Robotic Subsurface Mapping Using Ground Penetrating Radar

Herman Herman

CMU-RI-TR-97-19

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Roboticss

The Robotics Institute
Carnegie Mellon University
Pitsburgh, Pennsylvania 15213

May 1997

© 1997 by Herman Herman. All Rights reserved.

Abstract

The goal of our research is to develop an autonomous robot for subsurface mapping. We are motivated by the growing need for mapping buried pipes, hazardous waste, landmines and other buried objects. Most of these are large scale mapping problems, and to manually construct subsurface maps in these cases would require a significant amount of resources. Therefore, automating the subsurface mapping process is an important factor in alleviating these problems.

To achieve our goal, we have developed a robotic system that can autonomously gather and process Ground Penetrating Radar (GPR) data. The system uses a scanning laser rangefinder to construct an elevation map of an area. By using the elevation map, a robotic manipulator can follow the contour of the terrain when it moves the GPR antenna during the scanning process. The collected data are then processed to detect and locate buried objects. We have developed three new processing methods, two are volume based processing methods and one is a surface based processing method. In volume based processing, the 3-D data are directly processed to find the buried objects, while in surface based processing, the 3-D data are first reduced to a series of 2.5-D surfaces before further processing. Each of these methods has its own strengths and weaknesses. The volume based processing methods can be made very fast using parallel processing techniques, but they require an accurate propagation velocity of the GPR signal in the soil. On the other hand, the surface based processing method uses 3-D segmentation to recognize the shape of the buried objects, which does not require an accurate propagation velocity estimate. Both approaches are quite efficient and well suited for online data processing. In fact, they are so efficient that the current bottleneck in the subsurface mapping process is the data acquisition phase.

The main contribution of the thesis is the development of an autonomous system for detecting and localizing buried objects. Specifically, we have developed three new methods to find buried objects in 3-D GPR data. Using these methods, we are able to autonomously obtain subsurface data, locate and recognize buried objects. These methods differ from existing GPR data processing methods because they can autonomously extract the location, orientation, and parameters of the buried object from high resolution 3-D data. Most existing methods only enhance the GPR data for easier interpretation by human experts. There are some existing works in automated interpretation of GPR data, but they only work with 2-D GPR data. We implemented the three different methods and also tested them by building subsurface maps of various buried objects under different soil conditions. We also used these subsurface mapping methods to demonstrate an autonomous buried object retrieval system. In summary, we have developed a robotic system which make subsurface mapping faster, more accurate and reliable.

Acknowledgments

Foremost I would like to thank Anthony Stentz for his guidance and support for the last 5 years. He was always open to my ideas, gave thoughtful comments, and asked tough questions. In short, Tony is a great advisor.

I also would like to thank Andy Witkin for his guidance and support during the early years of my graduate career. What I learned during those years has proved to be invaluable in my research.

Thanks to Mike Heiler for introducing me to Ground Penetrating Radar and subsurface mapping. Without his help in setting up the testbed, many experiments would not have been possible. Thanks to James Osborn for his support and constructive comments on the subject of subsurface mapping.

Thanks to Sanjiv Singh for his collaboration on the subsurface mapping testbed. The demonstration of the autonomous buried object retrieval system would not have been possible without his help. Thanks to Mike Blackwell and Regis Hoffman for help with the laser scanner.

Thanks to Behnam Motazed, Jeffrey Daniels, Hans Moravec and Martial Hebert for agreeing to be on my thesis committee and for their thoughtful comments on the research.

Thanks to my officemates, Dave Wettergreen, Dimitrious Apostolopoulos, Mike Heiler, and Sanjiv Singh for their friendship and willingness to help on various occasions.

Thanks to Red Whittaker for founding the Field Robotics Center and providing the inspiration for many of us to do research in robotics. The support of many wonderful people in this organization has made the research considerably better.

Thanks to my father and mother. None of this would have been possible without their continuing love and support. I especially thank them for believing in me and teaching me that with hard work, I can reach all my goals in life. Thanks to my brother and sisters, Maman, Nany and Wawa, for their support and encouragement.

Finally, I would like to thank my wife Lingga for her love and support during the past couple of years, and Baylee for faithfully accompanying me during the long hours of work.

