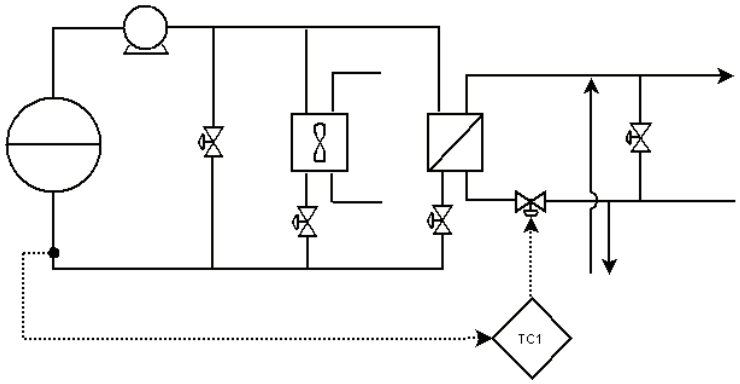

Figure 3.4: Region α control structure with u_4 controlling y_1.

Since the optimization was performed at steady-state, it was not concerned with how to get from one steady-state to another. In a dynamic model this is naturally important, and the control structure in Figure 3.4 may thus not be ideal because of its slow control. An alternative is to use the bypass valve u_1 for control of y_1, and the use the free DOF for optimization (one per line) for control of u_1 through input resetting (Figure 3.6). The steady-state gain between u_1 and y_1 is rather good, with k = 11.63, but – more importantly – the time constant is much smaller than between u_4 and y_1 (Figure C.1). The idea behind input resetting

is that one controller (K_2) is used for fast control by manipulating one input (u_2), while another controller (K_1) is tasked – on a larger time scale – with bringing u_2 back to its nominal value by manipulating another input (u_1). This can be illustrated as in Figure 3.5.

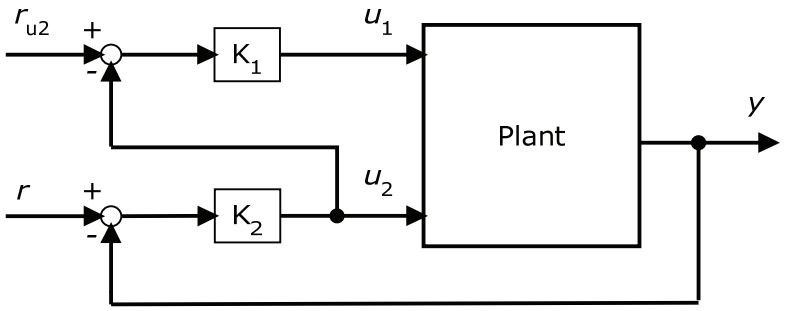

Figure 3.5: Input resetting structure (Skogestad and Postlethwaite, 2005).

Such a control structure is often implemented if u_2 is either too costly to be used over long periods of time, or if its power is limited, for instance due to valve saturation. In our case the motivation is that the Brobekk side bypass valve should be optimally non-varying.

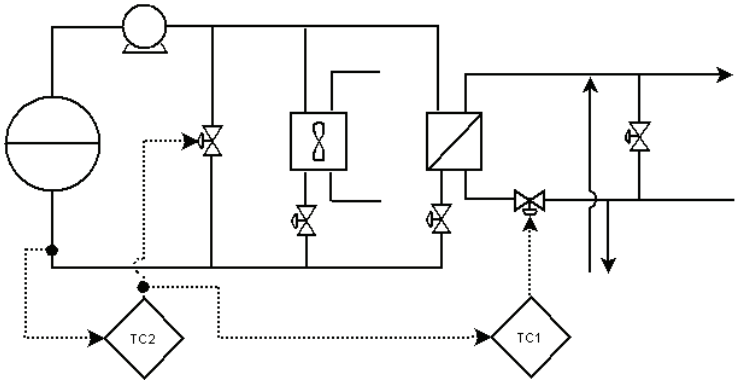

Figure 3.6: Region α control structure with u_1 controlling y_1, and u_4 controlling u_1.

3.4.2 REGION β

Region β is characterized by low to medium cold water temperatures, and low to high cold water flow rates. Based on reasoning similar to that in the previous region, u_5 – which has now replaced u_4 as an available MV – should be used to control y_1. However, in this region we have only one degree of freedom left for optimization, and the impact of this now becomes clear: There is only one valve u_5, but two heat exchanger lines. The question is therefore: Which y_1 to control? The answer is that our assumption of line symmetry makes it indifferent which one we choose. However, in reality, as soon as one of the furnaces experiences a disturbance, for instance in the flue gas temperature, the strategy of using u_5 for control of y_1 needs some modifications. Chapter 5.3 discusses this topic further.

From u_5 to y_1 the steady-state gain is only 3.217, and the response is also slow (Figure C.9). Based on this, it was chosen to test the input resetting strategy in this region as well. This involved the same reasoning as above, but once more our assumption of symmetry

allowed us to choose the line we desired. These two control proposals are illustrated in Figures 3.7 and 3.8. The behavior of these control structures is presented in Chapter 4.2.

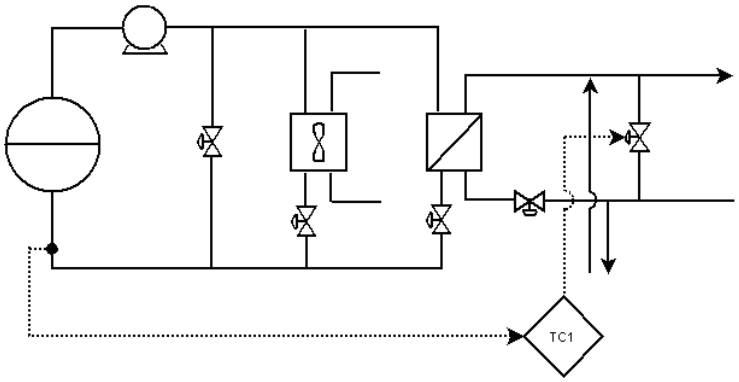

Figure 3.7: Region β control structure with u_5 controlling y_1.

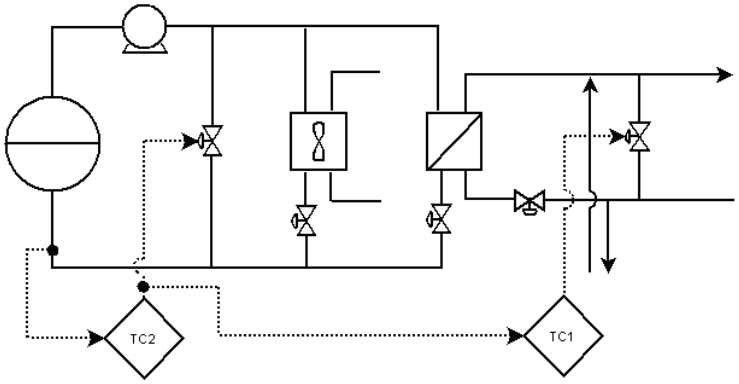

Figure 3.8: Region β control structure with u_1 controlling y_1, and u_5 controlling u_1.

3.4.3 REGION γ

Region γ is characterized by low to high cold water temperatures, and low to high cold water flow rates. This region has zero degrees of freedom available for optimization; only input u_1 is available for manipulation, and only output y_1 needs to be controlled. This is illustrated in Figure 3.9. Chapter 4.3 presents the behavior of this control structure.

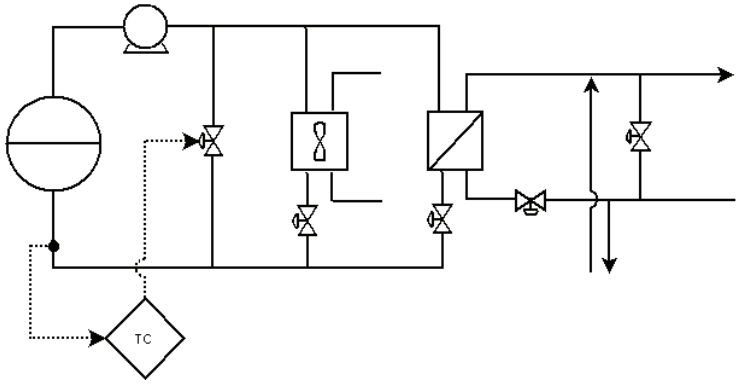

Figure 3.9: Region γ control structure with u_1 controlling y_1.

3.4.4 REGION δ

Region δ is characterized by medium to high cold water temperatures, and low to high cold water flow rates, a combination which requires the air coolers to be active. The manipulated variables available for control in this region are u_2 and u_7, and there are three candidate controlled outputs; y_1, y_2, and y_4. At the center point of the region we have the following steady-state gains (Table 3.4):

Table 3.4: Steady-state gains at the center point in region δ.

	u_2	u_7
y_1	41.22	-9.968×10^{-4}
y_2	-8.335	-3.980×10^{-4}
y_4	461.6	-3.450×10^{-2}

The reason for the low gains between u_7 and the three outputs is the low heat capacity of air relative to that of water. Purely based on Table 3.4 the recommendation is to pair u_2 with y_1, and u_7 with y_4. The same conclusion can be drawn from an RGA analysis (Bristol, 1966) on this system, as shown below.

An RGA analysis is a method for measuring the interaction between inputs and outputs under the assumption that all other loops than the one being analyzed are perfectly controlled. The RGA is defined as the element-wise product (also known as the Schur or Hadamard product) between the square steady-state gain matrix **G** and the transpose of the inverse of itself:

$$\mathbf{\Lambda} = \mathbf{G} \otimes \left(\mathbf{G}^{-1}\right)^{\mathrm{T}} \tag{3.4}$$

The RGA matrix has the property that all its rows and columns sum to one. Input-output pairing should be done on elements close to unity, as this ensures that we choose to control the output which the input in question affects the most. Ideally all other elements should be zero. In practice, however, this is very often not the case – a sign that interaction between the loops may occur. Elements larger than one preserve the dynamic response, but will also reduce the gain due to closed loop interactions (Kontogiannis et al., 1999; University of Edinburgh, 2008). Pairing on negative elements should be avoided, as this will give an extremely poor dynamic response (Kontogiannis et al., 1999).

As mentioned, Equation 3.4 requires that **G** is square, and it can thus not be applied directly to Table 3.4. We therefore divide this non-square system into three square sub-sets, and then apply Equation 3.4 to each of these. This is shown in Table 3.5. Alternatively, one could have used the generalized RGA or GRGA, which is an extension of the RGA to non-square systems (Chang and Yu, 1990; Skogestad and Postlethwaite, 2005), and in which the inverse of **G** is replaced by the Moore-Penrose pseudoinverse instead. However, in order to easily compare the different alternatives, the first method is chosen.

Table 3.5: RGA matrices for region δ.

	u_2	u_7
y_1	0.6638	0.3362
y_2	0.3362	0.6638

	u_2	u_7
y_2	0.6102	0.3898
y_4	0.3898	0.6102

	u_2	u_7
y_1	1.4783	-0.4783
y_4	-0.4783	1.4783

Since it is advised to pair on elements close to unity in the RGA matrix, we once more find that the previously mentioned pairing – u_2-y_1 and u_7-y_4 – is recommended. However, the RGA method above bases its recommendations purely on the steady-state gains, which corresponds to a very low frequency. We therefore do not know if this recommendation is also correct for higher frequencies. Thus, a frequency dependent RGA analysis (Skogestad and Postlethwaite, 2005) is performed. This is possible by replacing the elements in **G** in Equation 3.4 with transfer functions instead of the steady-state gains. A second interaction measure which will be used is the RGA number (Skogestad and Postlethwaite, 2005). For the diagonal pairing the RGA number is calculated like this;

$$\text{RGA number} = \left\| \Lambda(G) - \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \right\|_{\text{sum}} \tag{3.5}$$

whereas for off-diagonal pairing the following equation is used:

$$\text{RGA number} = \left\| \Lambda(G) - \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \right\|_{\text{sum}} \tag{3.6}$$

When evaluating the RGA numbers the rule is to choose the pairing which gives the smallest RGA number in the desired frequency region.

In Figures 3.10 to 3.12 the RGA matrices in Table 3.5 are represented as a function of frequency. The diagonal and off-diagonal RGA numbers are also included in the figures.

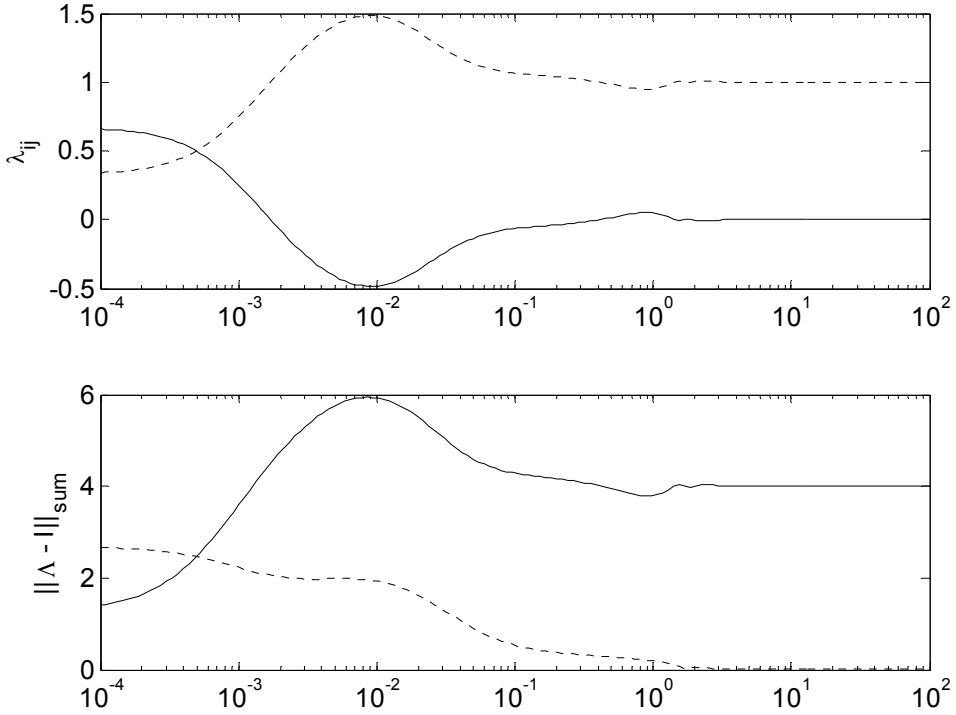

Figure 3.10: RGA elements (top) and RGA numbers (bottom) for the system $\{u_2, u_7, y_1, y_2\}$ as a function of frequency [rad/s]. The solid line represents diagonal pairing, while the dashed line represents off-diagonal pairing.

