Pacific Institute
for the Mathematical Sciences

Proceedings of the Sixth

PIMS Industrial Problem Solving Workshop
PIMS IPSW 6

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 Sixth Annual PIMS Industrial Problem Solving Workshop

Editor: J. Macki, University of Alberta

June, 2002

π

Foreword by the PIMS Director

The Sixth Annual PIMS Industrial Problem Solving Workshop was hosted by the PIMS at the University of British Columbia in Vancouver, May 27–31, 2002. More than 100 participants worked intensely on seven problems posed by industrial companies from across North America.

The problems were provided by Edmond Lou (Capital Health), Bruce McGee (McMillan-McGee Corp), Robin Parlee (Precix Advanced Cutting Technologies), Yeqi (Ritchie) He (RBC Financial Group), Edward Keyes (Semiconductor Insights), Kai Muenzer (Shell Canada), and Gary Billings (Talisman Energy).

PIMS looks forward to the Seventh Industrial Problem Solving Workshop which will be held at the University of Calgary next year.

Thank you to Jack Macki from the University of Alberta who edited these proceedings. I would also like extend my thanks to the other organizers Chris Bose (U. Victoria), Randy LeVeque (U. Washington), Huaxiong Huang (York U.), Marc Paulhus (U. Calgary), Manfred Trummer (SFU), and Ian Frigaard (UBC).

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

π

ii

Contents

$$\pi$$

π

List of Figures

π

List of Tables

Preface

In June of 2002, approximately forty-seven graduate students, twenty-one university professors and five industry representatives from five different companies/institutions assembled on the campus of the University of British Columbia.

A small subset of the academics had agreed to be *mentors*. They studied the problems posed by the industry representatives during the winter, and devised "training problems" for the students. The students honed their skills on these training problems during the preceding week at a camp. Then the students and mentors connected with the industry representatives and a group of interested academics (many of whom brought vital skills) for the IPSW.

The results were, as one would expect, mixed, ranging from spectacular to a small beginning. But all participants came away inspired by the experience, and the industry representatives were the beneficiaries of valuable new insights into their problems.

The problems were:

Stitching Integrated Circuit Images posed by Semiconductor Insights,
Price Pseudo-Variance and Pseudo-Volatility Swaps, posed by RBC Financial Group,
Electromagnetic Detection of Oil or Gas Deposits using Electromagnetic Waves posed by Shell Canada,
Seismic Inversion in Anisotropic, Inhomogeneous Media posed by Talisman Energy,
Mathematical Surface Matching to Diagnose Scoliosis posed by the Capital Health Region (Edmonton),
Resistance Monitoring in Well Sites posed by McGee-McMillan Corporation.

We hope the reports convey the intensity and pleasure of our participation in this workshop.

π

Acknowledgments

Sincere thanks:

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

π

Pimslips

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

- Ed, speaking with a colleague in the correct room for their project:
  "I saw him at breakfast and he asked if it was in this room"
  "What did you tell him?".

- Semiconductor Insights presentation (A slide of a huge integrated semiconductor board is being shown):
  "We are working with real data–no one said 'this is a picture of my aunt Minnie' ".

- Question: "Any progress for the graph theory group?"
  Answer (Jenny): "We have three approaches and they both work."

- Hugh, discussing models for seismic and electromagnetic detection:
  "The concept of 1D in geophysics is different from the mathematical definition of 1D..." .

- Alistair: "How many feet in a Newton?"

- Ian: "That is not constant enough".

- Brian: "Something that is elliptical is necessarily an ellipse."

- Jeff: "Going from 2D to 3D adds a whole new dimension to the problem."

- ....and the winner is..... Fabien Youbissi, on checking in at Gage Hall (remember, we are at the University of British Columbia):
  Receptionist: "Name please".
  Fabien: "Youbissi".
  Receptionist: "No, your name, not your university".
  Fabien: "Youbissi, Youbissi is my name" etc.....

π

Chapter 1

Stitching IC Images

Michael Baumann[1], Chris Bose[2], Lorraine Dame[2], Isabelle Dechene[3], Maria Inez Goncalves[2], Robert Israel[1], Edward Keyes[4] (presenter), Frithjof Lutscher[5], David Lyder[6], Colin Macdonald[7], Jenny McNulty[8], Joy Morris[9], Nancy Ann Neudauer[10], Benjamin Ong[7], Robert Piché[11] (editor), Arūnas Šalkauskas[12], Karen Seyffarth[13], Brett Stevens[14], Tzvetalin Vassilev[15], Brian Wetton[1], Fabien Youbissi[16]

1.1   Introduction

Errors in stitching (or mosaicing) of integrated circuit (IC) images cause errors in the automated generation of the schematic. This increases costs and introduces delays as engineers must step in with interactive computer tools to correct the stitching. Our goal was to develop automated image stitching methods that keep stitching errors under $F/2$, where the minimum feature size F is around 10 pixels.

Although image stitching software is used in many areas such as photogrammetry, biomedical imaging, and even amateur digital photography, these algorithms require relatively large image overlap. For this reason they cannot be used to stitch the IC images, whose overlap is typically less than 60 pixels for a 4096 by 4096 pixel image.

[1]University of British Columbia
[2]University of Victoria
[3]McGill University
[4]Semiconductor Insights Inc.
[5]University of Alberta
[6]King's University College
[7]Simon Fraser University
[8]University of Montana
[9]University of Lethbridge
[10]Pacific University
[11]Tampere University of Technology
[12]Sorex Software Inc.
[13]University of Calgary
[14]Carleton University
[15]University of Saskatchewan
[16]Laval University

IC images are currently stitched as follows. First, the offset between each image and its neighbours is determined by minimising an error function that measures the difference between features in overlapping image pairs. These offsets are then used to stitch the images together in some predetermined pattern. Errors tend to accumulate as the images are added sequentially to the grid.

In Section 1.2 we use algorithmic graph theory to study optimal patterns for adding the images one at a time to the grid. In the remaining sections we study ways of stitching all the images simultaneously using different optimisation approaches: least squares methods in Section 1.3, simulated annealing in Section 1.4, and nonlinear programming in Section 1.5.

1.2   Stitching using Graph Theory

Here we present a graph theoretic approach to the problem. This approach is purely local in the sense that we only use the fact that any given pair of adjacent images can be stitched together perfectly, i.e., all features match, due to the offset data from pairwise correlation. At present, these local correlations are used for stitching following a certain pattern. We propose different patterns which should give better global results. To that end, we formulate the problem in a graph theoretic way and introduce some error measures. We find and compare several alternative stitching patterns which have smaller errors than the one currently used. Finally, we extend the approach to include information about the reliability of the pairwise correlation data.

1.2.1   Notation

We assume the images are arranged in a rectangular shape of m rows and n columns, where $m \geq n$. Define the graph G by letting each image correspond to a vertex. Two vertices are joined by an edge if the corresponding images are adjacent. Hence, G is an $m \times n$ grid graph. We denote by $E(G)$ the edges of G, and for adjacent $u, v \in G$ we write $uv \in E(G)$.

We use the stitching pattern to define a *spanning tree* T of G, i.e., a connected subgraph without cycles. An edge of G is an edge of T if the corresponding images are stitched together perfectly according to the offset values. The current method of row-wise stitching corresponds to the row pattern given by the solid lines in Figure 1.1. Here, the images are stitched together row-wise and then the strips are stitched together along a "spine"

We assume that the error between u and v with respect to the true offset values increases with the distance between u and v in the tree T. To quantify this error, we define the following. For each $uv \in E(G)$,

$$
\begin{aligned}
d'_T(u,v) &\quad \text{denotes the distance between } u \text{ and } v \text{ in the tree } T; \\
d_T(u,v) &= d'_T(u,v) - 1; \\
S(T) &= \Sigma_{uv \in E(G)} d_T(u,v) \text{ (i.e., the sum of all the distances, in } T, \text{ between} \\
&\quad \text{adjacent vertices of } G); \\
M(T) &= \max_{uv \in E(G)} d_T(u,v).
\end{aligned}
$$

π
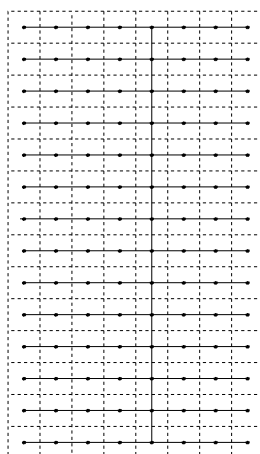
Figure 1.1: 14×8 **Row Pattern**, R

Note that we subtracted 1 from the distance in the tree to get d_T because we assume that pairwise stitching is free of error. We address the following optimization problems.

1. Find a spanning tree T of G that minimizes $S(T)$, the sum of the distances of the tree.

2. Find a spanning tree T of G that minimizes $M(T)$, the maximum of the distances of the tree.

3. Find a spanning tree T of G that minimizes $S(T)$, the sum of the distances of the tree, subject to $M(T)$ being as small as possible.

We calculate the total distance and the maximum distance for the row pattern R (Figure 1.1) of an $m \times n$ grid. It is optimal to centre the "spine" in the long direction. In this case $M(R) = 2\lfloor \frac{n}{2} \rfloor$ is as small as possible, and the total sum is given by:

$$S(R) = \begin{cases} (m-1)(n^2)/2 & \text{if } n \text{ is even} \\ (m-1)(n-1)(n+1)/2 & \text{if } n \text{ is odd} \end{cases} .$$

We note that this is the method currently employed.

1.2.2 Alternative Stitching Patterns

While the row pattern minimizes the maximum distance it does not minimize the total distance. In fact, one can replace a single edge of the tree R with a nearby edge and decrease the total sum. A pattern that reduces the total sum is the **Comb Pattern**, shown in Figure 1.2. This pattern is similar to the row pattern in that it has a long "spine", but it differs in that instead of long "lines" attached to the spine, it has "combs". It is again optimal in terms of $S(C)$ and $M(C)$ to centre the "spine" in the longest direction, as above. Then the maximum is $M(C) = 2\lfloor \frac{n}{2} \rfloor + 4 > M(R)$, but $S(C) < S(R)$. The values of $S(C)$ depend on the modulus of the parameters m and n, mod 3 and mod 2, respectively. Thus, there are six cases to consider; they are given in Table 1.1.

π
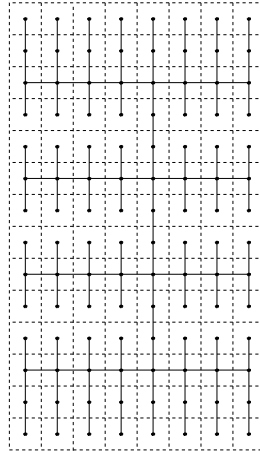
Figure 1.2: 14×8 Comb Pattern, C

m mod 3	n odd	n even
0	$\frac{1}{6}(m-3)(n-1)(n+17) + 4(n-1)$	$\frac{1}{6}(m-3)(n^2+16n-16) + 4(n-1)$
1	$\frac{1}{6}(m-4)(n-1)(n+17) + 8(n-1)$	$\frac{1}{6}(m-4)(n^2+16n-16) + 8(n-1)$
2	$\frac{1}{6}(m-5)(n-1)(n+17) + 12(n-1)$	$\frac{1}{6}(m-5)(n^2+16n-16) + 12(n-1)$

Table 1.1: Values for $S(C)$

The next method is to create a breadth-first search spanning tree which leads to what we call the **Breadth Pattern**, B. Let G_e be an $n \times n$ grid graph with n even and G_o be an $n \times (n-1)$ grid graph with n odd. Define an "almost square" grid graph to be a graph of the form G_e or G_o. We first define the breadth patterns B_e and B_o for these graphs. Label each vertex of the graph by the shortest distance to the perimeter of the grid. Figure 1.3 illustrates the 9×8 case. For the graph G_e, there are 6 vertices labeled $(n-3)/2$ in the centre and for the graph G_o, there are 4 vertices labeled $(n-2)/2$ in the centre. Begin with a spanning tree of these centre vertices as indicated in Figure 1.3.
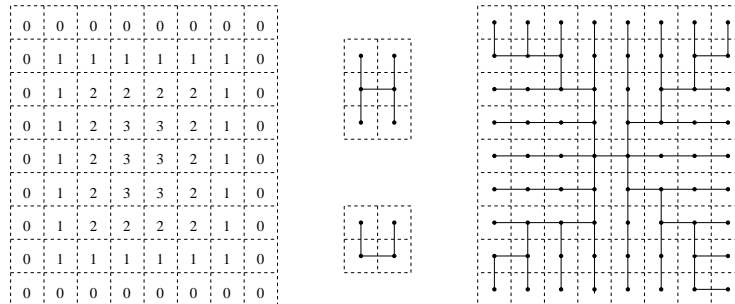
Figure 1.3: Breadth Pattern Construction

The procedure for growing the tree begins in the centre and radiates outwards. Suppose the tree has been grown to include all vertices labeled at least j, where $j > 0$. Add vertices labeled

$j - 1$ as described below.

- If a vertex x with label $j - 1$ is adjacent to a vertex y with label j, then y is unique. Add xy to the tree.

- If a vertex x with label $j - 1$ is **not** adjacent to a vertex with label j, then x has two neighbors labeled $j - 1$. Arbitrary choose one of these, say y, and add xy to the tree.

Figure 1.3 shows one possible breadth pattern for a 9×8 grid graph. We can easily calculate the sum of the distances for B_o and B_e in the almost square cases.

$$
\begin{aligned}
S(B_o) &= (n-1)(n+1)(2n-3)/6, \\
S(B_e) &= n(n+1)(n-1)/3.
\end{aligned}
$$

The breadth pattern, B, for any $m \times n$ grid graph is formed by placing half of the breadth pattern for the almost square grid on the top and bottom, and a row pattern in the middle. See Figure 1.4 for an example of this construction.
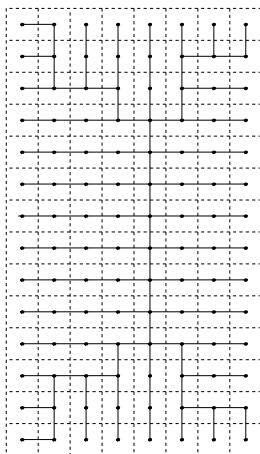
Figure 1.4: 14×8 Breadth Pattern, B

The maximum distance $M(B) = 2\lfloor \frac{n}{2} \rfloor$ is minimal, while the total sum is smaller than the row pattern sum. The values of $S(B)$ for the breadth pattern B of an $m \times n$ grid are given below.

$$
S(B) = \begin{cases} (3mn^2 - n^3 - 2n)/6 & \text{if } n \text{ is even,} \\ (3mn^2 - n^3 - 3m + n)/6 & \text{if } n \text{ is odd.} \end{cases}
$$

Comparing these patterns with the current row pattern, we see that the comb pattern improves the total sum at the expense of an increase in the maximum distance while the the breadth pattern improves the total sum while keeping the maximum distance at a minimum. We call a combination of these patterns a **Combed–Breadth Pattern** (Figure 1.5).

π
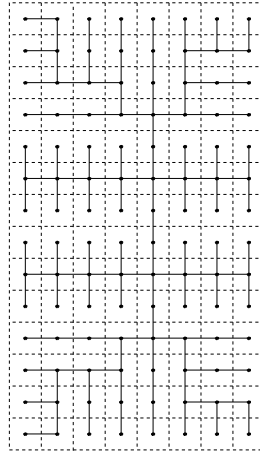
Figure 1.5: 14×8 Combed–Breadth Pattern, CB

1.2.3   Reliability

Not all the pairwise optimal offset values have the same confidence levels. We incorporate uncertainties by not allowing adjacent images to be stitched together if the corresponding offset values are highly uncertain. This is implemented by disallowing certain edges from the spanning tree.

This rule for excluding edges in the tree can be incorporated into breadth pattern. In Figure 1.6 the excluded edges are marked by grey barriers. Barriers were put where the uncertainty was higher than 7%. The tree is grown subject to the extra condition that it cannot grow across a barrier. We call this a **Barred Pattern**. One can also take the row pattern and the comb pattern and adjust them locally such that none of the edges of the spanning trees cross a barrier.
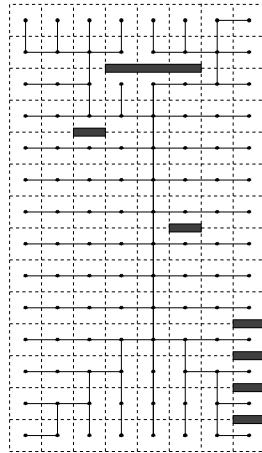
Figure 1.6: Barred Pattern, BA

In all cases, the numbers M and S cannot decrease. For this particular barred pattern, $M = 10$ for the modified row pattern and $M = 12$ for the modified comb and breadth patterns.

$$\pi$$

pattern	$S(\cdot)$	$M(\cdot)$
R (row)	416	8
C (comb)	234	12
B (breadth)	360	8
CB (combed-breadth)	258	12

Table 1.2: $S(\cdot)$ and $M(\cdot)$ values for a 14×8 grid

Also, the sum, S, increases by 4 for all three patterns. Table 1.2 contains calculations for Data set 1. In addition to the barred pattern in Figure 1.6 we also constructed a **Greedy Pattern** G, where we chose edges solely on the basis of confidence. The calculations for these patterns are also included in Table 1.2.

1.2.4 Data

Table 1.2 gives the $S(\cdot)$ and $M(\cdot)$ values for a 14×8 grid for the four tree patterns that we have defined.

Using the two sets of real data available, we evaluated the performance of each of these tree patterns. After stitching together the full image according to each pattern, we calculated the new pairwise relative offsets that each of these methods gave us. Then we calculated the differences both horizontally and vertically between these and the given "ideal" relative offsets. The larger a difference is, the more likely that it will create a significant misalignment that will lead to a problem in recreating the circuit. Thus we created three measures of the "badness" of each of our results.

The first measure was a sum of the absolute values of all of the horizontal and vertical differences between the relative offsets we obtained and the ideal relative offsets. This is denoted ABS ERROR. The second measure was chosen to highlight the fact that if a particular pair of images is placed (relative to one another) very far from their ideal relative alignment, this is worse than having a lot of images that are pairwise only slightly off their ideal alignment. Here we took the sum of the squares of all of the horizontal and vertical differences. This is labelled SQ. ERROR. The third measure took the threshold value of plus or minus 5 pixels as the worst tolerable error, and counted the number of values (horizontal or vertical) that were more than 5 pixels from their ideal relative alignment. This is labeled BAD VALUES in Tables 1.3, 1.4.

The barred and greedy patterns are present in Table 1.3 but not Table 1.4, since confidence data were only available for the first data set.

From this data, it appears that the breadth pattern produces the best results, although the barred pattern may improve on it in some ways. Similar to the barred pattern, one can modify the row and comb patterns to disallow certain edges. It seems that these highly regular patterns are quite sensitive to the precise locations of the barriers. If one changes an arbitrarily chosen edge in the spanning tree and replaces it in an optimal way, the breadth pattern performs much better than the row and comb patterns. This may be in part because there are many choices for constructing such a breadth pattern.

The greedy tree, although using high confidence values, resulted in extremely large values

π

pattern	ABS ERROR	SQ. ERROR	BAD VALUES
R	518	5092	32
C	513	5067	37
B	449	4229	30
CB	513	4897	39
BA	453	4309	27
G	721	7059	52

Table 1.3: Data Set 1

pattern	ABS ERROR	SQ. ERROR	BAD VALUES
R	569	6007	27
C	636	7560	34
B	517	5639	27
CB	619	7507	34

Table 1.4: Data Set 2

of M and S, and the effect of these is apparent in its poor performance. Of course, additional experimental data would be desirable, and it is likely that situations would arise in which different patterns provided the best performance, so one reasonable approach would be to maintain this as a small collection of trees to try on any given data set, and keep whichever final result is best. Some of these patterns could ultimately be discarded if additional experimental data showed them to consistently perform poorly.

We note that the breadth pattern continually outperforms the row pattern which is currently being used. In addition, the barred pattern may prove to be useful at a higher tolerance. Currently, a tolerance of 10% is used; however we chose 7% rather arbitrarily based on the one data set available. We also note that while the comb pattern has the minimum value of the total distances $S(C)$ it performs as well as the row pattern. Thus, it seems that the solution we seek will be the solution to optimization problem 3. Currently, the breadth pattern is the best candidate for a solution to that problem.

1.3   Least Squares

The pairwise offsets overdetermine the positions of images in the final grid. In this section we develop stitching solutions by simultaneously minimising the sum of the squares of the offset discrepancies.

1.3.1   Offset Equations

Let $(X_{i,j}, Y_{i,j})$ be the coordinates (in a global reference frame) for the northwest corner of the image in row i and column j (Figure 1.7). The grid of images has m rows and n columns.
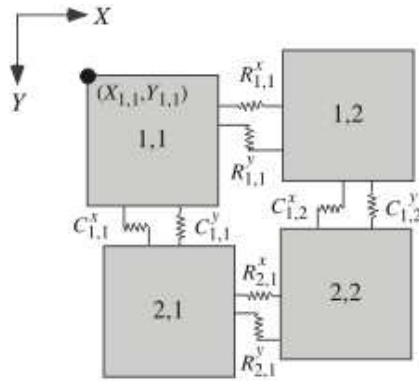
Figure 1.7: Coordinates and labelling of images in the grid. In least squares stitching the images can be interpreted as rigid plates connected by springs.

The x-component offsets $R^x_{i,j}$ between adjacent images in a row are described by the equations

$$X_{i,j+1} - X_{i,j} = R^x_{i,j} \quad (i = 1 : m, \; j = 1 : n - 1)$$

These $m(n-1)$ row stitching equations can be written compactly as

$$XB^T_n = R^x \tag{1.1}$$

where B_n is the $(n-1) \times n$ bidiagonal matrix

$$B_n = \begin{bmatrix} -1 & 1 & & & \\ & -1 & 1 & & \\ & & \ddots & \ddots & \\ & & & -1 & 1 \end{bmatrix}$$

The x-component offsets $C^x_{i,j}$ between adjacent images in a column are described by the equations

$$X_{i+1,j} - X_{i,j} = C^x_{i,j} \quad (i = 1 : m - 1, \; j = 1 : n)$$

These $(m-1)n$ column stitching equations can be summarised as

$$B_m X = C^x \tag{1.2}$$

The foregoing discussion can be repeated for y-components, leading to the stitching equations

$$YB^T_n = R^y, \quad B_m Y = C^y$$

These equations are of the same form as the equations (1.1–1.2) for the x-components, and so may be treated in the same way. The y-component equations are not coupled with the x-component equations, so they may be computed separately.

In order to use standard software, it is convenient to convert the stitching equations into matrix-vector form. The Kronecker product is a standard technique for this sort of conversion [1]. The matrix of image location x-components X is converted into a vector by stacking the columns,

$$\bar{x} = \text{vec}(X) = \begin{bmatrix} X_{1:m,1} \\ \vdots \\ X_{1:m,n} \end{bmatrix}$$

Converting the row stitching equations (1.1) gives

$$(B_n \otimes I_m)\bar{x} = \text{vec}(R^x)$$

where I_m is the $m \times m$ identity matrix and $\otimes$ is the Kronecker product. Similarly, the column stitching equations (1.2) are converted to

$$(I_n \otimes B_m)\bar{x} = \text{vec}(C^x)$$

The system of stitching equations is thus $\bar{A}\bar{x} = b$ where

$$\bar{A} = \begin{bmatrix} B_n \otimes I_m \\ I_n \otimes B_m \end{bmatrix}, \quad b = \begin{bmatrix} \text{vec}(R^x) \\ \text{vec}(C^x) \end{bmatrix}$$

Finally, we impose the condition $\bar{x}_1 = X_{1,1} = 0$, which fixes the global coordinate system's origin. This yields the system

$$Ax = b \tag{1.3}$$

where A is $\bar{A}$ with the first column deleted, and x is $\bar{x}$ with the first element deleted. The system (1.3) has full column rank.

1.3.2   Least squares

Equation system (1.3) has $(m-1)n + m(n-1)$ stitching equations for the $mn - 1$ unknowns in x. This system of equations is overdetermined whenever $m > 1$ and $n > 1$. Unless the offsets R^x and C^x happen to be consistent, any set X of image locations will give some nonzero error ("residual") in the equations.

One way to determine X is to seek the minimiser of the sum of squares of the residuals of equations (1.1–1.2) subject to the grounding condition $X_{1,1} = 0$. This is equivalent to the unconstrained minimisation of

$$(Ax - b)^T(Ax - b)$$

For a physical interpretation of the least squares formulation, imagine the grid as an array of rigid plates connected by elastic springs (Figure 1.7). A stitching equation that is satisfied corresponds to a spring that is neither stretched nor contracted. There are more springs than there are degrees of freedom, so the springs stretch or contract to find the configuration of minimum potential energy. This configuration is the least squares solution, and the residuals are the extensions in the springs.

π
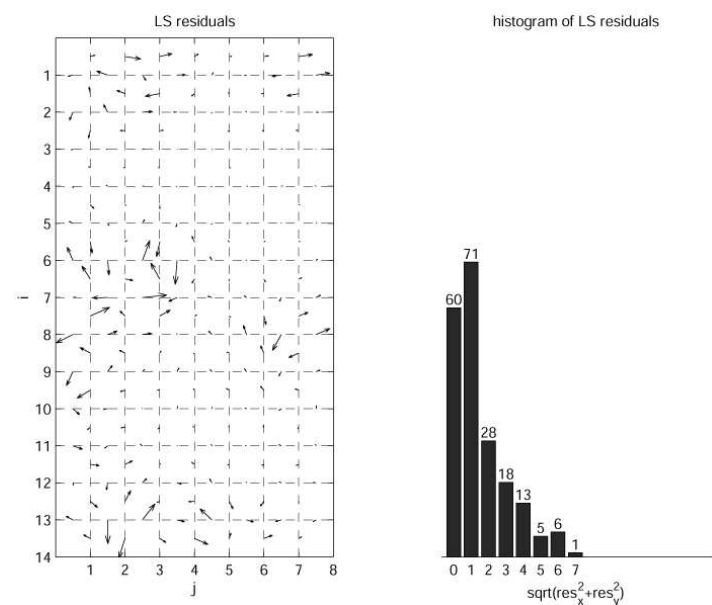
Figure 1.8: Residuals for least squares stitching

The least squares minimiser is the solution of the square system of "normal" equations

$$A^T A x = A^T b$$

The normal equation system is nonsingular because A has full rank, and it can be solved reliably and efficiently using standard sparse linear algebra software.

The Matlab script `LSstitching.m` [2] does least squares stitching of a 14 by 8 grid of images using offset data provided by the problem proposer. Computing time on a 400 MHz Macintosh is 0.02 s. The least squares solution's residuals are mostly 0–3 pixels, although a few seams have residuals of 4–7 pixels (Figure 1.8). The problem proposer did a visual inspection of the least squares solution and reported that all but one of the seams appears to be OK.

1.3.3 Constrained Least Squares

In order to reduce the maximum residual obtained by the least squares solution we impose the inequality constraints

$$-\alpha \le b - Ax \le \alpha$$

where α is the desired limit on the residual's absolute value. The smallest value of α that gives a feasible solution can be found by trial and error or by solving a linear programming problem. The constrained least squares problem can be solved using standard optimisation software packages.

In script `CLSstitching.m` [2] we used the Matlab Optimization Toolbox (v. 1.5.2) routine `conls`. For the 14 by 8 grid the smallest feasible α was 4 pixels for the x-components and 5 pixels for the y-components. The maximum residual thus has length $\sqrt{4^2 + 5^2} \approx 6$, which is only 1 pixel smaller than for the unconstrained least squares solution, and many more of the residuals are large (Figure 2), so this solution seems to be of little practical value.

π
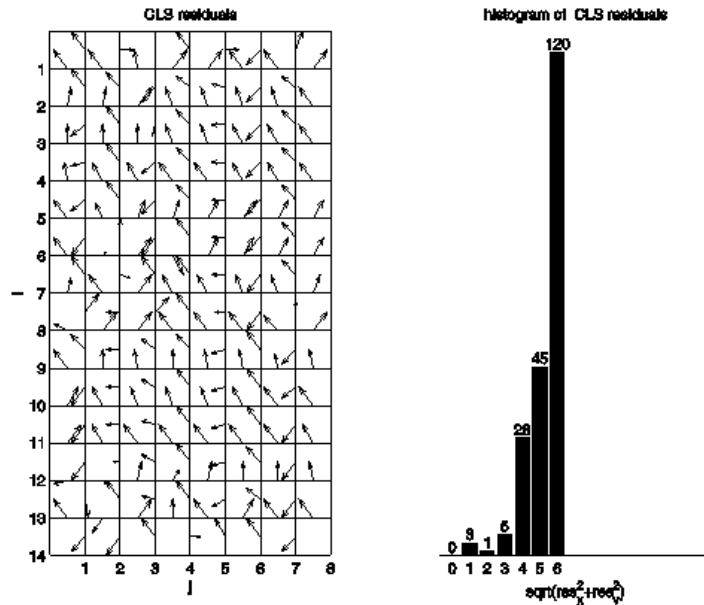
Figure 1.9: Residuals for constrained least squares stitching

1.3.4   Weighted Least Squares

Another approach to improving the practical quality of the stitching is to minimize the weighted sum of squares of the residuals,

$$(Ax - b)^T W^2 (Ax - b)$$

where the weighting matrix W is a nonsingular diagonal matrix. The diagonal elements of W can be interpreted as stiffness constants for the springs holding together the rigid plates in Figure 1.7. With $W = I$ all the springs have equal stiffness, and we revert to the original "unweighted" least squares problem. The weighted least squares problem can be solved by simply replacing A and b in the unweighted problem by WA and Wb.

One possible strategy for assigning the weights would be to give larger weights to those equations where the quality of the stitching is most sensitive to stretching or shear. For example, in parts of the images that contain a lot of parallel lines we want the lines to match up correctly at a seam, so we assign a large stiffness to the appropriate spring. A seam that contains no features would be assigned a small stiffness. It might be possible to derive this kind of sensitivity information automatically from the correlation data that were used by the problem proposer to compute the offsets; we would need more information about the images and the correlations to explore this possibility.

Another strategy would be to assign the weights using interactive graphics. Based on the solution of the unweighted least squares problem, a human operator would examine the seams where stitching quality is unacceptable, adjust those seams' stretching and shearing stiffnesses, and recompute the weighted least squares solution. This process could be iterated until (hopefully!) all the seams were acceptable. If the initial solution gives only a small number of unacceptable seams, this interactive adjustment would not be too laborious.

