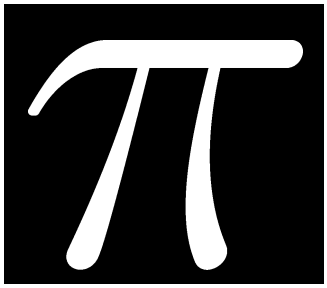

Pacific Institute for the Mathematical Sciences

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

Proceedings of the fourth

PIMS Industrial Problem Solving Workshop

PIMS IPSW 4

Co-sponsored by:

The Natural Science and Engineering Research Council of Canada

and

The Alberta Science and Research Authority

Editor: J. Macki, University of Alberta

Proceedings of the Fourth Annual PIMS Industrial Problem Solving Workshop

Editor: J. Macki, University of Alberta

Co-sponsored by:
The Natural Sciences and Engineering Research Council of Canada
The Alberta Science and Research Authority

June, 2000

Foreword by the PIMS Director

The Fourth Annual PIMS Industrial Problem Solving Workshop was hosted by the Department of Mathematical Sciences at the University of Alberta in Edmonton, May 29-June 2, 2000. For a full week more than 70 participants worked intensely on five problems posed by industrial companies from across North America.

Five problems were looked at this year. These came from *Stern Stewart and Co. (New York), McMillan-McGee Corp. (Calgary) and NEOppg Corp. (Houston), Imperial Oil Resources (Calgary), VisionSmart (Edmonton)* and *Michelin Tire.*

PIMS looks forward to the next Industrial Problem Solving Workshop which will happen at the University of Washington in Seattle next year.

Special thanks go to Jack Macki from the University of Alberta who edited these proceedings. I would also like to thank Marc Paulhus (University of Calgary) and the mentors Jeff deWynne (Oxford University), Dave Ross (Kodak Corporation, Rochester, NY), Huaxiong Huang (York University), Chris Bose (University of Victoria), John King (University of Nottingham) and Michael Lamoureux (University of Calgary).

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

Contents

List of Figures

List of Tables

Editor's Preface

From May 29 to June 3, 2000, approximately 68 graduate students and faculty members from universities in Canada, the United States, the United Kingdom worked intensely for a week with 11 representatives from industry on a set of industrial problems. This was the Fourth Pims Industrial Problem Solving Workshop, following successful workshops at the Universities of British Columbia, Victoria and Calgary. The five problems were selected from a large group solicited by Marc Paulhus, Industrial Facilitator for Pims based at the University of Calgary. The steering committee chose these four because they were tough problems which offered a chance of allowing major progress under a concentrated group effort. In all cases the committee's judgment was confirmed.

Format of the Workshop

The format of the workshop followed the by now standard formula, originally developed by the Oxford Study Group. The first day was devoted to presentation of the 5 problems in one-hour lecture/demonstrations. These presentations were made by the industrial representatives. In the afternoon, after the presentations were completed, the group split into 5 working groups, one for each problem. Each group had their own room for the next 4 days, and worked intensely on their chosen problem, aided by a designated facilitator. The duty of each facilitators was to keep the group focussed and resilient, and all of them did well. On the final day of the workshop, a designated speaker from the group presented their results to the assembled workshop participants, and each industrial representative responded with an evaluation of the progress made. Perhaps the best evidence of the value of the work done was the comment by one industrial representative that the results obtained should result in an annual saving of $500,000 (U.S.!). The workshop wrapped up with the Pimslips award (see further on) and a pizza party.

Acknowledgments

- The biggest thanks go to Marc Paulhus, who contacted companies and organizations, did preliminary evaluations, and guided us all through the complicated business of running a workshop.

- A big thanks to PIMS Director Nassif Ghoussoub for his visionary leadership, and local PIMS Director Bryant Moodie for his strong support. and Rex Westbrook of the University of Calgary, and Huaxiong Huang of York University were the steering committee and offered endless hours of their time and energy. Thank you.

- The 5 report writers devoted a staggering amount of their own time to putting the results of the workshop into a readable form. They were:

    - Huaxion Huang: Imperial Oil's "Many Wells" Problem;
    - Michael Lamoureux: The Michelin Tire Problem;
    - Kell Cheng: The Visionsmart Tennis Ball Problem;

- Tom Cottrell: The Stern-Stewart Delayed Incentives Problem;
- C. Sean Bohun: The McMillan-McGee Oil Recovery Problem.

- Lena Wang and Martine Bareil served as administrative assistants for the workshop. In fact, they carried out most of the behind the scenes work, and did a wonderful job under considerable duress.

Jack Macki, Editor,
Department of Mathematical Sciences,
University of Alberta,
Edmonton, Alberta T6G 2G1
jmacki@gpu.srv.ualberta.ca
Fax: (780) 492-6826, Phone: 492-5725

PIMSLIPS

When people are working intensely, they sometimes say things that, in retrospect, are amusing or downright hilarious. Here are some from IPSW4:

- *What I have told you so far is for a world that does not exist.* (Shane Stark)
- *Write it as a variable and we'll treat it as a constant.* (Anon.)
- Someone is writing on the board. Q. *Is that a 't' ?* Ans. *No, it's a 'j', 'j ' is 't'.* (Anon.)
- Glynis Carling is at the board, defining certain constants, T and E.
  Glynis: *Teei, Eeei, Oh.*
- *This is the three well problem ... well, well, well.* (Glynis)
- *Almost the same idea, but without the math ....* (Mufeed Mahmoud)
- In the course of work at the board, the expression $\frac{rU}{18QT\pi}$ came up.
  That's the Bill Clinton expression: 'Are you over 18, cutie pie?'. (Dave Ross)
- *If we use the prime 21 ... that's not a very good prime.*.* (Mike Lamoureux).
- *No! We are not being reckless, we are being wise!.* (Ranga Sreenivasan)
- Bruce McGee is leaving the board, Andrea is at the board, the number 3.5 is written on the board. Andrea: *Is there a point here?* Bruce: *No, I have no point.*
  Andrea: *No, is this 3 point 5?* Bruce: *Oh, yes, there is a point.*
- *Look, you cannot make a mathematical conclusion on my picture!* (Anon.)
- In the course of work at the board, the expression dv/dz came up many times. Canadians were saying *deevee deezed* and Americans were saying *devee dezee.*
  Dave Ross (American), fitting in: *dee veed dee zed.*
- *If all the f ... ing numbers were prime, this would be perfect.* (Miro)
- Things got a little testy in the oil recovery group. Glynis (from Imperial Oil) was trying to keep the group aware of how much steam is realistically available:
  Huaxiong Huang: *Why can't we solve this? ... because Shane is not willing to relax the constraints.*
  Shane Stark (also from Imperial Oil): *It's Glynis who won't let you have more.*
- *You're sounding like a reservoir engineer: 'I'll define the process and you solve it.'* (Shane)

∗ Winner of the PIMSLIPS award (a bottle of Pims Champagne).

The Problems and Their Industrial Presenters

- **Stern Stewart and Co. (New York)**
  Incentive-Alignment Contracts in a Principal-Agent Setting in the Presence of Real Options
  Consider a firm with shareholders (considered to be the Principal, as one entity) and management (the Agent). The Agent is compensated based on real increases in value. Should the Principal accept low increases now, based on estimates of considerable increases in the future (perhaps after the term of present management). Should the Principal insist on the highest possible short term increase, and compensate or punish the Agent accordingly? Given that the Principal's knowledge will not in general be as complete as that of the Agent, and that there is uncertainty involved at several levels, what is the best strategy for the Principal and the Agent to follow?

  Presenters: Dan Calistrate, Tom Cottrell, and Gordon Sick.
  Mentor: Professor Jeff deWynne, Oxford University.

- **McMillan-McGee Corp. (Calgary) and NEOppg Corp. (Houston)**
  Wellbore Modelling: Boundary Value Problems in the Recovery of Petroleum Fluids from an Oil Reservoir
  Placing a heater deep down in an oil well can decrease viscosity and increase production from the well. To accomplish this, an induction heater is lowered into the well. The problem is to find the steady state temperature distribution in the well, and thus to optimally position the heater and use it efficiently.

  Presenter: Bruce McGee
  Mentor: Dave Ross, Kodak Corporation (Rochester, N.Y.)

- **Imperial Oil Resources (Calgary)**
  Complex System Modelling: Application to Imperial Oil's Cold Lake Oil Sands Facilities
  At Cold Lake, Alberta, Imperial Oil uses a cyclic steam stimulation process to produce heavy oil from oil sands formations. Oil, water and gas are produced; the water and gas are recycled into the steam and oil treatment processes. There are 3200 wells still active in their field, so the problem is highly complex. The goal is to optimize the performance of this system, taking into account the interdependencies in the system and that there are varying time delays in the process.

  Presenters: Glynis Carling, Shane Stark (Imperial Oil Resources, Calgary)
  Mentor: Huaxiong Huang, York University

- **VisionSmart (Edmonton)**
  The Tennis Ball Problem
  Imagine that you wish to track the location of an object in space. As always in the real world, you are short of resources, so you choose to use just two cameras. Each camera produces a two-dimensional image, and these images are in fact composed of discrete pixels, say 640×480. The cameras watch the tennis ball (or missile) as it leaves its launcher and each acquires a series of frames at discrete time intervals (say 1/120 of a second for a ball at 100 mph). What the observer then has in hand is a series of images which *suggest* the ultimate trajectory of the tennis ball. The problems are:

  1. How do we account for the inevitable distortions in the cameras?
  2. How do we use our imperfect discrete date to best estimate the actual path?
  3. Can we take into account the spin of the ball?

  Pesenters: Dan Kenworth, Wolfgang (Visionsmart, Edmonton)
  Mentors: Chris Bose (University of Victoria); John King, University of Nottingham

- **Michelin Tire**
  General Statistical Design of Experimental Problem for Harmonics
  Tires are tested by subjecting them to a variety of forces and storing the response as a periodic wave form. Harmonic components of these waveforms are related to tire performance characteristics, such as noise and comfort, hence the control and reduction of the amplitude of certain harmonics is an important facet of design. Technicians may want to peform designed experiments during production to better understand how various processes and tire components affect the resulting harmonics. The problem deals with a specific technique for testing tires in production, more specifically, how this technique can be used with maximum efficiency.

  Presenter: Bill Mawby (Michelin, North Carolina)
  Mentor: Michael Lamoureux , University of Calgary.

Chapter 1

Optimal Strategy for Imperial Oil's Cold Lake Facilities

Brian Alspach,[1] Calin Anton,[2] Amos Ben-Zvi, [3] Kyle Biswanger,[4] Yonqiang Cao,[5] Glynis Carling,[6] Huaxiong Huang,[5] Dong Liang,[5] Margaret Liang,[4] Mufeed Mahmoud,[7] Jim Muldowney,[2] Reza Naserasr,[8] Marc Paulhus,[9] Cristina Popescu,[2] Miro Powojowski,[9] Bruce Rout,[8] Nikhil Shah,[10] Shane Stark,[6] Tzvetalin Vassilev[11]

Report prepared by Huaxiong Huang[12]

1.1 Introduction

Imperial Oil is Canada's largest producer of crude oil and a major producer of natural gas. It is also the largest refiner and marketer of petroleum products – sold primarily under the Esso brand – with a coast-to-coast supply network and a major supplier of petrochemicals. The Resources Division manages Imperial's natural-resource operations and is located in Calgary, Alberta. It is a major developer of Canada's vast reserves of oil sands through its operation at Cold Lake, Alberta, and its participation in Syncrude Canada.

At Cold Lake, Alberta, Imperial Oil uses a cyclic steam stimulation process to produce heavy oil from oil sands formations. High pressure steam is generated at central plant facilities. The steam is distributed through a pipeline system and injected into the reservoir at wells located

[1]University of Regina
[2]University of Alberta
[3]Queens University
[4]University of British Columbia
[5]York University
[6]Imperial Oil Resources, Calgary
[7]University of Western Ontario
[8]Simon Fraser University
[9]University of Calgary
[10]University of Waterloo
[11]University of Saskatchewan
[12]hhuang@yorku.ca

at a distance from the central plant facilities. Steam injection continues until the oil viscosity is such that the oil can be pumped to surface.

Oil, water and gas are produced during the production part of the cycle and are returned to the central plant facilities. Produced water is mixed with make-up water and is treated and reused as feed water for the steam generators. The produced gas supplements the purchased natural gas used as fuel for the steam generators. The produced oil is processed to remove sand and water. It is then diluted with a lighter hydrocarbon ("diluent") to meet pipeline viscosity specifications. The diluted bitumen is sold to refiners.

There are three central plant facilities at Cold Lake. A total of 3500 wells have been drilled to date of which 3200 are still active. Each well is associated with a single plant. The older wells have completed up to 10 cycles of steaming and production. New wells are periodically added. The duration of the steaming and production phases depends on the age of the well. The steaming phase lasts from 28 to more than 150 days. The production cycle lasts from 100 to 1600 days.

An important issue related to the operation of the Cold Lake system is whether the performance of the overall system can be optimized. Given the vast scale of the operation and the complicated interdependencies of the system, finding an optimal strategy for the entire operation seems to be an impossible task at the first glance. It was decided by the group during the workshop that solving a scaled-down version of the problem would be more productive. It was argued that a problem with many fewer wells could help the group to understand the full problem better and to identify suitable mathematical models.

The first part of this report summarizes the discussions and the models proposed for a four-well problem. With the help of our industrial participants from Imperial Oil, we have made suitable assumptions for the operation conditions and constraints and formulated the problem using several approaches. Results were obtained for some of the models. Because the models are nonlinear, one of the major difficulties associated with these models is that it becomes very expensive computationally to find the solution when the size of the problem grows. In other words, it becomes impractical when the number of wells reaches the level of the Cold Lake system.

The second part of the report discusses a new approach which was developed after the workshop by a smaller group of people. It turns out that the problem can be formulated as a linear programming (or linear optimization) problem in which the size of the problem is not related to the number of wells in the summary as in the non-linear approach, at least under special circumstances. Both continuous time formulation and the discrete time linear programming (LP) approach are discussed. The LP formulation can be viewed as a special discretization (maximizing sequence) of the continuous problem. Results using the LP approach are presented for a test case with 3200 wells at the beginning and another 6000 new wells added during the operation of the system, which is comparable to the size of the Cold Lake system. A number of simplifications are made for the test case. For example, we assumed that all the wells are the same and the performance of the wells is known. We have also assumed that the price of the oil and interest rate are given.

The organization of the rest of this report is as follows. In Section 2 a detailed description of the Cold Lake system is given, which is followed in Section 3 by the nonlinear models studied during the workshop. In Section 4, we introduce the linear models from the follow-up study

after the workshop. A discussion of possible directions for future study is given in Section 5.

1.2 Problem Description

The Cold Lake facilities can be viewed as a complex system. One of the challenges for the group is to identify the most important features of the system and to formulate the problem mathematically under reasonable assumptions.

After lengthy discussions among the group including the industrial partners, we made the following observations/assumptions.

- The oil production and steam consumption of each well is not known in general. Accurate prediction requires detailed modelling using fluid dynamics and statistical tools. However, on the scale of the Cold Lake operation, it is not unreasonable to model the production and steam consumption by a simple mathematical formula, based on the data from field operations. Typically, the production rate is high in the early stages and decreases as the wells become old.

- The actual cost of developing and maintaining wells during production cycles often varies, depending on many factors. For simplicity, we will use an empirical formula fitted to the field operation data.

- We recognize that wells are in general different. We assume that they can be categorized into different classes: from good wells with high production level and low cost to bad ones with low production.

- The steam is distributed from three plants to wells at various locations through the pipelines. A main physical constraint is the daily available steam. In reality, the steam available to wells also differs from one location to another. However, for simplicity, we will assume that the location factor can be ignored and the only constraint is the total daily available steam.

- The oil recovered from the wells (with water, sand and gas) is 'cleaned up', diluted, transported, and sold. The water is treated and recycled back to the steam plants as part of the water supply needed for generating the steam. There is a limit on the amount of water that can be used at a given time. However, this constraint is completely ignored. The price of the crude oil is obviously volatile. However, we will ignore the stochastic nature of the oil price and treat it as a constant.

- Another physical constraint is the daily oil production for the entire system (an upper bound) due to the ability in transportation and treatments. The constraint on each individual well is location dependent. However, we will ignore this and only impose a constraint on the total production level.

- Economic consideration provides a lower bound on the production level, which likely becomes important only towards the end of the entire operation. Again, a constraint will only be placed on the total production level.

- Finally, the unit lifting cost (ULC), defined as cost divided by the oil production is of practical concern, which appears as a constraint. We believe that if the right objective function is chosen, this constraint is likely satisfied automatically.

In light of the observations from field operation data and assumptions made above, we now introduce the daily production function for a typical well

$$\mathcal{P}(t) = 18e^{-0.0003t},$$

where $t \geq 0$ is the time (measured in days since the beginning of the operation of the well). Production is measured in m^3/day. Note that the production is averaged over the production and steam cycles of the well. In reality, a well only produces oil when it is in the production cycle. The averaged daily steam rate can be expressed as a function of time as

$$\mathcal{S}(t) = 115(t+1)^{-0.18} \tag{1.2}$$

m^3/day. Again, in reality, the steam is only required during the steam cycles. The averaged cost of maintaining normal production of a well is given as

$$\mathcal{C}(t) = 601(t+1)^{-0.105} \tag{1.3}$$

in dollars. All three are monotonically decreasing functions. Finally the cost of developing a new well is assumed to be a constant D_0. In addition, we also assume that the discount rate is given as a constant. With these assumptions we now formulate a four-well problem.

1.3 The Four-Well Problem

We consider a small oil field with only four wells, two existing wells and two new wells to be drilled. The question is when we should developed the remaining two wells so that some objective function can be maximized, subject to the constraints on the production level, unit lifting cost, and the steam consumption rate.

Three approaches (two of them related) which were discussed and formulated during the workshop will be presented here. We will discuss the semi-continuous approach first, followed by the discrete non-linear programming and the depth-first (dynamic programming) approaches.

1.3.1 Semi-continuous approach

The participants working on this approach are *Yonqiang Cao, Glynis Carling, Huaxiong Huang, Dong Liang, Margaret Liang, Jim Muldowney, Marc Paulhus and Shane Stark.*

We first assuming that all four wells are of the same class and a well can be terminated only once. Cases with different classes of wells can be treated similarly. The problem can be formulated as an optimization problem with the profit functional defined as

$$J(\vec{T}, \vec{E}) = \int_0^{T_\infty} \sum_{j=1}^{4} \left(p_0\mathcal{P}(t-T_j) - \mathcal{C}(t-T_j) - D_0\delta(t-T_j)\right) H(E_j - t)e^{-Rt}dt$$

where p_0 is the crude oil price. T_j and E_j are the starting and terminal time of well j, for $j = 1, 2, 3, 4$, or in vector form we have $\vec{T} = [T_1, ..., T_4]$ and $\vec{E} = [E_1, ..., E_4]$. $H(x)$ is the Heaviside function, i.e., $H(x) = 1$ when $x \geq 0$ and $H(x) = 0$ when $x < 0$. $\delta(x)$ is the Dirac delta function. It has the following properties: $\delta(x) = 0$ when $x \neq 0$ and $\int_{-\infty}^{\infty} \delta(x)dx = 1$. The objective is to find $\vec{T}$ and $\vec{E}$ such that the profit J is maximized subject to the following constraints.

1. The first constraint which needs to be satisfied is on the steam consumption,

$$\sum_{j=1}^{4} \mathcal{S}(t - T_j) \leq S^H \tag{1.5}$$

   for $0 \leq t \leq T_\infty$. S^H is the limit on the total daily steam consumption.

2. The constraints on the production level can be written similarly as

$$P^L \leq \sum_{j=1}^{4} \mathcal{P}(t - T_j) \leq P^H \tag{1.6}$$

   for $0 \leq t \leq T_\infty$. Here P^L and P^H are the lower and upper bounds for the total daily production.

3. Finally the constraint on the unit lifting cost is

$$\sum_{j=1}^{4} \mathcal{C}(t - T_j) \leq ULC^H \sum_{j=1}^{4} \mathcal{P}(t - T_j) \tag{1.7}$$

   for $0 \leq t \leq T_\infty$. Here ULC^H is a given upper bound.

Since the functions $\mathcal{P}$, $\mathcal{S}$ and $\mathcal{C}$ are monotonic, it is sufficient to enforce constraints before and after each T_i as

$$\begin{aligned}
\sum_{j<i} \mathcal{S}(T_i - T_j) &\leq S^H, \\
\sum_{j<i} \mathcal{P}(T_i - T_j) &\leq P^H, \\
\sum_{j\leq i} \mathcal{P}(T_i - T_j) &\geq P^L, \\
\sum_{j<i} \big(\mathcal{C}(T_i - T_j) - ULC^H \mathcal{P}(T_i - T_j)\big) &\leq 0, \quad \text{for} \quad i = 1, ..., 4.
\end{aligned} \tag{1.8}$$

The problem with two existing wells and two new wells was attempted using the Matlab Optimization toolbox and Monte-Carlo simulation. Matlab uses some standard methods for solving nonlinear optimization problems with constraints. The Monte-Carlo simulation is done

by generating the starting and ending time of each new well randomly during a given time interval using a uniform distribution random number generator. The profit associated with each realization is computed accordingly until one of the constraints is violated. Surprisingly (or perhaps not surprisingly), Matlab failed to generate any results. The Monte-Carlo simulation with 10,000 trials, on the other hand, takes less than a minute on a Dell laptop based on a 366 MHz Intel II CPU with 128M RAM.

1.3.2 Non-linear programming approach

The participant working on this approach is *Kyle Biswanger*.

We attempted to represent this problem as a linear program. In this formulation the problem is represented as a set of linear constraints, which form an n-dimensional polyhedron, and an objective function, which is optimized over the feasible region defined by the polyhedron. We can define the objective function in many ways, in this case we choose to maximize net present value. The feasible region is defined by two classes of constraints

1. those that are related to set limitations of the plant such as steam availability and daily operating cost; and
2. those that ensure the model behaves appropriately, for example ensuring that a well must start producing oil before it can be terminated.

Formally there is no distinction made between the constraints, however the distinction can help us understand the formulation. In defining the behavioral constraints, we encounter non-linearities. We have chosen to keep the non-linearities, and represent the problem as a non-linear program instead of a linear one. Non-linear programs are in general more difficult to solve. The formulation of the nonlinear program is achieve the objective

$$\max \sum_j \sum_i \frac{(Px_{ij} - c_{ij} - \psi_j)}{(1+r)^j} \tag{1.9}$$

subject to the following plant constraints

$$\sum_i x_{ij} \le Q_H^x, \quad \sum_i x_{ij} \ge Q_L^x, \quad \sum_i s_{ij} \le Q_H^s, \quad \sum_i \frac{c_{ij}}{x_{ij}} \le Q_H^c, \tag{1.10}$$

and the behavioral constraints

$$\begin{aligned} &\sum_j T_{ij} = 1, \quad \sum_j E_{ij} = 1, \quad \sum_{j=1}^{q} T_{ij} \ge \sum_{j=1}^{q+1} E_{ij}; \\ &q \in [1, m], \quad T_{ij} \ge 0, \quad E_{ij} \ge 0, \quad T_{ij}, E_{ij} \in [0, 1], \end{aligned} \tag{1.11}$$

where

$$\begin{aligned} x_{ij} &= \gamma_{ij} \sum_{k=1}^{m} A_k^x T_{i(j+1-k)}, \quad s_{ij} = \gamma_{ij} \sum_{k=1}^{m} A_k^s T_{i(j+1-k)}, \\ c_{ij} &= \gamma_{ij} \sum_{k=1}^{m} A_k^c T_{i(j+1-k)}, \quad \gamma_{ij} = 1 - \sum_{k=1}^{j} E_{ik}. \end{aligned}$$

In this formulation, the start times (T_{ij}) and end times (E_{ij}) are the decision variables. The definitions give relations, and are actually equality constraints on the problem. However, they relate the decision variables to variables that appear in the objective function, and in this sense are really nothing more than definitions. It is worth noting that when time is discretised, this approach is essentially the same as the continuous time model discussed earlier.

1.3.3 Depth-first approach

The participants working on this approach are

Brian Alspach, Calin Anton, Amos Ben-Zvi, Mufeed Mahmoud, Reza Naserasr, Cristina Popescu, Bruce Rout, and Tzvetalin Vassilev.

Again, we will explore the case of four wells, two of them working (one is in cycle nine and other in cycle nineteen of their producing life) and two to be drilled.

First we will make some definitions. We will define each well w_i as an ordered pair $w_i(c, y)$, where $c \in \mathcal{Z}$ is the class of the well (1-good, 2-medium, 3-poor) and $y \in \mathcal{Z}$ is the cycle of production in which the well is currently, $y = 0, 1, ..., 25$ (0 meaning that the well has not been developed yet and 25 meaning that the well has been abandoned, any other odd cycle number is a steaming cycle and any other even cycle number is a production cycle).

Let W be the set of all possible combinations (c_i, y_j), $i = 1, 2, 3$, and $j = 0, 1, ..., 25$, which is a 3×26 set. Further, we will define $\Lambda = \{w_1, ..., w_r \mid w_i \in W, i = 1, ..., r\}$. This must be regarded as the set of all possible combinations over the given set of wells. Then, the allowable set of wells can be defined by $\Lambda^* \subset \Lambda$ such that $\Lambda^* = \{w_1 ..., w_r \mid f_k \geq 0\}$. Here $f_k \geq 0$ represents our constraints, and as above $w_i \in W, i = 1, ..., r$. Let $a, b \subset \Lambda^*$, then we will define binary operation "$*$" on Λ^*: $* : \Lambda^* \times \Lambda^* \to \mathcal{Z}$,

$$a * b = t, \text{where } t = \begin{cases} 0, & \text{if } a = b \\ 1, & \text{if the transition } a \to b \text{ has one coupling line} \\ 2, & \text{if } a \to b \text{ has two coupling lines} \\ \vdots & \\ k, & \text{if } a \to b \text{ has k coupling lines} \end{cases}$$

Having this we can define another binary operation called "$s$uccession". We will say that b succeeds a if there is transition with length one, i.e. if they are immediate neighbours in the manifold. Formally:

$$B_a = \{b \mid b * a = 1\} \quad a, b \in \Lambda^*.$$

The main idea of the approach is to build a manifold of the states and transitions between them using their properties and constraints. Then we will find the best possible path between starting (initial) and target (final) position with respect to given objective function. Final states will be all these states for which we have no transition to an allowable state. For these states it also can be noted that they have no succeeding states, or for a final state z, $B_z = \emptyset$ holds. Building the manifold we will take advantage of the properties of the wells and constraints imposed in a way that will allow us to cut down infeasible transitions and consider only a relatively small number of choices at each step.

Now our problem is formulated as follows: we wish to travel from some state a (start state) to some state p (final point) while optimizing an objective function. At each step we will build the manifold and finally go through all possible paths.

For each well we are given: $\mathcal{C}_i-$ cost of production [\$/day], ULC_i^H- unit lifting cost [$\$/m^3$], $\mathcal{Q}_i^s-$ steam needed [m^3/day], and $\mathcal{Q}_i^x-$ oil produced [m^3/day]. We are also given the cost of drilling a new well, D_0. Initially we have $N = 4$ wells with two of them producing and two to be drilled.

The following constraints are imposed on our model:

1. Steam capacity: steam used (needed) at each moment must be at most as much as our producing capacity: $\mathcal{Q}^s = \sum_i \mathcal{Q}_i^s \leq S^H$ (an upper bound).

2. Production rate: quantity of oil produced at each moment can be within certain limits: $P^L \leq \mathcal{Q}^x \leq P^H$. Here $\mathcal{Q}^x = \sum_i \mathcal{Q}_i^x$ is our total production of oil and P^L and P^H are the bounds.

3. Unit Lifting Cost (ULC) must be less than a certain value, which is the value that gives us a profitable production. It is calculated by summing all production costs and dividing them by the total quantity of oil produced, i.e. $ULC = \sum_i (ULC_i \mathcal{Q}_i^x / \mathcal{Q}^x)$ and $ULC \leq ULC^H$.

An issue that simplifies the model is classification of wells. Wells are separated into three different classes according to their producing characteristics, for ease of modelling. Approximate differences between oil productions are: wells of class 1 produce roughly about 20% more oil than these of class 2 per working cycle, and these of class 3 produce about 20% less than class 2. We will assume that no more than two wells can be started at the same time.

An additional constraint imposed here is that the maximum number of wells which can be drilled at the same time is five.

With this scheme, we can now build the manifold of all states and find the optimal path on it, using a certain objective function. At each decision moment we need to generate the sequence of actions, which well to be started, developed or abandoned. On the time scale calculated we will consider certain points referred to as "decision moments". Each one of these moments is the point at which we evaluate the objective function and possible ways to act further. These moments can be the ends of steaming or producing cycles of wells that are in operation. We denote them by t_j.

We can choose different objective functions to evaluate and optimize. Comparison between the results obtained can add valuable insight to the solution of the problem. For example, we can choose the unit lifting cost with depreciation as the objective function. If we have an annual interest rate of r, and assume that ULC is the same during the steaming and production periods (costs spread over whole cycle), the following objective function can express our maximal benefit over the period criterion:

$$F = \sum_j \frac{365 C_j}{r} (e^{\frac{-r t_{j+1}}{365}} - e^{\frac{-r t_j}{365}})$$

Method	Total Profit (Million $)	$\vec{T}$ (days)	$\vec{E}$ (days)
MC	9.84	[-5370, -1110, 3164, 13082]	[498, 4347, 10227, 18179]

Table 1.1: Summary of Results for the Four-well Problem.

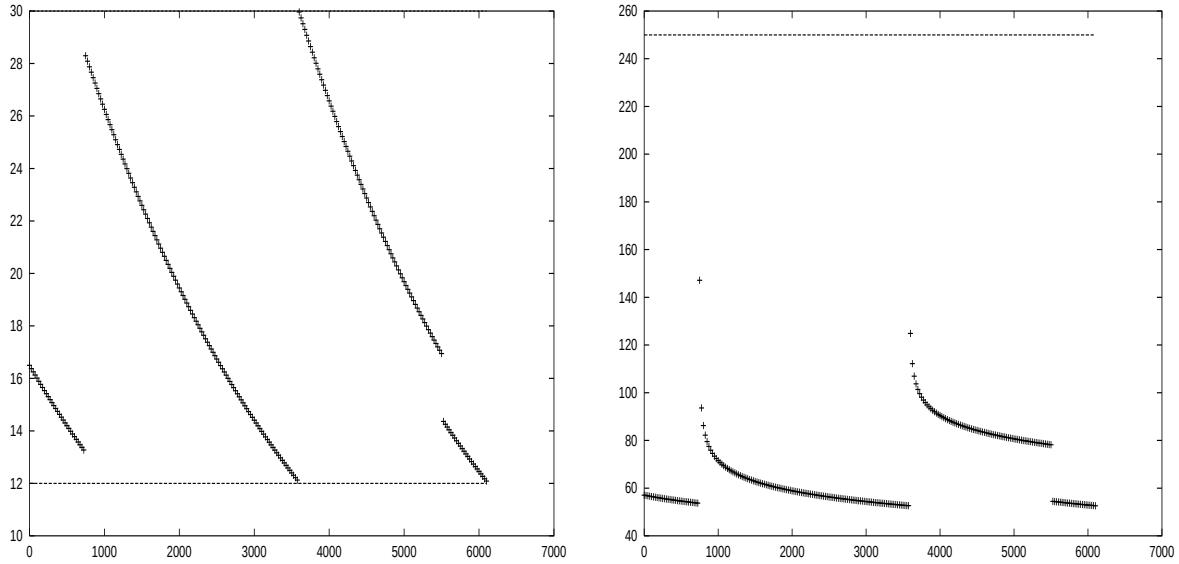

Figure 1.1: Time history of (a) the production level and (b) the steam consumption for the four-well problem.

1.3.4 Results

Attempts were made to solve the problems formulated above by each sub-group. However, due to the extremely limited time available during the workshop, only partial results were obtained.

Table 1.1 summarizes the results using the Monte Carlo simulation with a time step size of 25 days and one million realizations. The upper limit of steam consumption is $S^H = 250m^3$/day, while the daily production level bounds are 12 and 30 m^3. The discount rate is taken as $R = 0.05/365$ annually. The price of the commodity is $p_0 = 200\ \$/m^3$. The cost of developing a new well is $D_0 = \$500,000$ and the unit lifting cost bound is $ULC^H = 50\$/m^3$. The simulation takes about 30 minutes on a Dell 366 MHz laptop.

Figure 1.1 shows the time history of production level and steam consumption. It can be seen that the operation is regulated mostly by the upper bound of the production level for this particular example. We have assumed that the two existing well were drilled 5370 and 1110 days when the decision to drill two new wells are being made. For a copy of the code, please contact Huaxiong Huang [13].