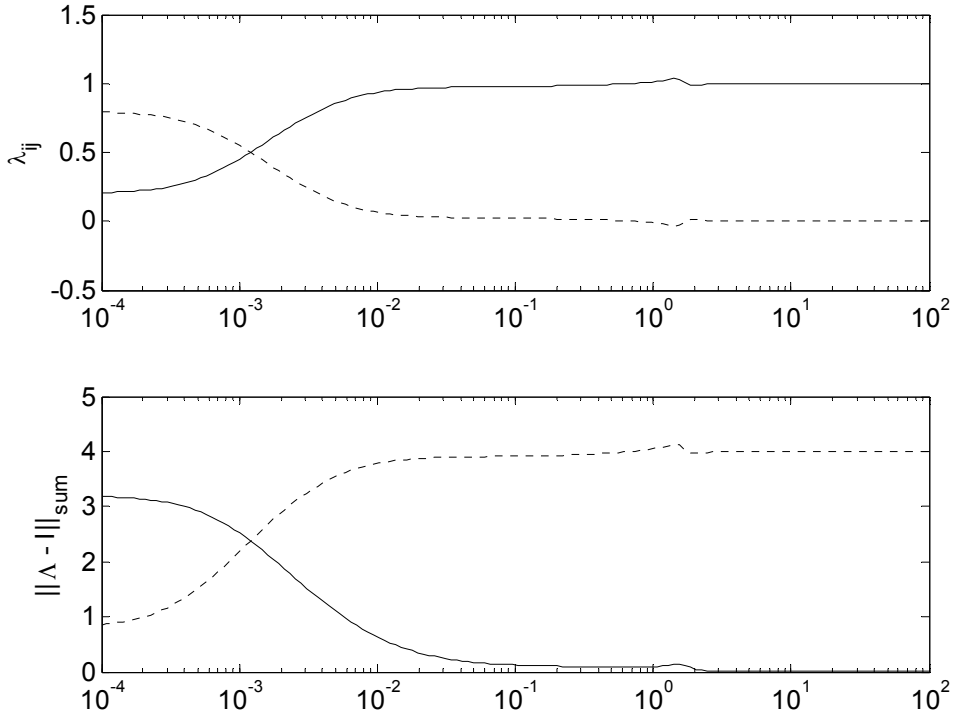

Figure 3.11: RGA elements (top) and RGA numbers (bottom) for the system $\{u_2, u_7, y_2, y_4\}$ as a function of frequency [rad/s]. The solid line represents diagonal pairing, while the dashed line represents off-diagonal pairing.

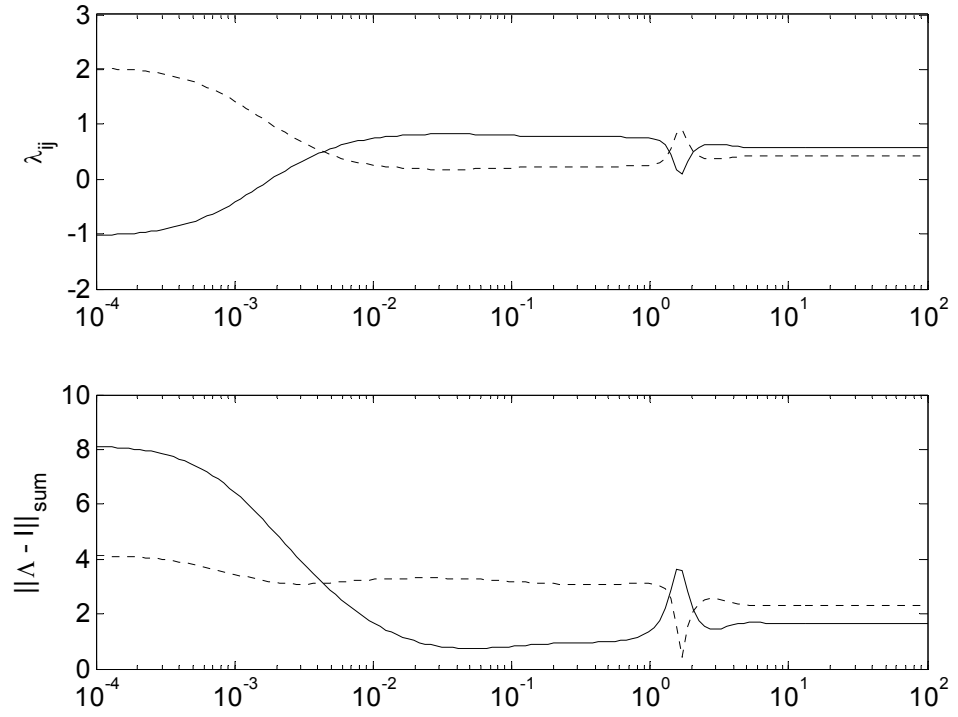

Figure 3.12: RGA elements (top) and RGA numbers (bottom) for the system $\{u_2, u_7, y_1, y_4\}$ as a function of frequency [rad/s]. The solid line represents diagonal elements or pairing, while the dashed line represents off-diagonal elements or pairing.

As the disturbances are of low frequency, we will seek to obtain good control of the plant at low frequencies. If we study the graphs in Figures 3.10 to 3.12 we can draw the following conclusions:

- Figure 3.10 reveals that neither diagonal nor off-diagonal pairing is an option for low frequencies, since the elements λ_{11} and λ_{21} are almost equal.

- Figure 3.11 seems somewhat more promising, with reduced interactions at low frequencies compared to Figure 3.10.

- Figure 3.12 – and specifically its off-diagonal pairing, u_2-y_4 and u_7-y_1 – seems to be the best alternative. Although this corresponds to pairing on elements larger than unity, this alternative avoids pairing on negative elements. In addition, this choice gives the lowest RGA number. Therefore, this pairing alternative is chosen for testing in Chapter 4.4. This choice is illustrated in Figure 3.13.

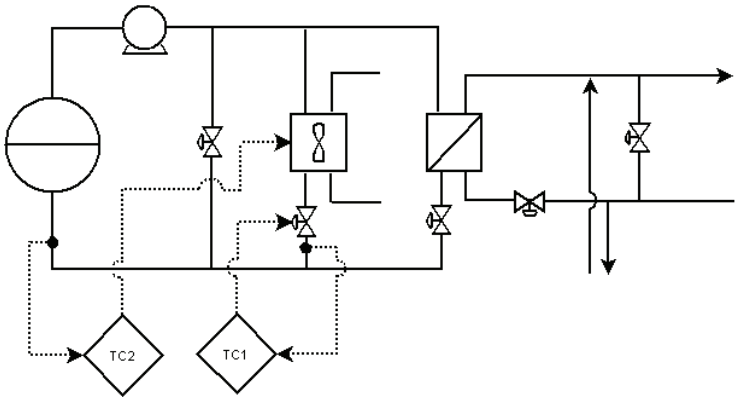

Figure 3.13: Region δ control structure with u_2 controlling y_4, and u_7 controlling y_1.

4. DYNAMIC SIMULATIONS

In order to test the proposed control structures for the four regions, region specific sequences of step disturbances were created. The disturbances were first passed through a filter (a first order transfer function with τ = 100 s) in order to remove high frequency components. The aim of the control structures was to keep y_1 within ± 1 °C, while the flow rate y_6 should be kept between 230 t/h and 260 t/h (Mordt, 2007).

Preliminary simulations showed that the flow dependent heat transfer coefficients (see Section 2.2.1) ideally also should have been dependent on the primary side inlet temperature and flow rate, and not only on the secondary side flow rate. The result was a small mismatch between the output values in the dynamic model and the steady-state model given the same input values. This was solved by adjusting the input values in the dynamic model slightly. Since the regions were defined based on where the inputs were at a constraint, this adjustment only resulted in the region borders being moved upwards about 1-1.5 °C, and did not invalidate the concept of operational regions.

4.1 REGION α

In region α the following disturbance sequence was used:

- d_1 from 67 °C to 65 °C
- d_2 from 840.8 t/h to 722.4 t/h
- d_2 from 722.4 t/h to 900 t/h
- d_1 from 65 °C to 70.5 °C
- d_1 from 70.5 °C to 65 °C, and d_2 from 900 t/h to 722.4 t/h

In Figure 4.1 only u_4 is used for control of y_1. This would be the ideal control strategy from a steady-state point of view, with a constant pump duty in lines 1 and 2. The maximum absolute deviation in y_1 is about 0.12 °C (middle, left), but the deviations endure for about 500 seconds each. This is due to the relatively large time constant between u_4 and y_1, as already discussed, as well as the slow dynamics of the disturbances themselves.

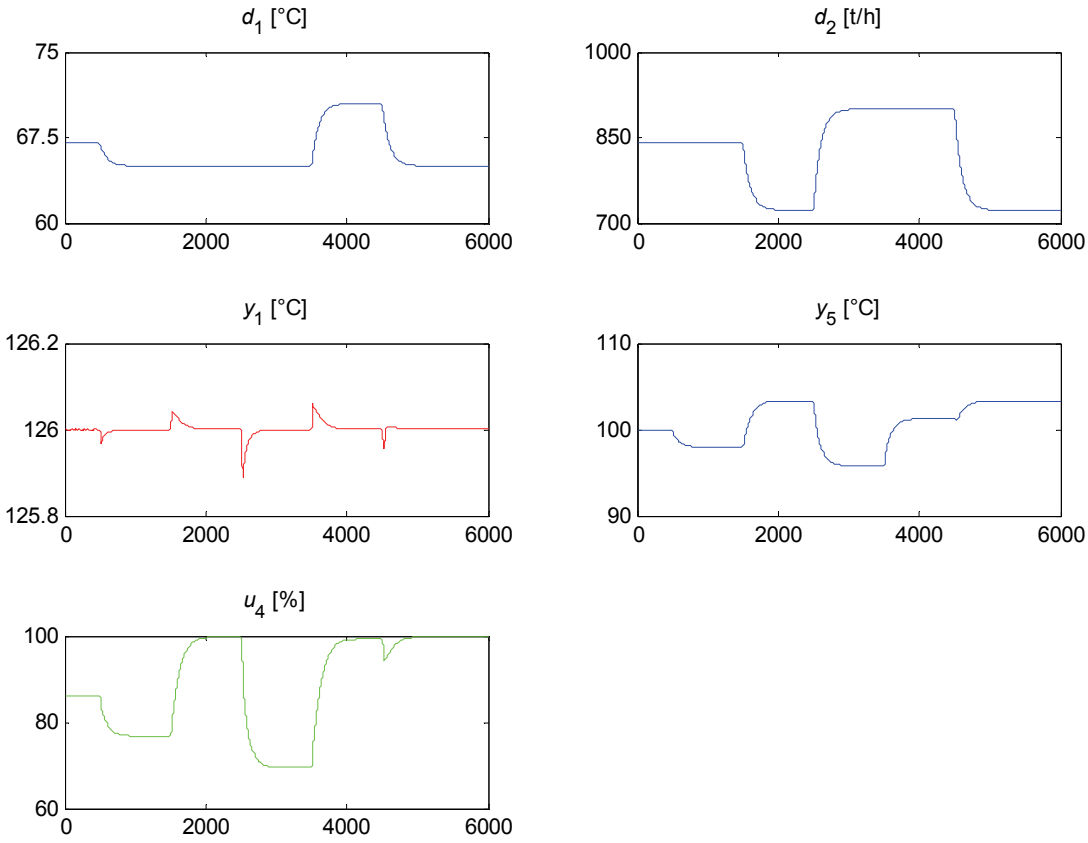

Figure 4.1: Closed-loop responses in region α when using u_4 to control y_1. The x axis represents time [s].

In Figure 4.2 input resetting is used. While u_1 takes care of the disturbances on a fast time scale, u_4 is tasked – on a slower time scale – with bringing u_1 back to its nominal value. This gave a significantly improved response in y_1 (middle, left), and u_1 was reset back to its optimal value after the transient. Based on these two results, it was chosen to use the input resetting control structure in region α.

