Erasmus Mundus Programme – SpaceMaster

Czech Technical University in Prague
Faculty of Electrical Engineering
Department of Control Engineering

Luleå Tekniska Universitet
Kiruna Space Campus
Department of Space Science

Diploma Thesis

Spatially Distributed Control: Heat Conduction in a Rod

Author: Christopher Rapson Year: 2008

Abstract

Spatially distributed control is a design principle for systems which include many repeating units with overlapping regions of influence. Transfer functions in both time and spatial dimensions are used to design a controller for applications from a paper mill to the next generation space telescope. This thesis is part of a project known as HeatAl to create a hardware system for demonstration experiments of spatially distributed control. The system consists of an array of heaters and thermometers along an aluminium rod.

Mathematical models of the system have been created using finite elements, PDE, state space and multi-dimensional transfer functions. A system identification experiment has produced comparable results to the model and given a value for the convection co-efficient. Ten different controllers have been designed from published literature and grouped by architecture – centralised (MIMO), decentralised (SISO) and spatially distributed. Simulation results favour simple PI or PID designs over all architectures. The models and controllers will be used in the ongoing HeatAl project.

Acknowledgements

Many thanks must go to Zdenek Hurak for being a never-ending stream of ideas and guidance. This thesis has been a great opportunity to learn from his experience. In addition I thank him and his wife for their permission to include rendered illustrations from their work.

Declaration of Authorship

I, , confirm that this work submitted for assessment is my own and is expressed in my own words. Any uses made within it of the works of any other author, in any form (e.g. ideas, equations, figures, text, tables, programs), are properly acknowledged at their point of use. A list of the references used is included.

Signed: …….

Date: ……

Table of Contents

1 Introduction 7
1.1 Spatially Distributed Control 7
1.2 Literature Review 9
1.3 Motivation 10
1.3.1 HeatAl Project 10
1.3.2 Thesis Work 11
1.4 Hardware Configuration 12
1.5 Applications 13
1.6 Space Applications 13
2 Modelling 15
2.1 General Heat Equation 15
2.2 System Parameters 16
2.3 Initial and Boundary Conditions 17
2.4 Finite Element Model 18
2.5 Conversion to State Space 19
2.6 Transfer Functions 20
2.6.1 Spatial Transfer Function 20
2.6.2 Temporal Transfer Function 23
2.7 Comparing Results 25
2.8 System Identification Experiments 28
2.9 Plate Modelling 30
2.10 Simulation conditions 32
3 Centralised Control Design 33
3.1 LQR 33
3.2 PI 36
3.3 $\mathcal{H}_\infty$ Optimal 37
3.4 Loop Shaping 41
3.5 MPC 43
4 Decentralised Control Design 46
4.1 PID 46
4.2 Control Design from Lunze 1992 - PI 47
4.2.1 By Mixed Sensitivity 48
4.2.2 By Considering a First Order Approximation 49
5 Spatially Distributed Control Design 53
5.1 FIR and PID 53
5.2 Control Design from Stein 2005 – Mixed Sensitivity 55
5.3 Control Design from Gorinevsky 2008 – Smoothing 58
5.4 Control Design from Gorinevsky 2006 – IIR 62
6 Implementing controller functions in C 63
7 Conclusions 64
7.1 Discussion 64

7.2 Future Work 64
7.3 Conclusions 65
8 References 66

Notational Conventions and Symbols

Symbol	Interpretation	Units
x	Continuous spatial dimension	m
k	Discrete spatial dimension	
v	Discrete spatial frequency domain variable	
t	Time	s
T_s	Sample Time	s
z	Discrete time frequency domain variable	
T	Temperature	^{o}Celsius or Kelvin
$A_\perp$	Cross-sectional area	m^2
C	Circumference	m
Q	Heat Energy	J
h	Convection coefficient	$W.m^{-2}.K^{-1}$
q_s	Heater power	$W.m^{-2}$
ρ	Density	$kg.m^{-3}$
κ	Thermal conductivity	$W.m^{-1}.K^{-1}$
c_p	Heat capacity per unit mass	$J.K^{-1}.kg^{-1}$

Glossary

FFT – Fast Fourier Transform
GUI – Graphical User Interface
LQ – Linear Quadratic
LTSI – Linear Time- and Spatially- Invariant
MEMS – Micro Electro-Mechanical Systems
MIMO – Multi Input, Multi Output
MPC – Model Predictive Control
PDE – Partial Differential Equation
SISO – Single Input, Singe Output

1 Introduction

This section outlines the concepts of spatially distributed control and the current relevant research efforts before explaining the role of this project. Applications for which spatially distributed control are used, or will be used in the near future are described to give a perspective for this technology in the real world.

1.1 Spatially Distributed Control

Spatially distributed control is a type of control system architecture which will be the main focus of this thesis. It is useful for systems which include a large number of identical subsystems spread over an array. The array may be arranged in one of several configurations, as shown below. If the system can be modelled as a repeating unit then a control system can be designed and implemented on a local scale, and afterwards duplicated over the complete array. A system which has this property is designated spatially invariant, and linearised models are given the name 'Linear Time- and Spatially- Invariant' or LTSI.

Figure 1 Possible configurations for a spatially distributed system: 1-D, 2-D cartesian and 2-D hexagonal [1].

A further condition for spatially distributed control is that there should be some overlap of the influence function between units of the system, that is, the controller array should be dense with respect to the physics of the system. It is typical for spatially distributed controllers to draw information from multiple sensors within the range of the influence function to determine the correct input at one actuator - MISO control. If the array does not fulfil this condition, then we have the much simpler scenario of a large number of independent systems.

Spatially distributed control may be seen as a subset of nD control. nD refers to the concept of controlling an output which is a function on two or more independent variables. For the case of an aluminium rod, the variables are space 'x' and time 't', with a desired output temperature T(x,t). Signals can be seen to propagate in time as usual, but also in space (conduction). One spatial and one time dimension may be the simplest in terms of visualisation – which is one of the main reasons for the choice in this project – but it is by no means the only possibility. For the author, the distinctiveness of nD control was made apparent when considering the work on repetitive control [2] using two time dimensions. Where the system used in this project can be effectively modelled as a MIMO

system, this would be inappropriate for dual time dimensions, or where there are two spatial dimensions – as for image processing.

Even where MIMO control is theoretically valid, spatially distributed control presents other advantages which have encouraged its ongoing development. The wiring required to route all sensor and actuator signals through a central processor can be prohibitive for large array sizes. Once a processor has accumulated all the sensor information, the calculations to determine actuator outputs will manipulate large matrices and take even more time. Spatially distributed control, on the other hand, scales to arbitrary sizes by considering only a small section of the whole system. In fact, usually the system is assumed to be infinite in all directions.

Aside from theoretical performance issues, spatially distributed architectures have other real-world advantages. If there is a failure at the central controller, the entire system is compromised whereas a spatially distributed system would normally be able to compensate using neighbouring controllers, so reliability is improved. Where signal noise is an issue, reducing the distance between sensors and the controller (or more specifically A/D converter) will help mitigate the disturbance [3].

There are also proponents of the opposite approach to MIMO, where the system is modelled as a collection of independent systems [4]. While this is valid in some cases, there are many more where actuators have a clear influence on the output at neighbouring outputs. In the case of Rogers' research [5] a coal excavator's position during one cycle affects the position during the next, and in a heated aluminium rod the laws of thermodynamics draw heat towards cooler regions, so that a single actuator can have a far-reaching effect. Ignoring these overlapping effects is clearly sub-optimal for those systems, and so the middle ground of spatially distributed control becomes an attractive option.

Finally, it should be noted that the terminology for this field is not standardised, since it is still relatively new and developing. 'Overlapping' or 'localised' control are also common. Spatially distributed may even be understood to mean the case of a large system where actuators and sensors are spread over a wide area even where the control system is centralised.

1.2 Literature Review

A large part of this project involves taking the theoretical designs proposed in journal papers and applying them to a real system to generate comparative results.

One prolific research team in this field is Dimitry Gorinevsky, Gunter Stein and Stephen Boyd. The focus of their publications seems to be an adaptation of well known theories to the spatially distributed field. As a point of difference to other authors, their stated goal is relatively simplified applicability. Through a range of assumptions they simplify both the human and computational elements of the design process. They show a strong preference for frequency-gridded linear-programming techniques which reach a solution quickly. By reducing the design process to a limited number of steps an engineer is more likely to implement the theory on a real system.

Some assumptions are justified within the paper but appear intuitively doubtful, such as the separability of a 2-dimensional function into independent spatial and temporal components [6] or that worst case scenarios for all three mixed sensitivity constraints occur at steady state, and hence transient dynamics can be ignored during the design process [7]. Such assumptions must be carefully considered by the engineer that will apply their methods.

Unfortunately, while the design methods are not complicated the explanations can be difficult to follow. Much effort was expended in deciphering unfamiliar notation which referred to quite standard methods, such as [7] where a form of mixed sensitivity design is derived without any mention of sensitivity functions, and where multiplicative error is included in a non-standard formulation. For the same paper, the derivation of a discrete spatial frequency model is a simple task made difficult, and notation is inconsistent even within the paper. Conversion from the frequency domain where controllers are designed and the spatial domain where they are applied requires truncation, and details such as these, which may be useful to an engineer, are skipped over entirely. Another paper which describes a μ-synthesis method [8] would have been interesting, but was not pursued for this reason. This thesis will attempt to address this difficulty by giving a step-by-step explanation of the procedure followed when generating these controllers.

Please note that these criticisms are made with sincere respect for the authors and their large body of work. Their methods have been very helpful in this project, and it will be seen in section 5 that the controllers give satisfactory results.

Several authors contend that the influence of neighbouring units can be regarded as disturbance inputs to a node, which is then modelled as an independent SISO system. From these papers, one published by J. Lunze and R. Abraham [4] has been selected for implementation within this thesis, and its performance can be seen in section 4.2. This paper also deals with a thermal control system with a one dimensional array of zones. The valid point is made that while the spatial influence is obvious at steady state, heat conduction takes time to reach an equilibrium. Over a short time span, the overlapping influence is actually small. The influence from the nearest neighbours is used during design of the controller, but not during its operation, and satisfactory performance is achieved.

Here also the explanations could be clearer – the method by which the stated design objectives (page 109) are converted to a set of all suitable controllers (120) is not mentioned at all. K_s is given as the static system model ($y=K_s.u$) which is taken from earlier equations relating to a state-space model, but is then derived as a constant matrix. A constant matrix implies that the system has an instantaneous response, but is more likely meant to be interpreted as the steady state gain of the system.

The issue of boundary conditions appears to be a common one among all spatially distributed control literature, where the problem is usually ignored by assumption. At best, the authors acknowledge that in the real system some heuristic tuning has been done to modify the control algorithm at the boundary nodes [4,8].

1.3 Motivation

1.3.1 HeatAl Project

N-Dimensional and spatially distributed control are topics of interest at the Control Department of the Czech Technical University. This thesis is part of one of their projects known as 'HeatAl'.

In the literature there is a notable scarcity of systems where spatially distributed control theory can be tested. Even simulation results are rare. Therefore the goal of the HeatAl project is to create a hardware system that can be used as a demonstration platform for experiments using spatially distributed control.

It should be noted that the dynamics of the system are relatively simple. Although this is determined in part by budget constraints, the lack of complexity should be seen as an advantage. nD control design is a step apart from standard approaches, and presents its own challenges that the engineer must become acquainted with. Even with only one spatial dimension the theories require a conceptual shift that should not be underestimated. Temperature control is one

of the standard control problems, and it is useful that the difficulties in this respect at least are well known. Furthermore, larger systems would prohibit the use of centralised control and the analysis would lose comparability.

1.3.2 Thesis Work

Within the context of the HeatAl project, it is the focus of this thesis to derive mathematical models of the system and design an assortment of controllers. The mathematical models provide an understanding of the system and are necessary prerequisites for controller synthesis, including controllers which will be designed in the future as the HeatAl project continues. Controllers are designed using centralised, decentralised and spatially distributed methodology. In all cases the designs are derived from published literature to provide legitimacy. It is beyond the scope of this work to suggest new control methods. The initial priority is to create a working system using proven designs to lay the groundwork for innovation.

By applying a range of different designs to the same system, comparisons of their performance can be drawn. The criteria by which they will be assessed include

- step response
- reference tracking
- disturbance rejection
- computational effort

Particularly when juxtaposed with MIMO controllers, for which the strengths and weaknesses of the various design methods are well established, this thesis seeks to provide a quantitative comparison for the new, lesser known methods. Where engineers have a choice between several potential control architectures, the kind of comparative study in this thesis could provide a useful starting point, although care should obviously be taken when extrapolating the results to other systems.

Beyond spatially distributed control and the HeatAl project there is a dual motivation for this thesis; in the development of 'standard' controllers the author has had the opportunity to implement knowledge and skills learnt during the SpaceMaster degree, in particular the Optimal and Robust Control course. The not-so-subtle difference between reading about a controller in a textbook and actually designing and implementing a real example has been a very useful experience. With a view to applying these controllers to a physical system, real world issues such as integrator windup, finite word length and actuator saturation will be considered during the design. This is another experience that is largely ignored in the classroom.

1.4 Hardware Configuration

The design and construction of the system has been the work of a concurrent project by Vaclav Klems, and a full description can be found in his report. What follows is a brief description of the most important components, since it will be useful during later sections if the reader has some conception of the plant that will be controlled.

The basic components of the system are a long, thin aluminium rod with a 1-dimensional array of heaters and sensors. The rod is a 1m long aluminium extrusion, with rectangular cross section 20x10mm. Heaters and sensors are placed in a pair as close as possible to each other, and each pair is referred to as a node of the system. In order to thermally isolate the rod, it will be suspended in a cradle by threads at either end. The only interaction with the environment is via convection with air at room temperature.

It is important for the principle of spatially distributed control that the nodes have overlapping influence functions. In a frequency context, the sampling rate should be at least double the natural system bandwidth, and preferably higher. Therefore the separation distance was balanced with the heaters' power output during design.

Actuators are controlled via a switching circuit. The strength of their heating is dictated by pulse width modulation from the microcontroller. Most of this report refers to power density, so 20W or 33,333Wm^{-2} corresponds to a 100% pulse width. Temperature sensors will also give their reading in terms of pulse width modulation. The microcontrollers communicate via UART with their nearest neighbours and with the central PC. MAX232 chips convert UART to RS232 standard. For the purposes of this project, one microcontroller can control multiple nodes; the separation will be maintained in software.

Another difficulty is obtaining a reliable power supply with constant voltage. The system can consume over 600W so a high power source is required. The final design uses several smaller, cheaper power supplies intended for computers.

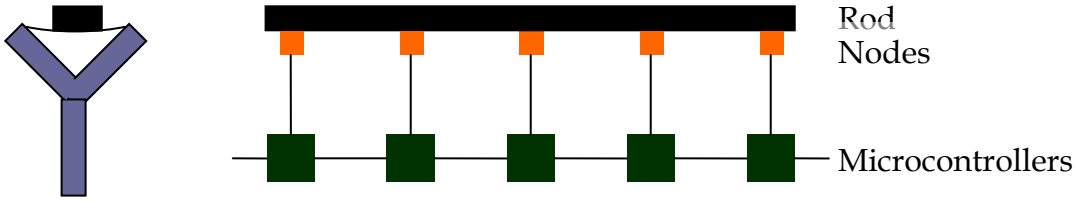

Figure 2 Diagram representing the basic system components.

1.5 Applications

Numerous applications of spatially distributed control exist already and this is only expected to increase as the potential is realised. Historically, the most common is in paper mills to control the thickness of paper across a large sheet [9]. As the paper passes over a drum, an array of sensors measures the thickness and adjusts the pressure to flatten it more or less, as required. More recently, heat treatment of silicon wafers has been controlled by spatially distributed controllers [10]. Doping in wafers is highly dependent on temperature, and must be accurately controlled in both time and space. Feedback from temperature sensors underneath a wafer are used to control an array of high-power infra-red lamps and achieve ramp rates up to $150^{o}Cs^{-1}$ with an even heating profile.

With the promising development of MEMS technology, even more possibilities are available. Micro-electro-mechanical systems can be manufactured on a single wafer so that each tiny processing unit has in-built connections to its nearest neighbours. This is obviously ideal for spatially distributed control since each controller has information from its nearest neighbours and operates entirely independently from the rest of the structure. Actuators, sensors and computational elements are all part of an integrated circuit, so wiring difficulties are eliminated. This technology will bring spatially distributed control into the fields of adaptive optics (see next section), dielectrophoresis and fluid flow control.

More obscure applications include high precision agriculture [3] where irrigation systems are controlled along a pivoting 400m sprinkler beam and meso-scale assembly using a 2-D array of electromagnetic actuators to induce random motion [11].

1.6 Space Applications

Here the term 'space' refers to the wider universe outside the earth's atmosphere, which is important to this project as a component of the SpaceMasters course. Typical for control systems, techniques developed on a basic aluminium rod can be adjusted and applied to cutting edge technology like the next generation of telescopes.

The quality of an image is highly dependent on the size of the mirror used in the telescope. A larger diameter will capture and focus more photons onto the CCD array. The size comes at a cost, however, in that it becomes progressively more difficult to maintain an accurate surface profile. At times a constant curvature may not even be desired - using techniques from adaptive optics, an actively controlled mirror can be adjusted to account for wavefront distortion.

In the case of space-based telescopes, volume and weight are premium constraints due to launch costs. Therefore a new generation of telescopes have been proposed with flexible mirrors that will deploy in space and require active control. Since the array will contain hundreds if not thousands of nodes, the problem becomes computationally infeasible for a central processor. This application appears to have been the motivating factor behind several research efforts [6,7]

Similarly for ground based radio-telescope antennas, with diameters over 30m it becomes difficult to keep a constant curvature over the entire surface. Passive control measures will not give satisfactory results in the presence of disturbances from pointing movements, wind loads and atmospheric disturbances. Numerous examples currently in the design phase include Cornell's Large Atacama Telescope to be installed in Chile and the competing CELT in California. All extremely large telescope designs decompose the mirror into a large number of segments, and at least these two are considering spatially distributed control as a potential control architecture. It is noted in [12] that any actuator of a mirror segment produces an equal and opposite force on the rest of the structure. Since a completely rigid frame would be prohibitively expensive, the actuators can be seen to have overlapping influence even where the mirror segments are separated. These interactions must be accounted for.

For arrays of many small antennas (shorter wavelength radiation) a spatially distributed control system is also a candidate architecture.

Several new missions use a number of small spacecraft, which are co-ordinated to fly in a formation. CLUSTER is perhaps the most successful example to date. There are several methods of controlling the group, such as leader-follower or where all satellites track an average trajectory plus their separation. At this time, it is not difficult for all satellites to communicate with all others in the formation, but as the numbers of satellites increase, this too becomes a potential application for spatially distributed control [9].

One application which is not mentioned in the literature is the thermal control system of a satellite. This is one of the most critical components of a mission due to temperature sensitivity of fuel and electronics, particularly since the temperature must be stabilised against both extremes of solar glare and deep space. The systems typically consist of a large number of nodes, with several sensors and a heater at each. Potentially, they could be combined to a network and controlled using a spatially distributed architecture. This would improve redundancy and improve performance by allowing the overlapping effects to be accounted for.

2 Modelling

In order to synthesise controllers for the system, a mathematical model of the internal physics is a prerequisite. The same model will be used in simulations to test the performance of the controllers. As presented here, the model has an entirely theoretical basis. A comparison was made with the real system once it reached a sufficient level of completion; see the system identification section 2.8. Code for the models can be found on http://dce.felk.cvut.cz/heat.

Finite element, state-space and transfer function based models are developed from the basic thermodynamics equations governing the system. The Roesser model is often used for n-D modelling and control, but is not applicable in this case since it assumes causality in all dimensions, and thermal systems propagate heat in both directions, that is, non-causally.

2.1 General Heat Equation

The starting point for the model is very general; the basic thermodynamics equations found in any textbook with some initial assumptions. The equation neglects mass flow and radiation effects. More contentiously, it only allows heat transfer in one spatial dimension. That is, the cross section of the rod is assumed to have a uniform temperature distribution such that there is no heat transfer in the y or z directions; the rod is 'thermally thin'. This is justified if the cross section is small enough that conduction is effectively instantaneous.

The equation below shows the rate of change of heat energy in a volume element ($A_\perp.\partial x$) of the rod is determined by a conduction term and a so-called generation function $g(x,t)$ which represents heat transfer across the boundary of the rod.

$$\dot{Q}(x,t) = \rho.c_p.\frac{\partial T(x,t)}{\partial t}.A_\perp.\partial x = \left(\kappa\frac{\partial^2 T(x,t)}{\partial t^2} + g(x,t)\right).A_\perp.\partial x \qquad (1.)$$

The generation function includes the heater power q_s and a convection term.

$$g(x,t).A_\perp.\partial x = q_s(x,t).C.\partial x + h.\left(T_{environment} - T(x,t)\right).C.\partial x \qquad (2.)$$

Note that this equation still refers to a slice $A_\perp.\partial x$, so that both the heater and convection must be modelled as operating around the entire circumference even though the heater will only occupy a part of this surface. Thus q_s (in $W.m^{-2}$) should be given as the heater power divided by the surface area of the slice.

Combining these equations gives the complete PDE

$$\frac{\rho.c_p}{\kappa}\frac{\partial T(x,t)}{\partial t} = \frac{\partial^2 T(x,t)}{\partial x^2} + \frac{C}{\kappa A_\perp}q_s(x,t) + \frac{h.C}{\kappa A_\perp}\left(T_{envuronment} - T(x,t)\right) \qquad (3.)$$

2.2 *System Parameters*

While the equations are general to many thermodynamic systems, the parameters are particular to this hardware configuration. The surrounding environment is air at room temperature, the rod is EN AW1350-F Aluminium and transistors are used as heaters.