Table of Contents

1. Robotic Subsurface Mapper 9

1.1. Introduction to Buried Object Mapping Problem 9

1.2. Research Objective 12

1.3. Technical Approach 13

1.4. Rationales 18

1.5. Applications of the Robotic Subsurface Mapper 19

2. Related Work 23

2.1. Subsurface Mapping 23

2.2. Automated Excavation and Buried Object Retrieval 26

3. Ground Penetrating Radar 29

3.1. Selection of Subsurface Sensors 29

3.2. GPR Data Collection and Data Format 37

3.3. Analysis of Different Antenna Array Configurations 39

3.4. GPR Data Visualization 43

3.5. 2-D Visualization 44

3.6. 3-D Visualization 50

3.7. GPR Data Processing and Interpretation 53

3.8. Example of GPR Data 55

4. Volume Based Processing 59

4.1. Overview of Volume Based GPR Data Processing 59

4.2. Background Noise Removal 60

4.3. 3-D Migration Using Coherent Summation 61

4.4. 3-D Migration using Reflector Pose Estimation 70

5. Surface Based Processing 95

5.1. Overview of Surface Based GPR Data Processing 95

5.2. Preprocessing 96

5.3. 3-D Segmentation 102

5.4. Surface Type Categorization 105

5.5. Parameter Estimation 108

5.6. Parametric Migration 110

5.7. Propagation Velocity Computation 111

5.8. Limitation of the Surface Based Processing 112

5.9. Analysis and Result 113

6. Subsurface Mapping Testbed 127

6.1. Testbed Description 127

6.2. Software Architecture 129

6.3. An Example of Mapping and Retrieval of a Buried Object 132

7. Conclusions 135

7.1. Conclusion 135

7.2. Contribution 137

7.3. Future Direction 138

8. References 139

Chapter 1. Robotic Subsurface Mapper

1.1. Introduction to Buried Object Mapping Problem

Over the last several decades, human has buried a large amount of hazardous waste, unexploded ordnance, landmines and other dangerous substances. During war periods, armies of different nations have buried millions of landmines around the world. A significant number of these landmines are still buried and active. They annually claim a significant number of innocent lives and maim many more people. Many people and companies have also improperly buried hazardous toxic wastes at various sites. As the containers that hold the toxic waste age, their conditions deteriorate, necessitating the need for retrieving and moving them to safer places. United Stated also has many military training sites which contains a huge number of unexploded ordnances. These sites can not be reused for civil applications until those unexploded ordnances are removed.

Finding these buried objects is hard because accurate maps, denoting where the objects are buried, rarely exist. In some cases, the lack of accurate map is deliberate, such as in the case of landmines. In some other cases, the maps are accidentally lost, as in the case of some buried hazardous waste sites. Even if the maps exist, vegetation or land-fill may have significantly transform the soil surface, rendering the maps unusable. Regardless of the reasons, remediators now need to find the precise locations of these buried objects.

Finding buried objects is also important in maintaining subsurface structures. We all have heard news of utility workers accidentally cutting phone or electrical lines. This happens because they do not know that there are buried utility structures at the excavation site. Even if a subsurface map of all buried structures exist, there is still a problem of registering the position of a buried structure in the map with its actual location in the world. It would be much easier if the workers could scan an area right before they excavate it and determine if there is any buried objects under it.

These are just a few examples which illustrate the growing need to find buried man-made objects. In some cases, remediation of these problems may also necessitate the retrieval of the buried objects. The scale of these problems is very large. We need to find and remove millions of landmines, and clean up hundreds of hazardous waste sites. Presently, the conventional remediation process is very time consuming and prone to errors. In order to understand the reason, let us examine how the remediation process is usually done.

Currently the conventional approach to this problem involves a sequence of manual operations. First, remediators scan an area using a sensor that could sense buried objects. The amount of data produced by the sensor is proportional to the size of the area. If the area is quite large, the sensor can easily generate immense amount of data. The sampling interval of the scanning process also affects the amount of data. As the sampling interval gets denser, the sensor produces higher resolution data for the same area, resulting in a larger amount of data. Experts then need to interpret the large amount of data in order to find the buried objects. This interpretation process could easily last for several days or weeks, depending on the size of the data. Once the experts find the buried objects in the data, they need to determine their actual locations in the world. The process of determining the actual location of an object from its position in the sensor data is called the registration process. Depending on how carefully the data are gathered, inaccurate registration will result in significant errors in the locations of the objects. As a result of these registration errors, the remediators would not find the buried objects at the computed locations. To minimize this problem, the position of the sensor during scanning is often measured by surveying equipments. The surveying activities result in a more accurate registration at the expense of longer scanning time.