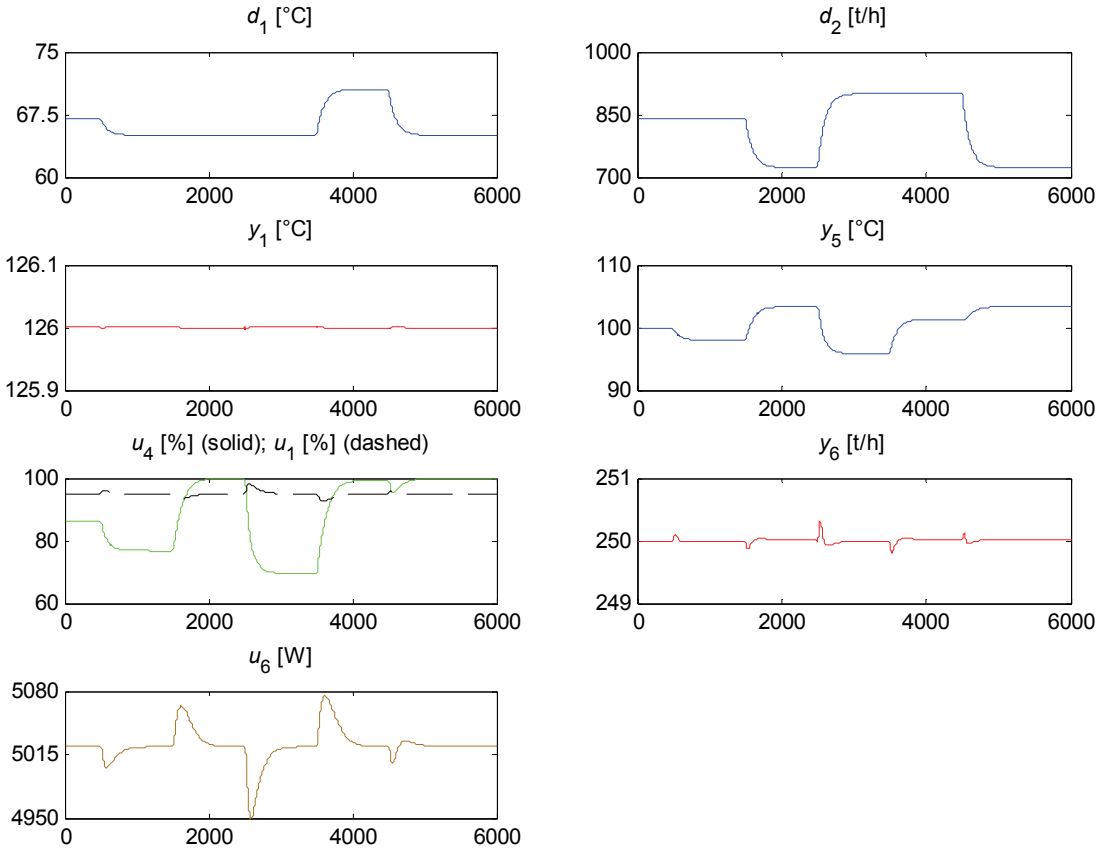

Figure 4.2: Closed-loop responses in region α when using u_1 to control y_1, u_4 to control u_1, and u_6 to control y_6. The x axis represents time [s].

From approximations of step responses using first order transfer functions, steady-state gains k and time constants τ_1 were obtained. The controllers were then tuned using the SIMC rules for PI controllers (Skogestad, 2003), where the controller gain and the integral time are calculated as follows:

$$K_c = \frac{\tau_1}{k\left(\tau_c + \theta\right)} \tag{4.1}$$

$$\tau_I = \min\left\{\tau_1, 4\left(\tau_c + \theta\right)\right\} \tag{4.2}$$

If the model has a known dead time, θ, then τ_c is the only tuning parameter. However, in the models obtained here, there was no dead time, and the sum of τ_c and θ was used as a tuning parameter instead.

The obtained controller settings are presented in Table 4.1. Controller FC manipulated u_6; controller TC/TC$_1$ manipulated u_4; and controller TC$_2$ manipulated u_1. Controller TC$_1$ was tuned using K_c and τ_I independently, as the steady-state gain k in Equation 4.1 was zero for the pairing u_4-u_1.

Table 4.1: Controller settings in the region α control structures.

Standard SISO control	$\tau_c + \theta$	K_c	τ_I
TC	3.54	-0.500	14.16
Input resetting	$\tau_c + \theta$	K_c	τ_I
TC_1	-	3.600	69.56
TC_2	0.16	0.500	0.635
FC	20.0	37.34	2.090

4.2 REGION β

Similar to region α, the proposed control structures in region β were tested against a sequence of five disturbances:

- d_2 from 700 t/h to 875 t/h
- d_1 from 70 °C to 77.4 °C, and d_2 from 875 t/h to 900 t/h
- d_1 from 77.4 °C to 72 °C, and d_2 from 900 t/h to 775 t/h
- d_1 from 72 °C to 65 °C, and d_2 from 775 t/h to 600 t/h
- d_2 from 600 t/h to 500 t/h

The results using the control structures in Section 3.4.2 are shown in Figures 4.3 and 4.4.

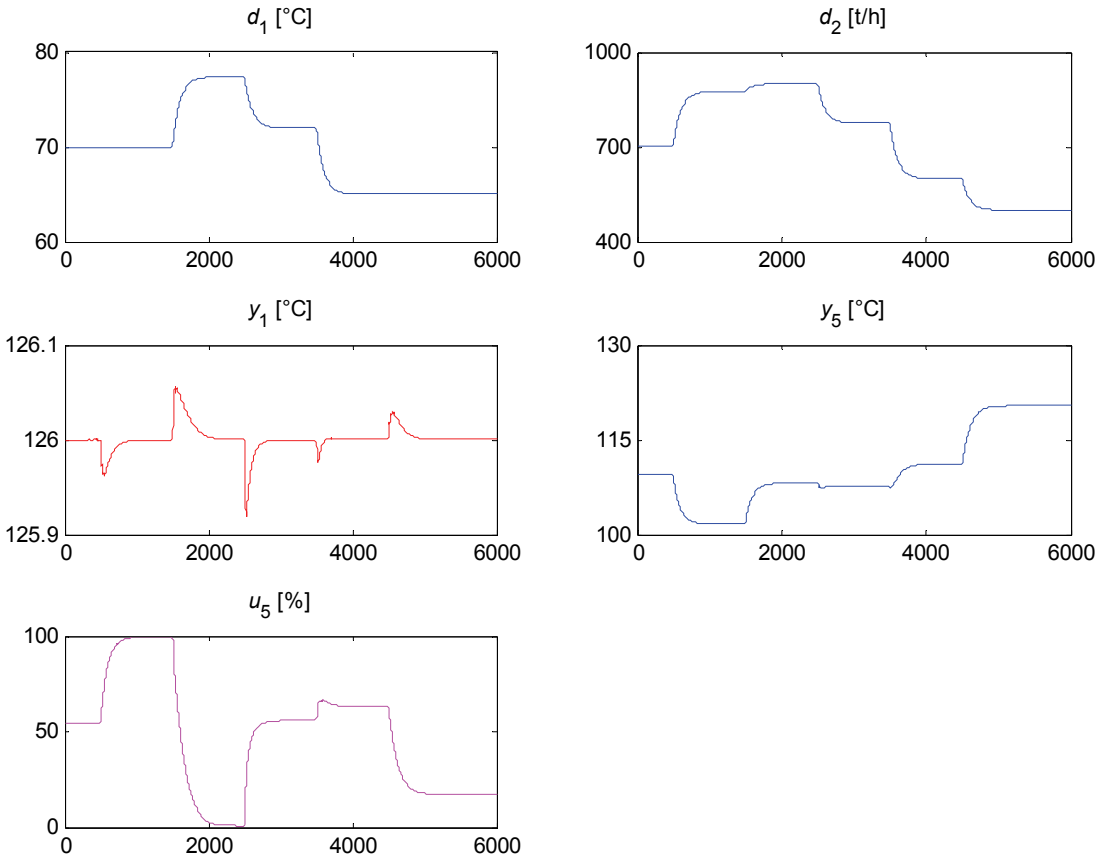

Figure 4.3: Closed-loop responses in region β when using u_5 to control y_1. The x axis represents time [s].

The response in y_1 is very similar to what is shown in Figure 4.1, and from a step response perspective this is what could be expected (compare Figures C.5 and C.9). Since none of the valves on the Brobekk side of the plant are manipulated, this control structure will be cost-efficient, but control can be improved by using input resetting (Figure 4.4).

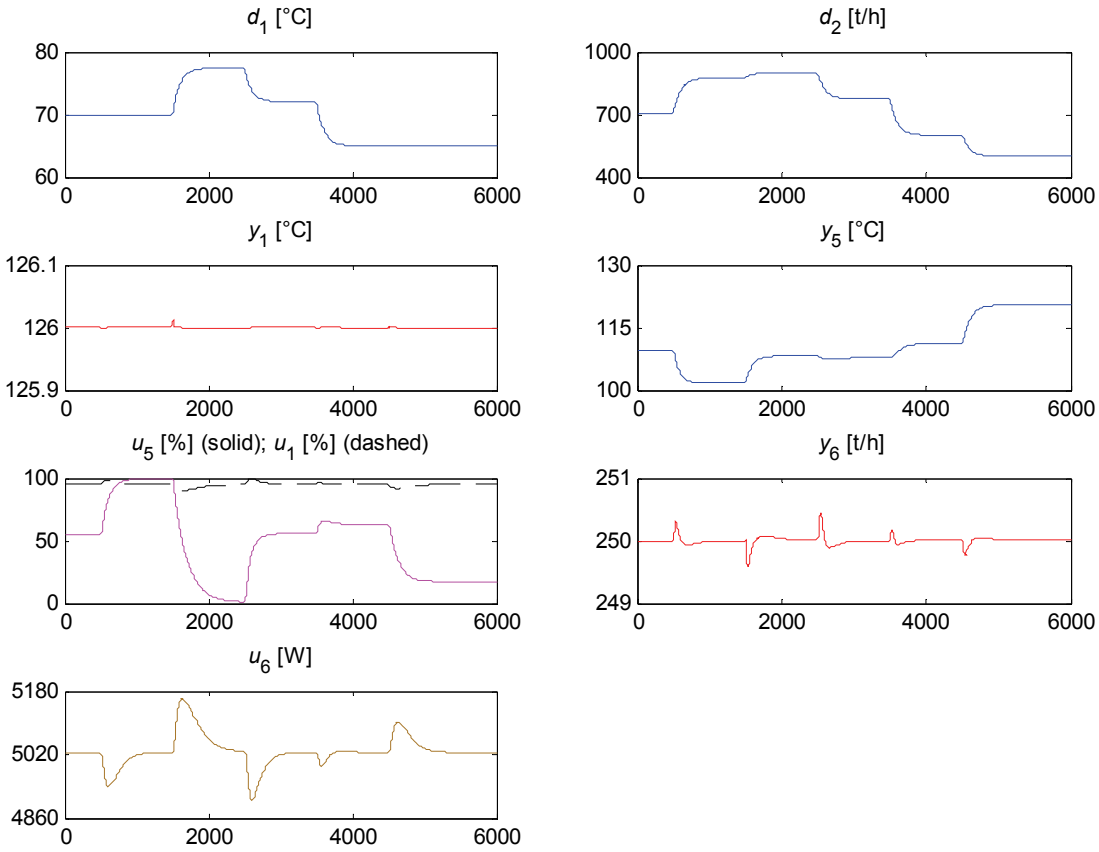

Figure 4.4: Closed-loop responses in region β when using u_1 to control y_1, u_5 to control u_1, and u_6 to control y_6. The x axis represents time [s].

When input resetting is being used, the control of y_1 is tightened, with very small deviations and faster settling times (middle, left). However, since u_1 is now being used for control, we experience deviations in y_6. Based on the observations from the two figures above we can thus conclude that it is most promising to use the input resetting strategy for region β. Below is presented the tuning parameters for the controllers in the two control structures. Controller FC manipulated u_6; controller TC/TC_1 manipulated u_5; and controller TC_2 manipulated u_1.

Table 4.2: Controller settings in the region β control structures.

Standard SISO control	$\tau_c + \theta$	K_c	τ_I
TC	7.89	2.500	31.54
Input resetting	$\tau_c + \theta$	K_c	τ_I
TC_1	-	-4.500	90.00
TC_2	0.16	0.500	0.635
FC	20.0	37.34	2.090

4.3 REGION γ

In the third region the following disturbance sequence was used:

- d_2 from 700 t/h to 900 t/h
- d_1 from 77.5 °C to 87.2 °C
- d_1 from 87.2 °C to 77.5 °C, and d_2 from 900 t/h to 700 t/h
- d_1 from 77.5 °C to 73.5 °C
- d_2 from 700 t/h to 500 t/h
- d_1 from 73.5 °C to 66.1 °C

The result from this simulation is shown in Figure 4.5. The control of y_1 is very good, but the control of y_6 is less tight, due to heavy valve usage. Table 4.3 shows the tuning parameters for the controllers. Controller TC manipulated u_1, while controller FC manipulated u_6.