[13] hhuang@yorku.ca

1.4 An Alternative Approach for the Cold Lake System

This part of the report is based on the work of *Huaxiong Huang, Marc Paulhus, Miro Powojowski, and Nikhil Shah.*

The main difficulty of extending some of the approaches studied during the workshop to a large problem is that it becomes very expensive to compute the solution as the number of wells increases. However, it has been observed that unless each well is absolutely unique, it is not necessary to distinguish well i from well j.

It turns out that this difficulty can be avoided if we formulate the problem differently when the wells are not very different from each other, or at least that they can be put into a small number of groups. In the second part of this report, we will discuss this new approach, first by presenting a discrete linear programming (LP) formula. It is followed by a continuous formulation with arbitrarily large number of wells. The discrete LP approach is re-derived as a special discretization of the continuous formula. The discrete problem is solved using an LP package and results are given for a test problem. For simplicity, we assume that all the wells are the same. A relaxation of this assumption will be discussed at the end of this section.

1.4.1 An LP approach

We will formulate the problem in such a way that the number of wells is not relevant to the size of the problem. Rather than solving for the starting and ending times of each individual well we will instead determine the number of wells in an age category any particular time in the future. This will allow us to set the problem up as straight-forward (although very large) linear programming problem that we can solve using a mixed-integer programming package written by Michel Berkelaar [2]. This package is freely available for non-commercial use.

Problem Definition

We will start by describing the problem we intend to solve. It has the following input file which we will describe in detail below.

```
Number of undeveloped wells:    6000
Time step (days):    100
1000
Abandon times (days):    300 500 800 1100 1600 2300 3200 4300 5600 7000
Starting well ages (days):    100 200 300 400 500 600 700 800 1000 2000
3200 4000 5600
Number of wells starting with those ages:    250 250 250 250 250 250 250
250 250 250 250 250 250
Max steam use (day):    120000
Max production rate (day):    42000
Min production rate (day):    9000
Discount rate (annual):    0.05
Price of commodity:    200
Well development cost:    500000
```

First is the number of potential wells that could be developed in this field, $N = 6000$. The "time step" is the decision interval for this project, in this case we will make decisions on the management of this project once every $T = 100$ days. The "total time length" is 11000 days (about 30 years). The times that we can abandon a well (presumably just before we steam it) are given as the vector of "abandon times". The abandon times are roughly the same as the steaming times given to us by Imperial Oil rounded to the nearest 100 days.

The next two vectors define the starting state for the field. The first vector defines the ages of the currently producing wells and the second vector defines the number of current wells of each of those ages. In this case we have 3250 active wells varying in age between 100 and 5600 days.

The "max steam use" $S^H = 12000$ is the maximum amount of steam that can be used on any single day. The "max production rate" $P^H = 42000$ and "min production rate" $P^L = 9000$ are bounds of the daily production rate of the entire project. The "discount rate" $R = 0.05/365$ is used to discount future cash flows. The "price of the commodity" $p_0 = 200$ is the market value of one cubic meter of bitumen and the "well development cost" $D_0 = 500000$ is the cost of drilling a new well.

In our model we are only considering a single class of wells, a so-called *class B* well as described during the workshop. If it has been x days since a well was originally developed, its production level $\mathcal{P}(x)$, steam consumption rate $\mathcal{S}(x)$ and cost function $\mathcal{C}(x)$ are given by (1.2)-(1.3). The effect of including more than one well class will be discussed later.

Method

Note that the time step defines a natural classification for the ages of the wells. If a well is of stage s then it is between $(s-1)T+1$ and sT days old. Define

$$a_s = (s-1)T + 1 \tag{1.12}$$

the minimum age of stage s well. In our example the oldest possible well is of age 7000 days (that is the highest abandon time) thus we have $C = 70$ different stages of well ages.

The time step also defines a discretization of time. Since the "total time length" is 11000 days we have $Y = 1099$ decision intervals over the life of the project. Define

$$d_t = (t-1)T + 1 \tag{1.13}$$

the starting time of decision interval t.

We are interested in solving for the number of wells of stage s in the decision interval t for all (t, s) which maximizes the profit function Equation (1.14). Thus we have 76930 decision variables. The variable labelled $z(t, s)$ is the number of wells of stage s in the t^{th} decision interval. We define $z(0, s)$ to be the starting profile of the field as defined by the input file.

The profit function can now be described as

$$\sum_{t=1}^{Y}\sum_{s=1}^{C} v(t,s)z(t,s) \tag{1.14}$$

where

$$v(t,s) = \sum_{k=0}^{T-1}(\mathcal{P}(a_s+k)p_0 - \mathcal{C}(a_s+k))e^{-R(d_t+k)} - \mathcal{D}(s)e^{-Rd_t} \tag{1.15}$$

which is simply the sum of the gross income (production times price) that a well of stage s will generate over a T day interval minus the operating cost of a well of stage s, all discounted.

$$\mathcal{D}(s) = \begin{cases} D_0, & \text{if } s = 1; \\ 0, & \text{otherwise;} \end{cases} \tag{1.16}$$

is the cost of developing a well, which only needs to be paid when a well is first developed. Note that the function $v(t, s)$ is non-linear but for any fixed pair (t, s) it is simply a constant. Thus Equation (1.14) is linear.

There are steam constraints that must be satisfied over every decision interval. That is for all t

$$\sum_{s=1}^{C} \mathcal{S}(a_s)z(t,s) \le S^H. \tag{1.17}$$

Note that $\mathcal{S}$ is decreasing and thus the constraint only needs to be checked at a_s. Moreover $\mathcal{S}(a_s)$ is a constant for a fixed s and thus Equation (1.17) defines 1099 linear constraints. Similarly we must include 1099 linear constraints for the upper bound on the production rate,

$$\sum_{s=1}^{C} \mathcal{P}(a_s)z(t,s) \le P^H \tag{1.18}$$

and 1099 linear constraints for the lower bound on the production rate,

$$\sum_{s=1}^{C} \mathcal{P}(a_s)z(t,s) \ge P^L. \tag{1.19}$$

There are also transition constraints. If a_{s+1} is an abandon time then for all t

$$z(t+1, s+1) \leq z(t, s) \tag{1.20}$$

since all wells of stage s must either be abandoned or become wells of stage $s+1$ the next decision interval. If a_{s+1} is not an abandon time then Equation (1.20) will be a strict equality. Also, to insure we do not develop more wells than are available, we must have

$$\sum_{t=1}^{T} z(t, 1) \leq N. \tag{1.21}$$

The solution of this problem is given after the discussion of the continuous formulation. The solution procedure is illustrated via a simpler example problem in the Appendix.

Note that for large problems we drop the integer declaration and solve it as an LP problem for real numbers, rounding the answers down to the nearest integer. If the total number of wells is large enough then this solution should be extremely close to the optimal integer solution. Indeed, this procedure relies on the fact that we are dealing with a large number of wells and will not work for a small number. We also note that the decision making time interval T can be altered as well. An interesting question is whether the value of maximum realized objective function will increase as T is reduced and approach an upper bound in the limit of continuous time. Another interesting question is whether allowing wells to be added/removed more freely instead of at specific time will affect the overall solution. These questions could be investigated using the LP model discussed here. However, to re-formulate the problem from the continuous point of view may also provide some useful insights. This is the rationale behind the continuous formulation to be discussed in the following.

1.4.2 A continuous time formulation

A linear optimization problem

We now re-formulate the problem as a linear continuous optimization problem with linear constraints. Define $n(t, a)$ as the number of wells at time t with age a. $n_0(t) = n(t, 0)$ is used to describe the number of new wells being developed at time t. Let $n_b(a) = n(0, a)$ denote the number of wells of age a at the beginning of the operation. We assume that the total number of wells $N = \int_0^{T-a} n_b(a)da + \int_0^{T_\infty} n_0(t)dt$ is given. Here T is the natural life span of wells (assumed to be the same for all wells) and T_∞ is the time of the entire operation of the oil-field.

Note now the production, steam consumption rate and cost are functions of age, $\mathcal{P}(a)$, $\mathcal{S}(a)$ and $\mathcal{C}(a)$, which are simply (1.2) - (1.3). The development cost of a well is given as $\mathcal{D}(a) = D_0\delta(a+)$. The total profit is a functional $J(n)$ defined as

$$J(n) = \int_0^{T_\infty} \int_0^{T} n(t, a)\mathcal{F}(a) \exp(-Rt)dadt \tag{1.22}$$

where

$$\mathcal{F}(a) = p_0\mathcal{P}(a) - \mathcal{C}(a) - \mathcal{D}(a)$$

is the profit of operating one age a well before discount and p_0 is the crude oil price. The objective is to find a function $n(t,a)$ such that $J(n)$ is maximized subject to the following constraints

$$\int_0^T n(t,a)\mathcal{S}(a)da \leq S^H, \tag{1.23}$$

$$\int_0^T n(t,a)\mathcal{P}(a)da \leq P^H, \tag{1.24}$$

$$\int_0^T n(t,a)\mathcal{P}(a)da \geq P^L, \tag{1.25}$$

$$\int_0^T n(t,a)\mathcal{C}(a)da \leq ULC^H \int_0^{T_a} n(t,a)\mathcal{P}(a)da. \tag{1.26}$$

Again S^H is the upper steam rate limit, P^H and P^L are the upper and lower production level, and ULC^H is the unit lifting cost.

Dynamic constraint on $n(t,a)$

We now discuss the dynamic constraint satisfied by $n(t,a)$. We note that the total number of wells is a conserved quantity except when they are shut down or removed from operation. Define $f(t,a) \geq 0$ as the number of wells of age a removed during a unit time interval at time t. Then we must have

$$\frac{\partial n}{\partial t} + \frac{\partial n}{\partial a} = -f(t,a). \tag{1.27}$$

This is a first order hyperbolic equation and it can be easily solved using the method of characteristics

$$n(t,a) = n(0,a-t) - \int_{a-t}^{a+t} f((\eta - a + t)/2, (\eta + a - t)/2)d\eta$$

for $a > t$ and

$$n(t,a) = n(t-a,0) - \int_{t-a}^{a+t} f((\eta - a + t)/2, (\eta + a - t)/2)d\eta$$

for $t \geq a$. Note that $n(t-a,0) = n_0(t-a)$ is unknown and $n(0,a-t) = n_b(a-t)$ is given. It can be observed as well that $n(t,a)$ is non-increasing function and the inequalities

$$n(t,a) < n_0(t-a) \text{ or } n_b(a-t)$$

hold when the wells are being removed along the characteristics, $\xi = a - t =$ constant. Otherwise, $n(t,a)$ is simply $n_0(t-a)$ or $n_b(a-t)$.

Discretization and numerical approximations

One approach for finding the extremum of a functional is to use direct methods. For example, one can use the Ritz method to search for an optimization sequence [1]. We now discuss a special case of the Ritz method where the continuous problem is approximated by a numerical

quadrature. We will show that it can be related with the LP formulation proposed earlier when a special discretization is used.

Suppose that a grid in $(t_r,\ a_s)$ is set up to cover the solution domain $[0,\ T_\infty]\times[0,\ T]$, with $t_r = t_{r-1} + T$, $a_s = a_{s-1} + T$, for $r = 1, 2, ...r_{\max}$ and $s = 1, 2, ..., s_{\max}$. The functional $J(n)$ can be re-written

$$\begin{aligned} J(n) &= \sum_r \sum_s \int_{t_{r-1}}^{t_r} \int_{a_{s-1}}^{a_s} n(t,a)\mathcal{F}(a)\exp(-Rt)dadt \\ &\approx \sum_r \sum_s T^2 \int_{t_{r-1}}^{t_r} n(t_{r-1}, a_{s-1}) \int_{a_{s-1}}^{a_s} \mathcal{F}(a)\exp(-Rt)dadt \\ &= \sum_r \sum_s \int_{t_{r-1}}^{t_r} T^2 n_{r-1,s-1} v_{r-1,s-1} \end{aligned}$$

where the grid function $n_{r,s}$ is the approximation of $n(t,a)$. This is equivalent to (1.14) in the LP formulation if $z(r,s) = T^2 n_{r,s}$ and $v_{r,s}$ is evaluated numerically using the Riemann sum with $\delta t = 1$ day. The constraints can be discretized similarly as

$$\sum_s n_{r,s}\mathcal{S}_s \le S^H, \quad \sum_s n_{r,s}\mathcal{P}_s \le P^H, \quad \sum_s n_{r,s}\mathcal{P}_s \ge P^L, \quad \sum_s n_{r,s}\mathcal{P}_s \le ULC^H \sum_s n_{r,s}\mathcal{C}_s.$$

Finally the dynamic equation can be discretized and replaced by the inequality $n_{r,s} \le n_{r-1,s-1}$ which is equivalent to $z(r+1, s+1) \le z(r,s)$.

Obviously other types of discretization can be used. For example, instead of using a grid function, we can use piecewise polynomials on each subinterval. From a numerical point of view, using grid functions fits into a finite difference framework while using piecewise polynomial approximation is a finite element method. From LP point of view, using grid functions can be interpreted as restrict adding/removing wells at specific times and using piecewise polynomials means that we will drill/abandon wells continuously with some restrictions.

Other related issues are the convergence of the discrete solution to the solution of the continuous problem (if it exists) in general and the speed of convergence associated with a particular discretization. However, addressing these issues properly is beyond the scope of this report. Instead, we now present the results obtained using this special discretization, which is equivalent to the LP formulation discussed earlier.

1.4.3 Solution

To solve the discrete LP problem, we wrote a program that takes an input file which defines a problem and writes it as an lp-file, which was submitted to lp-solver. The whole process takes about 30 minutes on a Pentium 400 machine with 256MB RAM. A copy of the code and the complete solution can be obtained from Marc Paulhus[14]. To summarize, we now have a linear objective function with 76930 decision variables, 1099 linear constraints for the steam requirements, 1099 linear constraints for each of the production bounds, and 76931 linear transition constraints. Remarkably we can solve this system. The optimal solution is an assignment to

[14] marc@ndtechnologies.com

the decision variables which is too much information to present here, but we will attempt to provide a brief glimpse into the optimal strategy for managing this project.

To this end we have provided 12 snapshots of what the field will look like at 1000 day intervals in Figures 1.2 and 1.3. These plots show the profile of the oil field at particular times in the future. The title of the plot indicates the future time displayed in that plot. The x-axis is the age of the wells in the project and the y-axis is the number of wells of each age.

- The *day 100* plot is almost identical to the starting profile we specified in the input file.
- The *day 1000* through *day 4000* plots still show the effect of the starting profile. Most of the wells that were active at the start are still active in 1000 days but nearly all of the starting wells have been abandoned by day 4000.
- The *day 5000* and *day 6000* plots show the middle evolution of the project. All wells in this window are abandoned at age 4300 days.
- The *day 7000* through *day 10900* plots show the winding down of the project. In this phase we allow wells to mature much more than in the earlier phases of the project. After 8300 days we have exhausted our supply of undeveloped wells and thus no new wells can be drilled.

Next we look at the constraints. Figure 1.4 shows the daily production rate over the life of the project and the daily steam requirements over the life of the project. We can see that with this particular input file the steam constraints are much more restricting than the production constraints.

Plotted in Figure 1.5 are the costs and profits over the life of the project. These plots show the daily profits and costs and do not include the one-time costs associated with drilling a new well. The magnitude of those development costs can be determined by looking at Figure 1.6a which shows the number of new wells started at future times.

Also plotted in Figure 1.6 are the number of wells abandoned when they reach the age of 4300 days. In total 4746 wells are abandoned before their life's end. Most of those (4496 to be precise) are abandoned when they reach the age of 4300 days. Finally we show the unit lifting cost as a function of time in Figure 1.7. Note how closely related this graph is to the profit graph (Figure 1.5b) which is to be expected.

A summary of some interesting statistics is presented in Table 1.2.

Before we finish this section, we note that adding more constraints should not be a problem as long as they are linear. If we expand our problem to include different classes of wells then the complexity of the problem increases dramatically. If we have m categories of wells there will be m times as many decision variables and m times as many transition constraints. Nevertheless, given enough time, computing power, and the proper software, this problem should be solvable for a realistic model of the Cold Lake project as long as $m \ll N$ (the total number of wells).[15] There are some very powerful commercial software packages available although we should stress how impressed we are with the power of Berkelaar's solver. See [3] and [4] for examples of good commercial LP solvers or see [5] for a more complete list.

[15] An interesting comparison on the linear and non-linear approaches can be made when $m = N$.

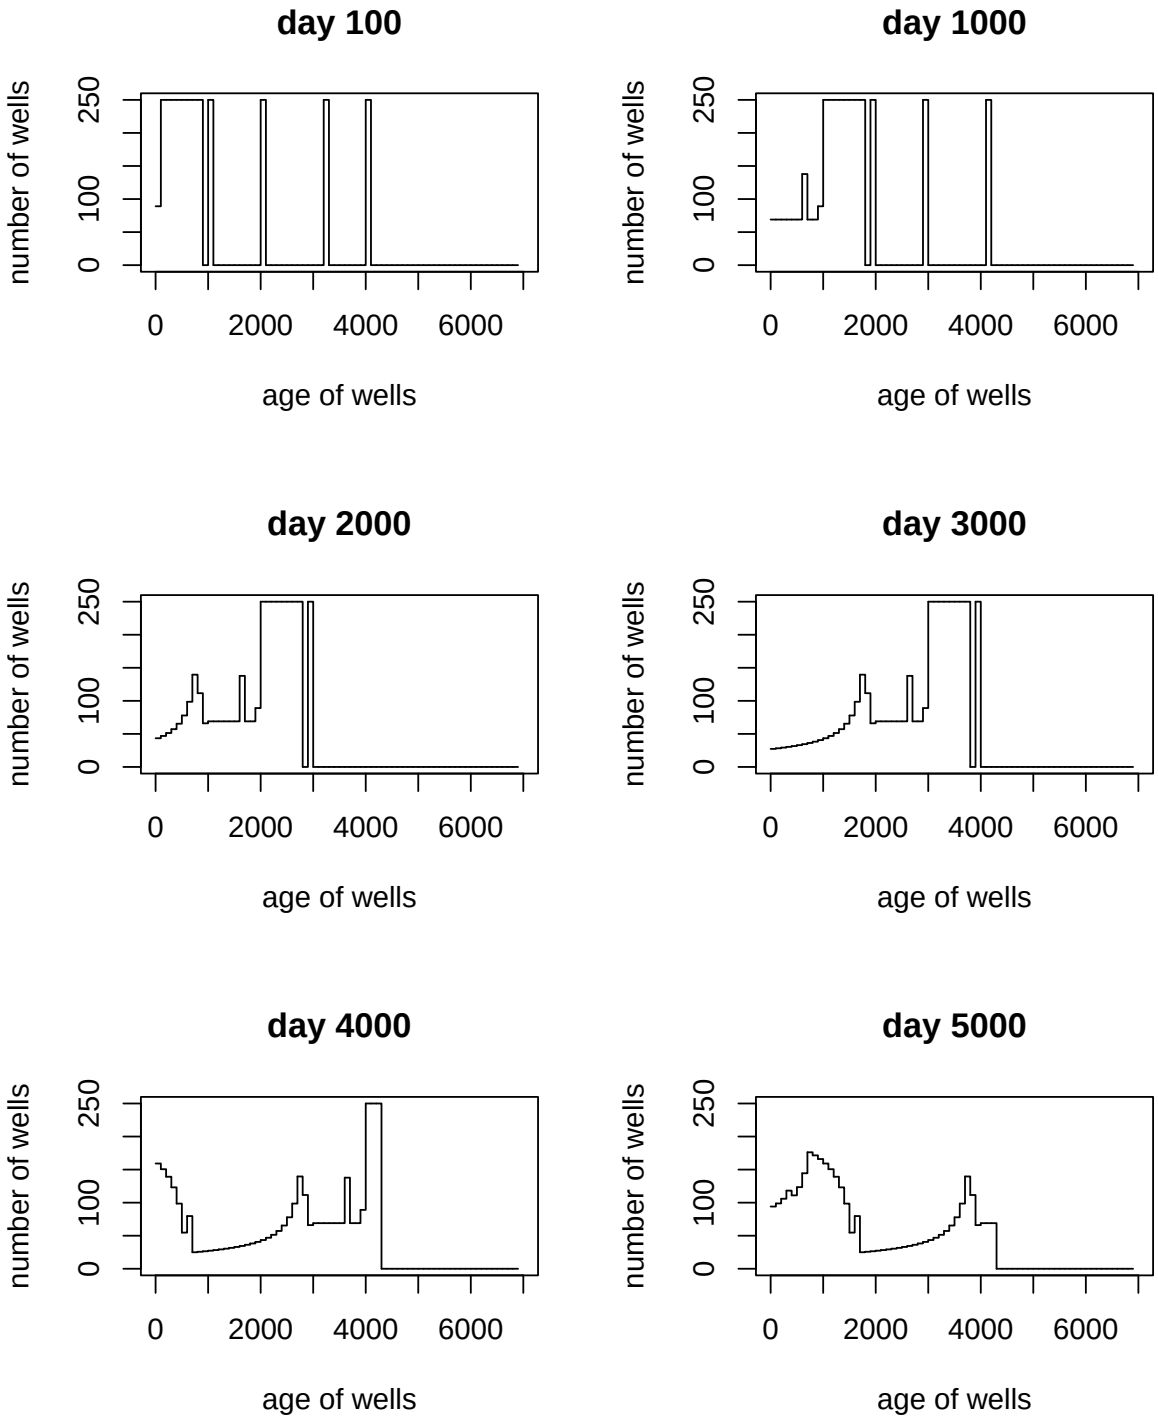

Figure 1.2: Snapshots of the field profile on days 100 through 5000

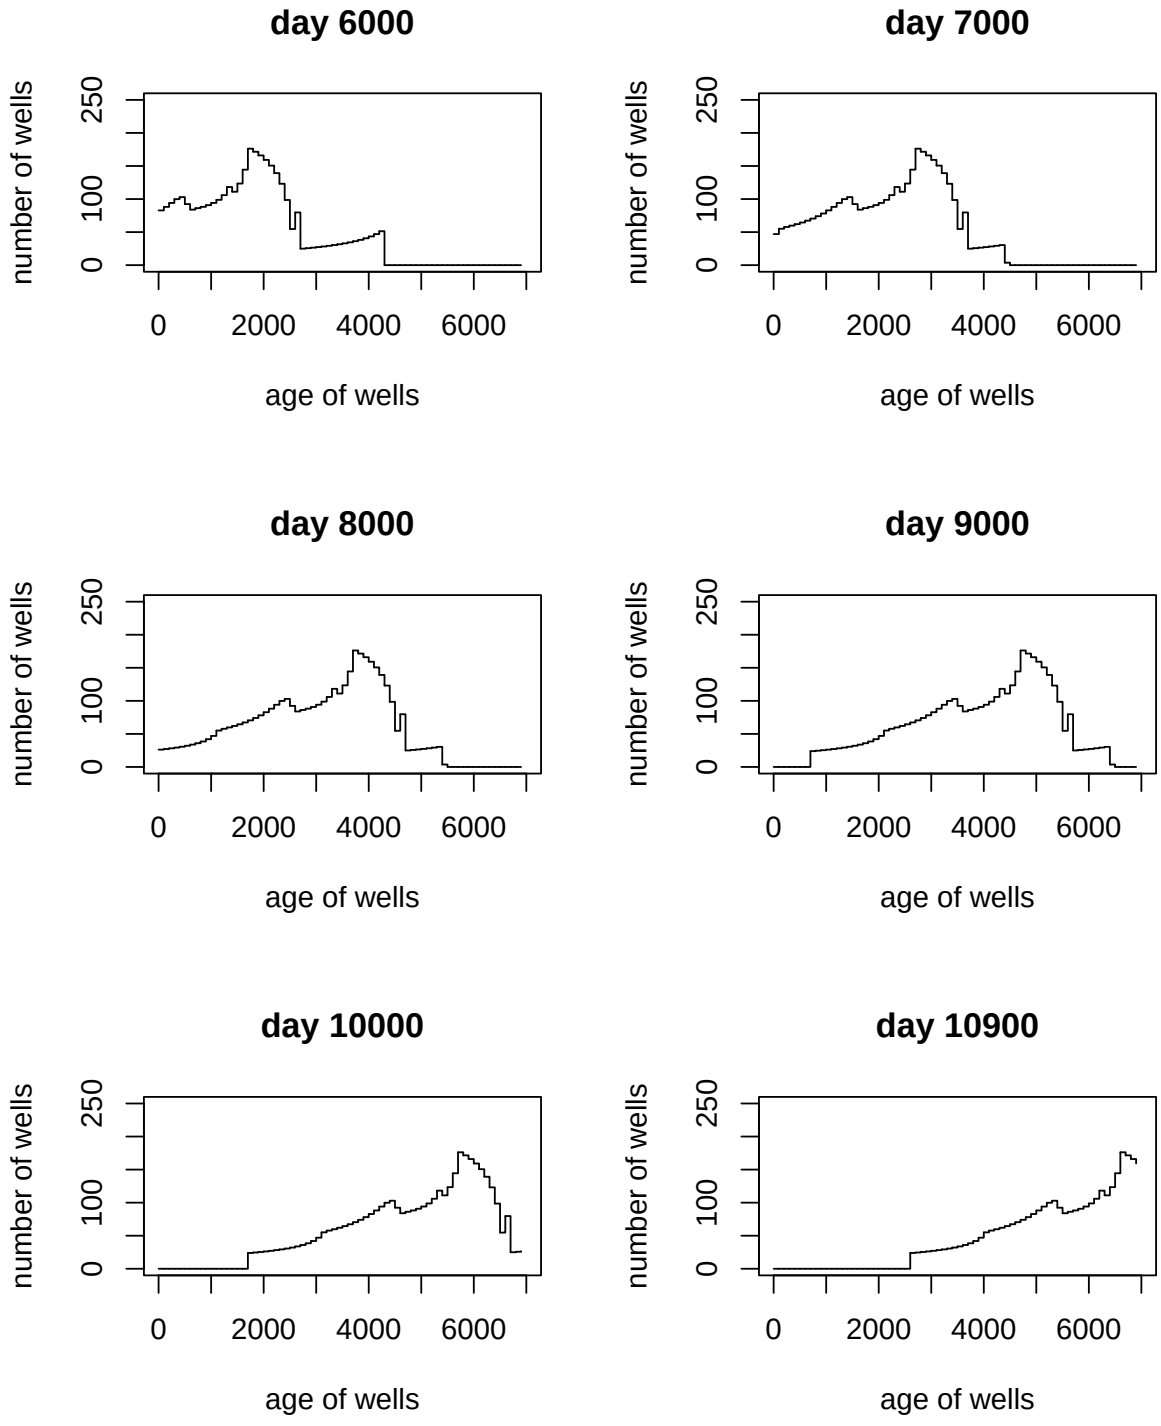

Figure 1.3: Snapshots of the field profile on days 6000 through 10900

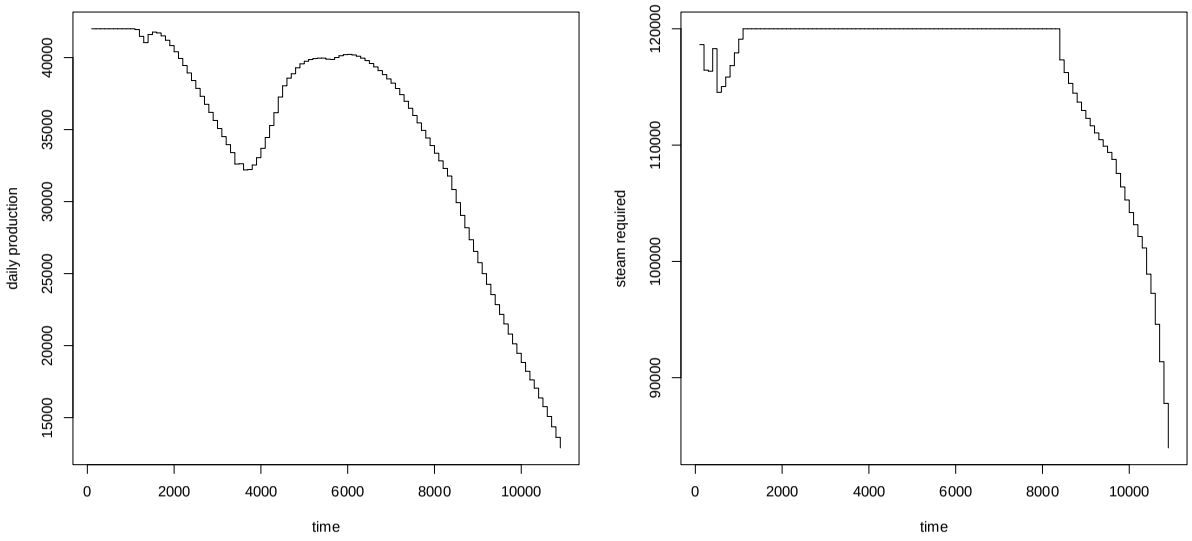

(a) (b)

Figure 1.4: The daily production versus time (a), and daily required steam versus time (b).

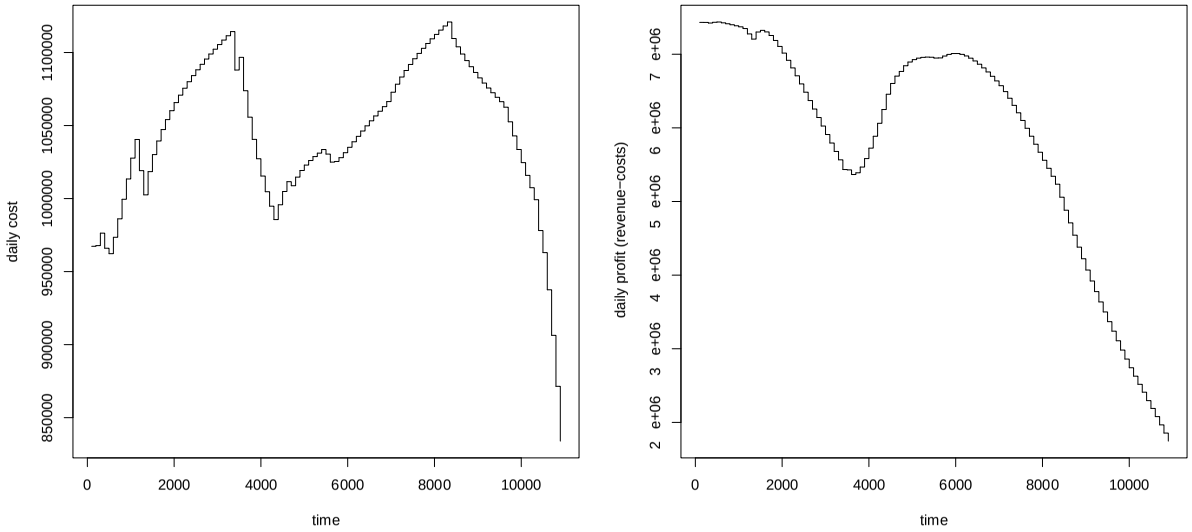

(a) (b)

Figure 1.5: The daily costs versus time (a), and daily profits versus time (b).

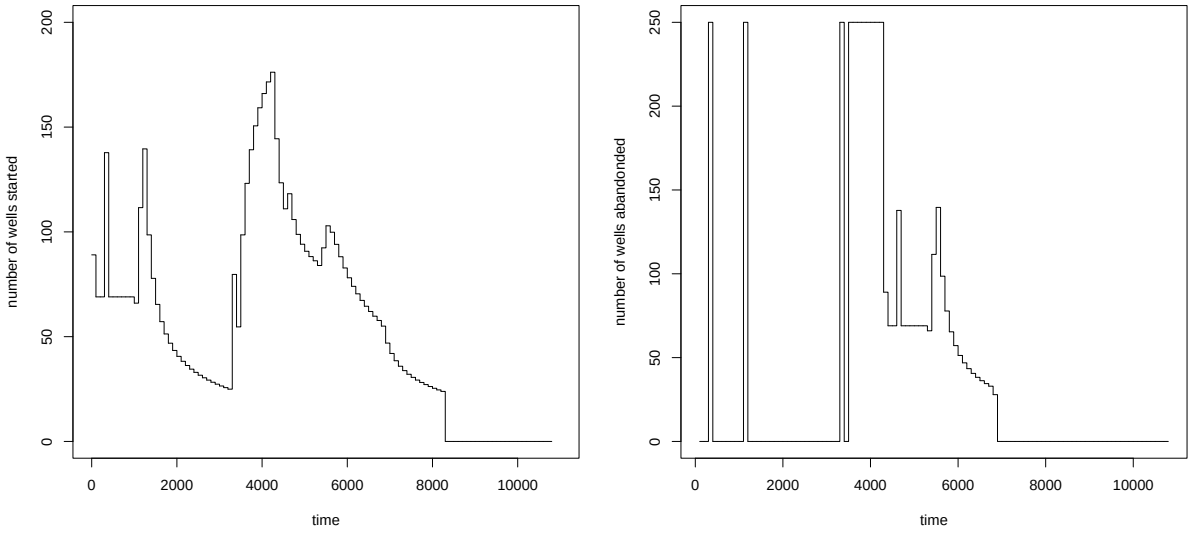

(a) (b)

Figure 1.6: The number of new wells started versus time (a), and the number of wells abandoned at age 4300 days versus time (b).

