Pacific Institute
for the Mathematical Sciences

http://www.pims.math.ca
pims@pims.math.ca

Proceedings of the third

PIMS Graduate Industrial Math Modelling Camp

Cosponsored by:

The Natural Science and Engineering Research Council of Canada

and

The British Columbia Information, Science and Technology Agency

FOREWORD BY THE PIMS DIRECTOR

The Pacific Institute for the Mathematical Sciences is committed to providing training for young mathematical scientists whether they are pursing careers in academia or in industry.

The **Graduate Industrial Mathematical Modelling Camp (GIMMC)** is one of two components of the annual PIMS Industrial Forum. The other component is the PIMS **Industrial Problem Solving Workshop** which takes place soon after the camp. GIMMC was conceived to give students the opportunity to learn about the modern methods of applied mathematics. It also gives them extensive training and helps prepare them for the Industrial Problem Solving Workshop.

At the workshop students work together in teams, under the supervision of invited mentors. Each mentor poses a problem arising from an industrial or engineering application and guides his or her team of graduate students through a modelling phase to a resolution.

The third GIMMC was held at Simon Fraser University, May 23–27, 2000. Forty-one graduate students came from North America came to SFU to work five mentors from industry. Almost all the students came from 16 universities across Canada, however one came from as far away as New York University. The five industrial mentors, who came from University of Minnesota, University of Southampton, Eastman Kodak, Rensselaer Polytechnic Institute and IBM, provided a wide range of interesting and challenging problems. It is my pleasure to announce that the programme was a huge success.

These proceedings contain the culmination of each teams work and they show how much can be achieved in a week of hard work.

I want to express my appreciation and gratitude to everyone involved in this workshop, in particular I wish to thank the organisers (Keith Promislow, Mary Catherine Kropinski, Sadika Jungic, Lindsay Hughes) and mentors (Rachel Kuske, Colin Please, David Ross, Donald Schwendeman, Brett Stevens). The great success of the first three years of GIMMC shows that we have much to look forward to in the future.

Dr. Nassif Ghoussoub, Director
Pacific Institute for the Mathematical Sciences

PREFACE

As preparations for the fourth Graduate Industrial Mathematical Modeling Camp (GIMMC) at University of Victoria are now well under way, its an appropriate time to reflect upon the success and direction of these workshops. There is no doubt that the GIMMC has grown in size from the initial one held at SFU in 1998. That one attracted 30 applicants and had a shoestring budget. The current edition has applications in the 100s and a permanent place in the PIMS budget. But the real measure of success of the GIMMC has to be the impact on the graduate students who have attended. Unfortunately, no formal records have been kept, but there is abundant anecdotal evidence: Math Pays Off!

Consider two students, Antonio (Tony) Cabal, a graduate student in applied math at University of Western Ontario who attended the 1998 GIMMC, and Tom Janiewicz, an undergrad at Simon Fraser University who attended the 2000 GIMMC. Tony worked on a problem mentored by David Ross (Eastman Kodak) whose goal was to model the of diffusion of surfactants in a thin flowing polymer known as a coating curtain. Such was the impression that Tony made upon David that when a position became available at Kodak later that year, David brought Tony in for an interview. Tony is now employed as a mathematical modeller with Kodak. As a member of the integrated materials and microstructures lab he develops and applies mathematical models of fluid mechanics and MEMS microactuators for ink jet printers. The crux of Tony's work involves the analysis and numerical solution of nonlinear PDEs. In addition, Tony has two patents pending for inventions which have grown out of his mathematical models! As David explains it "Tony is doing very well here, he is very good."

In Tom's case, he dove into the Catalytic Converter problem presented by Don Schwendeman from Rensselaer Polytechnic Institute. This problem is described in Chapter 5 of this Proceedings. A few months after completing the GIMMC, armed with his BS in applied math and the writeup of the Catalytic Converter problem, Tom interviewed at Universal Dynamics, a BC high-tech engineering/software firm. In Tom's words: "When I showed the interviewer the report on the catalytic converter from the workshop, he did not hesitate too long to offer me the job." He now works in the Brainwave group at Universal Dynamics with another programmer and two engineers on the mathematical underpinnings of a software system which controls manufacturing processes. Tom's work uses "control theory very intensely" and he includes Laplace Transforms, z transforms, and singular value decompositions among the mathematical techniques he has applied recently.

Tom has been eager to help establish contacts between Universal Dynamics and the PIMS universities; and perhaps to bring a problem to the GIMMC or the Industrial Problem Solving Workshop in the future. In this way Math and the GIMMC will continue to pay dividends for future students.

Keith Promislow and Mary Catherine Kropinski
Organising Committee
Department of Mathematics and Statistics
Simon Fraser University

Contents

Chapter 1

Catalytic Converter: A Simple Mathematical Solution to Understanding Operation

Participants: Donald Schwendeman (Mentor), Rozita Dara, Tomasz Janiewicz, Margaret Liang, Mohammad Oskoorouchi, Maurice Shevalier, Maikel Sianturi.

PROBLEM STATEMENT: A catalytic converter is used by automobiles for controlling emissions. It takes unburned gases, which can contribute to smog, and "burns" them. The "burning" involves a chemical reaction catalyzed by an inert metal located within the converter. The chemical reactions are temperature dependent and do not occur until the converter reaches a critical temperature.

In this workshop, the processes involved in a catalytic converter are examined, heat transfer, mass transfer, and the chemical reactions. A mathematical model of the converter is developed. The model is then used to simulate the converter, which is similar to the work done by Oh and Cavendishi [3].

1.1   Problem Description

When a car starts, the catalytic converter, which is cold, is exposed to hot gases. The converter is slowly heated in a non-uniform manner, during which time no conversion reaction occurs. When the converter reaches the critical temperature, the reaction occurs. Many of the reactions are exothermic and add to the heating of the converter. The converter then heats the exhaust gas passing over it which then contributes to the heating of the converter not yet at the critical temperature. The concentration of gases undergoing the reaction decreases as more of the converter is heated. Ideally, when the entire converter reaches the critical temperature this concentration will go to zero.

In this model many simplifications are made. One simplification involves the chemical species which undergo the catalytic reaction. In this model only CO is considered to react. The reaction involved is as follows:

$$CO + \frac{1}{2}O_2 \longrightarrow CO_2$$
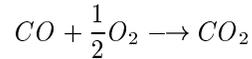
In this model a concentration of O_2 is assumed to be constant. Further the only source of CO_2 is assumed to be from the catalytic reaction, so the concentration of CO_2 is inversely related to CO.

The second major simplification involves the catalytic converter. It consists of many tubules embedded in a ceramic matrix. The ceramic is coated with platinum. Instead of considering thousands of tubules, an average tubule which is one dimensional in length is considered. It has a platinum coating on top and bottom of the ceramic. Through the center of the tubule the exhaust gas flows.

In the succeeding sections, a simple one dimensional model is developed. The model is then non-dimensionalized and solved both analytically and numerically for the gas and ceramic temperature, as well as the CO concentration.

1.2   Methodology of Solution

We derive the equations based on two principles, conservation of mass and conservation of heat. The variables we will attempt to solve for are,

c_g = Concentration of gas in free space.

c_s = Concentration of gas on the surface of the solid.

T_g = Temperature of gas inside the converter.

T_s = Temperature of the solid.

The *Equation for Conservation of Mass in Open Space* is

$$\underbrace{\frac{d}{dt}A\int_a^b c_g dx}_{i} = \underbrace{uAc_g^0}_{ii} - \underbrace{uAc_g^*}_{iii} + \underbrace{Pk\int_a^b (c_s - c_g)dx}_{iv},$$

where

(i) is the rate of change of mass of gas inside control volume $[a, b]$, A is the cross sectional area,

(ii) is the flux of gas at inlet a, u is the velocity of gas,

(iii) is the flux of gas at outlet b,

(iv) is mass transfer of gas to the surface of the solid, here P is the parameter of the area of the open space and k is the mass transfer coefficient.

Terms ii and iii can be written as a single integral, and since a and b are arbitrary we can eliminate the integrals. After some simplifications we end up with the following equation:

$$A\left(\frac{\partial}{\partial t}c_g + u\frac{\partial}{\partial x}c_g\right) = Pk(c_s - c_g). \tag{1.1}$$

After adding up all the contributions from various sources of heat, we will end up with the following equation:

$$\underbrace{\frac{d}{dt}A\int_a^b \bar{c}_g T_g \rho_g dx}_{I} = \underbrace{uA\bar{c}_g T_g^0 \rho_g}_{II} - \underbrace{uA\bar{c}_g T_g^* \rho_g}_{III} + \underbrace{Ph\int_a^b (T_s - T_g)dx}_{IV}.$$

where

(I) is the rate of change of heat inside control volume $[a,b]$, $\bar{c}_g$ is the specific heat of the gas, and ρ_g is the density of gas.

(II) is the heat flux at a, T_g^0 is the temperature of gas at a.

(III) is the heat flux at b, T_g^* is the temperature of gas at b.

and,

(IV) is the heat transfer to the surface of the solid, h is the heat transfer coefficient.

For reasons similar to equation (1.1) this equation simplifies to the following:

$$A\bar{c}_g \rho_g \left(\frac{\partial}{\partial t}T_g + u\frac{\partial}{\partial x}T_g\right) = Ph(T_s - T_g). \tag{1.2}$$

This is the equation for *Conservation of Energy in Open Space.*

The conservation of mass on the surface of solid is obtained by balancing the following two quantities

$$\underbrace{\int_a^b (A+B)\tilde{a}Rdx}_{*} = -\underbrace{\int_a^b \rho_g Pk(c_s - c_g)dx}_{**},$$

where,

* is moles of the gas generated by reaction on the surface from a to b, here $\tilde{a}$ is the area of platinum on the surface of the solid, B is the cross sectional area of the solid, and R is the reaction rate.

** is the mass transfer of gas from the surface.

After eliminating the integrals we end up with the *Equation for Conservation of Mass on Surface*:

$$(A+B)\tilde{a}R + \rho_g Pk(c_s - c_g) = 0. \tag{1.3}$$

The conservation of energy in the solid is represented by the following relationship

$$\underbrace{\frac{d}{dt}B\int_a^b \bar{c}_s \rho_s T_s dx}_{1} = \underbrace{DB\frac{\partial T_s}{\partial x}\Big|_{x=b}}_{2} - \underbrace{DB\frac{\partial T_s}{\partial x}\Big|_{x=a}}_{3} + \underbrace{\int_a^b Ph(T_g - T_s)dx}_{4} + \underbrace{\int_a^b Pqdx}_{5},$$

where

1 is the rate of change of heat inside the solid in the control volume in $[a,b]$, $\bar{c}_s$ is the specific heat of the solid, ρ_s is the density of the solid.

2 is the heat flux in solid at b, here D is the diffusion coefficient.

3 is the heat flux in solid at a.

4 is the heat transfer from solid to gas.

5 is the heat generation due to reaction and q is the heat energy due to reaction.

If we combine term 2 and 3 into a single integral, eliminate the integrals from the above equation we will have the equation for the *Conservation of Energy in the Solid*:

$$B\bar{c}_s \rho_s \frac{\partial T_s}{\partial t} = BD\frac{\partial^2 T_s}{\partial x^2} + Ph(T_g - T_s) + Pq \tag{1.4}$$

If we analyze the exhaust from the engine we can come up with the temperature and concentration of gas. These will give us insight into initial and boundary conditions. Thus T_g and c_g are known at

$x = 0$. Furthermore at any given time T_g is known as is c_g at the inlet. Since at any instance there is no heat flux to the surrounding environment since air is a good insulator, so $\frac{\partial T_s}{\partial x}|_{x=0,L} = 0$. Finally, since the converter is initially at room temperature $T_s|_{t=0}$ is known.

The next step in the analysis of the above four equations is *non-dimensionalization*. After considering four different time scales: the path length time scale the mass transfer time scale, the energy time scale, and the temperature built up time scale, we decided that the last is the most appropriate one for a consideration of the warm up problem. Upon scaling the variables, our equations transform into the following non-dimensional system:

$$u\frac{\partial}{\partial x}c_g = \alpha(c_s - c_g) \tag{1.5}$$

$$u\frac{\partial}{\partial x}T_g = \beta(T_s - T_g) \tag{1.6}$$

$$a\sigma c_s e^{\gamma T_s} = c_g - c_s \tag{1.7}$$

$$\frac{\partial}{\partial t}T_s = \delta\frac{\partial^2}{\partial x^2}T_s + \beta(T_g - T_s) + \alpha a\sigma c_s e^{\gamma T_s}, \tag{1.8}$$

where

$$\alpha = \frac{PkL}{Au_R} \qquad \beta = \frac{PhL}{A\bar{c}_g\rho_g u_R} \qquad \delta = \frac{Dt_R}{L^2\bar{c}_s\rho_s} \qquad \gamma = \frac{E\Delta T}{RT_s^2}.$$

Here L is the length of the converter, u_R is the velocity of the gas, E is the activation energy, and σ is a constant of the reaction of CO with O_2.

The initial and boundary conditions translate as follows:

- Since scaling eliminated time dependence in equations (1.5) and (1.6), we can drop $T_g|_{t=0}$ and $c_g|_{t=0}$.

- $T_g|_{x=0} = 0$.

- $c_g|_{x=0} = 1$.

- $\frac{\partial T_s}{\partial x}|_{x=0,L} = 0$.

- $T_s|_{t=0} = -1$.

1.3 Results

1.3.1 Analytical Solution

There are two stages for this problem. The first stage is the gentle heating of the converter, and the second stage is the reaction of chemical species.

In the heating stage, the temperature of the solid is almost independent of the location, so $\frac{\partial^2 T_s}{\partial x^2}$ is small. since δ is small, we can ignore the second derivative term in equation (1.8). Since γ is big and $T_s = -1$ initially, $e^{\gamma T_s}$ is small, so we can also cross out the exponential term in equation (1.8). Solving the modified equation (1.5)- (1.8), we get $c_g = c_s$ and they both decrease slowly. The temperature of the gas decreases as it moves down the converter and heats up the solid. As a consequence the temperature of the solid increases slowly.