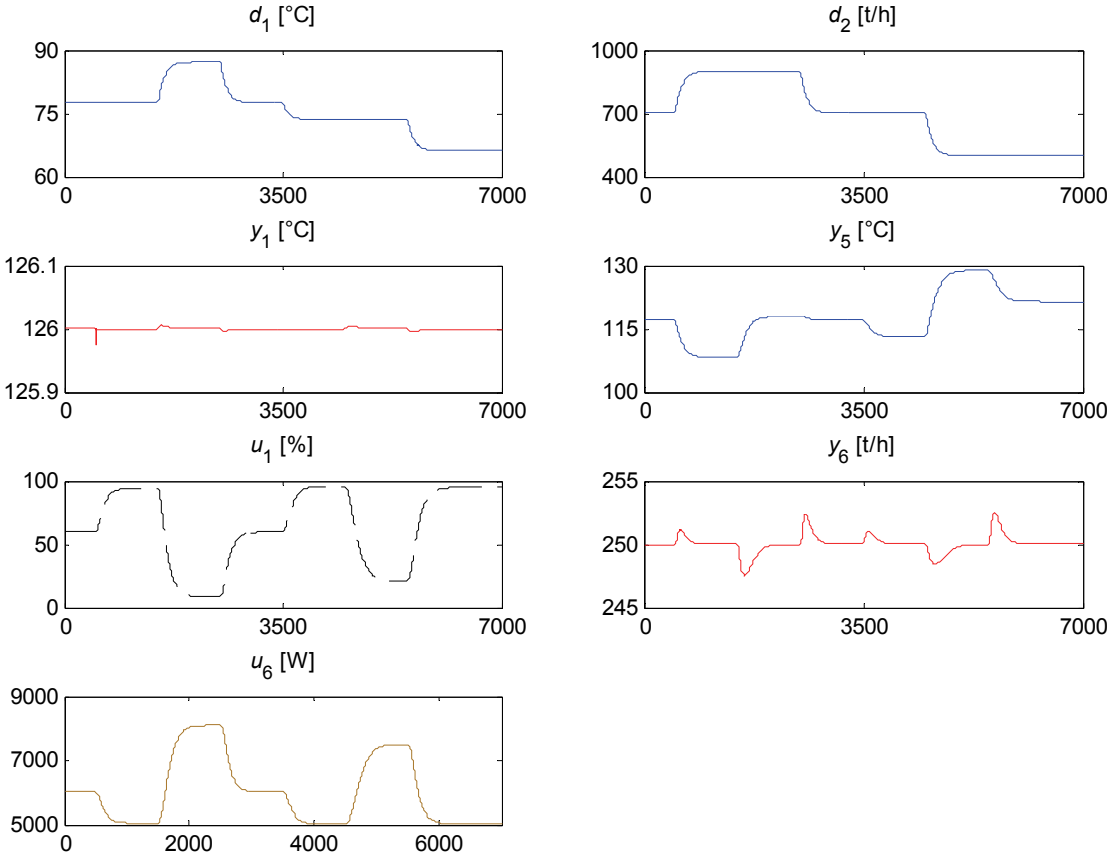

Figure 4.5: Closed-loop responses in region γ when using u_1 to control y_1. The x axis represents time [s].

Table 4.3: Controller settings in the region γ control structure.

Standard SISO control	$\tau_c + \theta$	K_c	τ_I
TC	0.16	0.500	0.635
FC	20.0	37.34	2.09

4.4 REGION δ

The candidate control structure from Section 3.4.4, namely to pair u_2 with y_4, and u_7 with y_1, was tested against the following disturbance sequence:

- d_1 from 86 °C to 90 °C, and d_2 from 740 t/h to 500 t/h
- d_1 from 90 °C to 77 °C
- d_1 from 77 °C to 90 °C, and d_2 from 500 t/h to 800 t/h
- d_2 from 800 t/h to 900 t/h
- d_1 from 90 °C to 88.2 °C

The results are shown in Figure 4.6. Both the first and the second disturbances are very large, as is reflected in y_1. The maximum deviations are, however, relatively small; only about ± 0.3 °C, and well inside what is allowed. Since u_2 is now used for control, there are larger than usual deviations in y_6. It can also be observed that y_5 peaks at 135 °C, which corresponds to the highest possible d_1 and the lowest possible d_2 (i.e. the point where the air cooler duty is at its maximum).

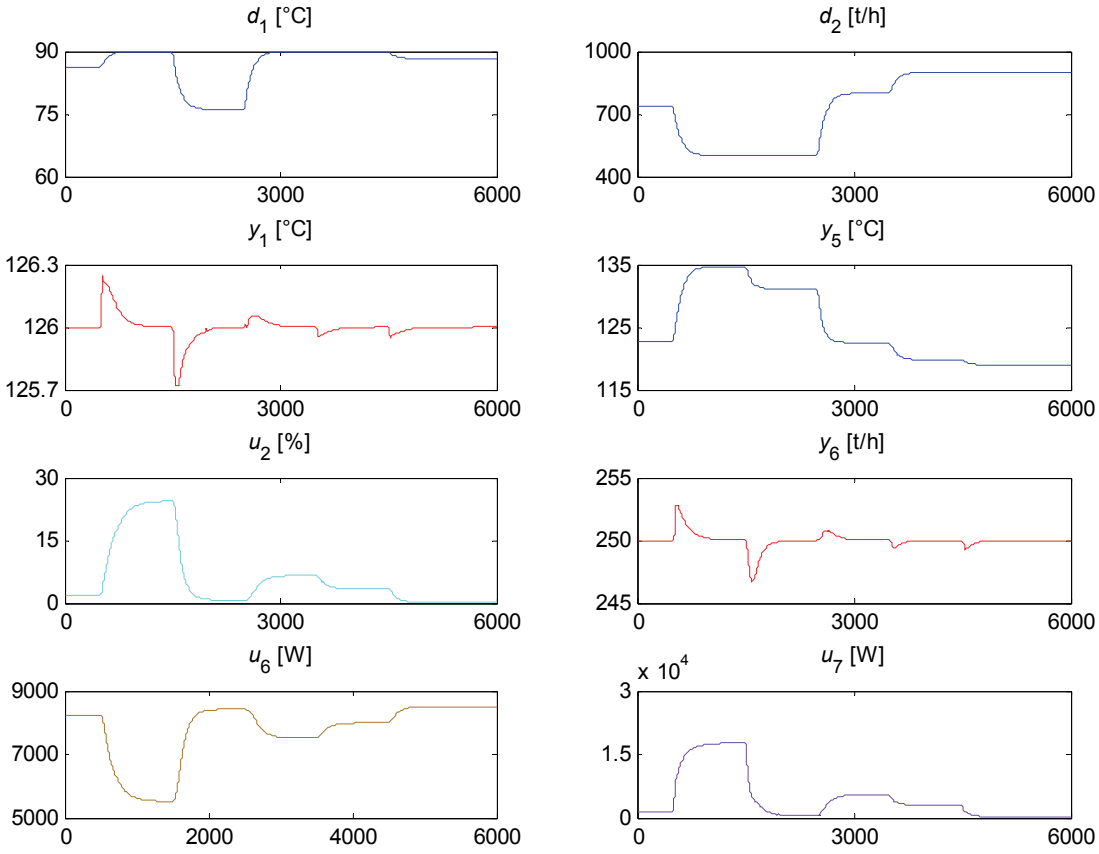

Figure 4.6: Closed-loop responses in region δ when using u_2 to control y_4, and u_7 to control y_1. The x axis represents time [s].

The controllers were found to perform satisfactorily with the tuning parameters given in Table 4.4. The large controller gain between u_7 and y_1 is due to the low steady-state gain, as discussed in Section 3.4.4. Note also that the $\tau_c + \theta$ tuning did not work for this controller. For the K_c used, this method recommended a τ_I of 109.2 seconds. This value, however, resulted in large oscillations in y_6, likely because of the interactions in this region. A value of 80 seconds for τ_I – still using the same K_c – worked much better. Controller TC_1 manipulated u_2; controller TC_2 manipulated u_7; and controller FC manipulated u_6.

Table 4.4: Controller settings in region δ control structure.

Standard SISO control	$\tau_c + \theta$	K_c	τ_I
TC_1	20.0	0.008	76.39
TC_2	-	-25000	80.00
FC	20.0	37.34	2.090

4.5 MULTI-REGION MODEL

4.5.1 DETAILS OF THE MULTI-REGION MODEL

The control structures that were chosen based on the simulations in the previous chapters were implemented into one joint Simulink model (Figures 5.7 and 5.8), and a Stateflow chart was tasked with the region identification. Only one line was implemented, as line symmetry was assumed during the optimization (see Section 3.3.1). This reduced the computational load significantly.

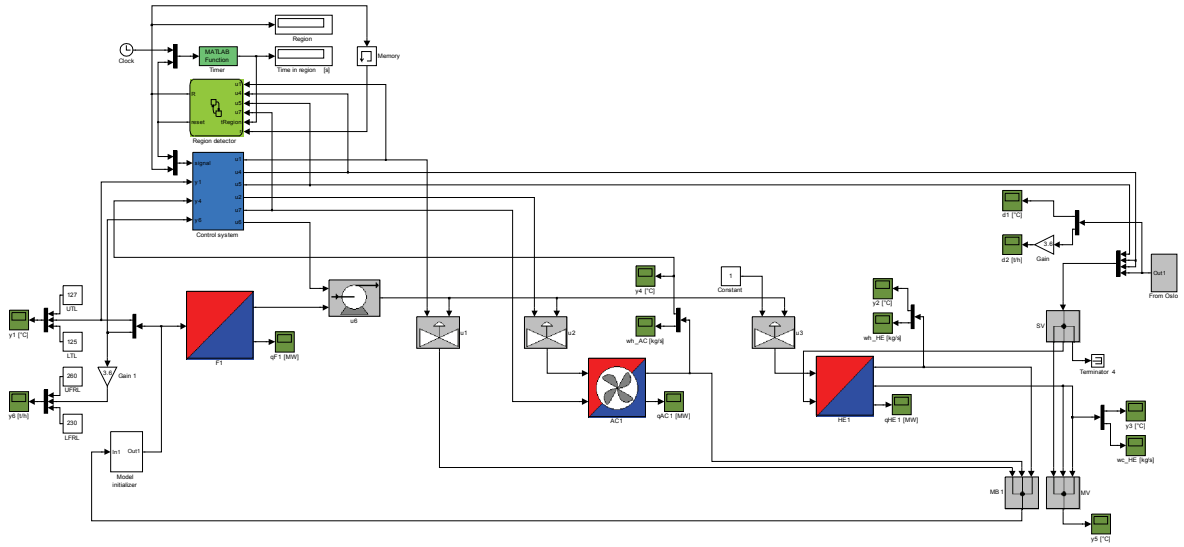

Figure 4.7: Simulink model of Brobekk incineration plant. Line symmetry is assumed. The model units are explained in the text.

The blue block in Figure 4.7 represents the control system; the light green block is the Stateflow chart; and the mint green block is a script which keeps track of how long the plant has been operating in its current region. The grey blocks are pumps, valves, splitters, and mixers; the red and blue blocks are the furnace, air cooler and primary heat exchanger, respectively; and the dark green blocks are signal plotters. The white blocks are constants, signal terminators, a model initializer, and gains. The Stateflow chart consisted of four states, representing the four operational regions, and six transitions, as shown in Table 4.5.

Table 4.5: Stateflow description of the six region transitions. tRegion is the amount of time that has passed since entering a region (in seconds), while r is the previous region number.

Transition	Command
$\alpha \rightarrow \beta$	[u4 = = 1 && tRegion >= 150 && r = = 1]
$\beta \rightarrow \gamma$	[u5 = = 0 && tRegion >= 150 && r = = 2]
$\gamma \rightarrow \delta$	[u1 = = 0 && tRegion >= 150 && r = = 3]
$\delta \rightarrow \gamma$	[u7 = = 0 && tRegion >= 150 && r = = 4]
$\gamma \rightarrow \beta$	[u1 >= 0.948 && tRegion >= 150 && r = = 3]
$\beta \rightarrow \alpha$	[u5 = = 1 && tRegion >= 150 && r = = 2]

The blue control system block – which is shown in closer detail in Figure 4.8 – contained the control structure itself. A signal from the Stateflow chart determined what controllers were to be active. During a transition from one region to another the integrator elements of the PI controllers were reset in order to reduce the amount of spikes or dips in the inputs.