A third strategy, which uses only the information given in the problem, is to assign small
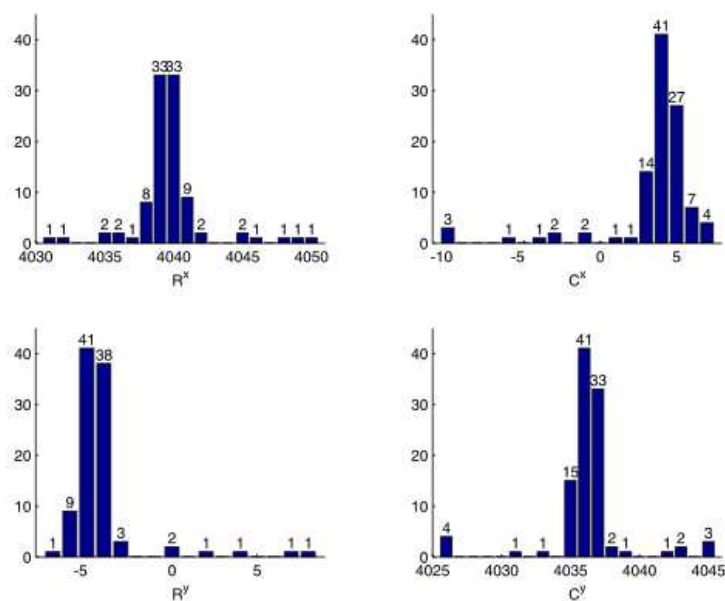
Figure 1.10: Histograms of offset values

weights to those equations whose offset value differs significantly from the mean value. The rationale is that these offsets are erroneous ("outliers"), because of some sort of measurement error, and should be ignored.

Applying this third strategy to the 14 by 8 grid, we assigned weights of $\sqrt{0.1}$ to those equations whose offset value differed by more than 3 pixels from the mean (Figure 1.10); the remaining equations' weights were 1.0. In the weighted least squares solution (`WLSstitching.m` [2]) the residuals of the outliers tends to increase (compared to their values in the unweighted least squares solution), but the remaining residuals mostly decrease (Figure 1.11). The problem proposer reports that with the weighted least squares solution all the seams appear to be OK, so this strategy seems to be promising.

1.3.5   Least Squares with Deformable Images

A way of reducing the residual at all the seams is to introduce more degrees of freedom by allowing the images to deform in various ways. We introduced deformability by dividing each image into four equal subimages. Theoretically, the sum of squares of the residuals scaled by the length of the seams of the subimages should decrease. Applying this approach to the 8 by 14 grid gave a maximum residual of 5 pixels, which is 2 pixels less than with the original least squares solution. We believe that more sophisticated deformation models, combined with weighting and more offset data (for example, measuring several x-offsets along an east-west seam), could bring all the seam residuals down to less than 1 pixel.

π
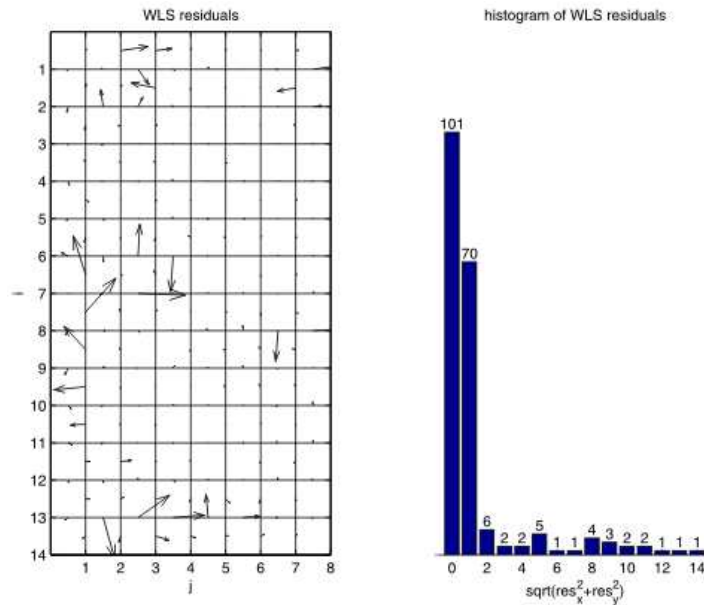
Figure 1.11: Residuals for weighted least squares stitching

1.4   Simulated Annealing

A more fundamental approach to stitching is to directly minimise an error function computed from the overlapping pixels in the entire image. In this section the minimisation problem is solved using the method of simulated annealing. Continuous minimisation methods for this problem are studied in section 1.5.

Simulated annealing is one of a family of meta-heuristics that are able to solve instances of hard optimization problems. These meta-heuristics take advantage of the ability to quickly define a local perturbation and then to recalculate the global objective function by a fast, local calculation. This family of meta-heuristics also includes *hill-climbing, Tabu search, genetic algorithms,* and *rising flood* techniques. The central problem is to avoid local extrema. Simulated annealing does this by analogy to cooling solids. Initially the setting is "hot", meaning that perturbations that worsen the objective function are sometimes allowed. We gradually "cool down" the setting, meaning that fewer bad moves are acceptable, as we hone in on better objective values. To avoid getting stuck in a local extremum, we periodically crank up the heat and again slowly cool. Repeating this leads to surprisingly good solutions.

We wrote a C++ simulated annealing program for image stitching. It is licensed under the GNU GPL and available for download [3]. The program has sophisticated error handling routines that identify well-formatted input data. Input parameters include minimum feature size, image translation bounds, temperature, number of reheating and recooling stages, and temperature thresholds.

We first defined an objective function that measured how far each image is placed from its ideal offsets using an error threshold defined by the minimum feature size. A seam whose error is less than this threshold in both coordinates contributes 0 to the objective function, otherwise it contributes 1. The program randomly selects an image to move and makes a small translation of

the image, locally recalculating the overall objective function by comparison to the ideal offsets for that single image. This objective function always rapidly diverged to the worst case (all boundaries bad) and was never able to improve. We believe the problem was the fact that the objective function had no way to recognize if a boundary was non-optimal but very close. It registers the same value as the boundary being far from a good fit. Thus once a boundary was bad, the vast majority of perturbations would leave the objective function unchanged leading to wandering on plateau in the state space.

We then defined an objective function that scored 0 if the alignment is within a feature size and scored a value that was proportional to how far away from correct the boundary was if it was not a good fit. This rapidly converges to a good layout.

Instead of doing merely local translations, it is straightforward to incorporate small rotational changes, shrinking and expanding changes, and slight skewing changes into this simulated annealing algorithm. This would require subroutines to rapidly recalculate alignment scores along changed borders between images. After performing the above generalized perturbations on a single image, we would compute the correlations of pairs of images to determine how the changed image correlates with adjacent images, and use that to measure the fitness of the image. If these routines are quick then the algorithm could address the very real concern that stack movement and other noise may produce skewed, scaled, or rotated images.

1.5 Nonlinear Programming

1.5.1 Cost Function

In this section we represent each image as a matrix $[u^{ij}]_{L_y \times L_x}$ of integers (corresponding to pixel grey scale values). There are target offset overlaps (S_x and S_y) which provide overlap regions to help align the images. Assuming that the target overlaps are achieved, the top left corner of each image define the reference grid (Figure 1.12). In practice however, errors during image acquisition modify these overlap regions. This can be quantified by associating an offset error $\vec{\Delta}^{pq} = (\Delta_x^{pq}, \Delta_y^{pq})$ with every image u^{pq}. The offset error is defined relative to the reference grid. The left overlap region between image u^{ij} and $u^{i,j-1}$ has width $\left(S_x - (\Delta_x^{ij} - \Delta_x^{i,j-1})\right)$ and height $\left(L_y - (\Delta_y^{ij} - \Delta_y^{i,j-1})\right)$. Similarly, the top overlap region between image u^{ij} and $u^{i-1,j}$ has width $\left(L_x - (\Delta_x^{ij} - \Delta_x^{i-1,j})\right)$ and height $\left(S_y - (\Delta_y^{ij} - \Delta_y^{i-1,j})\right)$.

It is also assumed that we have an upper bound on the offset errors, i.e., $|\Delta_x^{ij}| \le \beta_x$ and $|\Delta_y^{ij}| \le \beta_y$ where β_x and β_y are constants, or, an upper bound on the relative offset errors between every pair of neighbouring images. Let

$$\eta_x = \Delta_x^{i,j-1} - \Delta_x^{ij}, \tag{1.4a}$$

$$\eta_y = \Delta_y^{i,j-1} - \Delta_y^{ij}, \tag{1.4b}$$

$$\xi_x = \Delta_x^{i-1,j} - \Delta_x^{ij}, \tag{1.4c}$$

$$\xi_y = \Delta_y^{i-1,j} - \Delta_y^{ij}, \tag{1.4d}$$

$$\pi$$
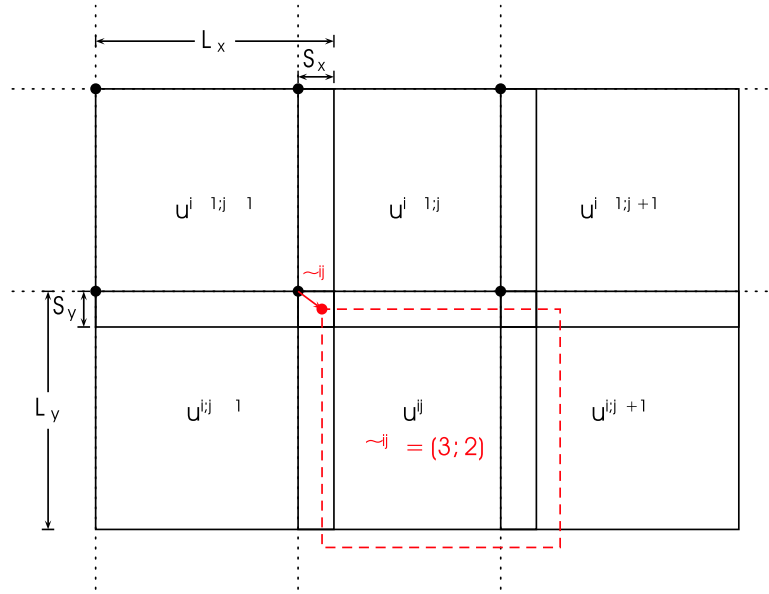
Figure 1.12: Image layout using reference grid approach. Solid lines show images at target offset overlaps (upper left corners are shown as $\bullet$), dashed box shows image u^{ij} displaced by offset error $\vec{\Delta}^{ij}$.

and thus

$$|\eta_x| \leq \beta_x, \quad |\eta_y| \leq \beta_y,$$
$$|\xi_x| \leq \beta_x, \quad |\xi_y| \leq \beta_y.$$

It is unclear from the problem specification whether the β_x and β_y bound the relative offset errors or the actual offset errors, but the following discussion and code handles both cases.

We calculate the cost function by summing the contributions from the top and left overlap regions for all images u^{ij}. The cost function contribution due to u^{ij} is

$$d^{ij} = \frac{\displaystyle\sum\sum_{\text{left overlap pixels}} \left|u^{ij} - u^{i,j-1}\right|^2}{\text{\# of left overlap pixels}} + \frac{\displaystyle\sum\sum_{\text{top overlap pixels}} \left|u^{ij} - u^{i-1,j}\right|^2}{\text{\# of top overlap pixels}},$$

$$= \text{cost\_contrib\_left}(u^{ij}, u^{i,j-1}, \eta_x, \eta_y)$$
$$+ \text{cost\_contrib\_top}(u^{ij}, u^{i-1,j}, \xi_x, \xi_y), \tag{1.5}$$

where the first term is 0 if $j = 1$ and the second term is 0 if $i = 1$. Note that d^{ij} depends explicitly on $\vec{\Delta}^{ij}$, $\vec{\Delta}^{i-1,j}$ and $\vec{\Delta}^{i,j-1}$. The total cost function is then calculated as

$$d = \sum_i \sum_j d^{ij}. \tag{1.6}$$

$$\boxed{\pi}$$

1.5.2 Minimization

At this point, we try minimizing the above cost function using Matlab's *fminsearch* over the space of $\vec{\Delta}$'s. However, we encounter an unexpected hurdle: the stitching of images is not a discrete problem whereas our cost function calculation assumes that the offset errors were integers. One may question the usefulness of resolving the offset errors to a "sub-pixel" level when feature sizes are $\mathcal{O}(10)$ pixels, however, finding a free black-box optimization code which generates and *uses* only integer values is difficult. We propose several fixes:

- round $\vec{\Delta}^{ij}$ when calculating the cost function.

- interpolate the image data on the sub-pixel level.

- interpolate our cost function.

As expected, rounding does not produce a convergent result. Interpolating the image is impractical for larger problems where realistically, a caching approach as discussed in Section 1.5.2 is required. We choose to implement interpolation of the cost function.

Cost Function Interpolation

Let cost_contrib_left() be a function that calculates the cost function contribution in the left overlap region between neighbouring images u^{ij} and $u^{i,j-1}$. Assuming that η_x and η_y are integers, the cost contribution is calculated as

$$\text{cost\_contrib\_left}(u^{ij},\ u^{i,j-1},\ \eta_x, \eta_y).$$

To accommodate fractional offset errors, consider

$$\eta_x = \lfloor \eta_x \rfloor + \delta_x, \tag{1.7a}$$
$$\eta_y = \lfloor \eta_y \rfloor + \delta_y, \tag{1.7b}$$

where $\vec{\delta} \in [0,1) \times [0,1)$, $\lfloor \cdot \rfloor = \text{floor}(\cdot)$ and $\lceil \cdot \rceil = \text{ceil}(\cdot)$.

We now build the lookup table for the cost function contribution at the 4 nearest neighbouring integer values of $\vec{\eta}$ by calculating

$$c_{\lfloor x \rfloor \lceil y \rceil} = \text{cost\_contrib\_left}(u^{ij},\ u^{i,j-1},\ \lfloor \eta_x \rfloor, \lceil \eta_y \rceil), \tag{1.8a}$$
$$c_{\lceil x \rceil \lceil y \rceil} = \text{cost\_contrib\_left}(u^{ij},\ u^{i,j-1},\ \lceil \eta_x \rceil, \lceil \eta_y \rceil), \tag{1.8b}$$
$$c_{\lfloor x \rfloor \lfloor y \rfloor} = \text{cost\_contrib\_left}(u^{ij},\ u^{i,j-1},\ \lfloor \eta_x \rfloor, \lfloor \eta_y \rfloor), \tag{1.8c}$$
$$c_{\lceil x \rceil \lfloor y \rfloor} = \text{cost\_contrib\_left}(u^{ij},\ u^{i,j-1},\ \lceil \eta_x \rceil, \lfloor \eta_y \rfloor). \tag{1.8d}$$

Our estimate, c of the cost function contribution is a weighted average of (1.8). Specifically, we use bilinear interpolation, by first calculating

$$c_{\lceil x \rceil} = (1 - \delta_y)c_{\lceil x \rceil \lfloor y \rfloor} + \delta_y c_{\lceil x \rceil \lceil y \rceil},$$
$$c_{\lfloor x \rfloor} = (1 - \delta_y)c_{\lfloor x \rfloor \lfloor y \rfloor} + \delta_y c_{\lfloor x \rfloor \lceil y \rceil},$$

and then calculating

$$c = (1 - \delta_x)c_{\lfloor x \rfloor} + \delta_x c_{\lceil x \rceil}.$$

A similar discussion applies for cost_contrib_up() using ξ_x and ξ_y.

π
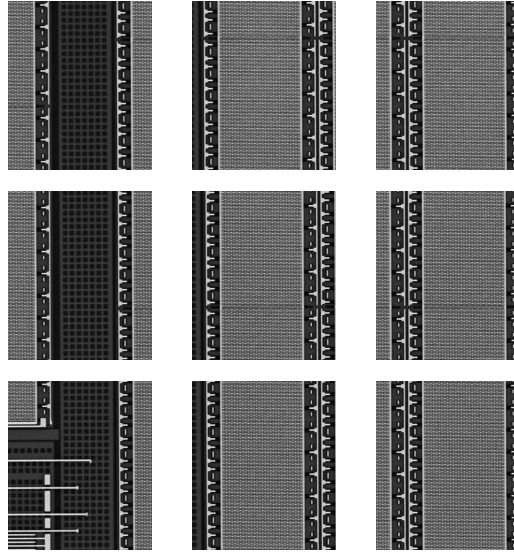
Figure 1.13: Test Case 1.

Cost Function Space

The set of images above was used as Test Case 1. The original image provided by Semiconductor Insights of a chip was artificially broken into nine 700×600 pieces with overlap regions of 100 pixels and random offset errors in the range of -6 to 6 in both the horizontal and vertical direction (Figure 1.13).

Unfortunately, given a random initial guess for the $\vec{\Delta}$'s, *fminsearch* could not find the global minimizer. We present two figures of the function space in question. Given two sets of offset errors say $\vec{\Delta}$ and $\vec{\Delta}_{\text{exact}}$, we define the radius r as $\|\vec{\Delta} - \vec{\Delta}_{\text{exact}}\|$. Figure 1.14 shows the cost as a function of radius for fixed directions $\hat{\theta}$ from the exact solution, i.e., define

$$\hat{\theta} = \frac{\vec{\Delta} - \vec{\Delta}_{exact}}{\|\vec{\Delta} - \vec{\Delta}_{\text{exact}}\|} \tag{1.9}$$

and plot cost as a function of r. Notice that none of curves are monotonic, in fact, a plot of cost as a function of radius for 1000 $\hat{\theta}$'s reveal that only one or two curves were monotone. The cost function space has a lot of local mimima; this presents problems for most optimization routines.

We also present the plots obtained by keeping most of the offset variables constant at the known exact errors while varying pairs of offset variables (Figure 1.15). One can clearly see the existence of many local minimas. Close to the exact solution however, the function spaces are very smooth (Figure 1.16). This provides some hope that an optimization scheme can indeed be used, provided a good initial guess is given.

Iterative Scheme

To provide a better initial guess for the optimization routine, we adopt an iterative approach to the problem. This involves stitching one additional image each iteration.

π
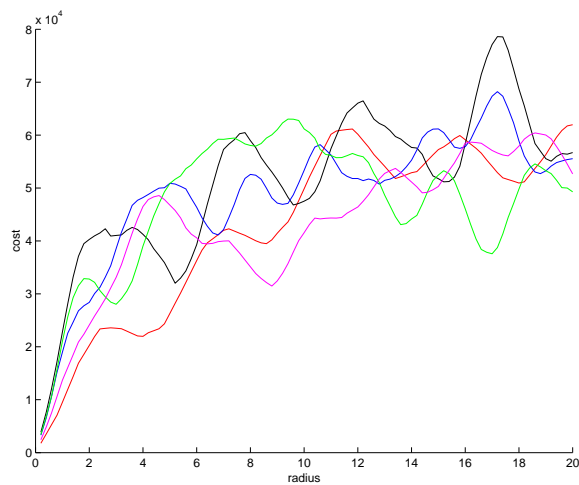
Figure 1.14: Cost as a function of radius for various directions in offset error space.
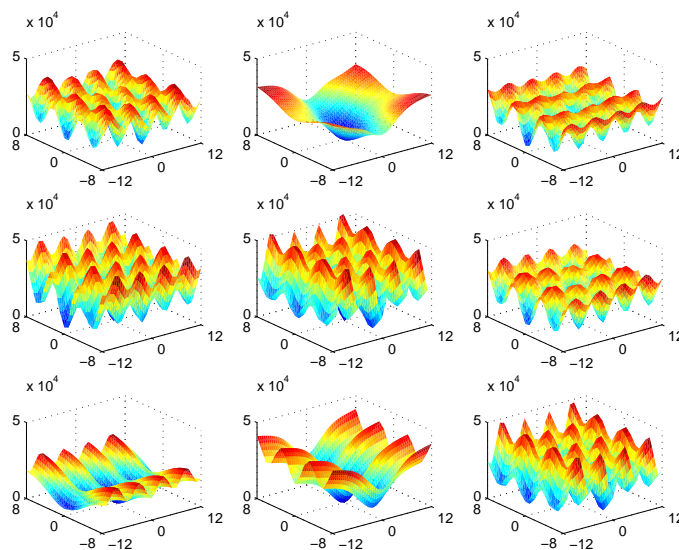
Figure 1.15: Multiple plots of the cost function space for arbitrary pairs of offset variables. (Other variables held at constant, correct minimizing values.)
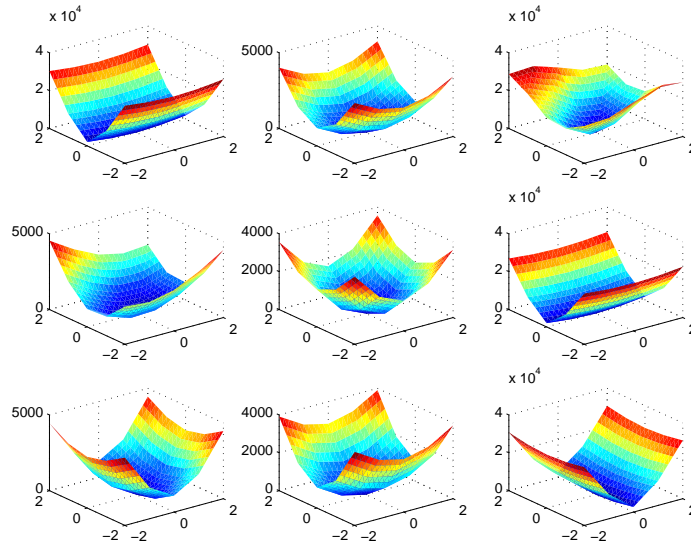
Figure 1.16: Multiple plots of the cost function space close to the global minimum. Also note the effect of bilinear interpolation.

We begin by choosing a particular ordering of the images, $(u^1, u^2, \ldots, u^K)$, $K = M * N$ where each u^k corresponds to a unique u^{ij} and u^k shares at least one common edge with $(u^1, \ldots, u^{k-1})$. A simple ordering is

$$
\begin{matrix}
u^1 & u^2 & u^3 \\
u^4 & u^5 & u^6
\end{matrix}
$$

More effective orderings are considered in Section 1.2.

The following pseudo-algorithm then illustrates the process of sequentially stitching one additional image per iteration:

1. Assume that $\vec{\Delta}^1 \equiv \vec{0}$ and set $k = 2$.

2. Stitch $u^1, \ldots, u^k$ assuming that the known $\vec{\Delta}^1, \ldots, \vec{\Delta}^{k-1}$ are fixed and only $\vec{\Delta}^k$ is allowed to change, i.e., find the optimal $\vec{\Delta}^k$ which stiches the images. (Note, we use a random initial guess for $\vec{\Delta}^k$)

3. Stitch $u^1, \ldots, u^k$ varying $\vec{\Delta}^1, \ldots, \vec{\Delta}^k$, using the values from step 2 as the initial guess.

4. Set $k = k + 1$ and return to step 2.

We acheived correct convergent results in approximately 90% of the trials, however, the convergence was very slow. On a 3×3 set of images (i.e., 14 offset variables), the code took over eight hours to complete on a 1.2 GHz Athlon MP.

π

Pre-Conditioner

In retrospect, it is not necessary and quite inefficient to take a random initial guess for $\vec{\Delta}^k$ in step 2, Section 1.5.2. Figure 1.16 shows that an initial guess within two pixels is likely sufficient for convergence.

A better approach is to do an exhaustive search over a finite set of values for $\vec{\Delta}^k$, and use the $\vec{\Delta}^k$ providing the smallest contribution to the cost function as the initial guess in step 3 of Section 1.5.2. A natural space of relative offset errors is $\{\vec{\Delta} \in \mathbb{Z}, |\vec{\eta}| \leq |\vec{\beta}|\}$ (where $\vec{\beta} = (\beta_x, \beta_y)$) if we have bounds on the offset errors. If we instead have bounds on the relative error between images, we can do something similar: $\{\vec{\eta} \in \mathbb{Z}, |\vec{\eta}| \leq |\vec{\beta}|\}$ where the inequalities are between u^{ij} and $u^{i-1,j}$ unless $i \equiv 1$, in which case the inequalities are between u^{ij} and $u^{i,j-1}$. An important point to realize is that although our formulation uses a fixed reference grid, the contribution to the cost function between neighbouring images only depends on the relative difference in offset errors between them.

Upon implementation of the pre-conditioner, our code became more robust, with a noticable increase in speed. We were actually quite surprised at how robust the code was. It was unclear if our results were attributed to the artificial construction of test images (cutting a single picture into multiple parts), so we tested our code on a variety of altered images, for example:

- After cutting up the image into smaller test images, each test image was saved as a jpeg with a quality rating of 20. Quality ratings for jpeg range from 0 to 100 where a higher number means less image degradation due to compression.

- Each test image was also subjected to a filter which added gaussian bluring (within a radius of 5 pixels).

- Each test image was subjected to a noise filter. This noise filter "jittered" the image by randomly perturbing some pixels a small distance[1]. This filter actually made the images look quite horrible and hard to optimize.

- The output from the pre-conditioner was perturbed by $|\epsilon| < 1$ to simulate the case when the pre-conditioner does not find the correct values.

Our code still found the global minimum for all the test cases, giving us confidence that it could work when real images are used. Lots of nasty attributes such as lens abberation, difference in contrasts, warping of images would be noticed in real images.

Caching

A co-author suggested that caching the cost function contribution for various $\vec{\Delta}$'s would provide significant speedup. This should translate to significant computational savings during our bilinear interpolation. Another important motivation is the difficulty in storing a 8×14 grid of 8k $\times$ 8k images in memory (7 Gb at one byte per pixel). There are two possible stages to implement caching: pre-caching the cost function contributions before the optimization step or caching during the optimization step (i.e., cache the cost contribution for each new integer $\vec{\Delta}$).

---

[1] "pick" filter in GIMP (www.gimp.org)

The former approach was taken because there is less book-keeping involved. It is likely though that caching during the optimization step will provide further speedup.

For every image u^{ij}, associate two matrices D_L and D_U which correspond to cached contributions from the left and top image respectively. As mentioned in Section 1.5.2, although our formulation uses the fixed reference grid, the contribution to the cost function from neighbouring images depends only on the relative difference in offset errors. We utilize the experimental constraint that the offset errors are bounded above by $\vec{\beta}$. Thus, D_L and D_U are matrices with $(2\beta_y+1)\times(2\beta_x+1)$ elements. Note, in practice, we cache a larger matrix to account for roundoff error.

For a 6×5 grid of 700×600 images (Test Case 2) with $\beta_x = 10$ and $\beta_y = 8$, caching took 7 minutes and optimization took 14 minutes. Note, our non-caching code for this problem was halted after 4 days. A rough estimate for caching the 14×8 problem described above is approximately nine hours. Note that the caching algorithm scales linearly with the number of images and linearly with the size of the overlap region. If the computation time is too long, the caching algorithm is easily parallelizable for a cluster type network.

1.5.3   Results

The optimization code encounters no difficulties fitting Test Case 2. A rough extrapolation suggests that the 14×8 problem with 240 variables would take about one day.

Matlab codes and more details on this work are available online[4]. The authors of this Section thank Jim Verner for proof-reading and constructive criticism and Jason Nielsen for help with extrapolation techniques.

1.5.4   Future Work

- One should *actually try* the fixed reference grid approach of sequentially fitting images. Global optimization may not be necessary.

- Combine the iterative scheme in Section 1.5.2 with a more effective ordering considered in Section 1.2.

- Employ a strategy whereby offset variables from images that have enough surrounding neighbours are removed from the optimization process. One could use this process to bound the number of parameters in the optimization.

- Throughout this discussion, we have assumed that the images have been pre-processed, e.g., the images have been rotated correctly. The pre-processing stage can be incorporated into the cost function.

- Recode our algorithm in a more efficient environment such as C or Fortran.

- The cost function is easily modified to fit two different layers of images. A first step would be to detect via-ups and via-downs in each image/layer. (A good segmentation algorithm based on levelsets or wavelets would probably do the job.) One could then calculate the cost function for each layer, sum them together to create the new cost function, and add

$$\pi$$

a penalty term to the new cost function if the vias connect to wires either above or below. The parameter space for two layers of 14×8 images can be handled by a desktop computer in a reasonable amount of time.

- This problem can be reformulated to incorporate the power of parallel processing. For example, the caching routine is easily parallelized.

- For large images, it is not necessary to read in the whole image when optimizing for the offset errors. In fact, less than 5% of the image is required assuming an overlap width of 100 on $8k \times 8k$ pictures. Incorporating this idea will relax memory constraints.

- Custom tailor a minimizer that pays attention to the pixel scale. It can use sub-pixel information if necessary but ideally should not because sub-pixel information is not important. A first attempt would probably involve a simple algorithm:

  1. Calculate an approximation to the Jacobian (finite differences).
  2. Calculate an approximation to the Hessian (finite differences).
  3. Calculate a search direction based on the Jacobian and Hessian (simple line search).
  4. Choose an appropriate step in the search direction - i.e., integer values.
  5. Evaluating the Hessian and Jacobian is probably going to be quite expensive, so it's probably better to "Freeze" the Jacobian and Hessian by returning to step 3 and recalculating a search direction.

- If the images were collected in a hexagonal pattern (i.e., acquisition of images for every other row is offset by $\frac{L_x}{2}$) then one would expect a better alignment due to increased connectivity.

- A more radical approach would be to incorporate (local) stitching into the circuit generation code. The code would follow traces, and whenever it encountered the edge of an image, it would simply do a local stitching with the neighbouring image and continue the trace onto that image.

1.6 Conclusions

The different solutions presented here require different levels of effort to be implemented, with corresponding different potential payoff. The graph theory solutions of Section 1.2 provide improvements on the patterns of sequential stitching currently used. The least squares codes of Section 1.3 fairly rapidly compute stitching solutions based on the offset data. The simulated annealing and nonlinear programming codes of Sections 1.4–1.5 are more flexible, because they can find optimal stitchings based on the image pixel data, but they are computationally much more demanding.

π

Bibliography

[1] James W. Demmel, *Applied Numerical Linear Algebra*, SIAM, 1997, Section 6.3.3.

[2] `ftp://ftp.cc.tut.fi/pub/math/piche/stitching/`

[3] `http://mathstat.carleton.ca/~brett/Research/Code/#ipsw02`

[4] `http://www.math.sfu.ca/~bwo/ipsw`

Chapter 2

Price Pseudo-Variance, Pseudo-Covariance, Pseudo-Volatility, and Pseudo-Correlation Swaps — In Analytical Closed-Forms

Andrei Badescu[1], Hammouda Ben Mekki[2], Asrat Fikre Gashaw[3], Yuanyuan Hua[4], Marat Molyboga[5], Tereza Neocleous[6], Yuri Petratchenko[5].

Report prepared by Raymond K. Cheng[7]
with assistance from Stéphan Lawi[7] and Anatoliy Swishchuck[3].

2.1 Introduction

This problem was brought to us by Ritchie (Yeqi) He of RBC Financial Group. It concerns the pricing of swaps involving the so-called pseudo-statistics, namely the pseudo-variance, -covariance, -volatility, and -correlation. These products provide an easy way for investors to gain exposure to the future level of volatility. Swaps, in this context, are forward contracts in which the underlying is an interest rate S_t. They are also path-dependent because the payoffs depend on the trajectory of S_t on some time period $[T_s, T_e] \subseteq [0, T]$, where T is the maturity date. The prefix *pseudo* is used to indicate that samples of the underlying are taken at discrete times $t_0, t_1, ..., t_n \in [T_s, T_e]$. In the usual complete market framework, the unique prices of these swaps can be computed as the mathematical expectation of the discounted payoffs under the measure in which the discounted rate process is a martingale. In this report, we present analytic

---

[1]University of Western Ontario
[2]University of Calgary
[3]York University
[4]McGill University
[5]Moscow State University
[6]University of Illinois
[7]University of Toronto

formulas for these expectations for the pseudo-variance and pseudo-volatility swaps, as requested by the problem statement. Also, we use Monte-Carlo simulation to experiment with a stochastic volatility model.