The values shown below are predicted from theory and are used for both simulation and controller synthesis. Some variation from the true values is expected. Therefore it is convenient at this stage to consider parametric uncertainty, which is later used to design robust controllers.

Table 1 System Parameters [13]

Symbol	Parameter	Minimum	Nominal	Maximum	Units
ρ	Density	2600	2700	2800	$kg.m^{-3}$
κ	Thermal conductivity	230	230	234	$W.m^{-1}.K^{-1}$
c_p	Heat capacity per unit mass	900	900	900	$J.K^{-1}.kg^{-1}$
q_s	Max Heater power	16	20	20	W/m^{-2}
l	Rod length	2	2	2	m
width	Rod width	0.019	0.02	0.021	m
height	Rod height	0.009	0.01	0.011	m
h	Convection co-efficient	2	10	25	$W.m^{-2}.K^{-1}$
$T_{environment}$	Ambient Temperature	15	25	30	^oC

The convection coefficient is particularly uncertain, since it depends on factors which are not constant such as temperature difference and others which are not well known, such as air density, viscosity and conductivity. The values given above are from a textbook, and give good agreement when the value (8.6795 $W.m^{-2}.K^{-1}$) is calculated using (4) [14].

$$h = 0.54.(\mathrm{Pr}.Gr)^{\frac{1}{4}} \qquad (4.)$$

where Pr and Gr represent the Prandtl and Grashof numbers respectively. The relationship holds where their product is between 10^3 and 10^9.

Including all uncertainties would lead to a very complicated and improbable model, so for robust design only thermal conductivity, heater power and the convection coefficient were varied. $T_{environment}$ is an input rather than part of the model, so this variation acts as a disturbance rather than true uncertainty. q_s represents the saturation limits of the control input, and can not be incorporated to any linear model. It was included in simulink models for testing controller performance.

2.3 *Initial and Boundary Conditions*

To solve a PDE requires some knowledge of conditions at the beginning of the experiment and at the ends of the rod. Initial conditions for all simulations were set at room temperature (25°C) along the length of the rod.

Boundary Conditions are more difficult to estimate. If Neumann conditions are used, the assumption is that the ends of the rod are perfectly insulated, or at least have a negligible convection effect compared with the length of the rod. For simulation of the central section of a long rod, this simplification is acceptable and would not be unusual in the nD literature.

Alternatively, Robin conditions can be used to model convection on the end surfaces. The author has chosen to use this type for the state space model, by enlarging the exposed surface area of the first and last slice to include the end face. Inaccuracies remain, however, in keeping the same Convection Coefficient, since natural convection occurs faster on vertical surfaces [14].

The characteristic width of the system spatial impulse response gives an indication of the boundary condition influence [7]. It is also useful to know the degree of overlap between actuators. The graphs below show the result of the central heater switched opposite to the other heaters on a long rod after 2000 seconds.

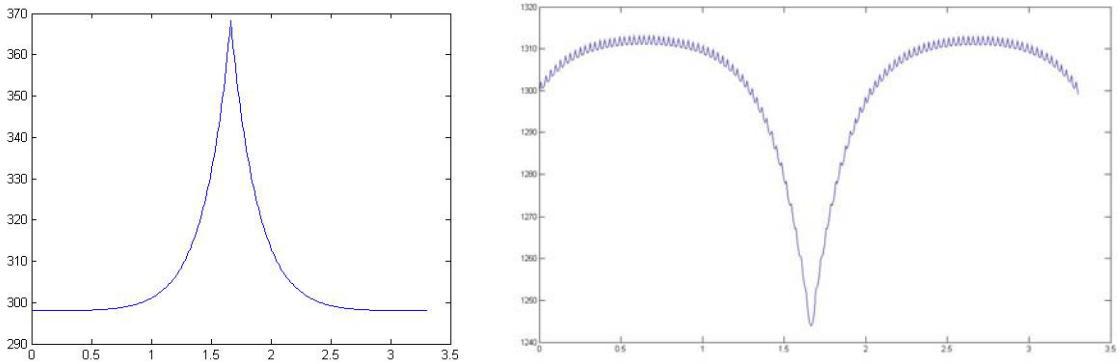

Figure 3 Influence function of one heater on a long beam.

The influence reduces by one order of magnitude at a point approximately 0.5m from the switched heater. (0.48 and 0.52m for the two scenarios above.) This suggests that the boundary conditions will have a large range of influence on this particular system and that there will be a large degree of overlap between actuation nodes. Note that the downwards curve at the outside of the right hand simulation represents the effect of the Robin conditions used in the simulation and cannot be seen as the effect of true boundary conditions.

2.4 *Finite Element Model*

The software package FEMLAB can used to create finite element simulations for a range of different systems, including thermodynamics. Here a three dimensional model of the rod is extruded to test the 'thermally thin' assumption.

The FEMLAB software generates a mesh on a model of the aluminium rod. Boundary conditions are set for convection on all surfaces, with the nominal convection coefficient. Several patches are embedded in one surface and here the boundary conditions are set as heat sources to simulate the heater array. A simulation with all heaters running for 10 minutes gives the following result.

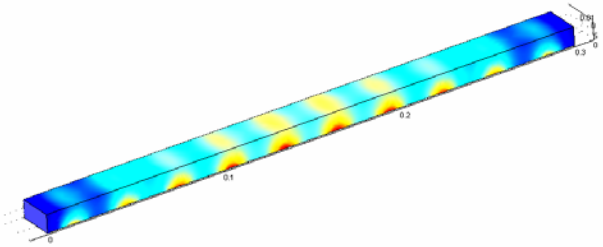

Figure 4 Surface temperature result from FEMLAB simulation. The temperature range was only 7K, from 892-899K so the colour variation from bottom to top is slightly misleading.

The temperature evolution in a slice of the rod is a good test of the thermally thin assumption. As can be seen below, there is virtually no delay in the heating effect from one side to the other. The cross section was taken at the edge of a heater, where the temperature gradient is steepest.

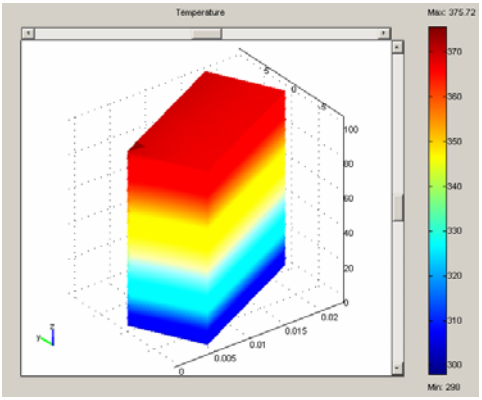

Figure 5 Temperature evolution in a cross-section of the rod. Time is the vertical axis, the heater is attached to the right face, and heat propagates through the section to the left.

2.5 *Conversion to State Space*

For control system design it is generally necessary to have some form of LTI model of the system. FEMLAB allows only constant input and the inputs for the PDE solver must be predefined before starting a simulation, whereas LTI models allow feedback to determine the inputs.

It has been decided to model the states as dense spatial samples in the x direction. The high density relative to the sensor array is necessary to capture the lower temperatures between the heaters and high spatial frequencies.

When equation (3) is separated into terms which depend on the states T(x,t) and the inputs $T_{environment}$ and q_s, and made explicit in the time derivative of the states, the A and B matrices of the state space system become apparent.

$$\frac{\partial T(x,t)}{\partial t} = \frac{1}{\rho.c_p}\left[\kappa\frac{\partial^2 T(x,t)}{\partial x^2} - \frac{h.C}{A_\perp}T(x,t) + \frac{h.C}{A_\perp}T_{envuronment} + \frac{C}{A_\perp}q_s(x,t)\right] \qquad (5.)$$

The second spatial derivative can be approximated by a discrete central difference, since the spatial dimension is non-causal. For the first and last states, the central difference becomes a forward or backward difference respectively.

The convection term is also modified for the first and last states. Here, the external area C is extended to include the surface area of the end face $A_\perp$. The correction by a factor of the spatial sample size Δ [metres] arises from a cancellation which occurs when dividing by ∂x between equations (1,2) and (3).

$$A = \frac{\kappa}{\rho.c_P\Delta^2}\begin{bmatrix} -1 & 1 & & & 0 \\ 1 & -2 & 1 & & \\ & & \ddots & & \\ & & 1 & -2 & 1 \\ 0 & & & 1 & -1 \end{bmatrix} - \frac{h}{\rho.c_P A_\perp}\begin{bmatrix} C+\frac{A_\perp}{\Delta} & & & 0 \\ & C & & \\ & & \ddots & \\ 0 & & & C+\frac{A_\perp}{\Delta} \end{bmatrix} \qquad (6.)$$

$$B = \begin{bmatrix} \frac{\alpha}{\rho.c_P A_\perp}\begin{bmatrix} C+\frac{A_\perp}{\Delta} & & & 0 \\ & C & & \\ & & \ddots & \\ 0 & & & C+\frac{A_\perp}{\Delta} \end{bmatrix}, & \frac{C}{\rho.c_P A_\perp} * Bheaters \end{bmatrix}$$

Bheaters refers to a matrix which relates each heater to the corresponding states. If the slice for a state corresponds to a heater then the element is 1, if there is an overlap then the element is some fraction, and if they do not coincide the element is 0.

Finally, the C matrix samples those states which are closest to the centre of each heater, representing temperature measurements from thermometers placed beside the heaters.

So the system is composed of large matrices whose structure is almost Toeplitz, except for the extremes. When these matrices are used to simulate the system or synthesise controllers there are many superfluous computations. This is one of the motivating factors for the spatially distributed control community, since where a smaller system which is representative of the much larger whole can be used, it will save much of the computation and can be scaled thereafter to arbitrary sizes.

2.6 Transfer Functions

Another possibility for an LTI system is a transfer function. For an nD system an independent function is derived for each dimension. The overall transfer function can be obtained as a convolution of this set, although this is generally inferior to the state space model and so was not considered here.

2.6.1 Spatial Transfer Function

The derivation of the spatial transfer function was motivated in part by its use in several publications by researchers of spatially distributed control [4,6,7,8]. To obtain the transfer function, a simulation was conducted with one heater switched on and the resulting spatial impulse response was measured after some time. The existing literature uses the steady state response, so the simulation was run for over 30 minutes.

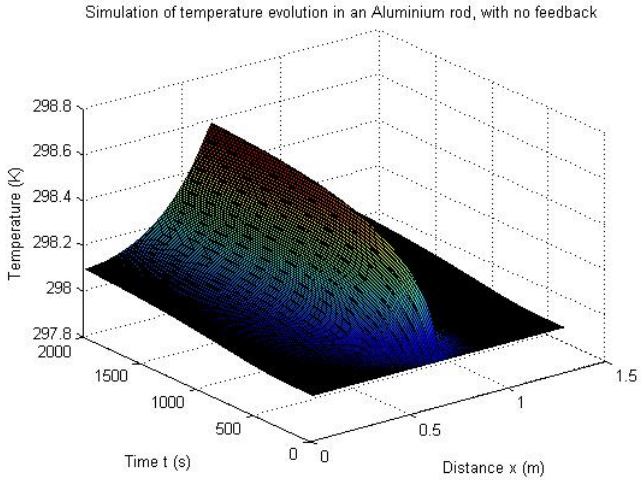

Figure 6 Spatial Impulse Response

The symmetrical, non-causal shape is clearly evident. It was soon realised, however, that the magnitude of the response from room temperature would be larger than at the operating temperature (estimated at 348K) since convective cooling is more powerful with a greater temperature difference. Therefore all heaters were switched on to maintain the temperature at 348K. The exact setting was determined from a feedback controlled simulation of the system. The central heater power was then raised by 1% of capacity. The resulting input function is shown below, and becomes too complex for an analytic derivation of the response.

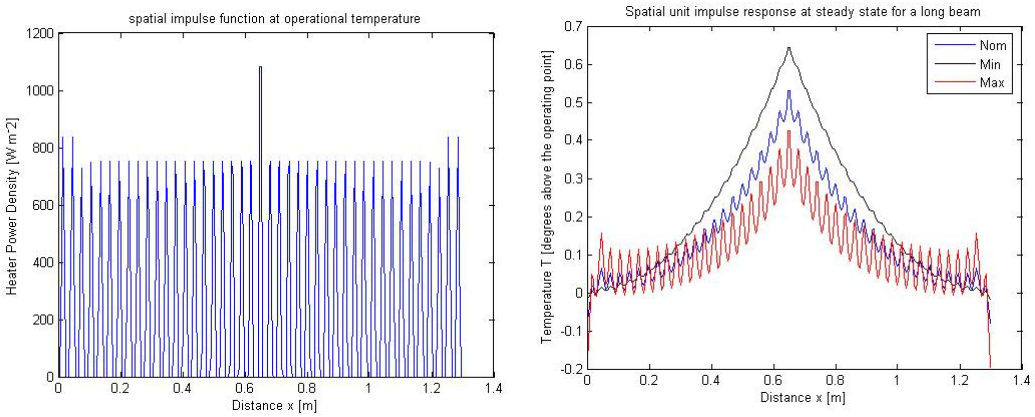

Figure 7 Spatial Impulse Input at x=0.65, with other inputs used to maintain the rod at the operating point. This input gave rise to the Spatial Impulse Response on the right.

Simulation with this input gives the steady state response shown, where the three lines represent the uncertain model with parameters selected for maximum, minimum and nominal cooling. The magnitude of the response is indeed smaller than before (0.53 compared with 0.58 in the nominal case).

From here it remains to convert the impulse response into a transfer function. The first step is to take a fourier transform. The graph below clearly shows a large DC component. The higher frequency peaks at 42.9 cycles per length and higher harmonics match the spacing of the actuators. Obviously, this is also the spatial sampling frequency of the temperature sensors. From the log-log plot of the same graph, the unity gain crossing (spatial bandwidth) can be found at approximately 7length^{-1} for the nominal case. Even for the worst case it is 11length^{-1}, a healthy margin below the nyquist frequency of 21.5.

To relate the per length frequency to the m^{-1} frequency, divide by the length of the rod, which in this experiment was 1.3m. (The 7length^{-1} bandwidth becomes 5.4m^{-1}.) Note that the highest unique frequency which could be measured in this simulation was 130 cycles per length, half the number of samples taken. Aliasing is clearly visible in the graph below.

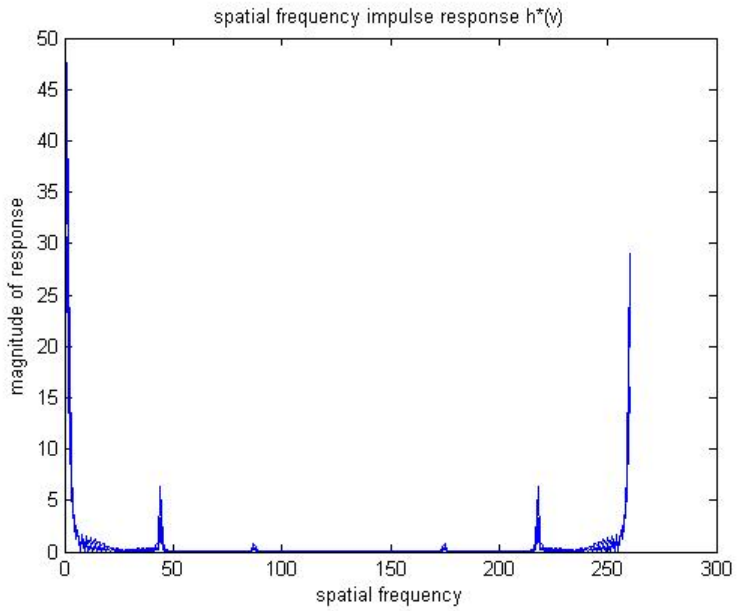

Figure 8 Spatial Impulse response fourier transform.

The spatial frequency components correspond directly to the coefficients of the transfer function in the discrete domain. In time the standard variable would be 'Z', for spatial co-ordinates 'V' has been chosen. These coefficients are of course complex, but when the function is evaluated for V on the unit circle (equivalent to the imaginary axis of the S domain, or a purely sinusoidal input) the result is real due to the symmetry of the function. The complex conjugate coefficients of V^{-1} and V^{1} will cancel where the magnitude of V is unity. The frequency is given by the argument of V as it progresses around the circle.

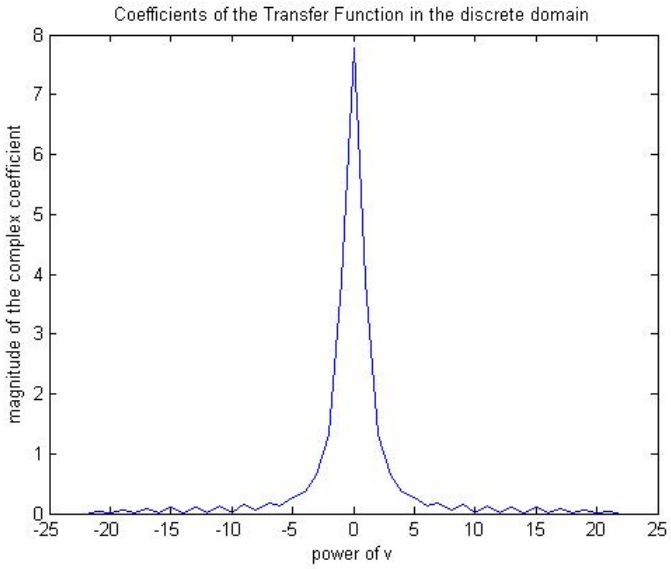

Figure 9 Magnitude of the coefficients of the Transfer Function in the discrete spatial domain

Truncated to 5 terms, the transfer function is

$$\ldots+(1.3+0.19i)v^{-2}+(-3.9-0.28i)v^{-1}+7.8+(-3.9+0.28i)v+(1.3-0.19i)v^{2}+\ldots \quad (7.)$$

with the same 0.03m sample spacing. When the experiment is repeated with a negative unit impulse, the following transfer function is derived

$$\ldots+(-1.3-0.19i)v^{-2}+(3.9+0.28i)v^{-1}-4.1+(3.9-0.28i)v+(-1.3+0.19i)v^{2}\ldots \quad (8.)$$

The magnitude is decreased in the central point, since the cooling process is slower than heating. The overall DC magnitude is reduced but all other frequency components remain the same.

2.6.2 Temporal Transfer Function

The temporal transfer function bears many similarities to the spatial transfer function and once more, the motivation stems from a journal article [4]. Using similar reasoning to before, the temporal transfer function can be found from the time impulse response. Again, the baseline input is raised to a point where the temperature would be maintained at 348K and for one sampling period all inputs are raised by a further 1% of their capacity.

The resulting output is shown below for the full length of the rod and an arbitrary point along the length is selected. The second graph shows the temporal evolution in profile. The magnitude by which the temperature is raised above the baseline at the sampling times corresponds directly to the coefficients of the transfer function in the discrete time domain.

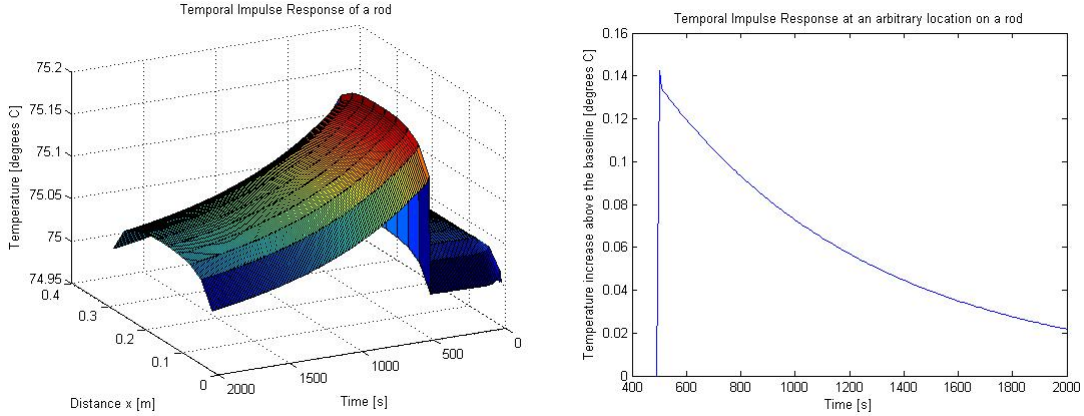

Figure 10 Temporal Impulse Response.

Thus the nominal system has a temporal transfer function of:

$$0.1425z^{-1}+0.1342z^{-2}+0.1323z^{-3}+0.1306z^{-4}+0.1289z^{-5}\ldots \quad (9.)$$

The process can be repeated for a perturbed system. Lunze [4] suggests that the worst case perturbed system is where one of the heaters is not functioning, and measuring the response at this location. Note that the system is no longer

spatially invariant, so it is not possible to create a uniform baseline temperature. The baseline input has been designed to minimise the error by increasing the heat applied on either side of the defective actuator, again by using a feedback simulation.