When the solid reaches its critical temperature, chemical reaction starts, and we reach the second stage. In this stage, $T_s > 0$, so we can not ignore the exponential term in equation (1.8). Inside the converter, the temperature does not change much before and after the reaction, so $\frac{\partial T_s}{\partial t} = 0$, and $\frac{\partial^2 T_s}{\partial x^2} = 0$.

Equation (1.8) becomes:

$$\beta(T_g - T_s) + \alpha a \sigma C_s e^{\gamma T_s} = 0.$$

Solving equations (1.5)- (1.8) gives:

$$c_g = \frac{\beta(1 - T_s)}{\beta - \frac{\sigma \alpha e^{\gamma T_s}}{1 + \sigma e^{\gamma T_s}}}.$$

Plot of T_s versus c_g is shown below:
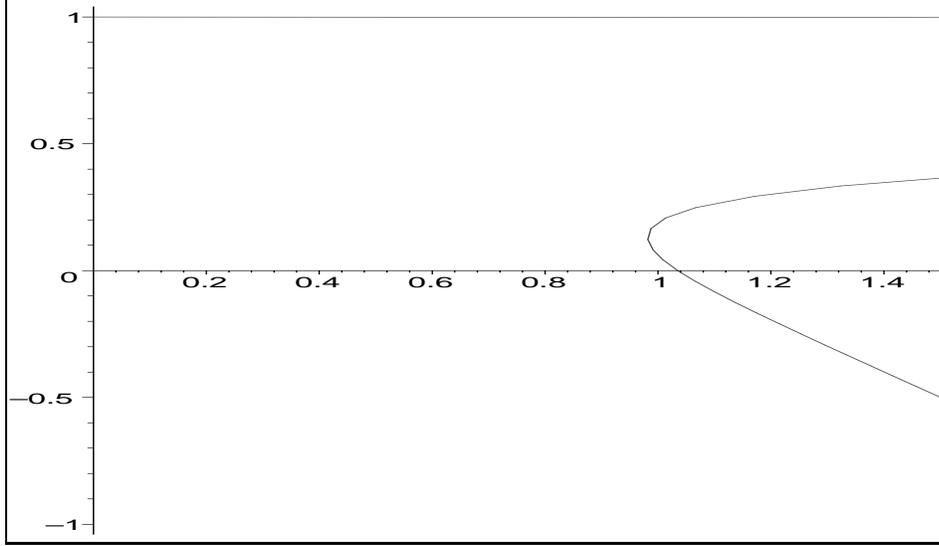
Figure 1.1: Temperature of the converter vs. concentration of pollutant

Equations (1.5) with (1.7) show that c_g is a decreasing function of location x. c_g equals 1 at x=0, and we must follow the curve in figure above with decreasing c_g. At the spot where the curve turns, the path must jump to the branch of the curve that is hot, where $T_s \approx 1$, which is also shown in the numerical analysis below.

1.3.2 Numerical Solution

The equations used in the numerical solution are equations (1.5)- (1.8). The solution obtained was the simplest and most straight forward. Equations (1.5) and (1.6) are ordinary differential equations but they have a c_s and T_s dependents. The c_s dependents of equation (1.5) is removed by substituting equation (1.7) and assuming T_s is known. Solving for c_g yields the following:

$$c_g = exp\left(\int_0^x \frac{\alpha}{u}\left[\frac{-a\sigma e^{\gamma T_s}}{1 + a\sigma e^{\gamma T_s}}\right] dx\right). \tag{1.9}$$

T_s is assumed to be known along discrete points on x. This equation is solved using the trapezoidal rule.
Equation (1.6) is also an ordinary differential equation with the following form:

$$T_g' + \frac{\beta}{u}T_g = \frac{\beta}{u}T_s.$$

The solution has a homogeneous and particular part which can be written as

$$T_g = e^{\frac{-\beta}{u}x}\int_0^x \frac{\beta}{u}T_s e^{\frac{\beta}{u}s}ds. \tag{1.10}$$

This equation is also solved using the trapezoidal rule.

Equation (1.8) is a partial differential equation. The first term on the right hand side is replaced by the central difference formula. This then converts it to an ordinary differential equation which is solved using a modified Euler Method.

The technique to solve the system of equations is as follows:

1. Equation (1.9) is solved for c_g using an initial value of T_s.

2. Equation (1.10) is solved for T_g using an initial value of T_s.

3. The new T_s is solved for using the modified Euler's Method.

This new T_s is then substituted into step 1 and the loop is repeated until $T_s|_{x=0} > 0.80$. The results of the simulations are shown in Figures 1.2, 1.3, and 1.4.
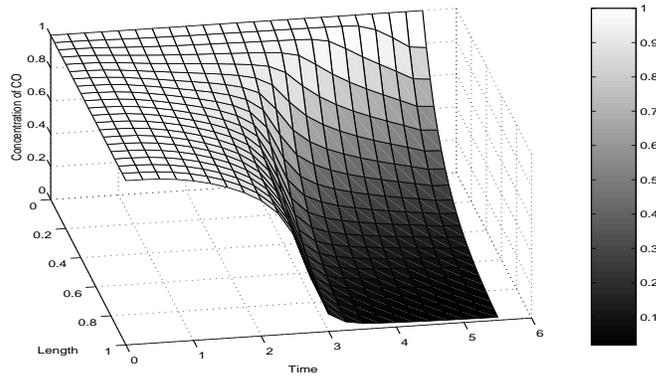
Figure 1.2: Temperature of the gas vs. the converter length

Figure (1.2) shows the temperature of the gas over the converter length as a function of time. Initially at $t = 0$ and $x = 0$ the gas has the condition of $T_g = 0$. Over the length of the converter the temperature drops to -1 at $x = 1$. This is expected since the heat from the gas is absorbed by the converter. As time progresses the temperature of the gas does not drop as much over the length of the converter. There is a point where the temperature of the gas increases. This is due to the conversion reaction occurring and the gas absorbs heat from the converter. As time continuous to progress the temperature of the gas increases until it is in equilibrium with the temperature of the converter.

Figure (1.3) shows the temperature of the converter over its length as a function of time. Initially at $t = 0$ the entire converter is at its initial temperature of -1. As time progresses the temperature of the converter increases due to heat absorption from the gas. There is a point in time where the temperature of the converter is greater than that of the gas. This is due to the onset of the conversion reaction. From this point in time on the temperature of the converter raises sharply due to more conversion reaction occurring. This temperature front travels towards the inlet located at $x = 0$ due to heat diffusion within the converter.

Figure (1.4) shows the concentration of CO over the length of the converter as a function of time. Initially the concentration of CO does not decrease, because the temperature of the converter is lower than the critical temperature. As time progresses the concentration of CO decreases as a function of length due consumption by the chemical reaction. At the end of the simulation the concentration of CO is 1 at the inlet and decreases to 0 at the outlet. This indicates that the converter is at optimum temperature resulting in the optimum conversion reaction occurring.
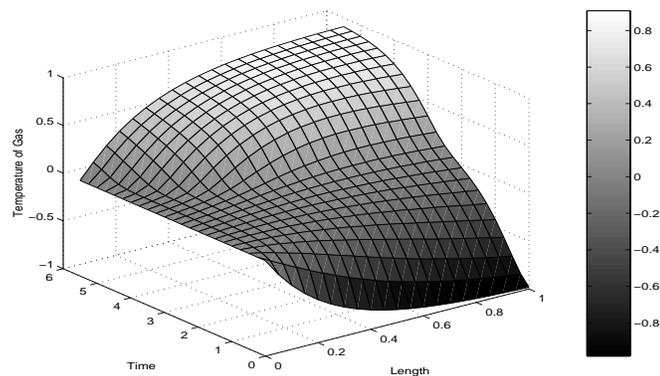
Figure 1.3: Temperature of the converter vs. its length
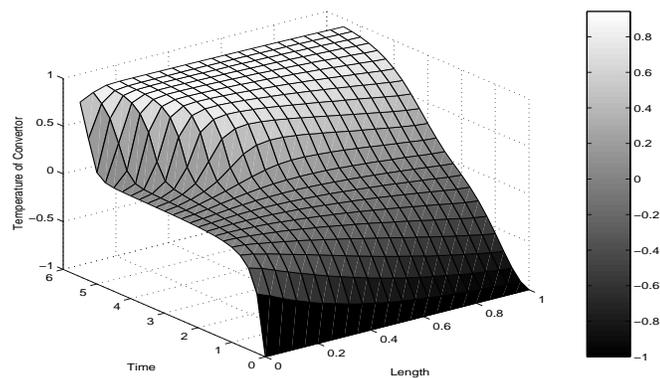
Figure 1.4: Concentration of CO vs. length of the converter

Chapter 2

Queue Compatible Gray Codes and Applications

Participants: Brett Stevens (Mentor), Paul Buskell, Paule Ecimovic, Cristian Ivanescu, Anamaria Savu, Abid Malik, Tzvetalin Vassilev, Boting Yang, Zhiduo Zhao.

PROBLEM STATEMENT: Our group treated the following aspects of Gray Codes: k-subset of an n set and the problem of the shortest circular covering n-word. In the process of investigating these problems, we encountered several interesting applications of Gray codes, some of which will be described below.

2.1 All Subset Covering Words

2.1.1 Formal Problem Statement

Problem: What is the shortest circular word on n-letters $\{1, 2, 3, \ldots n\}$ such that every subset of $\{1, 2, 3, \ldots n\}$ appears at least once as a sub-word (k consecutive letters for a size k subset).

$$f(n) \leq \sum_{k=0}^{n} k \binom{n}{k} = n2^{n-1}$$

For the above upper bound, it is loose according to the numerical results we have some approaches in improving that: e.g.

let $\sum_{k=0}^{n} k \binom{n}{k}$=S;

$\sum_{k=0}^{n} k \binom{n}{k+1}$=SS;

$\sum_{k=0}^{n-1} k \binom{n}{k}$=SA;

$\sum_{k=0}^{n-2} k \binom{n}{k}$=SB;

We have some results for these functions as below:

f(1)=1	f(2)=2	f(3)=3	f(4)=8	f(5)=13
S(1)=1	S(2)=4	S(3)=12	S(4)=32	S(5)=80
SS(1)=0	SS(2)=1	SS(3)=5	SS(4)=17	SS(5)=49
SA(1)=0	SA(2)=2	SA(3)=9	SA(4)=28	SA(5)=75
SB(1)=0	SB(2)=0	SB(3)=3	SB(4)=16	SB(5)=55

From the table, we get better upper bound function SB. That surpasses function S.

The function f(n) defines the length of the corresponding gray codes. e.g:

n	Gray code
1	1
2	12
3	1231
4	24123413
5	1234531425345

We have the known lower bound $\binom{n}{\lfloor \frac{n}{2} \rfloor}$.

2.1.2 Triangulational Gray Codes

In this section, we describe a kind of Gray code which is motivated by triangulations. Let S be a finite set of points in the Euclidean plane. A *triangulation* of S is a maximal straight-line plane graph whose vertices are the points of S. By maximality, each face is a triangle except for the exterior face, which is the complement of the convex hull of S. Without loss of generality, we can assume that all the points are

in the general position, that is, no three points are collinear. For each edge whose endpoints are in S, we can assign a number to it, for example, $1, 2, 3, \cdots, n(n-1)/2$. Let $\overline{n} = n(n-1)/2$. Each triangulation can be represented by a word of bits (i.e., numbers) $X^i = (x_1^{(i)}, x_2^{(i)}, \cdots, x_m^{(i)})$, where m is the number of edges in the triangulation. m is constant for each triangulation when S is given. Thus, we have a code word set $\mathcal{X} = \{X^1, X^2, \cdots, X^N\}$ where N is the number of all the triangulations of S. Let $T(S)$ be a triangulation of S, then an edge e of $T(S)$ is *flippable* if it is adjacent to two triangles whose union is a convex quadrilateral. So, the *flip* of e is an operation of removing e from $T(S)$ and replacing it by the other diagonal of the convex quadrilateral. In this way, we get a new triangulation $T'(S)$, and we say that $T'(S)$ is a flip of $T(S)$. It is well known that for any two X^i and X^j in $\mathcal{X}$ there exists a series of flips to transform X^i to X^j, say, $X^i \to X^{i_1} \to \cdots \to X^{i_j} \to X^j$. Thus, $(X^i, X^{i_1}, \cdots, X^{i_j}, X^j)$ is a kind of Gray codes.

The *triangulational Gray code* is a sequence $X^{i_0}, X^{i_1}, \cdots, X^{i_k}$ of distinct m-bit $\overline{n}$-ary code words such that adjacent words differ in exactly one bit (regardless of the position). The code may be defined by giving X^{i_0} and the transition sequence $T = (t_0, t_1, \cdots, t_{i_k-1})$, where t_l is the labelling number of the edge in which the code words X^{i_l} and $X^{i_{l+1}}$ differ.

An important problem in the triangulational Gray code is to compute the shortest distance between any two code words, where the shortest distance means the smallest number of flips needed to transform one word to the other. For this problem, we investigate a special case, that is, where S is a convex point set. We use the greedy algorithm to attack this problem.

The *one-vertex emission* triangulation is a triangulation each interior edge of which has the same endpoint. For each word in $\mathcal{X}$, there exists a triangulational Gray code such that this code can be transformed to a one-vertex emission triangulation. So, we have the following algorithm.

Algorithm (greedy strategy).
Step 0. Given two words X^i and X^j in $\mathcal{X}$.
Step 1. Select one vertex v in X^i which has the maximum degree.
Step 2. Computer the triangulational Gray code which transforms X^i to the v-vertex emission triangulation.
Step 3. Compute the triangulational Gray code which transforms the v-vertex emission triangulation to X^j.

We have observed that the approximation ratio of Algorithm is less than 2.

Triangulation in three dimensions is more complicated than that in two dimensions. A *3D triangulation* is a partition of the input domain, point set or polyhedron, into a collection of tetrahedra, that meet only at shared faces (vertices, edges, or triangles).