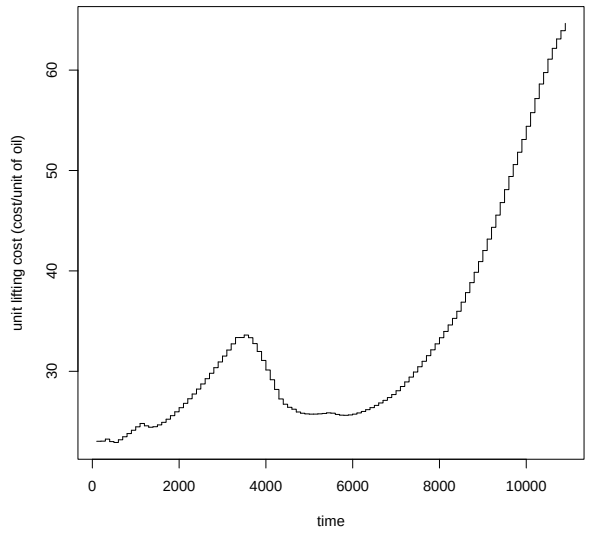

Figure 1.7: The unit lifting cost versus time

Table 1.2: The average, minimum, and maximum values of various characteristics over the life of the project.

	Average	Minimum	Maximum
Active Wells	3794	3089	4253
Daily Production	34145	12989	42000
Daily Steam Use	116297	83961	12000
Daily Costs	1044308	833941	1120901
Daily Profits	5784723	1745679	7437729
Unit Lifting Cost	43.78	22.91	64.66

1.5 Conclusion

The problem that Imperial Oil brought to the PIMS 4th Industrial Problem Workshop was to determine if there is a method to help manage the huge Cold Lake project. After exploring several approaches during and after the workshop, we feel that the answer is yes, as long as the it can be formulated as a linear programming problem in a manner similar to what we have described in the second part of this report.

The major advantage of the linear programming/linear optimization formulation over the non-linear approaches is that the complexity of the solution does not depend on the number of wells that are being modelled if the wells can be put into a small number of groups/classes. We believe that the linear approach is more suitable for a large project such as the Cold Lake system.

As we mentioned earlier, in this report we have assumed a constant value for the price of the commodity produced. An interesting and relevant complication to the problem would be to investigate the optimal strategy when that price is stochastic. A possible approach to this problem will be to use the techniques of *stochastic programming* but such an investigation is beyond the scope of this report.

Finally we note that another relevant issue faced by the managers of the Cold Lake facilities is how to deliver and distribute the steam from the three central plants to each individual well and the product from wells to the plants for treatment. This scheduling problem, which is not included in this report, will further complicate our mathematical models.

Acknowledgment. The author of the report (HH) would like to thank Kyle Biswanger and Mufeed Mahmoud for providing the materials for the nonlinear programming and the depth-first approaches. Special thanks to Marc Paulhus for the initial write-up of the linear programming formulation and many discussions during the preparation of this report.

1.6 Appendix - A Sample Problem for the LP Approach

To solidify the idea of using the LP approach, we look at a simpler problem's input and output files. In what follows $t_i s_j = z(i, j)$ in the notation of Section 4.

Input file:

```
Number of undeveloped wells:  3000
Time step (days):  100
Total time length (days):  400
Abandon times (days):  100 300 400
Starting well ages (days):  100 200 300
Number of wells starting with those ages:  400 400 400
Max steam use (day):  60000
Max production rate (day):  20000
Min production rate (day):  9000
Discount rate (annual):  0.05
Price of commodity:  200
Well development cost:  0
```

LP file:

/* Objective Function */

max: $311216\,t_1s_1 + 306425\,t_1s_2 + 298215\,t_1s_3 + 289585\,t_1s_4$
$+306981\,t_2s_1 + 302256\,t_2s_2 + 294157\,t_2s_3 + 285645\,t_2s_4$
$+302805\,t_3s_1 + 298144\,t_3s_2 + 290155\,t_3s_3 + 281758\,t_3s_4;$

/* Steam Constraints */

$115\,t_1s_1 + 50.11\,t_1s_2 + 44.27\,t_1s_3 + 41.17\,t_1s_4 \leq 60000;$
$115\,t_2s_1 + 50.11\,t_2s_2 + 44.27\,t_2s_3 + 41.17\,t_2s_4 \leq 60000;$
$115\,t_3s_1 + 50.11\,t_3s_2 + 44.27\,t_3s_3 + 41.17\,t_3s_4 \leq 60000;$

/* Lower production bound constraints */

$17.99\,t_1s_1 + 17.46\,t_1s_2 + 16.95\,t_1s_3 + 16.45\,t_1s_4 \geq 9000;$
$17.99\,t_2s_1 + 17.46\,t_2s_2 + 16.95\,t_2s_3 + 16.45\,t_2s_4 \geq 9000;$
$17.99\,t_3s_1 + 17.46\,t_3s_2 + 16.95\,t_3s_3 + 16.45\,t_3s_4 \geq 9000;$

/* Upper production bound constraints */

$17.99\,t_1s_1 + 17.46\,t_1s_2 + 16.95\,t_1s_3 + 16.45\,t_1s_4 \leq 20000;$
$17.99\,t_2s_1 + 17.46\,t_2s_2 + 16.95\,t_2s_3 + 16.45\,t_2s_4 \leq 20000;$
$17.99\,t_3s_1 + 17.46\,t_3s_2 + 16.95\,t_3s_3 + 16.45\,t_3s_4 \leq 20000;$

/* Transition Constraints */

$t_1s_1 + t_2s_1 + t_3s_1 \leq 3000;$
$t_1s_2 \leq 400;\ t_1s_3 = 400;\ t_1s_4 \leq 400;$
$t_2s_2 \leq t_1s_1;\ t_2s_3 = t_1s_2;\ t_2s_4 \leq t_1s_3;$
$t_3s_2\ et_2s_1;\ t_3s_3 = t_2s_2;\ t_3s_4 \leq t_2s_3;$

/* Declarations */

int $t_1s_1,\ t_1s_2,\ t_1s_3,\ t_1s_4,\ t_2s_1,\ t_2s_2,\ t_2s_3,\ t_2s_4,$
$t_3s_1,\ t_3s_2,\ t_3s_3,\ t_3s_4;$

Solution File:

```
Value of objective function:  942033225
```
$t_1s_1 = 95,\ t_1s_2 = 400,\ t_1s_3 = 400,\ t_1s_4 = 275,$
$t_2s_1 = 183,\ t_2s_2 = 95,\ t_2s_3 = 400,\ t_2s_4 = 400,$
$t_3s_1 = 262,\ t_3s_2 = 183,\ t_3s_3 = 95,\ t_3s_4 = 400$

Bibliography

[1] R. Courant and D. Hilbert, Methods of Mathematical Physics, Interscience Publishers Inc., New York, 1953, p. 175.

[2] *www.cs.sunysb.edu/∼algorith/implement/lpsolve/implement.shtml*

[3] *www.lindo.com*

[4] *www.cplex.com*

[5] *www-fp.mcs.anl.gov/otc/Guide/SoftwareGuide/index.html*

Chapter 2

Statistical Design of an Experimental Problem in Harmonics

Rita Aggarwala[1], Jahrul M. Alam[2], Md. Shafiqul Islam[3], Michael Lamoureux[1], Marc Paulhus[1], Miro Powojowski[1], Leila Rasekh[4]
Report prepared by M. Lamoureux

2.1 Abstract

The Michelin Tire Company requires its tires to be very uniform in order to provide a smooth, quiet ride. The design of the appropriate manufacturing technique leads to a problem in harmonic analysis, and to the problem of the design of a statistical experiment to accurately measure the harmonic components contributing to a measured force on the tire. Optimal designs were developed, as well as a number of useful Monte Carlo methods. MATLAB codes of the tests are provided.

2.2 Introduction

The Michelin Tire Company is interested in manufacturing tires that meet a certain level of uniformity. Generally speaking, the more uniform the tire, the smoother the ride and the quieter is the tire, in terms of rolling noise perceived by the passengers. Indeed, many of the contracts Michelin holds for tire production require that they manufacture tires that meet certain specified standards of uniformity.

Uniformity of a tire is evaluated by measuring the force exerted by the tire on a measuring device, as the tire is rotated 360 degrees about an axle. The force may be measured in the radial direction, in the axial direction, or both. This produces one or two force curves which describe the uniformity of the tire: a flat curve indicates a perfectly uniform tire. A non-flat

[1]University of Calgary
[2]University of Alberta
[3]Concordia University
[4]University of Guelph

curve can be expressed as a sum of sines and cosines, and the manufacturer is interested only in the magnitude of specific harmonics of this force curve. For instance, a contract may require that the manufactured tires have harmonics one through five with magnitudes below a certain threshold. Thus the problem of ensuring uniformity of the tires is more precisely a problem of minimizing, or reducing below a certain threshold, certain harmonic components of one or two force curves.

To understand how the non-uniformities in a tire appear, and how one may control them, it is necessary to understand the basic construction of a tire. Each tire is made up of a series of components, or sheets of material, which are wrapped one layer on top of another, while stretched out over a tire mold. For instance, the first layer would be an inner, airtight sheet of reinforced rubber which will form the inner air compartment for the tire. Next comes several layers of different types of rubber, then layers of cords and/or steel belting for reinforcement, more rubber, and eventually the treads are laid on top to finish off the tire.

There may be as many as twenty layers built up into the tire. As each sheet of material is wrapped around the mold, the ends are joined by glue or by melting, which results in a small bead of material at the join. These beads will lead to bumps, or non-uniformities in the final tire assembly. There are other factors as well that cause each layer to contribute some non-uniformities to the final product; the causes are not in question here, but they are indicated so it is understood that each layer somehow contributes to the uniformity or non-uniformity of the tire.

The technician has little control over what causes the particular bumps or non-uniformities in the layers that go into building the tire. One thing that the technician can control is the position (starting angle of assembly) for each layer going into the tire. For instance, the first layer may be aligned at 0 degrees relative to the mold, second layer rotated 15 degrees from the first layer, third layer rotated at 30 degrees, and so on.

The technician would like to choose this series of rotations in order to produce a tire that meets the uniformity requirements. This leads to a canonical problem of analysis and synthesis. One must first analyse the layers to see how each one contributes to the final uniformity, or non-uniformity, of the finished tire; then one must use this knowledge to design the finished tire that meets the specified uniformity requirements. The only freedom one has in both the analysis and synthesis is choosing the position of each layer.

Our goal in this problem is to solve the analysis problem. The task is to design an experiment where a technician can construct a series of test tires that will be used to determine, with statistical confidence, the contribution of each layer to the final profile, or force curve, of the generic tire. Once the contributions are known, the synthesis problem of finding the optimally smooth tire, or one that meets any specific uniformity conditions on the harmonic components, is relatively straightforward, and need not concern us here.

2.3 Methodology

It was clear early that this problem would require techniques of both Fourier analysis and statistics: Fourier techniques to describe the harmonic problem, and statistical techniques to deal with the problem of accurately measuring quantities that come from real experimental

measurements. A team of researchers was assembled that brought in expertise from both areas. It turned out that Monte Carlo methods would be useful as well, and appropriate simulations on the computer would be needed. MATLAB was chosen as the computational platform, as it combines the power of complex, matrix linear algebra required by the harmonic analysts with the statistical tools needed by the statisticians.

As the research progressed, our industrial collaborator from Michelin provided some very useful information based on their experience with manufacturing and analysing their tires. In particular, he indicated what types of statistical experiments had been performed, what tests looked promising, what constraints and costs could be expected in running the test. This information directed much of the research below, and details are indicated where appropriate.

A number of very promising, and even optimal, designs were obtained over the course of the workshop. Given the limited time available in a one week period, it was impossible to obtain complete characterizations for these designs, their robustness, and other features. With MATLAB, some good verifications were produced that are quite convincing. Thus, the results brought the analysis problem all the way to a solution, and our industrial collaborator has indicated it is a significant, valuable answer to their problem.

2.4 Simplifications

While the original problem involved a choice of one or two force curves to measure, it was quickly decided to work with just one force curve. This is only a minor simplification. The Monte Carlo methods developed for one curve can easily be extended to two curves. Also, in the optimal designs developed, the designs depends only on the harmonics studied under the test; thus, two curves can be included in the design by simply including the harmonics for both. Since the physical construction of test tires is expensive, it turns out this is a rather efficient way of minimizing the number of tires needed to test any number of force curves.

It was also noted that the force curves are in fact only sampled at 256 points, thus the problem of infinite dimensional spaces of curves is avoided. Indeed, a large simplification is obtained by noting a finite, discrete Fourier transform (which is easily done on the computer) reduces the harmonic problem to a natural setting.

In practice, the technician is not completely free in his choice of angles used to lay out components of the tire. In the development below, we assumed any angle could be chosen (though perhaps not all layers would be adjusted at the same time), and then one would use robustness results to see what happened when the choice of angle was restricted.

A statistical simplification is to assume the layers of the tire act independently, and additively, in producing the final force curve of the tire. This assumption is supported by the experience of the Michelin group of researchers.

Another major simplification was to get the whole team speaking a common language of Fourier series, least square approximation, and using MATLAB tools.

2.5 Fourier Series Formulation

The tire's characteristic force curve can be described as a continuous function on the interval $[0,1]$ with periodic boundary conditions. As the problem was originally posed, it was suggested that these periodic functions $f \in C[0,1]$ be expanded in terms of sines and cosine functions: that is, one writes

$$f(t) = \sum_{n=0}^{\infty} a_n \cos(2\pi nt) + b_n \sin(2\pi nt),$$

where a_n, b_n are real coefficients encoding the magnitude and phase of the corresponding harmonic. Algebraically, it is more convenient to use complex exponentials to expand the periodic function in a Fourier series, as

$$f(t) = \sum_{n=-\infty}^{\infty} c_n e^{2\pi int},$$

for complex coefficients c_n. Since f is real, there is some redundancy in this expansion requiring $c_{-n} = \overline{c_n}$, the complex conjugate of c_n, for each n.

More important is the redundancy due to the fact that the continuous function $f \in C[0,1]$ is observed only at finitely many points $t_k = k/256$, for $k = 0, 1, ..., 255$. The Fourier series expansion at these discrete points

$$f(t_k) = \sum_{n=-\infty}^{\infty} c_n e^{2\pi i \frac{nk}{256}}$$

collapses to a finite sum, because of the periodicity of the exponentials, so one may write a discrete Fourier expansion, with

$$f(t_k) = \sum_{n=-127}^{128} c_n e^{2\pi int_k}, \qquad t_k = 0, 1/256, 2/256, \ldots 255/256.$$

It is worth noting in passing that the Fourier coefficients c_n are quickly calculated using an FFT software routine, and so this formulation of the problem does not introduce any additional complexity in the problem.

In the problem at hand, a function f is considered the signature of a given tire component or layer that will affect the final force profile. If this component is rotated by an angle θ, the signature function f is shifted and the corresponding Fourier coefficients change. Introducing the notation S_θ for the shift operator, one obtains

$$\begin{aligned}
(S_\theta f)(t) &= f(t-\theta)) \\
&= \sum_n c_n e^{2\pi in(t-\theta)} \\
&= \sum_n (e^{-2\pi in\theta} c_n) e^{2\pi int} \\
&= \sum_n c_n^\theta e^{2\pi int}.
\end{aligned}$$

That is, the Fourier coefficients transform under the shift by θ as a linear transform $c_n \mapsto e^{-2\pi i n\theta} c_n$. Equivalently, the vector of coefficients c_n transforms as

$$\begin{pmatrix} \vdots \\ c_n \\ \vdots \end{pmatrix} \mapsto \begin{pmatrix} \vdots \\ c_n^\theta \\ \vdots \end{pmatrix} = \begin{pmatrix} \ddots & & \\ & e^{-2\pi n\theta} & \\ & & \ddots \end{pmatrix} \begin{pmatrix} \vdots \\ c_n \\ \vdots \end{pmatrix}$$

which is just multiplication by a diagonal matrix D_θ whose entries are complex exponentials.

2.6 The Component Problem

A tire is built up from a number of layers (tread, cords, airtight inner rubber, etc), usually on the order of $m = 20$ components. Each layer contributes a signature $f^k(t) \in C[0,1]$ to the observed force profile

$$F(t) = \sum_{k=1}^{m} f^k(t),$$

where the assumption (based on Michelin's experience) is that the contribution is additive, and thus each layer's contribution is independent of the others. The component functions f^k cannot be measured directly; however, a factory worker may modify the construction of the tire by changing the positioning of individual layers within the test tire. Each layer may be shifted independently by some angle θ. Applying a vector of shifts $\Theta = (\theta_1, \theta_2, \ldots, \theta_m)$, where θ_k is the rotation angle for k-th layer, gives an operation on the observed force profile as

$$(S_\Theta F)(t) = \sum_{k=1}^{m} (S_{\theta_k} f^k)(t).$$

In Fourier components, this becomes

$$\begin{aligned} \sum_n C_n^\Theta e^{2\pi i n t} &= \sum_k \sum_n (e^{-2\pi i n \theta_k} c_n^k) e^{2\pi i n t} \\ &= \sum_n (\sum_k e^{-2\pi i n \theta_k} c_n^k) e^{2\pi i n t}, \end{aligned}$$

and by equating terms in the Fourier expansion one obtains the transform directly on the coefficients as

$$C_n^\Theta = \sum_{k=1}^{m} e^{-2\pi i n \theta_k} c_n^k.$$

In particular, one observes there is no mixing of harmonics: that is, the n-th harmonic of the observed (transformed) force curve is a weighted sum of the n-th harmonics of the contributing layers.

The analysis problem is to determine the coefficients c_n^k from the observed C_n^Θ, using some choice of the vector of angles Θ. Since the observed spectra are real, it is enough to consider only non-negative n in determining the harmonics, and the constant term ($n = 0$) is irrelevant.

In practice, only a small values of n are of interest (eg. $n = 1, 2, \ldots, 5$), as these correspond to certain low frequency vibrations, but the coefficients must be determined for all layers (eg. $k = 1, 2, \ldots, 20$). Also note this is a statistical data problem, as the measured coefficients include measurement error and statistical deviations due to variations in the construction of these real tires.

2.7 The Linear Model

The problem is to determine individual coefficients c_n^k, for all layers $k = 1, \ldots, m$, from observations of the lumped coefficients C_n^Θ, where the experimental design involves choosing some appropriate vectors of angles $\Theta = (\theta_1, \theta_2, \ldots, \theta_m)$. The design also should find the number of vectors $\Theta^1, \Theta^2, \ldots, \Theta^R$ required to accurately determine the coefficients c_n^k. These coefficients must be determined for a range of harmonics, say $n = 1, \ldots, q$, and it will be convenient to design the experiments to work for all these harmonics simultaneously.

It is natural to group the coefficients into column vectors, as

$$\begin{pmatrix} c_1^k \\ \vdots \\ c_q^k \end{pmatrix} = \vec{c}^{\,k}$$

and

$$\begin{pmatrix} C_1^\Theta \\ \vdots \\ C_q^\Theta \end{pmatrix} = \vec{C}^\Theta.$$

The linear model encompassing all layers, and the range of harmonics, can thus be written in block form as

$$\begin{pmatrix} \vec{C}^{\Theta^1} \\ \vec{C}^{\Theta^2} \\ \vdots \\ \vec{C}^{\Theta^R} \end{pmatrix} = \begin{pmatrix} D_{\theta_1^1} & D_{\theta_2^1} & \ldots & D_{\theta_m^1} \\ D_{\theta_1^2} & D_{\theta_2^2} & \ldots & D_{\theta_m^2} \\ \vdots & & & \vdots \\ D_{\theta_1^R} & D_{\theta_2^R} & \ldots & D_{\theta_m^R} \end{pmatrix} \begin{pmatrix} \vec{c}^{\,1} \\ \vec{c}^{\,2} \\ \vdots \\ \vec{c}^{\,m} \end{pmatrix} + \vec{\epsilon}$$

where $\Theta^1, \ldots, \Theta^R$ is the choice of vectors of angles set in the experiment, the $D_{\theta_k^r}$ are $q \times q$ diagonal matrices with entries $e^{-2\pi i n \theta_k^r}$ on the diagonal, and $\vec{\epsilon}$ is the statistical measurement error.

More succinctly, the linear model is represented by $\vec{C}^{\vec{\Theta}} = D\vec{c} + \vec{\epsilon}$, with D a matrix in block form, each block a diagonal matrix as above. If these were real matrices, the solution via least squares is clear. It was a simple exercise, undertaken in the course of this workshop, to verify that even for complex matrices, the least square solution is obtained in a straightforward manner via computations with the usual complex inner product. Namely, one solves for $\vec{c}$ as

$$\vec{c} = (D^* D)^{-1} D^* \vec{C}^{\vec{\theta}},$$

where D^* indicates the complex conjugate transpose of the matrix D. Similarly, the variance estimates for the inversion will depend on the properties of matrix $(D^* D)^{-1}$.

Noting that D^*D is also in block form, it is convenient to permute rows and columns (essentially grouping terms by layers, rather than harmonics) to obtain a matrix $X = Perm(D)$ so that X^*X is in block diagonal form, with

$$X^*X = \begin{pmatrix} (Z_1) & & & 0 \\ & (Z_2) & & \\ & & \ddots & \\ 0 & & & (Z_q) \end{pmatrix},$$

where each $m \times m$ block (Z_n) has entries

$$(Z_n)_{jk} = \sum_{r=1}^{R} e^{2\pi i n(\theta_k^r - \theta_j^r)}.$$

This greatly simplifies the analysis, since each block (Z_n) may be examined separately. Notice each such block corresponds to a separate harmonic.

The problem becomes that of estimating the regression coefficients in the multiple regression model

$$\vec{C}^{\vec{\Theta}} = X\vec{c} + \vec{\epsilon}.$$

A standard assumption is that

$$Var(\vec{\epsilon}) = \sigma^2 I_{qm}$$

for some (unknown) value σ. It was pointed out that this assumption essentially says the R tires and q harmonics act independently, and different measurements have equal error; this may be a gross oversimplification worth further investigation. For instance, there may be some bias in the way the tires are constructed for the test, or trends reflected in the sequence in which the tires are built. On the other hand, the harmonics are orthogonal measures in a large dimensional space, and at least some of us were convinced that the first few harmonics would act independently, with similar measurement error. In any case, we proceed with this assumption.

The least-square estimator of $\vec{c}$ is

$$\hat{c} = (X^*X)^{-1}X^*\vec{C}^{\vec{\Theta}}$$

with variance

$$Var(\hat{c}) = (X^*X)^{-1}X^*Var(\vec{\epsilon})X(X^*X)^{-1} = \sigma^2(X^*X)^{-1}.$$

Thus the problem of finding an optimal design boils down to finding a matrix X such that X^*X is "good." For more general forms of $Var(\vec{\epsilon})$, the optimal condition is more complicated.

Some possible optimality conditions ("goodness" of X) include minimizing the determinant of the matrix $(X^*X)^{-1}$ (D-optimality), minimizing the spectral norm of $(X^*X)^{-1}$, or minimizing the maximum eigenvalue of $(X^*X)^{-1}$. It turns out these three conditions are equivalent, since the matrix (X^*X) has trace independent of the choice of angles (equal to mqR) and thus the minimum occurs when all eigenvalues are equal, and X^*X is R times the identity matrix. Generally speaking, the closer X^*X is to diagonal, the better.

Another optimality condition is to fix some vector w and minimize the variance $Var(w'\hat{c})$, which is a weighted sum of the entries of $\hat{c}$. This would be of interest to the manufacturer when some harmonics, or some layers, are deemed to be more important than others.

2.8 Numerical Experiments

Random sampling of the m-dimensional hypercube $[0, 2\pi]^m$ with a given sample size R (number of tires) produces random vectors of angles $\Theta^1, \ldots, \Theta^R$, which are exponentiated to produce random matrices X. The block diagonal structure of X^*X allows us to focus on one harmonic at a time. Moreover, with the components of Θ^r chosen uniformly in $[0, 2\pi]$, then the mod 2π part of multiples $n\Theta^r$ are also uniformly distributed. Thus the Monte Carlo designs work for all harmonics.

Software code was produced in MATLAB to generate these random matrices and search for a best solution, typically from a sample of 10,000 to 100,000 random matrices. Sample code is provided in the appendix. Plots were obtained to show how the performance of MC-best solution improved with increasing R. A $1/R$ dependence was easily observed, although even for large values of R, the MC-best solution was not at the theoretical best solution, where X^*X is a multiple of the identity.

This led to the following theoretical observation. Recall for only one (n-th) harmonic, that

$$(X^*X)_{jk} = \sum_{r=1}^{R} e^{2\pi i n(\theta_k^r - \theta_j^r)}.$$

Thus on the diagonal, $(X^*X)_{jj} = R$, while on the off-diagonal, as $R \to \infty$,

$$\begin{aligned} \frac{1}{R}(X^*X)_{jk} &= \sum_{r=1}^{R} \frac{1}{R} e^{2\pi i n(\theta_k^r - \theta_j^r)} \\ &\to \int_0^1 e^{2\pi i n t} dt \\ &= 0, \end{aligned}$$

where the random sum is simply an approximation to the integral. Thus

$$\lim_{R\to\infty} \frac{1}{R}(X^*X) = I_m,$$

which suggests that for large R one should see $(X^*X)^{-1}$ approximately equal to $\frac{1}{R} I_m$. Thus the Monte Carlo designs should be tending to this limit, although the convergence may be quite slow. In general, for uniform sampling, one expects

$$\sum_{r=1}^{R} \frac{1}{R} e^{2\pi i n(\theta_k^r - \theta_j^r)} \approx \frac{c}{\sqrt{R}}$$

for some constant $c > 0$.

A number of other numerical experiments were tried, including choosing random vectors of angles Θ where all but a few (say 4) of the angles were zero, and choosing angles from a discrete subset of $[0, 2\pi]$, say of 10 to 30 evenly distributed points. We tried to find interesting patterns in the resulting MC optimal designs, but did not see anything remarkable. However, these restricted MC designs have their use in practice, for instance when the operator constructing tires can only adjust a few layers at a time, or has only a limited precision in choice of angles. The $1/R$ behaviour was also noted in these designs. MATLAB code for these experiments is included in the appendix.

2.9 The Prime Method

Having explored the Monte Carlo method extensively, some ingenuity was required to find concrete patterns that would provide optimal methods without the MC search. An initial observation was that the matrix (X^*X) has off-diagonal terms which are sums of complex exponentials. It thus might be possible to arrange these exponentials to sum to zero. Indeed, note that

$$\sum_{r=1}^{R} e^{2\pi i(\frac{r}{R})} = 0,$$

hence if for each pair (j,k), the numbers $\{\theta_k^r - \theta_j^r\}_{r=1}^R$ are a permutation of the fractions $\{\frac{r}{R}\}_{r=1}^R$, then the matrix X is in its optimal form, with $(X^*X)^{-1} = \frac{1}{R} I_m$ exactly. This construction is then much better than the Monte Carlo method.

It turns out this can always be arranged when the number of layers m is a prime number, in which case the design size R is taken to be $R = m$. Here is an example of the matrix of angles for $m = 5$:

$$(\theta_j^r) = \frac{2\pi}{5} \cdot \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 2 & 3 & 4 \\ 0 & 2 & 4 & 1 & 3 \\ 0 & 3 & 1 & 4 & 2 \\ 0 & 4 & 3 & 2 & 1 \end{pmatrix} = \begin{pmatrix} \Theta^1 \\ \Theta^2 \\ \Theta^3 \\ \Theta^4 \\ \Theta^5 \end{pmatrix}.$$

Since $\mathbf{Z}/5$ is a field, it is easy to see that multiplication by n just permutes the rows of the matrix (θ_j^k) (when $n \neq 0 \bmod 5$). That is, the matrix $(n\theta_j^k)$ also gives an optimal form, so this first design works equally for all harmonics which are not multiples of 5.

In general, for m equal to any prime, choose the design matrix of angles to be $m \times m$ with entries

$$\theta_j^r = \frac{2\pi}{m}\{(j-1)(r-1) \bmod m\}.$$

Again, this design gives optimal X for any prime m and any harmonic n which is not a multiple of m. This is a result of the fact that $\mathbf{Z}/m$ is a field when m is prime, as shown in the following:

Theorem 1 *For integer m prime, n not a multiple of m,and design angles chosen* as

$$\theta_j^r = \frac{2\pi}{m}\{(j-1)(r-1) \bmod m\},$$

then the corresponding design matrix

$$X_{jk} = e^{n\theta_j^r}$$

is optimal.

Proof. The covariance matrix X^*X has entries $(X^*X)_{jk} = \sum_{r=1}^m e^{2\pi i n(\theta_k^r - \theta_j^r)}$ so along the diagonal, the exponentials are each equal to one, thus $(X^*X)_{jj} = m$. Off the diagonal, for any $j \neq k$, the exponentials have powers of the form $2\pi i(n(k-j)(r-1) \bmod m)$ for r in the range $1, 2, \ldots, m$. Since $n(k-j)$ is a non-zero element in the field $\mathbf{Z}/m$, multiplication by this element of the sequence $\{r-1\} = \{0, 1, \ldots, m-1\}$ simply permutes these elements of the field, so the

sum is over the m roots of unity. Hence each off-diagonal element of the matrix X^*X is zero. Thus X^*X is a multiple of the identity, and optimal.

In summary, this choice of design angles gives a powerful method both in that it gives explicitly an optimal solution, and that it also works for a range of harmonics.

There are some disadvantages to this method. First it is somewhat inflexible in the number of layers m, as m must be prime. This can be remedied by introducing "fake" layers to reach the next lowest prime. For instance, for 20 layers, just pretend there are 23 layers, three of which are virtual. Or group two insignificant layers and call them one – so 20 becomes the prime 19.

Second, it requires using as many angles as there are layers. That is, if one has 19 effective layers, each of the 19 layers must be set to various angles as the tires are constructed. This can be an expensive, if not impossible, construction in some tire plants. A partial solution is to use a method of blocking layers, as discussed in the section below.

Finally, there is the problem of setting angles exactly: the operator may have only limited accuracy on how precisely layer angles can be set during construction, and may have physical obstruction in choosing particular angles. This method expects the operator to freely choose the angles.

2.10 The Blocking Method

To avoid the problem of setting many angles for many layers, it is convenient to block off groups of layers and treat them as a single unit. Indeed, this block layer can be rotated by zero degrees – in effect, no rotation – so only the remaining layers need to be rotated.

There is much flexibility in this method, as one can choose how many layers to move, how to group them, and so forth. Rather than explore all the possible permutations, here is a simple example with 12 layers, and setting no more than 5 angles at a time. One way to proceed is to group the first 8 layers as one, and treat the last four independently, giving 5 effective layers. Next, group the first and last four as one, middle four treated independently. Last, group the final eight as one.