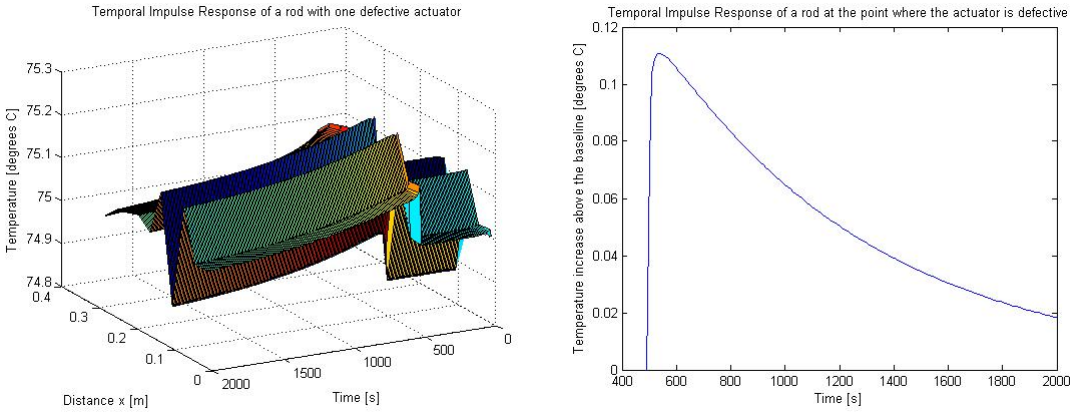

Figure 11 Temporal Impulse Response of the perturbed rod system.

The graphs show a smaller and smoother initial increase, followed by a similar exponential decay. By the end of the simulation the difference is much smaller. This implies that steady state error for the perturbed plant is small, which can also be seen on the following bode plot. The frequency responses of the nominal and perturbed plants diverge as frequency increases.

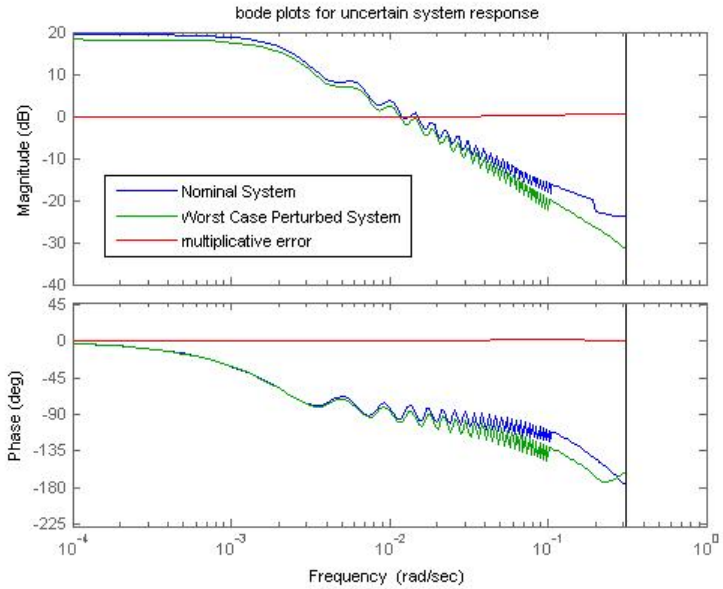

Figure 12 Time frequency response of the nominal and perturbed plant.

A similar result can be derived analytically by calculating the maximum cooling rate. At 75°C, this rate is $0.04^{o}Cs^{-1}$. By considering the amplitude and frequency of sinusoids which have this maximum ramp rate, a value for bandwidth of 7.038mHz is found. It can be seen that both these assessments correlate well with

the state space system when viewing the sigma plot below. Here the magnitude of the frequency response has a similar breakpoint at 10^{-3}Hz and a similar roll-off rate.

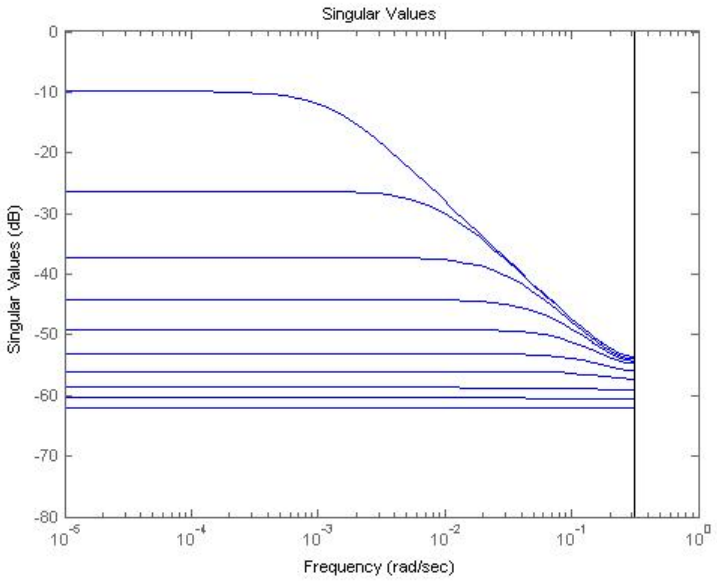

Figure 13 Plot of sigma values for the nominal system on logarithmic axes.

We can also test the system response to input sinusoids of varying frequency. He we see that a signal that should give an output sinusoid with an amplitude of 10K is too fast at 10^{-2}Hz, but achievable at 10^{-3}Hz.

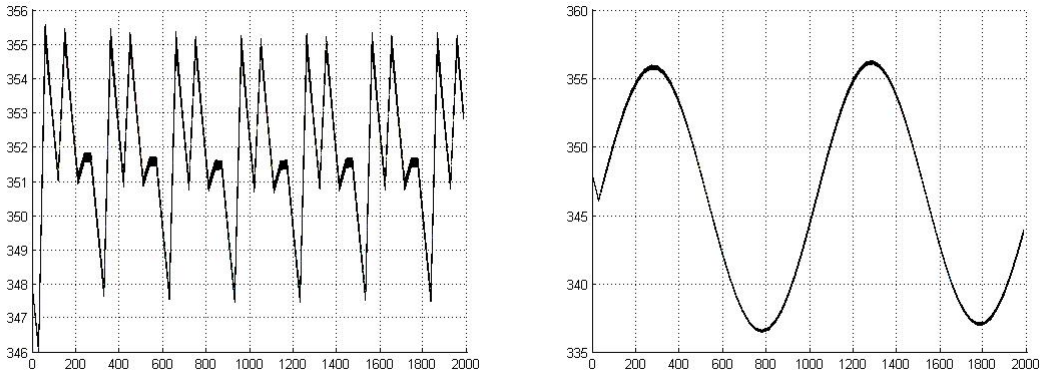

Figure 14 Response of the system at 0.01 (left) and 0.001Hz (right).

2.7 Comparing Results

Before the hardware was assembled and available for testing, it was useful to compare the different numerical simulations to confirm that they gave similar results for a range of scenarios. For these and other simulations a shortened rod of length 300mm was used. 10 nodes are spaced at 30mm intervals.

The first scenario chosen was to switch all heaters on for ten minutes. It can be seen that the general form of the solution is the same, and cooling between the nodes is in the region of one or two degrees for all models. The average temperature shows more variation – a difference of around 20 degrees. Since the result is 600 degrees above ambient temperature, this is actually a small difference – although it provides some motivation for using robust control design techniques.

The boundary conditions of the second kind used for the PDE model give a different shape to the result, which intuitively seems less likely than the domed shape of the others.

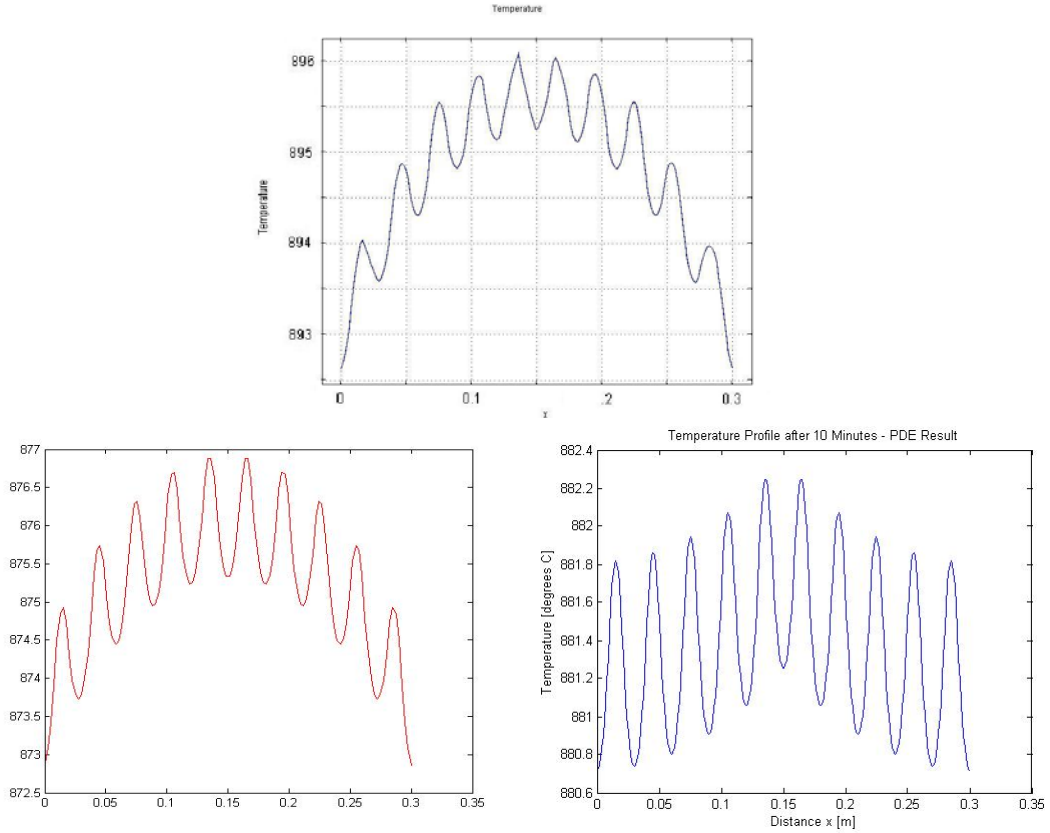

Figure 15 Results from the first simulation with constant heating from all inputs. Clockwise from the top left, the graphs are taken from FEMLAB, PDE and state space models.

The second scenario showed the evolution from initial conditions, where a sinusoidal temperature distribution was assumed at time 0. Such conditions can only be set for PDE and state space simulations, but for these two methods the results were almost identical. Only at the boundaries the PDE model shows less cooling

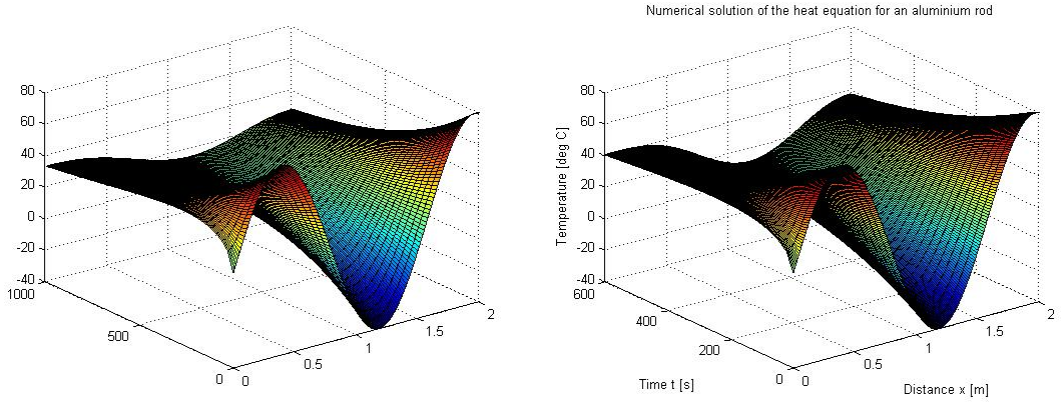

Figure 16 Evolution from initial conditions in state space (left) and PDE models.

In terms of the spatial resolution of the state space model, denser sampling gives better results. Lower resolution leads to mismatching of the heater position and the positions of the system states, and can give inaccurate results (see Figure 17 below). The model can however be manipulated to model the system using states corresponding exactly to the nodes, in which case high spatial frequency information is lost but the general trend is accurate. In hindsight this is probably the best model to use, since it would have reduced computation time for controller synthesis (except where this would have been impossible, see the discussion of $\mathcal{H}_\infty$ design in section 3.3, or where spatial frequency information was useful, see section 5.2). However, the author was wary of the inaccuracy shown by other low-resolution models and a model with a large number of states has been used for all controller synthesis methods.

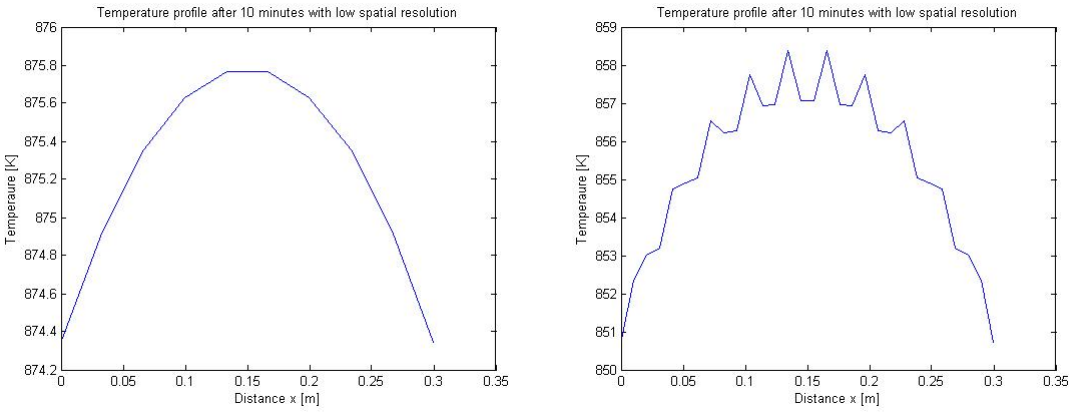

Figure 17 The first scenario is repeated with low spatial resolution. On the left the states' positions coincide with actuator placements, at right the resolution is 30 samples, evenly spaced along the 0.3m length of the rod.

2.8 System Identification Experiments

Unfortunately the complete system will not be ready before this report is submitted. However a section of the array has been assembled in a temporary configuration to test some elements of the model. In preparation for this test a few facts about the system have been discovered.

- Actual maximum power output is only 17W.
- The pulse width from the temperature sensors varies wildly.

The reduction in power output will only affect the controller design of the MPC, since the other controllers do not consider saturation. There may be some reduction in performance, but since the inputs hardly ever saturate positive this is not expected to be a big issue. Measurements are received at the microcontroller in the form of a pulse width modulated signal. The width of the pulse is seen to vary wildly over a larger than expected range. The duty cycle maintains a more consistent value, but the noise is still considerable. Therefore two times must be measured, instead of just one. The temperature is a function of the ratio between them.

For the test configuration, the rod is suspended between two desks with the maximum possible surface exposed to air. Only one heater is used, which is placed centrally on the rod and can be identified by the orange band on the picture below. An array of eight temperature sensors measures the spatial influence of the heater. The sensors are all placed to the same side of the heater, based on the assumption that the temperature distribution will be symmetrical. It was not possible to use a permanent glue to attach the sensors and actuators to the rod, so there is probably a reduction in heat transfer efficiency compared to the final system. Only one microcontroller is used with a multiplexer to interface to the sensor array. Temperature readings (in ratio form as explained above) are sent to a PC for analysis in Matlab. Air temperature was estimated from the readings prior to switching on the heater.

Table 2 System Identification test conditions

Power input (estimated)	12±2W
Sensor spacing	3cm
Sampling time	0.5s
Air temperature	26.7^{o}C

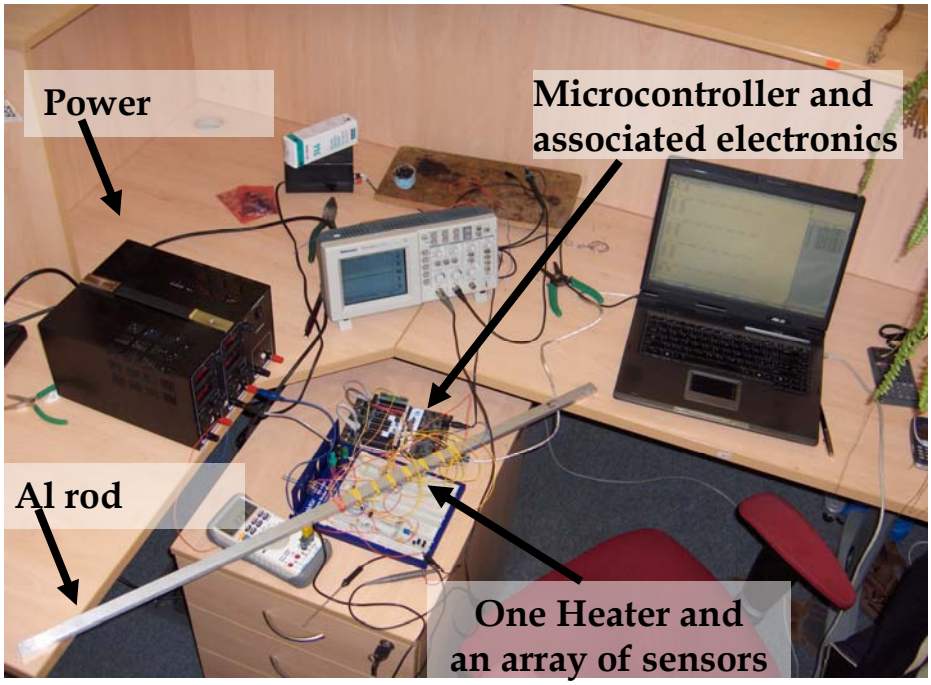

Figure 18 Temporary setup for system identification testing.

Due to time restrictions acceptable results were obtained for only one trial. The measurement noise is highly visible in the raw data – at constant temperature the average of the standard deviations for each sensor is 2.5°C. Therefore prior to analysis, the data is filtered along the time dimension with a low pass filter.

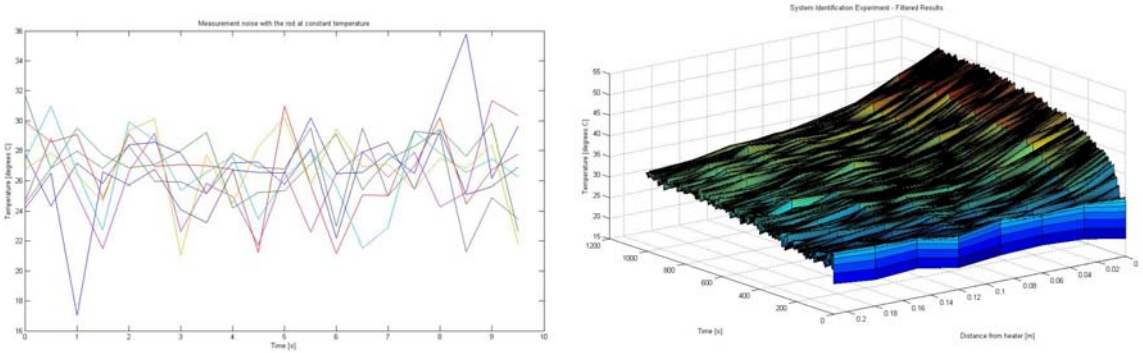

Figure 19 Noisy measurement data for a rod at constant temperature and filtered results from the system identification experiment.

Given the limited amount of data, the best match was sought only for the convection co-efficient *h*. Other parameters were assumed to be at their nominal values.

Simulations were conducted using a range of values for *h*. The accuracy was assessed based on the smallest overall squared error over all times and all positions, and the optimum was found at 31.

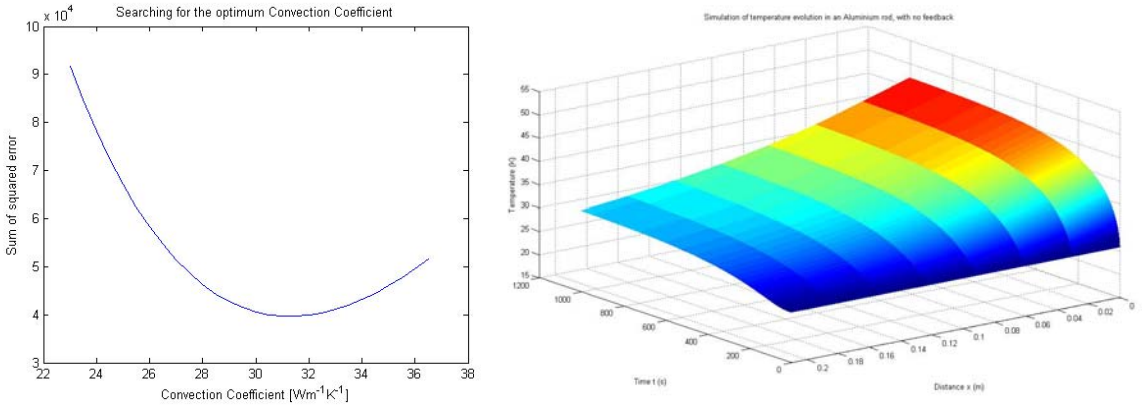

Figure 20 Finding an optimum value for *h* and using it in simulation. The results compare well with the experimental data.

This value is outside the expected range for *h* (2-25) which may affect some of the control designs. Being optimistic, it will be shown later that a major problem for most controllers is slow convection cooling at operating temperatures, so increasing the convection co-efficient is something of an advantage. Uncertainty analysis also shows that the scenario with low values of *h* gave the largest multiplicative uncertainty, so controllers which have been designed for multiplicative error will probably remain valid.

2.9 Plate Modelling

As a small aside to the main track of the project, it was decided to create a model of an aluminium plate to investigate how results could be extrapolated when an additional spatial dimension is added. A plate type surface has many more industrial applications, and with nodes scaling at a rate of n^2, the use of spatially distributed control becomes more advantageous.

Initially, the plate model was created in FEMLAB, by extending the rod and adding more heating patches. It was found however that large numbers of nodes generated a system for which the software was unable to generate a mesh. Therefore the array size was reduced to 4x4.