2.2   Definitions

Here we briefly describe the problem presented to us. The bibliography contains references on variance and volatility swaps. The market we consider consists of the strictly positive underlying rates $S_t^{(1)}$ and $S_t^{(2)}$ satisfying the stochastic differential equation (SDE)

$$\frac{dS_t^{(i)}}{S_t^{(i)}} = \mu_t^{(i)}\,dt + \sigma_t^{(i)}\,dW_t^i, \qquad t > 0, \qquad i = 1,\,2$$

and a numeraire N_t which is a zero-coupon bond $\mathrm{df}(t, T)$. Here $dW_t^{(i)}$ are standard Wiener processes (zero mean, unit variance per unit time) with correlation $\rho_t dt$. In the cases of the variance and volatility swaps, we drop the superscripts so for example we write $S_t = S_t^{(1)}$ and $\mu_t = \mu_t^{(1)}$, etc. We assume that the market is complete, so there exists a unique martingale measure Q with respect to N_t. First, let us define the following continuously realized (measured) statistics over an observation period $[T_s, T_e]$:

$$\Sigma_{(S)}^2(T_s, T_e) = \frac{1}{T_e - T_s}\int_{T_s}^{T_e}\sigma_\tau^2 d\tau, \qquad\qquad \text{(realized volatility-square)}$$

$$\Sigma_{(S^{(1)}, S^{(2)})}^2(T_s, T_e) = \frac{1}{T_e - T_s}\int_{T_s}^{T_e}\sigma_\tau^{(1)}\sigma_\tau^{(2)}\rho_\tau d\tau, \qquad\qquad \text{(realized volatility-cross)}$$

$$\sigma_{(S)}(T_s, T_e) = \sqrt{\frac{1}{T_e - T_s}\int_{T_s}^{T_e}\sigma_\tau^2 d\tau}, \qquad\qquad \text{(realized volatility)}$$

$$\rho_{(S^{(1)}, S^{(2)})}(T_s, T_e) = \frac{\int_{T_s}^{T_e}\sigma_\tau^{(1)}\sigma_\tau^{(2)}\rho_\tau d\tau}{\sqrt{\int_{T_s}^{T_e}\sigma_\tau^{(1)^2} d\tau}\sqrt{\int_{T_s}^{T_e}\sigma_\tau^{(2)^2} d\tau}}. \qquad\qquad \text{(realized correlation)}$$

Note that when $S_t^{(1)} \equiv S_t^{(2)}$, the correlation is equal to one and the volatility-cross coincides with the volatility-square. Here we only consider the simpler case where we approximate the above quantities by ones that are discretely sampled. Thus, let $T_s = t_0 < t_1 < ... < t_n = T_e$ be the sampling dates, and we define the log-return[1] for the underlying rate S as

$$X_i^{(k)} = \log\left(\frac{S_{t_i}^{(k)}}{S_{t_{i-1}}^{(k)}}\right), \qquad k = 1, 2, \qquad i = 1, 2, ..., n$$

and also denote the arithmetic mean by

$$\bar{X}_n^{(k)} = \frac{1}{n}\sum_{i=1}^n X_i^{(k)}, \qquad k = 1, 2.$$

---

[1]Here log denotes the natural logarithm.

π

Now, we can define the following realized pseudo-statistics:

$$\hat{\Sigma}^2_{(S)}(n; T_S, T_e) = \frac{n}{T_e - T_s}\left(\frac{1}{n-1}\sum_{i=1}^{n}(X_i - \bar{X}_n)^2\right), \qquad \text{(realized pseudo-volatility-square)}$$

$$\hat{\Sigma}^2_{(S^{(1)}, S^{(2)})}(n; T_s, T_e) = \frac{n}{T_e - T_s}\left(\frac{1}{n-1}\sum_{i=1}^{n}\prod_{k=1}^{2}(X_i^{(k)} - \bar{X}_n^{(k)})\right), \quad \text{(realized pseudo-volatility-cross)}$$

$$\hat{\sigma}_{(S)}(n; T_s, T_e) = \sqrt{\frac{n}{T_e - T_s}\left(\frac{1}{n-1}\sum_{i=1}^{n}\left(X_i - \bar{X}_n\right)^2\right)}, \qquad \text{(realized pseudo-volatility)}$$

$$\hat{\rho}_{(S^{(1)}, S^{(2)})}(n; T_s, T_e) = \frac{\sum_{i=1}^{n}\prod_{k=1}^{2}(X_i^{(k)} - \bar{X}_n^{(k)})}{\prod_{k=1}^{2}\sqrt{\sum_{i=1}^{n}\left(X_i^{(k)} - \bar{X}_n^{(k)}\right)^2}}. \qquad \text{(realized pseudo-correlation)}$$

Based on these pseudo-statistics, we can define the swaps, which are really forward contracts, by their payoffs at maturity date $T \geq T_e$:

$$V_{\text{var}}(T) = \alpha_{\text{var}} \cdot I \cdot \left[\hat{\Sigma}^2_{(S)}(n; T_S, T_e) - \Sigma^2_K\right], \qquad \text{(pseudo-variance swap)}$$

$$V_{\text{cov}}(T) = \alpha_{\text{var}} \cdot I \cdot \left[\hat{\Sigma}^2_{(S^{(1)}, S^{(2)})}(n; T_s, T_e) - \Sigma^2_K\right], \qquad \text{(pseudo-covariance swap)}$$

$$V_{\text{vol}}(T) = \alpha_{\text{vol}} \cdot I \cdot \left[\hat{\sigma}_{(S)}(n; T_s, T_e) - \sigma_K\right], \qquad \text{(pseudo-volatility swap)}$$

$$V_{\text{corr}}(T) = \alpha_{\text{corr}} \cdot I \cdot \left[\hat{\rho}_{(S^{(1)}, S^{(2)})}(n; T_s, T_e) - \rho_K\right]. \qquad \text{(pseudo-correlation swap)}$$

In the above, α_i are the converting parameters, $I = \pm 1$ is a long-short index, and Σ^2_K, σ_K and ρ_K are the strikes.

The problem is to price the above swaps. The general pricing formula is the mathematical expectation of the N_t-discounted claim under the unique martingale measure Q:

$$\begin{aligned} V_i(0) &= \text{E}^Q\left[\frac{N_0}{N_T}V_i(T)\right] \\ &= \text{E}^Q\left[\text{df}(0, T)V_i(T)\right] \\ &= \text{df}(0, T)\,\text{E}^Q\left[V_i(T)\right], \qquad i = \text{var, cov, vol, and corr.} \end{aligned}$$

Note that the payoffs are path-dependent. Our goal is to find analytic expressions for the above pricing problems.

2.3 Price Pseudo-Variance & Pseudo-Covariance Swaps

To find the price of a pseudo-variance swap, it is necessary to compute the expected value of $\hat{\Sigma}^2_{(S)}(n; T_S, T_e)$. First, notice that each X_i is independently and normally distributed with mean

$$\alpha_i = \int_{t_{i-1}}^{t_i}\left(\mu_t - \frac{\sigma_t^2}{2}\right)dt,$$

π

and variance

$$\beta_i^2 = \int_{t_{i-1}}^{t_i} \sigma_t^2 dt.$$

Now we compute

$$\mathrm{E}^Q \left[\hat{\Sigma}_{(S)}^2(n; T_S, T_e) \right] = \mathrm{E}^Q \left[\frac{n}{T_e - T_s} \left(\frac{1}{n-1} \sum_{i=1}^n (X_i - \bar{X}_n)^2 \right) \right]$$

$$= \frac{n}{T_e - T_s} \cdot \frac{1}{n-1} \mathrm{E}^Q \left[\sum_{i=1}^n (X_i - \bar{X}_n)^2 \right]$$

$$= \frac{n}{T_e - T_s} \cdot \frac{1}{n-1} \mathrm{E}^Q \left[\sum_{i=1}^n \left((\alpha_i + \beta_i Y_i) - \frac{1}{n} \sum_{j=1}^n (\alpha_j + \beta_j Y_j) \right)^2 \right],$$

where Y_i is a normal $N(0,1)$ random variable. Let $\bar{\alpha} := \sum_{k=1}^n \alpha_i$ be the mean. After some calculations, the above expression can be written

$$\mathrm{E}^Q \left[\hat{\Sigma}_{(S)}^2(n; T_S, T_e) \right] = \frac{n}{T_e - T_s} \cdot \frac{1}{n-1} \left(\sum_{i=1}^n (\alpha_i - \bar{\alpha})^2 + \mathrm{E}^Q \left[\sum_{i=1}^n \left(\beta_i Y_i - \frac{1}{n} \sum_{j=1}^n \beta_j Y_j \right)^2 \right] \right)$$

$$= \frac{n}{T_e - T_s} \cdot \frac{1}{n-1} \left(\sum_{i=1}^n (\alpha_i - \bar{\alpha})^2 + \frac{n-1}{n} \sum_{i=1}^n \beta_i^2 \right)$$

$$= \frac{n}{T_e - T_s} \cdot \frac{1}{n-1} \sum_{i=1}^n (\alpha_i - \bar{\alpha})^2 + \frac{1}{T_e - T_s} \sum_{i=1}^n \beta_i^2$$

$$= \frac{n}{T_e - T_s} \cdot \frac{1}{n-1} \sum_{i=1}^n (\alpha_i - \bar{\alpha})^2 + \frac{1}{T_e - T_s} \int_{T_s}^{T_e} \sigma_\tau^2 d\tau$$

$$= \frac{n}{T_e - T_s} \cdot \frac{1}{n-1} \sum_{i=1}^n (\alpha_i - \bar{\alpha})^2 + \Sigma_{(S)}^2(T_s, T_e).$$

We point out that a variance swap can be replicated by a portfolio of options, forwards and zero-coupon bonds [2].

As pointed out in the problem statement, the pricing of a pseudo-covariance swap can be reduced to the pricing of a pseudo-variance swap. More specifically, it was argued that the price of a covariance swap can be written

$$V_{\mathrm{cov}}(0) = \frac{1}{4} \alpha_{cov} \cdot I \cdot \mathrm{df}(0, T) \cdot \left(\mathrm{E}^Q \left[\hat{\Sigma}_{(S^{(1)} \cdot S^{(2)})}^2(n; T_S, T_e) \right] - \mathrm{E}^Q \left[\hat{\Sigma}_{(S^{(1)}/S^{(2)})}^2(n; T_S, T_e) \right] - 4 \Sigma_K^2 \right).$$

Hence we do not need to price the pseudo-covariance swap directly.

2.4   Price Pseudo-Volatility Swaps

To price the pseudo-volatility swap of strike σ_K, we need to compute the expected value of the realized-pseudo volatility $\hat{\sigma}_{(S)}(n; T_s, T_e)$. To get a second-order approximation, we use a Taylor

expansion of the square root function about some $\hat{\Sigma}_0$ close to the expected swap value:

$$
\mathrm{E}^Q\left[\hat{\sigma}_{(S)}(n; T_s, T_e)\right] = \mathrm{E}^Q\left[\sqrt{\hat{\Sigma}^2_{(S)}(n; T_S, T_e)}\right]
$$

$$
\approx \mathrm{E}^Q\left[\frac{\hat{\Sigma}^2_{(S)}(n; T_S, T_e) + \hat{\Sigma}_0}{2\sqrt{\hat{\Sigma}_0}} - \frac{(\hat{\Sigma}^2_{(S)}(n; T_S, T_e) - \hat{\Sigma}_0)^2}{8(\hat{\Sigma}_0)^{\frac{3}{2}}}\right]
$$

$$
= \frac{\mathrm{E}^Q\left[\hat{\Sigma}^2_{(S)}(n; T_S, T_e)\right] + \hat{\Sigma}_0}{2\sqrt{\hat{\Sigma}_0}} - \frac{\mathrm{E}^Q\left[\hat{\Sigma}^2_{(S)}(n; T_S, T_e) - \hat{\Sigma}_0\right]^2}{8(\hat{\Sigma}_0)^{\frac{3}{2}}}
$$

$$
= \frac{\mathrm{E}^Q\left[\hat{\Sigma}^2_{(S)}(n; T_S, T_e)\right] + \hat{\Sigma}_0}{2\sqrt{\hat{\Sigma}_0}} - \frac{\mathrm{Var}^Q\left[\hat{\Sigma}^2_{(S)}(n; T_S, T_e)\right] + \left(\mathrm{E}^Q\left[\hat{\Sigma}^2_{(S)}(n; T_S, T_e)\right] - \hat{\Sigma}_0\right)^2}{8(\hat{\Sigma}_0)^{\frac{3}{2}}}.
$$

Note that the expected value $\mathrm{E}^Q\left[\hat{\Sigma}^2_{(S)}(n; T_S, T_e)\right]$ is already computed in the previous section, so here we compute the variance:

$$
\mathrm{Var}^Q\left[\hat{\Sigma}^2_{(S)}(n; T_S, T_e)\right] = \mathrm{Var}^Q\left[\frac{n}{T_e - T_s}\left(\frac{1}{n-1}\sum_{i=1}^{n}(X_i - \bar{X}_n)^2\right)\right]
$$

$$
= \frac{1}{(T_e - T_s)^2} \cdot \frac{n^2}{(n-1)^2} \mathrm{Var}^Q\left[\sum_{i=1}^{n}(X_i - \bar{X}_n)^2\right]
$$

$$
= \frac{1}{(T_e - T_s)^2} \cdot \frac{n^2}{(n-1)^2}\left(\mathrm{E}^Q\left[\sum_{i=1}^{n}(X_i - \bar{X})^2\right]^2 - \left(\mathrm{E}^Q\left[\sum_{i=1}^{n}(X_i - \bar{X})^2\right]\right)^2\right).
$$

After some calculations, the first term in the bracket above can be written

$$
\mathrm{E}^Q\left[\sum_{i=1}^{n}(X_i - \bar{X})^2\right]^2 = \mathrm{E}^Q\left[\sum_{i,j=1}^{n}X_i^2 X_j^2\right] - \frac{2}{n}\mathrm{E}^Q\left[\left(\sum_{k=1}^{n}X_k^2\right)^2\sum_{i=1}^{n}X_i^2\right] + \frac{1}{n^2}\mathrm{E}^Q\left[\sum_{i=1}^{n}X_i\right]^4.
$$

Now, let us denote the coefficients by

$$
A := \mathrm{E}^Q\left[\sum_{i,j=1}^{n}X_i^2 X_j^2\right],
$$

$$
B := \mathrm{E}^Q\left[\left(\sum_{k=1}^{n}X_k^2\right)^2\sum_{i=1}^{n}X_i^2\right], \text{ and}
$$

$$
C := \mathrm{E}^Q\left[\sum_{i=1}^{n}X_i\right]^4.
$$

Then our formula for the variance can be written

$$
\mathrm{Var}^Q\left[\sum_{i=1}^{n}(X_i - \bar{X})^2\right] = A - \frac{2}{n}B + \frac{1}{n^2}C - \left(\mathrm{E}^Q\left[\sum_{i=1}^{n}(X_i - \bar{X})^2\right]\right)^2.
$$

π

Note that the last term has already been calculated in the previous section, so it remains to find formulas for $A, B,$ and C. Indeed, we spent much of our time to do this and here are the results:

$$A = \sum_i (\alpha_i^4 + 6\alpha_i^2\beta_i^2 + 3\beta_i^4) + 2\sum_{i<j}(\alpha_i^2 + \beta_i^2)(\alpha_j^2 + \beta_j^2),$$

$$B = \sum_i (\alpha_i^4 + 6\alpha_i^2\beta_i^2 + 3\beta_i^4) + 2\sum_{i<j}(\alpha_i^2 + \beta_i^2)(\alpha_j^2 + \beta_j^2) + 2\sum_{i<j}(3\alpha_i\beta_i^2 + \alpha_i^3)\alpha_j$$

$$+ 2\sum_{i<j}\alpha_i(3\alpha_j\beta_j^2 + \alpha_j^3) + 2\sum_{i<j<k}\alpha_i\alpha_j(\alpha_k^2 + \beta_k^2),$$

$$C = \sum_i (\alpha_i^4 + 6\alpha_i^2\beta_i^2 + 3\beta_i^4) + 4\sum_{i<j}(\alpha_i^3 + 3\alpha_i\beta_i^2)\alpha_j + 6\sum_{i<j}(\alpha_i^2 + \beta_i^2)(\alpha_j^2 + \beta_j^2)$$

$$+ 12\sum_{i<j<k}(\alpha_i^2 + \beta_i^2)\alpha_j\alpha_k + 4\sum_{i<j}\alpha_i(\alpha_j^3 + 3\alpha_j\beta_j^2) + 12\sum_{i<j<l}\alpha_j(\alpha_j^2 + \beta_j^2)\alpha_l$$

$$+ 12\sum_{i<j<l}\alpha_i\alpha_j(\alpha_l^2 + \beta_l^2) + 24\sum_{i<j<k<l}\alpha_i\alpha_j\alpha_k\alpha_l.$$

In the above, we have again used the notation that each random variable X_i follows a normal distribution with mean α_i and variance β_i^2.

Further, it can be shown that

$$\mathrm{E}^Q\left[\hat{\Sigma}_{(S)}^2(n; T_S, T_e)\right] \to \Sigma_{(S)}^2 \qquad \text{and} \qquad \mathrm{Var}^Q\left[\hat{\Sigma}_{(S)}^2(n; T_S, T_e)\right] \to 0,$$

as $n \to \infty$. This means that $\hat{\Sigma}_{(S)}^2(n; T_S, T_e)$ is an asymptotically unbiased and consistent estimator of $\Sigma_{(S)}^2(T_s, T_e)$.

2.5   Numerical Simulation

Following the idea of [1], we consider for the numerical simulations, a stochastic volatility model of the CIR type:

$$d\sigma_t^2 = \kappa(\theta^2 - \sigma_t^2)\,dt + \gamma\sigma_t\,dW_t^{(2)}. \tag{2.1}$$

Similar to the previous section, once discretized in time, X_t becomes a stochastic process taking random values, X_i, normally distributed with mean α_i and variance β_i^2, at each time step:

$$X_i = \log\left(\frac{S_i}{S_{i-1}}\right) \sim N(\alpha_i, \beta_i^2).$$

The parameters are related to the X_t process through the following equations:

$$\alpha_i = \mathrm{E}^Q[X_i] = \int_{t_{i-1}}^{t_i} \mu_t dt - \frac{1}{2}\mathrm{E}^Q\left[\int_{t_{i-1}}^{t_i} \sigma_t^2 dt\right],$$

$$\beta_i^2 = \mathrm{Var}^Q[X_i] = \mathrm{E}^Q\left[\int_{t_{i-1}}^{t_i} \sigma_t^2 dt\right].$$

In order to evaluate the expectation and the variance of the X_t process, two different ways are being used. The first one is the application of the Monte-Carlo technique directly on the discretized σ_t^2 process, whose behaviour is given by (2.1). The second one is the application of the Monte-Carlo technique on the discretized process $X_i = \log(S_i/S_{i-1})$.

The advantage of using the CIR model is that it yields an exact solution for the expectation and the variance of the X_t process (cf.[1]). Hence, the numerical results can be easily verified.

All simulations are generated over 1000 time steps. For the results in Table 2.5, only 100 scenarios are considered, since the idea of the simulations is more to give an illustration and improve our understanding of the processes than to provide a numerical proof. The values of the parameters are set to $\kappa = 10$, $\theta^2 = 0.2$, $\sigma_0 = 0.2$ and $\gamma = 0.75$. Note that $\gamma = 0$ yields a deterministic volatility (see Figure 2.5).

	Expectation	Variance
Exact	0.1281	0.0019
MC σ_t^2	0.1287	0.0022
MC $(\log S_t)^2$	0.1280	0.0021

Table 2.1: Exact and numerical results for the expectation and variance of X_t.

Although the relatively low number of scenarios, the numerical results match the exact solutions within some reasonable error range.

2.6 Another Approach

Here we describe another approach in finding the expectation and variance. The idea is to work with the bond-discounted rate process S_t^* in the martingale measure Q. Let $(\Omega, \mathcal{F}, \mathcal{F}_t, P)$ be a probability space with $\mathcal{F}_t$ being filtration. We define the discounted S_t process by

$$S_t^* = \frac{S_t}{B_t},$$

where B_t is the deterministic bond price at time t satisfying $dB_t = r_t B_t \, dt$. It is known that the unique martingale measure Q for the discounted rate process S_t^* is given by the Radon-Nikodym derivative

$$\frac{dQ}{dP} = \exp\left(\int_0^T \frac{r_t - \mu_t}{\sigma_t} dW_t - \frac{1}{2}\int_0^T \left(\frac{r_t - \mu_t}{\sigma_t}\right)^2 dt\right). \tag{2.2}$$

Under this martingale measure, the driftless process S_t^* satisfies $dS_t^* = \sigma_t S_t^* \, dW_t^*$, where W_t^* given by

$$W_t^* = W_t - \int_0^t \frac{r_s - \mu_s}{\sigma_s} \, ds, \qquad 0 \le t \le T$$

is a standard Wiener process on the probability space $(\Omega, \mathcal{F}, \mathcal{F}_t, Q)$. We also have

$$d\log S_t^* = -\frac{\sigma_t^2}{2} \, dt + \sigma_t \, dW_t^*.$$

$$\pi$$
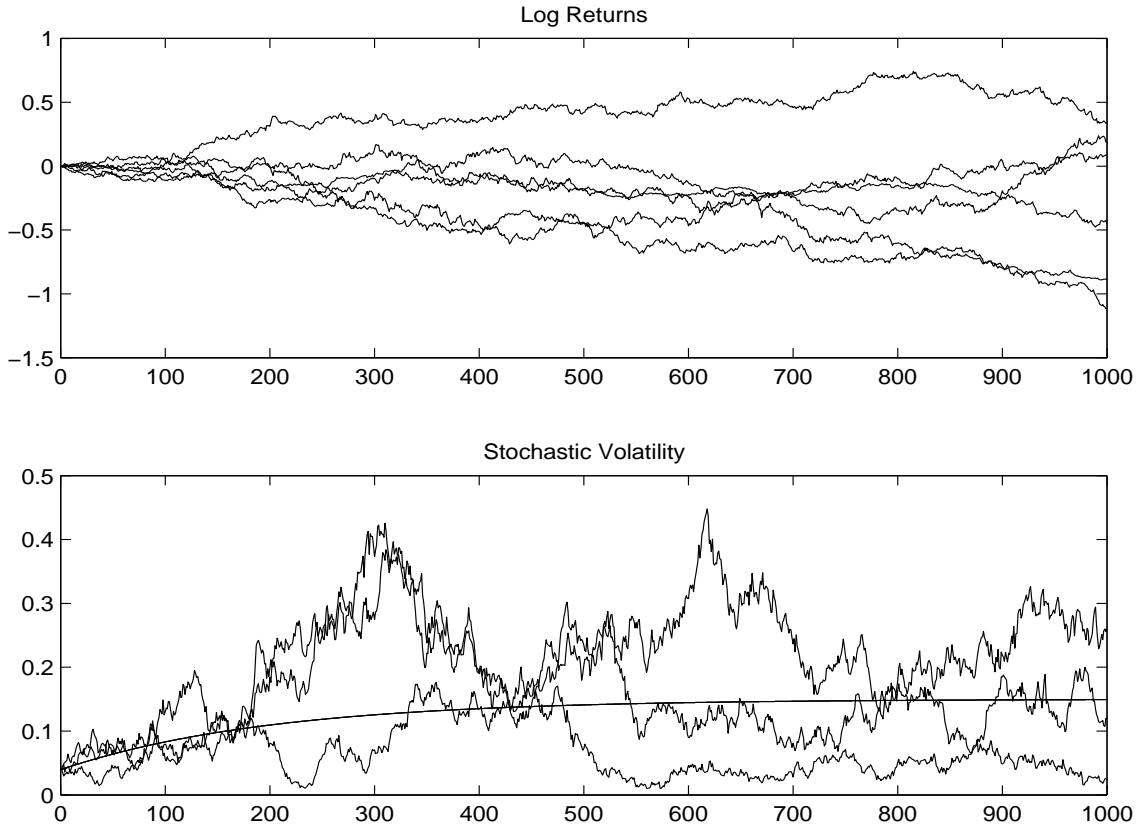
Figure 2.1: Monte-Carlo simulation, for the log-returns X_t and the stochastic volatility σ_t^2.

Then, the log-return for the discounted process S_t^* can be written

$$X_i^* := \log \frac{S_{t_i}^*}{S_{t_{i-1}}^*} = -\frac{1}{2}\int_{t_{i-1}}^{t_i} \sigma_t^2 \, dt + \int_{t_{i-1}}^{t_i} \sigma_t \, dW^*,$$

and it follows a normal distribution with mean

$$\alpha_i := \mathrm{E}^Q\left[X_i^*\right] = -\frac{1}{2}\int_{t_{i-1}}^{t_i} \sigma_t^2 \, dt$$

and variance

$$\beta_i^2 := \mathrm{Var}^Q\left[X_i^*\right] = \int_{t_{i-1}}^{t_i} \sigma_t^2 \, dt.$$

Thus, X_i^* has the representation $X_i^* = \alpha_i + \beta_i Y$, where Y is a standard normal $N(0,1)$. We note that

$$\mathrm{E}^Q\left[(X_i^*)^2\right] = \alpha_i^2 + \beta_i^2 \qquad \text{and} \qquad \mathrm{E}^Q\left[(X_i^*)^4\right] = \alpha_k^4 + 6\alpha_i^2\beta_i^2 + 3\beta_i^4.$$

Now, we compute the expected value of the pseudo-volatility-square, first neglecting the mean $\bar{X}_n^* := \frac{1}{n}\sum_{i=1}^{n} X_i^*$:

$$\mathrm{E}^Q\left[\hat{\Sigma}_{(S)}^2(n; T_S, T_e)\right] = \mathrm{E}^P\left[\hat{\Sigma}_{(S)}^2(n; T_S, T_e)M_T\right].$$

$\boxed{\pi}$

Here M_T is defined in the right-hand side of (2.2). If we use Taylor expansion for M_T, then we obtain the formula demonstrated below. We propose to calculate E^Q for the discounted process S_t^*, which simplifies some calculation in the case of S_t:

$$\mathrm{E}^Q\left[\hat{\Sigma}_{(S^*)}^2(n; T_S, T_e)\right] = \frac{n}{(T_e - T_s)} \cdot \frac{1}{(n-1)} \sum_{i=1}^{n} \mathrm{E}^Q\left[(X_i^*)^2\right]$$

$$= \frac{n}{(T_e - T_s)} \cdot \frac{1}{(n-1)} \sum_{i=1}^{n} \left(\alpha_i^2 + \beta_i^2\right)$$

$$= \frac{n}{(T_e - T_s)} \cdot \frac{1}{(n-1)} \sum_{i=1}^{n} \left(\frac{1}{4}\beta_i^4 + \beta_i^2\right).$$

If we do not neglect $\bar{X}_n^*$, then we obtain the more general expression

$$\mathrm{E}^Q\left[\hat{\Sigma}_{(S)}^2(n; T_S, T_e)\right] = \frac{n}{(T_e - T_s)} \cdot \frac{1}{(n-1)} \sum_{i=1}^{n} \left[\left(-\frac{1}{2}\beta_i^2 + \frac{1}{2n}\sum_{j=1}^{n}\beta_j^2\right)^2 + \frac{n-1}{n}\beta_i^2\right].$$

In the risk-neutral world, the process S_t satisfies the SDE

$$d\log S_t = \left(r_t - \frac{\sigma_t^2}{2}\right)dt + \sigma_t\, dW_t^*.$$

This means that under the measure Q the log-return X_i is normally distributed with mean

$$\alpha_i = \int_{t_{i-1}}^{t_i} \left(r_t - \frac{\sigma_t^2}{2}\right)dt$$

and variance

$$\beta_i^2 = \int_{t_{i-1}}^{t_i} \sigma_t^2\, dt.$$

In this case our formula becomes

$$\mathrm{E}^Q\left[\hat{\Sigma}_{(S)}^2(n; T_S, T_e)\right] = \frac{n}{(T_e - T_s)} \cdot \frac{1}{n-1} \sum_{i=1}^{n} \left[\left(\alpha_i - \frac{1}{n}\sum_{j=1}^{n}\alpha_j\right)^2 + \frac{n-1}{n}\beta_i^2\right],$$

which is the same as in Section 2.3.

To compute the variance, we need to find the the following expectation:

$$\mathrm{E}^Q\left[\left(\hat{\Sigma}_{(S)}^2(n; T_S, T_e)\right)^2\right] = \frac{1}{(T_e - T_s)^2} \cdot \frac{n^2}{(n-1)^2}\left(\sum_{i=1}^{n}\mathrm{E}^Q\left[(X_i^*)^4\right] + 2\sum_{i<j}\mathrm{E}^Q\left[(X_i^*)^2(X_j^*)^2\right]\right)$$

$$= \frac{1}{(T_e - T_s)^2} \cdot \frac{n^2}{(n-1)^2}\left(\sum_{i=1}^{n}(\alpha_i^4 + 6\alpha_i^2\beta_i^2 + 3\beta_i^4)\right.$$

$$\left. + 2\sum_{i<j}\left(\alpha_i^2\alpha_j^2 + \alpha_i^2\beta_j^2 + 2\beta_i^2\alpha_j^2 + 3\beta_i^2\beta_j^2 + 2\alpha_i\alpha_j\beta_i\beta_j\right)\right).$$

Since $\alpha_i = -\frac{1}{2}\beta_i^2$, the above expression can be simplified.

$$\pi$$

2.7   Conclusion

During the short week, our group has obtained closed-form formulas for the expectation and variance of the realized pseudo-volatility-square, as requested by the problem:

$$\mathrm{E}^Q \left[\hat{\Sigma}^2_{(S)}(n; T_S, T_e) \right] \qquad \text{(Question 1)}$$

$$\mathrm{Var}^Q \left[\hat{\Sigma}^2_{(S)}(n; T_S, T_e) \right] \qquad \text{(Question 2)}$$

Although these formulas are complicated, they can be (carefully) implemented on a computer to compute prices of pseudo-variance, pseudo-covariance, and pseudo-volatility swaps in real time. This is the advantage of formulas over (quasi-) Monte-Carlo simulation and is highly desired by market practitioners.

In theory, we should be able to obtain the answer for question three as well, that is, to compute the covariance of different combinations of

$$A_1(n) := \hat{\Sigma}^2_{(S^{(1)})}(n; T_S, T_e),$$

$$A_2(n) := \hat{\Sigma}^2_{(S^{(2)})}(n; T_S, T_e),$$

$$A_{12}(n) := \hat{\Sigma}^2_{(S^{(1)} \cdot S^{(2)})}(n; T_S, T_e), \qquad \text{and}$$

$$A_{1/2}(n) := \hat{\Sigma}^2_{(S^{(1)}/S^{(2)})}(n; T_S, T_e).$$

However, we did not have time to complete this task, so this can be part of future work. Also, we did not have the opportunity to verify the appropriateness or validity of our results with our industry representative Ritchie He of RBC Financial Group, since he was away for the last days of the week.

$$\pi$$

Bibliography

[1] Brockhaus, O. and Long, D., Volatility Swaps Made Simple, *Risk*, January 2000, pp. 92–95.