In three dimensional Euclidean space, a strictly convex hexahedron formed from five vertices can be triangulated in two ways: either as a pair of tetrahedra separated by a face, or as three tetrahedra surrounding an interior diagonal. A *3D flip* is one in which two (three) adjacent tetrahedra of the 3D triangulation form a strictly convex hexahedron, then one replace the tetrahedra by the other possible 3D triangulation of the hexahedron containing three (two) tetrahedra. The flip can be considered to be a face "flip", where one interior face is "flipped" for three interior faces or vice versa.

Similarly, we can define the 3D triangulational Gray code. The *3D triangulational Gray code* is a sequence $X^{i_0}, X^{i_1}, \cdots, X^{i_k}$ of distinct m'-bit $\overline{n'}$-ary code words such that adjacent words differ in exactly one bit (regardless of the position), where $\overline{n'} = n(n-1)(n-2)/6$ and m' can vary in words. So the length of the words may differ. However, if the length of the adjacent words is different, then the difference of the length is 1.

An important problem on the 3D triangulational Gray code is whether for two words in 3D, whether there must exist a 3D triangulational Gray code to connect them. For this problem, we just consider whether there exists a 3D triangulational Gray code which can transform a two-emission 3D triangulation to a one-vertex emission triangulation, Unfortunately, even so special case, we cannot obtain significant result. We are going to continue this work.

2.2 Universal cycles for k-subsets of an n-set

A universal cycle for k-subsets of $\{1, \ldots, n\}$ is a cyclic sequence of $\left(\begin{array}{c} n \\ k \end{array} \right)$ integers with the property that all subsets of $\{1, \ldots, n\}$ of size k appear exactly once consecutively in the sequence. As an example the word

$$1\,2\,3$$

contains $\{1\,2\}$ $\{2\,3\}$ and $\{3\,1\}$ only once i.e. all 2-subsets of $\{1, 2, 3\}$.

Problem: Given n and k is there any universal cycle and if there is how can we find it within a reasonable amount of time?

A necessary condition for the existence of the word is:

$$k \text{ divides } \left(\begin{array}{c} n-1 \\ k-1 \end{array} \right)$$

There are some trivial cases:

- $k = 1$ the universal cycle is $1\,2\,3 \ldots n$

- $k = n$ the universal cycle is $1\,2\,3 \ldots n$

- $k = n - 1$ the universal cycle is $1\,2\,3 \ldots n$

A nontrivial case is $k = n - 2$.
Result: We have established that is impossible to construct a universal cycle in this case even when the necessary condition is satisfied (i.e. n is odd)

Next we describe the ideas which led us to this result.
Assuming that such a word exists then the following must happen:

- Somewhere in the word, there is a length n subword that is the n-set $1\,2\,3 \ldots n$
  Proof: We can assume wlog that the universal cycle contains:

$$1\,2\,3 \ldots n-2\,x\,y\,z$$

  Since $2\,3 \ldots n-2\,x$ is a $(n-2)$-subset x must be $1, n$ or $n-1$. If x is 1 then the subset $1\,2\,3 \ldots n-2$ is repeated. So we may assume wlog $x = n - 1$. The allowed values for y are 1 or n. If y is n we are done. Otherwise $y = 1$ and the next position z is 2 or n. Continuing in this way if n does not appear we get a contradiction: $1\,2\,3 \ldots n-2$ appears twice. So n has to appear which implies that the cycle has to contain $1\,2\,3 \ldots n$

- If $m \leq k$ then a m-subset can appear at most:

$$\frac{1}{n-m+1-k} \left(\begin{array}{c} n-1 \\ k-1 \end{array} \right) \text{ times}$$

- We already know that the pattern $1\,2\,3 \ldots n$ appears somewhere inside of the cycle. If we prove that $2\,3 \ldots n-1$ are forced after $1\,2\,3 \ldots n$

position :	1	2	...	n	$n+1$	$n+2$	...	$2n-2$
number :	1	2	...	n	2	3	...	$n-1$

this will contradict the fact the $n-2$-subset $2\ 3 \ldots n-1$ appears just once and so the word can not exist.

Proof: We look for the numbers which can appear on positions $n+1\ n+2\ n+3 \ldots 2n-2$.

$$
\begin{array}{llllllllll}
position: & 1 & 2 & \ldots & n & n+1 & n+2 & \ldots & 2n-2 \\
number: & 1 & 2 & \ldots & n & ? & ? & \ldots & ?
\end{array}
$$

The first position in which 2 can appear is $n+1$, the first position on which 3 can appear is $n+2$,..., the first position in which $n-1$ can appear is $2n-2$.

So $n-1$ is forced to be in position $2n-2$ because otherwise $1\ 2\ 3 \ldots n-2$ are in positions $n+1\ n+2 \ldots 2n-2$, not necessarily in this order and so a subset is repeated. Also $n-2$ is forced to be in front of $n-1$ and so on. This is proved by the induction which follows.

Suppose that for a certain $i \geq 1$ we have the pattern, so n-i-1 is not in this position.

$$
\begin{array}{llllllllllllllll}
pos: & -i & . & -1 & 1 & . & n-i-2 & . & n-i+1 & . & n & n+1 & . & 2n-i-2 & 2n-i-1 & . & 2n-2 \\
num: & ? & . & ? & 1 & . & n-i-2 & . & n-i+1 & . & n & 1 & . & n-i-2 & n-i & . & n-1
\end{array}
$$

In the $-i, \ldots, -1$ position can be any i-subsequence of $\{n-i, \ldots, n-1, n, \}$. But any i-subsequence of $\{n-i, \ldots, n-1, n, \}$ can be joined with $\{1, 2, \ldots, n-i-2\}$ and this gives a $n-2$ subsequence which appears twice.

This relies on the fact that the case $k = n-1$ is trivial and an universal cycle is $1\ 2\ \ldots\ n$. To see this remove $\{1, \ldots, n-i-2\}$ from the pattern and get

$$n-i+1 \ldots n\ n-i \ldots n-1$$

which is acceptable for $k = n-i$ and the set $\{n-i, \ldots, n-1\ n\}$ q.e.d.

Once we know that a universal cycle does not exits we may ask what is the largest word which does not contain a $n-2$-subset twice. In general this word has the length $2n-3$ and is:

$$1\ 2\ 3 \ldots n\ 1\ 2\ 3\ \ldots n-3\ n$$

Another nontrivial case is $n-k=3$

The necessary condition for the existence of the word is $n = 1, 2 \pmod 3$. Hence when n is multiple of 3 the cycle does not exist. For the other values of n we could not prove or disprove that a universal cycle exists in general except the values shown below. The computer search shows some cases when a cycle does exist:

- $k = 4\ n = 7$
  The cycle which contains all the 5-subsets of $\{1, 2, 3, 4, 5, 6, 7\}$ is:
  1 2 3 4 5 1 2 3 6 4 1 2 7 5 3 1 6 7 4 2 5 6 3 7 4 1 5 6 2 7 3 4 5 6 7 1

- $k = 5\ n = 8$
  The cycle which contains all the 5-subsets of $\{1, 2, 3, 4, 5, 6, 7, 8\}$ is: 1 2 3 4 5 6 1 2 3 4 7 5 1 2 3 8 6 4 1 7 3 8 5 4 2 7 6 3 8 1 5 4 7 6 8 2 5 3 7 6 1 2 5 8 4 6 1 2 7 8 4 3 6 5 7 8

- $k = 7\ n = 10$ (Example found by Brad Jackson) The cycle with 5 fold symmetry which contains all the 7-subsets of $\{1, \ldots, 10\}$ is:
  2 3 4 5 6 8 1 2 3 5 6 7 10 2 5 8 9 3 7 10 4 5 9 2 3 4 7 8 10 2 3 5 7 1 2 6 10 4 7 10 1 2 4 ...

General results that we found are the following:

- Somewhere in word, there is a length $n-1$ word that is the $n-1$-set $1\ 2\ 3 \ldots n-1$
  Proof: is similar to that in case $n-k=2$

- A $n-4$-subset will appear at most twice inside of the cycle

Using the results obtained for $n-k = 1, 2, 3$ we can not say what happens for general n and k.

2.2.1 Example

Consider the circular word

$$1234531425345$$

Note that it contains as subwords (respecting the circular nature) all subsets of the 5-element set {1,2,3,4,5}, called the alphabet of the word.

Here is how I arrived at this word. First, I looked at the shortest words containing all subsets of a one, two, three, and four-element alphabet. These are, respectively:

```
1
1  2
1  2  3
1  2  3  4  1  3  2  4
```

Starting with the above word 1, write each word as a row in an array, followed underneath by the same word with everything shifted to the left by one space (remember that each is a circular word). Continue until as many rows as there are letters in each word have been added. The resulting arrays are:

```
1      1  2      1  2  3      1  2  3  4  1  3  2  4
       2  1      2  3  1      2  3  4  1  3  2  4  1
                 3  1  2      3  4  1  3  2  4  1  2
                             4  1  3  2  4  1  2  3
                             1  3  2  4  1  2  3  4
                             3  2  4  1  2  3  4  1
                             2  4  1  2  3  4  1  3
                             4  1  2  3  4  1  3  2
```

Note that each main anti-diagonal consists entirely of the highest letter in each alphabet, and the general "striped" appearance of the anti-diagonals.

Looking at successive pairs of arrays, we see that the larger array contains a portion of the smaller one, which is boxed below:

```
1      1  2      1  2  3      1  2  3  4  1  3  2  4
       2  1      2  3  1      2  3  4  1  3  2  4  1
                 3  1  2      3  4  1  3  2  4  1  2
                             4  1  3  2  4  1  2  3
                             1  3  2  4  1  2  3  4
                             3  2  4  1  2  3  4  1
                             2  4  1  2  3  4  1  3
                             4  1  2  3  4  1  3  2
```

This led me to see if an extension of the last array could produce a larger array containing a word on a five-letter alphabet with the desired property.

To understand my motivation, note that each of the "staircase" structures below:

```
1        1 2      1 2 3      1 2 3 4 1 3 2 4
         2 [1]    2 3 [1]    2 3 4 1 3 2 4 1
                  3 [1 2]    3 4 1 3 2 4 1 2
                             4 1 3 2 4 1 2 3
                             1 3 2 4 1 2 3 4
                             3 2 4 1 2 3 4 [1]
                             2 4 1 2 3 4 [1] 3
                             4 1 2 3 4 [1 3 2]
```
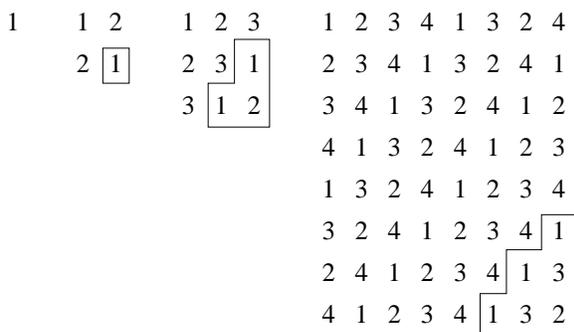
are simply formed by k-cycles on each k-alphabet acting on the original boxed areas. This led me to look for a 4-cycle that I could put in the following space in a larger array:

```
1 2 3 4 5          5
2 3 4 5          5
3 4 5          5
4 5          5
5          5
. . . . . . . . .
. . . . . . . . .
. . . . . . . . .
```
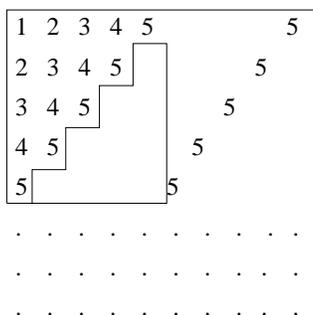
The striping phenomenon would then allow us to retrieve the word from the top row of the array.
I soon hit upon the following:

```
1 2 3 4 5 3 1 4 2 5
2 3 4 5 3 1 4 2 5
3 4 5 3 1 4 2 5
4 5 3 1 4 2 5
5 3 1 4 2 5
. . . . . . . . .
. . . . . . . . .
. . . . . . . . .
```
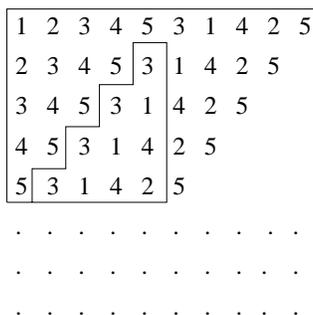
The top row as it sits only contains all of the 2-subsets of the 5-alphabet. It was found to lack the 3-subsets 1,4,5 and 2,3,5; adding 345 to the end gave the word shown at the beginning, which does contain all subsets of the 5-alphabet as circular subwords. In the course of our investigation, we computed that the lower bound for such a word turned out to be thirteen characters, rather than fifteen. Analogous work on finding a word on six letters containing all subsets of 6 letters continues.

2.2.2 Attempts

Initially, the problem of generating covering n-words and their substrings attracted our attention from the point of view of the following data structure motivated by binary gray coding of a given alphabet. Let $\{a_1, a_2, \ldots, a_m\}$ be an m-alphabet. Then, the following data structure will store all the k-words, where $k \in \{1, 2, \ldots, n\}$:

$$P^{ij}_{len(w_i)!}$$

$$P^{ij}_3$$

$$P^{ij}_2$$

$$P^{ij}_1$$ m-alphabet

$a_1\ a_2 .....\ a_m$

W
O w_1 c_{11}
R w_2 c_{ij}
D
S w_{2^m} $c_{2^m m}$
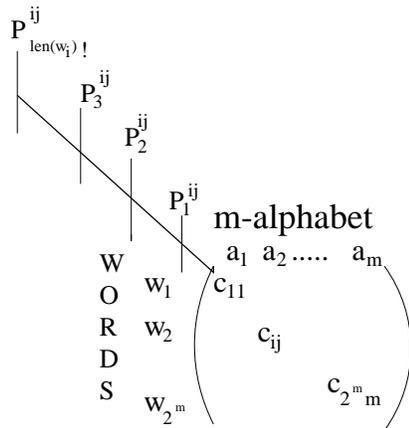
This matrix is three dimensional and contains the following information. The rows represent an ordering of the given m-alphabet, from the first character to the last. The rows represent k-words formed from characters of the given alphabet, as follows. The elements of this matrix are binary digits, with 1 in the column and row of the matrix if the given word contains the given character in the given row. Thus each row represents a word with a given sequence of the characters of the alphabet represented as a bit string. Along the third dimension are all the $k!$ permutations of a word of length k over the m-alphabet. Encoding the permutations of a given brought us to seek a binary Gray code for permutations. We considered canonical decomposition of permutations into transposition sequences unique to within a permutation of the natural ordering of the neutral element of the permutation group. Assuming each transposition could be applied at most once and in a unique sequence, we considered a least change ordering of binary transposition sequence codes representing a given permutation. This would be the binary Gray code representing a given permutation.