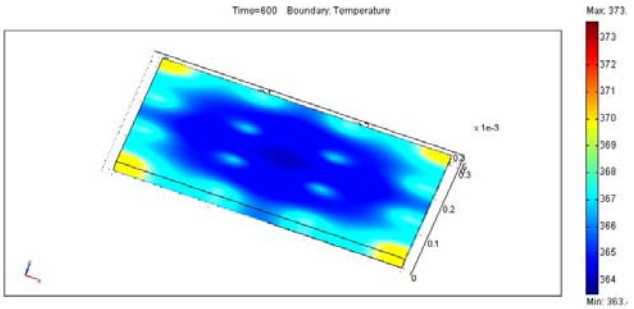

Figure 21 Temperature distribution in a plate after heating for 600s, using a finite element model.

An unexpected observation is that the edges became warmer than the centre. The extra convection at the edge does not compensate for the fact that heaters at the edge are surrounded by a smaller volume of aluminium, and so their heating is more concentrated. This effect was found to be reversed when the plate size is reduced, increasing the density of the heaters in the centre.

The state space matrices become considerably more complicated. There are many ways of indexing the states; here a raster sampling system is selected. Now, the spatial derivative in the conduction term must be two dimensional, and the nearest neighbours will not necessarily be those with the closest index. For instance the first point has neighbours with indices 2 and (N+1) where N is the number of samples in one direction. Conversely, the point with index N does not depend on point N+1 since they are on opposite sides of the plate. This means that the Toeplitz approximation is completely lost although the A matrix remains symmetric. There are now two types of boundary conditions – edges and corners. As explained previously for the rod case, these are modelled by allowing for extra surface area at these states.

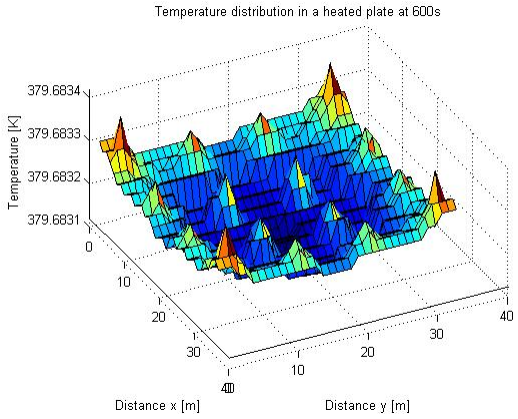

Figure 22 Temperature distribution in a plate after heating for 600s, using a state space model.

The resolution used for this simulation was approaching the memory limits for the computer available, so that using an array with more actuators requires reducing the spatial resolution. Similar to the rod, this results in difficulties mapping the heater influence to the system states, although the model can be modified for the case where each heater is matched directly with one state.

Simulations using this model appear accurate in a general sense, however the temperature variation predicted by the FEMLAB model is severely reduced. The reason for this is still unclear. This part of the project was not pursued further.

2.10 Simulation conditions

Controllers were tested in simulation using the nominal plant for their

- Step response
- Reference tracking for a wave of heat traversing the beam
- Disturbance rejection

For simulating the disturbances, estimates of the magnitude of process and measurement noises are made, since these are unknown prior to system identification. Worst case (6σ) uncertainty is expected to be up to

- ±5°C measurement error
- ±1W heating power
- 10°C variation in ambient temperature

Output disturbance is not considered as there is no conceivable source. If it exists it can always be modelled as input disturbance with a given transfer function G_d.

Input disturbance is modelled as sunlight on one part of the rod and a cool breeze on the other. Heater input disturbance and measurement noise are series of random numbers, with measurement noise filtered to favour low frequencies in both time and space. Each disturbance acts independently for 150s, and for the final period all three act simultaneously.

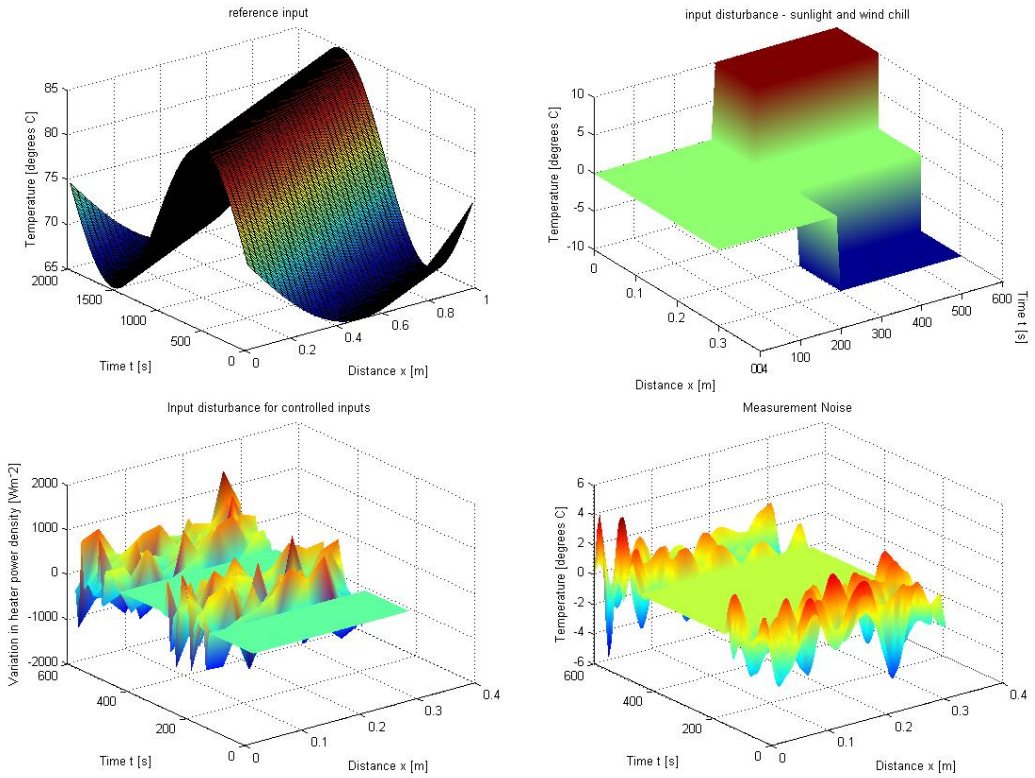

Figure 23 Sinusoidal reference input and three types of disturbance signals.

3 Centralised Control Design

Treating the rod as a MIMO system allows standard methods to be used to design a centralised controller. These designs provide a useful comparison to the spatially distributed methods, and an opportunity to implement the knowledge obtained during earlier courses of this masters program. The theory behind each controller will not be explained here but more information can be found in the references mentioned. Results will be presented graphically for each of the simulations mentioned together with a brief discussion.

A centralised controller is defined as one where information from all sensors are inputs to, and all actuator commands originate from a single regulator. In most cases the regulator is a PC. Measurements taken from each node are passed to the PC, a control signal is calculated and returned to the microcontroller.

3.1 LQR

To design Q, all states are equally important so a scalar multiplied by the identity matrix is appropriate. For R, there are two different types of input – the ambient temperature and the heater power. Since the ambient temperature can not be affected by the controller, this is harshly penalised. The control input is the power supplied to the heaters, and this part of the R matrix is relatively small. For practical implementations it may be more important to save energy, but here it is of secondary importance to the system performance. Care must be taken with scaling for the plant – where inputs can be up to 33,333 Wm^{-2} and outputs should not exceed 100°C.

The limiting factor for decreasing R is actuator saturation. Saturation can not be accounted for directly in LQR design, but should be considered indirectly via this weighting. A controller which requires too large or negative heating will obviously not perform as expected.

In the final design,

$$Q = I_1$$

$$R = \begin{bmatrix} 10^{-7} I_2 & 0 \\ 0 & 10^{20} I_1 \end{bmatrix} \qquad (10.)$$

where I_1 and I_2 are identity matrices with sizes given by the number of states and number of heaters respectively.

An LQR controller uses state feedback, so it must be implemented together with an observer such as a Kalman filter (or directly in an LQG formulation). The Kalman filter is designed to expect the disturbances mentioned above.

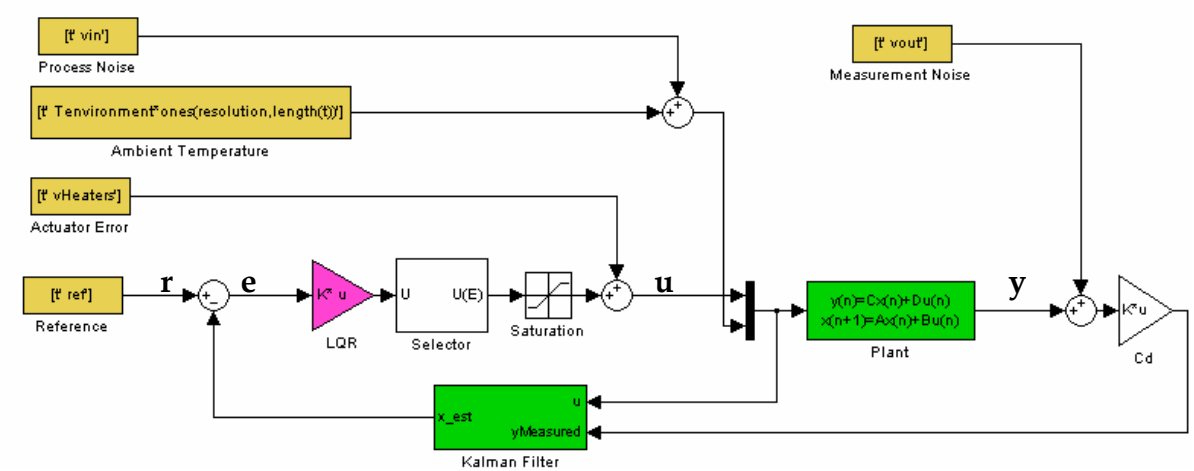

Figure 24 Simulink model of the LQR controlled system. For other controllers the plant, saturation and external inputs will all remain the same.

The step response of the LQR system shows a fast rise time and no overshoot, but the undiminishing steady state error is a concern. Particularly when this increases to 1.7°C with a perturbed plant.

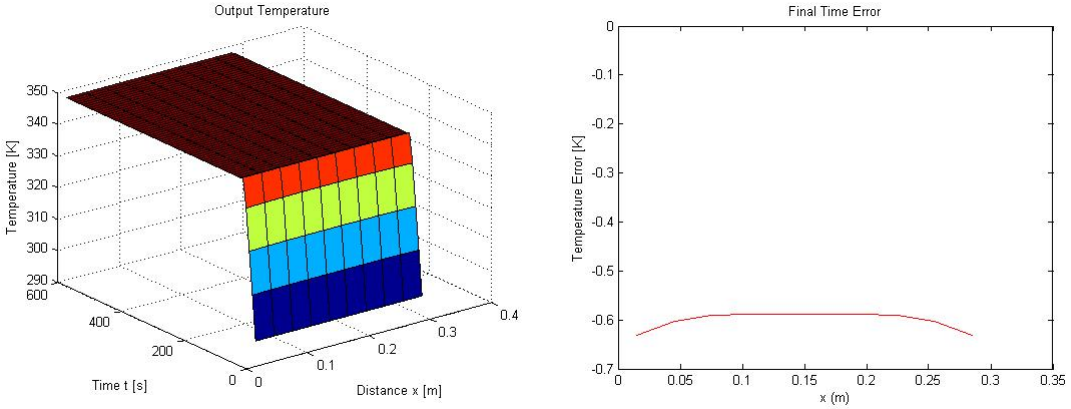

Figure 25 LQR step response.

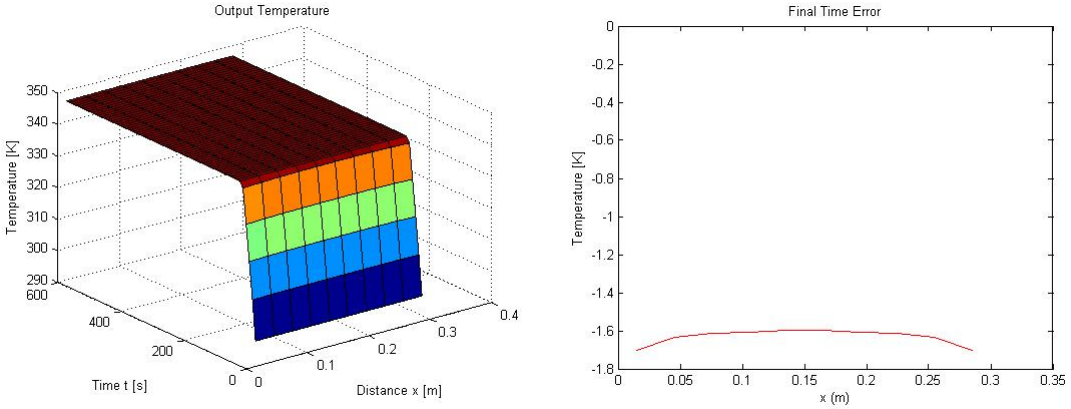

Figure 26 LQR step response simulation using a perturbed plant.

The system tracks a sinusoidal input quite closely, despite approaching the limits of both temporal and spatial bandwidth. Note that it is decreasing reference temperature which causes the largest error, since the actuators saturate low.

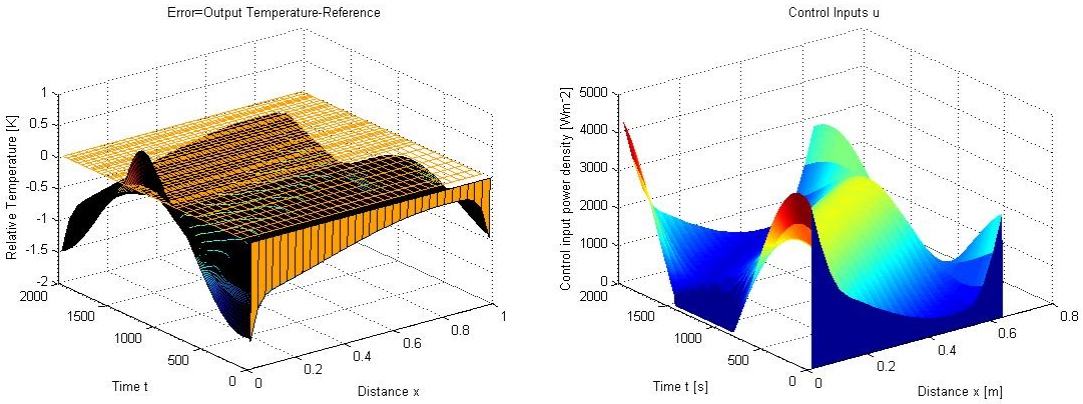

Figure 27 The reference tracking error and control input using an LQR controller.

When the system starts from the reference value, there is a sharp drop over the first sample period, since initially the controller receives zero error and commands zero input. By comparing the times where the three different disturbances are active, it appears that measurement noise has the largest influence. A gain of around 0.25 [no units] was calculated by comparing standard deviations of the noise source (in K) and the output (in K) during a separate simulation. Standard deviation was found from the root mean squared value after the DC component had been removed. Scaling is not considered, so for measurement noise the effect is a 0.25°C output change for each 1°C measurement error, up to the expected maximum of 10°C. So the maximum expected output fluctuation magnitude due to measurement noise is 2.5°C.

In the subsequent time period, the influence of the actuator error is visible but much smaller, with a gain of 0.64 [K/W]. Input disturbance has a smaller effect again, with a gain of only 0.02 [no units]. While the magnitude of the effect is small, it should be noted that changes to the ambient temperature may be effectively permanent and can increase the steady state error. Actuator measurement errors can be expected to average to zero, and will not shift the average value of the output.

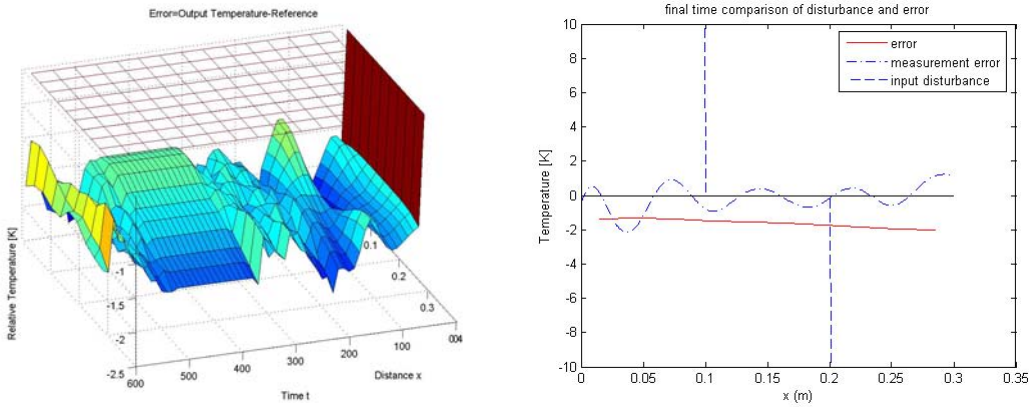

Figure 28 Temperature error during the simulation and at the final time step.

3.2 PI

The steady state error occurring in the LQ controller was thought to be problematic, despite the small magnitude of the error. An integral component was added to correct this, so in the extended design the integral of the error is multiplied by a constant K_i. The result is added to the instantaneous error and multiplied by the LQR gain K_{LQR}. The value of K_i was tuned heuristically to avoid large overshoots while reducing residual error as quickly as possible. The step response shows minimal overshoot and the average error after 600s is reduced by a factor of 5 for the nominal plant and a factor of 10 for the perturbed plant.

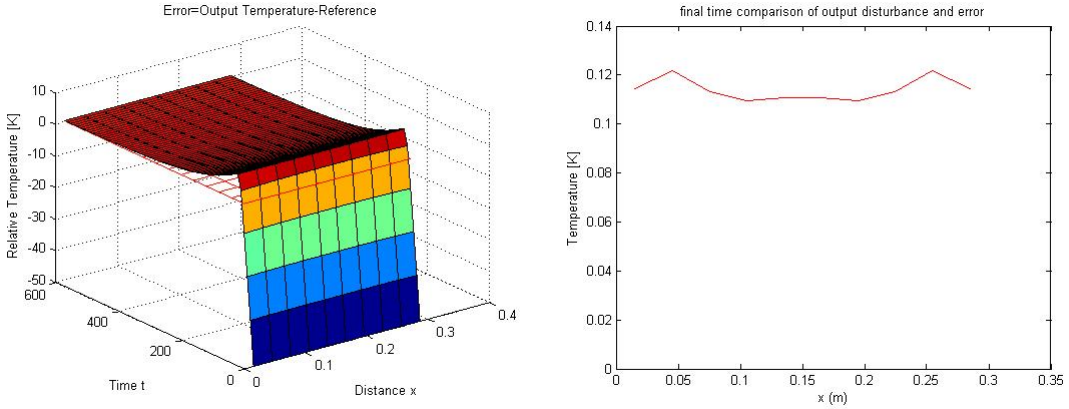

Figure 29 Step Response using a PI controller.

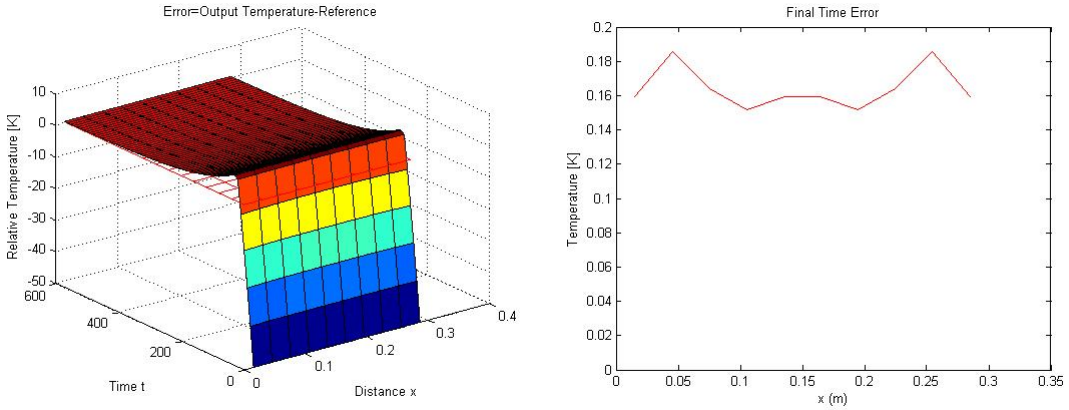

Figure 30 Step response using a PI controller and a perturbed plant.

The reference tracking response is similar to before, but shifted to a higher temperature with lower average error.

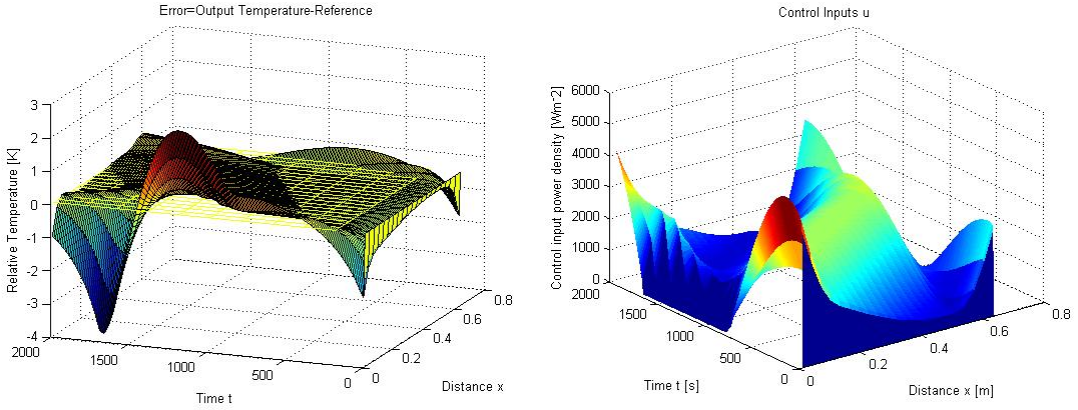

Figure 31 Reference tracking response with a PI controller. The error is shown on the left and the control input to the right.