[2] Demeterfi, K., Derman E., Kamal, M., and Zou J., A Guide to Variance Swaps, *Risk*, June 1999, pp. 54–59.

[3] Demeterfi, K., Derman E., Kamal, M., and Zou J., A Guide to Volatility and Variance Swaps, *The Journal of Derivatives*, Summer 1999, pp. 9–32.

[4] Hunt, P.J. and Kennedy, J.E., *Financial Derivatives in Theory and Practice*, John Wiley & Sons, 2000.

[5] Hunt, P.J., Kennedy, J.E., and Scott, E.M., Terminal Swap-Rate Models, Working paper, University of Warwick, 1996.

[6] Musiela, M., Rutkowski, M., *Martingale Methods in Financial Models*, Springer-Verlag, 1997.

Chapter 3

Determining Geological Properties by a Hybrid Seismic–Magnetotelluric Approach

Rita Aggarwala[1], Sean Bohun[2], Colin Farquharson[3], Hugh Geiger[1], Yuanyuan Hua[2], Huaxiong Huang[4], Michael Lamoureux[1], Jack Macki[5], Zheng Meng[6], Kai Muenzer[7], Doug Oldenburg[3], Robert Piche[8], Yongji Tan[9],

Report prepared by Hugh Geiger and Michael Lamoureux

3.1   Introduction

Of all the geophysical techniques, controlled-source reflection seismology is the preferred method for oil and gas exploration in the Canadian foothills, as it provides the best resolved images of structures at depth and can provide detailed information about subsurface rock properties such as elastic or acoustic impedance contrasts. Unfortunately, good quality seismic data cannot always be acquired over subsurface targets. The overburden can attenuate or scatter the seismic signal, preventing the signal from reaching the target zone and/or from returning to the detectors. The recorded signal might be swamped by noise, occasionally from external sources (such as production equipment), but more typically generated by interaction of the propagating wavefield with complicated near-surface conditions. In other cases, access to the surface over the target area is restricted and the resulting data set provides limited information about the target

[1]University of Calgary
[2]Pennsylvania State University
[3]University of British Columbia
[4]York University
[5]University of Alberta
[6]Rice University
[7]Shell Canada
[8]Tampere University of Technology
[9]Fudan University

and/or the overburden, both necessary prerequisites for high resolution imaging of impedance contrasts in the subsurface.

The controlled source audio magnetotelluric technique (CSAMT) is another viable geophysical tool for imaging subsurface structure. Unfortunately, electromagnetic waves propagating in conducting media satisfy the diffusion equation, and hence the best possible resolution of subsurface conductivity structure using the MT method is typically much worse than for seismic waves propagating in an elastic media. However, great strides have been made over the last decade in the acquisition, processing, interpretation and inversion of 2D and 3D MT data. It might be possible to collect high quality MT data in areas where the seismic data are of poor quality. In situ well logs that measure rock properties at the target and in the overburden show that the structural boundaries separating zones of different conductivity correspond exactly with major contrasts in elastic impedance. Thus, an inversion of the MT data constrained by a priori information obtained from seismic and well data or a joint inversion of both data sets together could provide a better image of the subsurface structure than the seismic alone.

Shell Canada posed a series of questions for consideration, all variations on a similar theme. These questions can be grouped in order from more general to more specific.

1. Could MT help to better define the geology, in comparison to using a regional geological model, in areas of poor seismic?

2. Can a geological model be produced by joint tomographic inversion of both data types?

3. Which would be favorable conditions for the application?

4. What kind of computing power is likely to be involved?

5. Given the recent advances in the acquisition, processing, interpretation and inversion of MT data and the associated improvement in lateral resolution, is there an emerging opportunity to combine MT with seismic in areas of complex geology?

6. What are the theoretical and practical issues involved in combining surface seismic and surface MT data by joint tomographic inversion in areas where seismic alone cannot sufficiently explain the subsurface in the Canadian Foothills.

7. Is common inversion of seismic and magnetotelluric data an option for gas exploration in the heavy structured thrust belt environment of the Canadian Foothills?

To properly address questions 4–7 would require an expert knowledge of 2D and 3D MT and seismic reflection methods, as well as software for forward modeling and inversion. Our group determined that these resources were not readily available, or if available could not be utilized in the short time span available. Therefore, we decided to focus our efforts on question 1–3. We limited our investigation to a simple 1D earth model where regional seismic and well logs suggest discrete layers, each with constant seismic velocity and constant electrical conductivity. In addition, the well logs provide rough estimates of velocity and conductivity for use as a starting point in the seismic and MT inversions.

To the best of our knowledge, there are no previous works that attempt a joint inversion of both MT and seismic data, although our joint inversion scheme closely follows the generalized

methodology proposed by Lines et al. (1988). Ellis (1999) shows that joint inversion of airborne EM data and pole-pole dc resistivity data can effectively reduce non-uniqueness in EM inversion. Sasaki (1989) shows that joint 2D inversion of MT and dipole-dipole resistivity can effectively improve the resolution compared to that from either of the separate inversions. More typically, a-priori information from seismic and well data are introduced as constraints on the lateral roughness of boundaries between abrupt changes in conductivity, and on the smoothness of conductivities within layers (Smith et al. 1999). Newman et al. (2002) have published the most comprehensive investigation to date of seismic-constrained MT inversion. In addressing a complex 3D subsalt imaging problem, they apply constraints that include fixed and floating boundaries as well as fixed values and/or lower bounds for conductivities above and within the salt. Routh and Oldenburg (1999) discuss the specific problem of 1D inversion for CSAMT, while Li et al (2000) examine 1D inversion for CSAMT with azimuthal anistropy.

In this paper, we focus more on the magnetotelluric problem, as it is somewhat less well-known. A deeper understanding of this model will be important to further studies in the feasibility of joint seimic-MT inversions.

3.2 Justification for 1D earth model

Although our 1D earth model is a gross simplification of the complex 2D and 3D geologic structure typical of the Canadian Foothills, an investigation of the joint inversion problem for 1D earth structure is necessary and useful. Most 2D and 3D MT inversion schemes require a reasonable starting model. The usual practice is to create the starting model from a profile of constrained 1D MT inversions. A robust procedure for the 1D inversion problem is unlikely to yield subsurface models suitable for oil and gas exploration, but might yield an intermediate model suitable for 2D and 3D MT inversions. A very similar argument applies to the seismic imaging problem, but it is worthwhile discussing in detail as there are important differences between the two.

Conventional seismic imaging is achieved by migrating the data. Migration is best described as a data processing step that convolves a space-time filter with the seismic data. The image can be created in a single pass, or in a series of passes that iteratively reshape the data towards the desired image. In short, then, the subsurface reflectivity model is obtained by migration, not inversion. However, one of the most important parameters controlling the behaviour of these filters is velocity. Hence, we require a reasonable velocity model of the subsurface, and it is towards this end that our inversion efforts are directed. In fact, one of the attractions of seismic migration is that a well-focused image of the subsurface can often be obtained with a very approximate velocity model (Geiger 2001). Thus, the output from our inversion process – a simple 1D layered velocity model – can be considered as an intermediate product in the quest for more detailed 2D and 3D velocity models, and these in turn as parameters controlling the behaviour of the migration operators.

π

3.3   Basic seismic wave theory

It is well known (e.g. Wapenaar and Berkhout, 1989) that acoustic and elastic wavefields satisfy the scalar wave equation. Although earth materials are elastic, most industry applications of the seismic reflection method only consider propagation of compressional waves, whose behaviour is adequately described by the acoustic wave equation for pressure. In an isotropic (i.e. material properties do not depend on direction of wave-propagation), source-free region, pressure P satisfies the homogeneous wave equation in the frequency domain:

$$(\nabla^2 + k^2)P = 0 \tag{3.1}$$

where $k = \omega\sqrt{\rho/K}$ is the wave number for frequency ω, density ρ , and adiabatic compression modulus K. Equation (3.1) is a scalar Helmholtz equation that describes the propagation of acoustic waves in real earth materials.

For a plane wave propagating along the z axis (positive down), equation (3.1) has solution

$$P = P^+ e^{ikz} + P^- e^{-ikz}. \tag{3.2}$$

If we express the time dependence of the harmonic wave for positive frequencies explicitly as $e^{-i\omega t}$, i.e.

$$P = P^+ e^{i(kz-\omega t)} + P^- e^{-i(kz-\omega t)}, \tag{3.3}$$

we see that the first and second terms on the RHS of equations (3.2) and (3.3) represent down-going waves (for z positive) and upgoing waves (for z negative), respectively. The phase velocity $\omega/k = \sqrt{K/\rho} \equiv \nu$ is only a function of the material properties and is independent of frequency, thus the acoustic wave propagation is non-dispersive. Note as well that, in a homogeneous media, propagating acoustic plane waves do not suffer attenuation. These properties are the fundamental reason for the superior resolution of the seismic methods.

3.4   Forward modelling of seismic reflection data

In essence, reflection seismology is a remote-sensing technique based on the principles of echolocation. An energy source at the surface of the Earth introduces sound waves that propagate through the various rock layers in the subsurface. Abrupt changes in the density and adiabatic compression moduli – which typically correspond to layer boundaries or structural discontinuities – impede the transmission of the sound waves. At each impedence constrast, or 'reflector, a small fraction of the propagating energy is reflected or diffracted and returns to the surface as a faint echo of the source. Receivers that measure acoustic pressure or particle motion detect these 'reflection events. The corresponding amplitudes are digitally recorded as time series known as 'seismograms. Because the signal propagates without significant dispersion or attenuation (other than geometrical spreading effects), an impulsive source will generate seismograms with well defined reflection events. The seismograms can be processed to create an image or representation of the subsurface.

A high-frequency approximation to the forward modeling operator (known as ray theory) can accurately describe wave propagation as a linear superposition of traveltime-delayed and

π

amplitude-scaled copies of the source signal, one for each raypath linking a given source with a subsurface reflector element and back to a receiver. Because the experiment is bandlimited, the infinitely many raypaths to all possible subsurface reflectors can be approximated by a finite subset (by discrete sampling of each subsurface reflectors). This subset of raypaths can be further reduced by the application of stationary phase, with the result that the vast majority of the recorded signal comes from specular reflections where the angles of incidence and reflection at a given reflector element are equal. For simple earth models such as the 1D layer case, the forward modeling problem for one source and receiver reduces to calculation of a single traveltime along a piecewise straight raypath. In most applications, the amplitude correction is ignored, or approximated by an estimate of the geometrical spreading along the raypath. We adopt this methodology to forward model the seismic data.

3.5   Inversion of seismic data

The fundamental inverse problem for seismic reflection consists of estimating acoustic or elastic impedance contrasts (angle dependent reflectivity) as a function of position (Gray, 1997 and Bleistein et al, 2001). However, as discussed above, the image is conventionally produced by migration (data processing) rather than via a true mathematical inversion. Our interest is in obtaining a velocity model that represents useful parameters for the migration model. In our model here, we examine the seismic data and pick travel times, which represent the arrival times for give seismic signals returning after relection from geological layers. These travel times are then used as inputs to the inversion problem for the velocity model.

3.6   Basic electromagnetic wave theory

Consider Maxwells equations in the frequency domain for an isotropic, source-free region with the constitutive relations $\vec{D} = \varepsilon\vec{E}$, $\vec{B} = \mu\vec{H}$, and $\vec{J} = \sigma\vec{E}$:

$$\vec{\nabla} \times \vec{H} = (\sigma - i\omega\varepsilon)\vec{E} \tag{3.4}$$

$$\vec{\nabla} \times \vec{E} = i\omega\mu\vec{H} \tag{3.5}$$

$$\vec{\nabla} \cdot \vec{H} = 0 \tag{3.6}$$

$$\vec{\nabla} \cdot \vec{E} = 0 \tag{3.7}$$

for frequency ω, magnetic permeability μ, dielectric permittivity ε, and electrical conductivity σ. The inverse of conductivity is known as electrical resistivity $\rho = \sigma^{-1}$, which should not be confused with density above.

In a homogeneous, isotropic, source-free region, both the electric field $\vec{E}$ and magnetic field $\vec{H}$ can be shown (e.g. Ward and Hohmann, 1988) to satisfy the homogeneous wave equation in the frequency domain

$$(\nabla^2 + k^2)\vec{E} = 0 \tag{3.8}$$

$$(\nabla^2 + k^2)\vec{H} = 0 \tag{3.9}$$

π

where k is the propagation constant or wave number, defined as $k^2 = i\omega\mu\sigma$ at low frequencies where conduction currents predominate over displacement currents. Equations (3.8) and (3.9) are vector Helmholtz

For a plane wave propagating along the z-axis (positive down), equations (3.8) and (3.9) have solutions

$$\vec{E} = \vec{E}^+ e^{ikz} + \vec{E}^- e^{-ikz}, \tag{3.10}$$

$$\vec{H} = \vec{H}^+ e^{ikz} + \vec{H}^- e^{-ikz}. \tag{3.11}$$

Expanding $k = \alpha + i\alpha$, where α is the positive real quantity $\sqrt{\omega\mu\sigma/2}$, and expressing the time dependence of the harmonic wave for positive frequencies explicitly as $e^{-i\omega t}$, we see that the first terms on the RHS of equations (3.10) and (3.11) represent downgoing waves (for z positive)

$$\vec{E}^+ e^{ikz} e^{-i\omega t} = (\vec{E}^+ e^{-\alpha z}) e^{i(\alpha z - \omega t)}, \tag{3.12}$$

while the second terms on the RHS of equations (3.10) and (3.11) represent upgoing waves (for z negative)

$$\vec{E}^- e^{-ikz} e^{-i\omega t} = (\vec{E}^- e^{\alpha z}) e^{-i(\alpha z + \omega t)}, \tag{3.13}$$

In both cases, the electromagnetic waves attenuate as they propagate. In addition, the phase velocity $\omega/\alpha = \sqrt{2\omega/\mu\sigma}$ is frequency dependent, even in common earth materials where conductivity and permeability are independent of frequency. Hence, electromagnetic waves propagating in materials with non-zero conductivity are dispersive. Dispersion and attenuation are the reasons for the lack of resolution of electromagnetic prospecting methods.

3.7   Forward modelling of magnetotelluric data

From the vanishing of the divergences (equations (3.6) and (3.7)), the fields lie completely in the x-y plane, so that

$$\vec{E} = (E_x, E_y, 0), \tag{3.14}$$

$$\vec{H} = (H_x, H_y, 0). \tag{3.15}$$

From the curl equation (3.5), the components of the two fields are related by

$$H_x(z) = -\frac{k_n}{\omega\mu}(E_y^{+n} e^{ik_n z} - E_y^{-n} e^{-ik_n z}) \tag{3.16}$$

$$H_y(z) = \frac{k_n}{\omega\mu}(E_x^{+n} e^{ik_n z} - E_x^{-n} e^{-ik_n z}) \tag{3.17}$$

where the subscript n indicates a specific layer in a layered media with conductivities σ_n. Equations (3.16) and (3.17) are only valid within the limits of the layer with index n.

In the Tikhonov-Cagnaird approach, the wave impedance is defined as

$$Z_{xy}(z) = \frac{E_x(z)}{H_y(z)} \tag{3.18}$$

$$Z_{yx}(z) = \frac{E_y(z)}{H_x(z)} = -Z_{xy}(z) \tag{3.19}$$

i.e. the ratio of orthogonal components of the electric and magnetic fields. Note that the wave impedance will be a complex quantity, and thus have both amplitude and phase.

For a homogeneous half-space, there are no upgoing waves. Substituting the downgoing x-component of equation (3.10) and the downgoing portion of equation (3.17) into equation (3.18) yields

$$Z_{xy} = \sqrt{\frac{-i\omega\mu}{\sigma}}. \tag{3.20}$$

Similarly, substituting the downgoing y-component of equation (3.11) and the downgoing portion of equation (3.16) into equation (3.19) yields

$$Z_{yx} = -\sqrt{\frac{-i\omega\mu}{\sigma}}. \tag{3.21}$$

For normally conducting rocks $\mu = \mu_0$, thus the observed impedance is dependent only on frequency and the electrical conductivity in the earth. It is not dependent either on the strengths of the individual components of the electromagnetic field or the measurement depth z. Thus measurements taken at the surface of the earth ($z = 0^+$) are sufficient to determine conductivity (or its inverse, resistivity). The resistivity of the half-space can be determined from the field by

$$\rho = \frac{1}{\sigma} = \frac{1}{\omega\mu}\frac{|E_x|^2}{|H_y|^2} = \frac{1}{\omega\mu}\frac{|E_y|^2}{|H_x|^2}. \tag{3.22}$$

For layered media, Tikhonov and Cagniard both showed (see e.g. Zhdanov and Keller, 1994, p262-269) that the impedance measured at the surface of the earth can be expressed as

$$Z_{1D} = \sqrt{\frac{-i\omega\mu}{\sigma_1}}C_N, \tag{3.23}$$

where C_N is the layered earth correction factor for the plane wave impedance for the layered earth as a whole. C_N is a function of the layer thicknesses and propagation constants in the subsurface, found by matching continuity conditions at each layer interface $z = h_n$ (both H_x and E_y, or H_y and E_x are continuous acrross layer interfaces).

Cagniard defined an apparent resistivity ρ_a in terms of the observed wave impedance at the earth's surface as

$$\rho_a = \frac{1}{\omega\mu}|Z|^2, \tag{3.24}$$

a quantity often used instead of the wave impedance Z_{1D}.

For the 1D magnetotelluric method, the ideal data set consists of measurements of the orthogonal horizontal components of the electric and magnetic field over a range of frequencies, and occasionally the vertical magnetic field (see discussion of tipper vector below). This allows determination of both amplitude and phase of the impedance. The phase of the impedance, although not been discussed above, is a valuable measurement because it can help constrain 1D inversions in the presence of noise and static shifts. However, it is often sufficient to have measurements of the apparent resisitivity over a range of frequencies, as we will assume here.

3.8   Inversion of magnetotelluric data

The fundamental inverse problem for MT consists of estimating conductivity as a function of position (Hohmann and Raiche, 1988; Whittall and Oldenburg, 1991). MT is considered a classical geophysical inverse problem. The MT inverse problem is non-linear and ill-posed. For any real subsurface distribution of conductivity and finite data set, the problem is underdetermined (more model parameters than data). In order that a solution of any sort can be achieved, the subsurface is simplified so that the problem is overdetermined (more data than model parameters). Thus there are a number of possible solutions, all of which might fit the data equally well. An additional contribution to non-uniqueness is the presence of random noise in the data, although repeatable geologic noise can manifest itself as static shifts (bias) in measured field components, resulting from near surface inhomogeneities (Wannamaker, 1999; Garcia and Jones, 2002), as well as fine-scale variations in resistivity-depth profiles as seen in well logs (White et al., 2001).

All MT inversions are constrained in some sense. The constraints might include careful preprocessing and/or the selection of subsets from the available data (Agarwal et al., 1993; Oldenburg and Ellis, 1993; Berdichevsky et al., 1998; Pedersen and Gharabi, 2000), choice of a particular class of models (e.g. smoothest model: Constable et al., 1987; Smith and Booker, 1988, or minimum blocky model: Fischer and Weibel, 1991), selection of a means for searching the model space to find the best fit model (Weaver and Agarwal, 1993; Grandis et al. 1999, Siripunvaraporn and Egbert, 2000), and the introduction of a priori bounds that limit or guide the fitting process (Whittall, 1986; Mackie and Madden, 1993; Smith et al. 1999; Newman et al. 2002) .

3.9   Practical considerations for 1D MT

In most cases, the earth cannot be adequately described by a 1D model. However, reasonable 1D inversions might still be desirable, if only to provide a good starting model for further 2D and 3D inversions. In addition, even in areas where a 1D model might be appropriate, near surface inhomogeneities distort the measured electric and magnetic fields and bias the calculation of impedances and apparent resistivities.

In areas of lateral inhomogeneity and/or anisotropy, the electric field in one direction may depend on magnetic field variations parallel to, as well as perpendicular to, its direction. The impedance is a tensor (see, e.g. Vozoff, 1991), given by

$$\begin{bmatrix} E_x \\ E_y \end{bmatrix} = \begin{bmatrix} Z_{xx} & Z_{xy} \\ Z_{yx} & Z_{yy} \end{bmatrix} \begin{bmatrix} H_x \\ H_y \end{bmatrix} \tag{3.25}$$

or, equivalently,

$$\mathbf{E} = \mathbf{ZH}. \tag{3.26}$$

In a 2D subsurface where geologic strike is parallel to the x-axis and perpendicular to the y-axis, the diagonal elements Z_{xx} and Z_{yy} are zero, and $Z_{xy} \neq -Z_{yx}$. Impedance Z_{xy} corresponds with an incident electric field along x, which is often referred to as the transverse electric (TE) mode. The E_x field of the TE mode induces no boundary or space charges and so there are

no boundary charge effects. Impedance Z_{yx} corresponds with an incident electric field along y, which is often referred to as the transverse magnetic mode or TM mode. The E_y field of the TM mode crosses boundaries, so that charges are induced, and anomalies persist to arbitrarily long periods. Thus the TM impedance is more sensitive to near-surface resistive structures than the TE impedance. Agarwal et al. (1993) examine the two common effective impedances which are rotationally invariant in the x-y plane:

$$Z_{ave} = \frac{1}{2}(Z_{xy} - Z_{yx}), \tag{3.27}$$

$$Z_{eff} = (Z_{xx}Z_{yy} - Z_{xy}Z_{yx})^{1/2}. \tag{3.28}$$

These are valid for any orientation of the acquisition geometry relative to geologic strike, although they conclude that the TE impedance is more suitable for 1D inversions in 2D media (e.g. to generate a starting model for 2D inversions).

Near lateral conductivity changes, $\vec{\nabla} \times \vec{E}$ has a vertical component, thus H_z is not zero and the $\vec{H}$ vector is tipped out of the horizontal plane. A vertical magnetic transfer function or tipper (A, B) can be defined as

$$H_z = AH_x + BH_y. \tag{3.29}$$

The tipper has both amplitude and phase and is frequency dependent. Given a 2D structure with strike, the tipper can help resolve ambiguity in strike and show which side of a contact is more conductive (Vozoff, 1991).

In general, near-surface inhomogeneities introduce galvanic distortion effects that can be described by distortion matrices that relate the regional electromagnetic fields to the measured (superscript D) distorted fields (see Garcia and Jones, 2002; Li et al., 2000):

$$\mathbf{E}^D = \mathbf{P}\mathbf{E} \tag{3.30}$$

$$\mathbf{H}^D = \mathbf{Q}_h\mathbf{E} = \mathbf{I} + \mathbf{Q}_h\mathbf{Z} \tag{3.31}$$

$$H_z^D = [(A, B) + \mathbf{Q}_z\mathbf{Z}]\mathbf{H} \tag{3.32}$$

where $\mathbf{I}$ is the identity matrix, and $\mathbf{P}$, $\mathbf{Q}_h$ and $\mathbf{Q}_z$ are real valued and frequency-independent distortion matrices:

$$\mathbf{P} = \begin{pmatrix} P_{xx} & P_{xy} \\ P_{yx} & P_{zz} \end{pmatrix}, \mathbf{Q}_h = \begin{pmatrix} Q_{xx} & Q_{xy} \\ Q_{yx} & Q_{zz} \end{pmatrix}, \text{ and } \mathbf{Q}_z = \begin{pmatrix} Q_{zx} & Q_{zy} \end{pmatrix}. \tag{3.33}$$

When the EM fields are measured in an arbitrary coordinate system relative to the geologic strike, the observed impedance $\mathbf{Z}^D$ is related to the desired regional 2D impedance $\mathbf{Z}_{2D}$ by

$$\mathbf{Z}^D = \mathbf{R}(\theta)(\mathbf{I} + \mathbf{P})\mathbf{Z_{2D}}(\mathbf{I} + \mathbf{QZ})^{-1}\mathbf{R^T}(\theta), \tag{3.34}$$

where $\mathbf{R}(\theta)$ is the Cartesian rotation tensor, and T denotes transpose. Theoretical and practical aspects of the decomposition are discussed in McNeice and Jones (2001), Garcia and Jones (2002), and Wannamaker (1999) (as well as references therein).

$$\pi$$

3.10   Workshop Model

At the IPSWorkshop, we choose to investigate a simple one-dimensional model, to get some idea of how the the joint seismic/magnetotelluric inversion might proceed. We consider a hybrid seismic-magnetotelluric method for determining geological properties using a one-dimensional three-layer model. We assume that the medium under consideration has three horizontal layers, each of them is homogeneous with constant layer thickness h_n, resistivity ρ_n and seismic velocity v_n, $n = 1, 2, 3$ from top down, and layer three has infinite thickness, i.e., $h_3 = \infty$.

3.10.1   Seismic method

Forward problem. Seismic signals are produced at certain location on the surface of the top layer ($z = 0$) and reflections of the signals from the interfaces of layers ($z = h_1$, h_2) are recorded at different locations. Suppose that the distance between the source and the jth receiver is $2x_j$, then $T_{n,j}$, the time for the signal to reach the receiver j, after nth reflection, can be computed as

$$T_{1,j} = \frac{\sqrt{x_j^2 + h_1^2}}{v_1}, \tag{3.35}$$

$$T_{2,j} = \sqrt{2\left(\frac{h_1}{v_1} + \frac{h_2 - h_1}{v_2}\right)\left(1 + \frac{2x_j^2}{h_1 v_1 + (h_2 - h_1)v_2}\right)} \tag{3.36}$$

where v_n are the seismic velocities in layer n.

Inverse problem. In practice, the situation is the opposite. We can record the travel time $T_{n,j}$ at location j and often the layer thickness and seismic velocity are not known. The question is how to find h_n and v_n based on the values of $T_{n,j}$.

   The solution of the inverse problem is typically non-unique and sometimes unstable. The main challenge is to find the correct solution or an approximation efficiently.

3.10.2   Magnetotelluric method

Forward problem. Assume that the electric-magnetic field satisfies the Maxwell equations and the traveling wave solution satisfies

$$\nabla \times \vec{E} + i\mu\omega\vec{H} = 0, \;\; \nabla \times \vec{H} - \sigma\vec{E} = 0 \tag{3.37}$$

where ω is the frequency of the wave, $\sigma = \rho^{-1}$ is the conductivity and μ is the permeability of the medium. The equations are equivalent to the Helmholtz equations for $\vec{E}$ and $\vec{H}$ when there are no sources/sinks inside the domain (i.e., both $\vec{E}$ and $\vec{H}$ are divergence-free) and in our one-dimensional setting they are

$$\frac{d^2 E_n}{dz^2} + k_n^2 E_n = 0, \tag{3.38}$$

$$\frac{d^2 H_n}{dz^2} + k_n^2 H_n = 0 \tag{3.39}$$

π

where E_n and H_n are the x and y components of the vector fields $\vec{E}_n$ and $\vec{H}_n$ in layer n, respectively and $k_n^2 = -i\sigma_n \mu \omega$.

The solutions for E_n and H_n can be obtained as

$$E_n = A_n \exp(ik_n z) + B_n \exp(-ik_n z), \tag{3.40}$$

$$H_n = \frac{k_n}{\omega \mu} \left(A_n \exp(ik_n z) - B_n \exp(-ik_n z) \right). \tag{3.41}$$

When H_1 (or E_1) is given, the coefficients A_n and B_n can be determined by matching conditions at each layer interface $z = h_n$ (both H_n and E_n are continuous across the layer interfaces). Note that $B_3 = 0$ since otherwise the solution will grow unbounded as $z \to \infty$. And each solution E_n and H_n can be expressed by H_1 (or E_1), the layer thickness h_n, the conductivities σ_n, and the frequency of the wave. In particular, the ratio of the electric and magnetic field strengths at the surface $(z-0)$ are uniquely determined by h_n and σ_n as

$$Z \equiv \left. \frac{E_1}{H_1} \right|_{z=0} = \frac{\omega \mu}{k_1} \frac{1 - \sqrt{\frac{\sigma_2}{\sigma_1}} R \tanh(ik_1 h_1)}{\sqrt{\frac{\sigma_2}{\sigma_1}} R - \tanh(ik_1 h_1)} \tag{3.42}$$

where

$$R = \frac{1 - K_{23} \exp(i2k_2 H_{21})}{1 + K_{23} \exp(i2k_3 H_{21})}, \quad K_{23} = \frac{k_2 - k_3}{k_2 + k_3}, \quad H_{21} = h_2 - h_1.$$

A more conventional quantity called the *apparent resistivity*

$$\rho_{app} = \left| \frac{i}{\omega \mu} Z^2 \right| \tag{3.43}$$

is normally used instead of the ratio Z.

Inverse problem. Again, often we can measure the apparent resistivity ρ_{app} corresponding to various frequencies. To obtain the geological information such as the layer thickness h_n and resistivity $\rho_n = \sigma_n^{-1}$ based on ρ_{app} is of more practical importance.

3.10.3   A Hybrid Seismic-Magnetotelluric Approach

In practice, there are situations when one set of data (seismic or magnetotelluric) is more difficult to obtain than others. It will be of interest to explore the feasibility of a hybrid approach which combines both data.

We consider the following approach for the inversion of seismic and magnetotelluric data by minimizing the objective function

$$J = \sum_{n,j} w_{n,j}^1 \left(T_{n,j}^m - T_{n,j}^c \right)^2 + \sum_j w_{n,j}^2 \left(\rho_{app}^m(\omega_j) - \rho_{app}^c(\omega_j) \right)^2 \tag{3.44}$$

where $T_{n,j}^m$ and ρ_{app}^m are measured traveling time for the seismic signal and apparent resistivity, $T_{n,j}^c$ and ρ_{app}^c are the computed values using (3.35), (3.36) and (3.43).

In the following, we present numerical experiments which illustrate the importance of the initial guess and the weight functions w^1 and w^2. In general the results suggest that we should always over-estimate the layer thickness for the initial guess. And it is necessary to use a staged procedure in order to find the correct solutions.

$$\boxed{\pi}$$

3.11   Numerical Experiments

To test the hybrid approach, we replace the measured data $T^m_{n,j}$ and ρ^m_{app} by the exact solutions obtained using (3.35), (3.36) and (3.43) based on the following parameter values: $h_1 = 400$, $h_2 = 4000$, $v_1 = 6332$, $v_2 = 4114$, $\rho_1 = 8000$, $\rho_2 = 40$, $\rho_3 = 8000$ and $\mu = 10^{-7}$. We assume that the seismic data is recorded at nine locations with distance between the source and receiver varying between 1000 and 5000 with equal increments:

$$2x_j = 1000 + \frac{j-1}{8}4000, \quad j = 1,...9.$$

The magnetotelluric data is recorded at the surface ($z = 0$) for 25 frequencies varying between 0.001 and 500 exponentially

$$\omega_j = 0.001 \left(\frac{500}{0.001}\right)^{\frac{j-1}{24}}, \quad j = 1,...25.$$

The minimization is carried out using the matlab routine `fminsearch` with maximum iteration number set at 2000, maximum number of function evaluations at 10000, convergence criteria for the function value at 10^{-8}. Table 3.1 listed the results for various computations with initial guess $h_1 = h_2 = 10000$, $v_1 = iv_2 = 10000$, $\rho_1 = 6000$, $\rho_2 = 600$, and $\rho_3 = 6000$. We have chosen equal weights for both the seismic and magnetotelluric data in the functional J in (3.44) with two parameters w^1 and w^2. For two of the computations ($w^2 = 100$ and 1000), the program is re-started using the results from the previous calculation as initial guess since some of the corrected values are obtained after the first attempt. It can be seen that the hybrid approach is effective when appropriate weights are chosen. Even though the first iteration does not yield all the correct solution, a second attempt does recover the rest of the solution for the case of $w^2 = 100$. However, for $w^2 = 1000$, the iteration diverges (second row from the bottom) and it converges only after taking the absolution value of the results from the first attempt as an initial guess (first row from the bottom).