Our purpose in retaining a binary coding scheme in the entire data structure was to enable a boolean comparison of the covering word w_1 with all its sub-words $w_2 \ldots w_{2^n}$ such that if the covering word in some boolean combination with all its sub-words yields 0 (a "collapsing function") then the covering word would contain all sub-strings made from its characters as sub-words, which would indicate a solution to the problem for a given word length.

Chapter 3

Optimal Design of a Micro-Electrical-Mechanical Systems Actuator

Participants: David Ross (Mentor), Kyle Biswanger, C. Sean Bohun, Lloyd Bridge, Leevan Ling, Dominique Noel, Simal Saujani, Daniel Spirn, Fridolin Ting.

PROBLEM STATEMENT: Fundamental to the design of an inkjet printer is precise delivery of ink from the printer to the paper. One proposed method is to manufacture a tiny beam of metal in such a way that when one end is heated, the beam bends thereby projecting a tiny volume of ink onto the paper.

A preliminary beam has been manufactured at Eastman Kodak with the overall dimensions 100μm $\times$ 20μm $\times 5\mu$m. This particular beam consisted of two materials, aluminium (Al) and silicon dioxide (SiO_2) in a ratio of 3:2. A voltage pulse of 10μs was applied to the beam heating it up to about 400K and resulting in a maximum rate of deflection of about 0.2ms^{-1}.

The problem set forth was to first model the beam described above in the hopes of understanding the underlying physics. The second goal was to generalize the model to design a beam with perhaps more layers that achieves a maximum deflection rate of at least 1ms^{-1}. Because of the nature of the fluid, the temperature of the beam must not exceed about 400K. In addition, the overall dimensions of the beam are required to be about the same as the preliminary beam discussed above. As a result, the only free parameters are the choice of materials for the beam and in which amounts they should be chosen.

3.1   Introduction

Solving the equations for the full beam/fluid flow, even numerically, is a formidable task, and now we shall proceed to simplify the model as much as possible. Of course, we intend to justify this process in the subsequent analysis.

The assumptions:

- We may treat the problem with one space dimension. Moreover, we shall neglect calculation of the flow field and opt to model the effect of the fluid on the beam with a parameterization scheme.

- The flow carries little fluid away (relative to the length scale of the beam) and so the convective term in temperature conservation equation may be dropped.

- We assume each layer of the beam to be homogeneous and the heating to be uniform; consequently, we expect a uniform temperature profile. Furthermore, we assume linear elasticity theory is sufficient to model the beam, and that boundary conditions may be applied at the initial (unstretched/contracted) positions. We shall also neglect the thermal expansion of the oxide.

These simplifications are implicit in what follows.

3.2   Heat Transport in the System

Of central importance in the modelling of this problem is the transport of heat from the Al into the SiO_2 and surrounding Isopar fluid. A current is supplied to the aluminium, which generates an amount of heat. Since the thermal expansion coefficient of Al is large with respect to SiO_2, the beam will bend. If we could determine the temperature of the aluminium as a function of time, we could approximate the displacement of the end of the beam and thus estimate the beam speed.

Listed below are some of the thermal properties of of Al, SiO_2 and the surrounding Isopar fluid. The density of a material is denoted as ρ while the specific heat and conductivity are denoted as c_v and k respectively.

Material	ρ (g cm^{-3})	c_v (J g^{-1}K^{-1})	k (J cm^{-1}s^{-1}K^{-1})
Fluid (Isopar)	0.77	2.1	1×10^{-3}
Silicon Dioxide	3.4	0.7	1.38×10^{-2}
Aluminium	2.7	0.5	2.31

With these values, the first question that we ask is can we disregard temperature variations in the oxide? If the temperature variations in the oxide layer are negligible, then thermally, we could simply model the beam as being made out of aluminium. The rule of thumb is that in time Δt, heat diffuses a length Δx given by the expression $\Delta x = (k\Delta t/\rho c_v)^{1/2}$. Hence for SiO_2, heat diffuses a length approximately $1\mu m$ in $5\mu s$. Since the depth of the SiO_2 is approximately $3\mu m$, we cannot disregard the temperature variations in SiO_2. Therefore, we must account for both materials.

The equations governing the heat flow are:

$$
\begin{aligned}
\rho_f c_{v\,f}\theta_t &= k_f \theta_{xx} & \text{(Fluid)} \\
\rho_{ox} c_{v\,ox}\theta_t &= k_{ox} \theta_{xx} & \text{(Silicon Oxide)} \\
\rho_{Al} c_{v\,Al}\theta_t &= k_{Al} \theta_{xx} + Q & \text{(Aluminium)}
\end{aligned}
\tag{3.1}
$$

where $Q = 5.35 \times 10^7$ Watts cm^{-3}. The boundary conditions are determined by the empirical fact that temperature is continuous and energy is conserved across the interface boundaries. These conditions imply

$$
\begin{aligned}
\theta(\text{interface}^-) &= \theta(\text{interface}^+) \\
k^- \theta_x(\text{interface}^-) &= k^+ \theta_x(\text{interface}^+)
\end{aligned}
\tag{3.2}
$$
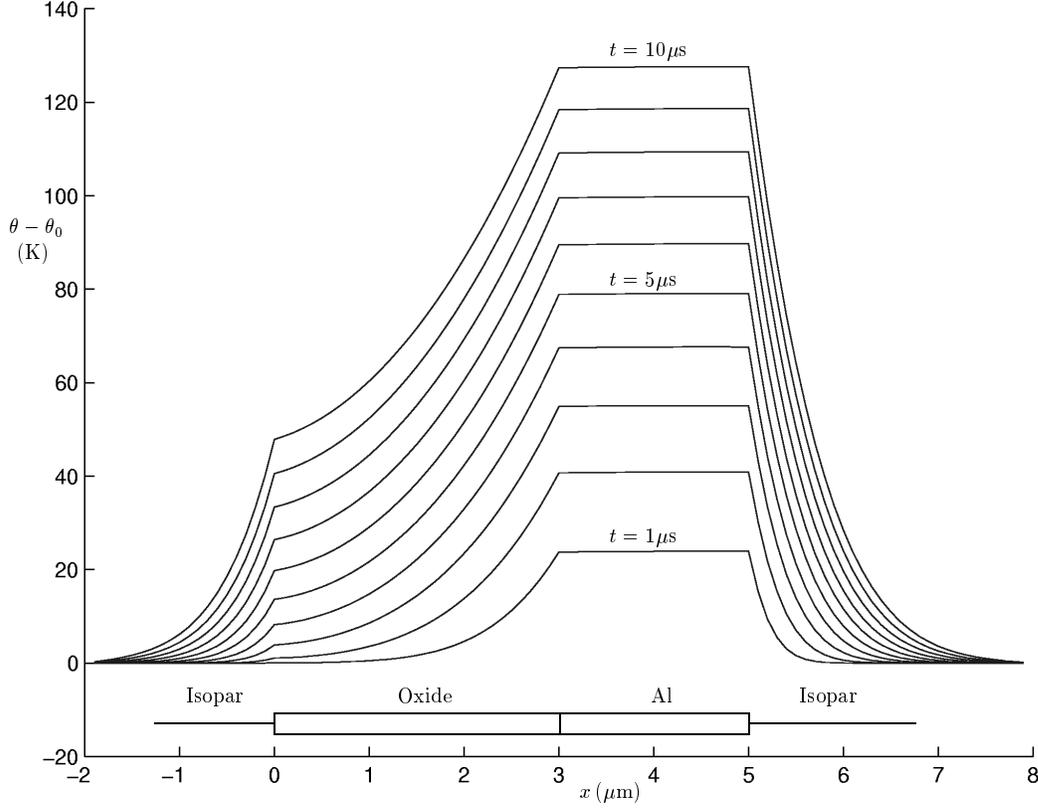
Figure 3.1: Temperature through a cross section of the beam as heat is applied.

at any interface. In addition, at infinity the system should be at room temperature so that $\theta(x = \pm\infty, t) = \theta_0 = 300\mathrm{K}$.

We decided to limit our scope to the one dimensional problem. Notice that the one dimensional heat flow equation in aluminium can be greatly simplified by integrating. As a result, we get

$$\rho_{\mathrm{Al}} c_{v\,\mathrm{Al}} \bar{\theta}_t = \frac{k_f \theta_x \Big|_c - k_{\mathrm{ox}} \theta_x \Big|_b}{L_{\mathrm{Al}}} + Q \tag{3.3}$$

where $\bar{\theta}_t$ is the rate of change of the average temperature of aluminium, L_{Al} is the width of the aluminium, and b and c denote boundaries oxide/aluminium and aluminium/fluid respectively. Since the conductivity k of aluminium is so high, it can be assumed that the temperature variation across the aluminium is zero. Hence the temperature of aluminium is spatially uniform. This fact greatly improves the efficiency of our numerical schemes. The numerical scheme we use did not make use of equation (3.3) but our results justify this approximation. The resulting temperature profile is displayed in figures 3.1 and 3.2.

3.3   Modelling the Beam

For the beam we consider a laminated beam with N layers labelled $\{1, 2, \ldots, N\}$ where layer j has a Young's modulus of E_j, a density of ρ_j and a thickness of $h_j - h_{j-1}$. With this notation, we take $h_0 = 0$ and $h_N = H$ the overall height of the laminated beam. When the beam is bent the surface outside the curve is stretched while the surface inside the curve is compressed. Internal to the beam there must be some surface which is neither stretched or compressed. This surface is known as the *neutral* surface. The location of this neutral surface, y_0, is found by summing the stress (force per unit area) in each of
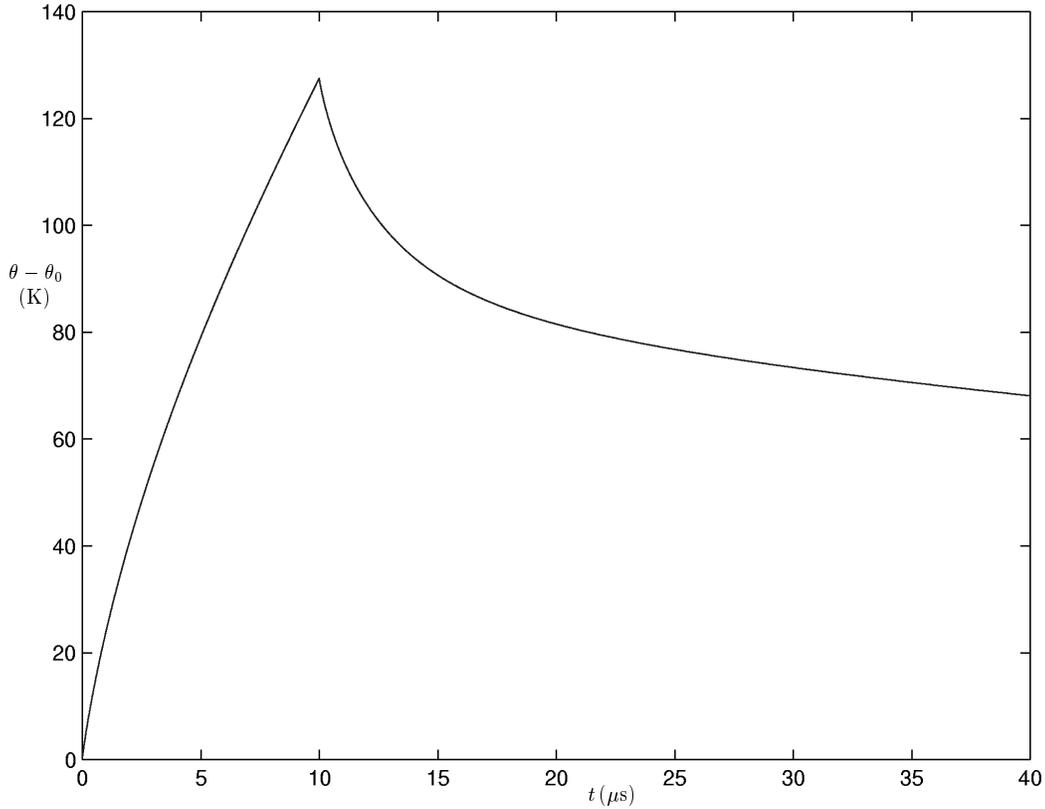
Figure 3.2: Temperature of the aluminium layer for a $t = 10\mu s$ heating pulse.

the layers and noting that the resultant stress is zero. This procedure gives

$$y_0 = \frac{\sum_{j=1}^{N} E_j(h_j^2 - h_{j-1}^2)}{2\sum_{j=1}^{N} E_j(h_j - h_{j-1})}. \tag{3.4}$$

It is interesting to note that if the Young's modulus was the same for all of the N layers then the above expression becomes a telescoping series and the neutral surface would lie at $h_N/2 = H/2$ which is the height of the centre of mass if the layers also all have the same density.