When disturbances are added, the effects are similar, although the average error attenuates to zero. The gains with respect to measurement noise, actuator error and input disturbance are 0.24, 1.13 and 0.03 respectively. The response to disturbances at the input is worse by a factor of 1.5-2 with respect to the pure proportional controller.

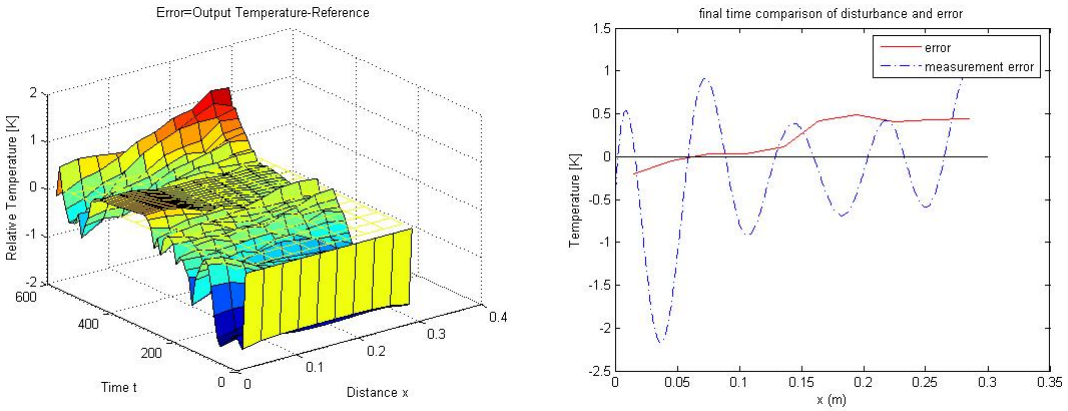

Figure 32 Disturbance rejection simulation for a PI controller.

3.3 $\mathcal{H}_\infty$ Optimal

There were several complications that made it difficult to synthesise an $\mathcal{H}_\infty$ controller. The `mixsyn` function in Matlab will only accept symmetric plants, that is having the same number of inputs and outputs. Intuitively having excess inputs should improve controllability, unfortunately the current algorithm does not allow for this. Considering the state space model, there are N+nH inputs, and only nH outputs. Where N is the number of states and nH is the number of nodes (heaters). So by including ambient temperature as a vector of inputs, the plant model has more inputs than outputs, but also more inputs than states. Any extra 'dummy' outputs would not be linearly independent, which is a requirement for both $\mathcal{H}_2$ and $\mathcal{H}_\infty$ control [15].

Matlab has an LMI routine for solving mixed synthesis which does not require symmetric plants, however the number of states in the model leads to an 'out of memory error' and the solution could not be reached.

Therefore the only possibility was to reduce the ambient temperature to a scalar input, assuming constant temperature along the beam. Then one additional output can be taken from an arbitrarily chosen state to make the plant appear symmetrical. Once the controller is designed, the entries which correspond to this output and to the ambient temperature input are removed.

W1 (Performance), W2 (input penalty) and W3 (model uncertainty) weighting filters have been selected as shown below. All filters are diagonal toeplitz, since the system is spatially invariant and all nodes should have the same weighting. W3 should be representative of the frequency dependent uncertainty in the model. Therefore the Frequency Response Data (FRD) models were generated from the nominal and two extreme plants. Each model consists of the gain from one of N+nH inputs to nH outputs at a range of frequencies. The multiplicative error compared to the nominal plant was calculated for each gain value.

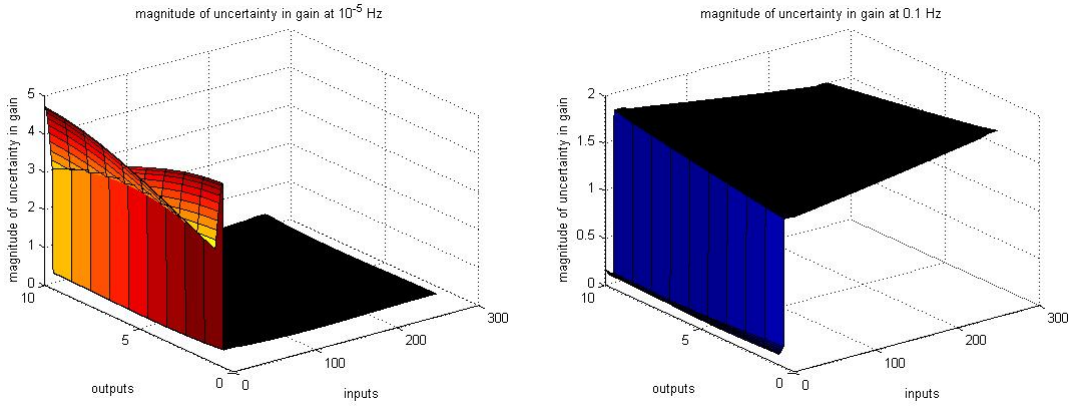

Figure 33 Magnitude of multiplicative uncertainty at low and high frequency. The first 10 inputs are the actively controlled heaters; the remainder correspond to air temperature.

The diagonal nature of the uncertainty can be seen, in that there is less uncertainty in the gain from an input to a nearby output. Of course, this is multiplicative uncertainty, since the magnitude of gain from an input to a distant output is very small.

Uncertainty in the gain from air temperature inputs follow a typical pattern from 10% at low frequencies to 150% outside the system bandwidth. In contrast, at low frequency the error arising from the heater inputs is large – nearly 500%, whereas at high frequency this decays to a more manageable 15%. Higher error at low frequency is unusual for a plant, and is probably due to non-linearities in the

model. At higher temperatures convective cooling is more effective, and low frequency inputs will tend to have a larger difference between their maximum and minimum temperatures. The convection co-efficient is the most uncertain parameter in the system, so coupled with temperature variation it could cause this degree of uncertainty.

Given this unusual situation, it was decided to design W3 based solely on the uncertainty with respect to air temperature inputs. It is hoped that the controller will be able to override uncertainty with respect to the heaters in a closed loop configuration, or at least that error caused by this uncertainty will result in faster changes in the control input, bringing it into the higher frequency region with lower uncertainty. Then W3 is modelled as a first order transfer function of identical magnitude between all inputs and all outputs. There are two possibilities which cover the system uncertainty, but option 'b' has the higher bandwidth so this will be used for controller synthesis.

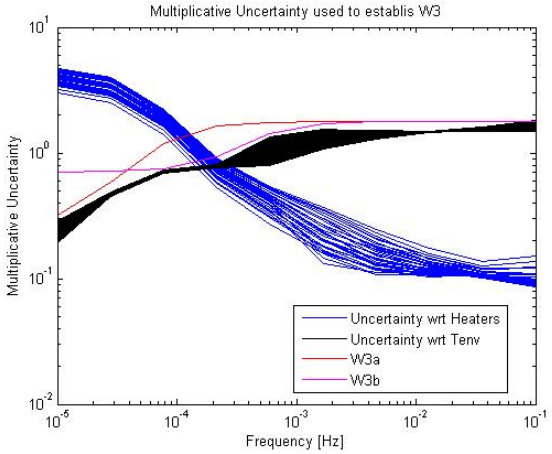

Figure 34 Multiplicative uncertainty as a function of frequency is used to design W3.

The W1 filter is designed to increase gain at low frequencies and at the same time allow performance to deteriorate at frequencies above the bandwidth, where the bandwidth is given by the unity crossing of W3. Therefore W1 for this system is set as a matrix with diagonal elements:

$$W1(i,i) = \frac{0.1s + 10^{-3}}{s + 10^{-5}} = \frac{0.1z - 0.09}{z - 0.9999} \qquad (11.)$$

Note that off-diagonal elements of W1 are set to zero so that performance is only required from an input to its corresponding output.

W2 has small constant elements (10^{-10}) on the diagonal, and zeros elsewhere.

The step response from the system is quite slow, but acceptable. Steady state error is small and appears to be vanishing. The speed of the response can be improved by increasing the bandwidth of W1, but then the stability is not

guaranteed. Of further interest is the anomalous response around the second node. Despite having a symmetric plant and symmetric filters as inputs, the synthesised controller seems to have non-constant gain across the system. This is most probably due to the fictitious output which was added and then removed.

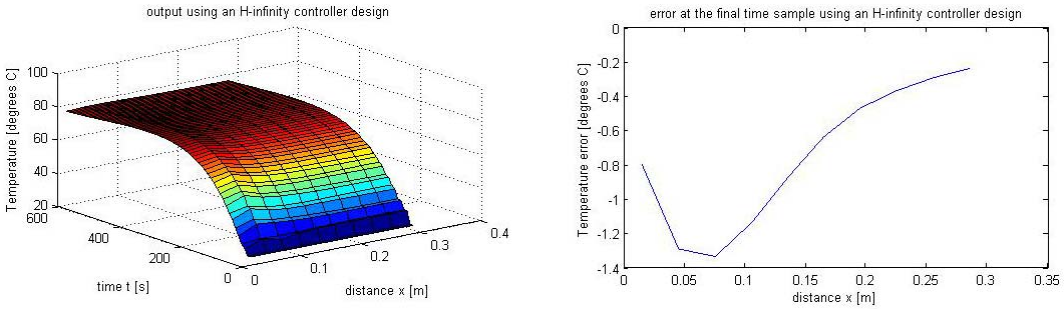

Figure 35 Step response using an $\mathcal{H}_\infty$ controller.

The reference tracking results are very bad, with the output consistently below the reference. The difference is up to 10.4°C. It is possible that the rate of change of the reference is simply beyond the bandwidth of the controller. Again there is an anomalous response, and still more perplexing is its disappearance in the second half of the simulation. Note that the length of the rod for this simulation is only 0.65m, in contrast with 1m used by the other controllers. A longer rod was found to have eigenvalues approaching the imaginary axis and caused computational conflicts in a ricatti equation within the `mixsyn` algorithm.

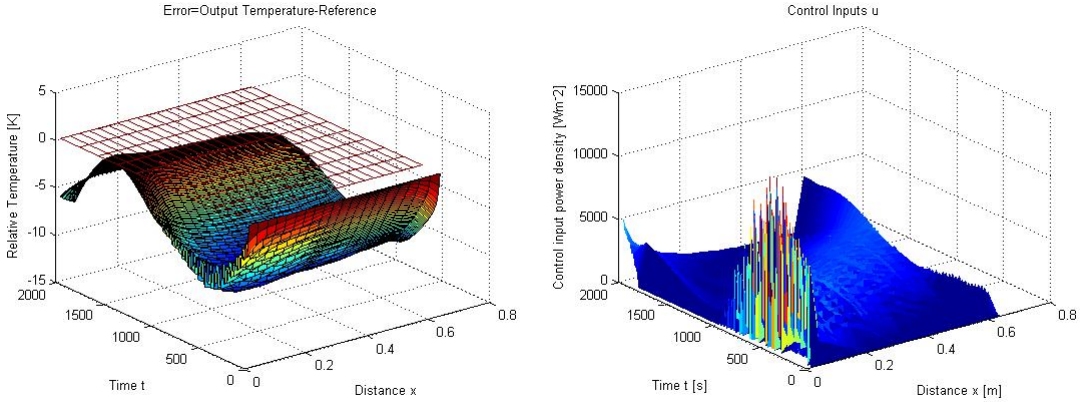

Figure 36 Reference tracking performance using an $\mathcal{H}_\infty$ controller.

While the performance of the system actually seems to improve when measurement noise is added, the overall performance is still much worse than the other controllers. Gains for measurement noise, actuator error and input disturbance were calculated at 1.22, 7.84 and 0.18 respectively.

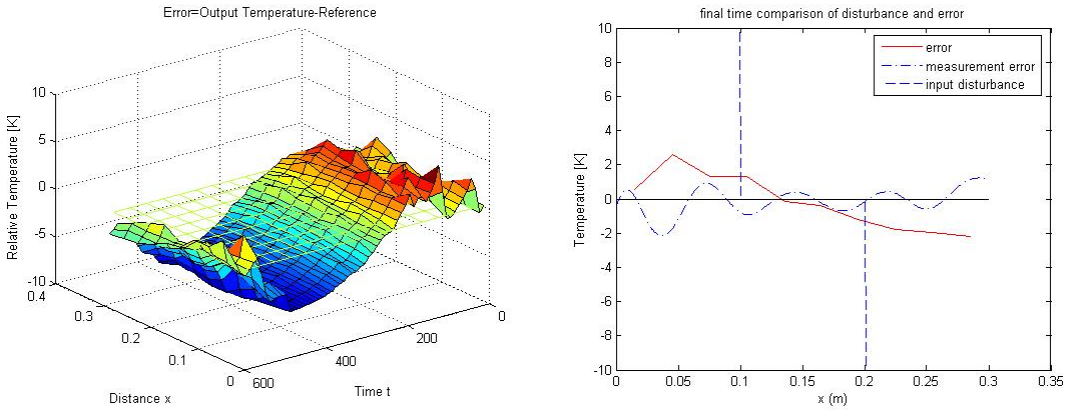

Figure 37 Results from a simulation to test disturbance rejection for an H-infinity controller.

The $\mathcal{H}_\infty$ controller caused many complications during the design process and the performance did not compensate for this. It must be said that the uncertainty considered in the design is large, and it is hoped that with further system identification experiments the magnitude can be reduced. It may also be possible to resolve the anomalous asymmetry in the controller, and so there are potential improvements for this controller that may justify further attention in the future.

3.4 Loop Shaping

The Matlab function `ncfsyn` performs Glover-McFarlane loop shaping and for this project the LQR controller is used as a starting point. The main difficulty was combining the LQR controller and Kalman filter to a single prefilter unit 'W1'. It was also found that the `ncfsyn` function in the robust control toolbox did not work as expected. Firstly, when either the plant or the prefilter is discrete it enters an infinite discretisation loop, therefore both systems are converted to continuous models before calling the function. Secondly, the calculation of CL (closed loop gain) requires a square plant and prefilter. As discussed in the section above it is a complicated process to convert the plant to a square form, and even more so for the combination of LQR and Kalman filter. Since the value of CL is irrelevant to the design, the function was modified to avoid this step. Finally, the algorithm gives a controller which should be implemented in positive feedback.

The standard comparison is between the open loop gains of the system before and after loop shaping. Here the graphs are dense because of the large number of inputs and outputs. The response for the LQR controller does not include the Kalman filter, since this is part of the feedback loop. The graphs show that the low frequency gain has been raised by an order of magnitude and the maximum roll-off rate at high frequencies is reduced. A low roll-off rate at the bandwidth frequency is associated with improved robustness [15].

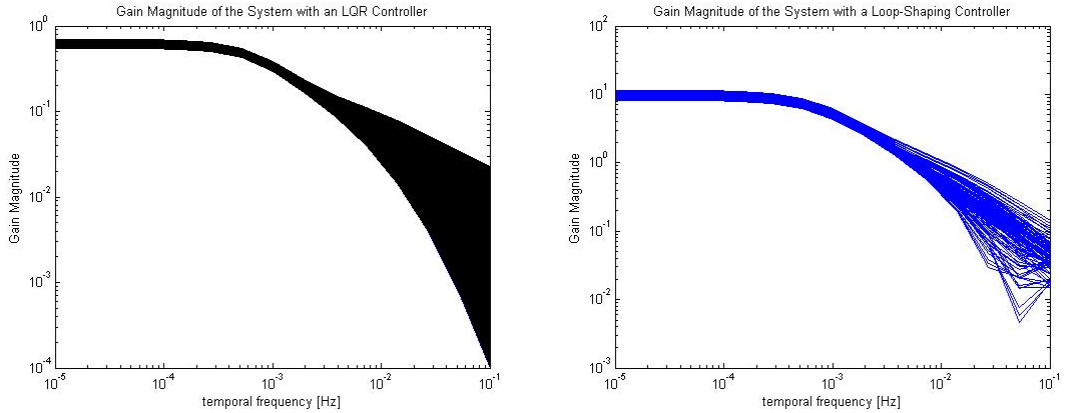

Figure 38 Open-loop gain 'L' for the system with LQR and Loop-shaping controllers.

Below is the step response using the synthesised controller. There is still some steady state error present, and overshoot leads to an increased settling time with respect to the pure LQR designs. The asymmetry of the response derives from the Kalman filter, although the underlying cause is unclear.

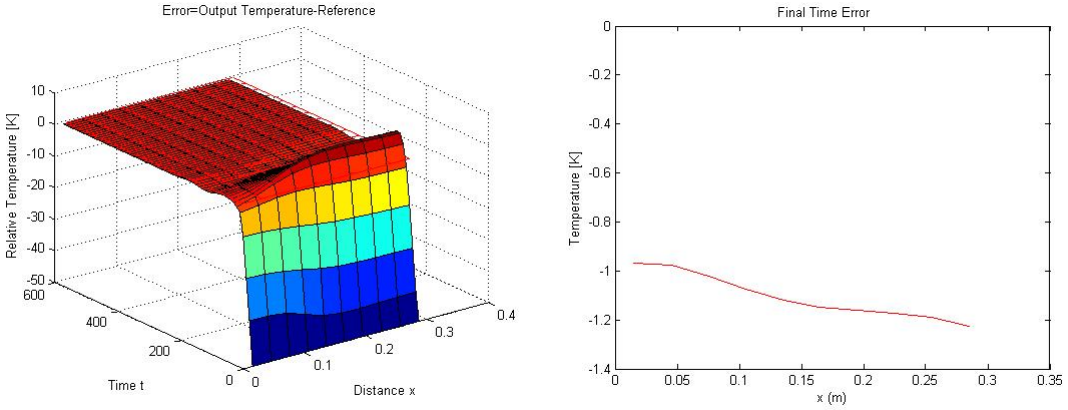

Figure 39 Step response using a Loop-shaping controller.

When the system is simulated with the sinusoidal reference signal, the temperature is consistently below the reference. The magnitude of the error is increased from the simulation with LQR.

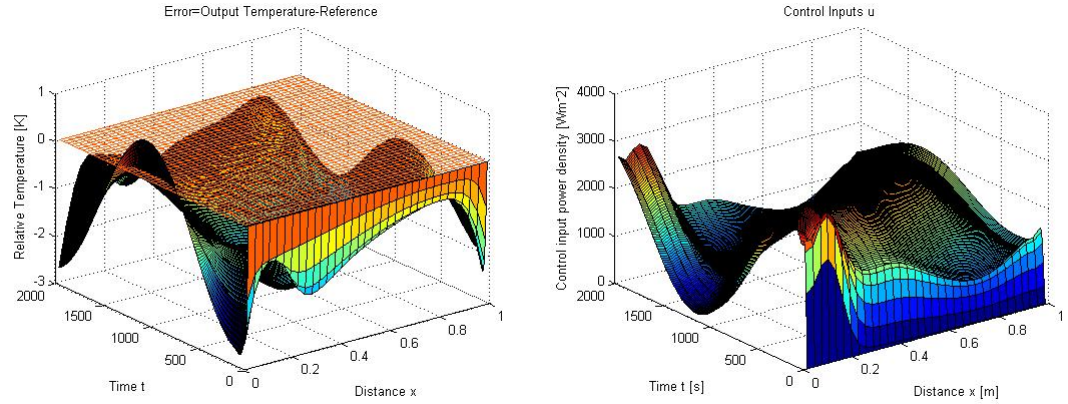

Figure 40 Reference tracking error and control input using a Loop-Shaping controller.

When disturbances are applied to the system, the response is similar to that observed previously. The gains with respect to measurement noise, actuator error and input disturbance are 0.38, 1.35 and 0.05 respectively, all of which have increased.

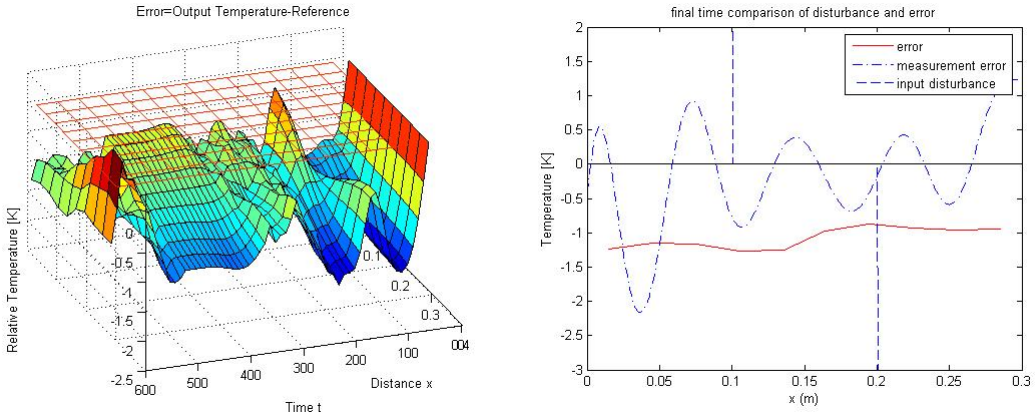

Figure 41 Attenuation of disturbances using the Loop-Shaping controller.

Therefore it must be concluded that the performance of the loop shaping controller in this application is somewhat disappointing, having diminished the performance compared to simpler controllers in all measured aspects except steady state error, where the improvement was minimal.

3.5 MPC

Model Predictive Control employs a markedly different control law to the other designs considered and so its advantages and disadvantages will be briefly explained here. At each sample time, a cost function is optimised with respect to an input trajectory, where the input and the corresponding response are predicted over several sample times. Therefore MPC is the only method which can explicitly include saturation limits as it calculates the trajectory. The main drawback is the computational complexity of optimising a predicted response at every sample time. For thermodynamic applications the relatively long time constants allow this luxury for a centralised system. So far the computations have been considered too complex for implementation on microcontrollers and therefore not suitable for distributed control, however some new methods have looked at ways to preprocess sections of the optimisation routine which may make this a possibility in the future [16,17]. While this is an interesting development, it will not be pursued in this thesis.