Thus, the groupings of layers looks like the following:

actual:	1	2	3	4	5	6	7	8	9	10	11	12
group a:	1	1	1	1	1	1	1	1	2	3	4	5
group b:	1	1	1	1	2	3	4	5	1	1	1	1
group c:	2	3	4	5	1	1	1	1	1	1	1	1

The corresponding design of vectors of angles is given by choosing the 5×5 blocks of the last section, using the prime method. For clarity, we can show this array in block form, with blanks

indicating the blocked zeros, giving this 15×12 array:

$$\theta_k^r = \frac{2\pi}{5} \begin{pmatrix} & & & & & & & & 0 & 0 & 0 & 0 \\ & & & & & & & & 1 & 2 & 3 & 4 \\ & & & & & & & & 2 & 4 & 1 & 3 \\ & & & & & & & & 3 & 1 & 4 & 2 \\ & & & & & & & & 4 & 3 & 2 & 1 \\ & & & & 0 & 0 & 0 & 0 & & & & \\ & & & & 1 & 2 & 3 & 4 & & & & \\ & & & & 2 & 4 & 1 & 3 & & & & \\ & & & & 3 & 1 & 4 & 2 & & & & \\ & & & & 4 & 3 & 2 & 1 & & & & \\ 0 & 0 & 0 & 0 & & & & & & & & \\ 1 & 2 & 3 & 4 & & & & & & & & \\ 2 & 4 & 1 & 3 & & & & & & & & \\ 3 & 1 & 4 & 2 & & & & & & & & \\ 4 & 3 & 2 & 1 & & & & & & & & \end{pmatrix}.$$

In fact the zero rows are not particularly useful, since there are plenty of zeroes elsewhere in the matrix. Eliminating the zero rows gives the following 12×12 matrix:

$$\theta_k^r = \frac{2\pi}{5} \begin{pmatrix} & & & & & & & & 1 & 2 & 3 & 4 \\ & & & & & & & & 2 & 4 & 1 & 3 \\ & & & & & & & & 3 & 1 & 4 & 2 \\ & & & & & & & & 4 & 3 & 2 & 1 \\ & & & & 1 & 2 & 3 & 4 & & & & \\ & & & & 2 & 4 & 1 & 3 & & & & \\ & & & & 3 & 1 & 4 & 2 & & & & \\ & & & & 4 & 3 & 2 & 1 & & & & \\ 1 & 2 & 3 & 4 & & & & & & & & \\ 2 & 4 & 1 & 3 & & & & & & & & \\ 3 & 1 & 4 & 2 & & & & & & & & \\ 4 & 3 & 2 & 1 & & & & & & & & \end{pmatrix}.$$

Exponentiating with $X_{kr} = e^{ni\theta_k^r}$, one finds the following covariance matrix with an elegant

block form:

$$(X^*X) = \begin{pmatrix} 12 & 7 & 7 & 7 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 \\ 7 & 12 & 7 & 7 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 \\ 7 & 7 & 12 & 7 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 \\ 7 & 7 & 7 & 12 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 \\ 2 & 2 & 2 & 2 & 12 & 7 & 7 & 7 & 2 & 2 & 2 & 2 \\ 2 & 2 & 2 & 2 & 7 & 12 & 7 & 7 & 2 & 2 & 2 & 2 \\ 2 & 2 & 2 & 2 & 7 & 7 & 12 & 7 & 2 & 2 & 2 & 2 \\ 2 & 2 & 2 & 2 & 7 & 7 & 7 & 12 & 2 & 2 & 2 & 2 \\ 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 12 & 7 & 7 & 7 \\ 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 7 & 12 & 7 & 7 \\ 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 7 & 7 & 12 & 7 \\ 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 & 7 & 7 & 7 & 12 \end{pmatrix}.$$

This is not an optimal matrix, but it is a reasonably good one whose performance is better than the MC designs found above. In general this method can be extended, using small primes to form $p \times p$ sub-design matrices which are used to tile the larger $R \times m$ full design matrix.

2.11 The GLP Method

The Good Lattice Point (GLP) method of Fang and Wang (Ref. [1]) uses a careful choice of lattice points in an m-dimensional hypercube to accelerate integration over a multidimensional Riemann sum. The basic insight is to look for angle combinations which will lead to sequences $\{\theta_k^r - \theta_j^r\}_{r=1}^R$ which will allow for fast convergence of the sum

$$\sum_{r=1}^{R} \frac{1}{R} e^{ni(\theta_k^r - \theta_j^r)} \to \int_0^1 e^{2\pi i n x} dx.$$

With the Monte Carlo method, random sampling of the hypercube produces random sequences on $[0, 1]$, but with a slow convergence of order $R^{-\frac{1}{2}}$. The GLP method will exhibit convergence at the faster rate of $R^{-1} \log^m(R)$. We tested the GLP method to see if we would obtain a good sequence of angles, and came up with a surprising conjecture.

First, let us recap the definition of a lattice point set and a GLP set, as discussed in reference [1]. Let $(R, h_1, h_2, \ldots, h_m)$ be a vector of integers satisfying

- $m < R$
- $1 \leq h_j < R$
- $h_j \neq h_k$, for all $j \neq k$
- (h_j, R) are coprime, for all j.

The *lattice point set* of the generating vector $(R; h_1, h_2, \ldots, h_m)$ is the set of vectors $\{(x_{r1}, \cdots, x_{rm}), r = 1, \cdots, R\}$, with values

$$x_{rj} = \text{frac}\left(\frac{2rh_j - 1}{2R}\right),$$

where "frac" denotes the fractional part of the given real number. If this set has the smallest discrepancy (defined in Fang and Wang), it is called a GLP set.

The principle behind GLP sets is that generating vectors can always be constructed so that a GLP set is created, whose points are uniformly distributed about the hypercube. Fang and Wang tabulate many different choices for a range of R and m, corresponding in our case to numbers of tires R in the experimental design, and number of layers m per tire.

In our tire example, the vector of angles are obtained from the lattice point sets by scaling by a factor of 2π, so $\theta_j^r = 2\pi x_{rj}$. We tested a number of the GLPs from the book to see how close they are to optimal, and found in every instance, they were exactly optimal. We have the following:

Conjecture 1 *Every GLP set produces an optimal design: that is,*

$$(X^*X)^{-1} = \frac{1}{R} I_m \text{ exactly.}$$

Moreover,

- *m can be chosen arbitrarily (not necessarily prime)*
- *R can be chosen arbitrarily (although prime a popular choice)*
- *the same design is optimal for all harmonics co-prime with R.*

In the workshop, there was not enough time to explore how GLP sets were constructed in the literature, so it was not clear to us how optimal designs were resulting from these choices. A quick review of work in the area indicates some number-theoretical results are being used to construct the charts of Fang and Wang. However, a simple examination shows the differences $x_{rj} - x_{tk} = \text{frac}(\frac{r}{R}(h_j - h_k))$, so as in the prime method, the sum in the covariance matrix will cycle around a subset of the R roots of unity. For a good choice of the h_j, this subset will always sum to zero. Thus, while this is short of a proof verifying the GLP method works, there is the basis for a useful technique, explored in next section.

2.12 The Simplified Lattice Method

The GLP method is in fact too sophisticated for the problem at hand, as all that is needed are roots of unity that sum to zero. However, it suggests the following technique, which we call the Simplified Lattice Method.

Theorem 2 *Fix an integer $m > 0$ and fix $\mathbf{N}$ a subset of $\{1, 2, \ldots\}$. Suppose $(R; h_1, h_2, \ldots, h_m)$ is a vector of integers satisfying*

- $m < R$
- *R is not a divisor of $n(h_j - h_k)$ for all $n \in \mathbf{N}, j \neq k$.*

Then the vectors of angles (scaled lattice points) defined by

$$\theta_j^r = \frac{2\pi}{R}(rh_j \bmod R)$$

gives an optimal design for all harmonics $n \in \mathbf{N}$. That is,

$$(X^*X)^{-1} = \frac{1}{R}I_m.$$

Proof. For harmonic n in the set $\mathbf{N}$ and $j \neq k$, the integer $n(h_j - h_k)$ is not divisible by R and hence the map $r \mapsto n(h_j - h_k)r \bmod R$ defines a endomorphism on the ring $\mathbf{Z}/R$ which has more than one element in its range, a subring of $\mathbf{Z}/R$. Thus when scaled by $2\pi i$ and exponentiated, one obtains some R' roots of unity, for some divisor $R' > 1$ of R. Thus the terms in the sum

$$\sum_{r=1}^{R} e^{2\pi nr(h_k - h_j)/R}$$

simply cycle around these R' roots of unity, and so sum to zero. Hence the off-diagonal terms of the covariance matrix X^*X are zero, the diagonal terms are R, and the optimal design is achieved.

These criteria are easy to fulfill in any situation of tires, as shown in the following.

Example. With m the number of layers in the tire, and $\mathbf{N}$ a finite set of harmonics, let R be any prime number strictly bigger than m and all integers $n \in \mathbf{N}$. Then the integer vector $(R; 0, 1, 2, \ldots, m-1)$ generates lattice points yielding an optimal design.

This example gives a method much like the original prime method described above. However, the number of layers m need not be prime, and one can select any finite set of harmonics, yet still obtain an optimal design. The number of tires R need not be prime: one could choose a composite number with some prime factor bigger than m and all n.

There remains the disadvantage that almost every layer on almost every tire must be set to a non-zero angle, and the angles must be set to accuracies on the order of $2\pi/R$. Blocking may be used as a partial solution to this problem.

2.13 Robustness

We considered two ways in which a study may be corrupted:

- a tire is lost;
- there is some errors in the angle set.

A lost tire amounts to deleting a row from the matrix X. We considered a numerical example with $R = 23$, by removing a row at random, and measuring the eigenvalues of the resulting suboptimal matrix. The result was still close to optimal (i.e. eigenvalues nearly constant), even with two or three rows removed. Indeed, these suboptimal designs were still much better than our Monte Carlo searches. This was enough to convince the team that the optimal design was fairly robust; however we did not have time to investigate this more fully.

Errors in the angle setting can be investigated by a Monte Carlo study where normal random variates are added to the angles. Again, the design appears to be quite robust to such departures, although we note higher harmonics are proportionally more sensitive.

MATLAB code for both these investigations is included in the appendix.

2.14 Tire Types vs Replicates

Introducing tire replicates of the same type (that is, tires with the same vector of angles Θ^r) can help reduce the variance of the estimators and may be more cost-effective. This would be the case if it is cheaper to manufacture a run of several tires of the same type in a given study.

The linear model becomes

$$\vec{C}^{\vec{\Theta}} = (X \otimes 1_n)\vec{c} + \vec{\epsilon},$$

where $\otimes$ denotes the Kronecker matrix product, and n is the number of replicates of each type of test tire. If one assumes the variance remains as $Var(\epsilon) = \sigma^2 I_{nR}$, then

$$\begin{aligned} Var(\hat{c}) &= ((X \otimes 1_n)^*(X \otimes 1_n))^{-1} (X \otimes 1_n)^* \sigma^2 (I_R \otimes I_n)(X \otimes 1_n) ((X \otimes 1_n)^*(X \otimes 1_n))^{-1} \\ &= ((X^*X \otimes 1_n^* 1_n))^{-1} \\ &= \frac{\sigma^2}{n}(X^*X)^{-1} \\ &= \frac{\sigma^2}{nR} I_R. \end{aligned}$$

Thus for good designs, the variance is inversely proportional to nR and it thus may be more effective to increase n and not R.

One must be careful to recognize that in the case of tire replicates, there are two different classes of error sources:

- errors specific to a particular type of tire (perhaps caused by inaccurate machine settings for the mold building one particular tire), and
- errors specific to individual tires or measurements.

The error structure is then of the form

$$Var(\vec{\epsilon}) = \begin{pmatrix} \gamma^2 + \sigma & \gamma^2 & \gamma^2 & \dots & \gamma^2 \\ \gamma^2 & \gamma^2 + \sigma^2 & \gamma^2 & \dots & \gamma^2 \\ \vdots & & & & \vdots \\ \gamma^2 & \gamma^2 & \gamma^2 & \dots & \gamma^2 + \sigma^2 \end{pmatrix}_{n \times n} \oplus \dots \qquad \text{R times},$$

or more concisely, $Var(\vec{\epsilon}) = \sigma^2 I_{nR} + \gamma^2 I_R \otimes (1_n 1_n^*)$, where σ^2 is a measure of the error in measurement, γ^2 the error in type. When both errors occur, repeating the calculation above shows the variance for the replicant test has the form

$$Var(\hat{c}) = (\frac{\sigma^2}{n} + \gamma^2)(X^*X)^{-1}.$$

Thus while increasing n, the number of tires in each replicant set, will reduce the effect of the σ^2 errors, beyond a certain point it becomes necessary to reduce $(X^*X)^{-1}$ to affect the other error sources.

2.15 Summary

Over the course of the workshop, our team has developed some concrete solutions to the analysis problem in building tires with maximal uniformity. We have developed Monte Carlo methods to approximate optimal designs, found a number of general, explicit constructions for optimal designs, and demonstrated a blocking technique that addresses some of the complexity issues in the optimal designs that are relevant to the industrial practitioner. We have also investigated the robustness of the optimal designs, and examined the effects of replicating tests to improve performance of the analysis. MATLAB codes for all these investigation have been included in the appendix of the report.

In addition, the techniques described would be useful in a variety of vibrational problems requiring the determination of the contributions to the harmonic components of a periodic signal.

2.16 Appendix

Over the course of the workshop, a number of short MATLAB scripts were produced to test some ideas, establish conjectures, run Monte Carlo methods, and generally explore ideas on the computer. Some of the more complete scripts are included here. Briefly, they are

- `stats.m` (Figure 2.1) Monte Carlo method for finding near optimal designs. A measure is provided of how the performance improves for increasing number of tires. Three optimality criteria are used.

- `stat1.m` (Figure 2.2) Another Monte Carlo method, with graphical output to demonstrate how the performance improves.

- `stat2.m` (Figure 2.3) Monte Carlo search for optimal designs, with only a limited number of layers adjusted at random – in this cases, only four layers may be moved for any one tire. Graphical output of performance.

- `minangle.m` (Figure 2.4) Chooses a near-optimal design using a Monte Carlo method, with limited precision on the angles. A display of the resulting angles is provided, to see if there are any useful patterns appearing in these random designs.

```
% file stats.m
% Monte Carlo method to find good designs, three different criteria
% See how performance improves as number of tires increase
zmean = [];
zstd = [];
for k = 20:5:80
    z = [];
    for j = 1:100,
       c = ones(20,1)/20;
       x = exp(2*pi*i*rand(k,20));
       y = inv(x'*x);
        z = [z; real(max(eig(y))), real(max(diag(y))), real(c'*y*c)];
     end
     zmean = [zmean;mean(z)];
     zstd = [zstd;std(z)];
end
zmean
zstd
```

Figure 2.1: Script stats.m

- `cyclic1.m` (Figure 2.5) Builds an optimal design using the prime method, and verifies that it is optimal.
- `Lattice.m` (Figure 2.6) Builds an optimal design using the GLP method, then randomly perturbs the design angles to simulate a technician with limited angle control. A test of robustness of the GLP optimal design.
- `removerow.m` (Figure 2.7) Beginning with an optimal design, randomly removes a row from the design matrix to see what happens to the performance of the design. Another test of robustness.

```
% File stat1.m
% Monte Carlo method to find good designs
% Some graphics output to help us see
zmean = [];
zstd = [];
zmin = [];
c = ones(20,1)/sqrt(20);
for k = 20:5:100
 z = [];
 for j = 1:100,
  x = exp(2*pi*i*rand(k,20));
  y = inv(x'*x);
   z = [z; real(max(eig(y))), real(max(diag(y))), real(c'*y*c)];
     end
   zmean = [zmean;mean(z)];
   zstd = [zstd;std(z)];
   zmin = [zmin;min(z)];
end
zmean
zstd
zmin
semilogy((20:5:100)' , zmean)
title('zmean')
pause
semilogy((20:5:100)' , zstd)
title('zstd')
pause
semilogy((20:5:100)' , zmin)
title('zmin')
```

Figure 2.2: Script stat1.m

```
% File stat2.m
% Monte Carlo method to find good designs, three different criteria
% Only four angles may be changed at a time
zmean = [];
zstd = [];
zmin = [];
c = ones(20,1)/sqrt(20);
for k = 25:5:100
    z = [];
    for j = 1:100
       x =[];
       for kk=1:k
          x = [x;exp(2*pi*i*rand(1,20).*(5>randperm(20)))];
       end
y = inv(x'*x);
   z = [z; real(max(eig(y))), real(max(diag(y))), real(c'*y*c)];
end
    zmean = [zmean;mean(z)];
    zstd = [zstd;std(z)];
    zmin = [zmin;min(z)];
end
zmean
zstd
semilogy((25:5:100)',zmean)
title('zmean, 4 angles')
pause
semilogy((25:5:100)',zstd)
title('zstd, 4 angles')
pause
semilogy((25:5:100)',zmin)
title('zmin, 4 angles')
```

Figure 2.3: Script stat2.m

```
%  File minangles.m
% This finds some good matrix, and the display what angles were chosen
CC = 5 ;  % the number of components (columns of X)
KK = 5 ;  % the number of tires (rows of X)
zmin = .1;  % the minimum value so far
zangles = [];% the angles at the minimum, so far
for j = 1:100000
   a = rand(KK,CC);
   x = exp(2*pi*i*a);
   y = inv(x'*x);
   z = real(max(eig(y)));
   if (z>zmin)
     j
     zmin = z
     zangles = a;
   end
end
'done - look at the plot!'
plot(zangles,'.')
title('angles by tire - max eigs')
pause
plot(zangles','.')
title('angles by component - max eigs')
```

Figure 2.4: Script minangles.m

```
%  File cyclic1.m  the cyclic (prime) method of choosing angles
%
%  p = a prime = number of components (including fakes) eg p = 19
%  h = which harmonic  1 <= h <= p-1
p = 19;
h = 1;
t = mod(h*(0:(p-1))'*(0:(p-1)) , p)/p; % cyclic choice of angles, scaled 0 to 1
x = exp(2*pi*i*t);
y = x'*x;
z = max(max(abs(y - p*eye(p))))
%  typically, z (the error) is size 10^{-15}
%  so we conclude x'*x = pI  (multiple of the identity matrix)
%  its inverse has eigenvalues 1/p
```

Figure 2.5: Script cyclic1.m

```
% File Lattice.m
% Computes a design via the GLP method, then introduces error in the angles
zmean = [];
zstd = [];
zmin = [];
c = ones(5,1)/sqrt(20);
h = [1,2,10,13,16];
l = 1:21;
a = (2*l'*h - 1)/(2*21);
t = a - floor(a);
%x = exp(2*pi*i*t)
%y = x'*x
%eig(y)
for k = 0:0.01:0.1
 z = [];
   for j = 1:100
   x = exp(2*pi*i*(t + k*randn(21,5)));
   y = inv(x'*x);
   z = [z; real(max(eig(y))), real(max(diag(y))), real(c'*y*c)]
   end
   zmean = [zmean;mean(z)];
   zstd = [zstd;std(z)];
   zmin = [zmin;min(z)];
end
zmean
zstd
zmin
plot((0:0.01:0.1) , zmean)
title('zmean')
pause
plot((0:0.01:0.1) , zstd)
title('zstd')
pause
plot((0:0.01:0.1) , zmin)
title('zmin')
```

Figure 2.6: Script Lattice.m

```
% File removerow.m
% This takes one of the nice matrices, removes a row at random
% and sees what happens to the eigenvalues
s = 5; % number of components
n = 21; % number of tires
h = [1,2,10,13,16];   % from the book
k = 1:n;
t = rem((2*k'*h-1)/(2*n),1);  % the angles, between 0 and 1
x = exp(2*pi*i*t);   % matrix
xp = x';
xq = xp(:,[1,2,4,5,6,7,10,11,12,13,15,16,17,18,19,20,21]);
y = xq*(xq');
    min(real(eig(y)))
```

Figure 2.7: Script removerow.m

Bibliography

[1] K.-T. Fang and Y. Wang, *Number-theoretic Methods in Statistics*, Chapman & Hall, London, 1994.

[2] T.W. Körner, *Fourier Analysis*, Cambridge University Press, Cambridge, 1988.

[3] R.C. St. John and N.R. Draper, *D-Optimality for Designs: A Review*, Technometrics, Vol. 17, No. 1, February 1975, pp. 15–22.

Chapter 3

The Tennis Ball Problem

Andrej Bona,[3] Chris Bose,[9] Kell Cheng,[3] Wolgan Engler,[1] Minglun Gong,[7] Cyril Guyot,[8] Cristian Ivanescu,[8] John King[5] Dan Kenway,[1] Nathan Krislock,[6] Claude Laflamme,[3] Abid M. Malill,[4] Suresh Pillai,[2] Anamaria Savu,[8] Fridolin Ting,[8] Satoshi Tomoda[3]

3.1 Introduction

Stereoscopic vision is a well-established phenomenon: biological evolution showed its utility in ancient times. In this workshop, we have examined some subtleties and limitations in applying this old concept to an entirely new application: with modern technology, we attempt to track the position of an early segment of a flying object, and then extrapolate its later trajectory.

The concept is easily described as two cameras peering into a region of interest, through which a tennis ball is tossed along some trajectory. With knowledge of two simultaneous measurements from the cameras, one should be able to triangulate a measured ball position. With a set of such measurements, iterated over some timing set, one should be able, in principle, to estimate the trajectory of that ball. With aid from the mechanics of the situation, moreover, one should be able to extrapolate to later times.

In practice, one is confronted by several difficulties. Some imaging difficulties stem from the optics, and others derive from computer image discretization. Some ball tracking problems are due to background geometry, and others arise in the mechanics of ball motion. Practicalities greatly expand the scope and complexity of the original concept. The workshop study group divided the base problem into the following clearly definable elements:

[1]VisionSmart
[2]University of British Columbia
[3]University of Calgary
[4]University of Guelph
[5]University of Notingham
[6]University of Regina
[7]University of Saskatchewan
[8]University of Toronto
[9]University of Victoria

1. Camera parameters, including camera type, frame rate, optics and placement.

2. Camera calibration methods and detection of bad calibration conditions.

3. Ball image detection and centroid estimation: lighting and background conditions and further complications from shadows and motion blur.

4. Mechanics of basic ball motion including gravity, air resistance and ball spin.

5. Algorithms using extra data where only one camera reports.

6. Direct algorithms on raw data from individual camera data.

7. Error assessment and data weighting.

The results from the workshop study follow.

3.2 Optimal Locations of the Cameras

Our main objective is to determine the best approximation of the centroids of the balls in trajectory by choosing the best possible placement of two cameras. In order to determine the optimal locations of the cameras, one needs to first consider what the least requirements are. Thanks to VisionSmart's experiments, we are give the minimum number of pixels required to compute for the centroid within reasonable errors. VisionSmart estimates this number to be eight. We must also know what is the maximum allowable distance d from the camera to the ball to satisfy this pixel requirement. See Figure 3.2 for the calculation of d.

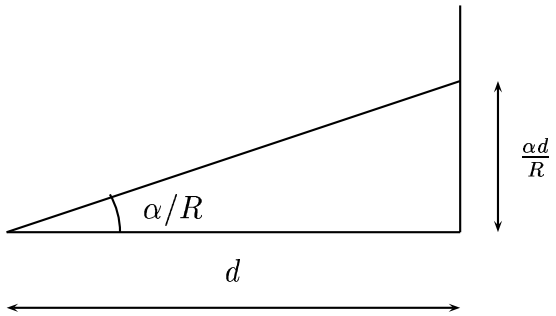

Figure 3.1: The maximum distance d

Let r be the radius of a tennis ball in meters, p the minimum number of pixels required on an image for the diameter of the tennis ball, α the lens visibility angle measured in radians , and R the resolution. In our case, these values are

$$r \approx 0.035, \text{ meters } p = 8, \text{ pixels } \alpha \approx 0.42, \text{ radians and } R \approx 500 \text{ pixels.}$$

At distance d, one pixel corresponds to approximately $\frac{\alpha d}{R}$. To guarantee p pixels on the image, we need $2r \geq p\frac{\alpha d}{R}$. So, $d \leq \frac{2rR}{\alpha p}$, that is, $d \leq \frac{2\times 0.035\times 500}{0.42\times 8} \approx 10.41m$. Hence, the maximum allowable distance from a camera to the tennis ball is around 10 meters.

Now that we have found the maximum effective distance from the camera to the ball, we can discuss the optimal locations for the cameras. One may first start by placing the two cameras opposite each other. A problem arises in this situation. The depth of the image is lost, which translates to difficulty in obtaining the ball trajectory. Instead of placing the cameras opposite each other, say, put them next to each other, pointing at the same direction. Again, the depth problem arises.

To solve the depth problem, one will need to place the cameras at some optimally chosen angle. In other words, the normal lines of the lenses form a certain angle other than multiples of 180 degrees. Note that we must intersect these lines; otherwise, we cannot recover the depth loss. So, let us try placing the cameras perpendicularly. That is, the directions of the two cameras form a right angle. It turns out that this is the optimal angle for getting the maximum number of pictures that provide the best approximation of the centroids.

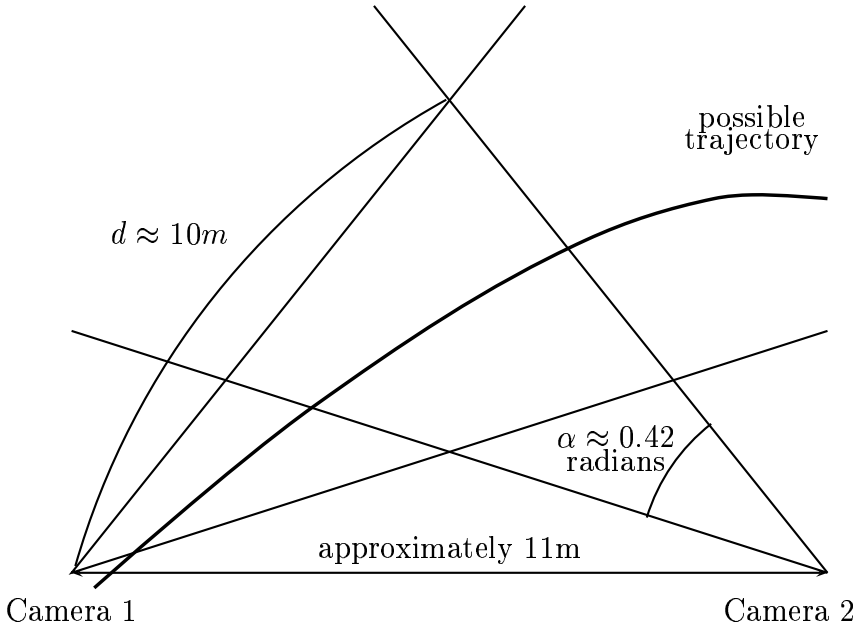

Figure 3.2: Camera positions viewed from above

Before proving the optimal angle is the right angle, we will sketch how we shall place the cameras relative to the tennis ball trajectory. With respect to the algorithm computing the trajectory described in Section 3.6, one of the cameras must capture the entire trajectory. The images from earlier parts of the trajectory provide more accurate information about the positions of the tennis ball, while the images from the later trajectory supply the effect due to spin of the ball and the gravitation. Thus, we place the cameras relative to the trajectory as in Figure 3.2.

To see the right angle is the optimal angle, consider Figure 3.2. In our case $\epsilon \approx 8.4 \times 10^{-4}$ radians ≈ 0.02 degrees and $m \approx 5.5$meters. If $\theta < \frac{\pi}{4}$, then the maximum possible error is

the distance of the two vertices which form a horizontal line segment. For $\theta > \frac{\pi}{4}$, the maximum possible error is the distance of the two vertices that form a vertical line.

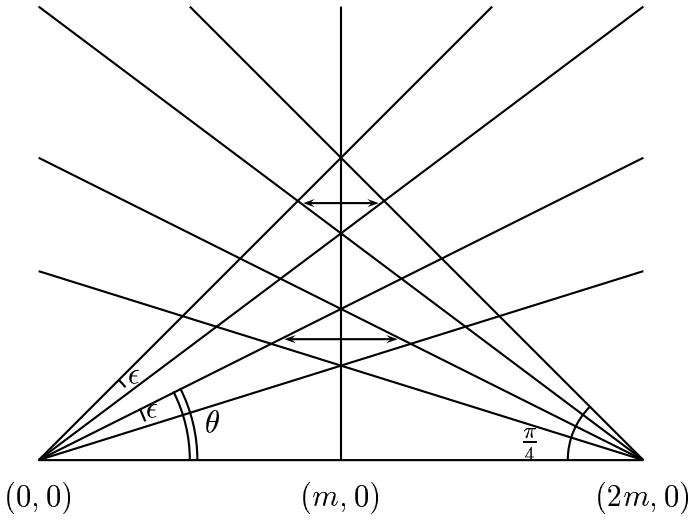

Figure 3.3: Camera relation

Without loss of generality, assume that $\theta < \frac{\pi}{4}$. We will find the x-coordinate of the left side point of the horizontal line segment. This point is the intersection of the line from the origin with slope $\tan(\theta)$ and the line with slope $-\tan(\theta - \epsilon)$ passing through the point $(2m, 0)$, that is, the intersection of the lines

$$\begin{aligned} y &= \tan(\theta) \cdot x, \text{ and} \\ y &= -\tan(\theta - \epsilon) \cdot x + 2m \tan(\theta - \epsilon). \end{aligned}$$

Thus, the x-coordinate of the required point is $\frac{2m\tan(\theta-\epsilon)}{\tan(\theta)+\tan(\theta-\epsilon)}$. It follows that the distance required is, by symmetry,

$$2\left(m - \frac{2m\tan(\theta-\epsilon)}{\tan(\theta)+\tan(\theta-\epsilon)}\right) = 2m\frac{\tan(\theta)-\tan(\theta-\epsilon)}{\tan(\theta)+\tan(\theta-\epsilon)}.$$

Thus, we need to show that $\frac{\tan(\theta)-\tan(\theta-\epsilon)}{\tan(\theta)+\tan(\theta-\epsilon)}$ is decreasing on $(0, \frac{\pi}{4})$.

By differentiating it, we get $\frac{-\tan(\theta)\sec^2(\theta-\epsilon)+\tan(\theta-\epsilon)\sec^2(\theta)}{(\tan(\theta)+\tan(\theta-\epsilon))^2}$ and we need to show that this is negative for all $\theta \in (0, \frac{\pi}{4})$. Since the denominator is non-negative, by simplifying the numerator, we get

$$\begin{aligned} & -\tan(\theta)\sec^2(\theta-\epsilon) + \tan(\theta-\epsilon)\sec^2(\theta) \\ = & -\frac{\sin(\theta)}{\cos(\theta)}\frac{1}{\cos^2(\theta-\epsilon)} + \frac{\sin(\theta-\epsilon)}{\cos(\theta-\epsilon)}\frac{1}{\cos^2(\theta)} \\ = & \frac{-\sin(\theta)\cos(\theta) + \sin(\theta-\epsilon)\cos(\theta-\epsilon)}{\cos^2(\theta)\cos^2(\theta-\epsilon)}. \end{aligned}$$

Since $-\sin(\theta)\cos(\theta) = -\frac{1}{2}\sin(2\theta)$ has negative values on $(0, \frac{\pi}{4})$, the original expression $2m\frac{\tan(\theta)-\tan(\theta-\epsilon)}{\tan(\theta)+\tan(\theta-\epsilon)}$ is decreasing on $(0, \frac{\pi}{4})$ as required.

Remark:

Although we have proved that the most accurate information is retrieved when the normal lines of the lenses intersect at a right angle, it maybe the case, in practice, that the closer the better. That is, the calculated maximum error may not occur in real life. To find this, we need to run more experiments. If this is indeed the case, then we would have to rearrange the orientation of the two cameras.

3.3 Calibration

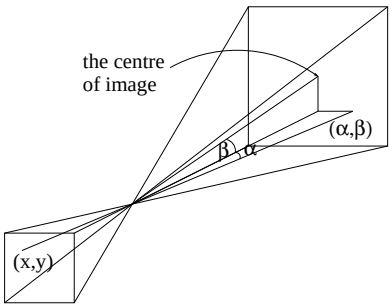

Figure 3.4: Optical Centre

To account for the lens distortion, we have to calibrate the data. We have to determine (α, β) as a function of (x, y), where (α, β) is an angular coordinate of a ray from the optical centre to a point on the grid, and (x, y) is a pixel coordinate of this point. This can be done by fitting a polynomial of sufficient degree, say four, through the points $((\alpha, \beta) \leftrightarrow (x, y))$. The second-degree polynomial is enough according to the VisionSmart's experiments. We chose four just to be on the safe side.

The next step is to fix the axis with respect to which we determined (α, β). This is done by mounting a pair of laser pointers on the camera in such a way that they are parallel to the axis and collinear with the reference axis. We can set the lasers in such a way that $x_1 = x_2$ (see Figure 3.3).

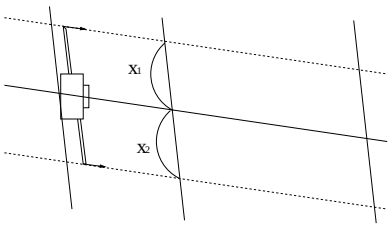

Figure 3.5: Calibration

This way, we can use the cameras in different places without recalibrating at the new location. Knowing these calibration polynomials, we can get the 3-dimensional coordinate of an object, in particular, the centroid of the tennis ball, by the standard surveyor's method.

3.4 Determining the Centroids (Lighting, Background, Shadows)

3.4.1 Determining the centroid of the ball

In the situation of a uniformly lit ball, we can compute the "centre of intensity" of the subtracted picture.

$$(x, y) = \frac{\sum w_{ij}(i, j)}{\sum w_{ij}},$$

where (x, y) is the coordinates of the centre and (i, j) is the coordinates of a pixel with intensity w_{ij}.

The major problem of finding the centroid of a ball lies in the shadowed portion of the ball. Improper lighting of the environment causes the ball to be bright on one side and dark on the other. (See Figure 3.4.1)

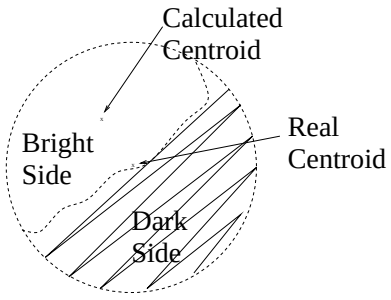

Figure 3.6:Shadow Image

If one were to calculate the centre of intensity of the ball in Figure 3.4.1, then one would find that it would be off the position of the real centroid. Of course, the centroid of the ball means the actual centre of mass of the ball.