3.3.1 Beam Equation

Having located the neutral surface, one can determine the beam equation for this laminated structure. This is accomplished by computing the moment in each of the N layers at two horizontal positions, $x = x_0$ and $x = x_0 + \Delta x$. The details of this derivation are simple yet tedious. The resulting beam equation is

$$\rho H u_{tt} + D u_{xxxx} = P$$

where

$$\rho H = \sum_{j=1}^{N} \rho_j(h_j - h_{j-1}), \qquad D = \frac{1}{3}\sum_{j=1}^{N} E_j(y - y_0)^3 \Big|_{h_{j-1}}^{h_j} \tag{3.5}$$

are the weighted density and the composite flexural rigidity respectively. and P is the external pressure. If $E_j = E \, \forall j$ then using (3.4) we find $D = EH^3/12$ as one would expect for a uniform beam of thickness H. The value of E and α for the various materials are listed below.

Material	E $(\mathrm{g\,cm^{-1}s^{-2}})$	α $(\mathrm{K^{-1}})$
Silicon Dioxide	6×10^{11}	$\simeq 0$
Aluminium	20×10^{11}	16×10^{-6}

3.3.2   Boundary Conditions

In order to be well posed, the equation for the beam requires a number of boundary conditions and initial conditions. The initial conditions are simply that the beam has no velocity and is not bent. That is, $u(x,0) = 0 = u_t(x,0)$.

There are four boundary conditions. Since the beam is fixed and clamped at the origin $x = 0$ we easily identify the conditions $u(0,t) = 0$ and $u_x(0,t) = 0$. In addition, the free end, $x = L$, does not experience any shear stress and as such, $u_{xxx}(L,t) = 0$.

The fourth boundary condition arises from the application of heat. Since the beam is laminated, each of the layers will expand at different rates when heated. This imbalance in the strains of the various layers creates a moment at the end $x = L$. We derive this temperature dependent moment next.

We first recall that the stress and strain are related by

$$\frac{F_j}{A_j} = E_j \frac{\Delta l_j}{l_j} \tag{3.6}$$

where E_j is the Young's modulus of the jth layer. Therefore a layer with $A_j = W(h_j - h_{j-1})$ will have $F_j = E_j W(h_j - h_{j-1})\Delta l_j/l_j$. The magnitude of l_j will depend on the layer. Before any heating takes place, each of the layers has a length denoted as l_0 and if we now heat the beam, each of the layers expands at a different rate. Let α_j denote the expansion rate of the jth layer so that $l_j = (1 + \alpha_j\theta)l_0$ is the amount the jth layer would have expanded at the temperature θ if it was not connected to the other layers. If we set l to be the mean amount of expansion of the beam as a whole after the various layers have expanded we have for the jth layer that

$$F_j = E_j W(h_j - h_{j-1})\frac{l - l_j}{l_j}.$$

However, these individual forces must cancel out so that $\sum_{j=1}^{N} F_j = 0$. Solving for l gives

$$l = \frac{\displaystyle\sum_{j=1}^{N} E_j (h_j - h_{j-1})}{\displaystyle\sum_{j=1}^{N} \frac{E_j}{l_j}(h_j - h_{j-1})}. \tag{3.7}$$

The quantity of interest is the the ratio $(l - l_j)/l_j$ and using the fact that even for temperatures on the order of 400K, $\alpha_j\theta \ll 1$ so using (3.7) gives the approximation

$$\frac{F_j}{A_j} = E_j \frac{l - l_j}{l_j} \simeq \theta E_j (\bar{\alpha} - \alpha_j) \qquad \text{where} \qquad \bar{\alpha} = \frac{\displaystyle\sum_{j=1}^{N} E_j \alpha_j (h_j - h_{j-1})}{\displaystyle\sum_{j=1}^{N} E_j (h_j - h_{j-1})}. \tag{3.8}$$

The moment generated by each layer satisfies $\theta E_j(\bar{\alpha} - \alpha_j) = E_j(y - y_0)u_{xx}(L)$. Multiplying by a factor of $(y - y_0)$ and integrating over the layers, one finds the total effective moment at the point $x = L$ to be

$$u_{xx}(L,t) = \frac{\theta(t)}{2D} \sum_{j=1}^{N} E_j(\bar{\alpha} - \alpha_j)(y - y_0)^2 \bigg|_{h_{j-1}}^{h_j} = \Gamma\theta(t) \tag{3.9}$$

which is linear with respect to the applied temperature.

3.4   Beam Fluid Interaction

Consider the following version of the beam equation that accounts to at least a first order approximation, for both the drag and the viscosity of the fluid

$$(\beta + \rho H)u_{tt} = -Du_{xxxx} - ku_t. \tag{3.10}$$

An expression for the natural frequency of the beam can be obtained by using separation of variables. Let $u(x,t) = F(x)G(t)$ and consider a slightly simplified version of the boundary conditions where the beam is not heated

$$u(x,0) = u_t(x,0) = u(0,t) = u_x(0,t) = u_{xx}(L,t) = u_{xxx}(L,t) = 0.$$

Under the separation, one obtains two expressions. For the spatial variable

$$F^{iv} - \frac{\lambda^4}{D}F = 0; \qquad F(0) = F'(0) = F''(L) = F'''(L) = 0$$

and for the temporal variable

$$(\beta + \rho H)G'' + kG' + \lambda^4 G = 0; \qquad G(0) = G'(0) = 0. \tag{3.11}$$

Focusing on the spatial equation, we find that

$$F(x) = A\left[\sin\left(\frac{\lambda x}{D^{1/4}}\right) - \sinh\left(\frac{\lambda x}{D^{1/4}}\right)\right] + B\left[\cos\left(\frac{\lambda x}{D^{1/4}}\right) - \cosh\left(\frac{\lambda x}{D^{1/4}}\right)\right]$$

where A and B are constants. The eigenvalues for λ arise from the boundary conditions at $x = L$. Computing the second and third derivatives at L leads to the compatibility condition

$$\begin{vmatrix} -\cos\xi - \cosh\xi & \sin\xi - \sinh\xi \\ -\sin\xi - \sinh\xi & -\cos\xi - \cosh\xi \end{vmatrix} = 0 \qquad \text{with} \qquad \xi = \frac{\lambda L}{D^{1/4}}.$$

This implies that the eigenvalues satisfy $1 + \cos\xi\cosh\xi = 0$ whose solutions are given by $\xi_0 = \pm 1.8751$ and $\xi_n \simeq \pm(2n + 1)\pi/2$ for $n \in \mathbb{N}$. The fundamental frequency and damping of the beam can now be determined by looking at the temporal equation.

3.4.1   Determining β and k

We observe from the experimental data available that, throughout its motion, the beam oscillates about some varying mean deflection. Not only this, but it is clear that, once the heat supply to the beam is turned off, the amplitude of these oscillations in the fluid decreases in time. Thus, into our model, we incorporate terms associated with a damped harmonic oscillator system, which will model the effect of the viscous fluid on the motion of the beam.

Since our model is one dimensional, we shall consider the free end of the beam, oscillating in one dimension in the fluid, as analogous to the mass in a mass-spring-dashpot system. For a mass m,

attached to the free end of a spring with spring constant c, and moving in a dashpot containing fluid with damping coefficient $2b$, the motion of the mass is governed by

$$m\ddot{x} + 2b\dot{x} + cx = 0. \tag{3.12}$$

Oscillatory solutions of this equation have the form

$$x(t) = Ae^{-bt/m} \sin\left(\sqrt{\frac{mc - b^2}{m^2}}t\right);$$

we identify the frequency of oscillation as $(mc - b^2)^{1/2}/m$ and the decay rate as b/m. Now, the frequency of oscillations in fluid appears constant, and was measured as 3.45×10^5 Hz. The appropriate data for frequency and damping calculations is summarized in the table below.

Fluid	Fundamental Frequency (MHz)	Amplitude at 15 μs (μm)	Amplitude at 35 μs (μm)
Air	0.484	0.27	0.26
Isopar	0.345	0.293	0.086

The decay rate, measured over the remaining time after $20\mu s$, is $b/m = 6.13 \times 10^4$. Following the separation of variables method we choose the fundamental mode $c = \lambda_0^4$, and so

$$\frac{\lambda_0^4}{m} - (6.13 \times 10^4)^2 = 4\pi^2 f_{\text{Isopar}}^2 = 4.70 \times 10^{12},$$

where $\lambda_0 = \xi_0^4 D/L^4$. That is,

$$m = \frac{\xi_0^4 D}{4.70 \times 10^{12} L^4} = 2.63 \times 10^{-12} \frac{D}{L^4} \quad \text{and} \quad b = 1.61 \times 10^{-7} \frac{D}{L^4}.$$

A comparison of (3.12) with the separation of variables (3.11) method yields

$$\beta = 2.63 \times 10^{-12} \frac{D}{L^4} - \rho H, \qquad k = 3.22 \times 10^{-7} \frac{D}{L^4} \tag{3.13}$$

as first approximations for the constants to be used in to match the given data. These numbers are later tuned to match the data as close as possible.

3.5 Results

As there were two goals in this project two cases were considered. The first case was a beam in the ratio of 2:3 of Al to SiO_2 in the Invar fluid. While in the second case, a ratio of 1:2 was chosen to maximize the coupling moment induced by the temperature. In this second case the beam is slightly thinner and therefore gets hotter for the same amount of energy input. The parameters for these two cases are summarized below. $T_{\max} = 397.3$K in case 2:3 and 402.4K in case 1:2 respectively.

Parameter	Case 2:3	Case 1:2
ρH	1.56×10^{-3}	9.50×10^{-4}
D	10.442	2.2542
Γ	4.755×10^{-2}	7.985×10^{-2}
Q	4.08×10^7	6.80×10^7
β	1.32×10^{-3}	1.32×10^{-3}
k	123	73.8
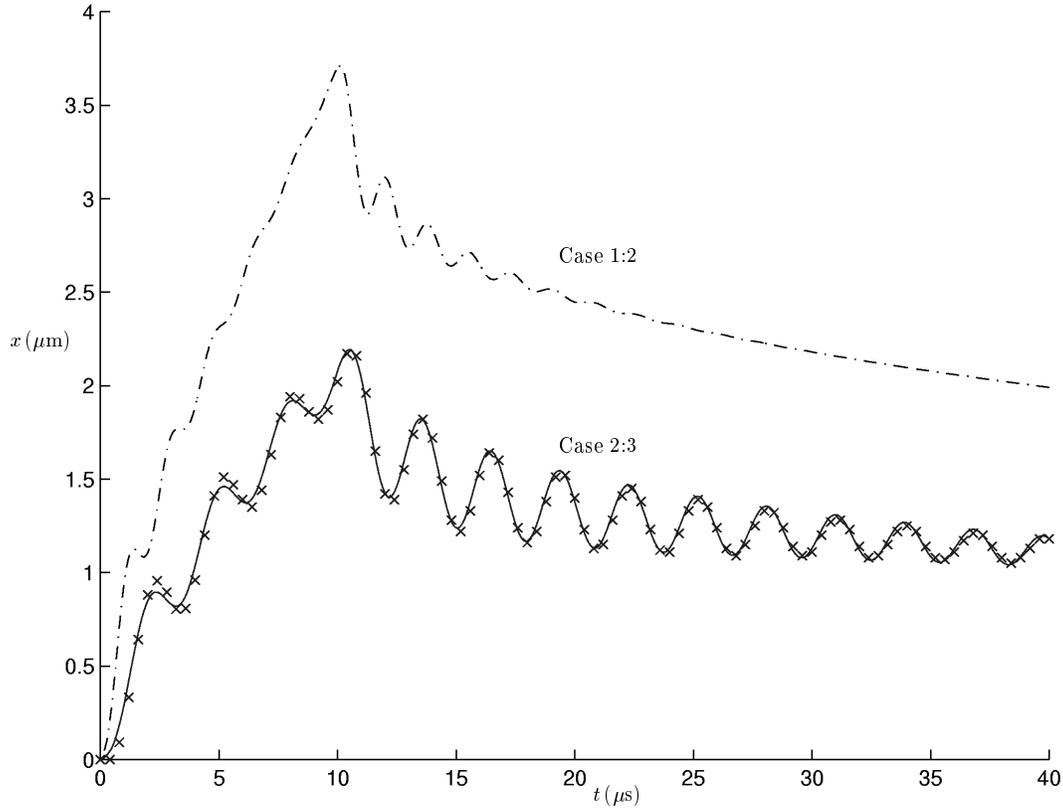
Figure 3.3: Position of the end of the beam with respect to time

The values of β and k are determined by matching the solution to the given experimental data in the case 2:3. Once these values are known, the same value of β is used in the case 1:2 as the same volume if Invar fluid is being accelerated in both cases. The value of k scales with the thickness of the beam.

For a given geometry a solution of the heat equation (3.1-3.2) determines $\theta(t)$. This time dependent temperature is then applied as a boundary condition for the beam equation (3.5), (3.8-3.10). Numerical solutions for the two cases are plotted below along with the experimental points. The agreement is astounding.

3.6   Conclusions and Directions

Our initial goal was to accelerate Isopar fluid to a speed of $1\mathrm{ms}^{-1}$ over $10\mu s$ using a beam that deflects when heated. Our first objective was to develop an appropriate physical model for the problem. The key simplifying assumptions included treating the problem as one dimensional, relying on the linear beam equation and neglecting the details of the fluid flow.

We were able to reproduce experimental results with high agreement. Furthermore, applying the theory, we were able to improve the speed of the fluid by a factor of 2.

Although we did not obtain our objective, we did make significant progress. The next step would be to consider more than two layers and possibly different materials. Despite the inherent difficulties, studying the two dimensional problem would be of interest. There's also evidence that an insulating layer would increase speed; this may increase the relaxation time beyond acceptable limits.

Chapter 4

Temperature Effects on a River or Estuary Due to the Construction of a Power Station

Participants: Colin Please (Mentor), Ibrahim Agyemang, Matthew Bolton, Samantha Carruthers, Irina Dinu, Shafiqul Islam, Jung Min Lee, Lila Rasekh, Sirod Sirisup, John Frederick Williams.

PROBLEM STATEMENT: It is expected that the construction of a power station on a river will have some pollution effects. We are particularly interested in the effect of the increase in temperature caused by the release of $2 \times 10^9 \, \mathrm{J \, s^{-1}}$ of heat from the power plant. The concern is that this increase in temperature may have possible ecological effects on the river.