To design the controller, the first iteration was made using the `mpctool` from the model predictive control toolbox. This allows all the design and simulation options to be set in a GUI, and it has an excellent tutorial. For quicker design

iteration, the procedure was later transferred to an m-file. Settings for the controller are shown in Table 3 below.

Table 3 Design parameters for a model predictive controller where N is the number of nodes.

Parameter	**Value**
Number of manipulated variables (u)	number of heaters
Number of measured disturbances ($T_{environment}$)	N
Prediction Horizon	10
Control Horizon	10
Actuator saturation limit – low	0
Actuator saturation limit – high	33,333
Actuator rate of change limit – low	-33,333
Actuator rate of change limit – high	33,333
Manipulated Variable weight	10^{-7}
Manipulated Variable rate weight	10^{-10}
Output variables weight	1

The step response of the MPC controller is the quickest of all responses, yet still has no overshoot.

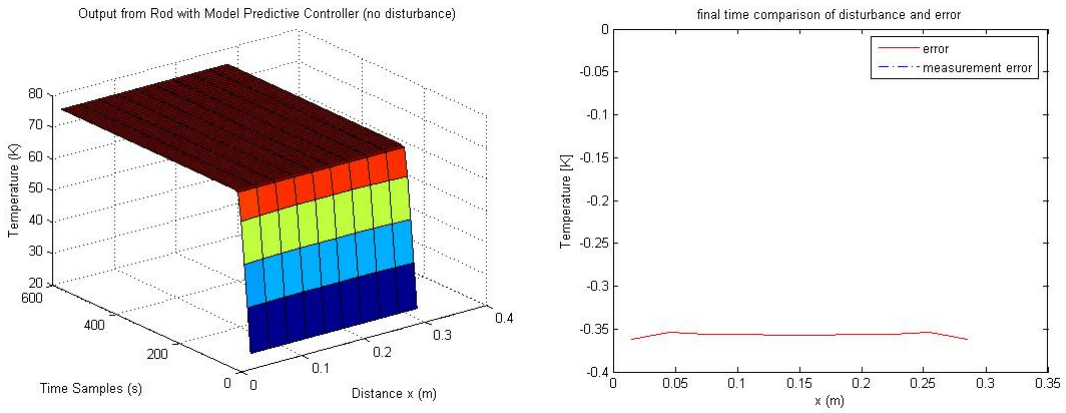

Figure 42 Step response using an MPC.

The MPC controller tracks the reference input very closely, apart from some anomalies at the beginning. The good performance is expected since the controller can foresee the changes in the reference value and adjust accordingly in advance. If the changes were faster or the horizon were longer the performance would be further improved, but even in this configuration the simulation took much longer than the other models. It is interesting to observe that the controller has commanded a negative input *u* at the beginning of the simulation, despite the saturation limits imposed as part of the MPC controller design.

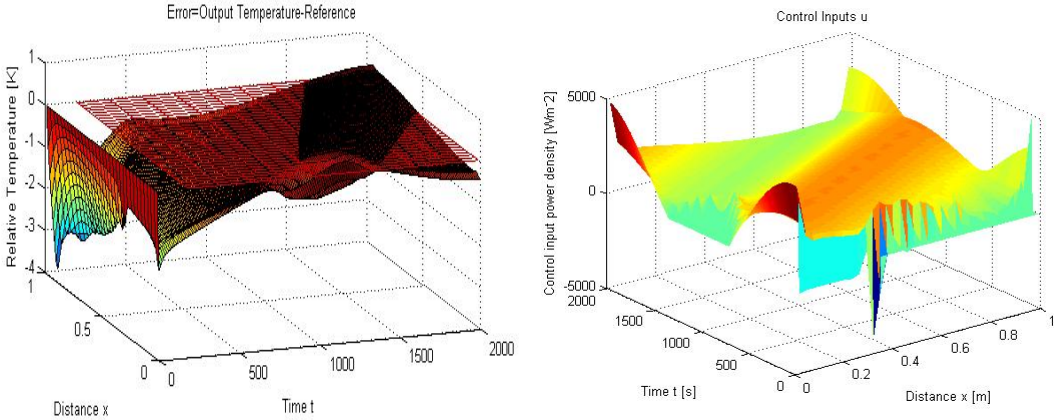

Figure 43 Reference tracking performance of the system using an MPC controller.

The MPC controller shows high susceptibility to measurement error, at the end of the simulation the output almost replicates the error values. The other noise sources have a similar effect to previous designs. Gains for measurement error, actuator error and input disturbance are 0.88, 1.93 and 0.04 respectively. The controller is not designed explicitly for robustness so it is not surprising that noise attenuation is not one of its strengths.

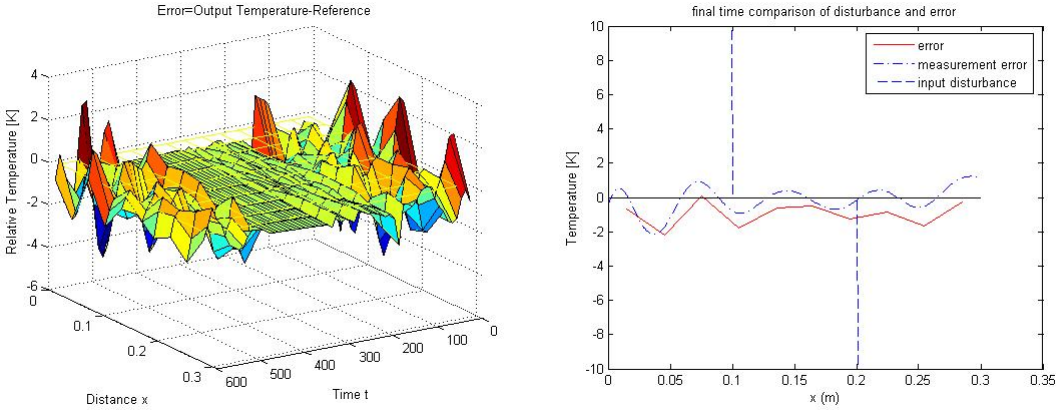

Figure 44 Disturbance rejection simulation results using an MPC controller.

As expected, a controller which can take account of the saturation of actuators and predict the future trajectory of the output performs well in the performance tests. The compromise comes in computation time, where one of the longer simulations takes 680 seconds for MPC and only 84 using LQR.

4 Decentralised Control Design

Decentralised control breaks the system down into a series of SISO systems, such that each microcontroller receives information from and has control over one node. Using this type of controller carries an implicit assumption that the cross-influence or overlap between nodes of the system is negligible. While this is true for some applications, adaptive optics is a common example where the overlap is of critical importance [6] and there are several others. One paper of interest by Lunze and Abraham [4] outlines a design methodology for a heat regulation system with similar dynamics to the one here. The paper claims that the performance objectives can be achieved by designing an array of robust SISO controllers and regarding overlap as disturbance, and further that this procedure can be generalised to a whole class of systems.

4.1 PID

The classical SISO controller is the well known proportional, integral and derivative controller. When tuning by Ziegler-Nichols, the stability limit was reached with K_p=4300 and an oscillation period driven by the sampling frequency. After some additional heuristic tuning to remove oscillations, the final design uses

K_p=2580
K_i=1720
K_d=107.5

at each node in the system.

The step response is fast with a small overshoot. There is a small persistent error which will eventually decay to zero.

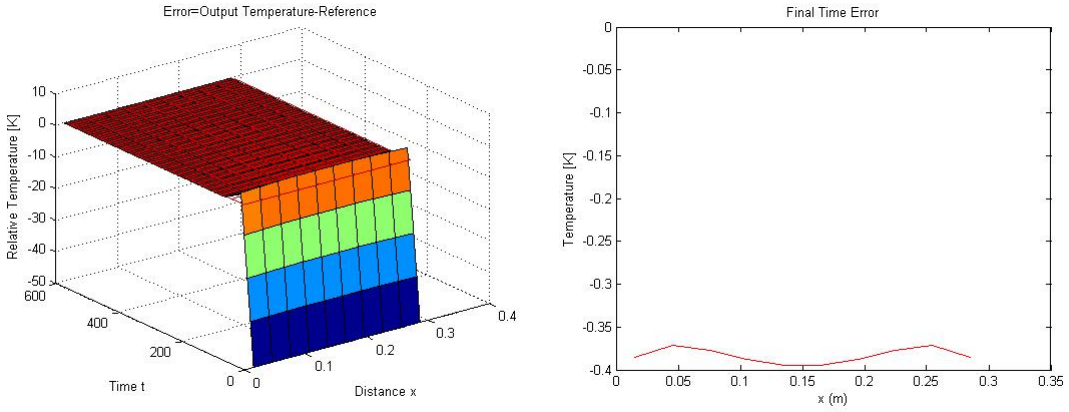

Figure 45 Step response using a decentralised PID controller.

The output only reaches the reference temperature during the time of maximum cooling. However the overall error is very small – the wave input is faithfully reproduced.

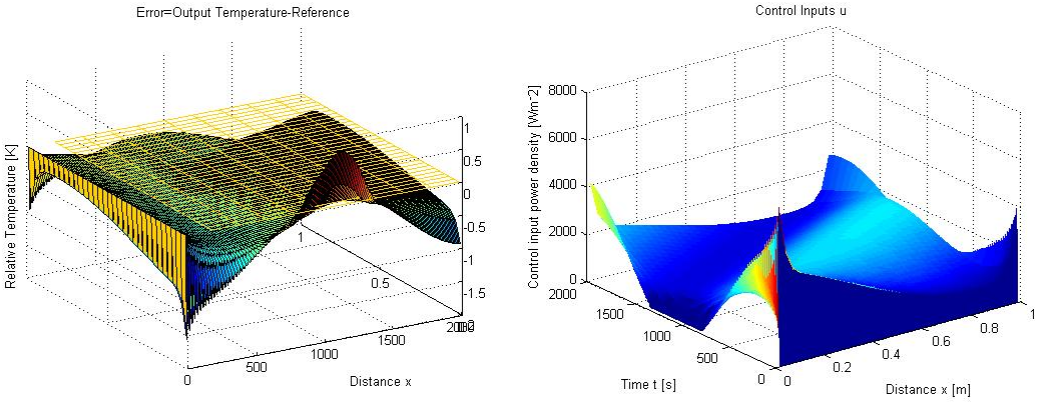

Figure 46 Reference tracking performance using a decentralised PID controller.

The PID controller has a good response to actuator error and input disturbance, with gains of 1.12 and 0.02 respectively. Measurement error still creates some issues, having a gain of 0.51. This can be seen easily on the graph where the deviations are largest at the beginning and end of the simulation where measurement noise occurs.

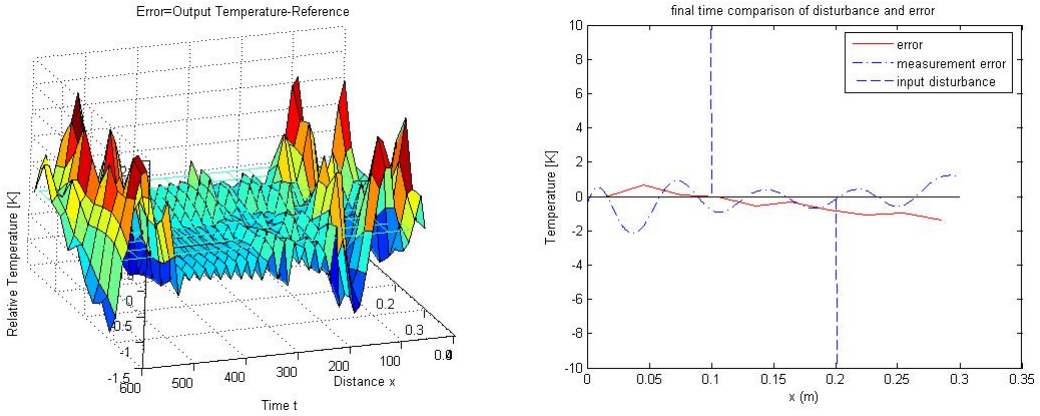

Figure 47 Disturbance rejection performance of a decentralised PID controller.

4.2 Control Design from Lunze 1992 - PI

The theoretical details of this paper are explained in the introduction of this report, see section 1.2. This section will deal only with the control design technique. While the explanation in the paper is extremely vague and difficult to follow, the authors list three design specifications

1. Closed Loop stability. Where arbitrary control stations are disconnected, this is the only requirement that must be fulfilled.
2. Step commands are followed and step disturbances are rejected.
3. Overshoot less than 2K, or follow a profile close to the operating point within 0.1K. Since the profile is not relevant to this application, only the overshoot requirement has been used.

Requirement 2 would be fulfilled by any controller with an integral component, which may be their motivation for choosing a PI control law. The first and third

requirement may be interpreted as a form of mixed sensitivity, and so this was the first technique that was tried.

4.2.1 By Mixed Sensitivity

To obtain the three filters for mixed sensitivity, the temporal transfer function was re-derived for a perturbed plant. As suggested in the paper, the worst case perturbation is one where a node of the system is disabled or malfunctioning. In this case, the heating at the node comes only from neighbouring nodes via conduction. Therefore there is a slight delay in the response, aswell as a reduced DC gain. Converting the impulse response to a transfer function and then to a frequency response data model gives the bode plot shown below, and evaluating the multiplicative error at each frequency gives the filter W3 for mixed synthesis. The filter is approximated by a first order function using the minimum and maximum errors.

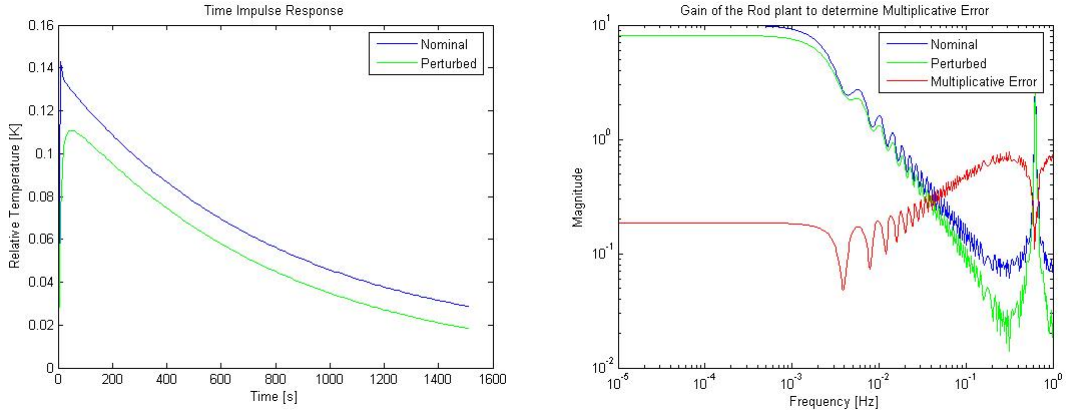

Figure 48 Response of the nominal and perturbed system are simulated to calculate multiplicative error.

Since actuator weights are not important, W2 is left blank. W1 is the tuning parameter, and is set to have a lower bandwidth than the limitations imposed by W3. At the same time, high DC gain is desired to minimise steady state error. It was found to be extremely difficult to design a robustly stable controller – the best gamma value achieved is 0.9841 and still there is a closed loop pole at 0.9999 – very close to instability.

When the system performance is tested, the response was found to be exceedingly slow. Even when the simulation is extended for several hours, some error remains. Therefore this design was rejected and not subject to further testing.

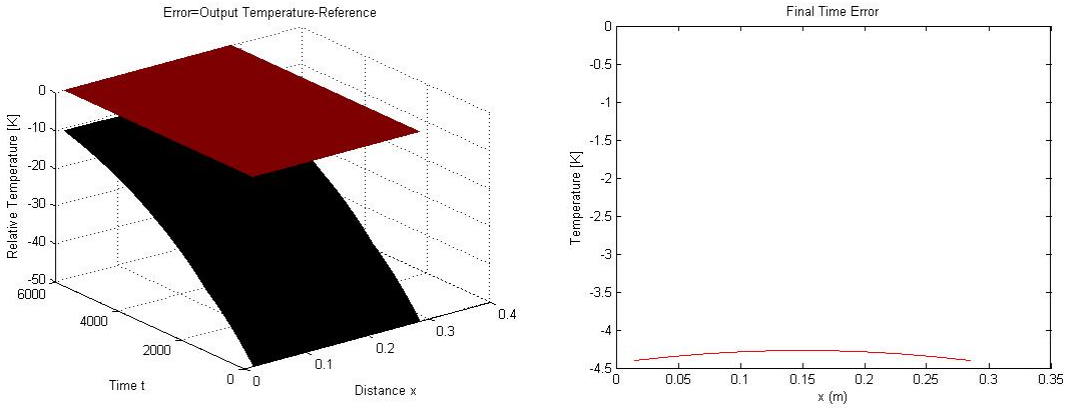

Figure 49 Step input response using a PI controller from a mixed sensitivity design.

4.2.2 By Considering a First Order Approximation

Later, after successive readings of the paper, the intention of the authors became clearer. They have approximated the dynamics at each node as a first order system. Control feedback is internalised to the Ā matrix, which becomes a second order system with the addition of the PI controller pole. Conduction '*s*' from neighbouring nodes is considered as an input whose magnitude is dependent on the output at neighbouring nodes, so that it resembles another form of feedback incorporating spatial influence.

$$\begin{bmatrix} \dot{x}_i \\ \dot{x}_{Ii} \end{bmatrix} = \begin{bmatrix} \left(-\frac{1}{\tau} - e_i \frac{k.k_P}{\tau} \right) x_i - e_i . \frac{k.k_I}{\tau} . x_{Ii} \\ x_i \end{bmatrix} + \begin{bmatrix} e_i \frac{k.k_P}{\tau} ref_i \\ -ref_i \end{bmatrix} + \begin{bmatrix} \frac{k}{\tau} s_i \\ 0 \end{bmatrix} \qquad (12.)$$

$$s = L.c'.x$$

where:

x_I is the integral of error
e_i is the functionality of the control stations (1 for functional or 0 for disabled)
k is the static gain of the system
τ is the time constant of the system

From this point, stability and overshoot conditions can be expressed in terms of the poles of the system. Alternatively, the paper suggests a method of calculating stability which uses the spatial influence information in a matrix '*G*'. The design procedure will be explained step-by-step below.

1. Derive the 'static gain matrix' K_s. K_s can be found from using the steady state spatial influence function coefficients (see section 2.6.1). For each row of the matrix, the function co-efficients are offset such that the peak influence is on the main diagonal.
2. Find the 'static gain constant' k by taking the inverse of the sum across one row of the full static gain matrix K_s.
3. Calculate the cross-influence function $\hat{L}$ from

$$\hat{L} = diag\left(\frac{1}{k}\right) - \hat{K}_s^{-1} \tag{13.}$$

where the 'hat' symbol indicates that the full matrix is truncated to 3x3, considering only the current zone and its nearest neighbours. The truncation should avoid the system boundaries.

4. Find the time constant τ by simulating the full system model cooling from a raised initial temperature and measure the time to change by 63%. The rate of change of temperature will be higher than this during heating.
5. Generate F and G matrices. $\hat{L}(2,2)$ represents the state evolution's dependence on its current value, and $\hat{L}(2,1)$ represents the dependence on its neighbours' current value. Since the influence function is symmetric, $\hat{L}(2,1) = \hat{L}(1,2)$.

$$F = \begin{bmatrix} -\frac{1}{\tau} - k_P \frac{k}{\tau} + \hat{L}(2,2)\frac{k}{\tau} & -k_I \frac{k}{\tau} \\ 1 & 0 \end{bmatrix}$$
$$G = \begin{bmatrix} \hat{L}(2,1)\frac{k}{\tau} & 0 \\ 0 & 0 \end{bmatrix} \tag{14.}$$

6. Repeat all steps except 4 for the perturbed system. The cooling time is independent of the functionality of heaters.
7. Perform a gridded search to find values of k_p and k_i which satisfy the stability and overshoot criterion:

$$\begin{aligned} &Stability: \quad real\left(eig\left(F,\ F-\sqrt{2}G,\ F+\sqrt{2}G\right)\right) < 0 \\ &\qquad stable\ for\ Nominal\ and\ Perturbed\ systems \\ &Overshoot: \quad \zeta_{system} < \zeta_{specifications} \\ &\qquad for\ the\ \text{sec}ond\ order\ closed\ loop\ systems \end{aligned} \tag{15.}$$

Since the overshoot criterion is not defined for the rod system, the 2K criterion from the paper has been used, giving a minimum damping ratio of 0.72.