3.12   Conclusions

Both seismic imaging and magnetotellurics are well-developed technologies with a rich research literature and history of development; in our one-week exercise, we could only expect to touch on some basic ideas of the methods and make some initial attempts to understand both methods and link them in a hybrid approach. With numerical experiments, we were able to demonstrate a basic understanding of both forms of inversion for a 1D earth model, and make some investigations on joint inversion using a weighed least squares technique. Even this simple method demonstrates promise for combining both seismic and MT data to give an effective inversion method. Weighting for the least squares is significant for balancing the contribution of data from the two methods. A more realistic experiment would involve real data from a 2D seismic and MT survey, but using the same methodology of least squares.

It is worth noting that an extensive literature on seismic and MT imaging methods exists, and a familiarity with this work is critical to making further progress on joint inversion. We include a representative sample of some key research articles in the bibliography below.

π

Table 3.1: *Effect of the weights w^1 and w^2 on the solution of the inverse problem using the hybrid approach. For comparison purpose, the seismic only ($w^2 = 0$) and magnetotelluric only ($w^1 = 0$) inversion are also included.*

w^1	w^2	h_1	h_2	v_1	v_2	ρ_1	ρ_2	ρ_3
1	0	-400	4061	6332	4455	-	-	-
0	100	0	3738	-	-	0	39	8745
1	10	-398	4062	6330	4455	8001	51	20387
1	100	400	4000	6332	4114	322	38	20285
		400	4000	6332	4114	8000	40	8000
1	1000	397.7	3998.6	6330	4114.1	-4256.3	40.1	7523.5
		399.2	3999.5	6331.3	4114	-2.143×10^{12}	40	7828.4
		400	4000	6332	4114	8000	40	8000

Bibliography

[1] Agarwal, A.K., Poll, H.E., and Weaver, J.T., 1993. *One- and two-dimensional inversion of magnetotelluric data in continental regions,* Physics of the Earth and Planetary Interiors, 81, pp. 155–176.

[2] Berdichevsky, M.N., Dmitriev, V.I., and Pozdnjakova, E.E., 1998. *On two-dimensional interpretation of magnetotelluric soundings,* Geophys. J. Int., 133, pp. 585–606.

[3] Bleistein, N., Cohen, J. K., and Stockwell, J. W. J., 2001. *Mathematics of multidimensional seismic inversion,* Springer-Verlag New York, Inc., Interdisciplinary Applied Mathematics, 6.

[4] Cagniard, L., 1953. *Basic theory of the magnetotelluric method of geophysical prospecting,* Geophysics, 18, pp. 605–535.

[5] Constable, S.C., Parker, R.L., and Constable, C.G., 1987. *Occams inversion: A practical algorithm for generating smooth models from electromagnetic sounding data,* Geophysics, 52, pp. 289–300.

[6] Ellis, R. G., 1999. *Joint 3-D electromagnetic inversion,* in "3D Electromagnetics" (Eds. B. Spies and M. Oristaglio), SEG Publ., GD7, Tulsa, USA, pp. 179–192.

[7] Fischer, G., and Weibel, P., 1991. *A new look at an old problem: magnetotelluric modelling of 1-D structures,* Geophys. J. Int., 106, pp. 161–167.

[8] Garcia, X. and Jones, A.G., 2002. *Decomposition of three-dimensional magnetotelluric data,* in Zhdanov, M.S. and Wannamaker, P.E. (eds), Proceedings of the second international symposium on three-dimensional electromagnetics, Elsevier, pp. 236–250.

[9] Geiger, H.D., 2001. *Relative-amplitude-preserving prestack time migration by the equivalent offset method,* Ph.D. thesis, University of Calgary.

[10] Grandis, H., Menveille, M., and Roussignol, M., 1999. *Bayesian inversion with Markov chains I. The magnetotelluric one-dimensional case,* Geophys. J. Int., 138, pp. 757–768.

[11] Gray, S. H., 1997. *True-amplitude seismic migration: A comparison of three approaches:,* Geophysics, 62, pp. 929–936.

[12] Hohmann, G.W. and Raiche, A.P., 1988. *Inversion of controlled-source electromagnetic data,* in Electromagnetic methods in applied geophysics, vol 1, theory (ed. Nabighian, M.N.), SEG Publ. pp. 469–503.

[13] Kaufman, A.A. and Keller, G.V., 1981. *The magnetotelluric sounding method,* Elsevier.

[14] Li, X., Oskooi, B., and Pedersen, L.B., 2000. *Inversion of controlled-source tensor magnetotelluric data for a layered earth with azimuthal anisotropy,* Geophysics, 65, pp. 452–464.

[15] Lines, L.R., Achultz, A.K., and Treitel, S., 1988. *Cooperative inversion of geophysical data,* Geophysics 53, pp. 8–20.

[16] Mackie, R.L. and Madden, T.R., 1993. *Three-dimensional magnetotelluric inversion using conjugate gradients,* Geophys. J. Int., 115, pp. 215–229.

[17] McNeice, G.W. and Jones, A.G., 2001. *Multisite, multifrequency tensor decomposition of magnetotelluric data,* Geophysics, 66, pp. 158–173.

[18] Newman, G.A., Hoversten, G.M. and Alumbaugh, D.L., 2002, *Three-dimensional magnetotelluric modeling and inversion : Application to sub-salt imaging,* in Three-dimensional electromagnetics: Proceedings of the second international symposium (Eds. Zhdanov, M.S. and Wannamaker, P.E.), Elsevier, pp. 127–152.

[19] Oldenburg, D.W. and Ellis, R.G., 1993. *Efficient inversion of magnetotelluric data in two dimensions,* Phys. Earth Planet. Inter., 81, pp. 177–200.

[20] Pedersen, L.B. and Gharibi, M., 2000. *Automatic 1-D inversion of magnetotelluric data: Finding the simplest possible model that fits the data,* Geophysics, 65, pp. 773–782.

[21] Rodi, W. and Mackie, R.L., 2001. *Nonlinear conjugate gradients algorithm for 2-D magnetotelluric inversion,* Geophysics, 66 pp. 174–187.

[22] Routh, P.S. and Oldenburg, D.W., 1999. *Inversion of controlled source audio-frequency magnetotellurics data for a horizontally layered earth,* Geophysics, 64, pp. 1689–1697.

[23] Sasaki, Y., 1989. *Two-dimensional joint inversion of magnetotelluric and dipole-dipole resistivity data,* Geophysics, 54, pp. 254–262.

[24] Siripunvaraporn, W. and Egbert, G., 2000. *An efficient data-subspace inversion method for 2-D magnetotelluric data,* Geophysics, 65, pp. 791–803.

[25] Smith, J.T. and Booker, J.R., 1991. *Rapid inversion of two- and three-dimensional magnetotelluric data,* J. Geophys. Res., 96, pp. 3905–3922.

[26] Smith, J.T. and Booker, J.R., 1988. *Magnetotelluric inversion for minimum structure,* Geophysics, 52, pp. 1565–1576.

[27] Smith, T., Hoversten, M., Gasperikova, E., and Morrision, F., 1999. *Sharp boundary inversion of 2D magnetotelluric data,* Geophys. Prosp., 47, pp. 469–486.

π

[28] Spichak V.V., Menvielle M. and Roussignol M., 1999. *Three-dimensional inversion of MT data using Bayesian statistics,* in "3D Electromagnetics" (Eds. B. Spies and M. Oristaglio), SEG Publ., GD7, Tulsa, USA, pp. 406–417.

[29] Tikhonov, A.N., 1950. *On determining the electric properties of deep layers of the earths crust,* Proc. (Dokaldy) Acad. Sci. USSR 83 2.

[30] Vozoff, K., 1991. *The magnetotelluric method,* in Electromagnetic methods in applied geophysics, vol 2, application, part B (ed. Nabighian, M.N.), SEG Publ. pp. 641–711.

[31] Wannamaker, P.E., 1999. *Affordable magnetotellurics: Interpretation in natural environments,* in Oristaglio, M. and Spies, B. (eds), Proceedings of the first international symposium on three-dimensional electromagnetics, SEG GDS 7, pp. 349–374.

[32] Ward, S.H. and Hohmann, G.W., 1988. *Electromagnetic theory for geophysical applications,* in Electromagnetic methods in applied geophysics, vol 1, theory (ed. Nabighian, M.N.), SEG Publ. pp. 131–311.

[33] White, B.S., Kohler, W.E., and Srnka, L.J., 2001. *Random scattering in magnetotellurics,* Geophysics, 66, pp. 188–204.

[34] Whittall, K.P., 1986. *Inversion of magnetotelluric data using localized conductivity constraints,* Geophysics, 51, pp. 1603–1607.

[35] Whittall, K.P. and Oldenburg, D.W., 1991. *Inversion of magnetotelluric data for a one-dimensional conductivity,* SEG geophysical monograph series 5.

[36] Yilmaz, O., 2001. *Seismic data analysis, processing, inversion, and interpretation of seismic data,* SEG investigations in geophysics.

[37] Zhdanov, M.S. and Keller, G.V., 1994. *The geoelectrical methods in geophysical exploration,* Elsevier.

[38] Zonge, K.L. and Hughes, L.J., 1991. *Controlled source audio-frequency magnetotellurics,* in Nabighian, M.N. (ed), Electromagnetic methods in applied geophysics, vol 2, application, part B, SEG Publ. pp. 713–809.

π

Chapter 4

A Seismic Inversion Problem for an Anisotropic, Inhomogeneous Medium

Dhavide Aruliah[1], Mikalai Birukou[2], Neville Dubash[3], Leslie Fairbairn[4], Jeff Grossman[5], Katherine Hegewisch[6], Namyong Lee[7], Robin Swain[8], Carlos Trenado[9], Industrial Contacts:Gary Billings[10], Michael Slawinski (Mentor)[8], Chad Wheaton (Mentor)[8],

4.1 Introduction

In the petroleum industry, many geoscientists rely on information provided by seismic data. Based on the interpretation of seismic data, geophysicists and geologists attempt to identify and target areas of economic importance, specifically hydrocarbon traps. We wish to extract from the data as much information as possible.

A key factor in the successful processing and interpretation of seismic data is knowing certain rock physics properties, in particular the elastic moduli of the section of earth through which seismic waves propagate. One means of attempting to quantify the rock physics properties is through a vertical seismic profile (VSP), wherein seismic sources are located on the surface and receivers are placed in a well-bore.

Traditionally, only traveltimes were measured at the receivers, but technological advances now allow measurement of the polarization angles (*i.e.*, particle displacement) as well. This provides an additional piece of information towards a practical means of accurately determining the physical parameters that characterize the medium.

[1]University of Toronto
[2]University of Waterloo
[3]University of British Columbia
[4]Simon Fraser University
[5]University of Calgary
[6]Washington State University
[7]Minnesota State University
[8]Memorial University
[9]New Mexico State University
[10]Talisman Energy
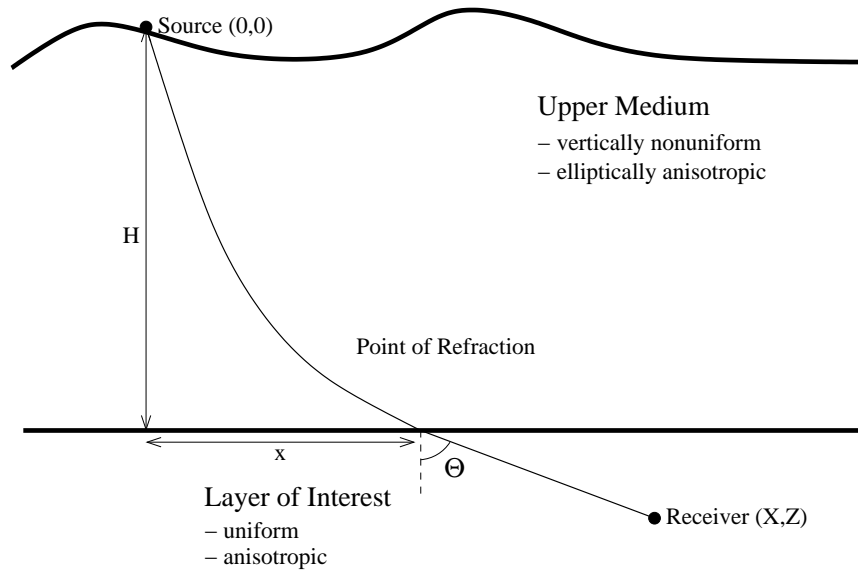
In this report, we consider the propagation of seismic waves through a medium that can be subdivided into of two distinct parts. The upper part is assumed to be azimuthally symmetric, linearly nonuniform with increasing depth, and the velocity dependance with direction consistent with elliptical anisotropy. The lower part, which is the layer of interest, is assumed to also be azimuthally symmetric, but uniform and nonelliptically anisotropic. Despite nonellipticity, we assume the angular dependance of the velocity can be described by a convex curve.

Our goal is to produce a single source-single receiver model which uses modern seismic measurements to determine the elastic moduli of the lower media. Once known, geoscientists could better describe the angular dependance of the velocity in the layer of interest and also would have some clues at to the actual material composing it.

4.2   Formulation of the Problem

We begin by assuming that the area of study consists of layered sedimentary rocks. Rotational symmetry in such a case effectively reduces problem to two dimensions. Since, in the case of rational symmetry, energy will propagate between source and receiver in a vertical plane, we can define our coordinate system such that our x and z coordinate axes contain the plane of propagation.

We start with a vertical cross section of the earth with a seismic source on the surface and a receiver located below. We define the horizontal offset between source and receiver to be X, and the vertical offset to be Z. That is, we define the source to be at the origin of our coordinate system, $(0,0)$, and the receiver to be at the location (X, Z). Furthermore, we assume that the subsurface is divided into two distinct media with the depth to the interface between them to be H. Finally, we define the horizontal offset between the source and the point of refraction of the seismic ray to be x.

π

4.2.1 Assumptions

We assume that the seismic source is of high frequency with a wavelength shorter than all characteristic lengths in the layers of interest. With this assumption, it follows, from asymptotic ray theory, that the signal travels along a ray, which is a trajectory of the stationary traveltime. Thus, we assume that the rays obey the variational principle of Fermat.

We also assume that the medium through which the seismic waves propagate is a perfectly linear elastic continuum. Hence the stress-strain relationships are described by Hooke's law.

4.2.2 Upper Medium

We consider the upper medium, $0 \leq z \leq H$, to be horizontally stratified $\{(i.e.,$ sedimentary rocks$)\}$ and both vertically nonuniform and elliptically anisotropic. Vertical nonuniformity is often the result of varying compaction of the subsurface and manifests itself as a linear increase in {vertical} velocity with depth. That is, $V_z(z) = a + bz$. In the case of elliptical anisotropy we can define a single parameter, χ, that describes the relationship between vertical and horizontal velocity[1]. Specifically,

$$\chi = \frac{V_x^2 - V_z^2}{2V_z^2}. \tag{4.1}$$

Combining the effects of both nonuniformity and anisotropy results in a ray velocity profile[2] of the form

$$V(z, \Theta; a, b, \chi) = (a + bz)\sqrt{\frac{1 + 2\chi}{1 + 2\chi \cos^2(\Theta)}} \tag{4.2}$$

where a, b, and χ are predetermined.

Then, by applying Noether's Theorem it can be shown that the ray form of Snell's Law leads to the expression for the ray parameter

$$p = \frac{2x}{\sqrt{(x^2 + (1 + 2\chi)\, z^2)\, (4a^2\, (1 + 2\chi) + 4ab\, (1 + 2\chi)\, z + b^2\, (x^2 + (1 + 2\chi)\, z^2))}}. \tag{4.3}$$

Applying the variational principle of Fermat, given the above assumptions, the traveltime function[3] of a ray traveling from a source at $(0,0)$ to a point (x, H) is

$$t_1(x, z; a, b, \chi) = \frac{1}{b} \ln \left| \frac{1 + \sqrt{1 - p(x)^2 a^2(1 + 2\chi)}}{1 + \sqrt{1 - p(x)^2(a + bz)(1 + 2\chi)}} \right|. \tag{4.4}$$

It has been assumed that the parameters a, b and χ, which characterize the upper medium, are known. These values are obtained by considering traveltimes from source to receivers located in the upper medium, and applying a best-fit method (*e.g.*, conjugate-gradient optimization) to the traveltime expression, t_1.

---

[1] Thomsen, L.,(1986) Weak Elastic Anisotropy, *Geophysics*, **51**, pp. 1954 - 1966.

[2] Affarwala, B., *et. al.*, (1997) Inversion of Anisotropic Velocity Parameter. *Proceedings of the first PIMS Industrial Solving Workshop*, PIMS, pp. 76-89.

[3] Slawinski, M., Wheaton, C., personal communication. 2002

$$\boxed{\pi}$$
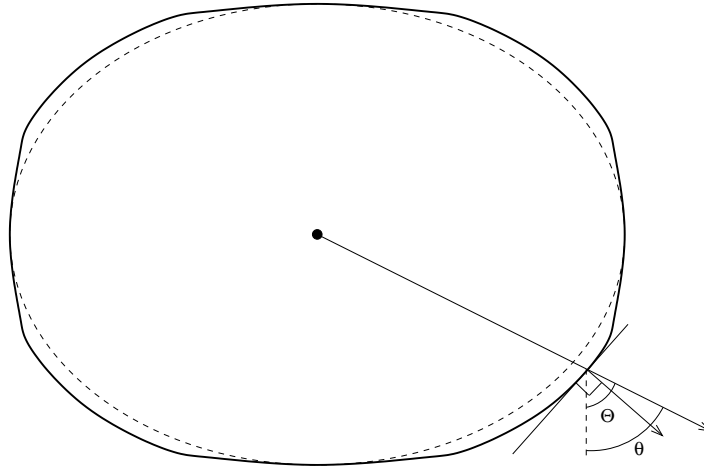
4.2.3   Layer of Interest

We consider the layer of interest, $H \leq z \leq Z$, to be uniform, perfectly elastic, and transversely isotropic. Uniformity implies that all rays will be straight lines from the refraction point to the receiver. Elasticity and transverse isotropy imply that under Hooke's Law, which give the relationship between stress, σ, and strain, ϵ, via the elastic moduli, C_{ij},

$$\sigma_i = C_{ij}\epsilon_j, \quad i, j \in [1, ..., 6],$$

there will be five independent elasticity constants. Furthermore, if we focus solely on compressional-wave data, this number further reduces to four. In other words, only $C_{11}, C_{13}, C_{33}, C_{44}$ will be independent and need to be solved for in this particular case.

In the equations above, we have made reference to the ray angle, Θ. However, the angle that is measured in field observations is not the ray angle. In fact, in the presence of anisotropy there are three distinct angles associated with wave propagation which must be carefully distinguished. The first is the ray angle, Θ, which is associated with the direction of energy propagation. In the case of uniform media where raypaths are straight, it is equal to the angle between the vertical and the radian to a point on the wavefront (see Figure 2). If we were to {then} construct a perpendicular to the wavefront at end of the radius, the angle between the perpendicular and the vertical would be the phase angle, θ. Thirdly, we have the displacement (or polarization) angle, ϕ, which describes the direction of particle displacement as the wavefront propagates. This last angle is obtained by measuring three orthogonal displacements at the receiver. It can be shown that both the ray and polarization angles, Θ and ϕ, respectively, are related to the phase angle by the following expressions[4],

$$\tan(\Theta) = \frac{\tan(\theta) + \frac{1}{v}\frac{\partial v}{\partial \theta}}{1 - \frac{\tan(\theta)}{v}\frac{\partial v}{\partial \theta}} \tag{4.5}$$

---

[4]Arnold, V.I., (1989) Mathematical Methods of Classical Mechanics (2nd ed.): Springer-Verlag, pp. 248 - 252.

π

$$\tan(\phi) = \frac{\rho v^2 - C_{44} \sin^2(\theta) - C_{33} \cos^2(\theta)}{(C_{13} + C_{44}) \sin(\theta) \cos(\theta)} \tag{4.6}$$

Our goal is to formulate and solve an inverse problem, in which a, b, χ, X, Z, T,and ϕ are the inputs and the elasticity constants, $C_{11}, C_{13}, C_{33}, C_{44}$, of the layer of interest are the outputs.

4.3 Method of Solution

4.3.1 Solving the Forward Problem

Before we attempt to tackle the inverse problem, we consider the forward problem. By starting with the forward problem we will be able to generate input values for the inverse problem while at the same time knowing what outputs the inverse should give. We begin by fixing the values of a, b, and χ for the upper medium. Also, we set the refraction point, x, the mass density of the lower layer of interest, ρ, the elastic moduli, C_{ij}, and the total traveltime, T,from source to receiver. Using the equations above, we are then able find the refraction point, x, and the polarization angle, ϕ.

First, we calculate the horizontal ray slowness, p, in the upper medium by substituting the refraction point (x, H) into equation (4.3)

$$p(x) = \frac{2x}{\sqrt{(x^2 + (1 + 2\chi)H^2)(4a^2(1 + 2\chi) + 4ab(1 + 2\chi)H + b^2(x^2 + 1 + (1 + 2\chi)H^2}}$$

and also the traveltime in the upper medium using equation (4.4)

$$t_1(x, z; a, b, \chi) = \frac{1}{b} \ln \left| \frac{1 + \sqrt{1 - p(x)^2 a^2(1 + 2\chi)}}{1 + \sqrt{1 - p(x)^2(a + bz)(1 + 2\chi)}} \right|.$$

By Snell's law, the tangential component of phase slowness, p, will be conserved across the interface. If θ and $v(\theta)$ are the phase angle and phase velocity, respectfully, for the lower layer we have

$$\frac{\sin(\theta)}{v(\theta)} = p.$$

The phase velocity[5], $v(\theta)$, is given by

$$v^2(\theta) = \frac{(C_{33} - C_{11}) \cos^2(\theta) + C_{11} + C_{44} + \sqrt{\Delta}}{2\rho} \tag{4.7}$$

with

$$\Delta = [(C_{11} - C_{33}) \cos^2(\theta) - (C_{11} + C_{44})]^2$$
$$- 4[C_{33}C_{44} \cos^4(\theta) - (2C_{13}C_{44} - C_{11}C_{33} + C_{13}^2) \cos^2(\theta) \sin^2(\theta) + C_{11}C_{44} \sin^4(\theta)].$$

---

[5]Daley, P.F. & Hron, F., 1977. Reflection and Transmission Coefficients for Transversely Isotropic Media, *Bull. Seism. Soc. Am.*, **67**, 661-675.

$$\boxed{\pi}$$

At this point we can solve for the phase angle and thereby the phase velocity in the lower medium. Once we have the phase velocity, we can solve for the ray angle using equation (4.5),

$$\tan(\Theta) = \frac{\tan(\theta) + \frac{1}{v}\frac{\partial v}{\partial \theta}}{1 - \frac{\tan(\theta)}{v}\frac{\partial v}{\partial \theta}}.$$

We then substitute the phase angle into the ray velocity expression, can find the ray velocity, V, using

$$V^2(\theta) = v^2(\theta) + \left(\frac{\partial v}{\partial \theta}\right)^2. \tag{4.8}$$

Since the layer of interest is uniform, the ray path is a straight line and the ray velocity in the direction of the ray angle will be constant. Thus the traveltime, t_2, in the layer of interest is

$$t_2 = \frac{\sqrt{(X-x)^2 + (Z-H)^2}}{V(\Theta)} \tag{4.9}$$

and the total travel time, T, is

$$T = t_1(x) + t_2(x).$$

Substituting the phase angle into expression (4.6) we get the polarization angle, ϕ,

$$\tan(\phi) = \frac{\rho v^2 - C_{44}\sin^2(\theta) - C_{33}\cos^2(\theta)}{(C_{13} + C_{44})\sin(\theta)\cos(\theta)}.$$

4.3.2   Solving the Inverse Problem

Next we attempt to solve the inverse problem directly by combining the expressions described above into a system of equations. Since the rays will be straight in the lower medium we can reexpress the equations above to get the following system of nonlinear equations,

$$\frac{X-x}{Z-H} - \frac{\tan(\theta) + \frac{1}{v}\frac{\partial v}{\partial \theta}}{1 - \frac{\tan(\theta)}{v}\frac{\partial v}{\partial \theta}} = 0 \tag{4.10}$$

$$T - \frac{\sqrt{(X-x)^2 + (Z-H)^2}}{V(\theta)} - t_1(x) = 0$$

$$p(x)v(\theta; C_{ij}) - \sin(\theta) = 0$$

$$\tan(\phi) - \frac{\rho v^2 - C_{44}\sin^2(\theta) - C_{33}\cos^2(\theta)}{(C_{13} + C_{44})\sin(\theta)\cos(\theta)} = 0.$$

Unfortunately we have the problem of an underdetermined system with six unknowns $(x, \theta, C_{11}, C_{13}, C_{33}, C_{44})$ but only four equations. However, the industrial contacts on the team tell us that reasonable estimates of C_{33} and C_{44} are not completely out of reach. Therefore we proceed with the added the restriction that C_{33} and C_{44} are both known.

π

This system can now be solved using either Jacobi Iteration or the Conjugate-gradient Method. Results

As a starting point for the forward problem, we used the following the parameter values which are typical in actual settings,

$$
\begin{array}{llll}
a = 2000 & C_{11} = 3.13e10 & Z = 1154.46 \\
b = 0.3 & C_{13} = 0.34e10 & X = 1057.03 \\
\chi = 0.2 & C_{33} = 2.25e10 & H = 700. \\
\rho = 2310 & C_{44} = 0.65e10
\end{array}
$$

Using the forward problem, we generated the corresponding values of T, ϕ. Then, we used our inverse problem code to try to regenerate the values for $x, C_{11}, C_{13}, C_{33}, C_{44}$. For exact values of the input values, T, ϕ, the system converged to the correct values of C_{ij}. Unfortunately however, for even small perturbations of the inputs (consistent with experimental error) a solution was often not possible and when solutions were found, they tended to be widely different from the true values. Due to the time constraints of the workshop we were not able proceed with an in-depth error analysis. However, it was noted that provided the experimental errors were random (as opposed to systematic) a collection of outputs obtained by considering numerous source-receiver setups, could be used to estimate the true C_{ij}'s. By taking the best fit for each set of estimates of C_{ij}, the results tended to be reasonably close to the true values.

4.4   Future Work

Through our work on this problem we feel that our progress was marked by "raising more questions than answers". As a starting point for future or continued work on this project we would suggest that an in-depth error analysis be performed. By doing so, hidden redundancies in the formulations and suitable initial/boundary conditions may be identified to combat the sensitivity to errors and non uniqueness of solutions.

More importantly, however, we feel that the problem would be best approached by considering not one, but two sources on the surface (Figure 3). By doing so this would result in eight equations in eight unknowns, negating the need for *a priori* knowledge of C_{33} and C_{44}.

π
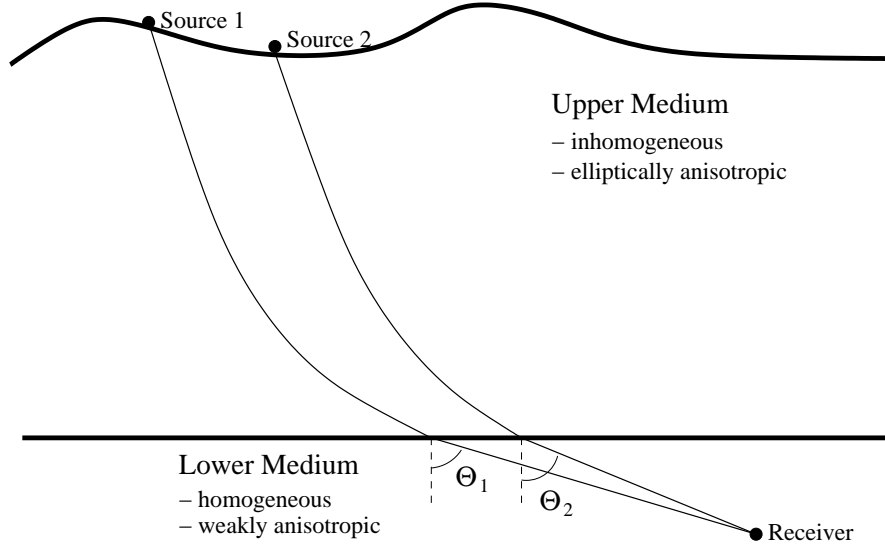
Adding a second source to equations (4.10) would give,

$$0 = \frac{X - x_1}{Z - H} - \frac{\tan(\theta_1) + \frac{1}{v_1}\frac{\partial v_1}{\partial \theta_1}}{1 - \frac{\tan(\theta_1)}{v_1}\frac{\partial v_1}{\partial \theta_1}} \tag{4.11}$$

$$0 = \frac{X - x_2}{Z - H} - \frac{\tan(\theta_2) + \frac{1}{v_2}\frac{\partial v_2}{\partial \theta_2}}{1 - \frac{\tan(\theta_2)}{v_2}\frac{\partial v_2}{\partial \theta_2}}$$

$$0 = T_1 - \frac{\sqrt{(X - x_1)^2 + (Z - H)^2}}{V_1(\theta_1)} - t_1(x_1)$$

$$0 = T_2 - \frac{\sqrt{(X - x_2)^2 + (Z - H)^2}}{V_2(\theta_2)} - t_2(x_2)$$

$$0 = p_1(x_1)v_1(\theta_1; C_{ij}) - \sin(\theta_1)$$

$$0 = p_2(x_2)v_2(\theta; C_{ij}) - \sin(\theta_2)$$

$$0 = \tan(\phi_1) - \frac{\rho v_1^2 - C_{44}\sin^2(\theta_1) - C_{33}\cos^2(\theta_1)}{(C_{13} + C_{44})\sin(\theta_1)\cos(\theta_1)}$$

$$0 = \tan(\phi_2) - \frac{\rho v_2^2 - C_{44}\sin^2(\theta_2) - C_{33}\cos^2(\theta_2)}{(C_{13} + C_{44})\sin(\theta_2)\cos(\theta_2)}$$

π

4.5 Sample File Used to Run Simulations

Throughout the coding of the equations presented above, a combination of Matlab^{TM} and Maple^{TM} code was used (depending on who owned the laptop we were using at the time). Following is the coding done using Maple ^{TM} to manipulate the input values as stated above. The first code was used to input values to be solved for in the inverse model, and generate data for the inverse model.

>restart:Digits:=40:

Forward Model where Cij's Known and Refraction Point (xr,zr)

Known.

>

># In this program we leave the Z of the receiver position a free parameter and solve for it in the process. We only set the horizontal offset of the receiver with respect to the refraction point, which we fix. The origin is assumed to be at the source.