The above series of steps is called subsurface mapping. It involves scanning an area using a sensor that can sense buried objects, and finding the buried objects in the sensor data. Once we have the subsurface map, then we can start addressing the buried object retrieval process.

The retrieval process also comes with its own set of problems. Many of the buried objects are hazardous to human and when they experience excessive forces, they can explode or, in the case of hazardous waste containers, their content can leak out. Heavy excavators are usually used to extricate these buried objects, and they can unintentionally apply excessive forces to the buried objects because the operator does not always know the location of the

excavator bucket with respect to the buried objects. The operator knows that the objects is located at a certain location or depth, but he/she has no information as to the relative location of the excavator bucket with respect to the buried objects. As a result, only a very thin layer of soil can be removed at a time to minimize the possibility of collision with the buried objects. Various mechanical solutions have also been tried. An example is an excavator that uses air jet to loosen the soil material and then remove the loose material using a vacuum suction device. Such mechanical solutions only work in some situations, depending on the soil type and the amount of soil that needs to be removed. A better solution would be to install a positioning device to determine the location of the bucket during excavation. Using the device, we can compute the relative position of the bucket with respect to the buried objects. Therefore, the possibility of collision is minimized, although the human operator is still exposed to possible danger from the buried objects.

We conclude that the current approaches to subsurface mapping and retrieval process are very time consuming and expose the remediators to unnecessary dangers. As a result, remediation of these buried objects is often a very costly proposition. We argue that we need to automate some of these steps. Automation could both lower the cost and shorten the time needed to perform the necessary steps. Our ultimate goal is to develop an autonomous robot that can find and retrieve buried objects. This robot would gather necessary subsurface data, interpret it, and retrieve the buried objects. We envision that the robot would consists of a computer controlled excavator, a subsurface sensor, and computers to control the mechanism and to interpret the sensor data. Using the robot, we would eliminate the operator exposure to danger from explosion or leaked chemical caused by collision with buried objects.

In this thesis, we address part of the problem, which is developing a robotic subsurface mapper to find buried objects autonomously. Specifically, we are developing algorithms that can automatically find buried objects using Ground Penetrating Radar (GPR) data. These algorithms are very important parts of the overall solution for autonomously finding and retrieving buried objects. We will show the development of the algorithms and the experimental results which prove their feasibility.

In the rest of this chapter we will explain the objective, technical approach, rationales and typical applications of our research. The issue of integrating subsurface mapping with buried objects retrieval will also be discussed in this chapter. In chapter 2, we will review existing works that have been done in the area of subsurface mapping and how they relate to our research. In chapter 3, we will explain the principle and operation of Ground Penetrating Radar (GPR), which is the sensor that we use to sense buried objects. Its limitations and different visualization techniques for its data will also be discussed. We will also examine the shortcomings of manual GPR data interpretation and the advantages of automatic GPR data interpretation. Chapter 4 and 5 will cover two different types of algorithms that we devel-

oped and implemented to automatically find buried objects in GPR data. Chapter 4 explains the first two algorithms, which is based on 3-D voxel processing, and chapter 5 explains the third algorithm, which is based on 3-D surface model based recognition. Each of these approaches has its own advantages and disadvantages, and in many cases they complement each other. In each chapter, we will also present our experimental results for each algorithm to prove the feasibility of our approaches. In Chapter 6 we will show our experimental test-bed and explain how we use it for our experiments in subsurface mapping and buried object retrieval. Finally, chapter 7 contains our conclusions and future directions for our research.