3.4.2 Suggested solution

1. Placing lights directly behind cameras
   Of course, we cannot put the lights directly behind the cameras, otherwise all the light will be blocked off by the cameras. Therefore, we suggest putting the lights just above or below the cameras.

   The reason why we want to do this is because of the following situation.

Consider Figure 1.

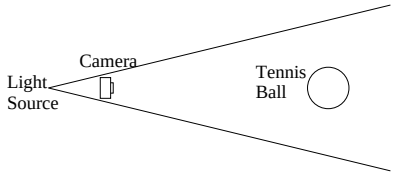

Figure 3.7: Light Source

Since the light is directly behind the camera, the reflected light will bounce directly back to the camera. This will certainly remove the shadow from the ball, and hence we will see a completely circular ball, not the shadowed version in Figure 3.4.1.

We can even improve this setup by using only green light for one light source and red light for the other light source. This can be easily implemented by putting a filter in front of the light sources. In addition, we can put a filter in front of the cameras so that: **each camera only sees the light reflected from its own illumination source.** All we can detect is green light and red light, and hence: **each camera image will be a full disc with centre approximately at the centre of intensity** allowing us to find the position of the centroid more easily.

2. Using black ball and white background
Lighting of the environment plays a major role in the production of the shadows on the tennis ball. As we saw above, this causes major problems in finding the centroid. We can reduce the difficulty of this problem of lighting by using a black ball and a white background. This is because the shadow on the ball is black. That is, the difference between the actual colour of the tennis ball and the shadow is small. Therefore, what we expect to see is a circular black ball in the image.

3. Tracing the shadow of the ball
As well as tracking the ball itself, we can track the shadow of the ball on the background. This way, we can get two estimates on the centroid of the ball. For instance, if the shadow travels on the ground, we can get another estimated xy-coordinate of the centroid of the ball, thereby improving the estimate for the centroid of the ball. Of course, we need to arrange the lighting so that there will be a shadow.

3.5 Dynamics

3.5.1 Formulation

Under appropriate assumptions, the trajectories can be taken to be governed by the sixth order system (in which $\hat{\boldsymbol{k}} = (0,0,1)$ and g is the acceleration due to gravity):

$$\begin{aligned} m\frac{d\boldsymbol{v}}{dt} &= -C_D\boldsymbol{v} + C_L\boldsymbol{\omega}\times\boldsymbol{v} - mg\hat{\boldsymbol{k}}, \\ I\frac{d\boldsymbol{\omega}}{dt} &= -C_A\boldsymbol{\omega}, \end{aligned}$$

where m and I are the mass and the moment of inertia of the sphere and $\boldsymbol{v}$ and $\boldsymbol{\omega}$ are its velocity and angular velocity. Asymptotic expressions (Rubinow and Keller, 1961) are available for the drag coefficient C_D, the lift coefficient C_L (corresponding to the Robins-Magnus effect) and the angular drag coefficient C_A in the limit of small Reynolds number; however, for the high Reynolds numbers of interest here, more phenomenological expressions are adopted, and we take

$$C_D = c_D|\boldsymbol{v}|, \qquad C_L = c_L, \qquad C_A = 0 \ \text{ (negligible angular drag)}$$

with c_D and c_L constant. These expressions follow from experimental fits (allowing c_D and c_L to depend on $a|\boldsymbol{\omega}|\big/|\boldsymbol{v}|$, where a is the radius of the sphere) as given by Stephanek (1988) for a tennis ball; we anticipate that $a|\boldsymbol{\omega}|\big/|\boldsymbol{v}|$ will be negligible for most cases of interest here. Taking initial conditions

$$\text{at } t = 0 \qquad \boldsymbol{x} = \boldsymbol{X}, \quad \boldsymbol{v} = \boldsymbol{V}, \quad \boldsymbol{\omega} = \boldsymbol{\Omega},$$

we thus arrive at

$$\begin{aligned} m\frac{d\boldsymbol{v}}{dt} &= -c_D|\boldsymbol{v}|\boldsymbol{v} + c_L\boldsymbol{\Omega}\times\boldsymbol{v} - mg\hat{\boldsymbol{k}}, \qquad \frac{d\boldsymbol{x}}{dt} = \boldsymbol{v}, \\ \text{at } t = 0 \qquad & \boldsymbol{x} = \boldsymbol{X}, \qquad \boldsymbol{v} = \boldsymbol{V} \end{aligned} \tag{3.1}$$

which is the system we will study.

3.5.2 Small-time expansion

From (3.1) we have as $t \to 0$ that

$$\begin{aligned} \boldsymbol{v} &= \boldsymbol{V} + \frac{1}{m}\Big(-c_D|\boldsymbol{V}|\boldsymbol{V} + c_L|\boldsymbol{V}|\boldsymbol{\Omega}\times\boldsymbol{V} - mg\hat{\boldsymbol{k}}\Big)t + O(t^2) \\ \boldsymbol{x} &= \boldsymbol{X} + \boldsymbol{V}t + \frac{1}{2m}\Big(-c_D|\boldsymbol{V}|\boldsymbol{V} + c_L\boldsymbol{\Omega}\times\boldsymbol{V} - mg\hat{\boldsymbol{k}}\Big)t^2 + O(t^3) \end{aligned} \tag{3.2}$$

The quadratic form (3.2) will be a good approximation provided

$$t \ll \frac{m}{c_D|\boldsymbol{V}|}\,, \qquad \frac{m}{c_L|\boldsymbol{\Omega}\times\boldsymbol{V}|}\,;$$

if the drag and lift terms are negligible, it will be adequate for all time but otherwise the non-linear system (3.1) should be solved for the later stages of the trajectory, rather than extrapolating the quadratic form (3.2). A best fit of the full data (using that from all the cameras at once) to the quadratic form (3.2) (using the known values of g and m) would enable the nine independent quantities from

$$\boldsymbol{X},\ \boldsymbol{V},\ c_D,\quad c_L\boldsymbol{\Omega}\times\boldsymbol{V} \tag{3.3}$$

to be estimated; alternatively (and preferably) c_D could be measured beforehand and the remaining eight quantities fitted ($\boldsymbol{X}$ may also be known *a priori*). Note that $\boldsymbol{\Omega}$ appears in (3.2) only as in the combination $c_L\boldsymbol{\Omega}\times\boldsymbol{V}$, so no information can be obtained in this way about the component of $\boldsymbol{\Omega}$ parallel to $\boldsymbol{V}$; moreover,measuring c_L beforehand would be of little benefit since only the combination $c_L\boldsymbol{\Omega}$ is relevant, with $\boldsymbol{\Omega}$ unknown *a priori*. In fact, it may in practice be best to neglect $c_L\boldsymbol{\Omega}$ in performing the fitting, at least in the first instance (if $\boldsymbol{v}$ deviates noticeably from $\boldsymbol{V}$, the lack of information on $c_L\boldsymbol{\Omega}\cdot\boldsymbol{V}$ may become a significant drawback). Fitting the data to the parameter set (3.3) (instead of using the data to estimate $\boldsymbol{x}(t)$ directly) would enable information from a single camera to be used at a given location (by projecting (3.2) onto the appropriate surface), so all the available data could be treated on the same footing and synchronization of the cameras would not be necessary (indeed, in principle a single camera could suffice); moreover, since (3.2)–(3.3) aim to incorporate in a systematic fashion the dynamics which determine the trajectories, it is hoped that errors resulting from the various measurements would to some degree be self-correcting (since consistency with the governing physics is being demanded), rather than accumulating.

3.5.3 Later times

The available evidence is that the drag terms in (3.1) are non-negligible (the lift (spin) terms may be small), so a parabolic fit will be inappropriate over the later stages of the trajectory. Thus (3.1) should be solved numerically until impact occurs, with the values of c_D, $\boldsymbol{X}$, $\boldsymbol{V}$ and $c_L\boldsymbol{\Omega}\times\boldsymbol{V}$ estimated as above. Evaluation of $c_L\boldsymbol{\Omega}\cdot\boldsymbol{V}$ would require more terms in the expansion (3.2) to be taken (these can be readily calculated) and used in the fitting; however, it is likely that the contribution of this component will be negligible, in which case it can be set to zero.

3.5.4 Recommendations

In summary, we propose:

(1) Fitting all the available camera data to the parameter set (3.3), or a subset thereof (whereby c_D and/or $\boldsymbol{X}$ are determined beforehand; it should be relatively straightforward to estimate c_D from separate experiments), the other fitting parameters being the distance of the sphere to the camera for each data point for each camera.

(2) Extrapolating to give the later stages of the trajectories by using the above estimates as parameter values and initial conditions in a rapid and straightforward numerical solution of the nonlinear system (3.1).

3.6 An Iterative Method for Computing Full Trajectories

In the previous section it was observed that the sequence of ball images from a single camera already contains more information than that which is used in the simple-minded triangulation approach for determining the trajectory of the ball. In this section we discuss an iterative algorithm which aims to reconstruct the 'missing coordinate' from the sequence of images obtained from one dominant camera which sees a long segment of the ball trajectory, while incorporating the triangulation data available in overlap with the second camera. Essentially we combine the original approach of VisionSmart, reworked to incorporate the full dynamics, with the projection method described in Section 3.5.

Briefly, at the k-th iteration of the routine, estimates are given on the 'missing coordinate' from each camera. Nonlinear regression against the parametric family of true dynamic trajectories is performed to produce a new table of pseudo-data for the trajectory. These computed pseudo-data are then compared with the camera images providing updated estimates on missing coordinates. These new estimates provide input for the next iteration. There are a number of reasonable options for stopping rules for the algorithm which we will discuss.

To begin, we assume that the following data are known

- Camera 1 tracks the entire trajectory (100 images approximately) (roughly) from behind. To establish the orientation, we assume Camera 1 is fixed in the yz-plane. The ball is launched in such a way that its trajectory lies in the positive orthant of $\mathbf{R}^3$.
- Camera 2 is fixed in the xz-plane and records at least 8 images from the first 20 images captured by Camera 1. Thus we have at least 8 images in the overlap for the two cameras
- Time $t = 0$ is defined to coincide with the first image in Camera 1.
- 'Error cones' on each data point obtained from all the images recorded by each camera overlap are assumed to be known from preprocessing of the data. We denote these as $\mathcal{C}_k(j)$, for the error cone associated with the k-th image in the j-th camera, $k = 0, 1, \ldots, 100,\ j = 1, 2$. The error cones are convex and compact (since we know the dynamics is restricted to lie within some big closed ball centered at the origin) and their vertex angles depend on the distance between adjacent pixels in the camera, the distance from the focal plane in the camera to the centre of the lens, and on the acuity of our centre of mass calculation from the pixelated images which are our raw data. Where the k-th image is missing from Camera 2, we may as well define $\mathcal{C}_k(2) = \mathbf{R}^3$.

Moreover, we adopt the assumption from the previous section that

- The spin vector $\mathbf{\Omega} =< \omega_1, \omega_2, \omega_3 >$ is assumed constant throughout the trajectory.

With these assumptions, we consider the full dynamic equations (3.1). Since the ball trajectories are very flat, we adopt the linearizing assumption that the speed $|\mathbf{v}|$ is nearly constant over the time of flight, so the drag term is well approximated by

$$c_D|\boldsymbol{V}||\boldsymbol{v}|. \tag{3.4}$$

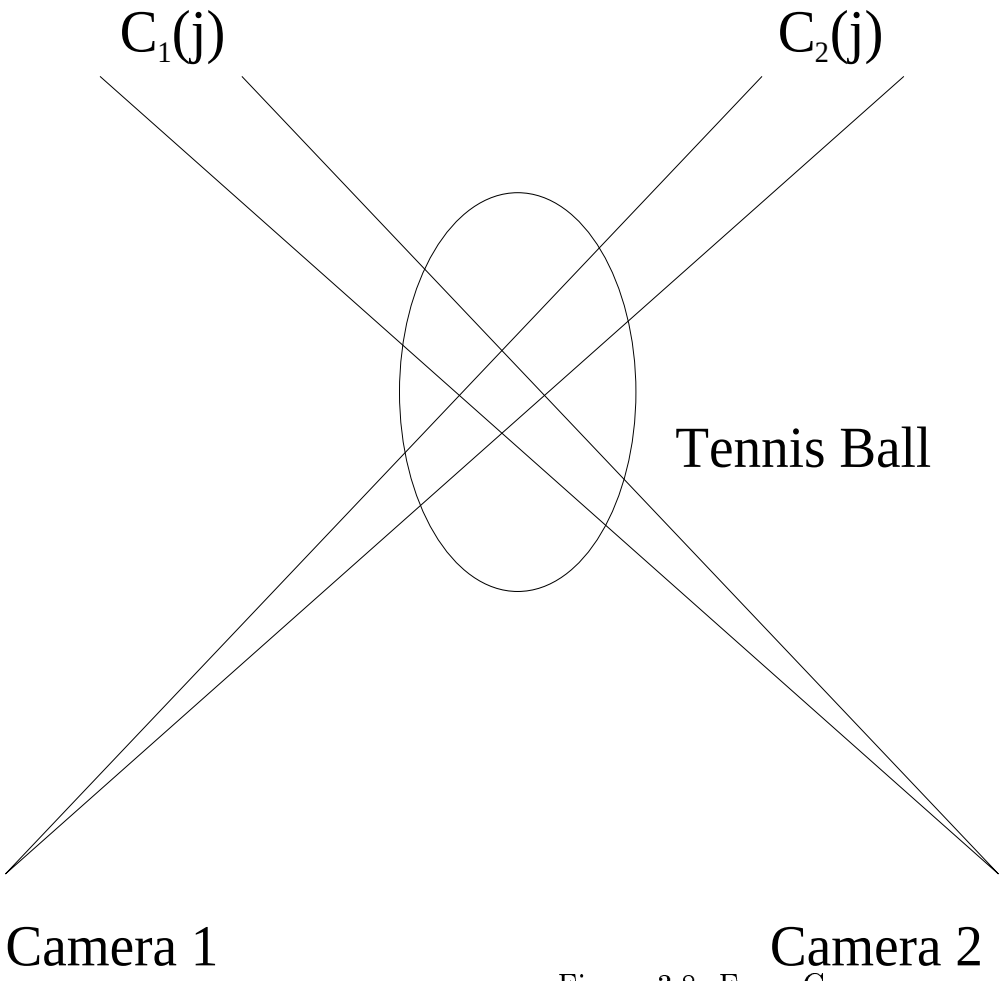

Figure 3.8: Error Cones

Recall that $\mathbf{V}$ is the ball's initial velocity.
In component form we are left with the equations

$$\begin{aligned}
\ddot{x}_1 &= -(\frac{c_D|\boldsymbol{V}|}{m})\dot{x}_1 - \frac{c_L}{m}\omega_3\dot{x}_2 + \frac{c_L}{m}\omega_2\dot{x}_3 \\
\ddot{x}_2 &= -(\frac{c_D|\boldsymbol{V}|}{m})\dot{x}_2 + \frac{c_L}{m}\omega_3\dot{x}_1 - \frac{c_L}{m}\omega_1\dot{x}_3 \\
\ddot{x}_3 &= -g - (\frac{c_D|\boldsymbol{V}|}{m})\dot{x}_3 + \frac{c_L}{m}\omega_1\dot{x}_2 - \frac{c_L}{m}\omega_2\dot{x}_1
\end{aligned}$$

which we write in compact form as

$$\ddot{\boldsymbol{x}} = m^{-1}A\dot{\boldsymbol{x}} - \boldsymbol{g}$$

where

$$A = \begin{bmatrix} -c_D|\boldsymbol{V}| & -c_L\omega_3 & c_L\omega_2 \\ c_L\omega_3 & -c_D|\boldsymbol{V}| & -c_L\omega_1 \\ -c_L\omega_2 & c_L\omega_1 & -c_D|\boldsymbol{V}| \end{bmatrix},$$

$$\boldsymbol{g} = \begin{bmatrix} 0 \\ 0 \\ g \end{bmatrix}$$

where c_D and c_L are the drag and lift coefficients (constants) respectively for the ball, m is its mass, and g is the constant acceleration due to gravity.

Since $A = -c_D|\boldsymbol{V}|I - c_L\Lambda_\Omega$ where it is expected that $c_D|\boldsymbol{V}| \gg c_L||\Lambda_\Omega||$, A is invertible and one easily solves the linear equation in closed form:

$$\boldsymbol{x}(t) = \boldsymbol{X} + mA^{-1}\boldsymbol{g}t - (e^{m^{-1}At} - I)(m^2A^{-2}\boldsymbol{g} - mA^{-1}\boldsymbol{V}) \tag{3.5}$$

We assume that the ball's mass m and the value of g is known (but see later discussion), however there remain 10 independent free parameters in this solution, represented by the initial position $\boldsymbol{X}$, velocity $\boldsymbol{V}$ the spin vector $c_L\boldsymbol{\Omega}$ and the drag coefficient c_D. VisionSmart's experience is that it is very difficult to estimate these parameters accurately enough from the few images in the camera overlap. For example, they typically reported a value of the gravitational constant g to be low by a factor of 2 when computed from their limited data! This is not too surprising since a quick calculation shows that the expected effect of gravitational acceleration (for example) is to introduce a curvature in the trajectory on the order of one pixel over the duration of the camera overlap. That is why we believe it is essential to incorporate the additional camera images in some way to improve accuracy, especially in the accelerative coefficients.

3.6.1 The first estimate on positions

The routine we recommend begins with a preliminary no-spin estimate on the missing data from the second camera. At this point we also use a textbook value for the drag coefficient for a sphere

$$c_D = 0.22D^2 \approx 0.00108$$

where $D = 0.07$ is the diameter of the ball. We adopt the no-spin condition $\boldsymbol{\Omega} =< 0, 0, 0 >$ and $g = 9.8m/s^2$. We are given the spatial data from the 8 images in the camera overlap, which we display as a table:

t=k	$x_1(k)$	$x_2(k)$	$x_3(k)$
5	$x_1(5)$	$x_2(5)$	$x_3(5)$
6	$x_1(6)$	$x_2(6)$	$x_3(6)$
$\vdots$	$\vdots$	$\vdots$	$\vdots$
12	$x_1(12)$	$x_2(12)$	$x_3(12)$

Using this data, along with equation (3.5) in a nonlinear least squares approximation we obtain our first estimate on the positions for all times $t = 0, 1, \ldots, 100$. Denote this new set of spatial data by

$$\hat{x}_1(k) = x_1|_{t=k}, \quad \hat{x}_2(k) = x_2|_{t=k}, \quad \hat{x}_3(k) = x_3|_{t=k}.$$

Now, it is possible that some of these computed data points are not consistent with the camera observations, so before we proceed with the next iteration we need to adjust these computed data as follows. The camera observation of the k-th image indicates the position of the centroid

of the ball at $t = k$ only within the error cone $\mathcal{C}_k$ emanating from the camera, centered on a ray through the centre of the camera lens and the position of the computed centroid on the camera focal plane. The true centroid lies somewhere inside $\mathcal{C}_k$. It may happen that our computed spatial data, for some $t = k$ fails to lie inside the error cone determined by the k-th image recorded by some camera. In that case we perturb our computed spatial data to lie in the error cone:

$$(x_1(k), x_2(k), x_3(k)) = (x_1, x_2, x_3)_{\min}$$

where $(x_1, x_2, x_3)_{\min}$ is the unique solution to

$$\text{minimize } |(x_1, x_2, x_3) - (\hat{x}_1(k), \hat{x}_2(k), \hat{x}_3(k))|$$

subject to $(x_1, x_2, x_3) \in \mathcal{C}_k$. For values of $k = 5, 6, \ldots, 12$ we use the constraint $(x_1, x_2, x_3) \in \mathcal{C}_k(1) \cap \mathcal{C}_k(2)$ (convex, cpct) where, again, $\mathcal{C}_k(j)$ is the k-th image error cone for the j-th camera. In this way we determine our first estimates on the 100 ball positions:

t=k	$x_1^{(1)}(k)$	$x_2^{(1)}(k)$	$x_3^{(1)}(k)$
0	$x_1^{(1)}(0)$	$x_2^{(1)}(0)$	$x_3^{(1)}(0)$
$\vdots$	$\vdots$	$\vdots$	$\vdots$
100	$x_1^{(1)}(100)$	$x_2^{(1)}(100)$	$x_3^{(1)}(100)$

as well as the first estimate on the drag parameter $c_D^{(1)}$.

3.6.2 The iteration loop ($n = 2, 3, \ldots$)

The input is the table of positions computed from the previous step:

t=k	$x_1^{(n)}(k)$	$x_2^{(n)}(k)$	$x_3^{(n)}(k)$
0	$x_1^{(n)}(0)$	$x_2^{(n)}(0)$	$x_3^{(n)}(0)$
$\vdots$	$\vdots$	$\vdots$	$\vdots$
100	$x_1^{(n)}(100)$	$x_2^{(n)}(100)$	$x_3^{(n)}(100)$

We should note at this point that there is non-uniform reliability of this data for geometric reasons. Since the cross sectional area of the error cone for each camera grows linearly in distance from the camera, the accuracy of the computed centroid decays linearly in distance from the camera. For the images $k = 13, 14, \ldots, 100$ the distance of the ball from camera 1 can be estimated by $x_1^{(n)}(k)$ so we should weight the k-th data point in the above table in proportion to $(x_1^{(n)}(k))^{-1}$. For simplicity we weight the data in the camera overlap uniformly as C_2 $k = 5, 6, \ldots 12$ and for $k = 0, 1, \ldots, 4$ uniformly as C_1. We should assume $C_1 < C_2$, and, again for simplicity we suggest $C_1 = \frac{C_2}{2}$. Finally, we need a probability distribution for the weights, which determines C_2:

$$(5/2 + 8)C_2 + \sum_{k=13}^{100} C_2(x_1^{(n)}(13))(x_1^{(n)}(k))^{-1} = 1$$

giving weights for the n-th iteration dataset

$$w^{(n)}(k) = \begin{cases} \frac{C_2}{2} & \text{if } 0 \leq k \leq 4 \\ C_2 & \text{if } 5 \leq k \leq 12 \\ C_2(x_1^{(n)}(13))(x_1^{(n)}(k))^{-1} & \text{if } k \geq 13 \end{cases}$$

Using the n-th data set, with the weights above, nonlinear least squares approximation on equation (3.5) where $\boldsymbol{V}$, $\boldsymbol{X}$, $c_L\boldsymbol{\Omega}$, c_D and g are the parameters to be fitted, yields the $(n+1)-$st estimate on the ball positions:

$$\hat{x}_1^{n+1}(k), \hat{x}_2^{n+1}(k), \hat{x}_3^{n+1}(k).$$

Again we use the hats to indicate that some of these computed data points may not be consistent with the camera images, so before we leave this iteration we perturb the data to lie in each camera's error cones (in exactly the same way which was done for the first iteration). This final step yields the output for the loop:

t=k	$x_1^{(n+1)}(k)$	$x_2^{(n+1)}(k)$	$x_3^{(n+1)}(k)$
0	$x_1^{(n+1)}(0)$	$x_2^{(n+1)}(0)$	$x_3^{(n+1)}(0)$
$\vdots$	$\vdots$	$\vdots$	$\vdots$
100	$x_1^{(n+1)}(100)$	$x_2^{(n+1)}(100)$	$x_3^{(n+1)}(100)$

along with estimates on the dynamic parameters $\boldsymbol{V}^{(n+1)}$, $\boldsymbol{X}^{(n+1)}$, $c_L\boldsymbol{\Omega}^{(n+1)}$, $c_D^{(n+1)}$ and $g^{(n+1)}$. Note that we are keeping the gravitational constant as an unknown parameter in the iterative procedure. A quick count gives 11 free parameters and 303 data points as input.

3.6.3 Stopping criterion

We suggest monitoring the value of $g^{(n)}$ and stopping when $g^{(n)} \approx 9.8$. It may happen that the spatial data stabilize before this is reached, in which case we should also stop, but then it is probably reasonable to question the initial data from the camera images. Of course any real implementation of this routine should have a maximal number of iterations fixed. The size of this cutoff could be determined by running the algorithm on various sample data sets.

3.6.4 Fine tuning

We have made some choices in the description above which should be experimented with in any implementation.

- The first iteration uses uniform weights for the data points coming from the camera overlap, when, in fact, geometric considerations give a non-uniform reliability to each point. The reliability could be estimated in order to apply a weighted least squares approximation, even at this first step.
- Similarly, there is quite a bit of flexibility in the weights profile $w^{(n)}(k)$ which might be used to advantage.

- It is possible to make estimates on the spin and drag effects from the camera 1 images only. These could be used in the first iteration in order to get more accurate first-run estimates on the positions of the ball for large k.

Bibliography

[1] Rubinow, S. I.; Keller, Joseph B. *The transverse force on a spinning sphere moving in a viscous fluid. Journal of Fluid Mechanics* **11** 1961 447 – 459.

[2] Stepanek, Antonin. *The aerodynamics of tennis balls – The topspin lob. The American Journal of Physics* **56**, 138 (1988).

Chapter 4

Designing Incentive-Alignment Contracts in a Principal-Agent Setting in the Presence of Real Options

Tom Cottrell[1] Dan Calistrate[1]

Abstract

We develop a model of incentive compensation for optimal upgrades supplied by an outsourced Information Technology department. We first consider the problem when the rate of technological development is certain and there are no information asymmetries between the parties. We extend this to allow private information between the principal and an agent acting as an external supplier of information technology upgrades. Based on the model in these simple circumstances, we then model uncertain technological improvement, where improvements evolve as geometric Brownian motion, and there is benefit to flexibility in the timing of the upgrade. We are aware of contracts, known as "evergreen upgrades" where a principal pays for upgrades at specified intervals. We find little support for such a contract in our model, and the loss of flexibility in the timing of upgrades is puzzling. The Stern-Stewart problem encourages us to consider just such instances, where contracts limit flexibility that it may be in the interest of both parties to retain. We conclude with a consideration of the wide range of future work needed in this area.

4.1 Introduction

The Stern-Stewart problem description provides a rich context for the development of incentive contracts in the presence of asymmetric information and the need for flexibility. The problem statement allows a broad range of problem definition and solution, and there is substantial scope for further work. For the purposes of this discussion, we will pursue the two general approaches to the problems suggested. These are:

[1]University of Calgary

- Computer upgrade contracts with outsourced IT

- Incentive contracts where there is need for flexibility

These problems share the difficulty that the principal requires private information held by the agent in order to make the appropriate trade-offs for an optimal decision. It is possible that the agent will be unable or unwilling to reveal this information. First, simply, it may be expensive for the agent to (credibly) relay the information to the principal. This would especially be the case for the first problem, where a firm that has decided to outsource information technology. At the time of the outsourcing decision, the principal may be well informed about the technology and its implications for application within the firm. Later, however, as the firm no longer retains the in-house expertise to evaluate technology, the firm may lack the ability to map emerging technological advancements to business objectives.

Second, it may not be in the agent's interest to reveal information in the most efficient way. Depending on the nature of the compensation agreement, the agent may have an incentive to delay or manipulate the way that verifiable information arrives to the client. The implications of this are described and developed below.

4.2 A Comment on Related Literature in Finance and Economics

This section is intended to provide a brief introduction to the kinds of issues raised in the Stern-Stewart problem, and to provide some reference to the academic work germane to these issues. While option pricing models as applied to real assets has received attention in theoretical work ([5], [18], [19], [9]) there remain persistent concerns that pragmatic difficulties with incentives and compensation have resulted in non-implementation of these models in corporate settings. There is research on the problem of agency generally, and particularly delegation from the headquarters out to divisions, in the development of capital budgeting rules (see [13], [1], [6], [14], [10], [11], [3], [12] and [17], where the latter two consider the effect of influence costs in capital budgeting). There is an extensive literature on the nature of incentive compensation, particularly for CEO's, and its effect on capital budgeting ([16], [7], [8]), and both theoretical and empirical work has expanded, particularly over controversial bonus and stock-option payouts to CEOs.. In this literature, Stern-Stewart's EVA model (e.g. [15]) plays an important role. Clearly, the Stern-Stewart problem requires significant background to understand the development of related models in finance and economics, and careful application of the results of these models to an understanding of the problem of option value in principal agent problems. In this short report, we can only begin to make progress by narrowing our scope. While a more thorough treatment is certainly warranted, we will have to settle in the time and space alloted for merely beginning the exploration.

4.3 Outsourced Upgrades

Expert advice on upgrade administration is an example of the more general problem of incentive contracts with asymmetric information. At issue is the requirement that for contracts to be enforceable, they must also be verifiable. If there is no need for contract enforcement, then the contract can be written broadly to allow the agent flexibility in the effort levels and timing of the implementation of the contract. If, however, it is in the interest of the agent to choose policies that are, at least potentially, at odds with the interests of the principal, then verifiability is essential for contract enforcement.

We first illustrate the difficulties with flexible contracts in the upgrade problem, and then model the considerations relevant to the optimal solution to the problem.

4.3.1 Upgrades under certainty with symmetric information

We employ a partial equilibrium model for optimal timing of upgrades. The model assumes that exogenous technological improvements lead to greater and greater computing power at constant prices. The firm's choice is to determine when to upgrade given this exogenous rate of improvement α, the firm's private adoption cost c, the contract price k and the firm's opportunity cost of investment δ. Upgrades occur at discrete times t_i, and increase the computing power of the firm by $\alpha\,(t_i - t_{i-1})$. On any interval where there is no upgrade $[t_i, t_{i+1})$, the benefit is:

$$\int_{t_i}^{t_{i+1}} e^{-\delta\, t} \alpha\ (t_i - t_{i-1})\, dt$$

So that, if there is no limit to the number of upgrades and the benefit, we have the following:

$$V_0 = \int_0^{\infty} \int_{t_i}^{t_{i+1}} e^{-\delta\, t} \alpha\ (t_i - t_{i-1})\, dt$$

Since the benefits of upgrading are exponential and the cost is linear, the firm will instantaneously upgrade for a reasonable set of parameters. For this discussion, however, we limit the upgrade to a discrete choice of upgrade times. We thus seek the optimal conditions for the first upgrade. For that first time interval, the value of the upgrade is measured by the improvement since time $t_0 = 0$, where we normalize the initial computing power to 0. If the cost of the upgrade is $k + c$, and there is value to delaying this outlay, the following equation represents the initial decision:

$$\max_t V_0 = e^{-\delta\, t} \left(\frac{\alpha\, t}{\delta} - (k + c) \right)$$

The principal maximizes the value of the firm by choosing the appropriate set of upgrade times $\mathbf{t}$. The first order condition provides the optimal for positive t given the non-negativity constraints on the remaining parameters.

$$t^* = \frac{\alpha + \delta^2\,(k + c)}{\alpha\, \delta}$$

Figure 4.1 illustrates the optimal initial upgrade times for discount rates of 15%, 20% and 30% as the exogenous rate of technological improvement increases, other parameters fixed. Higher

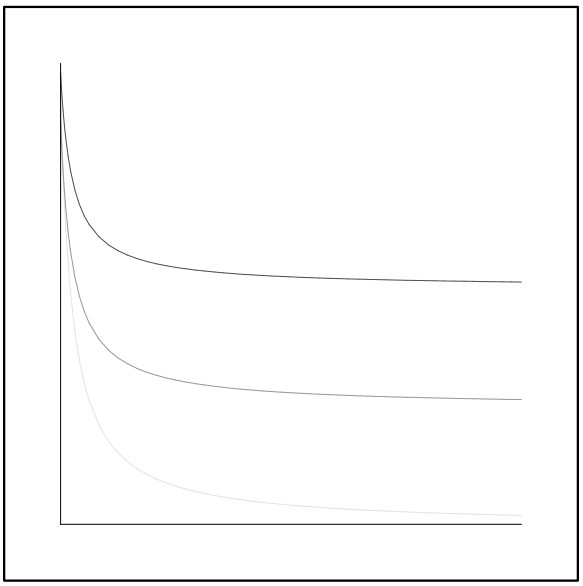

Figure 4.1: τ versus technology improvement α

discount rates lead to faster adoption rates. This is in part due to the parameterization where technological improvements overcome the benefits from delaying the risk of incurring the cost.

Figure 4.2 plots the optimal adoption times for a range of rates of exogenous technological improvement $\left(\frac{1}{2}, 1, 2\right)$ against the discount rate. Clearly, the optimal adoption is more sensitive to the discount rate than to the (linear) rate of technological improvement.

4.3.2 Outsourcing with full information

When both parties know the parameters and the nature of the pay-offs, it is straightforward to write an upgrade contract with an outsourced IT department. The optimal delay between upgrades, and the value of the technology at that time, is easily observed. Multiple upgrades in this framework would occur at multiples of t^*, so that $t_i = it^*$. This result is trivial in the full information case. With information symmetry, we expect there would be little value added by the outsourced IT department, and of course this is where the problem lies.