Three different situations are modelled.
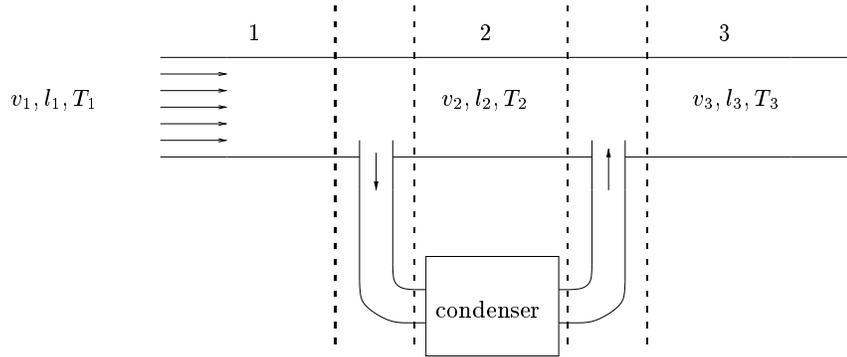
Figure 4.1: Idealized river with power station. $v_1 = 0.2\,\mathrm{m\,s^{-1}}$, $l_1 = 10\,\mathrm{m}$, $T_1 = 288\,\mathrm{K}$.

4.1 Introduction

We assumed that both river and estuary are channels with constant width. We also assume that without a power station the temperature remains constant and does not vary with depth. The density of the water is assumed constant because the temperature does not vary enough for it to be significant. Finally, the flow rate of the pipe to and from the power plant are taken to be 100 m^{3}s^{-1}.

Three situations are considered. The first is where we have a power station on a river and the intake pipe is upstream from the outflow. The second is where we reversed the position of the intake and outflow pipes. In the above instances, we considered the case where there is no heat loss to the surroundings and the case where there is a natural heat loss which is linearly dependent on the temperature. The final situation is the case of having a power station on an estuary. Here we must take into account that the velocity is time dependent, and the temperature is dependent on both time and space.

The following variable notation is used throughout.

$v :=$ the velocity of the river (*assumed plug flow - no depth dependence*)
$l :=$ the depth of the river
$\rho :=$ the density of water (*assumed constant* - $1000\,\mathrm{kg\,m^{-3}}$)
$w :=$ the width of the river (*assumed constant* - $100\,\mathrm{m}$)
$T :=$ the temperature of the river
$g :=$ the acceleration due to gravity ($9.81\,\mathrm{m\,s^{-2}}$)
$Q :=$ the heat energy added by the condenser ($2 \times 10^9\,\mathrm{J\,s^{-1}}$)
$E :=$ the flow rate of the water through the pipes ($100\,\mathrm{m^3s^{-1}}$)
$P :=$ the hydrostatic pressure
$c_p :=$ the specific heat of water ($4.186\,\mathrm{kJ\,kg^{-1}K^{-1}}$)
$h :=$ the surface heat transfer coefficient ($30\,\mathrm{J\,s^{-1}m^{-2}K^{-1}}$)

To model these situations we considered what was occurring across the three boundaries of the river (see figure 4.1) using the conservation laws of: mass, force, and energy. We assumed that there is no heat conduction in our model and the heated water is discharged vertically from the outlet.

4.2 Solution of Idealized Models

4.2.1 Case 1. Intake Upstream of Outflow

Note that the variable subscripts refer to the region from which thee corresponding quantities are taken. Also note that v_p refers to the velocity inside the pipe (assumed constant: $2\mathrm{m\,s^{-1}}$), and T_p is the temperature inside the pipe.

Mass Conservation

The amount of fluid entering any junction must equal the amount of fluid leaving the junction. This gives two relations,

$$\rho v_1 l_1 w \;=\; \rho v_2 l_2 w + \rho E \qquad \text{(1--2 junction)}$$

$$\rho v_2 l_2 w + \rho E \;=\; \rho v_3 l_3 w \qquad \text{(2--3 junction)}$$

Force Balance (Newton's second Law)

Balancing the forces at each junction also gives two expressions, where the the hydrostatic pressure is given by $P_i = (l_i - z)\,\rho g + P_{\text{air}}$.

$$w\left(\int_0^{l_1} P_1\,dz - \int_0^{l_2} P_2\,dz - \int_{l_2}^{l_1} P_{\text{air}}\,dz\right) \;=\; \rho w l_1 v_1^2 - \rho w l_2 v_2^2 \qquad \text{(1--2 junction)}$$

$$w\left(\int_0^{l_2} P_2\,dz - \int_0^{l_3} P_3\,dz - \int_{l_2}^{l_3} P_{\text{air}}\,dz\right) \;=\; \rho w l_2 v_2^2 - \rho w l_3 v_3^2 \qquad \text{(2--3 junction)}.$$

Conservation of Energy

Finally, equating the energy across any junction gives,

$$w v_1 l_1 \left(\tfrac{1}{2}\rho v_1^2 + \rho c_p T_1\right) - E\left(\tfrac{1}{2}\rho v_p^2 + \rho c_p T_1\right) \;=\; w v_2 l_2 \left(\tfrac{1}{2}\rho v_2^2 + \rho c_p T_2\right) \qquad \text{(1--2 junction)}$$

$$w v_2 l_2 \left(\tfrac{1}{2}\rho v_2^2 + \rho c_p T_2\right) + E\left(\tfrac{1}{2}\rho v_p^2 + \rho c_p T_p\right) \;=\; w v_3 l_3 \left(\tfrac{1}{2}\rho v_3^2 + \rho c_p T_3\right) \qquad \text{(2--3 junction)}$$

where the added heat from the power station is $\rho E c_p T_p = \rho E c_p T_1 + Q$.

Scaling

These nonlinear algebraic equations pose a formidable problem. In order to simplify the equations, appropriate dimensionless scalings were introduced in the hope that small parameters would be found that may be neglected. We set

$$l_i = (1 - \delta_i) l_1, \qquad v_i = \mu_i v_1, \qquad T_i = (1 - \beta_i) T_1$$

where $i \in \{2, 3\}$. In addition we identify the following dimensionless parameters:

$$\alpha_1 = \frac{E}{v_1 l_1 w} \sim \frac{1}{2}, \qquad \alpha_2 = \frac{v_1^2}{g l_1} \sim 4.1 \times 10^{-4}, \qquad \alpha_3 = \frac{v_1^2}{2 c_p T_1} \sim 1.7 \times 10^{-8},$$

$$\alpha_4 = \frac{v_p^2}{v_1^2} = 100, \qquad \alpha_5 = \frac{2Q}{E \rho v_1^2} 1.0 \times \; \sim 10^6.$$

Solution

With these scalings the equations become:

$$2 \;=\; 2\mu_2(1 - \delta_2) + 1\,,$$

$$1 \;=\; \mu_3(1 - \delta_3)\,,$$

$$1 - (1 - \delta_2)^2 \;=\; 2\alpha_2[1 - \mu_2^2(1 - \delta_2)]\,,$$

$$(1 - \delta_2)^2 - (1 - \delta_3)^2 \;=\; 2\alpha_2[\mu_2^2(1 - \delta_2) - \mu_3^2(1 - \delta_3)]\,,$$

$$C_1 \;=\; \alpha_3 \mu_2(1 - \delta_2) + \mu_2(1 - \delta_2)(1 - \beta_2)\,,$$

$$\mu_2^3 \alpha_3(1 - \delta_2) + (1 - \beta_2)(1 - \delta_2)\mu_2 + C_2 \;=\; \mu_3(1 - \beta_3)(1 - \delta_3) + \mu_3^3 \alpha_3(1 - \delta_3)$$

where $C_1 = 1 + \alpha_3 - \alpha_1\alpha_3 - \alpha_1$ and $C_2 = \alpha_4\alpha_1\alpha_3 + \alpha_1 + \alpha_5\alpha_1\alpha_3$. Recognizing the fact that many parameters in the system are so small that they are negligible, the system is easily solved to leading order:

$$\delta_2 = 0, \quad \beta_2 = 0, \qquad \mu_2 = 0.5,$$
$$\delta_3 = 0, \quad \beta_3 = -\alpha_1\alpha_3\alpha_5, \quad \mu_3 = 1.$$

Physically this means that the height of the river is essentially unchanged, the temperature increases by approximately $2.3°C$ at the outflow, and the velocity in region two is half the normal velocity of the river.

4.2.2 Case 2 - Reversed Flow

With an understanding of the important scalings in this problem we can repeat the analysis for the case where the inflow and outflow are reversed. From the analysis above, we can immediately write down the solution to the new problem.

$$\delta_2 = 0, \quad \beta_2 = -\alpha_1\alpha_3\alpha_5, \quad \mu_2 = 2,$$
$$\delta_3 = 0, \quad \beta_3 = -\alpha_1\alpha_3\alpha_5, \quad \mu_3 = 1.$$

Physically this means that the temperature increases only at the outflow and by the same amount as obtained earlier. This is because instead of the central section moving more slowly, now it moves faster. The fact that the recirculated water does not get hotter and hotter may seem counterintuitive at first glance. However, in this case the outflow mixes with the entire volume of the inflow diluting the amount of heat added. In summary this analysis has shown that the temperature downstream must increase by the same amount regardless of whether the outflow is upstream or downstream.

4.2.3 Case 3 - Heat Loss

With some insight into the problem, we can consider the more complicated case of accounting for heat loss. Assuming temperature loss is linearly proportional to the difference between the temperature at a position x and the natural temperature T_1, where the proportionally constant $h = 30 \, \mathrm{J\,s^{-1}m^{-2}K^{-1}}$, we have the energy equation:

$$\frac{dT}{dx} = \frac{h}{\rho l c_p v}[T_1 - T(x)]; \qquad T(x_0) = T_1$$

One can easily solve to find

$$T(x) = T_1\left[1 + e^{-\gamma(x-x_0)}\right]; \qquad x \geq x_0$$

where x_0 is the distance between the intake and outflow pipes and $1/\gamma = \rho l c_p v/h \sim 280$ km. Thus, the temperature decays exponentially as you go downstream with a decay length of approximately 280 km. This means that the temperature difference decays by a factor of e for every 280 km you go.

4.3 Flow in an Estuary

An estuary is a river system that is affected by large tides. To understand the flow in an estuary, we assume a river as in section 2 with the addition of a dam at one end. Through a weir in the dam, we imagine a constant rate of flow into the water. This models feeding into the river upstream from the tide without having to worry about changing water levels in the water network upstream. We know that the tide rises and falls periodically which gives us a time dependent boundary condition at one end of our estuary.

By considering conservation of mass and energy we have the shallow water equations

$$\frac{\partial l}{\partial t} + \frac{\partial(vl)}{\partial x} = 0, \qquad \frac{\partial(vl)}{\partial t} + \frac{\partial}{\partial x}\left(\frac{1}{2}gl^2 + v^2 l\right) = 0$$
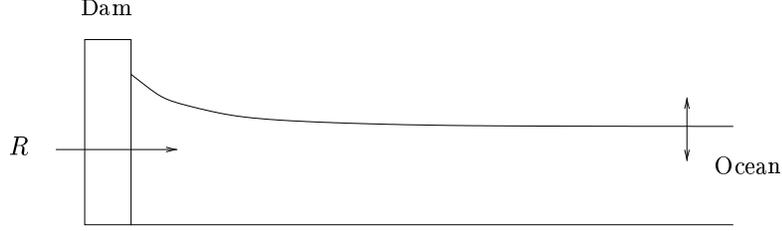
Figure 4.2: An idealized estuary.

with the boundary conditions

$$R = l(0,t)v(0,t)w\,, \qquad l(x_0,t) = l_0 + l_1\cos(\omega t)$$

where $R = 1400\,\mathrm{m^3 s^{-1}}$, $w = 100\,\mathrm{m}$, $l_0 = 10\,\mathrm{m}$, $l_1 = 1\,\mathrm{m}$, $\omega = 2\pi/12.4\,\mathrm{hrs}$, $x_0 = 100\,\mathrm{km}$.
In dimensionless form we then have

$$\frac{\partial h}{\partial \tau} + \epsilon\frac{\partial(uh)}{\partial y} = 0\,, \qquad \frac{\partial(uh)}{\partial \tau} + \epsilon\frac{\partial}{\partial y}(u^2 h + \alpha_1 h^2) = 0$$

with

$$u(0,\tau)h(0,\tau) = \alpha_2\,, \qquad h(1,\tau) = 1 + \epsilon\cos\tau$$

and where $\tau = \omega t$, $l_0 h = l$, $x_0 y = x$, $v_0 u = v$,

$$v_0 = \frac{\omega l_1 x_0}{l_0} \sim 1.4\,\mathrm{ms^{-1}}\,, \quad \epsilon = \frac{l_1}{l_0} = \frac{1}{10}\,, \quad \alpha_1 = \frac{l_0 g}{2v_0^2} \sim 24.8\,, \quad \alpha_2 = \frac{R}{l_0 w v_0} \sim 1\,.$$

Assuming that

$$h = 1 + \epsilon h_1(y,\tau) \quad \text{and} \quad u = 0 + \epsilon u_1(y,\tau)\,,$$

we get

$$\frac{\partial h_1}{\partial \tau} + \epsilon\frac{\partial u_1}{\partial y} = 0 \quad \text{and} \quad \frac{\partial u_1}{\partial \tau} + 2\epsilon\alpha_1\frac{\partial h_1}{\partial y} = 0\,.$$

Cross-differentiating we obtain the wave equation for the correction to the height

$$\frac{\partial^2 h_1}{\partial \tau^2} - c^2\frac{\partial^2 h_1}{\partial y^2} = 0$$

with $c^2 = 2\alpha_1\epsilon^2$ and boundary conditions

$$h_1(1,\tau) = \cos\tau, \qquad \frac{\partial h_1(0,\tau)}{\partial y} = 0\,.$$

Detailed analysis of this equation, and recalling our physical scalings suggests that a reasonable flow in the estuary generated by the tidal motion is of the form

$$v = uv_0 = v_1 + v_0\cos(\omega t), \qquad l = l_0 h = l_0\,.$$

The spatial variation has been neglected as it varies over a scale much larger than we are considering. Although the height truly does vary in time, it does so slowly and relatively little and is thus not the most important effect.
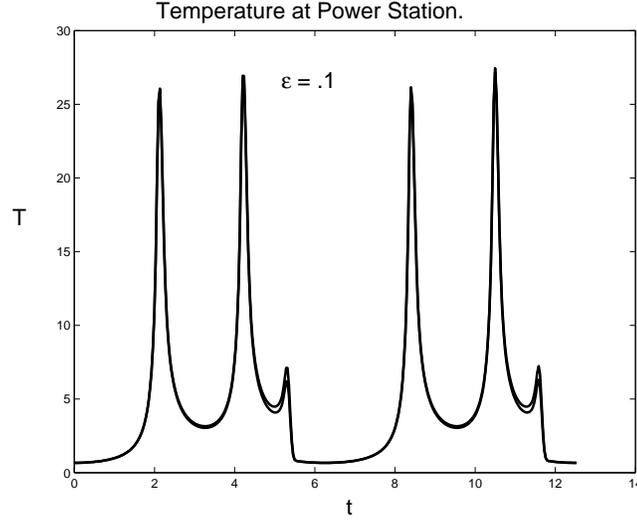
Figure 4.3: Temperature variation in the idealized estuary.