8. Choose K_P and K_i from the range solutions which fulfil the criterion.

The figures below show the process of choosing PI gains. Any value of K_P above 1.7×10^4 achieves stability, and the upper bound on K_i depends on the choice of K_P. The choices are confirmed in sisotool, although this ignores the contribution of spatial influence. Higher gains than expected are necessary to move the locus away from the imaginary axis to achieve robust stability, and to bring it close to the real axis and reduce overshoot. The 'x' marks the final controller selection:

$K_P = 2.5 \times 10^4$
$K_i = 2.5 \times 10^6$

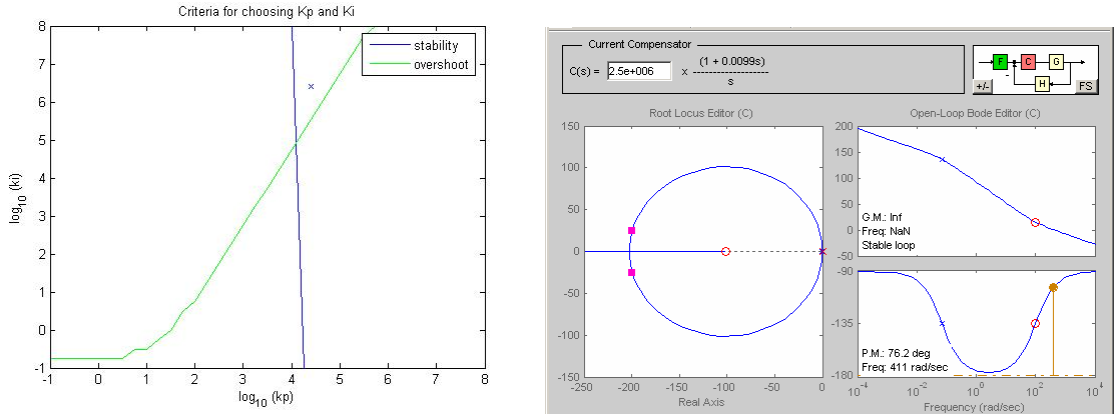

Figure 50 Designing a PI controller, on the left are the criteria from the paper, with the final selected point marked. The results are confirmed using sisotool.

When the controller is implemented with a step input, it can be seen to rise very quickly to the reference value, with a small initial overshoot. Thereafter, however there is a large, unacceptable oscillation. This demonstrates the dangers of designing a controller in the continuous domain and ignoring saturation limits. When the temperature falls below the reference, the actuators are switched on at maximum power, but cannot be switched off again for ten seconds - the temperature rises dramatically. When the controller wishes to reduce the temperature, the actuators cannot provide negative heating and saturate low. The steady-state error is not relevant here, so the commanded plant input shows the problem with saturation.

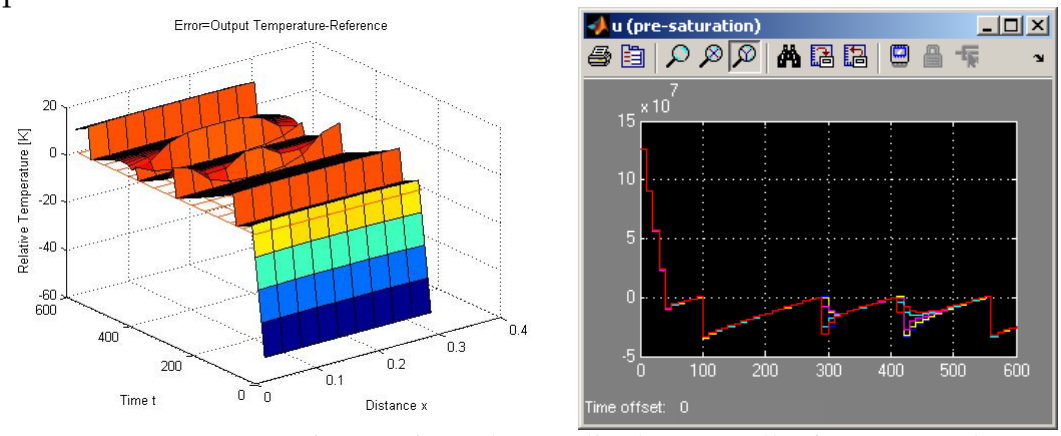

Figure 51 Response to a step input using a decentralised PI controller from a paper by Lunze. 'Pre-saturation u' is the command from the controller, which is often beyond the plant's capability.

When the system is commanded to follow the sinusoidal reference, the performance is surprisingly good, although it is still troubled by large oscillations.

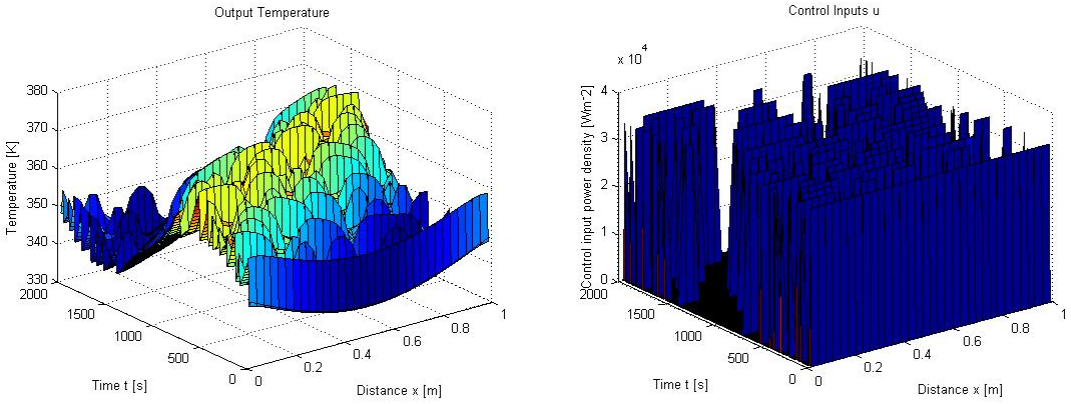

Figure 52 Reference tracking performance using a decentralised PI controller from a paper by Lunze.

The figures below show that the system performance is not adversely affected by noise, in fact, it appears to be improved for the case of measurement error. If the detected temperature distribution has a random element, only some sensors detect that the temperature is below the reference. Then only those actuators are switched on, and the rod temperature does not rise as drastically as before. The gains for measurement error, actuator error and input disturbance are found to be 1.34, 23.8 and 0.37 respectively, although the large values reflect more the oscillatory nature of the controller response than the effect of noise.

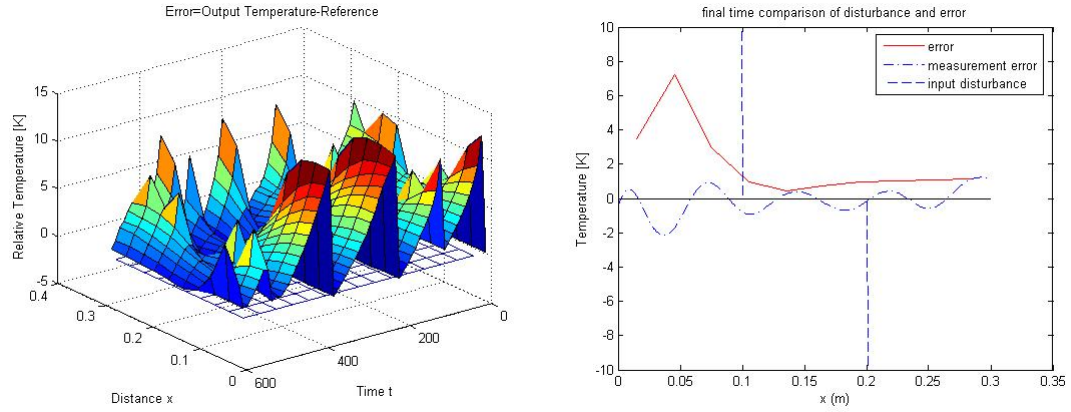

Figure 53 Disturbance rejection performance using a decentralised PI controller from a journal paper.

5 Spatially Distributed Control Design

There are numerous methods for designing spatially distributed controllers in the literature. A few have been chosen, primarily for their conceptual simplicity, which is appropriate for the author's level of experience in this field, and as the first controllers to be implemented on the rod hardware. Once these are operational, more complex designs can be developed in the future.

5.1 FIR and PID

This concept is taken from a paper which has spatial IIR filters as its main topic [18]. By combining several measurements from a region to give a singe weighted average value, standard SISO techniques can be used to govern the plant input.

The values for the non-causal FIR filter were taken from a truncation of the LQR controller designed in section 3.1 and the PID is tuned by Ziegler-Nichols. The FIR component is scaled for unity gain. For the boundary conditions, the filter must be modified since no measurements can be taken from beyond the boundary. Here the LQR is consulted again and the diagonal and two neighbouring values are taken and scaled to create a causal or anti-causal FIR as appropriate.

The LQR feedback equation for u_i is given by

$$\ldots k_{i-2}T_{i-2} + k_{i-1}T_{i-1} + k_iT_i + k_{i+1}T_{i+1} + k_{i+2}T_{i+2} + \ldots$$

the FIR filter becomes

$$\left[\frac{k_{i-1} \quad , \quad k_i \quad , \quad k_{i+1}}{k_{i-1} + k_i + k_{i+1}} \right] \tag{16.}$$

The PID values are

$K_p = 2646$

$K_i = 4.41$

$K_d = 110.25$

The step response shows a small overshoot which decays towards zero steady state error. This situation where the temperature at the nodes is above the reference is actually favourable, since the spaces in between the heaters have a lower temperature.

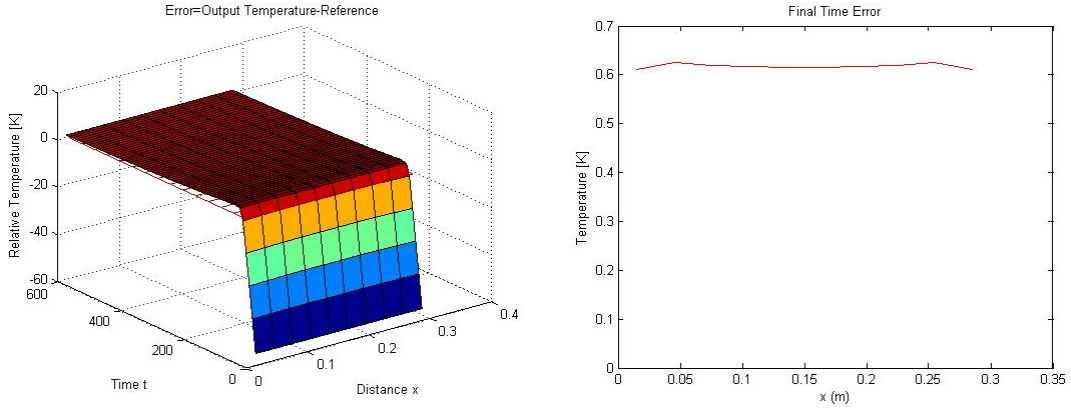

Figure 54 Step response of the system using a spatial FIR filter and PID controller.

The sinusoidal reference is successfully tracked with a consistently small error. As expected, the largest errors are at the boundaries where the infinite length assumption breaks down.

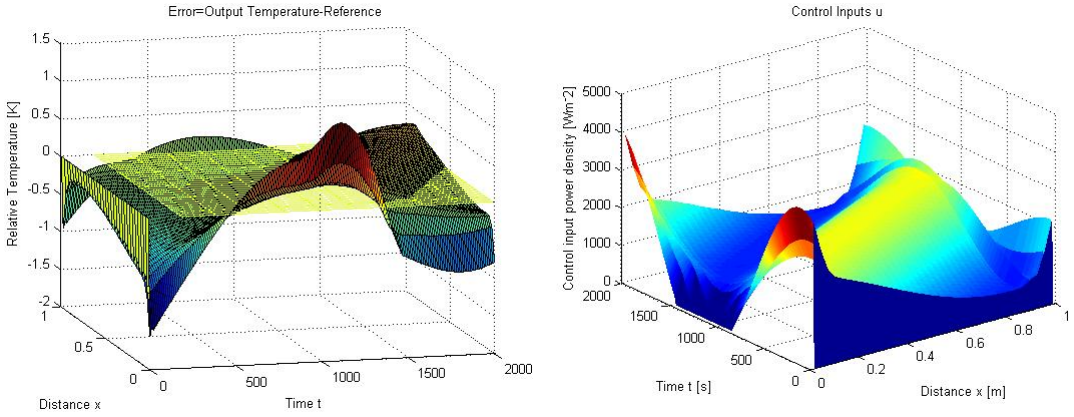

Figure 55 Reference tracking performance using a spatial FIR filter and PID controller.

All disturbances are well attenuated, perhaps because the localised structure of the controller does not allow them to propagate beyond a certain region of influence. The maximum error is small; less than 2K. The gains for measurement noise, actuator error and input disturbance are 0.43, 1.00 and 0.02 respectively.

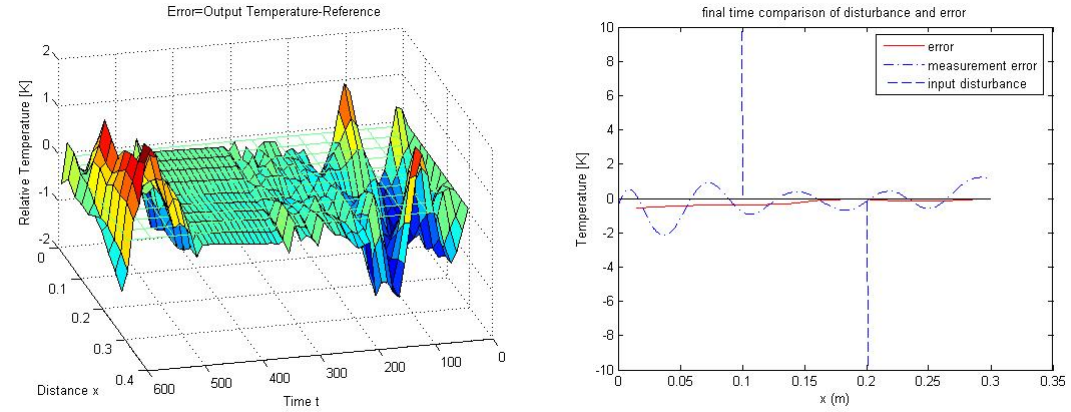

Figure 56 Disturbance rejection simulation results using a spatial FIR filter and PID controller.

5.2 Control Design from Stein 2005 – Mixed Sensitivity

This paper develops a method for designing spatially localised controllers. Transient behaviour of the system is assumed to be trivial relative to the steady state spatial response. In this way he is able to ignore the time response and focus on the spatial dimension(s). The justification given is that the design objectives (equation 18) have worst case conditions at steady state (when the laplace variable s=0). Also the adaptive optics system considered has powerful actuators which dominate the dynamics of the mirror surface. The design proceeds as follows:

1. The structure of the control law is set as

$$Plant: \quad output = Hu + disturbance$$

$$controller: \left(\frac{du}{dt}\right)_m = u_{cmd,m} - \sum_{o=-\infty}^{\infty} a(m-o) \qquad (17.)$$

$$u_{cmd,m} = -\sum_{o=-\infty}^{\infty} k(m-o).(output_o - reference_o)$$

   Where k is the feedback gain and a is integrator leakage. Setting a=0 gives a PI response, which can become unstable for perturbed plants.
2. Recall the spatial impulse response $h_k(x)$ and the related transfer function $h^*(v)$ from section 2.6.1 for the nominal plant and perturbed models for conditions which give maximum and minimum cooling.
3. From the impulse response, derive the static gain matrix H (identical to K_s in the paper by Lunze). Calculate the maximum singular value $\bar{\sigma}$(H) which is interchangeable with $\|h\|$ for large array sizes.
4. Calculate the multiplicative error for the spatial impulse response as a function of x. The maximum error is denoted Δ.
5. Evaluate the transfer function $h^*(v)$ for v gridded from 0-π on the unit circle. This gives the response of the system to input with spatial frequencies from constant to infinite frequency.
6. Select design parameters p_0^*, p_1^*, p_2^*, W_k, W_a for the linear program. Optionally, the paper explains how p_{1max}^* and p_3^* can add further constraints to the solution.

Robust Stability :

$$a^* + (h^* \pm \Delta\|h\|)k^* \geq p_0^*$$

Robust Small Deformation Errors (Nominal Performance) :

$$a^* + (h^* \pm \Delta\|h\|)k^* \geq p_1^* \quad \text{valid only where } h^* - \Delta\|h\| > 0 \qquad (18.)$$

Robust Bounded Actuator Displacements :

$$a^* + (h^* \pm \Delta\|h\| - p_2^*)k^* \geq 0$$

Minimisation Criterion

$$\int (W_k * k^* + W_a * a^*) dv$$

Subject to

$$a^* \geq 0$$

7. Solve the system of equations 18 using the Matlab `linprog` function to find a^* and k^*. Note that the inequality signs must be switched by multiplying by -1.
8. Take inverse fourier transforms of a^* and k^* and truncate to the desired length. Since a and k should be symmetric filters, the co-efficients can be taken in either direction starting from the first value (going backwards implies taking the first value, then the last, then the second-last and so on). Only the first value, which becomes the central filter co-efficient, is unique.
9. Implement a and k in the control law from step 1. At the boundaries, where the spatial sums in equation 17 become impossible, the value at the boundary is duplicated as both T_m and T_{m-1} (or T_{m+1} depending on which end). This corresponds to an imaginary extension to the rod which has the same temperature as the last node.

Is was found that for many combinations of parameters the solution does not converge at all gridded frequencies, and that convergence was not always linearly dependent on each variable. An acceptable controller was achieved with the following values:

$p_0^* = 10^{-6}$
$p_1^* = 0.025$
$p_2^* = 15$
$W_k = 10^{-9}$
$W_a = 10^4$

p_0^* is very small, and only serves to maintain stability at all frequencies. p_1^* represents the desired minimum temporal bandwidth – this condition is only imposed within the spatial bandwidth. p_2^* is technically a penalty on the energy for actuators. In this system we do not wish to penalise energy for the heater, but this serves to discourage negative inputs which are impossible. Therefore higher

values of p_2* reduce overshoot. W_k and W_a penalise the control gain and leakage terms. Stein recommends a small W_k, while a large W_a will ensure that leakage – and therefore steady state error – is kept small. The parameters can be made dependent on the spatial frequency, but this has not been attempted yet.

There is an inherent danger in transferring the system dynamics to the frequency domain, since the fourier transform assumes an infinite domain. This is one of the common assumptions for spatially distributed control, but where it does not hold there will be windowing effects. The same problem applies for the inverse fourier transform, particularly since the designed controllers must be truncated to a small number of co-efficients for practical implementation.

The system responds relatively slowly to a step input, with a larger than desired overshoot. There is some persistent error, whose magnitude is directly related to the leakage term a*.

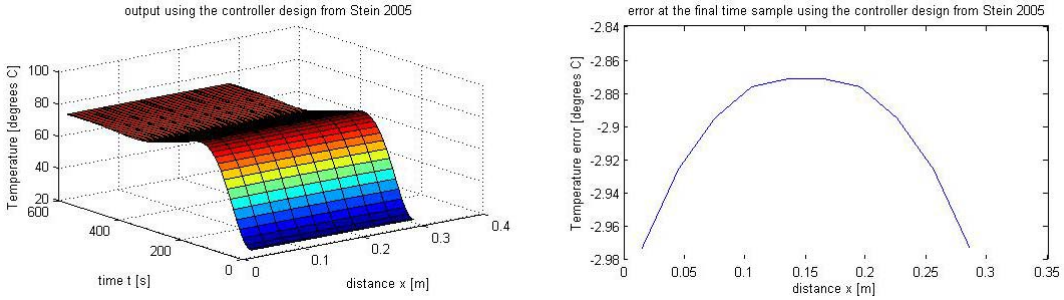

Figure 57 Step response using the controller from a paper by Stein.

A large error is observed across the whole length and time of this simulation. The controller is too slow to track this reference input. Other designs with faster reactions were trialled, but the step response overshoot and settling times were unacceptable.

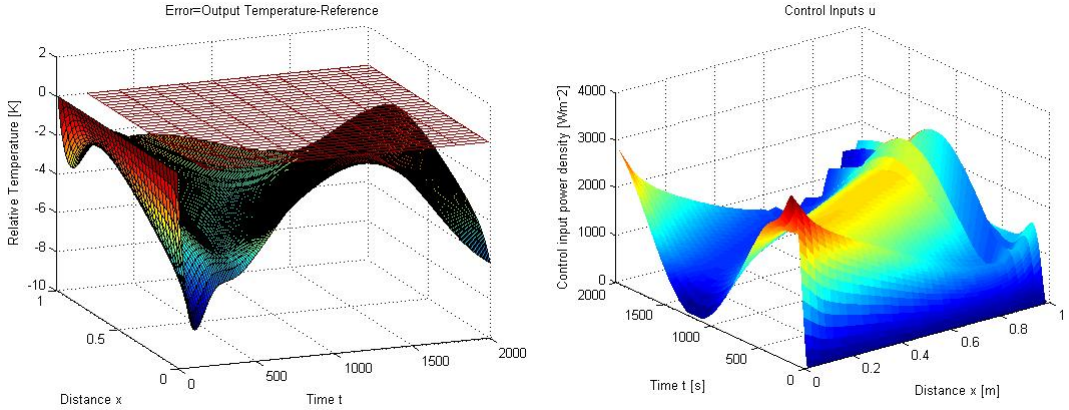

Figure 58 Reference tracking simulation results using the controller from a paper by Stein.

The controller appears to reject disturbance fairly well, although the persistent error has increased. Gain for measurement error, actuator error and input

disturbances are 0.56, 4.20 and 0.09 respectively. The actuator error gain is concerning, although visually the effects do not appear as prominent as this would indicate.

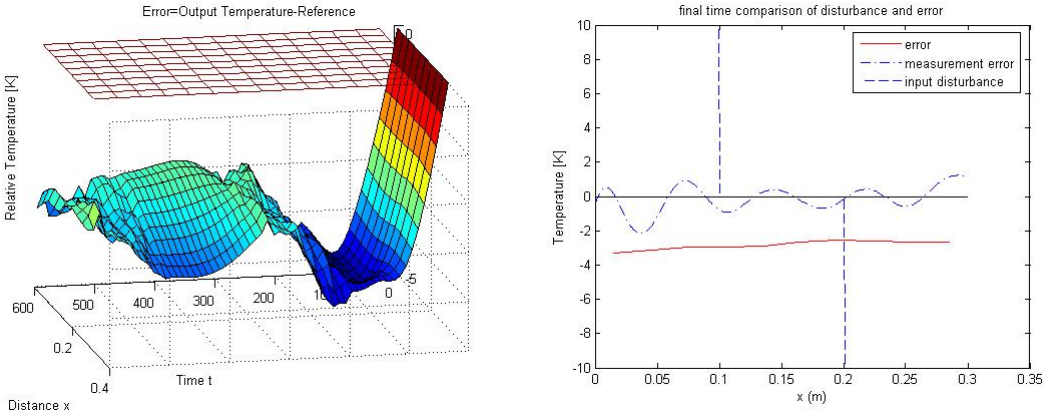

Figure 59 Disturbance rejection performance using a controller from a paper by Stein.