4.3.3 Agency with asymmetric information on upgrades

In the outsourcing decision, the IT supplier becomes an agent of the firm, with a separation of information on model parameters. Once a firm has decided to outsource information technology functions, it is no longer in a position to map the value of the technology improvements to its business needs, so that it can no longer know α. The agent will not know the principal's private adoption cost c or its discount rate δ. The timing of the upgrade will depend upon the incentives in the outsourcing contract.

The principal offers the agent a menu of contracts in order to induce the agent to identify the right time for the upgrade. If this is a one-shot game, then the optimal contract is simply

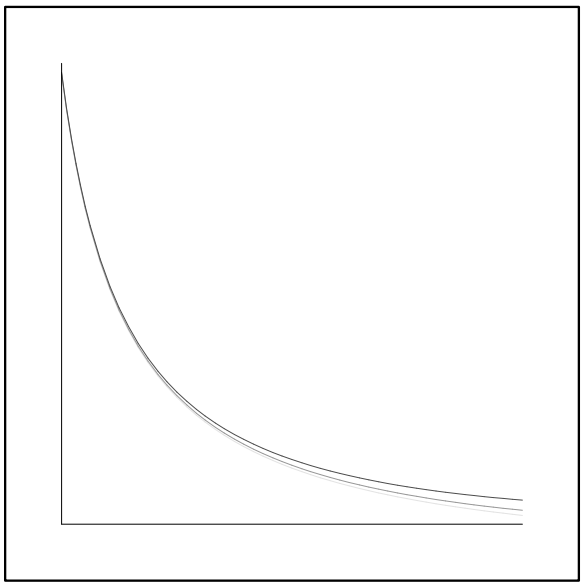

Figure 4.2: τ versus discount rate δ

a fraction of the benefit that the principal receives for the upgrade ($k = \lambda V$), a result that is standard in agency theory. The principal and agent maximize, respectively:

$$\max_t V_p = e^{-\delta t}(1-\lambda)\left(\frac{\alpha\, t}{\delta} - c\right), \quad \max_t V_a = e^{-\delta t}\lambda\left(\frac{\alpha\, t}{\delta} - c\right)$$

and the optimal upgrade time is the same for each:

$$t^* = \frac{c\delta^2 + \alpha}{\delta\,\alpha}$$

The payment made is

$$k = \lambda\,\frac{\alpha}{\delta^2}$$

the value of which, at the time of the initial contract, is a benefit to the agent of:

$$k = e^{\left(1-\frac{c\delta^2}{\alpha}\right)}\lambda\frac{\alpha}{\delta^2} > \bar{U}$$

There are certain constraints on the contract price k. First, λ cannot be too 'small'. That is, it must compensate the agent for whatever time, effort and materials are required in the execution of the contract($> \bar{U}$). Second, this assumes that principal and agent have the same discount rates. If this is not true, the agent's optimal exercise will differ from the principal, as illustrated in figure 4.3.

If the contract price is a percentage of the benefit, an agent with a discount rate higher than the principal will prefer to upgrade sooner, and agents with a lower discount rate will upgrade later. This introduces a constraint into the optimal contract. One way to compensate for this is to allow the principal to stipulate a certain level of quality in the contract. For this to be credible,

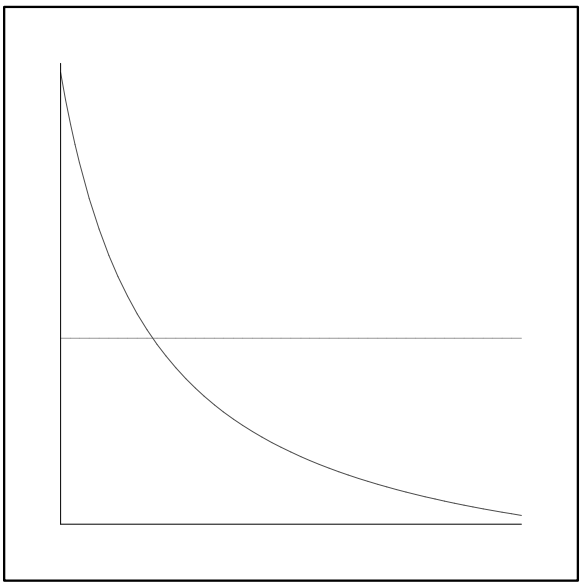

Figure 4.3: Optimal upgrade time versus δ_a holding δ_p fixed

we require that the principal observe the quality level at time t. This is problematic since, if the principal could observe quality, there would be no information asymmetry. Alternatively, we may assume that the principal does observe the value $\alpha\, t$ with some lag Δ_t. In this case, the contract price k_p^*, defined as the value of the contract at the principal's optimal exercise, can be set aside in escrow (perhaps pending suitable verification by an external party) and $k_p^*\, e^{(\delta_a\, \Delta_t)}$ conveyed after Δ_t.

More realistically, we seek an incentive compatible contract that satisfies the constraint for optimal upgrade. We define the relationship between the discount rates as $\delta_a = \delta_p + \epsilon$. Since the agent is likely to be more risk averse than the diversified principal, ϵ is assumed to be some positive constant. When this is the case, the agent must be paid a bonus for waiting. In this case, the optimal exercise for the agent is $t_a^* = \frac{\delta_a c \delta_p + \alpha}{\delta_a \alpha}$ and the time that the agent must be induced to wait is $t_p^* - t_a^* = \frac{1}{\delta_p} - \frac{1}{\delta_p + \epsilon}$. Thus the benefit to the agent at time 0 is

$$k_a^* = e^{\left(\frac{\delta_a}{\delta_p}\left(1 - \frac{c{\delta_p}^2}{\alpha}\right)\right)} \lambda \frac{\alpha}{\delta_p \delta_a} > \bar{U}$$

which represents a premium over the case where discount rates are symmetric. If we allow $\epsilon < 0$, the principal must then pay a bonus to the agent for early exercise. Figure 4.4 illustrates the amount of the incremental bonus (paid at the time of the upgrade) compared to the ratio of discount rates. The bonus is zero when the discount rates are equal. When the agent's discount rate is lower than the principal, the agent must be paid a bonus to exercise early. When the agent's discount rate is higher than the principal's, a bonus must be paid to delay exercise. This amounts to an 'information rent' for the agent.

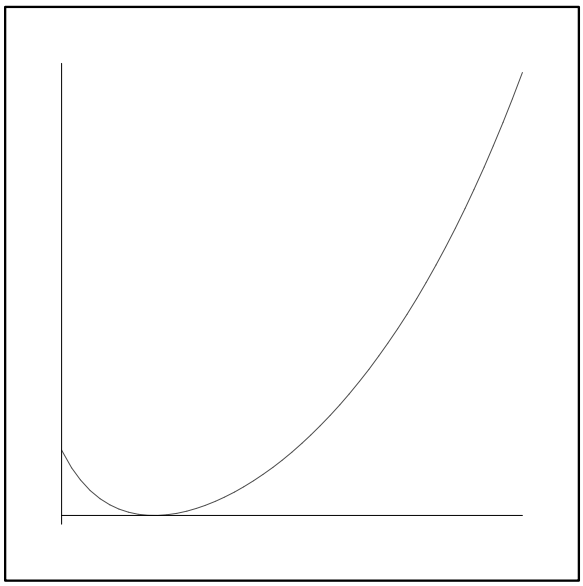

Figure 4.4: Agent's information rent for δ_a holding δ_p fixed

4.3.4 Multiple quality dimensions

The Stern-Stewart problem description considers the determination of an optimal policy where there are multiple quality dimensions. For this simple game under certainty, this introduces two[2] difficulties. First, in order to be verifiable, the contract must stipulate how the agent should make trade-offs in the various technologies. For example, suppose that there is a dramatic, unexpected improvement in display technology, and the principal values this highly. How should the agent structure the upgrade contract to offset a slower than anticipated improvement in CPU performance? This will be accomplished much more easily in the case of a linear than a non-linear relationship. More importantly, to the extent that the optimal upgrade time is incentive compatible, the complexities of stipulating precise relationships in the contract can be avoided.

Second, the more detail that is specified, the easier it is to verify compliance with the contract. However, the transaction cost of writing such a contract is likely to require significant time and expense, and this may exceed the value of avoided enforcement. Increased verifiability of contract breach increases the cost of writing the contract while presumably decreasing the cost of contract enforcement.

Multiple dimensions in measuring the quality of the upgrade make the contract more complex in the case where value to the firm cannot be modeled parametrically. If value can be measured, then a linear contract as introduced above will give the right incentives. Nevertheless, the model constraints encountered earlier resurface here. The fact that quality increases along several dimensions does not, by itself, change the incentive compatibility requirements, although it makes writing the contract more complex.

[2]This would become considerably more complex when we allow technological improvements to evolve stochastically, in which case implementation would involve multivariate stochastic processes.

4.3.5 Multiple upgrades

Placing the upgrades in a repeated game framework, instead of the one-shot game presented above, also will not solve the fundamental problem. In a repeated game framework, verifiability of the contract with delay Δ_t will lead to the same kinds of outcomes as the one-shot game. Writing a single contract to cover several upgrades will have the same characteristics, solution and inefficiencies as the one-shot game.

Writing the contracts in sequence introduces new problems. In the one-shot game, the agent's revelation of the optimal upgrade time in the first period identifies the value of α to the principal. However, knowing that the principal will infer α, the agent has an incentive to modify upgrade times as long as it is privately beneficial. If there are other aspects to the outsourcing agreement, however, the agent may provide the principal with optimal times in order to retain future business prospects. To the extent that the agent's 'reputation' for truthful upgrades affects these business prospects, it may be in the agent's interest to reveal α early on.

This latter point relates to the Stern-Stewart problem more generally, in that much of the concern over the agent's private incentives are less important in the context of an on-going relationship with the principal. So, for example, if the agent has a number of contracts with the principal, not just over upgrades, there are incentives for truthful revelation if the principal will eventually observe product quality. The greater the lag Δ_t the less important is this incentive, but it remains a potential rationale for outsourcing a number of activities to the agent.

4.3.6 Summary points on upgrades under certainty

This section has presented a simple upgrade problem in the context of constant, linear exogenous quality improvements. The optimality conditions in this setting are straight-forward, and allow a simple contract to cover the transaction. In contrast, asymmetric information significantly complicates the model. Differences in risk preferences between the risk-averse agent and the risk-neutral principal impose constraints on the model. While we have model-led a constant λ and adjusted the value of payment, it would be helpful to consider optimal values of λ for a range of discount relationships $\frac{\delta_a}{\delta_p}$. While more work on the simple model could certainly extend it, we will devote our efforts instead to the development of a similar problem where the upgrade occurs under uncertainty.

4.4 Upgrades Under Uncertainty

The simple model of linear technological innovation in the prior section allows for straight-forward solution of the problems in information asymmetry. In this section, we turn to the more complex, and more to the point, more in line with the Stern-Stewart problem, model of technological evolution in computing as a diffusion[3] process. Assume that the value of the

[3] A more sophisticated model may be developed that includes jump processes for technological discontinuities and dramatic changes in the various components of a computer system. These jumps may be either positive or negative shocks. An example of a positive shock is the development of new algorithms for data compression that improves the quality of the transfer. A negative shock would be an announcement that the developer of a certain computer system would no longer produce or maintain hardware.

underlying technology evolves according to the following equation of motion:

$$dP = \alpha P + \sigma P dW$$

Where P is the value, α the drift and σ the volatility of the technology process.

If there is a perfectly competitive market to supply the technology, then the value at time t, denoted P_t, and the cost of acquiring the technology at that time, denoted K_t, will be equal. As long as this is the case, the principal can upgrade according to whatever schedule it prefers, exchanging the cost of adoption for the value of the technology. In perfectly competitive markets, a myopic upgrade policy will do.

A perfectly competitive market structure appears to trivialize the problem, but this result assumes that the effects of information asymmetry are competed away in a market with free entry and exit. The Stern-Stewart problem description identifies concern over the loss of the option to delay. In order for there to be value to delay, it must be that the current technology is not appropriately priced to the firm. For example, the assumption that adoption at a fixed date is suboptimal implies that the asset is overpriced in comparison to the firm's own valuation of expected availability of future quality. If this valuation were the same for all adopters, the technology would not trade until it had declined to this lower price.

If trade is based on the firm's private adoption cost, then a range of prices may exist in the market. We may model the strike price as the contract price plus the firm's private adoption cost: $K_t = k_t + c_t$, where k is the price on delivery of the contract and c is the adoption cost. The firm's objective is to gain maximum value from the contract:

$$\max V_t = P_t - K_t$$

4.4.1 Information and transaction costs

The Stern-Stewart description of an evergreen upgrade, however, suggests a deeper problem, subject to transaction uncertainty. The principal pays a sum at time t_0 for the upgrade of computer systems at a future date, denoted T. At T, the agent receives k_T to provide a computer that is "upgraded to the latest standards" for the principal. In terms of financial contract, this is similar to a forward contract, usually written for standardized commodities. In contrast to standard forward contracts, however, the evergreen contract does not specify the quality of the product at the date of delivery, and so is particularly subject to *ex post* opportunism. For example, establishing "latest standards" in computing is problematic, and is especially of concern to the principal if the agent is facing liquidity constraints as the upgrade date approaches. [4]

[4] Why anyone would enter such a contract is puzzling to us, yet this is precisely the kind of agreement made at some university computing centers. We have heard it argued that this contract allows the upgrade provider the ability to secure the necessary resources at the time of the upgrade. The implication is that giving the agent commitments to timing allows coordination benefits that compensate for the lack of flexibility in timing. At issue is an estimate of the principal's loss due to the agent's inability (due to lack of resources) to upgrade versus the loss of option value due to the contract commitment. Without some estimate of the probability and magnitude of these losses, we are not in a position to comment on the merits of the "evergreen" contract. However, suitable extensions to the model should provide some reasonable estimate of the trade-offs.

If the model were played as a one-shot game, the principal would risk contracting with an unscrupulous agent where moral hazard would result in the agent purchasing the minimum acceptable "latest standard" computer, and pocketing the difference. In a repeated game framework, however, the agent may supply the requisite quality if future profits depend upon contract compliance. These results are standard in game theoretic models of repeated interaction, such as the cooperation observed in repeated "prisoner's dilemma" games (see [2]).

In this framework, however, there is concern over the loss of flexibility in the timing of the upgrade decision. There are two potential concerns. One is that the evergreen program stipulates the timing of the upgrade well in advance of realizing the value of it. Clearly, flexibility in the timing of the upgrade is important if technological evolution is uncertain. In the case of upgrades under certainty, we saw that the contract could be written for any number of upgrades with a specified lag between upgrades. When there is uncertainty because technological innovation is a diffusion process, the principal may prefer to delay an upgrade when the contracted time arrives. That is, the European option would add value to a pre-specified contract time. But second, an American option would add even greater value than the European option, since it provides greater flexibility about the timing of exercise. It is a standard result in option pricing that the greater exercise flexibility of the American option implies that the European option is a lower bound for its value. Adding these flexibilities to the contract could only increase the value of the upgrade decision.

4.4.2 Asymmetric information and value uncertainty

We next present a model of the delegation of the upgrade decision under asymmetric information. The agency concerns under value uncertainty in an outsourced model are not fundamentally different from the decision presented above. Exercise at the optimal will have similar value and verifiability constraints, and the fact that P evolves stochastically, while complicating the picture, does not appear to fundamentally alter the nature of the contract. That is, if we allow the agent to share in the profitability of the contract, the agent will have incentives to execute the upgrade contract at the time that is optimal to the principal. For this reason, we next consider how the upgrade must be structured with an in-house IT department. This analysis will provide at least some insight into what it may make sense to outsource the upgrade decision, since, as we shall see, there are important difficulties in making the upgrade efficiently. In this context, then, principal refers to the headquarters of the organization, and agent refers to the IT department.

We make the following assumptions about information held by the principal and the agent. 1) At the time that an opportunity becomes available, both the principal and the agent are aware of the initial value of the upgrade, and the drift and volatility. These are denoted P_0, μ and σ. 2) After hiring the agent, it is assumed to be too expensive for the principal to continue monitoring the value of the upgrade; P_t is in only the agent's information set. The upgrade is undertaken when the agent recommends, and this is denoted time T. The value of the upgrade at exercise, P_T is known by both the principal and agent. The model assumes that both the principal and the agent are aware of the optimal conditions to exercise the option, but only the agent has private information on whether those conditions are satisfied. The agent's task is to monitor the value process, and announce to the principal when to upgrade.

4.4.3 Payoffs

The upgrade date, denoted T, is not known at the start of the contract. As before, subscripts a refer to the agent, p to the principal. The model assumes that the principal discounts at the risk-free discount rate during the wait-time, and the agent's discount rate is greater than that: $\delta_a >> \delta_p$.

Prior to exercise, the value of the upgrade P_t evolves stochastically according to geometric Brownian motion. The principal's cost at the time of the upgrade is $K = k_t + c_t + b$. The terms are as before, though we have added a bonus b, and we assume for simplicity that k, c are fixed over the relevant time horizon. The value of the option to develop the opportunity is denoted $V = \max(P - K, 0)$. There is an additional constraint to pay the IT employee a salary s. Under these conditions, the expected value of the upgrade at the start of the contract is defined as:

$$\pi_p = e^{(-\delta_p T)} (V - K) - \frac{s\left(1 - e^{(-\delta_p T)}\right)}{\delta_p}$$

4.4.4 Simple model of agent's pay-offs

It is useful to understand the timing decision if it is made inside the firm. Constraints on compensation require that the agent be paid a salary s plus a bonus b at exercise. We define the agent's objective function as the utility from the compensation package. The agent's risk-aversion is characterized by a higher discount rate over the project development horizon. Under these conditions, the risk-aversion is embedded in the agent's demand for a higher premium on future cash flows, making the agent more impatient than the risk-neutral principal. We seek the right compensation package to offer the agent, since the principal seeks someone who will make the exercise decision consistent with the principal's preferences.

$$\pi_a = e^{(-\delta_a T)} b + \frac{s\left(1 - e^{(-\delta_a T)}\right)}{\delta_a}$$

At this point, the compensation does not assure that the agent will act in the principal's interest, or even that the agent will want to enter into the contract. If the principal offers a choice between a fixed bonus or a continuous salary, the agent will choose either to exercise immediately and get the bonus ($T = 0$), or never to exercise and collect the salary ($T = \infty$). This will be a function of the agent's discount rate δ_a, the bonus b and the salary s. The agent will take the salary in perpetuity unless $\frac{s}{\delta_a} < b$, in which case, the agent will take the bonus immediately.

In either case, the return to the principal will not depend on optimal exercise, and thus the agent offers no value. The principal has several options. First, she may offer a bonus that is a linear function of her optimal, denoted V_p^*. If salary is set $s = 0$, the agent will have an incentive to maximize the value of the option. The optimal exercise for the agent is denoted V_a^*. Under these constraints, the model assumes certain regularity conditions that yield the following results. First, the optimal value of the option is a function of the volatility of the process (σ), the drift in the process (μ), the discount rate (ρ), and a convenience term that represents the difference between the discount rate and the drift (δ). There are also restrictions

on the parameters for solutions finite in T.

$$\begin{aligned} dV &= \mu V\, dt + \sigma\, V\, dz \\ \delta &= \rho - \mu \\ \mu &< \rho \end{aligned}$$

The following value-matching and smooth-pasting conditions are optimal:

$$A(P^*)^{\beta} = \frac{P^*}{\delta} - K$$

$$\beta\, \mathrm{A}(P^*)^{(\beta-1)} = \frac{1}{\delta}$$

In order to satisfy these, β must solve the following quadratic:

$$\frac{1}{2}\sigma^2\, \beta\, (\beta - 1) + (\rho - \delta)\, \beta - \rho = 0$$

and if $\beta > 1$:

$$\beta = \frac{1}{2} - \frac{\rho - \delta}{\sigma^2} + \sqrt{(\frac{\rho - \delta}{\sigma^2} - \frac{1}{2})^2 + 2\,\frac{\rho}{\sigma^2}}$$

So, for given K, ρ, α, σ we can find an optimal exercise condition.

$$P^* = \frac{\beta\, \delta\, K}{\beta - 1}$$

Optimal exercise will differ according to risk-preferences, modeled as differences in discount rates ρ. Define δ strictly positive as follows:

$$\delta_p = \delta_p - \mu - s > \delta_a = \delta_a - \mu > 0$$

$$\beta_p = \frac{1}{2} - \frac{\mu}{\sigma^2} + \sqrt{(\frac{\mu}{\sigma^2} - \frac{1}{2})^2 + 2\,\frac{\delta_p}{\sigma^2}}$$

$$\beta_a = \frac{1}{2} - \frac{\mu}{\sigma^2} + \sqrt{(\frac{\mu}{\sigma^2} - \frac{1}{2})^2 + 2\,\frac{\delta_a}{\sigma^2}}$$

Clearly, $\beta_p < \beta_a$ since $\delta_p < \delta_a$, and, applying this to V_a and V_p, the following are the optimal strike prices:

$$V_a = \frac{\left(1/2 - \frac{\mu}{\sigma^2} + \sqrt{\left(\frac{\mu}{\sigma^2} - 1/2\right)^2 + 2\,\frac{\delta_a}{\sigma^2}}\right) K}{-1/2 - \frac{\mu}{\sigma^2} + \sqrt{\left(\frac{\mu}{\sigma^2} - 1/2\right)^2 + 2\,\frac{\delta_a}{\sigma^2}}}$$

$$V_p = \frac{\left(1/2 - \frac{\mu}{\sigma^2} + \sqrt{\left(\frac{\mu}{\sigma^2} - 1/2\right)^2 + 2\,\frac{\delta_p}{\sigma^2}}\right) K}{-1/2 - \frac{\mu}{\sigma^2} + \sqrt{\left(\frac{\mu}{\sigma^2} - 1/2\right)^2 + 2\,\frac{\delta_p}{\sigma^2}}}$$

These results demonstrate that $V_p > V_a$ since $\delta_p > \delta_a$, implying that the agent will exercise prematurely in comparison to the principal's optimal. Figure 4.5 depicts the divergence between the principal and the agent's optimal exercise as a function of the volatility of the project.

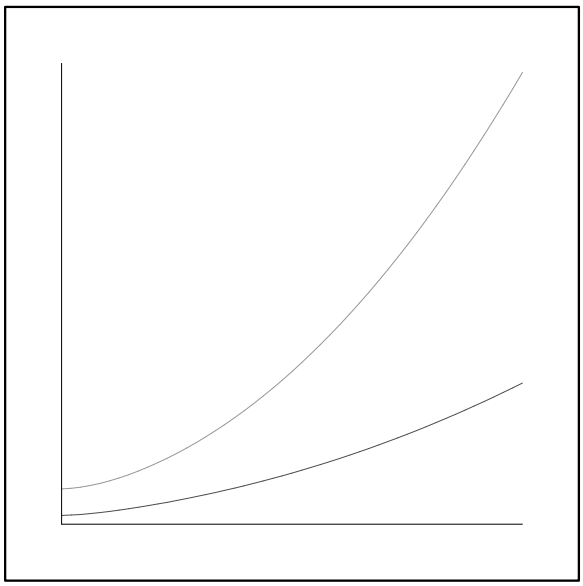

Figure 4.5: Principal and Agent value versus σ

4.4.5 Results and implications

If the principal is restricted to the choice of salary plus bonus, the optimal compensation contract must account for both the incentive to reveal information and the reservation wage of the agent. As the model implies, the bonus must be constructed as a function of the difference between the discount rates, the volatility of the underlying price, and the drift of the project. If the bonus is constructed as $\lambda (V_p - V_a)$ with $\lambda > 0$ the agent may exercise at the appropriate time, but a similar constraint as under uncertainty must compensate the agent for 1) the expected wait and 2) the difference in discounting. The solution is not fundamentally different than the bonus under certainty. The incentive must give the agent sufficient reason to wait until the price process reaches the smooth-pasting condition, and for this we require a combined salary and bonus. If s compensates the agent for precisely the reservation wage of waiting, then the bonus[5] will be adequate to induce the agent to wait.

If the principal can impose a sufficient penalty on the agent for early exercise, and if the value of the project becomes common knowledge at exercise, then the agent can be induced to exercise at the principal's optimal point. Without the ability to charge a penalty, the bonus must satisfy the constraint that the discounted value of the bonus plus the periodic compensation is sufficient to induce the wait. By relating the bonus to the principal's optimal strike, we give the agent an incentive to maximize the bonus.

[5]In a linear compensation scheme with only a bonus payment that is a constant proportion of the value at exercise, this would be resolved.

4.4.6 Complications

These results assume that the principal would never choose to abandon the upgrade. While the march of technological progress has lead to decade after decade of upgrades, this need not continue to be the case. At some point, there may be no rationale for the principal to hire the agent to monitor technology. In such a case, the principal would like the agent to abandon monitoring, but, because the agent is compensated for waiting, the agent may have no incentive to inform the principal to abandon. Under the kinds of diffusion processes that we have examined, the principal may form expectations about when to expect an announcement of optimal upgrade time, however given sufficient volatility in the process, there may be significant uncertainty and room for inefficient delay. Alternatively, in a jump-diffusion process, sudden increases or decreases in the value of the upgrade may make the need sudden abandonment hard to detect.

In such a situation, the principal must pay the agent an abandonment bonus. We leave the construction of such a bonus for later work, but in recognizing the need for such a bonus, we determine at least one of the benefits of a linear compensation scheme. If the agent's compensation is based on a share of the profits at the time of the upgrade, both parties will abandon at the optimal time. This leads to inevitable difficulties in incentives if the costs are shared by the parties, however. In a debt-like arrangement, where the principal pays for project continuance, the agent has an incentive to avoid abandonment on the chance that things may turn around. In an equity-like arrangement, where both parties share the costs of continuation, the agent's incentives will be aligned with the principals.

4.5 Summary and Conclusions

The results from this preliminary model are suggestive of why firms may outsource the IT function. The models under certainty demonstrate efficiencies in contracts that allow sharing. There are a number of practical reasons why these contracts may be hard to write or enforce. We have emphasized the difficulties that the principal has in monitoring the agent, and only referred to the problems that the agent may have in observing the principal's private valuation of the technology. Verifiability is an important problem for both players in this game.

The Stern-Stewart problem description allows for a much more sophisticated treatment of the problems in flexibility and contracts. While the treatment of the literature in this report is brief, there is a great deal more in economic models that allow for renegotiation of contracts. For example, some models allow the principal to sell the business to the agent (see [4]). This will clearly make more sense in other contexts than upgrades, but it does suggest an approach to the development of models where contingent claims allow for flexibility in contract enforcement. While there is much more to do, this report provides a framework for beginning, and it is hoped that much of the work referenced here can be applied to extensions of the models.

Bibliography

[1] Antle, R. and Eppen, G., *Capital rationing and organizational slack in capital budgeting*, Management Science, **31** (1985), 163-174.

[2] Axelrod, Robert, *The Evolution of Cooperation*, Basic Books, New York, NY, 1984.

[3] Bryan, Stephen Hwang, LeeSeok and Lilien, Steven , *CEO stock-based compensation: An empirical analysis of incentive-intensity, relative mix and economic determinants*, Journal of Business, **73**, no. 4 (2000), 661-693.

[4] Demski, Joel S. and Sappington, David E.M., *Resolving double moral hazard problems with buyout agreements*, RAND Journal of Economics, **22**, no. 2 (1991), 232-240.

[5] Dixit, Avinash K. and Pindyck, Robert S., *Investment under Uncertainty*, Princeton University Press, Princeton, NJ, 1994.

[6] Dybvig, P. H. and Zender, J. F., *Capital structure and dividend irrelevance with asymmetric information*, Review of Financial Studies, **4** (1991), 201-219.

[7] Fudenberg, Drew, Holmstrom, Bengt and Milgrom, Paul, *Short-term contracts and long-term agency relationships*, Journal of economic theory, **51** (1990), 1-31.

[8] Gibbons, Robert and Murphy, Kevin J., *Does executive compensation affect investment*, The continental Bank Journal of Applied Corporate Finance, **5**, no. 2 (1992), 99-109.

[9]

[10] Harris, Milton and Raviv, Artur, *The capital budgeting process: incentives and information*, Journal of Finance, **51**, no. 4 (1996), 1139-1174.

[11] Harris, Milton and Raviv, Artur, *Capital budgeting and delegation*, Journal of Financal Economics, **50**, no. 3 (1998), 259-289.

[12] Meyer, Margaret, Milgrom, Paul and Roberts, John, *Organizational prospects, influence costs and ownership changes*, Journal of economics and management strategy, **1** (1992), 9-35.

[13] Myers, S. C. and Majluf, N. S., *Corporate financing and investment decisions when firms have information that investors do not have*, Journal of Financial Economics, **13** (1984), 187-221.

[14] Persons, J. C., *Renegotiation and the impossibility of optimal investment*, Review of Financial Studies, **7** (1994), 419-449.

[15] Rogerson, William P., *Intertemporal cost allocation and mangerial investment incentives: A theory explaining the use of Economic Value Added as a performance measure*, Journal of Political Economy, **105**, no. 4 (1997), 770-795.

[16] Rosen, Sherwin, *Authority, control and the distribution of earnings*, Bell Journal of Economics, **13** (1982), 311-323.

[17] Scharfstein, David S. and Stein, Jeremy C., *The dark side of internal capital markets: divisional rent-seeking and inefficient investment*, NBER Working Paper Series, **W5969** (1997), 1-39.

[18] Sick, Gordon, *Real options*, Finance (R.A. Jarrow, V. Maksimovic and W. T. Ziemba, eds.), Chap. 21, Elsevier, Amsterdam, 1995, 631-690.

[19] Trigeorgis, Lenos, *Real Options : Managerial Flexibility and Strategy in Resource Allocation*, MIT Press, Cambridge, MA 1996.

Chapter 5

Electromagnetic Wellbore Heating

Ibrahim Agyemang[1], Matthew Bolton[2], Lloyd Bridge[2], Samantha Carruthers[3], Tom Janiewicz[4], Margaret Liang[2], Leevan Ling[4], Bruce McGee[5], Andrea McPhee[6], Dominique Noel[1], David Ross[7], Maurice Shevalier[8], Sirod Sirisup[9], Daniel Spirn[10], Ranga Sreenivasan[8], Jef Williams[4].

Report prepared by C. Sean Bohun[3].

5.1 Introduction

This paper is concerned with the recovery of petroleum fluids from an oil reservoir using electrical energy. By its very nature this problem must deal with both the equations that describe the fluid flow as well as the heat flow equations. In general, the oil in the wellbore is very viscous with the consequence that the fluid moves slowly. As a result, the amount of oil collected in a given time is quite small. To increase the production rate of the well, the oil's velocity needs to be increased. One method of accomplishing this is by heating the fluid using an electromagnetic induction tool (EMIT). The simple principle behind the EMIT is that it heats the fluid thereby decreasing its viscosity and increasing its velocity. This method of increasing the production rate of a given wellbore is currently being utilized with the generalization that for wells of several hundred metres in length, several EMIT regions are placed in the wellbore at intervals of about one hundred metres. So that they are all supplied sufficient power, these EMIT regions are connected by a cable surrounded by a steel housing.

We are interested in developing a mathematical model of this problem with the ultimate

[1]University of Alberta
[2]University of British Columbia
[3]University of Victoria
[4]Simon Fraser University
[5]McMillan-McGee Corporation
[6]University of Toronto
[7]Eastman Kodak
[8]University of Calgary
[9]McGill University
[10]New York University

goal of investigating analytically the relationship between the temperature of the EMIT and the production rate of the oil.

One approach to this problem is to write out the full system of coupled partial differential equations that relate the temperature and the velocity flux and then to solve them numerically with an expensive computational fluid dynamics (CFD) program. Indeed, this method has been used in the past [4] and it will be used to test the accuracy of our simplified model in the absence of experimental data. The purpose of *this* paper is to carefully analyse each of the physical processes in this system and by making some basic assumptions, to derive a simple set of equations that still captures the main features of the system modelled with the CFD code.

This paper is organised in the following way. Section 5.2 describes the overall geometry of the problem and establishes the coordinates used to describe the model. At this point the problem is broken into three subproblems: i) the flow of fluid in the reservoir, ii) the flow of fluid in the wellbore and iii) the generation of temperature from the heat sources in any EMIT regions. Parts i) and ii) result in a second order ODE for the oil flux for a fixed viscosity. From part iii) it is found that the temperature of the fluid is inversely proportional to the velocity. Fluid that moves slowly past an EMIT region will absorb more heat than the same amount of fluid that moves quickly past an EMIT. As a result, slowing the fluid velocity increases the temperature and therefore decreases the viscosity. This viscosity is used in parts i) and ii) thereby closing the system of equations.

Part i) is described in section 5.3, where a relationship between the axial changes in the fluid flux and the pressure in the wellbore is derived. The details of part ii) can be found in section 5.4 where a relationship for the velocity and the pressure from the Navier-Stokes equations is obtained by averaging over the radius of the wellbore. Under the assumptions made, the pressure is found to be related to the radius of the wellbore by a form of Poiseuille's law.