># Enter in the input values and solve for receiver location, total traveltime, and polarization angle at receiver

>a:=2000: b:=0.8: chi:=0.3: zr:=800: roe:=2310: # zr=depth of interface/refraction point.

>C11:=3.13*10^10: C13:=0.34*10^10: C33:=2.25*10^10: C44:=0.65*10^10:

&&&&&&&&&&&&&&&&&&&&&&&&&&&&&&&&&

>X:=2600: xr:=1154.46:   # X=receiver depth, xr= refraction point.

&&&&&&&&&&&&&&&&&&&&&&&&&&&&&&&&&

>

># Traveltime in the upper medium

>p:=2*xr/sqrt((xr^2+(1+2*chi)*zr^2)*(a^2*(4+8*chi)+4*a*b*(1+2*chi)*zr+b^2*(xr^2+(1+2*chi)*zr^2))):

>t1:=(1/b)*ln((a+b*zr)/a * (1+sqrt(1-p^2*a^2*(1+2*chi)))/(1+sqrt(1-p^2*(a+b*zr)^2*(1+2*chi)))):

>

># Calculate the phase velocity and its derivative with respect to phase angle, theta

>Delta:=((C11-C33)*cos(theta)^2-(C11+C44))^2-4*(C33*C44*cos(theta)^4-(2*C13*C44-C11*C33+C13^2)*cos(theta)^2*sin(theta)^2+C11*C44*sin(theta)^4):

>v:= sqrt((C11+C44+(C33-C11)*cos(theta)^2+sqrt(Delta))/(2*roe)):

The Phase Angle for the ray traveling down from the refraction point (xr,zr):

will have problems if p*v >1. arcsin will then be complex (ie, Shadow Zone.) When you fix the refraction point, snell's law must have a solution for theta (a phase angle must exist). Since you are using p, conserved from upper medium, combined with v(theta) for the lower medium, if your refraction point is such that your angle of incidence is beyond the critical angle, snell's law will not have a solution in the lower medium because there will be NO TRANSMITION.

>eqn:=theta=arcsin(p*v):

>theta:=fsolve(eqn,theta,0..Pi/2);

>

># Solve for the Ray Angle from refraction point to receiver at (X,Z).

>rayang:=arctan((tan(theta)+1/v*dv)/(1-tan(theta)/v* dv)):

>rayangle:=evalf(%):

>

π

>\# Calculate the polarization angle along the ray in the lower medium

>phi:=arctan((roe*v^2-C44*sin(theta)^2-C33*cos(theta)^2)/((C13+C44)*sin(theta)*cos(theta))):

>

>\# Solve for the Ray velocity

>rayvelocity:=sqrt(v^2+dv^2):

Solve for the traveltime in the lower medium, t2.

——

Start by fixing the offset of x-coordinate of the receiver position from the refraction point, xr, than calculate the traveltime and z-position of receiver.

>xo:=X-xr:

——

>w:=xo/sin(rayangle):

>t2:=w/rayvelocity:

>Z:=zr+xo*cot(rayangle):

>T:=t1+t2:

==

Data to be used for the inverse problem is as follows,

C33, C44 Values;

>C33:=C33; C44:=C44;

Receiver location,

>X,Z:=X,Z;

Total Traveltime...

>TravelTime:=T;

And Polarizaton at the receiver is....

>phi:=phi; evalf(phi*180/Pi);

==

Inverse Code Should Give,

>xr:=xr; C11:=C11; C13:=C13;

The second code that is presented is a formulation of the inverse problem which relies on Maple^{TM}'s "fsolve" numerical solver to determine the values of the elastic moduli.

π

Inverse Model for Traveltimes and Polarization Angles for 1 Source / 1 Receiver Setup.

```
>

># List the Input Data for the inverse problem.

>restart:

>Digits:=40:

>a:=2000: b:=0.8: chi:=0.3: zr:=700: roe:=2310:

>C33:=2.25*10^10: C44:=0.65*10^10:

>X := 1057.0356:

>Z := 729.57398:

>T1 := .48321603:

>phi1 := 1.24456:

>p1:=2*xr1/sqrt((xr1^2+(1+2*chi)*zr^2)*(a^2*(4+8*chi)+4*a*b*(1+2*chi)*zr+b^2*(xr1^2+(1+2*chi)*zr^2))):

>Delta1:=((C11-C33)*cos(theta1)^2-(C11+C44))^2-4*(C33*C44*cos(theta1)^4-(2*C13*C44-C11*C33+C13^2)*
cos(theta1)^2*sin(theta1)^2+C11*C44*sin(theta1)^4):

>v1:=sqrt((C11+C44+(C33-C11)*cos(theta1)^2+sqrt(Delta1))/(2*roe)):

>dv1:=diff(v1,theta1):

>t1a:=(1/b)*ln(abs((a+b*zr)/a*(1+sqrt(1-p1^2*a^2*(1+2*chi)))/(1+sqrt(1-p1^2*(a+b*zr)^2*(1+2*chi))))):

># System of Equations to solve for C11,C13,C33,C44,xr1,theta1

>eqn1:=v1*p1-sin(theta1)=0:

>eqn2:=T1-sqrt(((X-xr1)^2+(Z-zr)^2)/(v1^2+dv1^2))-t1a=0:

>eqn3:= tan(phi1)=( (roe*v1^2-C44*sin(theta1)^2-C33*cos(theta1)^2)/((C13+C44)*sin(theta1)*cos(theta1))):

>eqn4:= (X-xr1)/(Z-zr)=((tan(theta1)+1/v1*dv1)/(1-tan(theta1)/v1*dv1)):

>sys:=[eqn1,eqn2,eqn3,eqn4]:

>sols:=fsolve({eqn1,eqn2,eqn3,eqn4},{xr1,C11,C13,C33,C44,theta1});

>assign(sols);

>xr1:=xr1; C11:=C11; C13:=C13; C33:=C33; C44:=C44;
```

$\boxed{\pi}$

Bibliography

[1] Affarwala, B., *et. al.*, (1997) Inversion of Anisotropic Velocity Parameter. *Proceedings of the first PIMS Industrial Solving Workshop*, PIMS, pp. 76-89.

[2] Arnold, V.I., (1989) Mathematical Methods of Classical Mechanincs (2nd edition): Springer - Verlag

[3] Cerveny, V., (2001) Seismic Ray Theory: Cambridge University Press.

[4] Daley, P.F. & Hron, F., (1977) Reflection and Transmission Coefficients for Transversely Isotropic Media, *Bull. Seism. Soc. Am.*, **67**, 661-675.

[5] Slawinski, M., (2002) Waves and Rays in Elastic Continua, in print.

[6] Thomsen, L.,(1986) Weak Elastic Anisotropy, *Geophysics*, **51**, pp. 1954 - 1966.

Chapter 5

Mathematical Surface Matching of Maps of the Human Torso

Rita Aggarwala[1], Ibrahim Agyemang[2], Erik Andries[3], Melanie Beck[4], Elaine Beltaos[2], Alicia Canton[5], Olivier Dubois[4], Charles Fortin[4], Ana Granados[5], Maria Landry[6], Edmond Lou[7], Matthias Mueck[8], Jeff Orchard[9], Marc Paulhus[10], Adam Webber[11],

Report prepared by Jeff Orchard[12] and Adam Webber[13]

5.1   Introduction

Scoliosis is a fairly common condition that usually appears during adolescence. It involves mild to severe deformations of the spine. Doctors and patients are interested in measuring the progression of the disease. Internal changes in spine angles and twists can be seen from X-rays, but surface (appearance) changes are often of primary concern to adolescent patients. Measuring the amount of surface change can be a difficult problem. Natural growth, changing body shape, and progression of the disease all factor into forming the shape of the patient's torso. Separating these effects, and extracting the change due to scoliosis alone, is challenging. In an effort to quantify the changes due to scoliosis, Capital Health acquires three-dimensional (3D) surface scans of a subject's back every 3 to 6 months using a laser scanner [1]. The problem addressed

[1]University of Calgary
[2]University of Alberta
[3]University of New Mexico
[4]McGill University
[5]University of Washington
[6]University of British Columbia
[7]Capital Health
[8]University of Toronto
[9]Simon Fraser University
[10]Pacific Institute for the Mathematical Sciences
[11]University of Victoria
[12]jjo@sfu.ca
[13]awebber@math.uvic.ca

in this paper is how best to use this data to obtain useful measures for both the doctor and patient.

There are two separate approaches to extract useful information from the data. The first involves making direct measurements on each individual scan and comparing these measurements over time. The second approach is to compare successive images of the back, noting the differences between them. The first approach is simpler, and some measurements of this type have already been implemented by Capital Health. However, these measures are extracted from a 2-dimensional (2D) representation of the data. It is hoped that measures that make use of the 3D data may be more informative. The second approach involves aligning two successive 3D back surfaces to each other. Capital Health currently implements this alignment operation manually [1], and thus measurements taken from these alignments are not completely reproducible due to human error.

Two main issues were taken into account in this project: the reproducibility of the results, and stringent time constraints. In particular, user influence should be removed from each step of the data processing, and results should be obtained within three minutes of acquiring the scan.

The report begins with a description of the data collection, followed by a description of the data processing required to align two back surfaces. A section is devoted to calculating the cosmetic score, a measure of deformity of the back. The paper concludes with a few suggestions for improvements on data collection and use.

5.2 Experimental design and data collection

Patients with scoliosis are examined every 3 to 6 months. Laser surface scanning is a part of this examination, and the entire procedure takes only a few minutes.

Before the scan takes place, several steps are followed. Patients are dressed in a gown and their shoes are removed. The patient is asked to stand normally within a standard frame apparatus, their arms extended to 90° holding onto two poles. The frame has four adjustable guides: one on each hip, and one aligned with the front of each shoulder. The function of the guides is simply to assist the patient in holding the same position during the scan. Although the guides are not used as a measuring device, the nurse records the guide settings.

The laser scanner is approximately 2 metres away from the patient. Once the patient is positioned in the frame, the nurse adjusts the field of view of the camera by selecting one of three zoom settings. The normal setting is 0.72, with 0.6 and 0.8 being used occasionally. To minimize patient motion due to breathing, the patient is asked to take a deep breath and hold it during the scan. The scan takes 0.6 seconds, so patient movement within a single scan is minimal.

For some scans, the nurse places 5 fiducial markers on the patient's back before the patient is placed in the frame: one on the middle of the back of the neck, one on each scapula (at the inferior tip of the shoulder blade), and one on each lower back dimple (the slight indentations on either side of the lumbar spine). Since the markers are placed on easily-identifiable physiological landmarks, we expect their placement to be consistent.

Each scan results in the collection of two different sets of data. The first is a TIFF image,

a digitized photograph of the patient's back. This image contains colour information as well as position information. The second set of data is a cloud of 3D data points (approximately 200×200) that corresponds directly to the TIFF image. Although the scanner is designed to detect the exposed back and arms of the patient, it also tends to detect several other objects within the scanner's range. Automatically stripping the raw data of these extraneous structures was one of the main challenges faced by our group.

5.3 Cropping

The first processing step on the data is to select the portion of data that is relevant for analysis. Each back image (here "image" refers to a set of 3D coordinates) not only includes points describing the surface of the patient's back, but also some points corresponding to the frame, the patient's head, clothing, etc. These points are to be removed by a *cropping* procedure described below.

Our cropping procedure involves only the first two coordinates of each point. As such, it is a 2D cropping procedure. We will denote the first two coordinates by x and y, respectively. Figure 5.1 illustrates an example of an initial (uncropped) image viewed in the xy-plane.
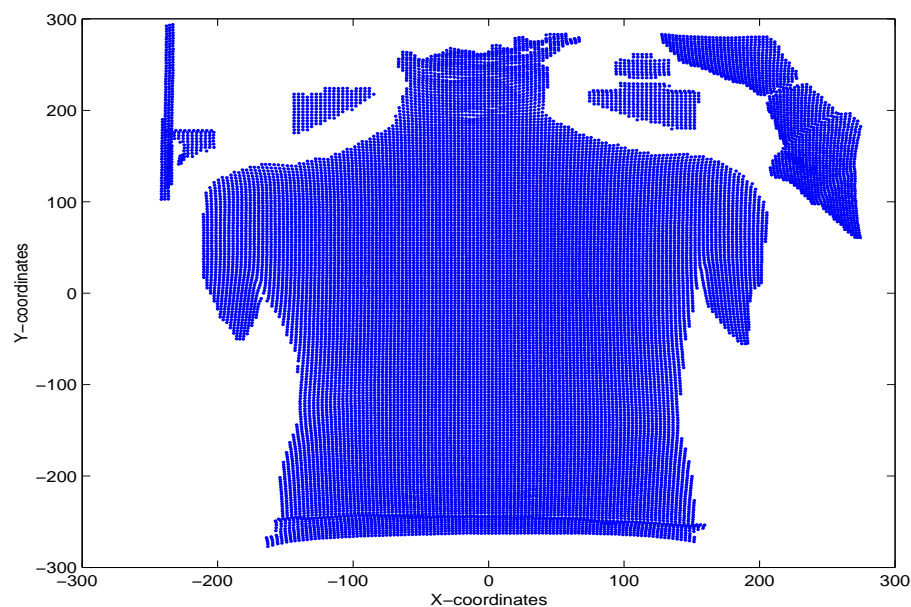
Figure 5.1: Surface data before cropping. This broken 3D surface is projected onto the xy-plane.

The first step of the cropping algorithm is to eliminate the clouds of points which are not connected to the body of the patient. Two points are not connected if there is a large enough space between them where there is no data point. Hence, we need to define what *large enough* means. Let $Y_{\max}$ be the largest Y-coordinate among the points in the data set. Similarly, define $Y_{\min}$, $X_{\max}$ and $X_{\min}$. Now define

$$\epsilon_x := \frac{X_{\max} - X_{\min}}{130} \quad \text{and} \quad \epsilon_y := \frac{Y_{\max} - Y_{\min}}{130} \, .$$

π

We will say that two sets of points are disconnected from each other if there is an empty rectangle of size $\epsilon_x \times \epsilon_y$ between them.

Although the image may not always be centered, the axis $x = 0$ always runs through the back. We search for a point (x, y) which satisfies

$$0 \leq x \leq \epsilon_x , \ Y_{\min} \leq y \leq Y_{\min} + \epsilon_y .$$

Once we find one, we look to the right of this rectangle to see if we can find a point (x, y) which satisfies

$$\epsilon_x \leq x \leq 2\epsilon_x , \ Y_{\min} \leq y \leq Y_{\min} + \epsilon_y .$$

We continue moving to the right by steps of size ϵ_x until we find a rectangle containing no data points. When this occurs, we eliminate all the points which are further to the right of this empty rectangle. The y-value is incremented by ϵ_y, and the procedure is repeated on the next line (i.e. $Y_{\min} + \epsilon_y \leq y \leq Y_{\min} + 2\epsilon_y$). Lines are processed one at a time until $Y_{\max}$ is reached. Then, we repeat all the previous steps, this time moving to the left of $x = 0$, i.e for the points satisfying $x \leq 0$. Figure 5.2 shows the dataset from figure 5.1 after cropping.
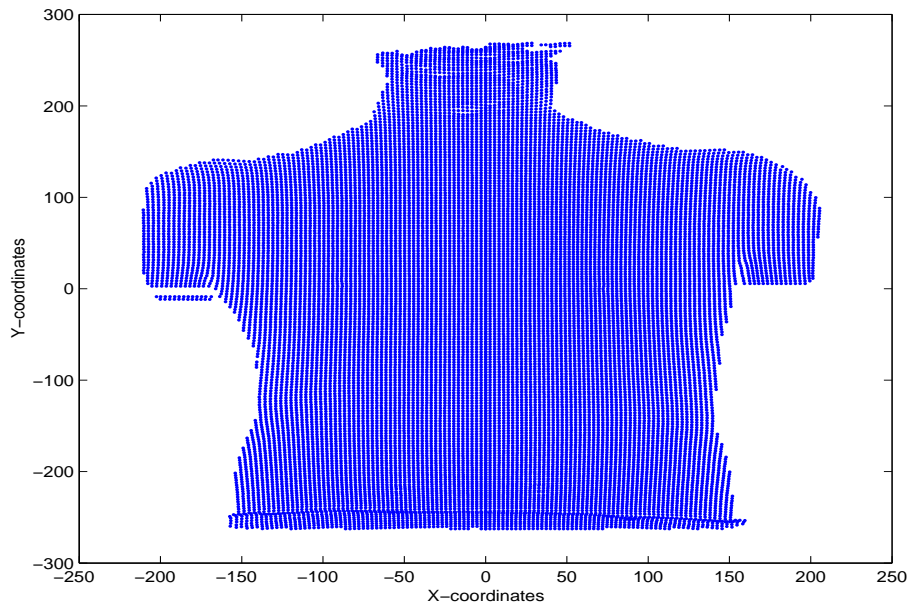
Figure 5.2: Surface data after the first cropping stage. The remaining piece is a single connected component.

At this point, we consider the dataset to be a connected cloud of points. However, it still contains points corresponding to the neck of the patient. To eliminate these points, we simply clear all rows that are shorter than 1/3 of the width at the shoulders. This finalizes our cropping procedure. Figure 5.3 illustrates the final result for our example.
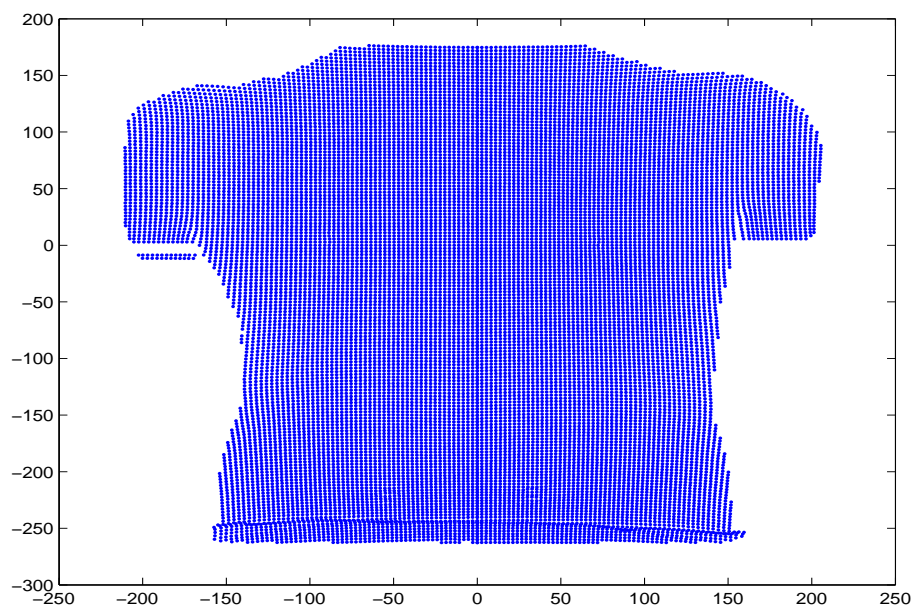
Figure 5.3: Surface data after the final cropping stage. The protruding neck is no longer present.

5.4 Initial Surface Alignment

Recall that our analysis of the progression of scoliosis involves the comparison between two scans acquired at different times, typically 3 to 6 months apart. Once the two images are cropped, we need to align them to allow direct comparison. A good initial guess is essential to get convergence of the alignment fine-tuning phase (see section 5.5). In each alignment phase, we seek an $\mathbb{R}^3 \to \mathbb{R}^3$ coordinate transformation that best aligns one image to the other. To simplify the problem, we consider only a subset of all possible affine maps, along with translations. This simplification is valid because the degrees of freedom for patient motion are limited by the effects of normal growth, the positional frame, and the nature of scoliosis. The set of affine maps that we will consider for the initial alignment are those that are a composition of rotations in the xy-plane (about the z-axis), dilatations for changes in growth and width (scaling along the y-axis and x-axis respectively), and translations in directions x, y and z.

Keeping one of the images fixed, the second is transformed via this affine map to best fit the first image. To establish an initial estimate of the appropriate affine map, we can use the coordinates of corresponding points from each of the two images. That is, each location on the patient's back is sampled in both scans. However, the coordinates of this point may be different in each scan. The affine map is chosen to best match these coordinates.

Since the patient is standing upright and positioned in a fixed frame, it can be assumed that the rotation in the xy-plane is very small, and therefore can be approximated by a y-axis shear in the direction of the x-axis. This simplification enables one to represent the affine map (using

homogeneous coordinates) with the matrix,

$$M = \begin{pmatrix} \nu_1 & \tilde{a} & 0 & t_1 \\ 0 & \nu_2 & 0 & t_2 \\ 0 & 0 & 1 & t_3 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

where ν_1 and ν_2 are the x and y scaling parameters, t_1, t_2 and t_3 are translation parameters, and $\tilde{a} = a\nu_1$, where a is the shear coefficient.

The system of linear equations given by the matrix M can be rewritten as

$$\begin{pmatrix} \nu_1 & \tilde{a} & 0 \\ 0 & \nu_2 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} + \begin{pmatrix} t_1 \\ t_2 \\ t_3 \end{pmatrix} = \begin{pmatrix} x' \\ y' \\ z' \end{pmatrix}.$$

Since the unknown parameters are $\nu_1, \nu_2, \tilde{a}, t_1, t_2$ and t_3, we transform this system into

$$\begin{pmatrix} x & 0 & y & 1 & 0 & 0 \\ 0 & y & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \tilde{a} \\ t_1 \\ t_2 \\ t_3 \end{pmatrix} = \begin{pmatrix} x' \\ y' \\ z' - z \end{pmatrix}.$$

We have six unknowns, so we need at least six equations. Since the third row of the matrix does not depend on (x, y, z), it can appear only once (otherwise, a singular system will result). Hence, three reference points are needed to estimate all the motion parameters. If we denote by (x_1, y_1, z_1), (x_2, y_2, z_2) and (x_3, y_3, z_3) the coordinates of these three points in the first scan, and (x_1', y_1', z_1'), (x_2', y_2', z_2') and (x_3', y_3', z_3') the coordinates of the three corresponding points in the second scan, we obtain the following system:

$$\begin{pmatrix} x_1 & 0 & y_1 & 1 & 0 & 0 \\ x_2 & 0 & y_2 & 1 & 0 & 0 \\ x_3 & 0 & y_3 & 1 & 0 & 0 \\ 0 & y_1 & 0 & 0 & 1 & 0 \\ 0 & y_2 & 0 & 0 & 1 & 0 \\ 0 & y_3 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \tilde{a} \\ t_1 \\ t_2 \\ t_3 \end{pmatrix} = \begin{pmatrix} x_1' \\ x_2' \\ x_3' \\ y_1' \\ y_2' \\ y_3' \\ z_1' - z_1 \end{pmatrix}.$$

This system is over-determined with seven equations for six unknowns. However, Matlab automatically returns the linear least squares solution. Thus we obtain directly the six parameters that give the best initial fit.

To get a good initial guess, it is very important that the three points be located far apart. One option for locating corresponding points in a pair of scans is to locate the fiducial markers. A Matlab program was developed to automatically detect the position of the markers located on the dimples. The markers used for the scans available to us were (bolt) washers, taped onto the
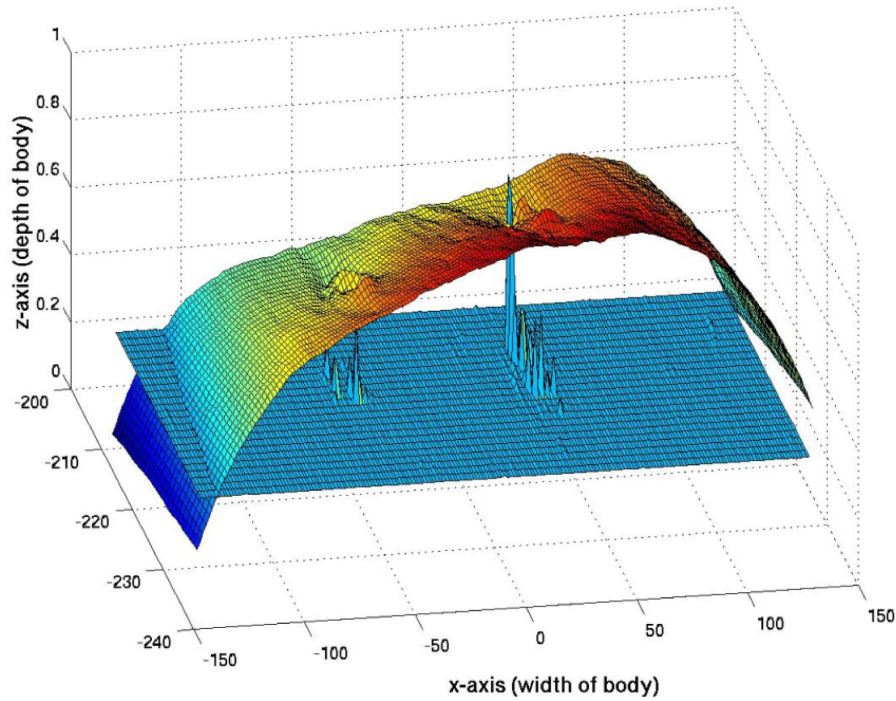
Figure 5.4: Automatic marker detection on the lower back. The spikes in the nearly-flat surface indicate the location of the markers on the back surface (also shown).

patient's skin. These markers are difficult to detect in the 3D surface by visual inspection since they appear as minuscule Heaviside steps in the z direction. However, since they are attached in regions where the torso surface is relatively smooth, they can be detected easily by applying a two-dimensional second-order finite difference stencil,

$$\Delta z(x, y) = z(x - h, y) + z(x + h, y) + z(x, y - h) + z(x, y + h) - 4z(x, y) \ ,$$

to the regions on the torso where the markers are likely to be found. The first difference eliminates a linear trend, while the second difference eliminates a quadratic trend in the data. Since the washers manifest themselves in the 3D cloud of data as annular Heaviside functions,

$$f(x, y) = H\left(r_2 - (x - x_0)^2 - (y - y_0)^2\right) - H\left(r_1 - (x - x_0)^2 - (y - y_0)^2\right) \ ,$$

the second-order differencing eliminates the local quadratic curvature of the torso but does not eliminate the washer-induced Heaviside *roughness*. Hence, by plotting $|\Delta z(x, y)|$ across the relevant regions, the presence of the markers is indicated by large spikes in curvature amplitude. Figure 5.4 shows a portion of the lower back surface, along with $|\Delta z(x, y)|$.

Although the above method for finding corresponding points appears to be effective, points were selected manually for this preliminary study. Two points were selected in the outer part of the shoulders, and one point was selected in the middle of the lower back (between the two dimples).

To illustrate the initial fitting, figure 5.5 shows the initial alignment of two surfaces for the same patient, acquired a few months apart.
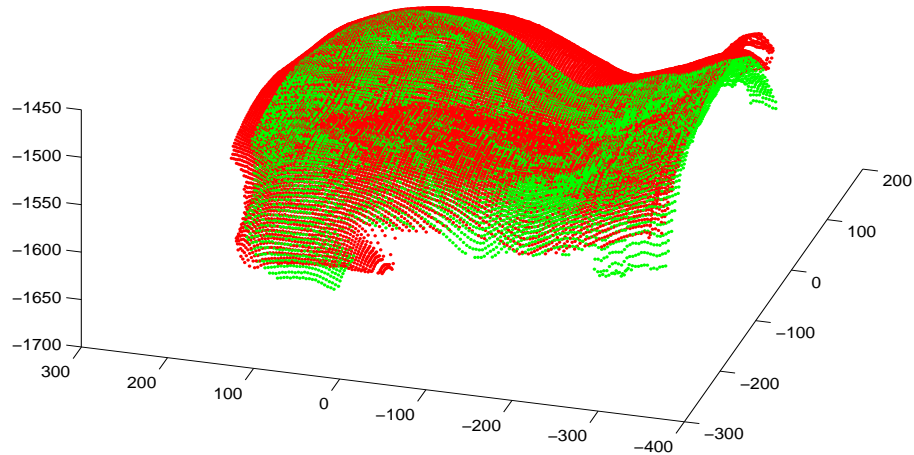
π

Figure 5.5: Initial alignment of two scans acquired a few months apart. The red surface was aligned to the green surface using the 3-point method outlined in section 5.4.

5.5 Final Surface Alignment

Fine-tuning of the alignment between the two back surfaces is done by adjusting the alignment parameters to minimize a given cost function. Such automatic surface matching has been demonstrated on teeth [2]. Some implementation decisions were made based on the feasibility of performing the optimization on the computers available to us during the workshop.

The automatic re-alignment (registration) method consists of 3 main parts: transformation, cost function, and optimization scheme.

5.5.1 Transformation

The human body is capable of moving into many different positions. The progression of scoliosis accounts for some particular kinds of motion. However, the objective of registration is to remove the effects of body position that are not due to scoliosis. The manual re-alignment process was used to correct a certain subset of typical mis-alignments [1]. In particular, manual alignment allowed the operator to scale, rotate and translate the surfaces. Once the scale factors, rotations and translations are chosen, applying the transformation to a cloud of 3D data points is trivial. However, comparing two clouds of data points that are not sampled at the same locations is quite difficult. Initially, we compared the surfaces by resampling one surface onto the same grid as the other using Matlab's `griddata` function. The `griddata` function performs linear interpolation using a Delaunay triangulation. This method proved to be too slow, since the automatic optimization method required many such comparisons.

Instead of treating the mesh as a general 3D cloud of data points, we converted it to a gray-level image. Such an image uses gray-level to indicate the height of the surface over

the corresponding (x, y) point. Since the surface has a bounded xy domain, a value of zero is assumed anywhere that the surface was not defined. One of the main differences between representing the data as an image instead of a point cloud is that the point cloud data need not be a function, since it can have two points with the same (x, y) coordinates, but different z values. An image can only represent a function. This does not appear to be a major restriction since the original data is range data (which has to be a function), and the alignment rotations tend to be rather small. The advantage to using the image representation is that there are very efficient resampling techniques for regularly gridded image data.

Similar to the subset of affine maps used in section 5.4, the set of transformations used in the final alignment algorithm has 6 parameters (although they are not the same 6 parameters as in section 5.4): translation along the x- and y-axes, scaling along the x- and y-axes, rotation about the z-axis, and rotation about the x-axis. The rotation about the z-axis, the translations and the stretches are all simple 2D image manipulations. However, the rotation about the x-axis is not. Since we expect the x-axis rotation to be quite small, we can approximate it with a shear in the z-direction along the y-axis. This operation is equivalent to adding a y-gradient to the image. All of these operations are essentially 2D, and can be done quickly using built-in Matlab routines.

In this 2D image-based transformation framework, we still use `griddata` to obtain the initial image representation from the cloud of points. All subsequent manipulations are done on the 2D image. One issue, however, is that `griddata` assumes the xy-domain of the surface of interest is the convex hull of all the points. This is a natural assumption to make in general, but does not apply in this context. Figure 5.6 shows the image representation of a torso surface. Notice that the region inside the line joining the waist and the outer edge of the arms is not treated as background, even though it is clearly not part of the back (see figure 5.3 for comparison). This is a property of the interpolation method used in `griddata`. Further pre-processing to insert background points in that region may solve the problem. However, for this study, the full convex hull was used.