4.4   Temperature Variation on an Estuary

In our model of an estuary we will make all the same assumptions as above for the geometry and physical properties of the estuary and power plant. From our analysis of the algebraic case and of the previous modelling of the estuary, we also assume that $v = v_1 + v_0 \cos(\omega t)$ in the river and that the height of the river is constant. We will take the mean velocity to be very small, that is $\frac{v_1}{v_0} = \delta \ll 1$. Writing down the dimensionless energy balance with heat loss as we had done earlier, we have

$$\frac{\partial \beta}{\partial \tau} + u(\tau)\frac{\partial \beta}{\partial y} = -\gamma \beta \,,$$

where

$$u = \cos(\tau) + \delta \ \ \text{for} \ \ -y^* < y < 0$$
$$u = \cos(\tau) + \mu_2\delta \ \ \text{for} \ \ 0 < y < y^* \,,$$

with conservation across the outlet by the equation

$$\delta\left[(\mu_2 - 1) + \left(\mu_2\beta(y_-^*, \tau) - \beta(y_+^*, \tau)\right)\cos\tau\right] + \alpha_1\left[1 + \beta(0,\tau) + \left(\beta(y_-^*,\tau) - \beta(y_+^*,\tau)\right)\cos\tau\right] = 0 \,.$$

Although in the no loss case we have a simple-looking advection equation, we are confounded from finding an analytical solution by the complicated velocity term and the nonlocal jump condition across the outlet. Instead we consider a numerical solution.

To solve the problem numerically we have assumed that the source and outlet are at the same location. This is reasonable due to the scale - the separation is usually less than 1 km but we are interested in a scale of hundreds of kilometres. For these computations we used an explicit up-winding first-order finite difference method.

This lets us answer the questions posed by our mythical engineer, what is the temperature distribution in space and how hot does the water near the station get? As one would expect, we get an oscillatory solution slowly drifting down river and decaying slowly when heat loss is considered. The two important parameters in this problem are the heat loss coefficient h and the mean river velocity v_0.

4.5   Conclusions

A model of the temperature distribution in an idealized 1-D river has been constructed. As a first study, mass, force and energy conservation arguments were used to obtain the steady state solution for constant river velocity. This model was then refined to include heat loss such that an estimate for the length scale of temperature decay downstream of the power station could be obtained. The problem of a tidal estuary (with heat loss) was addressed, and the resulting nonlinear advection equation was solved numerically. As one would expect, when the mean flow velocity is small the temperature distribution is localised and the peak temperature is high; when it is large the temperature is spatially more spread out, and the peak temperature is lower. The numerical model provides a method to calculate the important quantitative information required to assess the environmental impact of the power station on the estuary.

Chapter 5

Optimal Policies for Disk Controllers

Participants: Rachel Kuske (Mentor), Nicola Costanzino, Bruce Rout, Calin Anton, Cristina Popescu, Leonid Mocofan, Amir Sepasi, Nathan Krislock, Zhihui Xue.

PROBLEM STATEMENT: The problem we consider is one of trying to maximize the amount of information processed by a system consisting of a CPU and a RAID disk controller. We envision a situation where the information to be processed consists of write data, read data, read requests and a small amount of other miscellaneous jobs. We make the distinction between read data and read requests because a request to retrieve data from, say, the hard drive or the cache is very small in size, while the actual data that is retrieved and read may be very large. On the other hand, the data to be written is accompanied with a write request that is very small in size compared to the average size of data that is to be written. For this reason, we don't distinguish between a stream that consists of write data accompanied by a write request, and a stream consisting of just write data. The CPU processes these requests by sending them (and any associated data) to the disk controller. Since we assume that our CPU can only process serially, the job of the disk controller is to manage access to the hard disks, such that the total time the CPU spends processing information is reduced from the time it would take without the controller.

Our job is to first propose a simple model for the system in question and identify the important and necessary parameters. Once this is done, we consider a few cases where we try to optimize the processing ability of the system by tuning the parameters of the disk controller.

5.1   Methodology of Solution

Below is a schematic of the system we consider in this report.

CACHE

X_w

λ_w f_w μ_w Y_w g_w Z_w γ_w

HD

CONTROLLER

CPU

X_r W_r

λ_r f_t μ_t g_r Z_r γ_r

f_r μ_r

Y_r
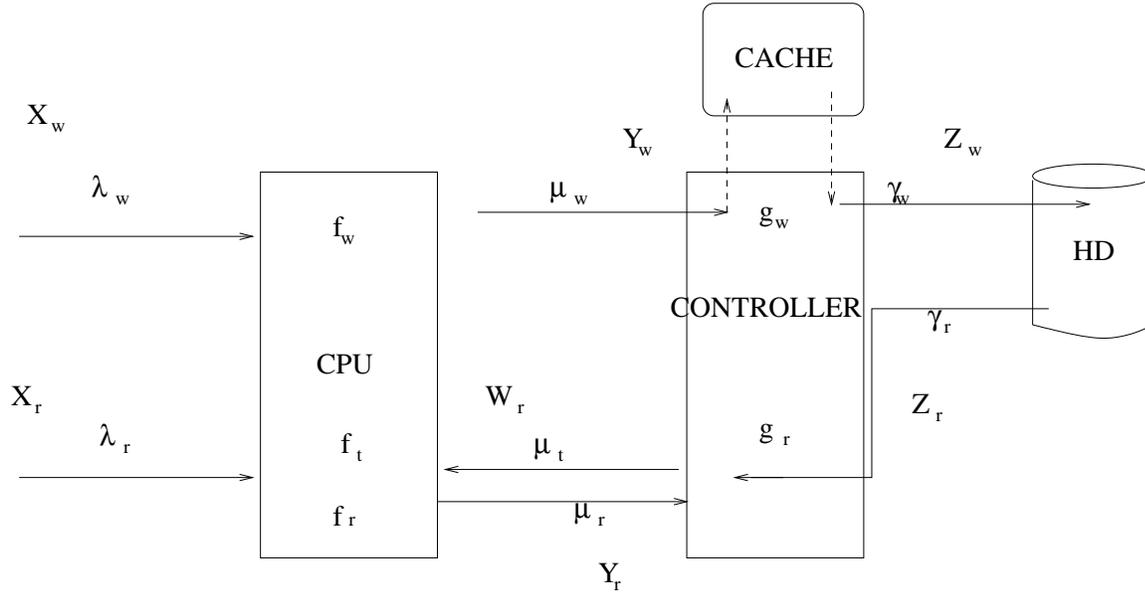
We describe this system using a fluid model, which assumes that the data is flowing in and out at a fairly regular rate. This approach has been very successful in situations where there is more-or-less uniform usage, or that the time scale in which we monitor the flow of information is large compared to the time interval between the discrete batches being sent to the CPU. Using this approach, we can model the data and requests as a fluid-like steam quantities, and describe everything in terms of rates with units of blocks per unit time. The parameters of the model are defined as follows:

λ_w : rate of incoming write data
λ_r : rate of incoming read requests
X_w : amount of write data yet to be processed by CPU
X_r : amount of read requests yet to be processed by CPU
μ_w : rate CPU can write data to cache (Megabytes per second)
μ_r : rate CPU can process read requests
f_w : fraction of time CPU spends writing data
f_r : fraction of time CPU spends writing write requests
Y_w : amount of write data in the cache
Y_r : amount of read data in the cache
d : ratio between average read data size and average read request size
γ_w : rate controller processes writes to HD
γ_r : rate controller processes read requests
g_w : fraction of time controller spends writing data to HD
g_r : fraction of time controller spends processing read requests
Z_w : amount of write data written to HD
Z_r : amount of read data written to cache
μ_τ : rate CPU can process read data
f_τ : fraction of time CPU spends reading data
f_o : amount of time controller spends dealing with other jobs

W_r : amount of read data processed by CPU
$Y_{\max}$: maximum data the cache can hold
T : the amount of time so that the system is completely processed

Consider a situation as in the previous schematic. For the analytical portion of this report, we will work in the regime where the cache never gets full. Using the notations above we arrive at the following model of a CPU- controller system:

$$
\begin{aligned}
\dot{X}_w(t) &= \lambda_w - \mu_w f_w & \dot{X}_r(t) &= \lambda_r - \mu_r f_r \\
\dot{Y}_w(t) &= \mu_w f_w - \gamma_w g_w(t) & \dot{Y}_r(t) &= \mu_r f_r - \gamma_r g_r(t)\, d \\
\dot{Z}_w(t) &= \gamma_w g_w(t) & \dot{Z}_r(t) &= \gamma_r g_w - f_\tau \mu_\tau \\
& & \dot{W}_r(t) &= \mu_\tau f_\tau
\end{aligned}
\tag{5.1}
$$

where we set

$$
f_w = \begin{cases} \tilde{f}_w & \text{if } Y_r + Y_w + Z_r < Y_{\max} \\ 0 & \text{otherwise}, \end{cases}
$$

and

$$
f_\tau = \begin{cases} \tilde{f}_\tau & \text{if } Z_r > 0 \\ 0 & \text{otherwise}. \end{cases}
$$

This form for f_w ensures us that the cache is never full, and likewise the choice for f_τ ensures us that we only spend time reading from the cache if there is data in the cache to be read. Clearly, we need to impose some conditions on the state variables in order that we can pick out optimal g_w and g_r. For the analytical portion of this report, we will consider the case where we never fill up the cache. Under this assumption we have the constraints

$$
\begin{aligned}
Y_r > 0, \qquad Y_w > 0, \qquad Z_r > 0, \qquad Y_r + Y_w + Z_r < Y_{\max} \\
f_r + f_w + f_o + f_\tau = 1, \qquad g_r + g_w = 1.
\end{aligned}
\tag{5.2}
$$

We pick out parameters g_r and g_w such that the average throughput $\mathcal{E}$ is maximized, where our control $\mathcal{E}$ is given by

$$
\mathcal{E} := \frac{1}{T} \int_0^T [Z_w(g_w)(t) + W_r(g_r)(t)]dt
\tag{5.3}
$$

5.2   Results

We analysed three cases. The first case is one in which the disk controller has a fixed ratio for the dedication percentages g_w and g_r. The second case is one in which we assume that the dedication percentages vary with time in a way that is proportional to the amount of data in the controllers cache. Finally, we numerically consider a stochastic generalization of the first case.

5.2.1   CASE 1

We first consider maximizing the throughput over all constant values of g_r and g_w, that is, we maximize $\mathcal{E}$ over the set

$$
\mathcal{A} := \{g_w \in [0,1]\}.
$$

In this case, a quick look at (5.3) convinces us that $\mathcal{E}$ is maximized by maximizing g_w. However, the constraints limit the size of g_w. Analyzing the requirements of (5.2) yields:

$$Y_r > 0 \Longrightarrow 1 - g_w < \frac{\mu_r f_r}{\gamma_r d} \in (0,1]$$

$$Y_w > 0 \Longrightarrow g_w < \frac{\mu_w f_w}{\gamma_w} \in (0,1]$$

$$Z_r > 0 \Longrightarrow g_w > \frac{\mu_\tau f_\tau}{\gamma_r} \in (0,1].$$

Hence our optimal policy for this case is to set the disk controller dedication percentages to

$$g_w = \min\left\{ \frac{\mu_w f_w}{\gamma_w}, 1 - \frac{\mu_\tau f_\tau}{\gamma_r} \right\} \qquad g_r = 1 - g_w. \tag{5.4}$$

Example 1:

As a demonstrative example, we choose $\lambda_w = 1, \lambda_r = 1/20, \mu_w = 2\mu_r = 20, f_w = 1/2, f_r = 1/400, d = 1/100, \gamma_w = 4/3, \gamma_r = 8/3, \mu_{tau} = 8/3$ and $f_\tau = 1/4$. These parameters satisfy all the inequalities and tell us that the optimum dedication percentages for the disk controller for this strategy is $g_w = 3/4$ and $g_r = 1/4$.