Finally, section 5.5 details the derivation of part iii), the temperature equations. This derivation is complicated by the fact that there are four radial regions of the radial problem to consider; EMIT, casing, reservoir and wellbore with the first three forming the boundary conditions for the heat equation in the wellbore region. Furthermore, there are three axial regions: EMIT region, cable region, and a region where there is neither EMIT nor cable. Section 5.6 pulls the results of sections 5.3, 5.4 and 5.5 together and section 5.7 illustrates the analytical solution of the resulting model in the simple situation when no heat is applied to the oil.

In section 5.8, we discuss numerical results of the simplified model with respect to the results predicted by the CFD code. On comparison, we find considerable qualitative agreement between the two models which is quite remarkable considering their relative complexities. These aspects are further discussed in the final section of the paper.

5.2 Geometry

Figure 5.1 depicts the overall geometry of the problem. A horizontal cylindrical well extends from $z = 0$ to $z = L$. Fluid flows radially into the well from the surrounding media and is drawn out with a pump which is located at $z = L$ where a fixed pressure of P_P is maintained.

At $z = 0$, where the end cap is situated, the motion of the flow is radially inward (no horizontal flow at this point). As one increases in z, the action of the pump comes into effect

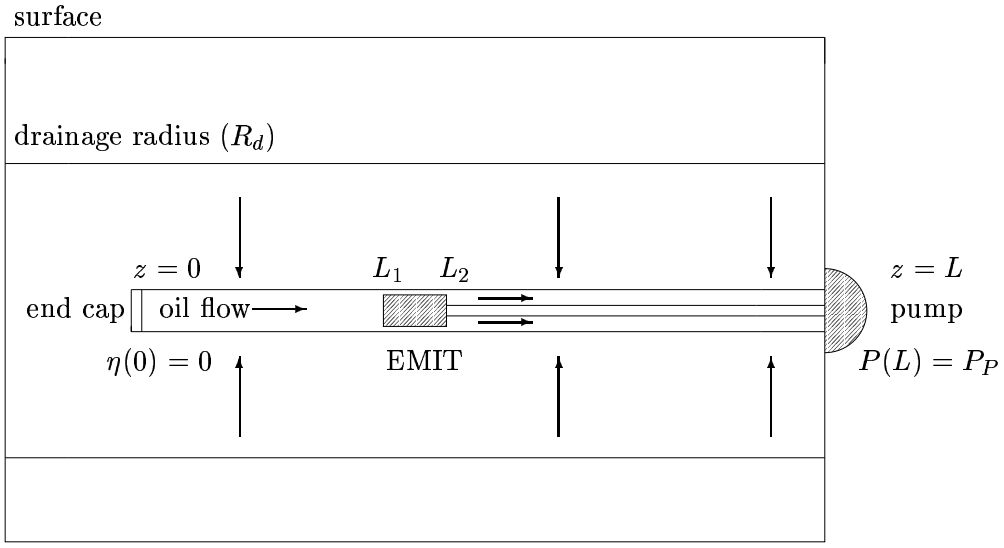

Figure 5.1: Overall geometry for the horizontal wellbore problem.

and imparts a horizontal component to the fluid flow.

This figure shows only one EMIT region which extends from $z = L_1$ to $z = L_2$. It is in these EMIT regions that the oil is heated. Power is supplied to the EMITs through a cable housing resulting in three different regions. Starting at the pump we have a cable housing region that extends to the first EMIT. If there are other EMIT regions then they must also be joined with cable housing and eventually, after the last EMIT region, the wellbore is open with no impediment to the horizontal flow.

5.3 Axial Velocity: Darcy's Law

Once the horizontal well is drilled, fluid seeps from the surrounding region into the wellbore. Once inside the wellbore, the fluid is drawn out with a pump that maintains a fixed pressure at one end of the well. The rate at which the fluid seeps into the wellbore is a function of the pressure differential and the viscosity of the fluid. Indeed, the flow rate (volume/time) of the fluid into a segment of the wellbore of length Δz is given by the expression [2]

$$q(z) = \frac{2\pi k[P_R - P(z)]}{\mu_o \ln(R_d/R_c)}\Delta z \tag{5.1}$$

where k is the permeability of the reservoir, P_R is the reservoir pressure, $P(z)$ is the pressure inside the wellbore at the axial position z, μ_o is the viscosity at the ambient temperature T_a, R_d is the drainage radius and R_c is the outer radius of the casing.

Since we are assuming that we are at a steady state, we make the assumption that the radially flowing fluid remains unheated until it reaches the inner radius of the casing at which point it instantly becomes heated to the temperature of the fluid at that particular z position. Consequently the viscosity in expression (5.1) will remain as μ_o even once the temperature of the wellbore is increased.

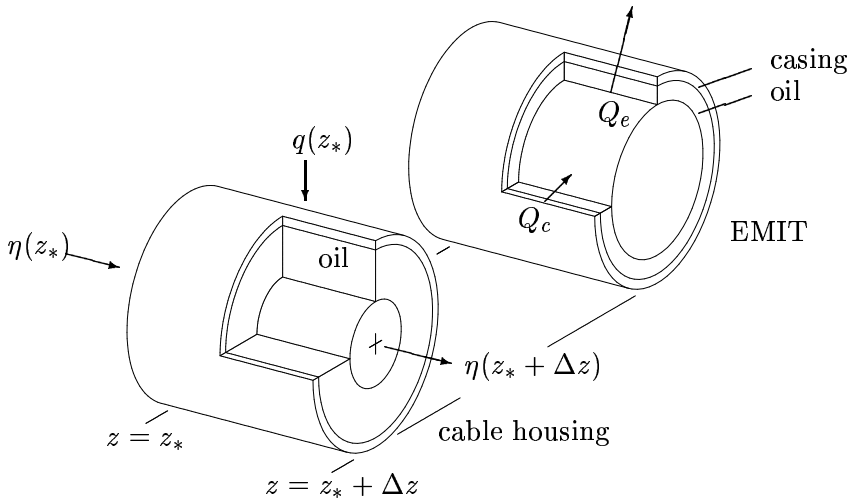

Figure 5.2: An infinitesimal section of the wellbore for the EMIT or cable housing regions.

Using equation (5.1) one can find an expression for the average axial velocity of the fluid, $\bar{v}(z)$. Let R_z denote the inner radius of the wellbore which could be zero, the radius of the EMIT tool, R_e, or the outer radius of the electrical housing, R_h. Using this definition of R_z, let $\eta(z) = \pi(R_w^2 - R_z^2)\bar{v}(z)$ which is the flux in the wellbore. The advantage of using $\eta(z)$ rather than $\bar{v}(z)$ is that $\eta(z)$ is a continuous function whereas the velocity $\bar{v}(z)$ is not.

Figure 5.2 shows an infinitesimal annular section of the wellbore of length Δz. At $z = z_*$ the axial flux is $\eta(z_*)$ while the radial flux is given by expression (5.1). By the conservation of mass, these two components combine to give the axial flux at $z = z_* + \Delta z$. In other words,

$$\eta(z_*) + \frac{2\pi k[P_R - P(z_*)]}{\mu_o \ln(R_d/R_c)}\Delta z = \eta(z_* + \Delta z).$$

Rearranging terms and letting $\Delta z \to 0$ gives the expression

$$\frac{d\eta}{dz} = \frac{2\pi k[P_R - P(z)]}{\mu_o \ln(R_d/R_c)}; \qquad \eta(0) = 0. \tag{5.2}$$

The boundary condition $\eta(0) = 0$ just expresses our approximation that the axial fluid velocity is zero at the end of the wellbore. Since $P(z) < P_R$ throughout the wellbore, $d\eta/dz > 0$ which is consistent with having the fluid flux increase as it approaches the pump located at $z = L$.

5.4 Axial Pressure: Navier-Stokes

A relationship between wellbore pressure and flow velocity is obtained from the Navier-Stokes equations for an incompressible[11] viscous fluid,

$$\rho\frac{\partial \vec{v}}{\partial t} + \rho(\vec{v}\cdot\nabla)\vec{v} = -\nabla P + \mu\Delta\vec{v}. \tag{5.3}$$

[11] A fluid is said to be incompressible if the velocity satisfies $\nabla \cdot \vec{v} = 0$.

Again refer to figure 5.2 where we consider an arbitrary yet constant cross section. Assume a steady fluid flow inside the wellbore that propagates in the $\hat{\mathbf{k}}$ direction. Assuming that the flow is radially symmetric, we have $\vec{v} = v(r)\hat{\mathbf{k}}$ and the continuity equation is automatically satisfied. Resolving (5.3) into the r, θ and z directions gives $\partial P/\partial r = 0 = \partial P/\partial\theta$ and

$$\mu \frac{1}{r}\frac{\partial}{\partial r}\left(r\frac{\partial v}{\partial r}\right) = \frac{\partial P}{\partial z}. \tag{5.4}$$

The first two conditions on the pressure imply that $P = P(z)$. As a consequence, the RHS of (5.4) is a function of z alone while the LHS is a function of r alone. The only way that this can be so is if the pressure is constant over the cross-section of the wellbore. This implies that for our annular domain $R_z < r < R_w$ and $z_* < z < z_* + \Delta z$ we must solve

$$\mu \frac{1}{r}\frac{\partial}{\partial r}\left(r\frac{\partial v}{\partial r}\right) = \frac{\Delta P}{\Delta z}; \qquad v(R_z) = 0 = v(R_w)$$

where $\Delta P = P(z_* + \Delta z) - P(z_*) < 0$ and we have imposed a no slip condition on the boundary. The general solution for the velocity distribution as a function of radius in this case is found to be

$$v(r) = \frac{1}{4\mu}\frac{\Delta P}{\Delta z}\left[r^2 - R_w^2 - \frac{R_w^2 - R_z^2}{\ln(R_w/R_z)}\ln\left(\frac{r}{R_w}\right)\right].$$

For regions in which there is no EMIT tool or tubing ($R_z = 0$), this reduces to the familiar parabolic flow profile

$$v(r) = \frac{1}{4\mu}\frac{\Delta P}{\Delta z}(r^2 - R_w^2).$$

In order to find the average flux of oil at any fixed value of z one needs to compute the average of this radial velocity. Computing this average results in

$$\overline{v} = \frac{2}{(R_w^2 - R_z^2)}\int_{R_z}^{R_w} v(r)\, r dr = \frac{1}{8\mu}\frac{\Delta P}{\Delta z}\left[\frac{R_w^2 - R_z^2}{\ln(R_w/R_z)} - \frac{R_w^4 - R_z^4}{R_w^2 - R_z^2}\right].$$

If we now let $\Delta z \to 0$, rearrange terms and use the definition of η this becomes

$$\frac{dP}{dz} = -\frac{8\mu}{\pi}\left[\frac{\ln(R_w/R_z)}{(R_w^4 - R_z^4)\ln(R_w/R_z) - (R_w^2 - R_z^2)^2}\right]\eta(z); \qquad P(0) = P_P \tag{5.5}$$

where P_P is the pressure maintained by the pump at $z = L$. If one allows $R_z \to 0^+$ then equation (5.5) reduces to

$$\frac{dP}{dz} = -\frac{8\mu}{\pi R_w^4}\eta(z)$$

which is the popular *Hagen-Poiseuille* [1] equation. We have now reached the point where we can consider what happens as the wellbore is heated.

Data	Symbol	Value
Wellbore Properties		
Outer Casing Radius	R_c	69.85 mm
Inner Casing Radius	R_w	63.50 mm
EMIT Radius	R_e	50.80 mm
Housing Radius	R_h	30.1625 mm
Left Edge of EMIT	L_0	495 m
Right Edge of EMIT	L_1	505 m
Wellbore Length	L	1,000 m
Reservoir Properties		
Permeability[12]	k	10,000 mD
Ambient Viscosity[13]	μ_o	15,000 cP
Drainage Radius	R_d	100 m
Reservoir Pressure[14]	P_R	5,000 kPa
Producing Pressure	P_P	500 kPa
Thermal Properties		
Fluid Heat Capacity	ρC_f	2.8×10^6 J m^{-3}K^{-1}
Casing Power	Q_c	795.8 kW m^{-3}
EMIT Power	Q_e	79.58 kW m^{-3}
Ambient Temperature	T_a	30°C

Table 5.1: Input Data for the example calculations.

5.5 Including the Temperature

For the purposes of this discussion, we assume that the wellbore has a steady state temperature distribution that is a function of r and z alone. As well, to simplify the expressions, we will take the far field temperature to be zero. We also assume that thermal conduction in the z direction is negligible. The primary sources of heat are the EMIT regions and the casings around them and it is assumed that the heat production is uniform. Since the reservoir and casings are porous,[15] heat is convected radially in them. In the wellbore, oil flows along the axis and therefore heat is convected axially in the wellbore.

Consider the casing. In this region, the total rate of heat flow is the sum of the flow due to heat conduction and the flow due to the movement of fluid in the radial direction. The result of this is that the heat flux in the casing is given by

$$\vec{\Phi}_c = \left(\lambda_c \frac{\partial u}{\partial r} + \rho C_f v_r u \right) \hat{\mathbf{r}}. \tag{5.6}$$

We have denoted the temperature as $u(r, z)$, the radial speed of the fluid as v_r, the conductivity

[12] 1 darcy = 9.86923×10^{-13} m^2.
[13] 1 centipoise = 1×10^{-3} kg m^{-1}s^{-1}.
[14] 1 pascal = 1 kg m^{-1}s^{-2}.
[15] The casing actually has holes drilled into it for the transport of oil.

of the casing as λ_c and ρC_f for the fluid heat capacity.

For any closed region the total heat energy produced must equal the total heat energy lost. Let Q_c be the amount of heat energy produced in the casing per volume per unit time. By choosing a cylindrical region of radius r and length Δz expression (5.6) implies that

$$2\pi\Delta z\left(\lambda_c r\frac{\partial u}{\partial r}+\rho C_f r v_r u\right)+2\pi\Delta z\int Q_c\, r dr=0 \qquad \text{or} \qquad \frac{1}{r}\frac{\partial}{\partial r}\left(\lambda_c r\frac{\partial u}{\partial r}+\beta u\right)+Q_c=0 \tag{5.7}$$

where $\beta=\rho C_f r v_r$. A similar line of reasoning yields expressions for the reservoir, wellbore and EMIT. These equations are summarized below:

$$\text{Reservoir:} \qquad \frac{1}{r}\frac{\partial}{\partial r}\left(\lambda_r r\frac{\partial u}{\partial r}+\beta u\right)=0 \tag{5.8}$$

$$\text{Wellbore:} \qquad \frac{1}{r}\frac{\partial}{\partial r}\left(\lambda_w r\frac{\partial u}{\partial r}\right)=\frac{\rho C_f}{\pi}\frac{\partial}{\partial z}(\eta u) \tag{5.9}$$

$$\text{EMIT:} \qquad \frac{1}{r}\frac{\partial}{\partial r}\left(\lambda_e r\frac{\partial u}{\partial r}\right)+Q_e=0. \tag{5.10}$$

We are not interested in resolving the details of the radial temperature distribution in the wellbore. Rather, we only care about the axial variations of the mean temperature. This permits a simplification. First define the mean temperature over the wellbore cross-section at z to be

$$T(z)=\frac{2}{(R_w^2-R_z^2)}\int_{R_z}^{R_w}u(r,z)\,r dr$$

as we did with η in the derivation of equation (5.5). Recall that R_z is the inner radius and depends on z. With this definition, the equation for the wellbore (5.9) can be integrated resulting in the expression

$$\rho C_f(R_w^2-R_z^2)\frac{d}{dz}(\eta T)=2\pi\lambda_w r\frac{\partial u}{\partial r}\Big|_{R_z}^{R_w}. \tag{5.11}$$

The thermal flux in the wellbore is given by

$$\vec{\Phi}_w=\lambda_w\frac{\partial u}{\partial r}\hat{\mathbf{r}}+\frac{\rho C_f\eta u}{\pi r^2}\hat{\mathbf{k}}.$$

This must be continuous at the interfaces. Hence in the radial direction for the wellbore/casing and wellbore/EMIT interfaces one has respectively

$$r\lambda_w\frac{\partial u}{\partial r}\Big|_{R_w}=\left(r\lambda_c\frac{\partial u}{\partial r}+\beta u\right)\Big|_{R_w} \qquad \text{and} \qquad r\lambda_w\frac{\partial u}{\partial r}\Big|_{R_z}=r\lambda_e\frac{\partial u}{\partial r}\Big|_{R_z}.$$

It remains for us to evaluate these fluxes.

By solving (5.10) we find that

$$\lambda_e r\frac{\partial u}{\partial r}\Big|_{R_z}=-\frac{1}{2}Q_e R_z^2. \tag{5.12}$$

Furthermore, by solving (5.8) one finds that the temperature has the general form

$$u(r,z) = k_1 + k_2 r^{-\beta/\lambda_r}$$

where k_1 and k_2 are constants. Because the far field temperature is zero, one must have $T(\pm\infty, z) = k_1 = 0$. As a result,

$$\lambda_r r \frac{\partial u}{\partial r} + \beta u = \beta k_1 = 0. \tag{5.13}$$

So the thermal flux in the reservoir and, in particular, the thermal flux through the casing/reservoir interface is zero. Using this as a boundary condition one can integrate the expression for the casing (5.7) to find that

$$\left(\lambda_c r \frac{\partial u}{\partial r} + \beta u\right)\Big|_{R_w} = \frac{1}{2} Q_c (R_c^2 - R_w^2). \tag{5.14}$$

Collecting equations (5.11)-(5.14) gives the final result that

$$\rho C_f \frac{d}{dz}[\eta(z)\, T(z)] = \pi \left[Q_c(z)(R_c^2 - R_w^2) + Q_e(z) R_e^2\right]; \qquad T(0) = 0. \tag{5.15}$$

The heat sources Q_c and Q_e have been written as functions of z. If one is not in an EMIT region, these functions are simply zero. As a result, the RHS of (5.15) piecewise constant; nonzero only where an EMIT is located. Integrating (5.15) gives the result that

$$T(z) = \begin{cases} 0; & 0 \le z < L_0 \\ \dfrac{\Omega}{\eta(z)} \dfrac{z - L_0}{L_1 - L_0}; & L_0 \le z < L_1 \\ \dfrac{\Omega}{\eta(z)}; & L_1 \le z \le L, \end{cases} \qquad \Omega = \frac{\pi[Q_c(R_c^2 - R_w^2) + Q_e R_e^2](L_1 - L_0)}{\rho C_f}. \tag{5.16}$$

From the values in table 5.1 we find that $\Omega \sim 3.10 \times 10^{-2}$ $\mathrm{m^3 s^{-1} K}$. Notice that if $\eta(z)$ were constant then the temperature would increase monotonically as one moved from $z = 0$ to $z = L$. However, since $\eta(z)$ actually increases as one moves toward the pump, the temperature of the fluid must decrease once it passes an EMIT region.

The temperature affects the velocity and the pressure through the viscosity. This viscosity is given empirically in units of thousands of centipoise through

$$\log_{10} \mu(T) = -3.002 + \left(\frac{453.29}{303.5 + T}\right)^{3.5644}. \tag{5.17}$$

Hence, one can see that an increase of 100°C can result in a decrease in viscosity of about three orders of magnitude. One final point is that the far field temperature should correspond to the ambient viscosity μ_0. Since $\mu_0 = 15000$ cP we associate the far field temperature of zero with the ambient temperature of $T_a = 30$°C.

5.6 Velocity, Pressure, Temperature Summary

In the axial direction, the rate of change of $\eta(z) = \pi(R_w^2 - R_z^2)\bar{v}(z)$ is governed by Darcy's Law and in our approximation it is assumed that the fluid is not heated until it reaches the wellbore. For the axial pressure, the Navier-Stokes equations are solved for an annular region by assuming that the fluid is incompressible. When we apply heat to wellbore, it is this fluid in the wellbore that is heated and not the fluid in the surrounding region. Therefore by combining the equations (5.2) and (5.5) we can summarize the problem for $\eta(z)$ as

$$\frac{d^2\eta}{dz^2} - \frac{16k\Gamma(z)}{\ln(R_d/R_c)}\frac{\mu(T(z))}{\mu_o}\eta = 0; \qquad \eta(0) = 0, \quad \frac{d\eta}{dz}\Big|_L = \frac{2\pi k(P_R - P_P)}{\mu_o \ln(R_d/R_c)},$$
$$\Gamma(z) = \frac{\ln(R_w/R_z)}{(R_w^4 - R_z^4)\ln(R_w/R_z) - (R_w^2 - R_z^2)^2}$$

where the temperature as a function of z is given with the expression (5.16) and the viscosity μ is given explicitly by (5.17). Recall that $\mu_o = \mu(0)$. So we see that under the assumptions, this problem reduces to finding the solution of a second order nonlinear *ordinary* differential equation for $\eta(z)$.

Once we have solved for $\eta(z)$, relation (5.16) determines the temperature profile and consequently the viscosity as a function of z. The pressure as a function of z is given by expression (5.2) with the result that

$$P(z) = P_R - \frac{\mu_o}{2\pi k}\ln\left(\frac{R_d}{R_c}\right)\frac{d\eta}{dz}. \tag{5.18}$$

Notice that since η increases monotonically as one approaches the pump, the slope of η is positive definite. Consequently, the pressure decreases monotonically as one approaches the pump. One can also extract the average velocity and the production rate at the pump from η. The velocity is given by $\bar{v}(z) = \eta(z)/(\pi(R_w^2 - R_z^2))$ and the volume flux at the pump is $\eta(L)$.

5.7 An Illustrative Example

The overall characteristics of our model can be extracted by studying the case when there are no EMIT regions. In this case one has $Q_c(z) \equiv 0$ and $Q_e(z) \equiv 0$ and from equations (5.16) and (5.17), $\mu(T(z)) = \mu(0) = \mu_o$. In addition, $R_z = 0$ so that $\Gamma(z) = 1/R_w^4$ and the equation for $\eta(z)$ reduces to

$$\frac{d^2\eta}{dz^2} - \frac{16k}{R_w^4\ln(R_d/R_c)}\eta = 0; \qquad \eta(0) = 0, \quad \frac{d\eta}{dz}\Big|_L = \frac{2\pi k(P_R - P_P)}{\mu_o \ln(R_d/R_c)}.$$

The explicit form of the solution is easily verified to be

$$\eta(z) = \frac{\pi R_w^4}{8\mu_o}\frac{P_R - P_P}{L_{\text{crit}}}\frac{\sinh\gamma z}{\cosh\gamma L}; \qquad L_{\text{crit}}^2 = \frac{R_w^4}{16k}\ln\left(\frac{R_d}{R_c}\right) = \frac{1}{\gamma^2}.$$

From expression (5.18) the corresponding pressure is

$$P(z) = P_R - (P_R - P_P)\frac{\cosh\gamma z}{\cosh\gamma L}.$$

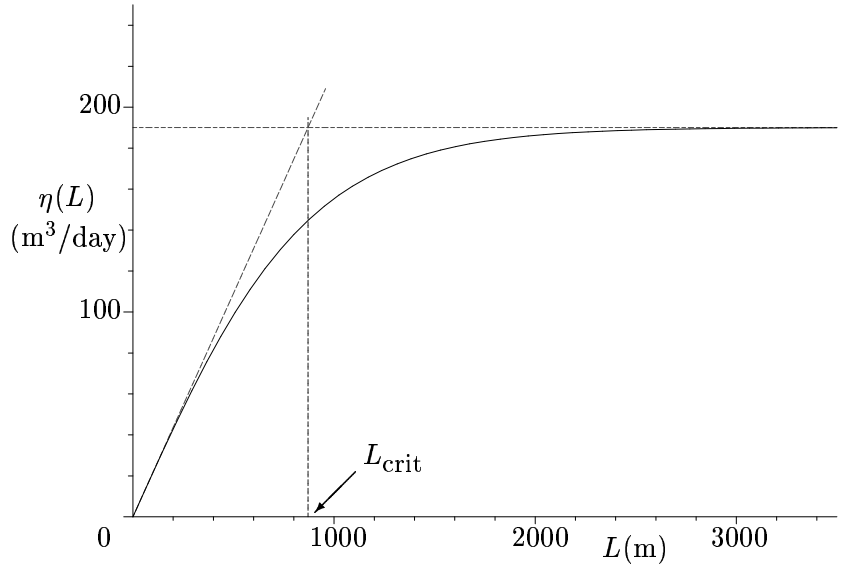

Figure 5.3: The production rate at the pump as a function of the length of the wellbore.

The total production rate at the pump is given by the revealing expression

$$\eta(L) = \frac{\pi R_w^4}{8\mu_o}\frac{P_R - P_P}{L_{\mathrm{crit}}}\tanh\gamma L \sim \begin{cases} \dfrac{\pi R_w^4}{8\mu_o}\dfrac{P_R - P_P}{L_{\mathrm{crit}}}; & L \gg L_{\mathrm{crit}} \\ \dfrac{\pi R_w^4}{8\mu_o}\dfrac{P_R - P_P}{L_{\mathrm{crit}}}\dfrac{L}{L_{\mathrm{crit}}}; & L \ll L_{\mathrm{crit}} \end{cases}$$

which is depicted in figure 5.3 for the data given in table 5.1 where $L_{\mathrm{crit}} = 865$ m and the maximum production rate is 191.32 $\mathrm{m}^3\mathrm{day}^{-1}$. What is immediately apparent is that, without heating, drilling a horizontal well beyond the critical length will not yield any significant increase in production.

5.8 Results

As we stated in the introduction, the agreement between the qualitative results of the simplified model and those predicted by the CFD code are quite remarkable. This is especially true in light of their respective computational costs. Figure 5.4 illustrates the results for the data described in table 5.1.

Only the pressure, temperature and final production rate are easily available from the CFD code. Because of this, only the pressure and temperature curves in figure 5.4 have a dashed line. To solve the nonlinear ODE described in section 5.6 two different method were employed; a shooting method and a method of successive over relaxation (SOR). Whenever they are discernible, the solution from the shooting method is a solid line while the SOR solution is a dashed dot line. We begin our discussion with the pressure curve.

Because of the boundary condition $P(L) = P_P$, all of the pressure curves intersect at $z = L$. At $z = 0$ the CFD code predicts that $P_{\mathrm{cfd}}(0) = 4.07 \times 10^3$ kPa while the solution of the simplified model gives $P_{\mathrm{sm}}(0) = 3.44 \times 10^3$ kPa. Despite the fact that our model tends to

underestimate the results from the CFD code, the amount of pressure drop across the EMIT region is predicted correctly. $P_{\text{sm}}(z)$ can be made to match $P_{\text{cfd}}(z)$ by artificially decreasing the outer casing radius R_c however this would in turn decrease the predicted production rate. These observations indicate that the differences in the pressure predicted with the simplified model and the CFD code are greatest where the fluids tend to form a boundary layer on the outer casing wall. The last curve in this plot is solely for comparison purposes. It is the pressure curve for the example described in section 5.7 where there is no EMIT region.

Comparing the temperature curves, there is a distinct difference in the shape of the two curves. However, the temperature from the CFD code is just the temperature at a fixed radius rather than an average over the radius of the wellbore for a given value of z. The peak temperature in the EMIT region is faithfully reproduced, but the rate at which the fluids lose heat is larger in the CFD model. Consequently, the surface temperature of simplified model is larger, $T_{\text{sm}}(L) = 37.4^o\text{C}$, than that of the CFD model, $T_{\text{cfd}}(L) = 33.1^o\text{C}$. Curves for the viscosity and the velocity could not be compared with the CFD model as these quantities were not directly accessible.

For the production rate refer to figure 5.5. As can be seen, the simplified model underestimates the production rate at the pump. In fact, $\eta_{\text{cfd}}(L) = 187.6$ m^{3}day^{-1} while $\eta_{\text{sm}}(L) = 115.2$ m^{3}day^{-1}. This can also be explained with the formation of a boundary layer. In the full model the boundary layer for the cool fluids to the left of the EMIT region would extend further into the wellbore than in the region to the right of the EMIT where the fluid is heated. This relative difference in the effective outside radius of the casing due to a boundary layer of varying thickness would tend to boost the overall production rate.

In general, the simplified model seems to reproduce the overall characteristics of the full model described by the CFD code. Moreover, it does this with a comparatively low computational cost. One of the primary benefits of this reduction in the computational cost is that it allows one to run a number of numerical experiments cheaply and *on site* in the vicinity of the wellbore itself. As an example, for a given set of wellbore characteristics, we can quickly iterate the simplified model to search for the position of the EMIT which maximizes the production rate at the pump. Because of the inherent nonlinearity, it is not at all clear that this position should simply be the midpoint $z = L/2$. While there will be a unique location that yields a global maximum for the production rate, it is not at all clear that there can not be other local maxima or minima. From figure 5.5 we know that for the geometry in table 5.1 that $\eta_{\text{sm}}(L) = 115.2$ m^{3}day^{-1} when the EMIT is located at $L/2$. The complete curve where the EMIT is allowed to move from one end of the wellbore to the other is shown in figure 5.6. One finds that the greater the distance between the EMIT and the pump, the greater the production. However at about $z = 0.40L$ the rate of increase in production begins to flatten out. Moving the EMIT tool 100 m for an increase of only 1% in production rate at the pump hardly seems worth the effort. Placing the centre of the EMIT at 400 m yields a production rate of at the pump of $\eta_{\text{sm}}(L) = 116.6$ m^{3}day^{-1}. From our observations of the two models, the production rate predicted by the CFD code will of course be greater than this, but more importantly, it should be greater than $\eta_{\text{cfd}}(L) = 187.6$ m^{3}day^{-1} obtained when the EMIT region is located at $L/2$.

5.9 Conclusions and Recommendations

We present a mathematical model for the flow of fluids in a horizontal well in which the well has one or more regions that are electrically heated from an external source. By making some basic assumptions, we develop a nonlinear second order *ordinary* differential equation for the volume flux $eta(z)$ as a function of distance along the wellbore. From the volume flux, both the average radial temperature and the pressure can be extracted. This model succeeds in capturing many of the features observed in the solutions presented with an expensive computational fluid dynamics (CFD) program. Our model tends to underestimate both the pressure and the production rate however, it retains the overall structure of these quantities. The large reduction in computational cost when using the simplified model allows one to quickly run a series of numerical experiments to see the effects of changing various parameters. One such experiment is considered and it is found if the wellbore extends from $z = 0$ to $z = L$ with the pump at $z = L$ then the EMIT should be placed at about $z = 0.40L$ to maximize the production. While the production rate could be increased further by moving the EMIT closer to $z = 0$, this would only increase the production rate a mere 1-2%.

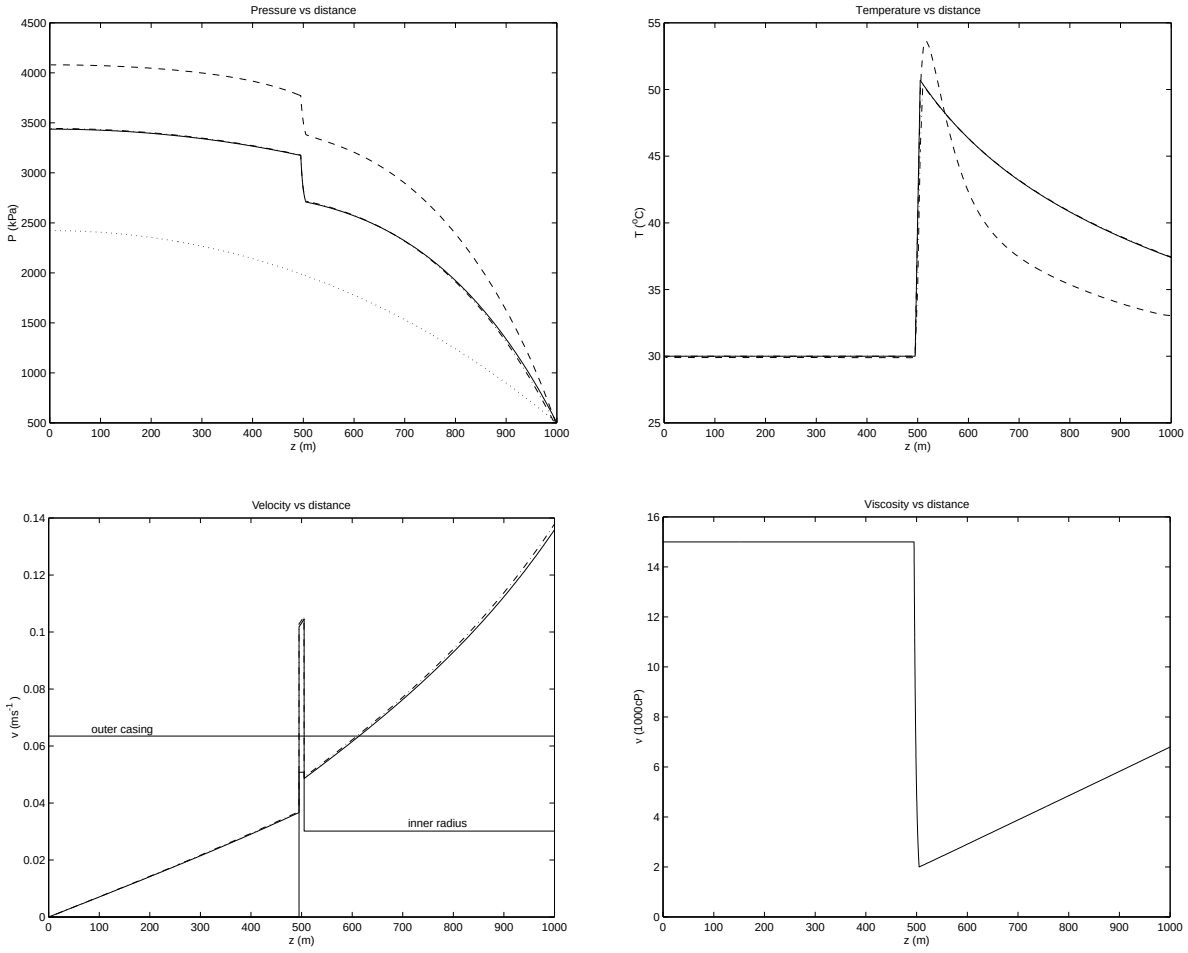

Figure 5.4: The four figures are the pressure, temperature, velocity and viscosity as a function of distance along the wellbore. Only the pressure and temperature for the CFD code was available. These are the dashed lines in the respective plots. More than one method was used to solve the simplified model. Where they are distinguishable, the shooting method solution is a solid line where the SOR method is indicated with a dashed dot. A longitudinal section of the wellbore in indicated in the plot of the velocity.