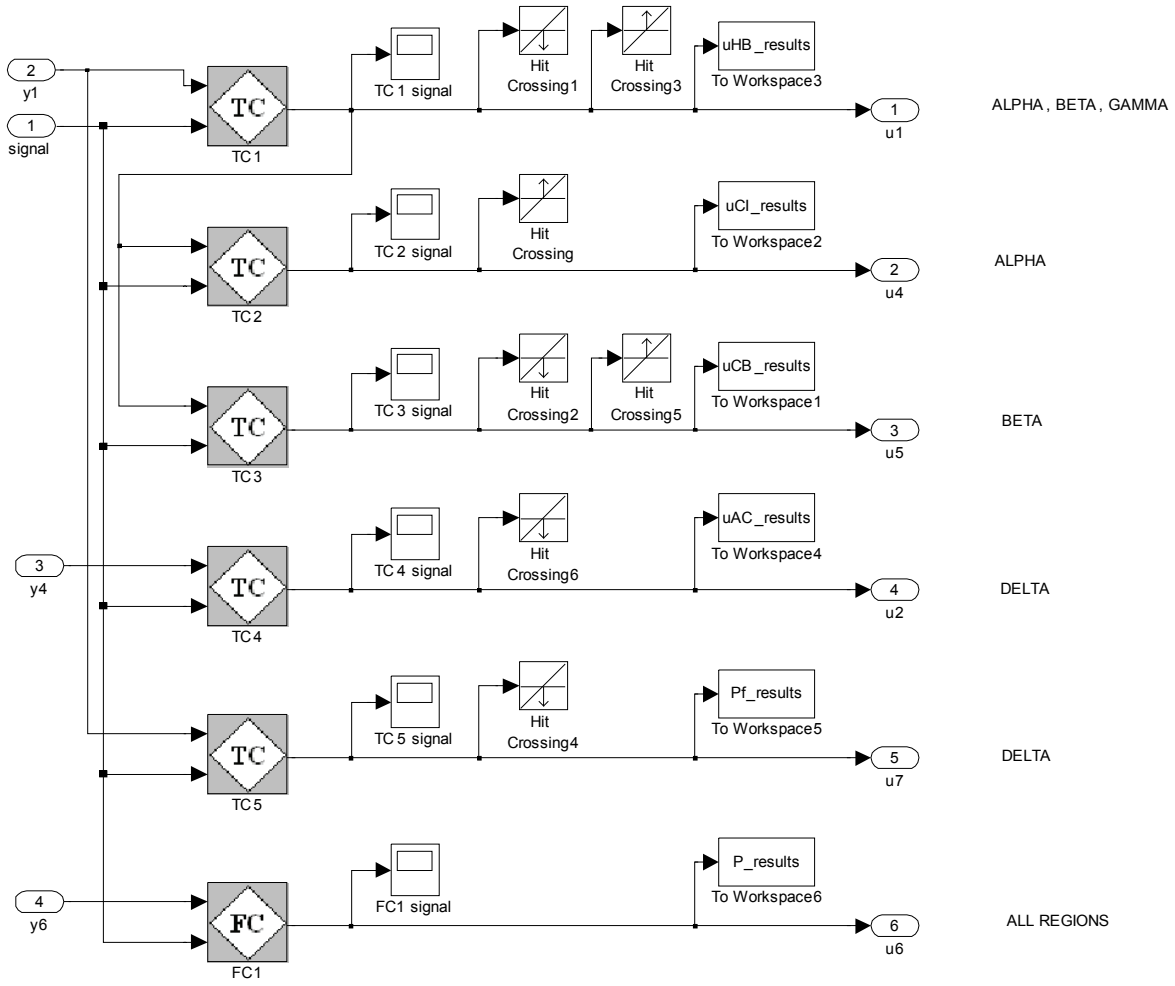

Figure 4.8: The proposed control structure for Brobekk incineration plant.
The grey blocks are the PI controllers. The `Hit Crossing` blocks are included to ensure that the variable-step solver computes time steps just before and after the transitions for increased reliability.

5.6.2 SIMULATION RESULTS

In Chapters 4.1 to 4.4 it was shown how the single-region control structures behaved for various disturbances when applied to the dynamic model. In this chapter we will look at the transitions between the control structures. In the first example (Figure 4.9) we start in the center point of region α and make a step in d_1 from 67 °C to 74 °C, which brings the plant operation into region β.

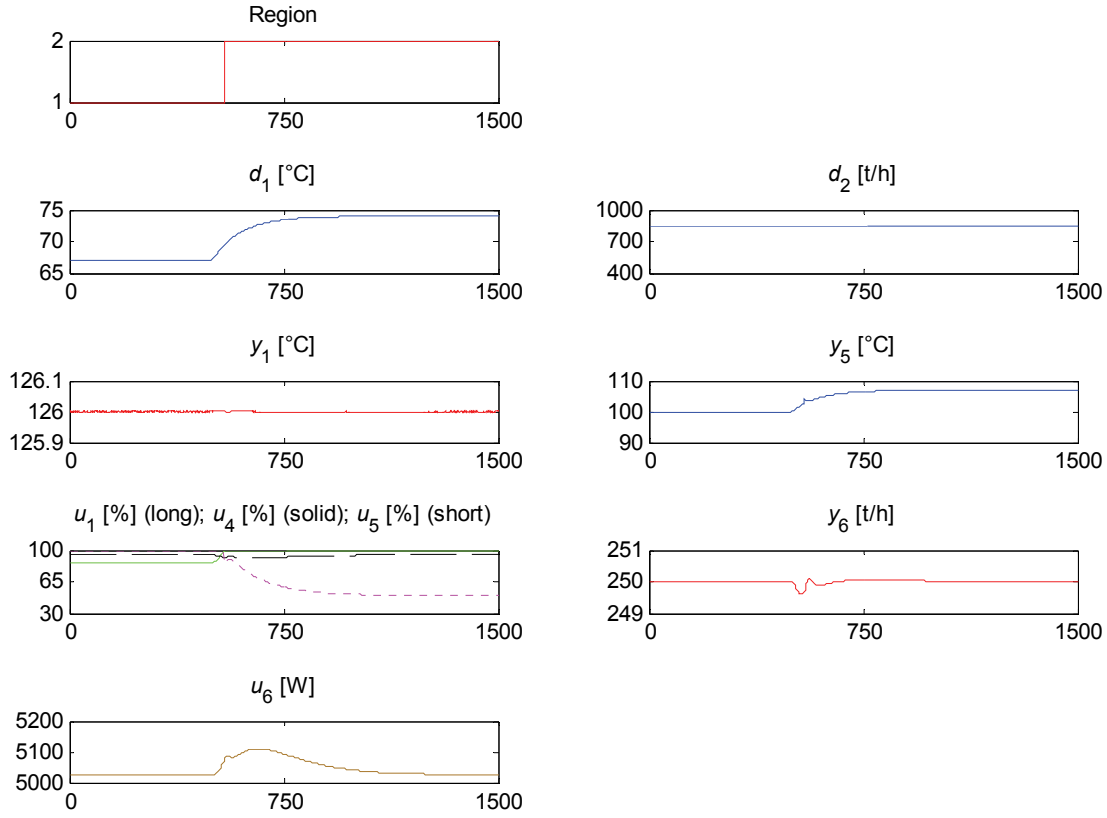

Figure 4.9: Transition from region α to region β after a step in the Viken side inlet temperature.

The transition from region α to region β went smoothly, and the step had virtually no effect on y_1. However, some noise can be seen. The smooth transition is due to the fact that the control structures in these two regions are similar; both use input resetting on u_1. Therefore, as u_4 saturates (green curve), u_5 takes over control of u_1 by starting to close (magenta curve). The result is tight control of y_1 and very little use of u_1 (black curve), the latter resulting in only small deviations in y_6.

Figure 4.10 shows an example of a step which leads to a transition between region β and region γ; from its nominal value of 70 °C, d_1 is increased by 11 °C, and we end up relatively close to the border towards region δ. During the transition a small, but rapid, spike occurs in u_1, causing the flow rate to increase. Because of this, the upper limit for y_6 (i.e. 260 t/h) is exceeded by 5 t/h. This situation also results in a spike in y_1, since the bypass is transiently more open than it should.

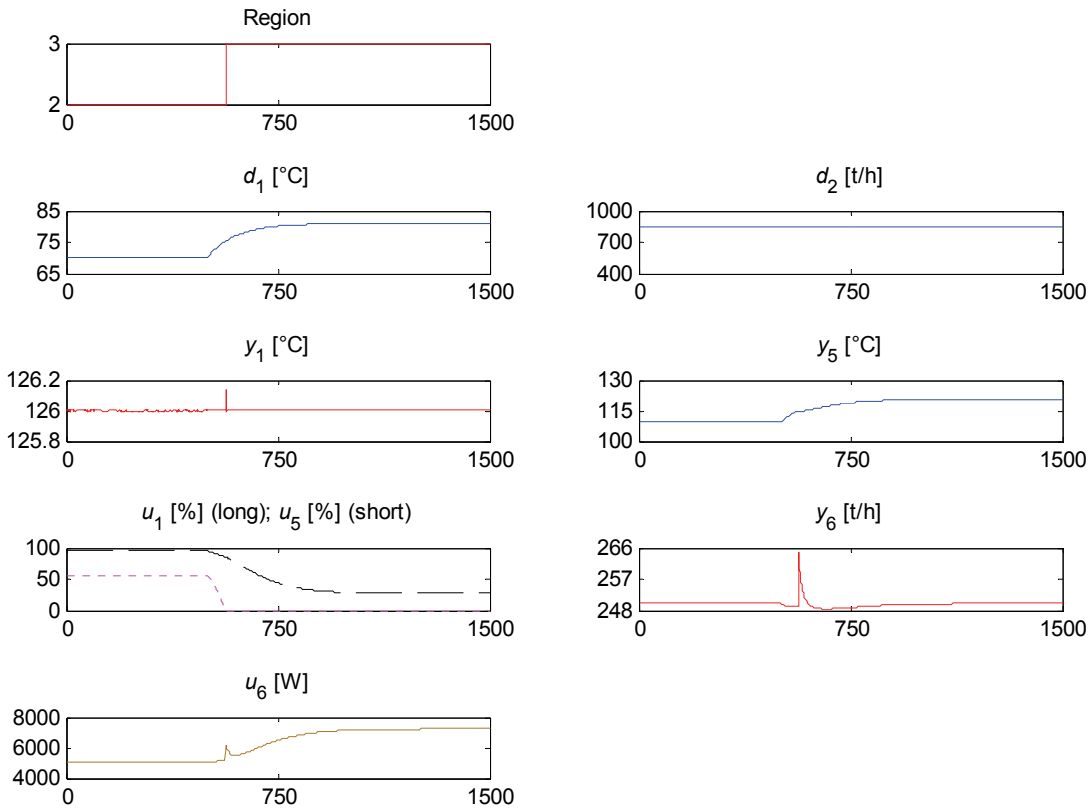

Figure 4.10: Transition from region β to region γ after a step in the Viken side inlet temperature.

The four last transitions, from β to α, from γ to β, from γ to δ, and from δ to γ, all gave unsatisfactory results and will not be presented here. As an example, from region γ to region δ large and sustaining oscillations in y_1 were observed, and from region γ to region β a steady-state offset was observed.

5. DISCUSSION

5.1 OPTIMIZATION

The steady-state optimization showed that the plant operation space could be divided into four distinct regions, defined by where sets of the manipulated variables were at constraints. The assumption of line symmetry was regarded as valid due to three phenomena: First, the two furnaces have identical steady-state heat duties. Second, it is counter-productive to distribute the cold water from Oslo unevenly between the lines, as the split flows are joined again before leaving the plant. And third; neither the temperature nor the flow rate of the flue gas in the furnaces was considered as a disturbance.

Due to the goal of the optimization, the temperature in the flow towards Oslo/Viken Fjernvarme would necessarily vary. It was assumed that Viken Fjernvarme can add additional heat to its district heating network if it needs more than the 32 MW produced at Brobekk. However, Viken Fjernvarme has expressed concerns regarding the temperature towards Oslo being too low (below 105 °C) in a part of the operational space. If some slack is implemented on the furnace inlet temperature during the optimization, it can receive water with a somewhat higher temperature than in this study.

5.2 SINGLE-REGION CONTROL STRUCTURES

REGION α

Based on two promising control structures it was chosen to use an input resetting strategy in region α, with u_1 controlling y_1, and u_4 controlling u_1. This gave very good results, with very low deviations in y_1, as well as fast settling times. In u_6 the difference between the lowest and highest duties was only 130 W, implying a very small load on the pump.

According to the optimization, u_1 should not be used for control, but rather be kept constant. This strategy was also tested, i.e. with only u_4 controlling y_1. Due to a slower response, the control was slower and with increased deviations.

REGION β

In region β the situation was similar to in region α, only with u_4 replaced by u_5. This reduced the number of degrees of freedom available for optimization from two to one. Since there are two lines we are technically lacking one degree of freedom. The strategy of using input resetting was only possible due to the fact that there were no disturbances in the furnaces, and the assumption of perfect symmetry. However, for a real application, one possibility would be to use input resetting for only one line – thus accepting sub-optimality – and use a SISO control loop for the other line. On the other hand, such a strategy will fail if the valve which is not reset goes into saturation.

An example may clarify the situation: Assume that only furnace 1 experiences an increased flue gas temperature, which, in turn leads to an increase in $y_{1,1}$. The bypass $u_{1,1}$ starts to close, and – if input resetting is being used in line 1 – u_5 will begin to close. This

sends more water through the u_4 valves, and $y_{1,1}$ will eventually be brought back to its set point. However, in line 2 – where there has been no disturbance in the furnace – the situation is different. Since there is no need for more cold water, $y_{1,2}$ gets lower. Bypass valve $u_{1,2}$ starts to open, but will quickly reach its limit, due to its high nominal value in this region (~ 94 %). Thus, $y_{1,2}$ will not be brought back to its set point, and will instead stabilize at a lower temperature.

In the real plant it would be advisable to control a combination of the two outputs, e.g. $½(u_{1,1} + u_{1,2})$. This will ensure that the controller sees the impact of varying u_5 on both lines. Simultaneously, the inputs $u_{1,1}$ and $u_{1,2}$ are used to control $y_{1,1}$ and $y_{1,2}$, respectively. When symmetry of the two lines is assumed, then all values and results from the previous optimization hold for the system with both lines.

Generally, the results from the simulation were quite similar to those from the input resetting control structure in region α, with very small deviations in y_1. The pump load was also almost the same as in region α.

REGION γ

Region γ turned out to have zero degrees of freedom for optimization. With only u_1 available for manipulation, and only y_1 left to be controlled, the choice of control structure was a standard SISO loop. The simulations showed very good results for y_1, but with larger deviations in y_6. The pump duty varied between 5000 W and about 8250 W, but with relatively slow dynamics. However, how fast the pump can rev up or down must be taken into account in a possible real application of this control structure.

REGION δ

The forth region turned out to be the most complicated, with two manipulated variables available for control (u_2 and u_7) and three potential controlled variables (y_1, y_2, and y_4). A frequency dependent relative gain array analysis was conducted in order to determine the pairing. Since the control would be aimed for low frequencies, it turned out that the recommended pairing was u_2-y_4 and u_7-y_1. Subsequent simulations showed that the deviations in y_1 were also here well inside the ± 1 °C band. This region also showed large variations in u_6, as well as in u_7.