1.2. Research Objective

Our research objective is the following:

> ***"Develop an intelligent robotic subsurface mapper that can autonomously detect, locate and compute geometric parameters of buried objects"***

We will refer to the problem of detecting, locating and computing geometric parameters of buried objects as the problem of "Subsurface Mapping". The geometric parameters of a buried object include its size, shape, and 3-D orientation

To achieve the objective, the robotic subsurface mapper must be able to satisfy the following requirements:

1. **Scan the soil surface and build a terrain elevation map to guide sensor placement.**
2. **Scan the subsurface using a subsurface sensor.**
3. **Detect, locate and measure buried object from the subsurface sensor data.**
4. **Display the subsurface data and the buried object to the operator for notification and confirmation.**

If retrieval of the buried objects is necessary then we need two additional requirements:

5. **Expose the object by excavating the soil above the object.**
6. **Retrieve the object.**

This thesis concentrates on requirement 1, 2, 3 and 4 which address the problem of subsurface mapping for buried objects. One of our success criteria is to shorten the time needed for subsurface mapping, so the solution to these requirements should be faster than manual approaches.

We are not directly addressing the requirements 5 and 6 for retrieval of buried objects, but we take into account the fact that the subsurface mapping process must be able to be integrated well with the retrieval process. In fact, we found that the ability to remove the soil above the buried objects actually makes the subsurface mapping problem easier. We will show how we can get a more accurate subsurface map by repeatedly removing a layer of soil above the objects and rescan the buried objects.

Although we are developing an autonomous system for building subsurface map. It is possible to keep a human operator in the loop for safety and high-level planning. Instead of gathering and interpreting the data, the operator acts as a supervisor for the intelligent robot. In critical operation involving unexploded ordnance retrieval, an expert operator can help guiding the sensing and retrieval process so the fuse of the ordnance can be located. In cases where the system fails to correctly interpret the subsurface data or fails to plan an appropriate automated excavation, an expert operator can also help the system by providing additional information.

1.3. Technical Approach

In this section, we will address our technical approach to subsurface mapping. We will describe the sensor that we use to sense buried objects and how we process the output of the sensors to find the buried objects. We also address the very important issue of integrating subsurface mapping with buried object retrieval process by proposing a technique that can be used to create a more accurate subsurface map during the retrieval process.

1.3.1. Subsurface Mapping

Our technical approach to the subsurface mapping problem uses a robotic mechanism to move the sensor and automated data processing techniques to find the buried objects in the sensor data. In our work, the robot consists of an industrial manipulator as the scanning platform for the sensor, a scanning laser rangefinder to map the terrain shape, and a sensor mounted at the end effector to sense buried objects. First, the robot maps the terrain shape using its scanning laser rangefinder. It uses the elevation map to guide the manipulator so it can follow the contour of the terrain during scanning. The robot also tags the sensor data with the sensor position, so we will be able to accurately register the position of a buried object in the sensor data to its actual location.

We use Ground Penetrating Radar (GPR) as the robot's subsurface sensor. GPR is one of the most versatile sensors for sensing buried objects. It can sense nonmetallic as well as metallic objects under various soil conditions. When it scans a 2-D area, it produces a 3-D volume data. With suitable processing, this 3-D volume data can provide us with the location, orientation, size and shape of the buried objects. Currently, most GPR data are interpreted manu-

ally by human experts. They examine the data to find the buried objects, and compute their location, orientation and shape. This is a very time consuming process and prone to interpretation errors. We suggest that a better solution would be to automate the interpretation process. To achieve this, we have developed and implemented three new algorithms that can automate the process of finding buried objects in GPR data, and computing their location, orientation, size and shape. These algorithms are based on 3-D computer vision methods, and they reduce the GPR 3-D volume data into a few object's parameters.

Two of these algorithms directly process the volume data to find the buried objects. We call this approach, "Volume Based Processing". To further accelerate the execution times of the algorithms, we modified one of the algorithm so it can be run on multiple processors. Due to the local nature of the computation, the 3-D data can be split up into smaller pieces and each pieces can be computed on different processor. So by adding additional processors, we can reduce the execution time of the algorithm. This is true until the number of processors becomes large enough that the communication between the processors become a bottleneck. In our experiment we use as many as 10 processors to run our algorithm without experiencing communication bottleneck.

The third algorithm reduces the 3-D volume data into a series of possible objects' surfaces and then uses model based recognition techniques to determine if any of these surfaces belongs to a buried object. We call this approach "Surface Based Processing". This approach is much less sensitive to the problem caused by the soil inhomogeneity, since it finds the objects by detecting their shapes. The shapes appear similar under various soil conditions.