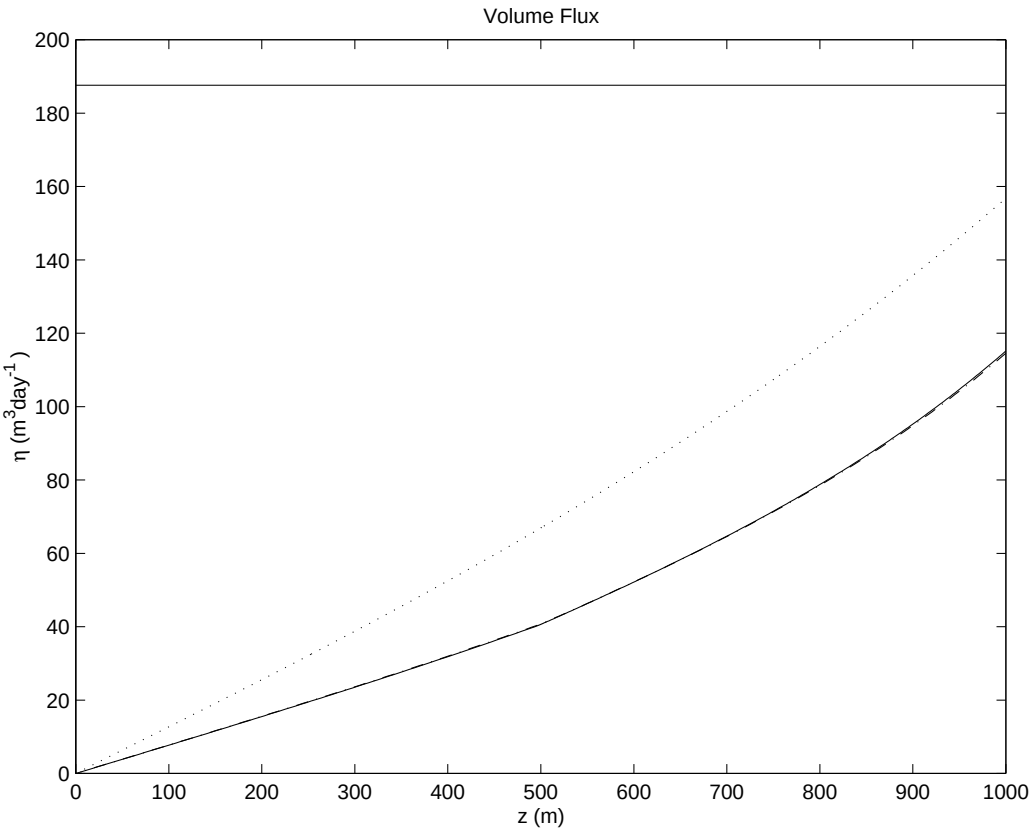

Figure 5.5: The production rate $\eta_{\mathrm{sm}}(z)$ as a function of distance along the wellbore for the simplified model. Only the production at the pump $\eta_{\mathrm{cfd}}(L)$ was available from the CFD code. The dotted curve is the production rate for the case of no EMIT described in section 5.7.

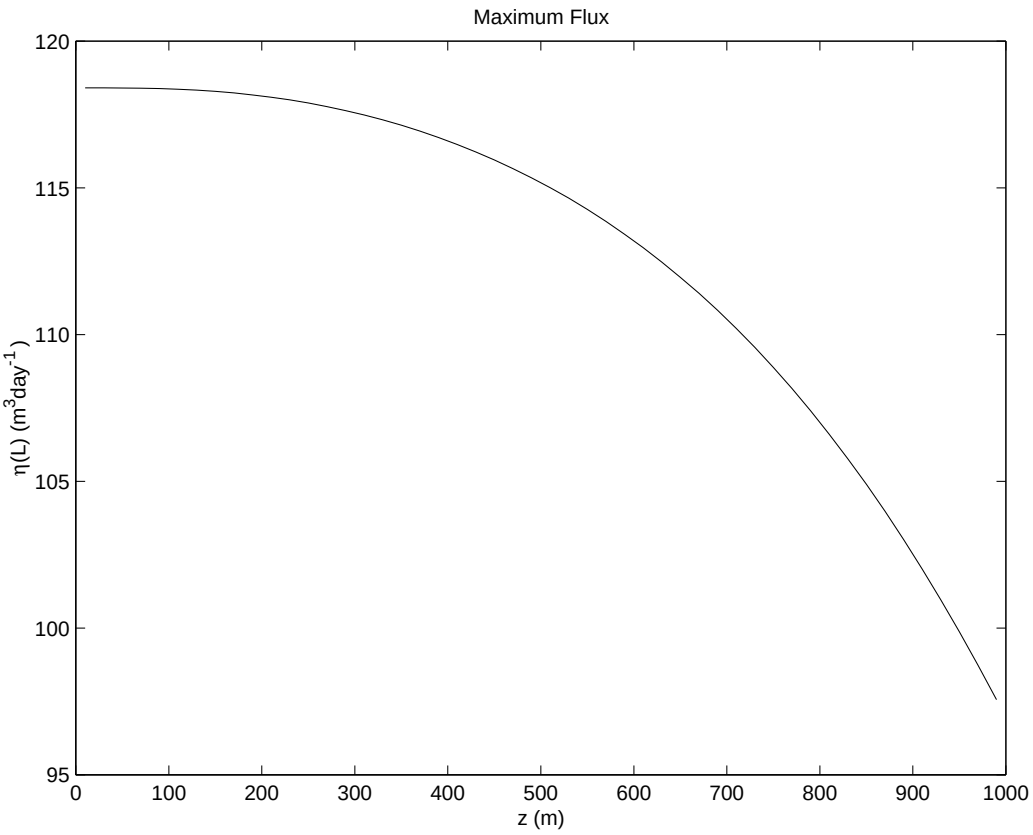

Figure 5.6: The production rate at the pump $\eta_{\mathrm{sm}}(L)$ as a function of the location of the centre of the EMIT.

Bibliography

[1] Landau, L.D. & Lifshitz, E.M., (1959). *Fluid Mechanics.* Addison-Wesley Series in Advanced Physics. Pergamon Press Ltd.: Reading Massachusetts. pp. 56-57.

[2] Hubbert, K.M., (1956). *Darcy's Law and the Field Equations of the Flow of Underground Fluids.* Petroleum Transactions, AIME, Vol. 207, pp. 222-239.

[3] Myint-U, T., (1973). *Partial Differential Equations in Mathematical Physics.* American Elsevier Publishing Company Inc.: New York.

[4] Vinsome, P.K.W., McGee, B.C.W., Vermeulen, F.E. & Chute, F.S, (1994). *Electrical Heating.* Journal of Canadian Petroleum Technology, Vol. 33, No. 9, pp. 29-35.

Appendix A

List of Participants

Rita Aggarwala	University of Calgary	rita@math.ucalgary.ca
Ibrahim Agyemang	University of Alberta	agyemang@math.ualberta.ca
Jahrul Alam	McMaster University	
Brian Alspach	University of Regina	alspach@uregina.ca
Calin Anton	University of Alberta	asy@cs.ualberta.ca
Mansour Assaf	University of Ottawa	massaf@site.uottawa.ca
Yagoob Aziz Zanjani	Frayand Roshd Company	zandjani@iasbs.ac.ir
Amos Ben-Zvi	Queens University	benzvi@chee.queensu.ca
Kyle Biswanger	University of British Columbia	kbiswan@coe.ubc.ca
Sean Bohun	University of Victoria	bohun@math.uvic.ca
Matthew Bolton	University of British Columbia	mattb@math.ubc.ca
Andrej Bona	University of Calgary	bona@math.ucalgary.ca
Chris Bose	University of Victoria	cbose@math.uvic.ca
Lloyd Bridge	University of British Columbia	lloyd@math.ubc.ca
Paul Buskell	University of Manitoba	umbuske0@cc.UManitoba.CA
Yongqiang Cao	York University	yqcao@pascal.math.yorku.ca
Samantha Carruthers	University of Victoria	sam@math.uvic.ca
Kell Cheng	University of Calgary	khfcheng@math.ucalgary.ca
Robin Clysdale	University of Calgary	raclysdale@yahoo.com
Nicola Costanzino	Simon Fraser University	ncostanz@cs.sfu.ca
Tom Cottrel	University of Calgary	cottrell@ucalgary.ca
Rozita Dara	University of Guelph	rozita@orca.cis.uoguelph.ca
Jeff DeWynne	Oxford University	dewynne@maths.ox.ac.uk
Irina Dinu	University of Alberta	andri29@yahoo.com
Paule Ecimovic	York University	paule@yorku.ca
Minglun Gong	University of Saskatchewan	mig937@mail.usask.ca
Cyril Guyot	University of Toronto	cguyot@fields.utoronto.ca
Huaxiong Huang	York University	hhuang@yorku.ca
Shafiqul Islam	Concordia University	msislam@alcor.concordia.ca
Cristian Ivanescu	University of Toronto	cristian@math.toronto.edu
Tomasz Janiewicz	Simon Fraser University	tjaniewi@sfu.ca
John King	University of Nottingham	John.King@nottingham.ac.uk

Nathan Krislock	University of Regina	krislock@cs.uregina.ca
Claude Laflamme	University of Calgary	laf@math.ucalgary.ca
Michael Lamoureux	University of Calgary	mikel@math.ucalgary.ca
Dong Liang	York University	dliang@math.yorku.ca
Margaret Liang	University of British Columbia	mliang@physics.ubc.ca
Shanea Liew	Queens University	liewcn@chee.queensu.ca
Leevan Ling	Simon Fraser University	lling@cs.sfu.ca
Mufeed Mahmoud	University of Western Ontario	mahmomm@engga.uwo.ca
Abid Malik	University of Guelph	grads205@snowhite.cis.uoguelph.ca
Andrea Mcphee	University of Toronto	mcphee@math.utoronto.ca
Jung Min Lee	University of Alberta	jungmin@math.ualberta.ca
Leonid Marian Mocofan	University of Alberta	leonid@cs.ualberta.ca
Reza Naserasr	Simon Fraser University	naserasr@cecm.sfu.ca
Dominique Noel	University of Alberta	noeldominique@hotmail.com
Mohammad Oskoorouchi	McGill University	oskoorom@management.mcgill.ca
Marc Paulhus	University of Calgary	paulhusm@math.ucalgary.ca
Suresh Pillai	University of British Columbia	pillai@physics.ubc.ca
Christina Popescu	University of Alberta	popescu@math.ualberta.ca
Miro Powojowski	University of Calgary	powoj@math.ucalgary.ca
Lila Rasekh	University of Guelph	lrasekh@uoguelph.ca
David Ross	Eastman Kodak Company	ross@kodak.com
Bruce Rout	Simon Fraser University	bbrout@sfu.ca
Erin Rozman	Douglas College	paulanderin@home.com
Simal Saujani	University of Toronto	simal@mam.physics.utoronto.ca
Ana Maria Savu	University of Toronto	ana@math.toronto.edu
Ranga Screenivasan	University of Calgary	srs@ucalgary.ca
Amir Sepasi	University of British Columbia	asepasi@coe.ubc.ca
Maurice Shevalier	University of Calgary	maurice@earth.geo.ucalgary.ca
Maikel Sianturi	University of Manitoba	umsiantu@cc.UManitoba.CA
Sirod Sirisup	McGill University	sirisup@math.mcgill.ca
Daniel Spirn	New York University	spirn@cims.nyu.ed
Fridolin Ting	University of Toronto	fting@math.toronto.edu
Satoshi Tomoda	University of Calgary	tomoda@math.ucalgary.ca
Tzvetlatin Vassilev	University of Saskatchewan	tsv552@mail.usask.ca
Niall Whelan	McMaster University	whelan@physics.mcmaster.ca
John F. Williams	Simon Fraser University	jfwillia@cs.sfu.ca
Zhihui Xue	McMaster University	xuez@mcmaster.ca
Boting Yang	Memorial University of Newfoundland	boting@cs.mun.ca
Zhiduo Zhao	University of Western Ontario	zdzhao@csd.uwo.ca

Appendix B

The Organization of an Industrial Problem Solving Workshop

Introduction

History

The "Study Groups" with Industry started in Oxford in 1968, with the aim of creating a mutually beneficial link beween researchers in industry and academic applied mathematicians. The formula of setting aside a week for intensive study of a few real world problems proved so successful that the Study Group became an annual fixture in many countries around the world. The success of such workshops can be attributed to a number of reasons, for example:

- they foster contacts between academicia and industry, sometimes leading to mutually useful collaborative programs,
- they have the potential to lead to challenging new areas of research, and research in these areas may lead to improvements in Canadia's ability to compete in a high-technology world,
- they are and excellent source of relevant research topics for graduate students and academics,
- they allow companies to meet and evaluate prospective employees.

Then there are the obvious reasons: the company needs to have a problem solved, and academics experience the pleasuresof working with enthusiastic colleagues.

Some IPSW's have a daily summing up of the day's progress– including triumphs and frustrations.

It is important to try to have at least one company representative present for the entire workshop, to answer any questions that may arise. It this is not possible, the absolute minimum is to have a representative present on Monday and Friday, and easy to contact during the week.

How to Contact Industry

Personal contacts are clearly the best way to get in touch with company. These may come about from previous workshops (get in touch with previous organizers) or, in the case of PIMS, from the industrial coordinators at various sites and the industrial working seminars which PIMS regularly hosts. Many academics have their own personal contacts, particularly from former graduate students and colleagues who are working in industry. As a group's experience with industrial modelling increases, so does its reputation and number of contacts.

To develop a new contact, a letter, fax or phone call to a company's research department or a senior engineer may work. This communication should outline

- how the meeting is set up,
- the benefits for the company and how much (or little)it will cost them.
- the cost may be a major issue with them, so stress the long- and short-term benefits from past ipsw's,
- be sure to stress problems from similar industry treated in previous workshops,
- try to line up people from industry who have enjoyed previous workshops who are willing to serve as references,

A selection of problems from previous workshops is given at the end of this Appendix.

Screening the Problems

Once problems start to come in, they require filtering. Industrial folks may not befamiliar with recent academic work, so the problems they submit may be already solved, trivial, reasonable, or unsolvable (however, an unsolvable problem might be of interest, i.e.a different approach might be suggested). Presuming the coordinator is not an expert on every subject, s/he should try to enlist expert help in deciding which category fits a given problem. It is essential that a coordinator have access to a group of experts who can assist. Screening can also be carried out by searching the literature, using a science citation index,

Even if a problem appears trctable to the coordinator, it must be suitable for the participants. This is a two-way street, however, since experts will be attracted to pariticipate in the IPSW if they see a problem which interests them.

The number of problems chosen depends on the number of participants. Typically, ten people per problem is a realistic number. Bear in mind that only a small proportion of the participants will be knowledgable in a given field,. Many people are attending to learn new topics or through curiosity.

Report/Proceeding

By the Thursday of the workshop it should have been decided by each problem group who will make Friday's presentation on the results obtained. Each group will also decide who will be responsible for coordinating the preparation of the final written report for publication in the proceedings. The actual writing is often done by several different people, and the full problem report assembled (in LaTeX) by this person. Two months after the IPSW is the deadline for this task. The problem reports are then assembled into the final report document by a member of the organizing committee. It is strongly recommended that the budget include an item for using a qualified technical typist to prepare the final LaTeXdocument.

Organization

Most requirements for an IPSW are similar to those for a normal academic meeting. the following items highlight a number of differences that should be kept in mind.

Location

An IPSW must be held near a university, since library and computer access are essential. A large room with full high-tech presentation capabilities is essential for the first and last days, so everyone can assemble to hear the initial presentations and final reports. For the daily work, small rooms are preferred, preferably located in proximity to one another to allow for easy circulation of people who want to keep up with more than one problem. These rooms should have boards and overheads, and need to be close to rooms with computer access. When booking rooms, make it clear to the University official involved that these requirements are **essential, and room booking should not changed without consultation.**

Accommodation

Academics should be housed close together, allowing the discussions to keep going into the evening (and frequently in the bar). Note, industrial people will probably prefer hotel accommodaton over university dorms; these should be close to the university so the representatives are not "left out".

Fees

Most meetings require companies to pay a fee, to help defray costs for the academics. The fees from 5-8 companies are unlikely to cover travel and accommodation costs for everyone. Typical charges range from $1600 to $4000. Govenment and other types of research grants can also be useful; emphasizing the benefits to be gained from interaction between universities and industry, such as increased competitiveness, can help in obtaining funds. Academics, and their universities, may also be able and willing to pay their own costs.

Facilities

The library should be accessible for all participants to consult journals and books and make xerox copies. It is essential to arrange full borrowing privileges and simple access to copying facilities (e.g., by providing copy cards). Local academics may be able to assist

Computer access is a must, preferably ona 24-hour basis. This includes: access to sophisticated software like MatLab (with all the toolboxes), linear programming packages, SPSS and other stats packages, Maple and Mathematica, full wordprocessing including TeX, high speed number crunching computers with ability to write routines in Fortran and C, telnet, full printing capabilities. Also, everyone needs remote access to their email accounts. These facilities must be in reasonably close proximity to the problem discussion rooms.

Food and Drink

Accommodation and meals should be provided for all participants. Coffee and tea breaks are invaluable for regular summing up, and to get honest opinions and fresh ideas from people who are working on other problems. It is useful to provide a list of restaurants, pubs and bars in the vicinity of the workshop rooms and the accommodation, so the discussions can continue in the evening.

References

Useful world-wide contacts on study group meetings/ipsw's:

Canada: Dr. Nassif Ghoussoub,*Director*, Pacific Institute for the Mathematical Sciences , Vancouver, BC
UK: Dr. John Ockendon, *OCIAM Research Director*, University of Oxford.
U.S. : Dr. Ellis Cumberbatch, *Professor of Mathematics*, Claremont Graduate School, Pomona, California.
U.S.: Dr. Donal Schwendeman, *Associate Professor of Mathematics*, Rensselaer Polytechnic Institute, Troy, NY

Web Sites:

http://www.pims.math.ca
http://macserver.maths.mu.oz.au/misg/
http://www.maths.bath.ac.uk/ESGI97
http://www.maths.ox.ac.uk/ociam (see newsletter in particular)
http://www.math.rpi.edu/Faculty/Schwendeman/Workshop/MPI2001/home.html
http://www.indmath.uni-linz.ac.at
http://www.mat.dtu.dk/ECMI
http://www.siam.org

A Selection of Problems from Past Workshops

IPSW-1: August, 1997, University of British Columbia

- Inversion for Anisotropic-Velocity Parameter (Petro-Canada)
- Fingerprint Identification (Kinetic Sciences Inc.)
- Modelling Bronchial Epithelial Lesions (BC Cancer Labs)
- Optimally Cutting Logs (MacMillan Bloedel)
- Modelling Stress Intensityin a Thermoroll (MacMillan Bloedel)
- Measuring the Stress Intensity of a Composite Vessel (Powertech Labs)
- Evaluating Computational Methods in Fuel Cell Systems (Ballard Powersystems)
- Rapid Thermal Processing of Semiconductors (Vortek Industies Ltd.)
- Optimal Strategies for Electricity Trading (Powerex Ltd.)

IPSW-2, June, 1998, University of Calgary

- An Optimal Strategy for Maintaining Excess Capacity (Boeing Corp.)
- Inventory Optimization Using a Renewal Model for Sales (Boeing Corp.)
- A Problem in Petroleum Reservoir Simulation (Computing Modelling Group Ltd.)
- On Seismic Imaging: Geodesics, Isocrons, and Fermat's Principle (The Geomechanics Group Ltd.)
- Trip Wire Detection for Land Mines (ITRES Research Ltd.)
- Torsion in Multistrand Cables (Powertech Labs Inc.)
- Automatic Detection of Egg Shell Cracks (VisionSmart Inc.)

IPSW-3, June, 1999, University of Victoria

- Contaminant Transport in Municipal Water Systems (Charles Howard and Associates Ltd.)
- Quality Control for Multi-variable Problems (Chemex Laboratories)
- Modelling Batch Interfaces (Enbridge Inc.)
- Efficient Portfolio Selection (Merak Projects Ltd.)
- Dynamics of Large Mining Excavators (RSI Technologies)
- Classification of Chemical Compound Pharmacophore Structures (Searle Corp.)

Appendix C

Detailed Organizational Guidelines

As a result of the June, 2000 workshop at the University of Alberta, the organizers prepared a detailed list of "to-do's" for the secretaries who actually carry out the organizational details.

To Do List for IPSW:

Ask organizers for structure of the workshop/camp e.g. When ? (should be from Sunday, during which there is only a welcome reception, up to the next Sunday (cheaper flights)), how many students, how many speakers, what financial aid is available for anyone, who books the flights (me, MMC or participants?), what is the budget and the breakdown? Do we pay for the participants' accommodation? Note the differences between 'participants' and 'supported students'.

Establish what the budget is and the allocation of funds - ASAP. Create binder with 9 dividers: To do, Contact Info, Social Events, Room booking (classroom and labs), Accommodations, Flight, Emails, Budget, and one spare. Call hotels/dorms to find suitable place regarding price, quality, services and availability. Negotiate with manager for better price and additional services for 'special visitors'. Book hotel (block booking) based on results.

Try to get some free (complimentary) passes to local attractions. Consider if there are posters, brochures and/or other advertising medium to create to send to the mass mailing list of all universities and international contacts. Prepare conference poster, brochure, schedule of fees, etc.

As early as possible: Find how much we can pay facilitators and then, send letters of invitation to them. Include a deadline for acceptance of invitation. Acknowledge acceptance of facilitator invitation.

Create Registration form and if needed online registration Put all information on the Web. (coordinate with PIMS main office).

Contact person responsible for MMC (math modeling camp) and keep in touch with them. If they book the flights, make sure they do so ahead of time so you can get the itineraries to plan shuttles. Ask them for their database of people (Supported students) attending the camp and whom they know is coming to the IPSW. All the flights should be very much at the same time and should arrive no later than 3 PM on the Sunday before the workshop, so that everybody can attend the first reception.

Send conference advertisements (poster, brochures, etc) to major universities, individuals, CMS membership lists, CAMS membership lists, as appropriate.

Create master mailing list of interested participants and supported students including address - for labels, name tags, mail outs, and check off list. The database should include everybody's affiliation, phone numbers and fax number as well since this information will be needed later for the computer IDs. Ask for meal requirements.

Book facilities for special functions (e.g. getaways, Banquet Hall, Faculty Club).

Send out first information package to interested participants and supported students, if not on the web. Print out website, what conference is about. This should be done only if asked by the participants/students.

Remind conference organizers to determine serious/eligible participants and cross out others from Master list.

Check out possible dinner places - faculty club, or catered - early as possible. Ask if coffee breaks are catered or done by us. Make arrangements for functions (i.e. Chose menu, purchase necessary liquor permits, PA, etc.) Note: Are there to be any head tables , etc.? If so, place name card may be desirable.) There should be a reception taking place on the Sunday before the IPSW. This is where everybody gets their registration packages with the nametags, the per diems and sign the travel claims. A banquet is often organized for the Thursday of the IPSW week. As for the presentation day, (the Friday), a pizza party could take place.

Get meal requirements ahead of time (vegetarian, etc.)

Send letter informing facilitators and industrial representatives. (should include reference to audio-visual equipment requests, computer needs, etc.); letter should specify receipt of registration fee or outline registration procedure as applicable; could include extra registration brochure.

Send immigration letter to all participants coming from outside Canada. Letter should specify that either: the recipient is not receiving any remuneration for some or all of his/her expenses. Return information from this letter (i.e. Flight times, port of entry, etc) should be submitted to the Administrative Officer (Academic) who will, in turn, submit it to the appropriate immigration authorities at the ports-of-entry. Sample letter are available from the Administrative Officer (Academic). To be done only if asked to do so by participants.

Book classrooms through the appropriate person. Book room for lounge —make clear that this is what the room is to be used for and that some furniture moving will be involved. Arrange for getting the keys later on.

Get everybody's arrival and departure dates and times, the sooner the better, including meal requirements, from MMC.

Mark what time each flight is arriving and departing on check off list. Once completed, give to hotel/accommodation. Book all the rooms for accommodation and do university paperwork if required. Send hardcopy to hotel making sure make/keep a list of the names being paid for by the university. (Usually, if the students stay in a hotel, they are at least two to a suite). Local residents are not provided with accommodation or per/diem. They also have to provide for their own transportation to and from the airport, if they are students returning from the MMC.

Draft conference schedule.

Order maps and campus information from Public Affairs .

Book computer lab room and classrooms Finalize. Send letter to all faciliators informing them of their duties. urther information about the conference, the city and surroundings, a map, etc. (similar info package than for students and participants).

Send information package to all other participants/students including preliminary conference schedule, Local information (Official Visitor's Guide), a map, map of the airport, transit information, how they get to their hotels, the shuttle schedule (see below) etc. If not enough time to send it individually, prepare packages and send to MMC site to be distributed during the camp.

Book shuttle van or car from Vehicle Pool. Arrange for shuttle van driver. Or book chartered shuttle from airport, enough vans for all arrivals. Make schedule according to check off list of flight arrival. Send schedule to the students via e-mail or include it in info package.

Order necessary supplies for conference such as: Conference brochures, name badges, campus maps, city maps, notebooks and pens (if desirable), etc. Also order name inserts, holders, disks, paper, transparencies, pens, etc.

Arrange for audio-visual technician (Audio-Visual Services may have some suggestions.) If needed. Arrange for coffee/tea/refreshment services during conference period. Decide where take place .

Determine detailed computer/programming needs - e.g., statistical tool box, maple, mathematica, matlab. Arrange through university Computing Systems for E-mail access for conference participants. Is access available 24 hours a day in a computing lab close to the conference rooms? Get computer IDs forall participants. In order to give us the temporary passwords, they will generally need a database including full name, address, and phone number of each out of campus participants and students.

Make nametag inserts (should include initials, surname and home institution). Print passwords and login ID from timetabling on small stickers and stick behind each appropriate nametags. Put nametags in small box in alphabetical order for registration.

Prepare and submit visiting speaker payment forms . Process Travel advances. If they get per diem money, ask organizer to go to bank and get small bills (loonies, toonies, 20$ and 50$) according to how much each individual would receive. Money is handed out at initial wine and cheese.

Contact and send letter to Transportation Canada for arrangements regarding a booth and/or shuttle van parking at the Airport. (Not to be done if you have booked a chartered shuttle).

Book rooms (offices) for facilitators from Val. Get the appropriate keys.

Make necessary signs. These could include: airport sign (for meeting point, saying IPSW), conference registration sign, directional signage on campus to conference sight, program change sign, etc. Arrange for signs to be laminated if desired.

Ask speakers and organizers if need for projector or computer screen projector system. Book necessary audio-visual equipment. Make necessary arrangements for storage and implementation of the equipment.

Invite local people to key events– Dean of Science, Academic/Research VPs.

Ask how many vegetarians, and establish RSVP time - two weeks in advance.

Send out second information package including updated schedule, shuttle pick up schedule, the instructions for them to call and book own departure shuttle, etc. if not already in first

package. If no time, send an e-mail with all the updated info or send to MMC site for distribution.

Arrange for registration desk set-up (if some registration is to occur prior to the actual start of the conference). Make sure have the appropriate space for a registration table to be set up at the first reception. Confirm function arrangements. Check when firm numbers are required and determine needs - podium, screen, overhead projector, meal requirements etc.

Create check off list for each event - attendance for dinner, pizza party, receptions, welcome parties, per diem given, travel claims signed in alphabetical order.

Finalize and Organize refreshments for breaks.

Call Sky Shuttle for arranging Chartered shuttles for pick up from the airport . The return shuttle will be booked by the students themselves.

Prepare social tickets. Each ticket could include: the ticket number, the cost of the tickets, the functions it includes, the locations and time of each function.

Send a fax to Building Services if IPSW falls on a holiday eg. Victoria Day or if event is over the weekend. Security may have to be turned off and the doors may have to stay unlocked. Get keys for every room and all labs booked, library, and outside entrance doors.

Prepare any cheques for visiting speakers out of conference account. Make sure have everybody's full mailing address, if not, prepare sheet to be filled while giving away per/diem and cheques during registration.

Set dinner events - order food. Finalize details with Liquor choices and order. Buy liquor permit. Inform the necessary people of the firm numbers for functions. Confirm details.

Confirm access/programs for computers - account ID's . Get copy cards and library cards available. Replenish copy cards.

Prepare registration packages including the relevant information for each individual (see below). Registration package could also include: per/diems, abstract book from the year before, problems description, conference program,city/campus information, advertising, notebook, pen, name badge, social tickets (If being held off-campus, the invitation should probably also include an area map), etc. Be sure to prepare an adequate number of "blank" registration packages and nametags for last-minute drop-in participants. Have any drop-in registrants complete a form giving their name and home/institution address to be later added to the mailing list. Place in hotel/accommodations room or distribute at first function if take place before the workshop starts. For all those participants receiving cheques, include mention that they can get assistance for cashing their cheque.

Make an explanation sheet on how to use the computer passwords and where they can use it. Put in registration packages.

Ask any accompanying spouses if they would like to have a nametag as well or prepare ahead of time.

When distribute registration packages, let them know we will select one or two persons per industrial problem responsible for photocopy cards and a set of keys (in exchange for an ID to ensure they will return all items), and let everybody know we will have keys for the rooms booked (labs and classrooms). If want to borrow books from library, have to come see us to get a library card and will have to return it right away. Make a list of who gets which copy card and keys.

Provide a box where participants can recycle their name badges after the conference is over. This is to be said when handling out registration packages.

Update Budget sheet.

Post all necessary signs and send the updated schedules to the hotels/accommodation for them to post it.

Make a box including pads of paper, pens, overhead transparencies and pens, white board pens, chalk, as well as any announcement to be made in the morning (to be updated every day).

During IPSW

Each morning, check if rooms are opened and if overhead projectors are working etc. and bring box.

Take pictures! – of each group working on the problems, as well as of any recipient of a prize . Get a bottle of PIMMS (which is an English liquor) for the PIMSlip award. To be awarded on the Friday of the presentation.

Make sure everyone signs their travel reimbursements forms and provides air tickets and boarding passes etc.

Prepare a press release before the presentation for the Friday and send to local media. Prepare and post up-to-date conference schedule on a day-to-day basis. Revise list of participants as required and make copies of the list available to all participants by the second-to-last or last day of the conference. (Keeping in mind FOIPP) Keep track of long-distance calls made by participants and their charges. These charges should be paid by the participants to the conference in cash.

Keep track of:

- Printing costs
- Supply costs
- Support staff salaries
- Postage costs
- Long-distance telephone call costs
- Facilitator costs
- Social Function costs
- Incoming Funds - grants, registration fees, including membership fees, if any
- Laminating costs
- Shuttle van costs
- Coffee breaks cost

After Event

Any money received as payment of participants' long distance calls should be reimbursed to the appropriate person.

Prepare and submit to relevant people or organizations a final conference budget.

Send letter of thanks to facilitators, industrial reps, and helpers.

Ensure all invoices/per diems/travel reimbursements have been paid.

PIMS Contact Information

email: pims@pims.math.ca

http://www.pims.math.ca

Director's Office
Pacific Institute for the Mathematical Sciences
Room 200, 1933 West Mall
University of British Columbia
Vancouver BC V6T 1Z2
Canada