The relative gain array analysis revealed that some interactions could be expected, especially for other pairing combinations than the one chosen. This is probably due to the fact that u_2 and u_7 must be used simultaneously; if for instance u_2 opens to fast – before the fan has started or gained enough revolutions – this valve will function more as a bypass rather than a supplier of cooled water. For the chosen pairing, however, this behavior was not seen.

5.3 MULTI-REGION CONTROL STRUCTURE

The experienced difficulty with the transition from one region to another is contributed to the resetting of the integrator part of the PI controllers. Since all controllers manipulating a certain input measure the corresponding output regardless of whether the controller is active or not, the integral part may integrate up the error and give a completely wrong signal immediately after the controller is activated (i.e. when a transition has taken place). The idea is that if the integrators are reset properly, this problem will be removed, as the first signal from the newly activated controller equals the last signal from the previously active controller (given that they manipulate the same input). However, this has not proven an easy task, and it seems that this is an area for further study.

In the model, valve saturation dictated when to switch from one region to another. An alternative strategy would be to measure the disturbances and compare these against a look-up table of region versus Viken side inlet temperature and flow rate, or against a curve fitting of the operating points along the region boundaries.

6. CONCLUSION

Based on the results from the steady-state optimization, the operational space of the plant could be divided into four distinctive regions. Four separate decentralized control structures were then created. The two first regions, named α and β, were controlled using an input resetting strategy, while the third and fourth region, named γ and δ, were controlled using regular SISO control loops. The optimization also showed that the Brobekk side heat exchanger valve should always be fully open, and not be used for control. At the time of writing of this report, however, this valve is a vital part of the real plant's control structure. Also, the optimization revealed that the Brobekk side bypass valve is never fully open, indicating that it has sufficient capacity. In the real plant, however, this valve is planned to be replaced with a valve with larger capacity. A revision of the input/output pairing in the rest of the plant may render this upgrade superfluous.

Implementation of the control structures in the dynamic Simulink model gave very good results when no region transitions were taking place. The overall maximum deviation in the furnace inlet temperature – encountered in region δ – was only 30 % of the allowed deviation. In regions α, β, and γ, the deviations in the same output were virtually immeasurable.

Implementation of transitions between the control structures revealed that only the transition from region α to β was performing as expected. The other five transitions gave varying degrees of undesired spikes and dips in the input variables.

Overall, it is concluded that the proposed control structure is able to control the part of the plant studied in this report to within its specified limits in a simple and robust way. However, this requires that Energigjenvinningsetaten is given access to the three valves on Viken Fjernvarme's side of the plant. This strategy will ensure that Viken Fjernvarme always receives only 32 MW, thus efficiently removing the problem of cooling of the plant. If Viken Fjernvarme needs more energy than this, it needs to use its own heaters.

Based on this work, the following topics seem interesting for further research:

1. The steady-state optimization was considering the plant from the point of view of Energigjenvinningsetaten. The question is: How different will the results be if Viken Fjernvarme is considered more important than now – or even *the* most important party?

2. Region β will need a different control structure as soon as either of the furnaces experiences a disturbance. Whatever this structure may look like, it will have to consider both lines simultaneously, as the Viken side bypass valve affects both lines if manipulated.

3. The transitions between the different control structures were not performing as desired due to the integrator resetting in the PI controllers. A consistent way to perform this resetting, regardless of which region one is entering, should be developed.

REFERENCES

Bristol, E. H. (1966), *On a New Measure of Interaction for Multivariable Process Control*. IEEE Trans. Auto. Control, AC-11, 133-134.

Chang, J.-W.; Yu, C.-C. (1990), *The relative gain for non-square multivariable systems*. Chem. Eng. Sci., 45(5), 1309-1323.

Hertzberg, T. (2008), *Lecture notes*.

Kontogiannis, E.; Munro, E.; Impram, T. S. (1999), *Frequency Domain Control Structure Design Tools* in *Progress in system and robot analysis and control design*, Springer, Berlin.

Mathisen, K.W.; Morari, M.; Skogestad, S. (1993), *Dynamic Models for Heat Exchangers and Heat Exchanger Networks*, ESCAPE 3, Graz, Austria.

Mordt, H. (2007), *Private correspondence*.

RenoSam; Rambøll (2006), *The most efficient waste management system in Europe: Waste-to-energy in Denmark*.

Roald, A. S. (2001), *Kjemiteknikk I*, revised edition.

Skogestad, S. (2000), *Plantwide control: the search for the self-optimizing control structure*. J. Proc. Cont., 10, 487-507.

Skogestad, S. (2003), *Simple analytic rules for model reduction and PID controller tuning*. J. Proc. Cont., 13, 291-309.

Skogestad, S.; Postlethwaite, I. (2005), *Multivariable Feedback Control – Analysis and Design*, 2nd ed., John Wiley & Sons, Ltd.

The MathWorks, Inc. (2008), *Stateflow 7.1: Design and simulate state machines and control logic*. http://www.mathworks.com/products/stateflow/.

Tomlab Optimization (2008), *TOMLAB/MINOS v6.0*. http://tomopt.com/tomlab/products/minos/

University of Edinburgh (2008), *The ECOSSE Control HyperCourse*. http://eweb.chemeng.ed.ac.uk/courses/control/restricted/course/index.html

APPENDIX

A. ABBREVIATIONS

AC	Air cooler
DHN	District heating network
DOF	Degree of freedom
EGE	Energigjenvinningsetaten, Oslo kommune
GRGA	Generalized relative gain array
HE	Heat exchanger
IT	Information technology
LM	Logarithmic mean
NTU	Number of transfer units
RGA	Relative gain array
SISO	Single input, single output
WIP	Waste incineration plant

B. LIST OF SYMBOLS

Latin symbols

Symbol	Description	Unit
A	area	m^2
c_p	specific heat capacity	$J\ kg^{-1}\ K^{-1}$
d	disturbance	...
$\mathbf{D}$	heat exchanger transfer function matrix	-
$\mathbf{G}$	steady-state gain matrix	...
h	heat transfer coefficient	$W\ m^{-2}\ K^{-1}$
k	gain	...
K_V	flow factor	$m^3\ h^{-1}\ bar^1$
N	number of cells	-
p	pressure	Pa
P	power	W
q	conducted heat	W
Q	volumetric flow rate	$m^3\ h^{-1}$
r	set point	...
S	NTU method variable	-
t	time	s
T	temperature	°C
u	input; normalized valve position	...; -
U	overall heat transfer coefficient	$W\ m^{-2}\ K^{-1}$
V	volume	m^3
w	mass flow rate	$kg\ s^{-1}$
y	output	...
z	fraction of air velocity remaining in diffusor	-

Greek symbols

Symbol	Description	Unit
α	number of transfer units	-
β	heat capacity flow rate ratio	-
γ	NTU method variable	-
δ	NTU method variable	-
ε	Flow direction coefficient	-
$\mathbf{\Lambda}$	RGA matrix	-
η	efficiency; number of transfer units	-
ρ	density	kg m^{-3}
τ	time constant	s

Superscript symbols

Symbol	Description
c	cold side; cold
h	hot side; hot
T	transpose
w	wall side

Subscript symbols

Symbol	Description
°	reference state
c	controller
i	index variable; in
I	integral
j	index variable
o	out
p	constant pressure

C. SELECTED STEP RESPONSES

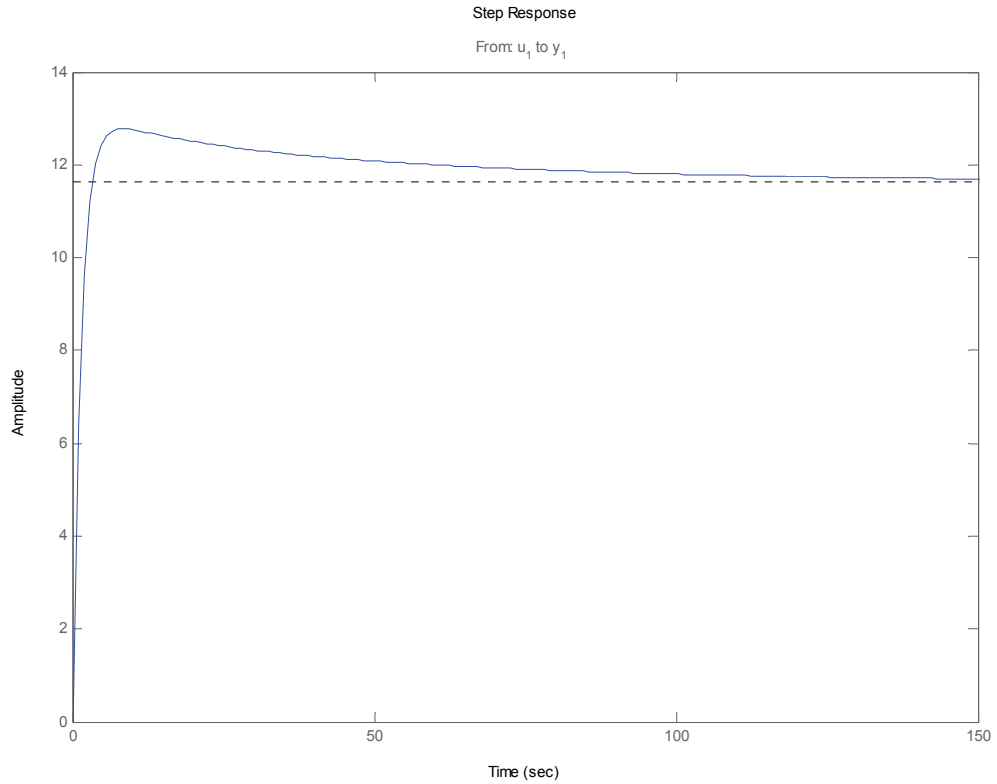

Figure C.1: Step response from u_1 to y_1.

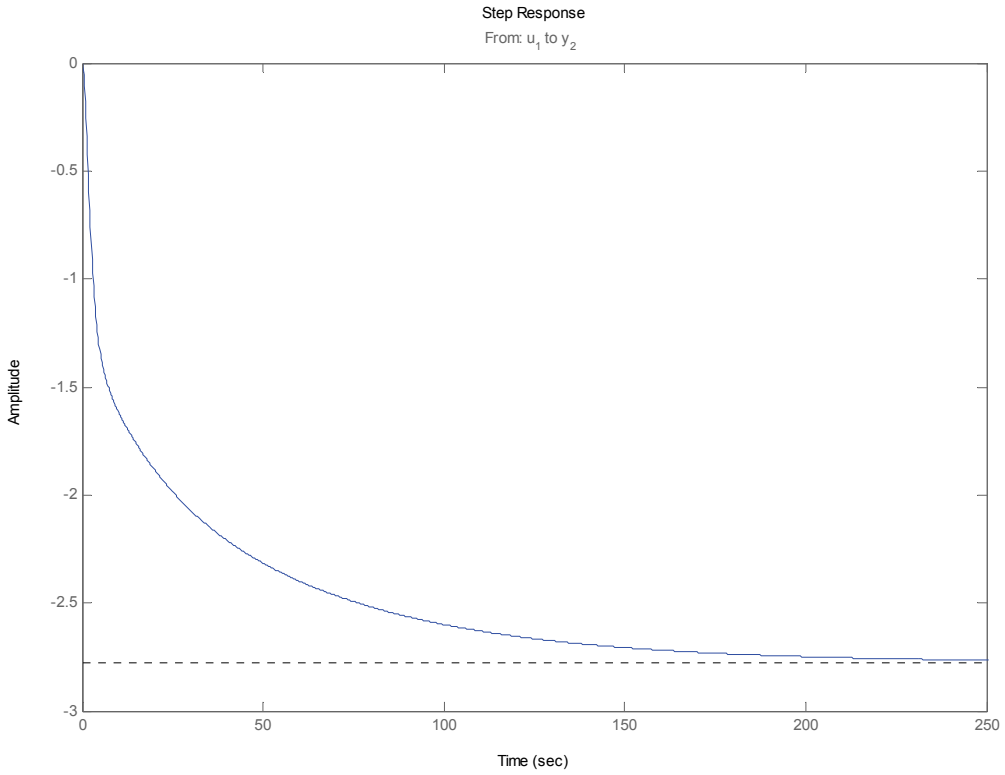

Figure C.2: Step response from u_1 to y_2.