5.2.2   CASE 2

We now turn our attention to the case where we set

$$g_w = \alpha(Y_r(\alpha) + Y_w(\alpha) + Z_r(\alpha))(t), \qquad \alpha > 0 \tag{5.5}$$

that is, the percentage of time dedicated to process the write request is directly proportional to how full the cache is. This is a reasonable ansatz because the fuller the cache is, the more time we should dedicate to emptying it out (recall that we are working under the requirement that the cache never gets overflowed). For this case we have the nonlinear program

$$\max \frac{1}{T} \int_0^T [Z_w(g_w)(t) + W_r(g_r)(t)]dt$$

over all g_w such that

$$g_w := \alpha(Y_r + Y_w + Z_r)$$

satisfies (5.2). In this case we must still solve the ODE system (5.1), but this time we substitute our ansatz (5.5) for g_w. This leads to the system

$$\dot{\mathbf{v}} = A\mathbf{v} + \mathbf{b}$$

where

$$A = \begin{pmatrix} -\alpha\gamma_w & -\alpha\gamma_w & -\alpha\gamma_w \\ -\alpha\gamma_r d & -\alpha\gamma_r d & -\alpha\gamma_r d \\ -\alpha\gamma_r & -\alpha\gamma_r & -\alpha\gamma_r \end{pmatrix},$$

$$\dot{\mathbf{b}} = \begin{pmatrix} \mu_w f_w \\ -\gamma_r d + \mu_r f_r \\ \gamma_r - f_\tau \mu_\tau \end{pmatrix}.$$

and

$$v = \begin{pmatrix} Y_w \\ Y_r \\ Z_r \end{pmatrix}.$$

Example 2:

Equations for the full nonlinear program (i.e. without putting in any values for the parameters) are horrendously long and unintelligible, so for simplicity we consider a particular example for the parameter values given in the previous case. For this the solutions to the ODE's are

$$Y_r(t;\alpha) = 0.04848t - \frac{1}{\alpha}(0.002877 - 0.0002877e^{-4.00\alpha t})$$

$$Y_w(t;\alpha) = \frac{0.09381}{\alpha}(4.5298\alpha t + 1.53347(1 - e^{-4.00\alpha t}))$$

$$Z_r(t;\alpha) = -0.47343t + \frac{0.287717}{\alpha}(1 - e^{-4.00\alpha t}).$$

With some computations, we see that the problem

$$\max \frac{1}{T} \int_0^T [Z_w(g_w)(t) + W_r(g_r)(t)]dt$$

subject to (5.2) is equivalent to

$$\max_\alpha \{\alpha(Y_r(t;\alpha) + Y_w(t;\alpha) + Z_r(t;\alpha))\}$$

under the same constraints. The feasible set for α is given through

$Y_r > 0 \implies \alpha \in (-1.360755/T, 0) \cup (0.76365 \times 10^{-8}, \infty)$
$Y_w > 0 \implies \alpha \in \mathbf{R}$
$Z_r > 0 \implies \alpha < 0.5365675/T.$

which implies that $\alpha \in (-1.360755/T, 0) \cup (0.76365 \times 10^{-8}, 0.5365675/T)$. After solving the nonlinear program we find that the optimal value of α is

$$\alpha = 0.5365675/T, \tag{5.6}$$

and thus putting this value of α in the formula for g_w, we get

$$g_w(\alpha t, T) : \quad = \quad 0.5365675(-0.4249499666\alpha t + 0.5356811791 - 0.5356811791T\ e^{-2.144839153\frac{\alpha t}{T}}$$

$$+ \quad 0.1748382854(2.430531537\frac{\alpha t}{T} + 1.533466752 - 1.53346755e^{-2.144839153\frac{\alpha t}{T}})T)/T.$$

Representing the g_w as a function of t and T we get the pictures from figure (5.1), where Dt represents the total time the experiments is run for. As it can be seen, g_w tends to stabilize at a value around 0. The relatively small value of g_w can be interpreted as a proof that the disk controller is kind of "intelligent".

5.2.3   CASE 3

Here we consider a case where the input data has an associated stationary distribution. For the simulations, a standard poisson distributed write and read input data stream was randomly generated having
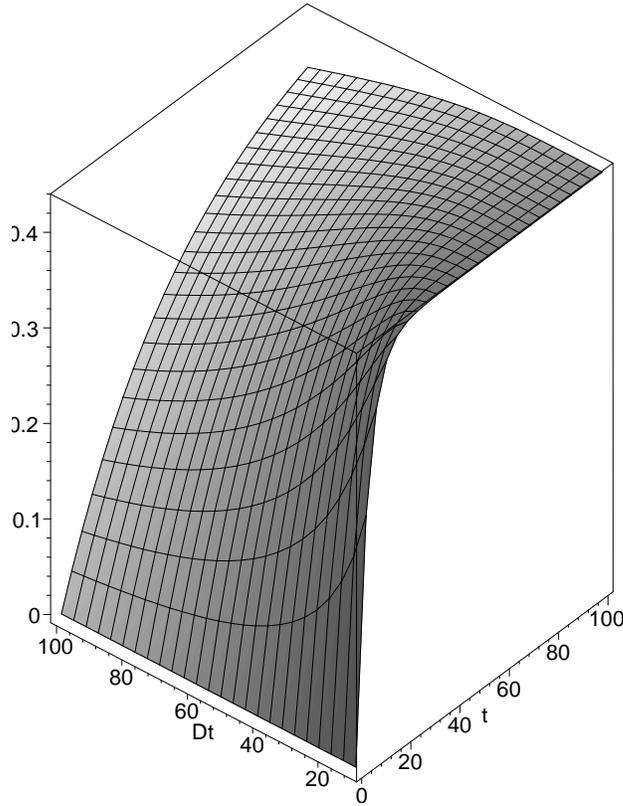
Figure 5.1: g_w as a function of t and T

an averages λ_w and λ_r respectively. For the parameters, we used the sample values found in Examples 1 and 2.

The model assumed a disk controller buffer size of 4 Megabytes for our simulations taking time steps of 3.75 seconds before clearing the buffer. The first figure (5.2) shows the distribution of data being input to the CPU and the "output" stream as read requests. The second figure (5.3) shows the states of the controller buffer. It can be seen that the buffer is being cleared each time step under our parameters but is being nearly fully utilized throughout the run. The intermittent stream represents data sent to the buffer before it is cleared by the controller (in the middle of time steps). The data lines across the bottom of the graph represent amounts of data in the cache at the end of each time step.

The third figure (5.4) shows expected amounts of input and output data. In the model the input data is in a buffer and has to wait for the CPU to send it to the controller's cache. The output stream represents the expected read requests serviced in each time step.

It can be seen from this figure that over 11 hours of simulation the parameters predict the system can keep up to the demands of reads and writes to the hard drives provided the CPU has access to an input buffer of 180 Megabytes. The expectation of output hits a maximum of 40 Megabytes of read requests. Here we have assumed the read request size is an average of 1/100 of the actual output data stream.

The model makes certain assumptions. It first fills the cache with incoming data. It then decides how much data to input by checking that it does not overflow the cache. The model also calculates the maximum amount of data that the CPU can write to the cache as $\mu_w, f_w, \Delta t$. The model puts into the cache whatever is less, the amount to fill the cache or the maximum it can write in the time allotted. After the data during that time step has been handled, the CPU then loads the incoming data requests into the cache, again after checking that the cache doesn't overload and deciding on the least value of either how much data it has to handle as read requests, filling the buffer or how much data the CPU
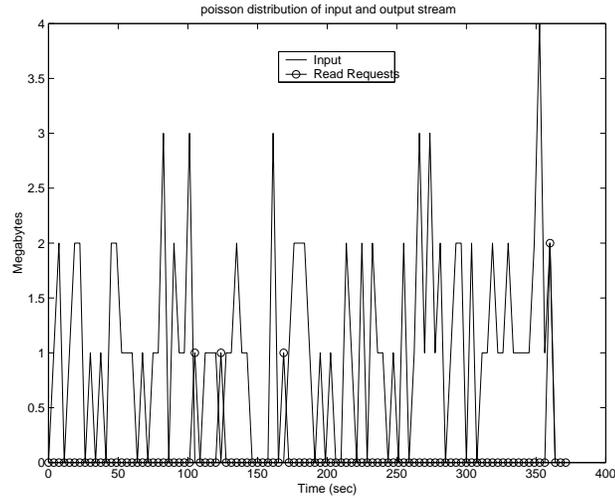
Figure 5.2: The distribution of data being input to the CPU and the "output" stream as read requests.

can handle, namely $\mu_r, f_r, \Delta t$.

After the CPU has loaded the buffer as much as it can, the model switches its attention to the controller. The controller first empties data from the cache and writes it to the hard drive. The model decides on the least value between the amount of write data in the cache and the maximum amount of data it can possible handle in the time step, namely $\gamma_w, g_w, \Delta t$. Once the controller has tried to clear the cache as much as it can it then sends data to the CPU. Here the model has the controller decide between how much it has in read requests in the cache to process and the maximum amount of read requests it can handle. The model then also takes into account the maximum amount of data that the CPU can output and decides between all three to determine the amount of data to put in the output stream.

The CPU's input buffer is reduced by the amount of data sent to the cache and the amount of data expected is lowered by the amount of data in the output stream.
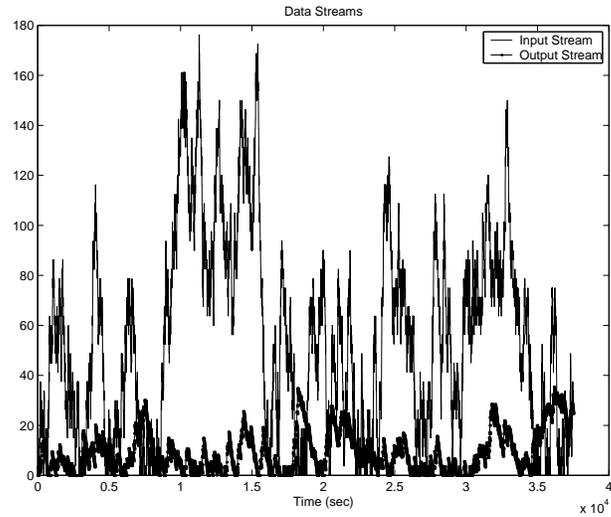
Figure 5.3: The states of the controller buffer.
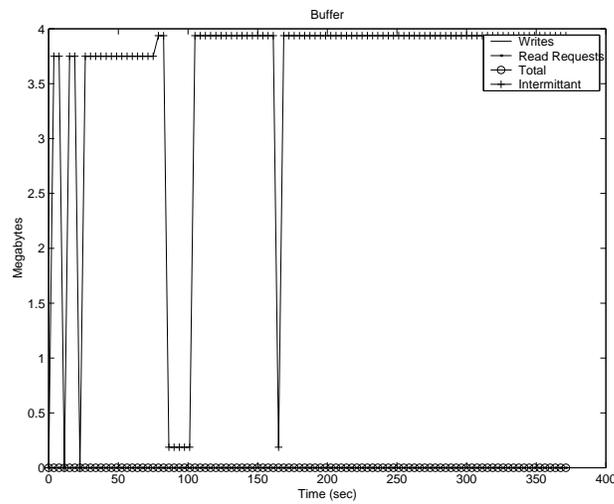
Figure 5.4: The expected amounts of input and output data.

Bibliography

[1] Greetings from the Online Encyclopedia of Integer Sequences. http://www.research.att.com/cgi-bin/access.cgi/as/njas/sequences/eishis.cgi.

[2] Glenn Hurlbert. On universal cycles for k-subsets of an n-set. *SIAM*, 7(4):598–604, 1994.

[3] Se H. Oh and James C. Cavendish. Transients of monolitic catalytic converter: Response to step changes in feedstream temperature as related to controlling automobile emissions. *Ind. Eng. Chem. Prod. Res. Dev.*, 21(1):29–37, 1982.

List of Participants

Organising Committee

Keith Promislow	Simon Fraser University
Mary Catherine Kropinski	Simon Fraser University
Sadika Jungic	Simon Fraser University
Lindsay Hughes	Pacific Institute for the Mathematical Sciences

Mentors

Rachel Kuske	University of Minnesota
Colin Please	University of Southampton
David Ross	Eastman Kodak
Donald Schwendeman	Rensselaer Polytechnic Institute
Brett Stevens	IBM

Students

Nicola Costanzino	Simon Fraser University
Ronald Haynes	Simon Fraser University
Tomasz Janiewicz	Simon Fraser University
Leevan Ling	Simon Fraser University
Bruce Rout	Simon Fraser University
Liu Tang	Simon Fraser University
John Frederick Williams	Simon Fraser University
Ibrahim Agyemang	University of Alberta
Calin Anton	University of Alberta
Irina Dinu	University of Alberta
Jung Min Lee	University of Alberta
Leonid Mocofan	University of Alberta
Dominique Noel	University of Alberta
Cristina Popescu	University of Alberta
Kyle Biswanger	University of British Columbia
Matthew Bolton	University of British Columbia
Lloyd Bridge	University of British Columbia
Margaret Liang	University of British Columbia
Amir Sepasi	University of British Columbia
Maurice Shevalier	University of Calgary
C. Sean Bohun	University of Victoria
Samantha Carruthers	University of Victoria
Shafiqul Islam	Concordia University
Mohammad Oskoorouchi	McGill University
Sirod Sirisup	McGill University
Zhihui Xue	McMaster University
Boting Yang	Memorial University of Newfoundland
Daniel Spirn	New York University
Rozita Dara	University of Guelph
Abid Malik	University of Guelph
Lila Rasekh	University of Guelph
Paul Buskell	University of Manitoba
Maikel Sianturi	University of Manitoba
Nathan Krislock	University of Regina
Tzvetalin Vassilev	University of Saskatchewan
Cristian Ivanescu	University of Toronto
Simal Saujani	University of Toronto
Anamaria Savu	University of Toronto
Fridolin Ting	University of Toronto
Zhiduo Zhao	University of Western Ontario
Paule Ecimovic	York University

PIMS Contact Information

email: pims@pims.math.ca

http://www.pims.math.ca

- **Director: N. Ghoussoub**
  Admin. Asst: J. Burian
  Phone: 604-822-3922
  Fax: 604-822-0883
  email: directorpims.math.ca

- **SFU-Site Director: R. Russell**
  email: sfu@pims.math.ca

- **UAlberta-Site Director: B. Moodie**
  email: ua@pims.math.ca

- **UBC-Site Director: D. Rolfsen**
  email: ubc@pims.math.ca

- **UCalgary-Site Director: M. Lamoureux**
  email: uc@pims.math.ca

- **UVic-Site Director: F. Diacu**
  email: uvic@pims.math.ca

- **UWashington-Site Director: T. Toro**
  email: uw@pims.math.ca