Synthesising this controller was a complicated process involving sampled spatial frequency domains, linear programming and robust control theory. It appears that the performance of the controller is somewhat disappointing and does not justify the effort invested.

5.3 Control Design from Gorinevsky 2008 – Smoothing

The design in this paper follows similar lines to Stein's method (Gorinevsky and Stein are collaborators, the reference here is to the first author for each paper). Both propositions use a temporal PI controller with an allowance for spatial influence and optimise the controllers using linear programming. Once again, the design is intended to be relatively simplified, foregoing many of the complications associated with 'modern control'. The emphasis is on a procedure which can be followed by most engineers and will perform to specification in most scenarios.

The controller uses a temporal PI controller with a spatial operator K to "compensate for spatial interaction effects in the plant". To address problems with stability at certain spatial frequencies a smoothing operator S is also included. The resulting control law is

$$u = z^{-1}.u - z^{-1}c(z^{-1}).K(V)(output - ref) - z^{-1}S(V)u \qquad (19.)$$

To synthesise the controller, the author proceeded as follows (several steps will be recalled from the previous method):

1. Recall the spatial impulse response $h_k(x)$ and the related transfer function $h^*(v)$ from section 2.6.1 for the nominal plant and perturbed models for conditions which give maximum and minimum cooling.

2. From the impulse response, derive the static gain matrix H (identical to K_s in the paper by Lunze). Calculate the maximum singular value $\bar{\sigma}$ (H).
3. Calculate the multiplicative error for the spatial impulse response as a function of *x*. The maximum error is denoted Δ.
4. Evaluate the transfer function *h*(v)* for *v* gridded from 0-π on the unit circle. This gives the response of the system to input with spatial frequencies from constant to infinite frequency.
5. Use the temporal PI controller $c(z^{-1})$ with gains taken from the distributed PID controller (see section 5.1).
6. Approximate $G(z^{-1})$ by a first order transfer function, as explained in section 4.2.
7. Choose design parameters. Unfortunately the interpretation for some of the parameters is not clear in the paper, so some choices are arbitrary or heuristic. The following values have been used for the final design:

Table 4 Design parameters used for the controller design from a paper by Gorinevsky 2008.

Parameter	Symbol	Value
bound on plant output y	d_o	300°C
bound on plant input u	u_o	maxHeat
bound on temporal steady state error	e_o	0.5
bound on multiplicative error	∂_o	$\Delta\|h\|$
bound on uncertainty in the controller	∂_c	0
bound on uncertainty in the smoothing function	∂_s	0
coefficient of the modal loop gain inside the active region	α_p	100
coefficient of the modal loop gain outside the active region	α_o	1
coefficient of the modal smoothing gain S inside the active region	β_p	100
coefficient of the modal smoothing gain S outside the active region	β_o	1
constant in the linear approximation of the domain inside the active region	γ_p	0.7
constant in the linear approximation of the domain inside the active region	γ_o	0.85

8. The parameters are inputs to a linear program:

Minimise

$$(1 - S(v)^T.x - k_i.G(v).K(v)^T.x)^2$$

Subject to :

Bounded Output

$$-S(v)^T.x - k_i.G(v)^T.K(v)^T.x \le \frac{d_o}{u_o}K(v)^T.x \le S(v)^T.x + k_i.G(v)^T.K(v)^T.x$$

Bounded Steady State Error (*only for* $v <$ *Spatial Bandwidth*)

$$-S(v)^T.x - k_i.G(v)^T.K(v)^T.x \le \frac{d_o}{e_o}S(v)^T.x \le S(v)^T.x + k_i.G(v)^T.K(v)^T.x$$

Steady State Robustness

$$-S(v)^T.x - k_i.G(v)^T.K(v)^T.x \le \delta_o.K(v)^T.x \le S(v)^T.x + k_i.G(v)^T.K(v)^T.x$$

Robustness to controller uncerta int *y*

$$-S(v)^T.x - k_i.G(v)^T.K(v)^T.x \le \delta_c.K(v)^T.x \le S(v)^T.x + k_i.G(v)^T.K(v)^T.x$$

Robustness to smoothing function uncerta int *y*

$$-S(v)^T.x - k_i.G(v)^T.K(v)^T.x \le \delta_s.K(v)^T.x \le S(v)^T.x + k_i.G(v)^T.K(v)^T.x$$

Linear approximation to the domain D

$$\alpha_p k_i.G(v)^T.x + \beta_p S(v)^T.x + \gamma_p \le 0 \quad , \quad v < \text{Spatial Bandwidth} \qquad (20.)$$

$$\alpha_o k_i.G(v)^T.x + \beta_o S(v)^T.x + \gamma_o \le 0 \quad , \quad v > \text{Spatial Bandwidth}$$

where x is a vector of the filter coefficients $[\kappa_{o,...}\kappa_N,\sigma_{o,...}\ \sigma_N]$. $K(v)$ and $S(v)$ are the frequency dependent gains of each filter, but for this design they were taken as constants.

9. Take the inverse fourier transform of the filter coefficients $K(v)^T[\kappa_{o,...}\kappa_N]$ and $S(v)^T[\sigma_{o,...}\ \sigma_N]$ and truncate to the desired length – in this case only one nearest neighbour on either side is considered.

The step response shows a rather large overshoot and ongoing oscillations. The small final time error is merely lucky timing.

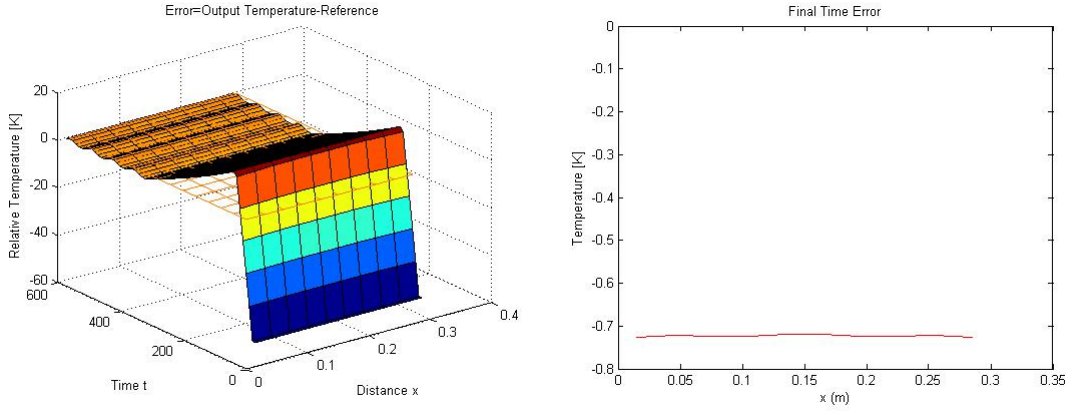

Figure 60 Step response using the controller design from Gorinevsky 2008.

In contrast, the reference tracking performance is very good. The oscillations in the output are fast but small enough to be acceptable and the average error is small for the whole simulation.

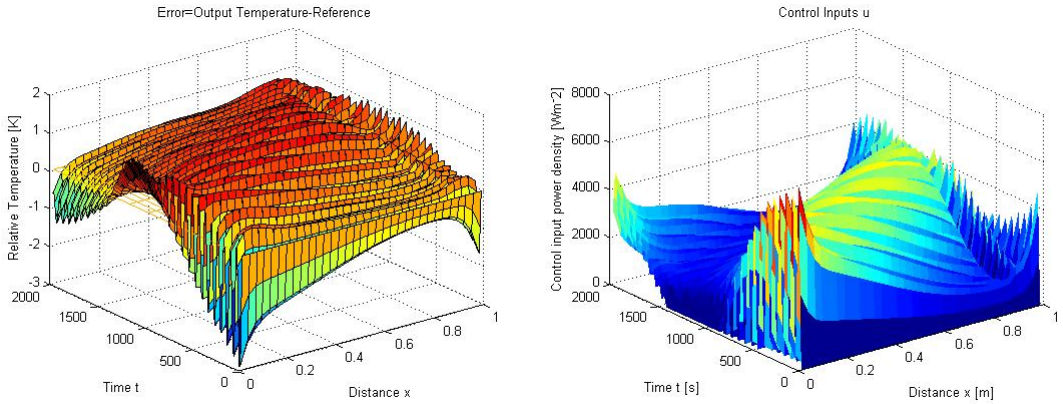

Figure 61 Reference tracking performance simulation using the controller design from Gorinevsky 2008.

Given that the controller response will oscillate even for the nominal plant, there were no high expectations for the disturbance rejection performance. Surprisingly, only uncertainty in the actuators leads to a bad response – the gain for this disturbance is 3.49. For measurement error and input disturbance the gains are 0.67 and 0.07 respectively.

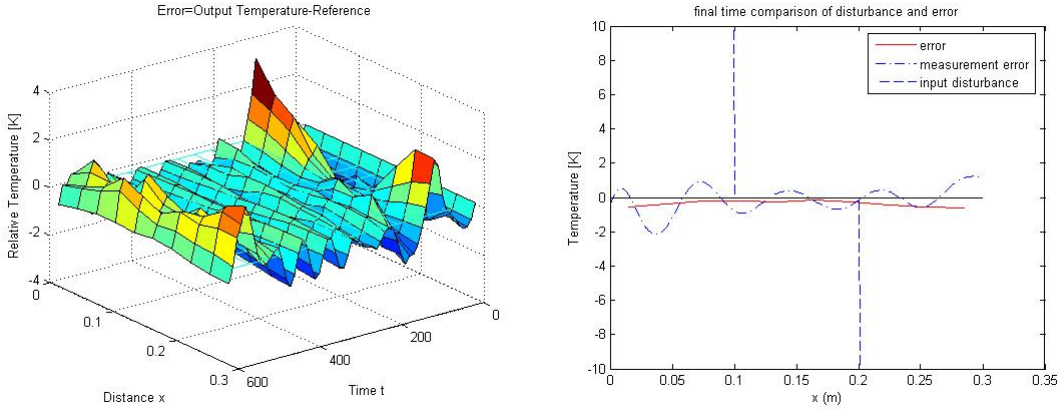

Figure 62 A simulation to test the disturbance rejection performance of the controller from Gorinevsky 2008.

It was found that attempts to improve the tuning of the controller often had no effect or led to instability. In this sense, it has been difficult to reach a good design, although more time to reach a better understanding of the design parameters may help to improve it.

5.4 Control Design from Gorinevsky 2006 – IIR

Similar to the FIR controller designed earlier, it is equally possible to use an IIR filter to take account of the spatial influence of the plant. IIR filters have an improved response compared to FIR filters of the same order, but they also present several difficulties when implemented on a non-causal dimension. Stable non-causal functions with infinite influence must decay in all directions. For implementation, it is typical to divide each non-causal dimension into causal and anti-causal components, where the causal poles lie inside the unit circle and the anti-causal poles must be outside. This paper refers to another iterative method, where the IIR function

$$\sum_{m=-M}^{M} b_m z^{-m} y = \sum_{l=-L}^{L} a_l z^{-l} x \tag{21.}$$

is approximated by

$$y_{n+1} = y_n - B(z) y_n + A(z) x_n \tag{22.}$$

n is the iteration number, and the complete array is evaluated at each iteration. The paper claims that a scaling factor can be found such that the iterations will converge, but this has not been the case for this author, and so the simulation using an IIR controller hangs on an infinite loop.

The co-efficients for A and B are designed using a linear program. First they must be split into a frequency dependent component *c* and constants *p*. For B,

$$\begin{aligned} &B\left(e^{i\omega}\right) = c_b{}^T(\omega) p_b \\ &c_b(\omega) = \begin{bmatrix} 1 & 2\cos\omega & \dots & 2\cos M\omega \end{bmatrix}^T \\ &p_b = \begin{bmatrix} b_0 & b_1 & \dots & b_M \end{bmatrix}^T \end{aligned} \tag{23.}$$

then the problem can be solved for all frequencies simultaneously.

$$\begin{aligned} &\textit{Minimise} \\ &\left(c_b{}^T . p_b - 1\right)^2 \\ &\textit{Subject to}: \\ &\textit{Passband Gain} \\ &\left(1 - r_p\right) c_b{}^T . p_b \le c_a{}^T . p_a \le (1 + r_p) c_b{}^T . p_b \\ &\textit{Stopband Attenuation} \\ &- r_s c_b{}^T p_b \le c_a{}^T . p_a \le r_s c_b{}^T . p_b \end{aligned} \tag{24.}$$

Using a three-tap filter and r_p=0.5, r_s=0.25 gives

p_a=[-0.1729x10^{-16} 0.2798x10^{-16}]
p_b=[-0.3458x10^{-16} 0.5595x10^{-16}]

Time constraints have prevented the design from being taken any further and resolving the convergence issue. This is particularly unfortunate in light of the

success of the FIR based controller, as the IIR performance should be similar and better, however the implementation is difficult due to the non-causal nature of the system.

6 Implementing controller functions in C

Decentralised and spatially distributed controllers are to be implemented on the Atmega microcontroller. While most of the functionality of these controllers falls under the design and build part of the project carried out by Vaclav Klems, the control algorithm is a separate function translated from the designs created in Matlab. The functions for the various controllers are interchangeable, having the same input and return types. All the differences between the controllers can be resolved internally, so the function is called as shown:

```
unsigned int u=controller( unsigned int *temperatures,
                           unsigned int *ref)
```

The array of temperatures is taken from sensor readings and the array of reference values are received from a governing controller. This form of polymorphism could even be extended to centralised controllers, where the function would obtain '`u`' by sending and receiving data from a central computer.

The constants T_s, `maxHeat` and the controller parameters K_p, K_i, K_d, and `FIR` are declared inside the function. Equally, they could be saved in memory during initialisation, but the savings in processing time would be small and it would add complexity to the main program. The function needs access to memory to store variables between cycles. For differentiating and integrating, `lastIntermediateVariable` and `integralError` must be stored as global variables. Where `intermediateVariable` is a scaled version of the current error (difference between the measured value and reference). Integrator windup has been avoided by pausing the integrator when the actuators are saturated.

The Atmega162 is an 8-bit microcontroller and does not use floating point numbers. Where better than integer precision is desired, numbers must be converted to a fraction of two integers. For instance FIR coefficients (which are restricted to between zero and one in this case) are multiplied by a factor of 100 when they are implemented and the magnitude is corrected later. Dividing operations reduce the precision since all results are rounded to an integer, therefore these operations are left as late as possible in the function to mitigate compounding of rounding errors.

The code for the functions can be found along with all the Matlab functions on http://dce.felk.cvut.cz/heat

7 Conclusions

7.1 Discussion

It is unfortunate that the rod system has not been completed in time to test the control designs on a real system, however some valuable insights can be gained from the simulation results. Of the advantages offered by spatially distributed control (scalability, parallel computational efficiency, robust to a single point failure) only the reduction in computational demands during the design phase was directly evident during this thesis. Spatially distributed controllers have simplified control laws and relatively few design variables which are then repeated over the whole array. Synthesis of loop-shaping and $\mathcal{H}_\infty$ controllers consumed considerable computing time, and MPC was worse still. Processing time can not be accurately assessed from the spatially distributed simulations since computation which would be shared must be completed by one processor. Hopefully, when the hardware system is completed, this and other advantages will become apparent.

Equally, the disadvantages related to boundary conditions have been ignored in this thesis, and the extra work required to connect so many nodes has been left to Vaclav Klems. The primary disadvantage of spatially distributed control which has been a factor for this thesis is the steep learning curve that applies to most control engineers. Where modern MIMO control methods are well known, an engineer designing a spatially distributed control system must first familiarise themselves with new techniques. Some concepts are not obvious, and much more effort is expended in this phase of a project. This may change if spatially distributed control becomes accepted and standardised, but not for several years.

The most overwhelming impression from the controllers in this project is that the simplest designs have performed best. LQR, PI and the FIR-PID controller have performed best across all three tests. If we return to the comparison of spatially distributed and classical control, the issue of simplicity cuts both ways. nD systems theory is conceptually and mathematically complex, but pays off in computationally efficient designs and design procedures.

No evidence of the theoretical inferiority of decentralised controllers was found in this work.

7.2 Future Work

The first priority for future work is to complete the hardware for the HeatAl project. Once this is in place, the goal should be to improve the model and assess parameters accurately in order to reduce uncertainty.

More simulations should be conducted using the perturbed plants. It is expected that the designs which cater explicitly to model uncertainty (which have generally shown poor response in the tests so far) will outperform the others in these scenarios. Obviously, the controllers should ultimately be trialled on the hardware to confirm the simulation results (or otherwise).

Some of the controllers designed within this project could be further developed. The IIR filter should be completed and the anomalous asymmetries in the $\mathcal{H}_\infty$ and loop shaping controllers should be investigated.

The concept of adapting MPC control to a microcontroller mentioned in [17] has potential for future spatially distributed applications.

7.3 Conclusions

Mathematical models have been derived to describe the thermodynamics of a long thin rod subject to heating by an array of actuators. Different models use finite element methods, partial differential equations, state space and n-D transfer function theory. There is good agreement between all models.

A system identification experiment showed a value for the convection co-efficient which is outside the expected range, however this is thought to be beneficial for the performance of most controllers.

Ten different controllers have been designed, grouped by their control architecture into centralised, decentralised and spatially distributed types. Simulations have tested their step response, reference tracking performance and robustness to various types of disturbance. In all tests and for all types of controllers, simple PI or PID controllers have performed best. The responses of more complicated modern designs were in general too slow.

This is only the beginning of the HeatAl project, and once the hardware is complete the groundwork will hopefully have been laid for future experiments. At CVUT this will be the first spatially distributed system on which researchers can perform real tests. It is hoped that the hardware, models and lessons learnt can be of assistance as they refine techniques and work towards some of the applications mentioned below.

Relevant code for the models and controllers will be placed on the internet at http://dce.felk.cvut.cz/heat

8 References

1 Rendered images provided by Zdenek Hurak.

2 Skelton, R.E.; Iwasaki, T.; Karolos, M.: *A Unified Algebraic Approach to Linear Control Design.* CRC Press 1997.

3 Wall, R.W.: *Agriculture Irrigation System Control and Data Communication for Real-time Variable Water Application*. Retrieved from www.ee.uidaho.edu/ee/digital/rwall/echelon1.pdf

4 Lunze, J.; Abraham, R.: *Modelling and Decentralized Control of a Multizone Crystal Growth Furnace*. International Journal of Robust and Nonlinear Control, Vol.2, 1992

5 Rogers, E. *Long Wall Coal Cutting*. Chapter 1-1-1 of:
Galkowski, K.; Wood, J: *Multidimensional Signals, Circuits and Systems.* CRC Press 2001.

6 Gorinevsky D.; Boyd S.; Stein G.: *Design of Low-bandwidth Spatially Distributed Feedback*. IEEE Transactions on Automatic Control, Vol.53, No.2, 2008

7 Stein, G.; Gorinevsky, D.: *Design of Surface Shape Control for Large Two-Dimensional Arrays*. IEEE Transactions on Control Systems Technology, Vol.13, No.3, May 2005

8 Gorinevsky, G.; Stein, G.: *Structured Uncertainty Analysis of Tobust Stability for Multidimensional Array Systems*. IEEE Transaction on Automatic Control, Vol.48, No.8, 2003

9 D'Andrea, R.; Dullerud, G.: *Distributed control design for spatially interconnected systems*, IEEE Transactions on Automatic Control, Vol.48, No.9 2003

10 Franklin, G.; Powell, J.D.; Emami-Naeini, A.: *Feedback Control of Dynamic Systems*. Prentice Hall 2006.

11 Golem. *Bio-Inspired Assembly of Meso-Scale Components.* Retrieved from http://www.golem-project.eu/

12 MacMartin, D.G.: *Control Challenges for Extremely Large Telescopes*. Smart Structures and Materials 2003: Industrial and Commercial Applications of Smart Structures Technologies.

13 eFunda. *Engineering Fundamentals.* Retrieved from http://www.efunda.com/materials/alloys/alloy_home/show_alloy_found.cfm?ID=AA_1350&prop=all&Page_Title=%20Metal%20Alloys%20Keyword%20Search%20Results

14 Earle, R.L.: *Convection Heat Transfer*. NZIFST 1993 Retrieved from http://www.nzifst.org.nz/unitoperations/httrtheory6.htm

15 Skogestad, S.; Postlethwaite, I.: *Multivariable Feedback Control – Analysis and Design*. John Wiley and Sons Ltd, 2005

16 Emami-Naeini, A et. al., *Modeling and Control of Distributed Thermal Systems*. IEEE Transactions on Control Systems Technology, Vol.11, No.5. 2003

17 Bemporad, A.; Borrelli, F.; Morari, M.: *Model Predictive Control Based on Linear Programming – The Explicit Solution*. IEEE Transactions on Automatic Control, Vol.47, No.12, 2002

18 Gorinevsky, D.; Boyd, S.: *Optimization-Based Design and Iplementation of Multidimensional Zero-Phase IIR Filters*. IEEE Transactions on Circuits and Systems –I: Regular Papers, Vol.53, No.2, 2006