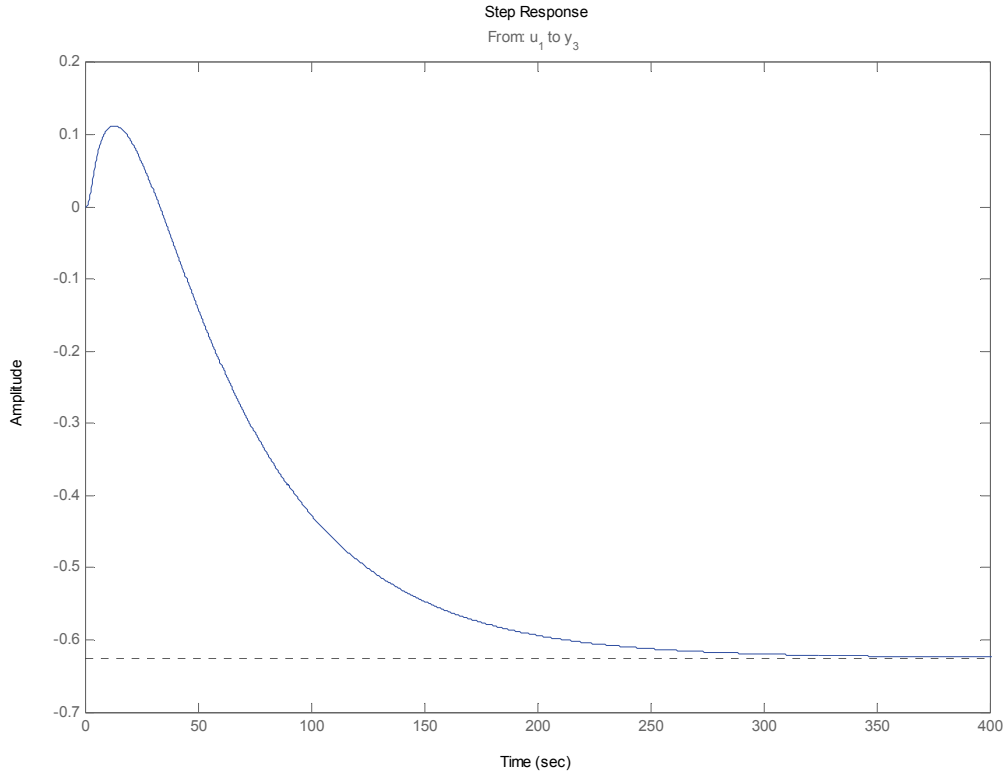

Figure C.3: Step response from u_1 to y_3.

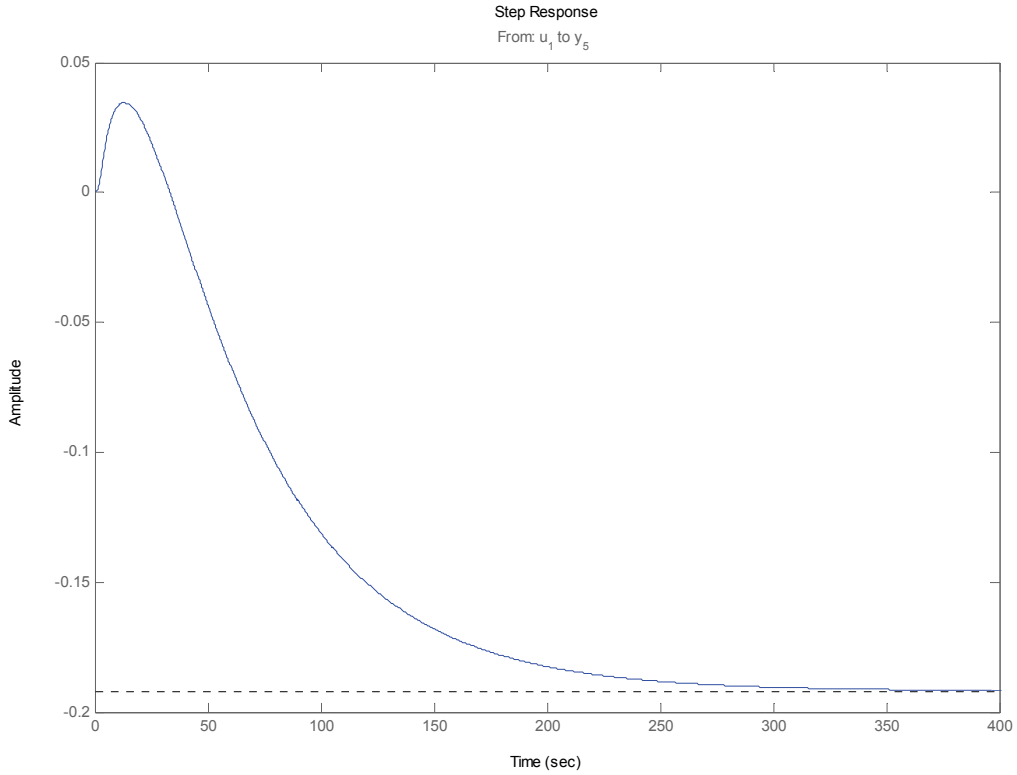

Figure C.4: Step response from u_1 to y_5.

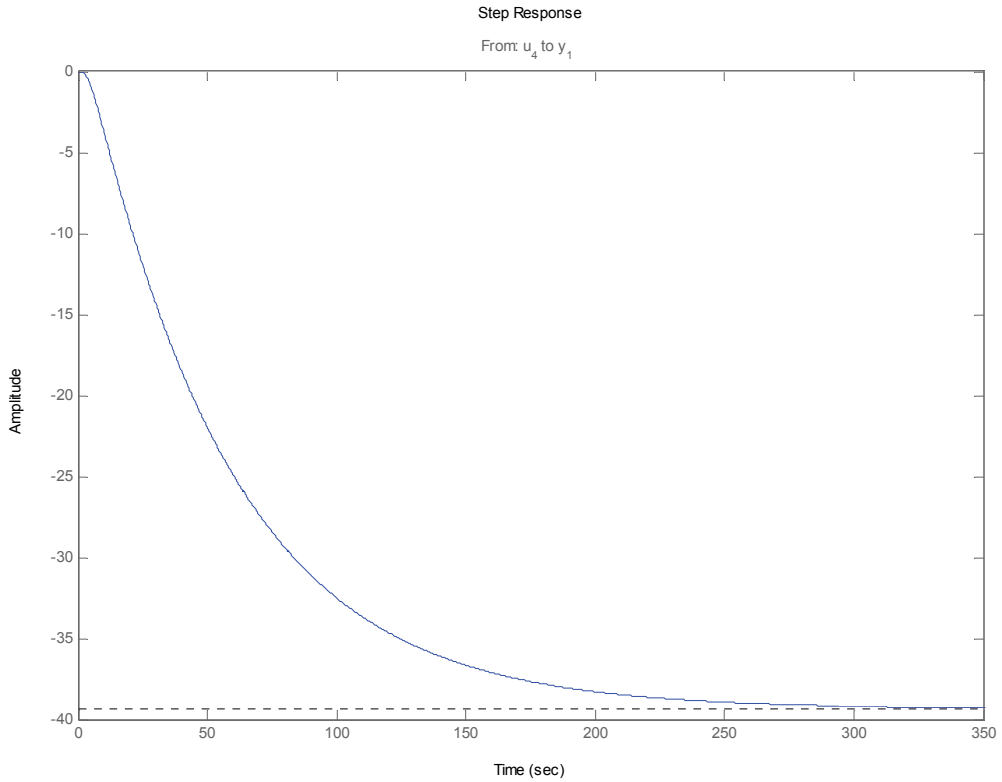

Figure C.5: Step response from u_4 to y_1.

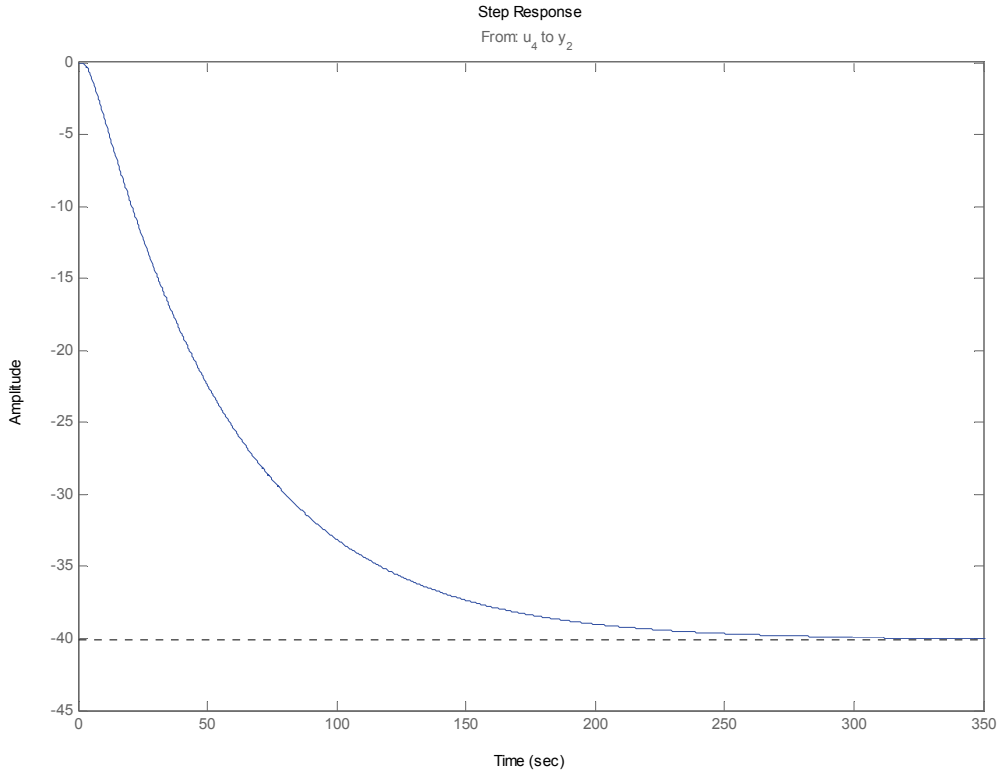

Figure C.6: Step response from u_4 to y_2.

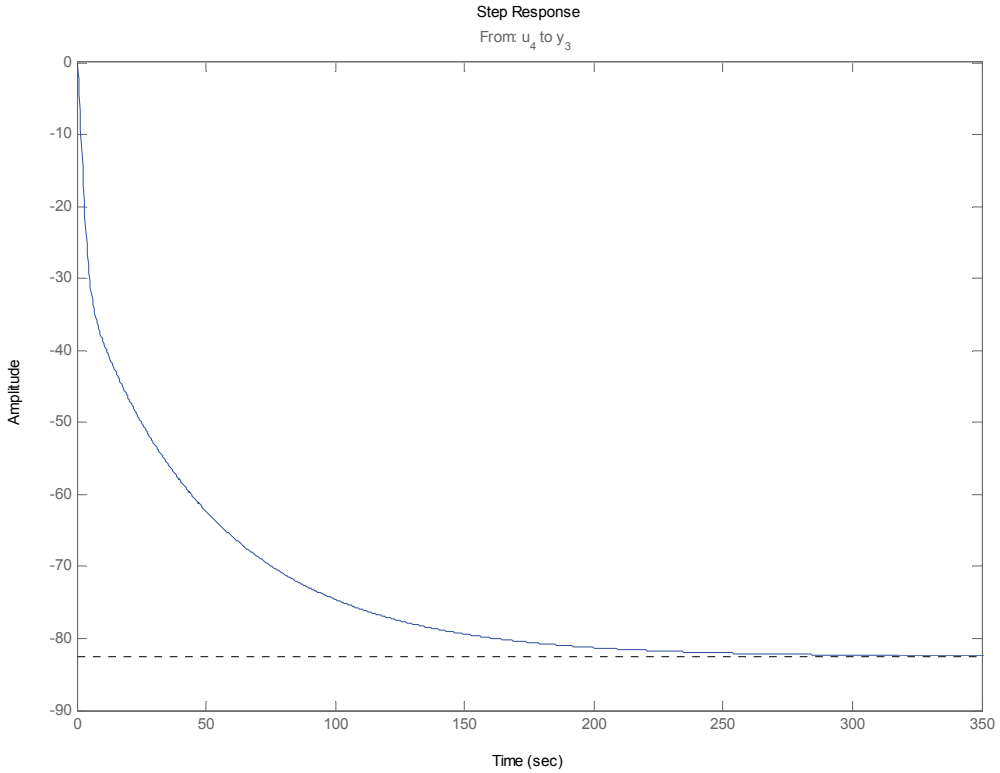

Figure C.7: Step response from u_4 to y_3.

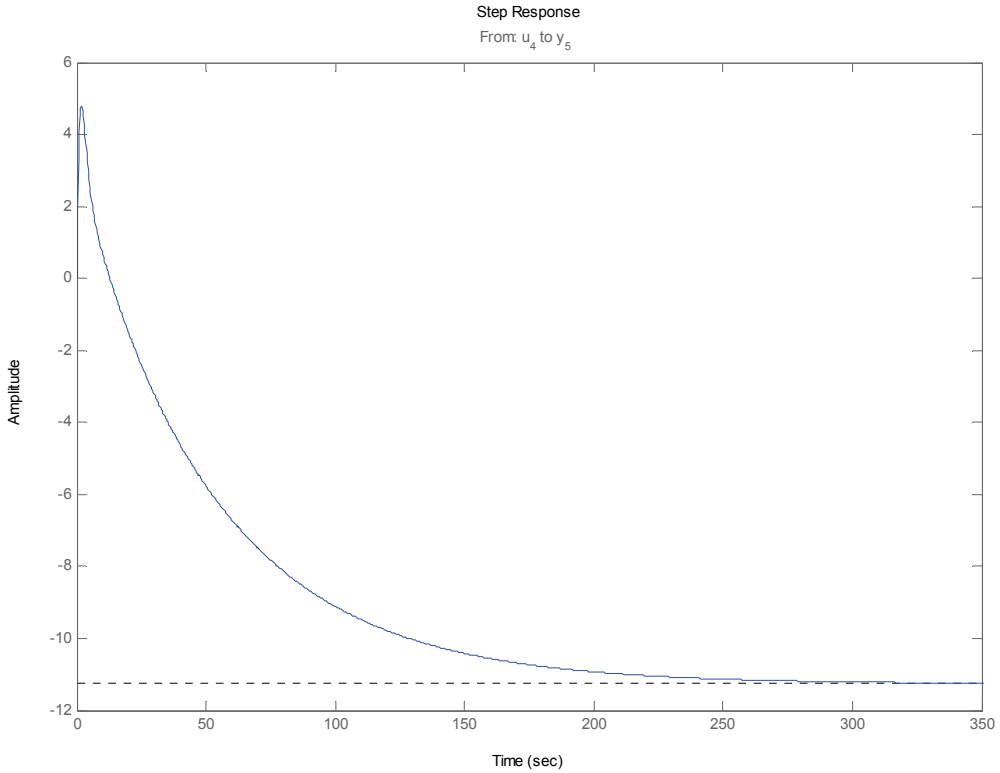

Figure C.8: Step response from u_4 to y_5.