5.5.2   Cost Function

A measure of goodness-of-fit must be established in order to perform automatic alignment. For this study, we chose to use the sum of squared residuals between the image representations of the two back surfaces. There are many other cost functions that can be utilized, but none others were implemented for this report.

It should be noted that the sum of squares cost function is sensitive to outliers. The original range data was a cloud of points with z-values between 1400 and 1700. Using these values directly in the gray-level images caused the background to have a very different value from the foreground. As a result, non-overlapping regions between the convex hulls of the two images resulted in extreme outliers, causing very large errors. Thus, the fit was dominated by the amount of overlap. Since the background value of zero was chosen arbitrarily, choosing a different background value could reduce this effect. Equivalently, translating both surfaces so that their minimum z-value is mapped to zero vastly reduces the drastic impact of the distant background. This practice was adopted for our study.
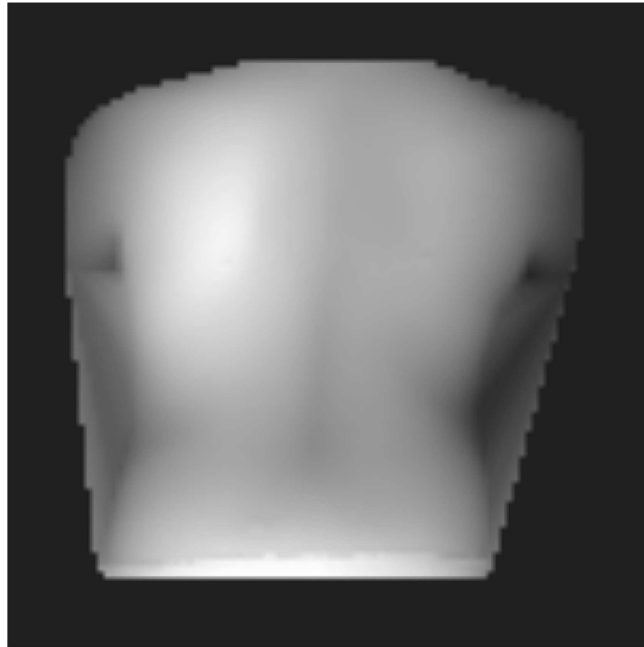
π

Figure 5.6: Image representation of back surface data. The image has been slightly enhanced to show the nonzero regions between the arms and the waist.

5.5.3   Optimization

The procedure for finding the optimal set of 6 motion parameters was off-loaded to Matlab's optimization toolbox. In a first attempt, we fed the cost function to the unconstrained minimization procedure, `fminsearch`. However, because of the number of parameters, the optimization was very slow and did not always converge to the expected minimum.

To restrict the search, it is easy to find bounds for the motion parameters, and so we used the routine `fmincon` instead. Since we have already performed an initial alignment, the optimal value for each parameter in the final adjustment should be very small. For example, we required that the translations be less than half the range of the data in each direction. The rotations should be limited by 45 degrees. As for the shear corresponding to the forward bending of the patient, we have derived bounds by assuming that the depth of each point should not change by more than the z-range of the data. These bounds are quite loose and could be made tighter.

Figure 5.7 shows the difference between the two back images before and after final alignment. The degree of misalignment is noticeably reduced after the final alignment phase.

5.6   Cosmetic Score

Although the severity of cosmetic deformation may not always be highly correlated with the severity of spinal deformation, quantization of the cosmetic severity is sought by patients making adjustments to living in their society. Investigation into the relationship between the cosmetic deformation and the spinal deformation is not a focus of the present study, though it may prove
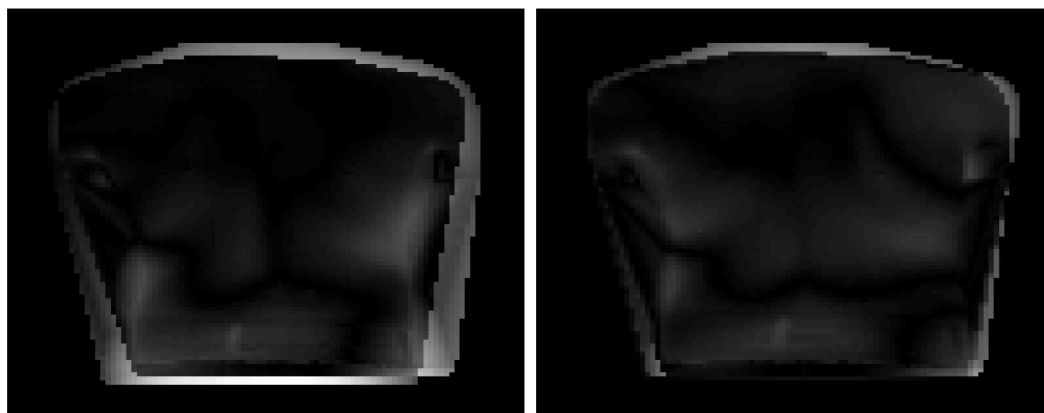
Figure 5.7: Difference images before (left) and after (right) final alignment. The lighter regions show areas of misalignment.

to be useful in future work.

A single, dimensionless number known as the *2-dimensional cosmetic score* is currently calculated as a weighted average of:

- the (scaled) angle between a line joining the two scapula bones (where markers are placed) and a line parallel to the ground

- the (scaled) difference between the angles that the right and left shoulder silhouettes make with a line parallel to the ground

- the waist asymmetry index calculated as a function of several different distances measured from markers to a vertical center line (also determined from markers).

A high cosmetic score is associated with more severe cosmetic deformation.

The dimensionless property of the cosmetic score is appealing, as it is thought to be invariant to an individual's growth over time and positioning in the frame apparatus, and thus the surface alignment issues discussed in previous sections are not of concern here. The simple score is based on the measurement of seven locators on the back, and can be automated easily using, for example, reflective markers that can be detected once cropping has been completed.

The current 2D cosmetic score does not take into account twisting or related depth information available through the 3D laser scan. Generalizing the 2D cosmetic score to three dimensions, while preserving the invariance property of the score, is simple. Possible measurements that may be included in the calculation of a 3D cosmetic score are:

- the twist angle between the line connecting the scapula markers and the line connecting shoulders

- the twist angle between the line connecting the scapula markers and the line connecting dimples in the lower back

- the twist angle between the line connecting shoulders and the line connecting dimples in the lower back.

Since marker position can be obtained in three dimensional Euclidean space, these angles can be measured easily and incorporated into a generalized 3D cosmetic score. Furthermore, automated marker localization will make the calculation of this score fast and reproducible. Additional measures of importance may be considered with clinical input from medical experts familiar with the disease.

Perhaps a more interesting and revealing exploration is a comparison between the current 2D score and a generalized 3D score. A "perfect" 2D score can obviously be associated with an imperfect 3D score. Also of interest is be the predictive ability of the final cosmetic score. Determining the formula for the score requires not only determining which measurements are most important in describing the severity of cosmetic deformation, but also verification by an experiment that can incorporate errors associated with operator variation, repeated measurements over time, and so on. Finally, an investigation comparing the cosmetic scores for scoliosis patients with normal (scoliosis-free) and post-operative individuals will be informative and important in determining the utility of the ideas currently in use as well as those being presented here.

5.7   Conclusions

We have shown that, in principle, the data processing required to crop and align the surfaces can be done automatically, and within the 3-minute time window, as required. However, there are simple procedural modifications that can make this processing more robust and practical.

The scanner zoom settings are not standardized between visits. In the future, the scanner settings should be recorded and then reused in successive visits, minimizing the need for the alignment algorithm to accommodate for large scale changes.

The markers should be designed so that they are more easily detected in the surface data. Markers with a deeper cross-section, or with reflective properties could be used to create a very distinct and noticeable bump on the 3D surface. Once these markers can be found automatically, efficiently, and robustly, they should be used to automate the initial alignment procedure.

In addition to what was demonstrated in this study, further study using more degrees of freedom in the motion model may yield important results. A particular subset of motion parameters (eg. a twisting along the spine, or y-axis) could be used to model the deformation caused by scoliosis, thereby enabling the alignment algorithm to quantify the progression of scoliosis with motion parameters. It should be noted, however, that the particular types of motion used in this study could be implemented on the 2D image representation of the back data. A deformation such as twisting could be more difficult to implement on the 2D images.

Data from various physiological features and the location of markers might be useful for calculating a 3D cosmetic score that represents the severity of observable deformity. This 3D score will likely contain more specific information than the currently used 2D score. A study should be conducted to evaluate various forms of the 3D score and determine what particular mixture of measurements best encapsulates the deformity, and is least susceptible to measurement/operator error.

π

Bibliography

[1] D. C. Berg, D. L. Hill, J. V. Raso, E. Lou, T. Church, M. J. Moreau, J. K. Mahood. Using three-dimensional difference maps to assess changes in scoliotic deformities. *Medical & Biological Engineering & Computing*, Vol. 40, February 2002, pp. 1-6.

[2] H. L. Mitchell, R. G. Chadwick. Mathematical shape matching as a tool in tooth wear assessment - development and conduct. *Journal of Oral Rehabilitation*, Vol. 25, 1998, pp. 921-928.

Chapter 6

Resistance Monitoring

Lyudmyla Barannyk[1], Matthew Bolton[3], Roger Donaldson[3], Ian Frigaard[3], Jeff Gilmore[4], Ying Han[5], Huaxiong Huang[6], Heather Lehr[7], Bruce McGee[13], Anuj Mubayi[8], Tim Myers[9], Anthony Pierce[3], Bidhan Roy[10], Tetyana Segin[1], Quinze Zou[11]

Report prepared by Rex Westbrook[12] & C. Sean Bohun[2].

6.1 Introduction

MacMillan McGee Corp. is involved in the removal of soil contaminants using a process in which a set of electrodes is placed in the soil and the soil is heated by passing a current through it. The electrical conductivity of the soil, and hence the heat generated, is enhanced by pumping heated water in at some of the electrodes and removing it at a well (or set of wells). The heated water also provides some convected heating, which in the case of soil, far exceeds the conductive heating.

   The problem considered was that of estimating the temperature field in the region of interest, the contaminated region, using measurements of electrical potential and current and also of temperature, at accessible points such as the wells and electrodes and the soil surface. These measurements in conjunction with the field equations for the electrical potential and temperature provide a mathematical "Inverse Problem". There is a known empirical relationship between

[1]New Jersey Institute of Technology
[2]The Pennsylvania State University
[3]University of British Columbia
[4]Simon Fraser University
[5]McGill University
[6]York University
[7]University of Texas at Austin
[8]University of Texas at Arlington
[9]University of Cape Town
[10]University of Illinois
[11]University of Washington
[12]University of Calgary
[13]McMillan-McGee Corporation

temperature and electrical conductivity so that knowledge of one at any given place and time also determines the other.

On the timescale considered, essentially days, the equation for the electrical potential is static. At any given time the potential V satisfies the equation $\nabla \cdot (\sigma \nabla V) = 0$. Time enters the equation only as a parameter since σ is temperature and hence time dependent.

The problem of finding σ when both the potential V and the current density $\sigma \partial V / \partial n$ are known on the boundary of the domain is a standard inverse problem of long standing. See for example Kohn and Vogelius [4]. It is known, with some restrictions on smoothness, that a unique solution σ exists in spaces of dimension $n \geq 2$. (In one dimension only the harmonic mean of σ can be determined). It is also known that the problem is ill posed and hence that an accurate numerical solution will be difficult especially when the input data is subject to measurement errors. There is much literature on both theoretical and numerical aspects of the problem but we were not able to find a recent comprehensive survey.

In Section 6.5 we examine a possible method for solving the electrical inverse problem which could possibly be used in a time stepping algorithm when the conductivity changes little in each step.

Since we are also able to make temperature measurements there is also the possibility of examining an inverse problem for the temperature equation. There seems to be much less literature on this problem, which in our case is essentially, a first order equation with a heat source. (We neglect thermal conductivity, which is small compared with the convection). Combining the results of both inverse problems might give a more robust method of estimating the temperature in the soil.

Removing soil contaminants using the combined techniques of electrical heating and fluid injection changes the medium in a predictable fashion. On occasion, this method of contaminant removal breaks down and the properties of the medium itself can suffer a sudden change. It is important to be able to detect not only if this situation arises, but also where in the medium the method breaks down.

The equations for the problem are given in the following section.

6.2   Equations

The basic equations of our problem are an energy equation, a quasi-static electrical equation and a porous medium equation. The electrical equation for the potential V is:

$$\nabla \cdot (\sigma \nabla V) = 0 \qquad (6.1)$$

where σ is the electrical conductivity. The equations for the fluid flow, D'Arcy's law and incompressibility, are:

$$\vec{u} = -\frac{k}{\mu} \nabla p, \quad \nabla \cdot \vec{u} = 0 \qquad (6.2)$$

where $\vec{u}$ is the fluid velocity, p is the pressure, μ is the kinematic viscosity and k is the permeabilty. For the energy balance with Joule heating as a source term we have [2]:

$$\frac{\partial}{\partial t}(\rho_a c_a T) + \vec{u} \cdot \nabla(\rho_w c_w T) = \nabla \cdot (\kappa_a \nabla T) + \sigma \, |\nabla V|^2 \qquad (6.3)$$

π

$\rho_a c_a$	2.5×10^6 J/m^3/°C
$\rho_w c_w$	4.2×10^6 J/m^3/°C
κ_a	1 J/s/m/C
T_0	100°C
T	$10 - 100$°C
V	$100 - 200$ volts
σ	$0.1 - 0.01$ /Ω/m
$|\vec{u}|$	0.3 m/day
ϕ	0.3
S_w	$0.5 - 1$
a	3 /°C

Table 6.1: Some typical numerical values for this problem. Note that 1 watt sec = 1 Joule and that for this problem, a typical length is 10 m. A typical time for the process is 60 days.

where T is the temperature, ρ is the density, c is the specific heat, κ is the thermal conductivity, with the subscripts a and w denoting average and water respectively. The average values are computed as $\rho_a c_a = \phi \rho_w c_w + (1 - \phi)\rho_s c_s$, $\kappa_a = \phi \kappa_w + (1 - \phi)\kappa_s$ where the subscript s is for soil and ϕ is the volume fraction of water.

In addition to these equations we have an empirical relation between the conductivity σ and the temperature T

$$\sigma = \sigma_w \frac{\phi^{\alpha_n}}{P_\alpha} S_w^2 f(T) + (1 - \phi)\sigma_s \tag{6.4}$$

where S_w is the water saturation, σ_s is a constant depending on the soil type, and α_n and P_α are empirically obtained constants. The function $f(T)$ is found to be linearly dominated and is taken to be $f(T) = 1 + a(T - T_0)$. Typically f increases by a factor of three for a 100°C increase in temperature from T_0. Some typical numerical values, supplied by Bruce McGee, are indicated in Table 6.1.

We now scale the energy equation (6.3) by writing

$$\vec{x}' = \vec{x}/L, \quad t' = t/\tau, \quad \vec{u}' = \vec{u}/U, \quad \sigma' = \sigma/\sigma_0, \quad V' = V/V_0, \quad T' = (T - T_0)/\triangle T$$

to obtain the scaled equation (dropping primes):

$$\lambda \frac{\partial}{\partial t}T + \vec{u} \cdot \nabla(T) = \varepsilon \nabla \cdot (\nabla T) + \gamma \sigma |\nabla V|^2 \tag{6.5}$$

where λ, ε, γ are constants given by

$$\lambda = \frac{\rho_a c_a L}{U \tau \rho_w c_w}, \quad \varepsilon = \frac{\kappa_a}{U L \rho_w c_w}, \quad \gamma = \frac{\sigma_0 V_0^2}{U L \rho_w c_w \triangle T}.$$

If we take $L = 10$ m, $\tau = 20$ days, $U = 0.3$ m/day, $V_0 = 100$ volts, $\sigma_0 = 0.1$ /Ω/m and $\triangle T = 20$°C we find,

$$\lambda \simeq 1, \quad \varepsilon \simeq 0.7 \times 10^{-2}, \quad \gamma \simeq 1/3.$$

π

Also with these scalings $f(T) = 1 + \alpha T$ with $\alpha = a\triangle T$ and we write $\sigma = \widetilde{\sigma} f(T)$ with $\widetilde{\sigma}$ and α being $O(1)$.

Boundary conditions for the problem are that the potential or current flow is given at all electrodes and at the soil surface. Both being zero at infinity. For the fluid, pressure or velocity is given at input and output wells. For the temperature we assume a known temperature at input wells.

The small value of ε indicates the dominance of convection over conduction in the equations and neglecting conduction gives a first order equation for T. This means that the temperature along the characteristics is governed by either an initial temperature or a fluid temperature at a fluid input. If these temperatures are not compatible, thermal shocks could appear. The small conduction would smooth these shocks in an internal boundary layer of thickness $\varepsilon^{1/2}$. Boundary layers might also occur at output wells but we avoid these by assuming that the temperature at output wells is that of the emerging fluid. In other words we do not apply temperature boundary conditions at output wells.

6.3   A one dimensional example

We consider a one-dimensional problem with two electrodes that coincide with the input and output wells. As boundary conditions we take constant current at the electrodes and a given temperature at the input well. We will assume constant scaled ambient temperature zero as the initial condition. In the one-dimensional case the velocity vector $\vec{u}$ becomes a constant which we take to be 1. We also assume that $\widetilde{\sigma}$ and α are constant.

The potential equation:

$$\frac{\partial}{\partial x}\left(\sigma\frac{\partial V}{\partial x}\right) = 0 \tag{6.6}$$

may be integrated to give $I = \sigma\partial V/\partial x = I(t)$ a function of time only. With the given boundary condition I is constant in both space and time. To find V we must first solve for T to obtain σ. V may then be found by another integration.

With I constant, the scaled energy equation (6.5) becomes:

$$\frac{\partial T}{\partial t} + \frac{\partial T}{\partial x} = \frac{I^2}{3\widetilde{\sigma}(1 + \alpha T)}. \tag{6.7}$$

This equation may be solved by the method of characteristics. We solve the equation on the domain $t > 0$, $0 < x < L$. The boundary conditions are $T = 0$ on $t = 0$, (ambient physical temperature T_0) $T = g(t)$, $t > 0$ on $x = 0$. The characteristic equations are:

$$\frac{dt}{ds} = 1, \ \frac{dx}{ds} = 1, \ \frac{dT}{ds} = \frac{K}{1 + \alpha T} \ \text{ where } \ K = \frac{1}{3}\frac{I^2}{\widetilde{\sigma}}. \tag{6.8}$$

These equations must be solved in two distinct regimes. Fluid already in the soil will be initially at temperature $T = 0$. This fluid moves with unit speed and at any particular time, since all the fluid will have been subjected to the same amount of heating, the temperature will be the same. The appropriate initial conditions for a particle initially at x_0 are, $t = 0$, $T = 0$,

$x = x_0$. Since the input well is at $x = 0$ there will be a front at $x = t$, and this solution will apply for $x > t$. Fluid that has been injected at the input well at time t_i will enter the system with a temperature $T = g(t_i)$, and its temperature will increase as it moves away from the well. This fluid will remain behind the front and boundary conditions are $t = t_i$, $T = g(t_i)$, $x = 0$. The solution is:

$$T(x,t) = \begin{cases} \dfrac{1}{\alpha}(\sqrt{1 + 2\alpha K t} - 1), & x > t \\ \dfrac{1}{\alpha}\left[\sqrt{(1 + \alpha g(t - x))^2 + 2\alpha K x} - 1\right], & x < t. \end{cases} \tag{6.9}$$

If $g(0) \neq 0$, T will be discontinuous on $x = t$. A smooth solution may be obtained by including the small conductivity term and introducing an interior layer by means of a substitution $\xi = (t - x)/\varepsilon^{1/2}$. Even when $g(0) = 0$ the derivatives of T will be discontinuous so that a boundary layer will again be required in order to obtain a smooth function T. The change in T, however, will only be of order $\varepsilon^{1/2}$ in this case.

For a specific example we assume fluid is injected at a constant temperature $T_1 > 0$ so that $g(t) = T_1$ for $t > 0$. The potential difference V between the electrodes is given by $V = I \int_0^L 1/\sigma \, dx$. This and the temperature at $x = L$ can be measured. The potential difference is given by:

$$V(L,t) = \begin{cases} \dfrac{L - t}{\sqrt{1 + 2\alpha K t}} + \dfrac{1}{\alpha K}\left[\sqrt{(1 + \alpha T_1)^2 + 2\alpha K t} - (1 + \alpha T_1)\right], & 0 < t < L \\ \dfrac{1}{\alpha K}\left[\sqrt{(1 + \alpha T_1)^2 + 2\alpha K L} - (1 + \alpha T_1)\right], & t > L. \end{cases} \tag{6.10}$$

The temperature at $x = L$ is given by:

$$T(L,t) = \begin{cases} \dfrac{1}{\alpha}(\sqrt{1 + 2\alpha K t} - 1), & 0 < t < L \\ \dfrac{1}{\alpha}\left[\sqrt{(1 + \alpha T_1)^2 + 2\alpha K L} - 1\right], & t > L. \end{cases} \tag{6.11}$$

Measurement of these quantities would allow the calculation of the only unknown quantities α and $\tilde{\sigma}$. It would also give some indication of whether the assumption that α and $\tilde{\sigma}$ are constant is a reasonable one.

The quantities above are calculated using $\alpha = 0.4$, $K = 1.25$, $L = 1$ and $T_1 = 0.5$. Since $T_1 \neq 0$ there is a temperature front along the characteristic $x = t$. In Figure 6.1 we show the overall resistance which with constant current is just V. We see that V becomes constant after the front reaches $x = L$. In Figure 6.2 we show the temperature as a function of both time and position. The movement of the front is clear.
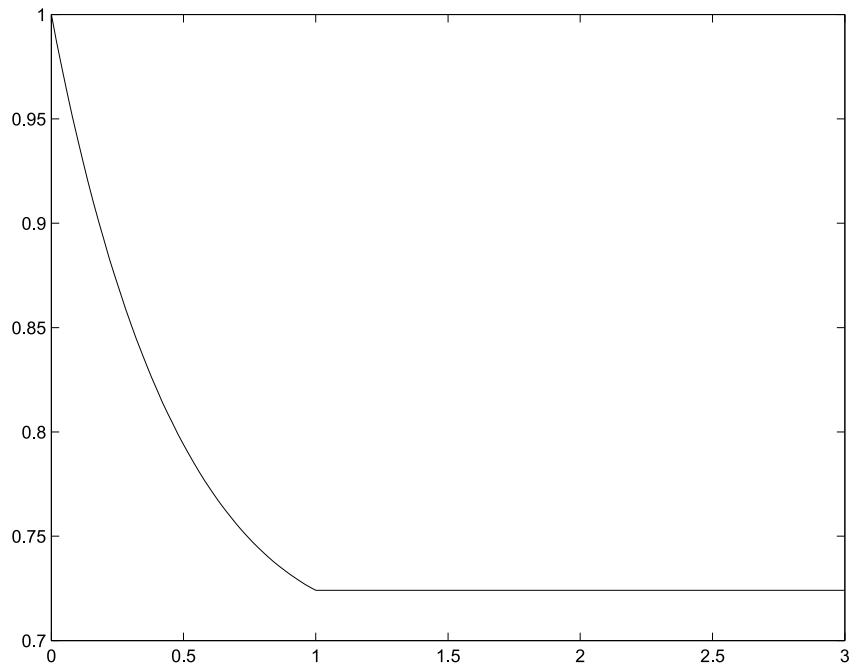
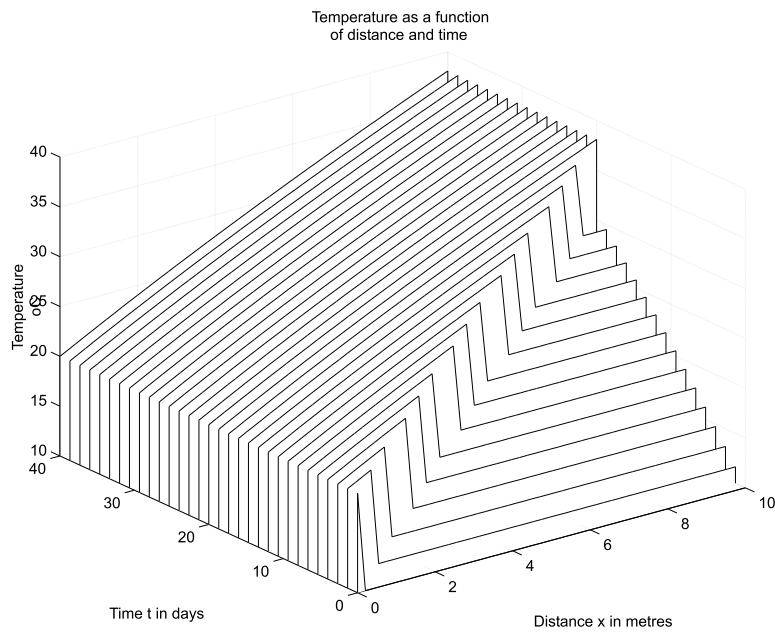$$\pi$$

Figure 6.1: Overall resistance.

Figure 6.2: Temperature as a function of space and time.

If we prescribe the potenal V at the electrodes the current I becomes a function of time which is determined by integrating the potential equation (6.6) with respect to x from $x = 0$ to $x = L$. If we take the potential difference between the electrodes to be 1, we find

$$I(t) = 1/\int_0^L \frac{1}{1 + \alpha T(x,t)}\, dx. \tag{6.12}$$

The equation (6.7) then becomes non-local and an analytic solution is no longer possible. The equation can however still be solved numerically. With constant $\widetilde{\sigma}$ the results are similar to those for the case of constant current.

In some numerical examples we look at the effects which occur in the current flow and temperature at the well (both measurable quantities) due to an anomaly in the conductivity where we assume that $\widetilde{\sigma}$ has a piecewise change over a region beginning at a specified time. We find that in this simple case of a single anomaly measurements at the boundary are enough to show the time at which the anomaly occurs, its location and its magnitude. With more than one anomaly information may be obtained but there is a loss of uniqueness.

Numerical solutions have been found for an anomaly occurring at time 30 days over the region between 4 and 5 metres. In Figure 6.3 we show graphs of the overall resistance and in Figure 6.4 the output temperature for the case with and without an anomaly which is an increase in scaled conductivity $\widetilde{\sigma}$ from 1.0 to 1.5. Notice that the overall resistance drops at the time the anomaly occurs. The temperatures in the cases with and without anomaly also start to separate at this time. There are also noticeable changes in the temperature at the times that the characteristics through the ends of the anomalous region reach the output well. Since the speed of propagation along the characteristics is known this enables us to locate the anomaly. The percentage change in the resistance at the jump allows us to estimate the percentage change in $\widetilde{\sigma}$.

In Figures 6.5, 6.6 we show similar graphs with an anomaly which is a decrease in scaled conductivity $\widetilde{\sigma}$ from 1.0 to 0.7.
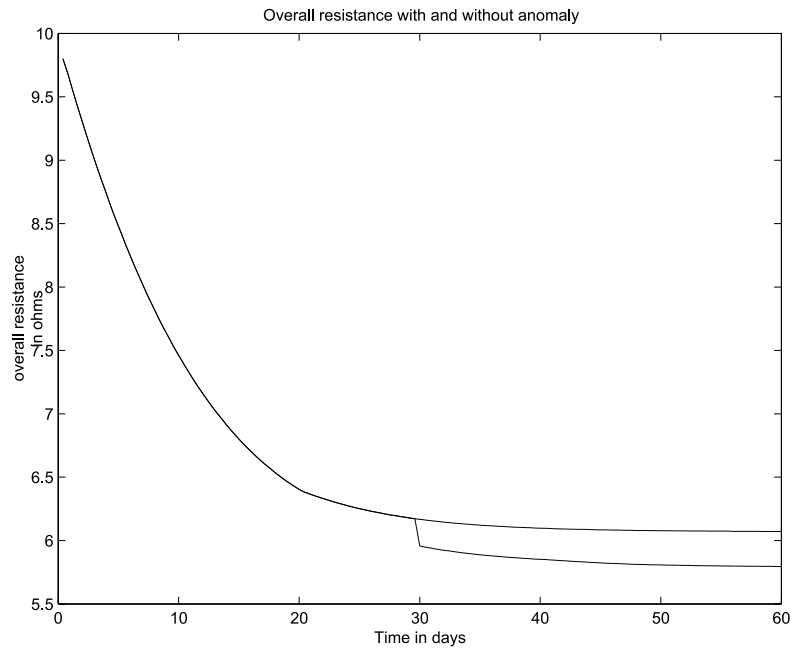
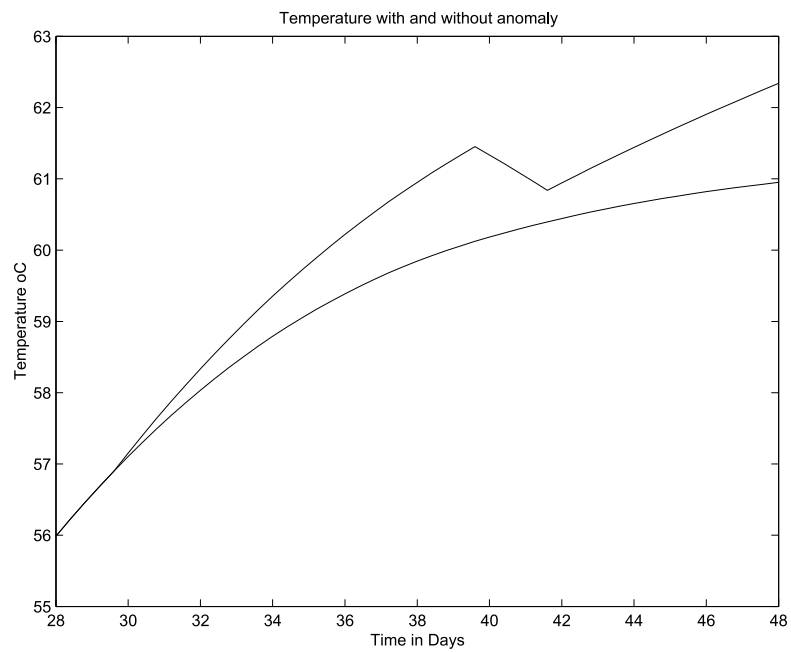Figure 6.3: Overall resistance anomalous increase in conductivity.

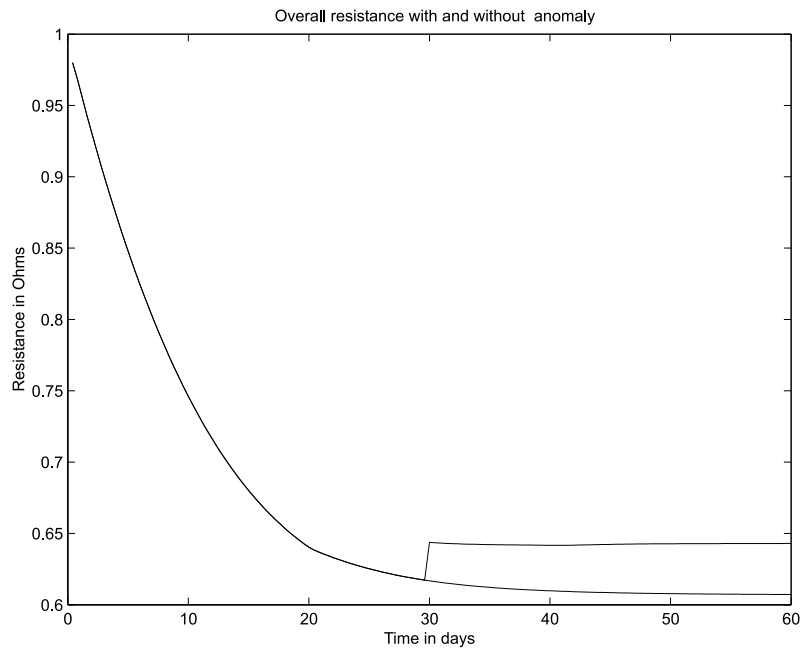Figure 6.4: Temperature (Zoomed). Anomalous increase in conductivity.

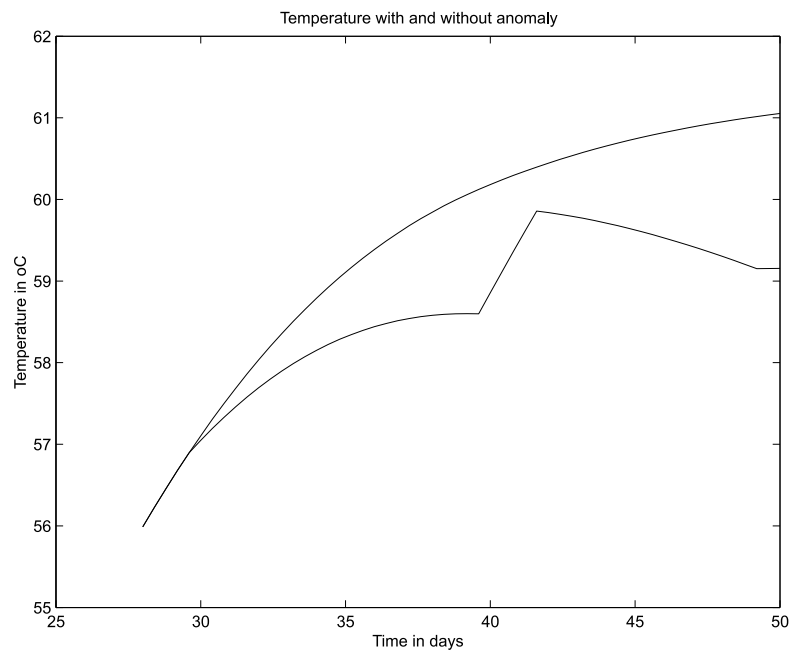Figure 6.5: Resistance. Anomalous decrease in conductivity.

Figure 6.6: Temperature (Zoomed). Anomalous decrease in conductivity.

6.4   The three dimensional case

We make here a few remarks extending the one-dimensional case. These have not been developed but could be the basis for further analysis. The equations for the fluid flow are not coupled to the other equations and the fluid flow is essentially determined by the boundary conditions at the wells. These are Dirichlet, if the pressure head is given, or Neumann, if the velocity is given.

In the case where the permeability, viscosity and electrical conductivity are constant, the pressure head and the electrical potential both satisfy the Laplace equation, if the boundaries and boundary conditions match (i.e. both Dirichlet or both Neumann) the fluid streamlines and the current streamlines will coincide. These streamlines are also characteristic curves for the energy equation and since they will begin and end at a well the energy equation (6.3) may be integrated along a streamline to give:

$$T_2 - T_1 = \int_{r_1}^{r_2} \frac{I^2(r)}{\sigma} \frac{1}{\rho_w c_w u(r)} \, dr \tag{6.13}$$

where r is the distance along a streamline and $u(r)$ is the fluid speed along the streamline. Subscripts 1 and 2 refer to the wells at the beginning and end of the streamline. The temperature T_1 can be controlled, sudden changes in T_2, which should be local due to the small thermal conductivity, would indicate an anomaly along the corresponding streamline (characteristic).

Both the compressibility and potential equation can be integrated along a stream tube to give constant volumetric flow $Q = \int_{\partial A} \vec{u} \cdot \vec{n} \, dA$ and constant total current flow $I = \int_{\partial A} \sigma \nabla V \cdot \vec{n} \, dA$ along a stream tube (here ∂A is a surface perpendicular to the stream tube). These also can be monitored at the wells.

In conditions of steady state temperature one can use the incompressibility equation (6.2) and the potential equation (6.1) to reduce the energy equation (6.3) to:

$$\nabla \cdot \left(\vec{u} T - \frac{1}{\rho_w c_w} \sigma V \nabla V \right) = 0. \tag{6.14}$$

Integrating this along a stream tube we see that

$$\int_{\partial A} \left(\vec{u} T - \frac{1}{\rho_w c_w} \sigma V \nabla V \right) \cdot \vec{n} \, dA = C \tag{6.15}$$

is also constant along a thin stream tube. At a well this is essentially

$$QT - \frac{1}{\rho_w c_w} V I.$$

By monitoring these quantities at the two wells we should be able to detect the occurrence of anomalies in the conductivity. Collating this with the temperature measurements might allow us to pinpoint any anomalies in both time and space.

6.5   The inverse potential problem.

As stated in the introduction, by making a series of impedance measurements about the boundary of the well site an image of the conductivity distribution can be determined. This is complicated by the fact that not all of the boundary is available for measurement. If however we

concentrate on sudden changes in conductivity, a detailed image is not required and what follows may be considered rather as a first step in their detection.

6.5.1   Solution of the forward problem and a discussion of the resolving power

To simplify the inverse problem we consider a vertical section of soil of thickness w and cylindrical cross section of radius r_0. The domain is surrounded with N equally spaced electrodes and we assume that a given electrode can either supply current or monitor voltage. Electrodes extend the width of the section and do not overlap ensuring that the width of an electrode, $2a$, satisfies the bound $2aN < 2\pi r_0$. For a given set of electrodes there are many possible applied current patterns. Our assumption will be that a current of I is applied between two of the electrodes and the remainder of the electrodes monitor the resulting voltage at the boundary of the domain.

The forward problem is given a conductivity, $\sigma(r, \theta, z)$, to determine the potential $\varphi(r, \theta, z)$ satisfying

$$\nabla \cdot (\sigma \nabla \varphi) = 0 \tag{6.16}$$

and the boundary condition

$$J(\theta) = \sigma \left.\frac{\partial \varphi}{\partial r}\right|_{r=r_0} = \begin{cases} \dfrac{I}{2aw}, & \alpha - \frac{a}{r_0} \leq \theta \leq \alpha + \frac{a}{r_0} \\[2mm] -\dfrac{I}{2aw}, & \beta - \frac{a}{r_0} \leq \theta \leq \beta + \frac{a}{r_0} \\[2mm] 0, & \text{all other } \theta. \end{cases} \tag{6.17}$$

The angles α, β are the angular positions of the electrodes. An estimate of the sensitivity of the potential to changes in the conductivity can be determined by assuming the conductivity is piecewise constant

$$\sigma = \begin{cases} \sigma_0, & r_1 \leq r \leq r_0 \\ \sigma_1, & 0 \leq r < r_1. \end{cases} \tag{6.18}$$

Since both the potential and the current are continuous at $r = r_1$ one has the conditions

$$\varphi|_{r=r_1^-} = \varphi|_{r=r_1^+}, \qquad \sigma_1 \left.\frac{\partial \varphi}{\partial r}\right|_{r=r_1^-} = \sigma_0 \left.\frac{\partial \varphi}{\partial r}\right|_{r=r_1^+}. \tag{6.19}$$

Solving (6.16) with (6.18)-(6.19) with the constraints that φ is periodic in θ and constant with respect to z yields

$$\varphi(r, \theta) = \varphi_0 + \sum_{k=1}^{\infty} f_k(r)[a_k \cos(k\theta) + b_k \sin(k\theta)]$$

where

$$f_k(r) = \begin{cases} \left(1 + \dfrac{\sigma_0 - \sigma_1}{\sigma_0 + \sigma_1}\right)\left(\dfrac{r}{r_1}\right)^k, & 0 \leq r < r_1 \\[3mm] \left(\dfrac{r}{r_1}\right)^k + \left(\dfrac{\sigma_0 - \sigma_1}{\sigma_0 + \sigma_1}\right)\left(\dfrac{r_1}{r}\right)^k, & r_1 \leq r \leq r_0. \end{cases} \tag{6.20}$$
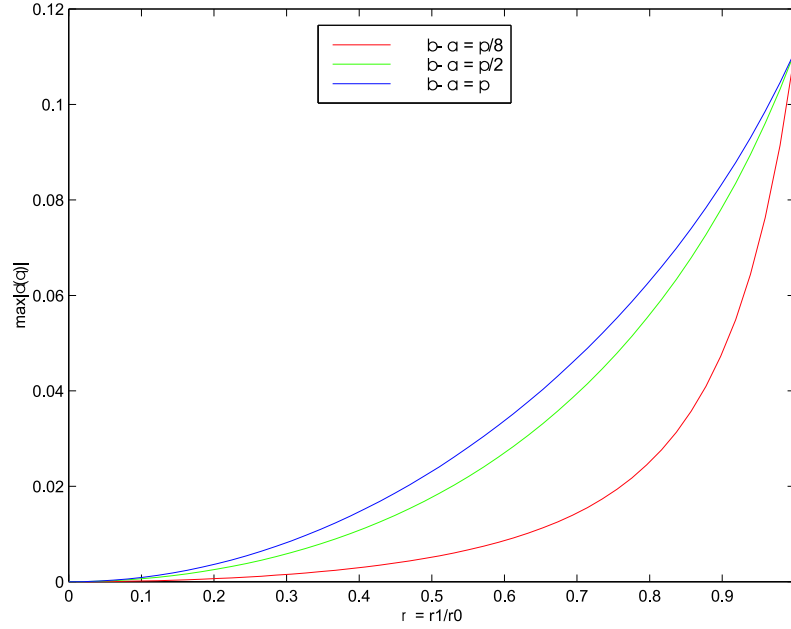
π

Figure 6.7: Resolving power for various configurations.

Applying the boundary condition (6.17) yields the following expression for the potential on the boundary of the domain

$$\varphi(r_0, \theta) = \varphi_0 + \frac{Ir_0}{\sigma_0 aw} \sum_{k=1}^{\infty} \frac{1 + \mu\rho^{2k}}{1 - \mu\rho^{2k}} \frac{1}{k^2} \sin\left(\frac{ka}{r_0}\right) [\cos k(\theta - \alpha) - \cos k(\theta - \beta)]$$

where $\mu = (\sigma_0 - \sigma_1)/(\sigma_0 + \sigma_1)$ and $\rho = r_1/r_0$.

Denote $\overline{\varphi}$ as the potential at $r = r_0$ when $\sigma_1 = \sigma_0$. Using this definition one can define the resolving power as

$$\delta = \frac{\varphi - \overline{\varphi}}{\overline{\varphi}} = \frac{\displaystyle\sum_{k=1}^{\infty} \left[\frac{1 + \mu\rho^{2k}}{1 - \mu\rho^{2k}} - 1\right] \frac{1}{k^2} \sin\left(\frac{ka}{r_0}\right) [\cos k(\theta - \alpha) - \cos k(\theta - \beta)]}{\displaystyle\sum_{k=1}^{\infty} \frac{1}{k^2} \sin\left(\frac{ka}{r_0}\right) [\cos k(\theta - \alpha) - \cos k(\theta - \beta)]}.$$

Figure 6.7 plots $\max |\delta(\theta)|$ as a function of ρ in the case of $N = 16$ electrodes. In addition, $a = 0.1$, $r_0 = 1$ and $\sigma_1/\sigma_0 = 0.9$. Illustrated are cases where the current is applied at adjacent electrodes ($\beta - \alpha = \pi/8$), electrodes at right angles ($\beta - \alpha = \pi/2$) and opposite electrodes ($\beta - \alpha = \pi$). Clearly with an anomaly centered at $r = 0$ current applied at opposite electrodes is the most sensitive.

6.5.2   An iterative technique for solving the inverse problem

The method implemented for solving the inverse problem is to use the sensitivity theorem usually attributed to Geselowitz [3]. For a given region of conductivity σ an application of current I

applied between electrodes A and B generates a potential field $\varphi(\sigma)$ inside the volume. The potential between electrodes C and D is denoted φ_{CD}. By applying the same current between electrodes C and D one obtains the potential ψ_{AB} between electrodes A and B. One may note that

$$\phi_{CD}(\sigma) = \frac{1}{I} \int_{\Omega} \sigma \nabla \phi(\sigma) \cdot \nabla \psi(\sigma) \, d\Omega = \psi_{AB}(\sigma).$$

If the conductivity of the region changes to $\sigma + \Delta\sigma$ then the change in the transfer impedance between electrodes A, B and C, D is

$$\Delta Z = \phi_{CD}(\sigma + \triangle\sigma) - \phi_{CD}(\sigma) = \phi_{CD}(\sigma + \triangle\sigma) - \psi_{AB}(\sigma) = -\frac{1}{I} \int_{\Omega} \Delta\sigma \, \nabla\psi(\sigma) \cdot \nabla\phi(\sigma + \Delta\sigma) \, d\Omega. \tag{6.21}$$

To solve the inverse problem one assumes a conductivity distribution σ and solves the corresponding forward problem obtaining φ. Taking the difference between the calculated $\varphi(\sigma)$ and the measured $\varphi(\sigma + \triangle\sigma)$ yields ΔZ. In principle we can use this information to solve for $\Delta\sigma$ and find the conductivity distribution.

Consider a mesh $\{\Omega_k\}_{k=1}^{M}$ consisting of M elements each having its own conductivity σ_k. In this way one may approximate the conductivity distribution as a finite sum of characteristic functions

$$\sigma = \sum_{j=1}^{M} \sigma_j \chi(\Omega_j), \qquad \chi(\Omega_j) = \begin{cases} 1, & (x,y) \in \Omega_j \\ 0, & \text{otherwise.} \end{cases}$$

We refer to the collection of conductivities, $\{\sigma_j\}_{j=1}^{M}$, as the vector $\vec{\sigma} \in \mathbb{R}^M$. The corresponding discrete version of the sensitivity theorem (6.21) becomes

$$\Delta Z = -\sum_{j=1}^{M} \Delta\sigma_j \int_{\Omega_j} (\nabla\varphi \cdot \nabla\psi) \, d\Omega_j. \tag{6.22}$$

In this way ΔZ refers to the difference in voltage between electrodes C and D when the input electrodes are A and B and the conductivity changes from σ to $\sigma + \Delta\sigma$.

During the inversion process a number of different current injection patterns are applied. If for example there are E electrodes and current is applied at adjacent electrodes then there are $E-1$ patterns and therefore a total of $(E-1)\times(E-1)$ readings.[1] This collection of measurements on the boundary over the class of chosen patterns generates the vector $\Delta\vec{Z} = \vec{\zeta} - \vec{\varphi} \in \mathbb{R}^N$. Here $\vec{\zeta}$ is the vector of measured voltages and $\vec{\varphi}$ is the vector of the calculated potentials for the current conductivity distribution σ. The final form of the (6.22) is

$$\Delta\vec{Z} = S\Delta\vec{\sigma} \tag{6.23}$$

where the S is an $N \times M$ matrix whose elements are

$$S_{ij} = -\int_{\Omega_j} (\nabla\varphi_i \cdot \nabla\psi_i) \, d\Omega_j \tag{6.24}$$

---

[1]This assumes the capability of measuring the voltage at the injection electrodes.

π

the result in cell j obtained when applying pattern i. Current applied to the electrode pair A, B generates φ_i while current applied at C, D generates ψ_i. Because of symmetries $N = E(E-1)/2$. One of the difficulties with the reciprocity theorem is that $S^T S$ is ill conditioned and may be singular if $N \leq M$. Another problem encountered in using the reciprocity theorem is that one cannot use a basic FEM package to solve the forward problem. Direct access to the basis functions of each cell is required to compute S.

Because the problem is ill posed, Tikhonov regularization is used to determine the conductivity. With this method, rather than solving (6.23), one attempts at the kth iteration to minimize the functional

$$f(\Delta\vec{\sigma}_k) = \left\| S\Delta\vec{\sigma}_k - (\vec{\zeta} - \vec{\varphi}_k) \right\|^2 + \lambda \left\| F\vec{\sigma}_k \right\|^2 \qquad (6.25)$$

which indicates a trade off with obtaining a solution to $S\Delta\vec{\sigma}_k = \vec{\zeta} - \vec{\varphi}_k$ and not letting $\|F\vec{\sigma}_k\|$ get too large. Typically F is chosen to penalize solutions with properties that one wants to avoid. In our case, we want to avoid solutions with discontinuities so F is a smoothing operator. Since the current implementation is a Laplacian operator, F is an $M \times M$ matrix with $F_{ij} = -1$ if node i is connected to node j ($i \neq j$) and $F_{ii} =$ number of elements connected to node i. If one lets the regularization parameter $\lambda \to 0$ one obtains a generalized solution of $S\Delta\vec{\sigma}_k = \vec{\zeta} - \vec{\varphi}_k$. We now detail the algorithm used to generate the solution.

Step 1: The first step is to determine an appropriate estimate for a homogeneous conductivity. This is accomplished by setting the conductivity of each cell to one and generating the solutions of the forward problems

$$\nabla \cdot (\nabla \varphi_0) = 0, \qquad \left.\frac{\partial \varphi_0}{\partial \mathbf{n}}\right|_{\partial\Omega} = \vec{J}$$

where each of the applied current patterns, $\vec{J}$, are used in turn. The effective homogeneous conductivity σ_0 is chosen to minimize

$$\left\| \vec{\varphi}_0 \frac{1}{\sigma_0} - \vec{\zeta} \right\|^2. \qquad (6.26)$$

Solving (6.26) gives the least squares solution

$$\sigma_0 = \frac{\vec{\varphi}_0 \cdot \vec{\varphi}_0}{\vec{\varphi}_0 \cdot \vec{\zeta}}.$$

Every cell is set to this initial conductivity giving the constant initial vector $\vec{\sigma}_0 = \{\sigma_0\}_{i=1}^M$. The iteration parameter $k = 1$.

Step 2: The next step is to generate perturbations in the conductivity. At this point we solve the forward problems

$$\nabla \cdot (\sigma_{k-1} \nabla \varphi_k) = 0, \qquad \left.\sigma_{k-1} \frac{\partial \varphi_k}{\partial r}\right|_{\partial\Omega} = J;$$

for the $E-1$ current patterns and generate the vectors $\vec{\varphi}_k$ by appending the collection the boundary values at the electrodes for each of the applied patterns (as done in Step 1). For
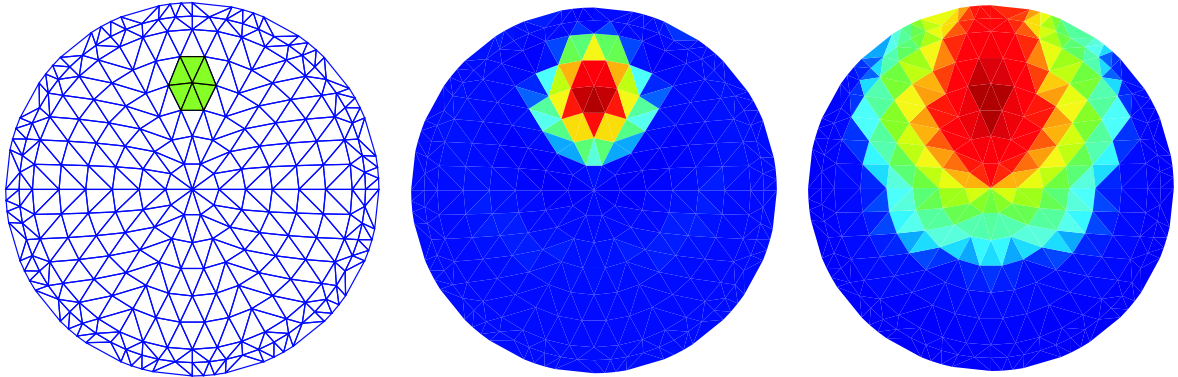
$$\pi$$

Figure 6.8: Initial anomaly and reconstructed solutions for two values of λ.

all subsequent steps the forward solution in all of the cells is required and not just those corresponding to electrodes. Using the forward solutions one computes the sensitivity matrix, S, from expression (6.24). At this point $\Delta\vec{\sigma}_k$ is chosen to minimize (6.25) and as a result, satisfies

$$\left(S^T S + \lambda F^T F\right) \Delta\vec{\sigma}_k = S^T(\vec{\zeta} - \vec{\varphi}_k) + \lambda F^T F \vec{\sigma}_{k-1}.$$

Step 3: The updated conductivity is set to $\vec{\sigma}_k = \vec{\sigma}_{k-1} + \Delta\vec{\sigma}_k$ and we check the stopping criterion. If $\|\Delta\vec{\sigma}_k\| \geq \epsilon$ then $k \mapsto k+1$ iterate S, $\vec{\sigma}_k$, $\vec{\varphi}_k$ and return to Step 2. Otherwise the loop terminates.

For the numerical simulations we chose $N = 16$ electrodes and a homogeneously conductive region of $\sigma = 1$ with a small conductive anomaly with $\sigma = 0.9$. The mesh has $M = 279$ nodes. Figure 6.8 displays the original anomaly and two reconstructed solutions corresponding to $\lambda = 1 \times 10^{-4}$ and $\lambda = 10$ respectively.

Clearly the cylindrical geometry is an idealization of the actual geometry. However, the promising results of the reconstructed solutions shown here indicate that detecting sudden large-scale changes in conductivity is certainly feasible. Moreover, for a given configuration of electrodes, estimates can be obtained for the size of the smallest detectable anomaly of a given strength.

6.6 Summary

Effort was focussed on the detection of sudden localised anomalies in the conductivity and progress was made on a simple one dimensional model for the coupled problem and on a two dimensional model of the uncoupled potential problem. These investigations showed some promise and could be developed further.

π

Since substantial literature exists on the potential problem, further work should begin with a thorough literature search. The MacMillan McGee problem may present some novel aspects. Since measurements will only be made on an actual decontamination process, multiple current inputs will not be available. Measurements will also only be available at electrodes, wells and the soil surface so that the effect of boundary conditions at "infinity" should be considered. This all adds to the illposedness (see Cheney, Issacson and Newell [1])

The use of temperature as well as potential measurements for the inverse problem might also prove fruitful. This would require an efficient algorithm for the numerical solution of the three dimensional coupled forward problem.

Changes in overall resistance similar to that found in the one dimensional example have been observed in practice and it would be interesting to measure temperatures at the output wells to see whether effects predicted by the one dimensional case do actually occur and whether they are spatially localised.

π

Bibliography

[1] Cheney, M., Isaacson, D. & Newell, J.C. (1999). *Electrical Impedance Tomography*, SIAM Review, 41, 1, pp 85-102.

[2] Fowler, A.C. (1997). *Mathematical Models in the Applied Sciences*, Cambridge University Press: Cambridge.

[3] Geselowitz, D.B. (1971). *An application of electrocardiographic lead theory to impedance plethysmography*, IEEE Transactions on Biomedical Engineering, BME-18, pp 38-41.

[4] Kohn, R.V. & Vogelius, M. (1984). *Identification of an unknown conductivity by means of measurement at the boundary*, In Inverse Problems, (Ed. McLaughlin, M.D.) SIAM-AMS Proceedings, 14, pp. 113-123.

Appendix A

List of Participants

Ibrahim Abba	University of British Columbia	abba@chml.ubc.ca
Rita Aggarwala	University of Calgary	rita@math.ucalgary.ca
Ibrahim Agyemang	University of Alberta	ibagyemang@hotmail.com
Erik Andries	University of New Mexico	andriese@ahpcc.unm.edu
Andrei Badescu	University of Western Ontario	andre_badescu@hahoo.com
Niu Baoning	Queen's University	niu@cs.queensu.ca
Lyudmyla Barannyk	New Jersey Institute of Technology	llb1313@njit.edu
Michael Baumann	University of British Columbia	michael@baumann.ca
Melanie Beck	McGill University	beck@math.mcgill.ca
Elaine Beltaos	University of Alberta	ebeltaos@math.ualberta.ca
Hammouda Ben Mekki	University of Calgary	hammouda@math.ucalgary.ca
Petra Berenbrick	Simon Fraser University	
Gary Billings	Talisman	GBillings@talisman-energy.com
Mikalai Birukou	University of Winnipeg	mbirukou@ io.uwinnipeg.ca
Sean Bohun	Penn State	csb15@psu.edu
Matt Bolton	University of British Columbia	mattb@math.ubc.ca
Chris Bose	University of Victoria	cbose@math.uvic.ca
Chris Budd	University of Bath	
Alicia Canton	University of Washington	canton@math.washington.edu
Bo Young Chang	University of British Columbia	boyoung@math.ubc.ca
Jiuping Chen	University of British Columbia	jchen@geop.ubc.ca
Raymond Ka Wai Cheng	University of Toronto	kwrcheng@math.utoronto.ca
Lorraine Dame	University of Victoria	ifdpanda@math.uvic.ca
Isabelle Dechene	McGill University	dechene@math.mcgill.ca
Aruliah Dhavide	Field Institute	dhavida@fields.utoronto.ca
Roger Donaldson	University of British Columbia	rdonald@math.ubc.ca
Neville Dubash	University of British Columbia	dubashmech.ubc.ca
Oliver Dubois	McGill University	dubois@math.mcgill.ca
Leslie Fairbairn	Simon Fraser University	lfairbai@sfu.ca
Colin Farquharson	University of British Columbia	farq@eos.ubc.ca
Halasan Florina	University of British Columbia	halasan@math.ubc.ca
Charles Fortin	McGill University	fortin@math.mcgill.ca
Ellen Fowler	University of British Columbia	ellenfowler@telus.net

Ian Frigaard	University of British Columbia	frigaard@math.ubc.ca
Guojun Gan	York University	gjgan@mathstat.yorku.ca
Asrat Fikre Gashaw	York University	agashaw@yorku.ca
Hugh Geiger	University of Calgary	geiger@geo.ucalgary.ca
Jeffrey Gilmore	Simon Fraser University	jrg@sfu.ca
Maria Inez Goncalves	University of Victoria	inez@math.uvic.ca
Ana Granados	University of Washington	granados@math.washington.edu
Jeff Grossman	University of Calgary	grossman@geo.ucalgary.ca
J. Ying Han	McGill University	ying@math.mcgill.ca
Katherine Hegewisch	WSU Pullman	khegewisch@wsu.edu
Yuanyuan Hua	McGill University	hua@math.mcgill.ca
Huaxiong Huang	York University	hhuang@yorku.ca
Baoguo Jiang	University of Western Ontario	bjiang3@uwo.ca
Edward Keyes	Semiconductor Insights	edward@semiconductor.com
Pedram Khademsameni	Simon Fraser University	pedram@sfu.ca
Michael Lamoureux	University of Calgary	mikel@math.ucalgary.ca
Maria Landry	University of British Columbia	maria@math.ubc.ca
Stephan Lawi	University of Toronto	slawi@math.toronto.edu
Namyong Lee	Minnesota State University	nlee@krypton.mnsu.edu
Heather Lehr	University of Texas at Austin	heather@math.utexas.edu
Qingguo Li	Simon Fraser University	qlib@sfu.ca
Edmond Lou	Capital Health	edmond_lou@shaw.ca
Jingxiang Luo	University Saskatchewan	jil072@cs.usask.ca
Frithjof Lutscher	University of Alberta	flutscher@math.ualberta.ca
David Lyder	King's University College	dlyder@kingsu.ab.ca
Colin Macdonald	Simon Fraser University	cbm@math.sfu.ca
Jack Macki	University of Alberta	jmacki@gpu.srv.ualberta.ca
Bruce McGee	McMillan-McGee Corp.	mcgee@mcmillan-mcgee.com
Jenny McNulty	University of Montana	mcnulty@mso.umt.edu
Alexander Melnikov	University of Alberta	melnikov@ualberta.ca
Zheng Meng	Rice University	zhmeng@rice.edu
Marat Molyboga	Moscow State University	moliboga@yahoo.com
Joy Morris	University of Lethbridge	morris@cs.uleth.ca
Miguel Moyers	University of British Columbia	mamgmex@math.ubc.ca
Anuj Mubayi	University of Texas at Arlington	mubayi@exchange.uta.edu
Matthias Mueck	University of Toronto	mueck@math.toronto.edu
Kai Muenzer	Shell Canada	Kai.Muenzer@shell.ca
Tim Myers	University of Cape Town	myers@maths.uct.ac.za
Tereza Neocleous	University of Illinois	neocleou@stat.uiuc.edu
Nancy Ann Neudauer	Pacific University	nancy@pacificu.edu
Eva-Marie Nosal	University of British Columbia	eva-marie@math.ubc.ca
Arian Novruzi	University of British Columbia	a.novruzi@math.ubc.ca
Douglas Oldenburg	University of British Columbia	doug@eos.ubc.ca
Benjamin Ong	Simon Fraser University	bwo@sfu.ca
Jeff Orchard	Simon Fraser University	jjo@cs.sfu.ca

π

Gergely Orosi	University of British Columbia	orosi@interchange.ubc.ca
Anita Parkinson	University of British Columbia	AParkinson@coe.ubc.ca
Marc Paulhus	University of Calgary	paulhusm@math.ucalgary.ca
Antony Peirce	University of British Columbia	peirce@math.ubc.ca
Yuri Petratchenko	Moscow State University	ygpetratyandex.ru
Robert Piche	Tampere University of Technology	robert.piche@tut.fi
Nicolas Robidoux	University of New Mexico	nrobidoux@pims.math.ca
Bidhan Roy	University of Illinois	bcroy@uiuc.edu
Daniel Ryan	University of British Columbia	dpryan@math.ubc.ca
Arunas Salkauskas	Sorex Software	arunas@sorexsoftware.com
Tetyana Segin	New Jersey Institute of Technology	tms3@njit.edu
Tingting Shu	University of Alberta	tshu@ualberta.ca
Victoria Skornyakova	University of Alberta	vs2@ualberta.ca
Michael Slawinski	Memorial University	mslawins@mun.ca
Brett Stevens	Carleton University	brett@math.carleton.ca
Robin Swain	Memorial University of Newfoundland	rswain@math.mun.ca
Jacobus Swarts	University of Victoria	cobus@math.uvic.ca
Anatoliy Swishchuck	York University	aswishch@mathstat.yorku.ca
Yongi Tan	Fudan University	
Carlos Trendado	New Mexico State University	ctrenado@nmsu.edu
Tzvetalin Vassilev	University of Saskatchewan	tsv552@mail.usask.ca
Adam Webber	University of Victoria	awebber@math.uvic.ca
Rex Westbrook	University of Calgary	westbroo@ucalgary.ca
Brian Wetton	University of British Columbia	wetton@math.ubc.ca
Chad Wheaton	Memorial University	
Huamei Yin	University of Alberta	huamei@ualberta.ca
Fabien Youbissi	University of Laval	fyoubi@mat.ulaval.ca
Qingze Zou	University of Washington	qzouatuw@u.washington.edu

π

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

Pacific Institute for the Mathematical Sciences