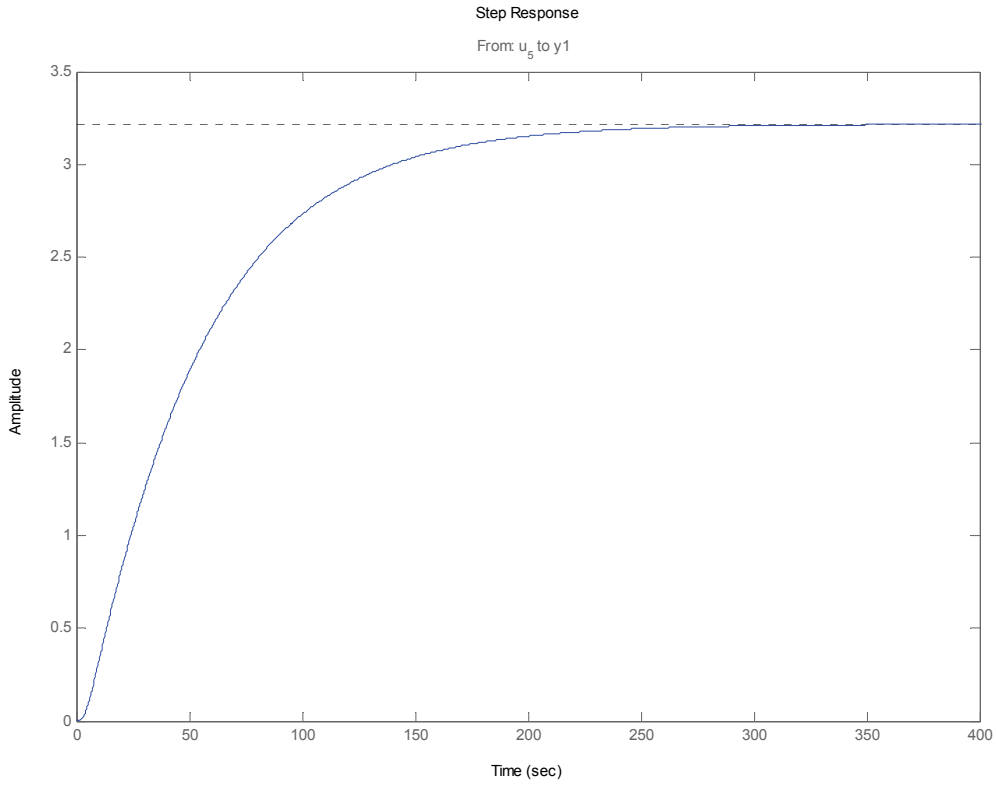

Figure C.9: Step response from u_5 to y_1.

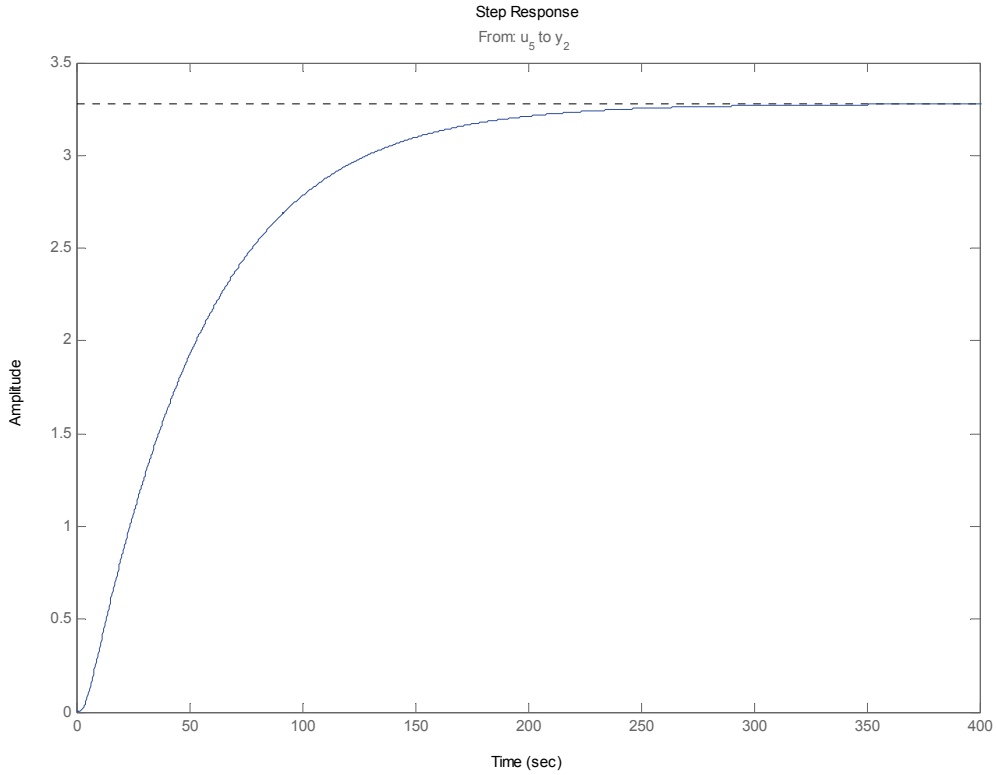

Figure C.10: Step response from u_5 to y_2.

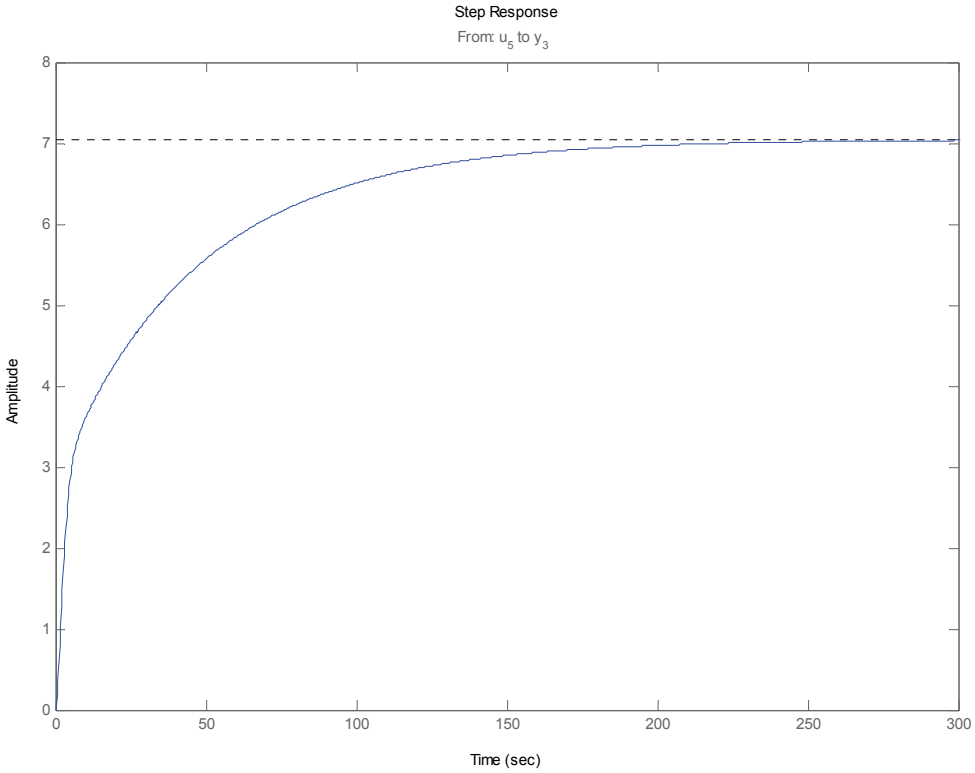

Figure C.11: Step response from u_5 to y_3.

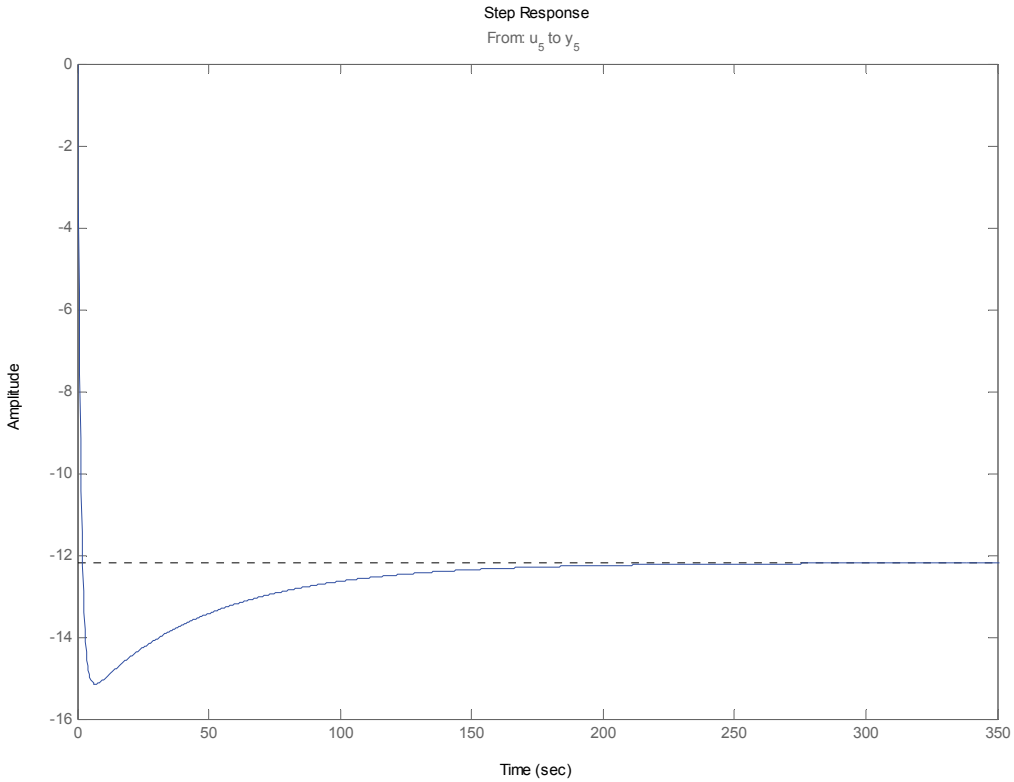

Figure C.12: Step response from u_5 to y_5.

D. VALUES OF MODEL PARAMETERS

The tables below give the values of the parameters that were used in the steady-state and dynamic models. Note that the values of h^h and h^c were obtained from curve fitted data in the steady-state model; the values presented below were only used in the dynamic model.

General parameters		
ρ^h	912.892	kg m^{-3}
ρ^c	959.049	kg m^{-3}
ρ^w	8030	kg m^{-3}
ρ^h_{air}	0.9285	kg m^{-3}
ρ^c_{air}	1.161	kg m^{-3}
c_p^h	4321.56	J kg^{-1} K^{-1}
c_p^c	4213.84	J kg^{-1} K^{-1}
c_p^w	5030	J kg^{-1} K^{-1}
c_p^{air}	1018.5	J kg^{-1} K^{-1}

Heat exchanger parameters		
A	74.6	m^2
V^h	0.2667	m^3
V^c	0.3492	m^3
V^w	0.3730	m^3
N	10	-
h^h	9796.09	W m^{-2} K^{-1}
h^c	9796.09	W m^{-2} K^{-1}

Furnace parameters		
UA	52250.6	W K^{-1}
T_i^{air}	1000	°C
w^{air}	19.89	kg s^{-1}

Flow factors		
$K_{V,u1}$	83.303	m^3 h^{-1} bar^{-1}
$K_{V,u2}$	260.32	m^3 h^{-1} bar^{-1}
$K_{V,u3}$	260.32	m^3 h^{-1} bar^{-1}

Air cooler parameters		
A	74.6	m^2
V^h	0.2667	m^3
V^c	0.3492	m^3
V^w	0.3730	m^3
N	10	-
h^h	9796.09	W m^{-2} K^{-1}
h^c	9796.09	W m^{-2} K^{-1}

Pump parameters		
p_o	15	bara
η	0.9	-

Fan parameters		
T_i^{air}	10	°C
Δp	1000	Pa
A_{duct}	3.14	m^2
z	0.6	-
η	0.9	-

E. OPERATIONAL CONSTRAINTS

Below is shown a list of operational constraints in the part of the Brobekk plant studied in this work.

Inputs	
u_1	0-100 %
u_2	0-100 %
u_3	0.55-100 %
u_4	0.25-100 %
u_5	0-100 %
u_6	0-75 kW
u_7	0-30 kW
u_8	0-75 kW

Outputs	
y_1	126 °C (steady-state model); 125-127 °C (dynamic model)
y_2	110-126 °C
y_3	90-149 °C
y_4	-
y_5	90-135 °C
y_6	250 t/h (steady-state model); 230-260 t/h (dynamic model)

Disturbances	
d_1	65-90 °C
d_2	500-900 t/h