Using these algorithms, along with automated data gathering, the robot can automatically build the subsurface map of buried objects. The steps that we describe above is illustrated in Figure 1. As shown in the figure, the subsurface map produced by our algorithms, contains

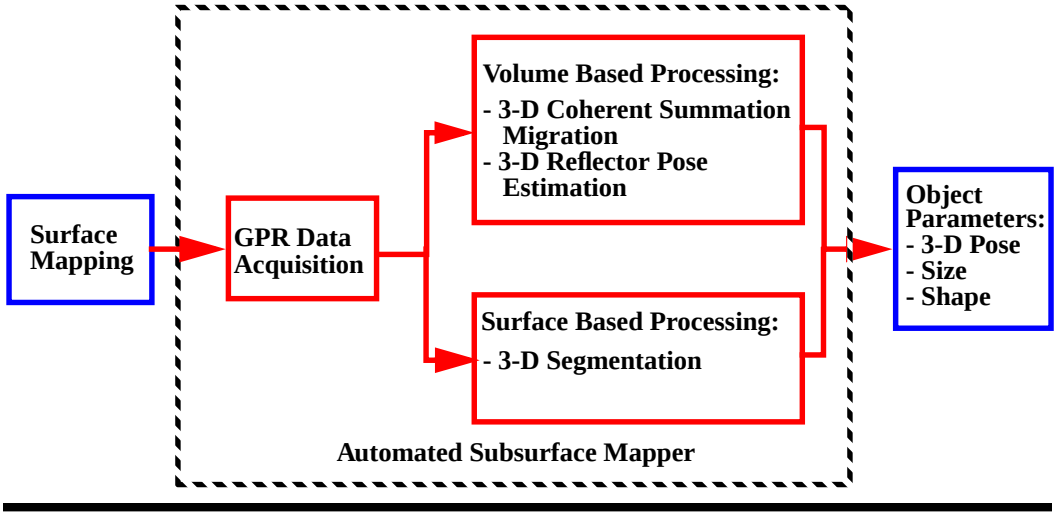

Figure 1: Proposed approach for autonomous subsurface mapping

some parameters that are previously very hard to get. For example, our automated algorithms can easily compute the object's 3-D orientation from the 3-D GPR data. In order to obtain the same information using manual techniques would be very time consuming because multiple sections of the 3-D data must be examined to compute the 3-D orientation of a buried object.

1.3.2. Integration of Subsurface Mapping and Buried Object Retrieval

In some cases, subsurface mapping is not enough, we also need to retrieve the buried objects. During the retrieval process, it is much more important to have a highly accurate subsurface map. Error in the position estimate of the object may cause collision between the excavator bucket and the buried object. The acceptable error in the position estimate of the object depends on the distance of the excavator bucket and the buried object. When the excavator bucket is digging far away from the buried object, even a large relative error in the position estimate on the object is acceptable. As the excavator removes layers of soil above the object and gets closer to the object, we need to have a more accurate estimate on the position of the object.

Our solution to this problem uses repeated "Scan and Dig Cycle". During each cycle, the robot rescans the area, regenerates the subsurface map and removes a layer of soil. After every cycle, the robot gets closer to the buried object and there are less soil between the sensor and the object. Since soil inhomogeneity is one of the main source of error, less soil between the sensor and the object translates to a smaller error in the position estimate of the object. As a result we can gradually improve our position estimate of the buried object.

Figure 2 illustrates this concept. The robot consists of a computer controlled excavator with a subsurface sensor attached to its bucket. It moves the bucket in order to scan an area using the sensor. Our algorithms then process the scanned data to detect and locate the buried objects. After an object has been located, the robot would remove a layer of soil above the object and rescan the are to improve the estimate on the object's location. It continually repeat this "Sense and Dig Cycle" until the object is very close to the surface of the soil (Figure 2d). At this point it will retrieve the object.

The removal of soil serves multiple purposes. First, it needs to be done for the robot to retrieve the buried object. Second, it enables the sensor to get a better scans of the object by getting closer to it, thereby improving the accuracy of the subsurface map. Finally, by comparing the scans gathered before and after removal of each layer of soil, we can obtain a better estimate of the soil parameters. As far as we know, this thesis is the first work which addresses both issues of automatically processing 3-D GPR data to find buried objects and integrating the mapping process with the soil removal to improve the estimate on the parameters of the buried object and soil.

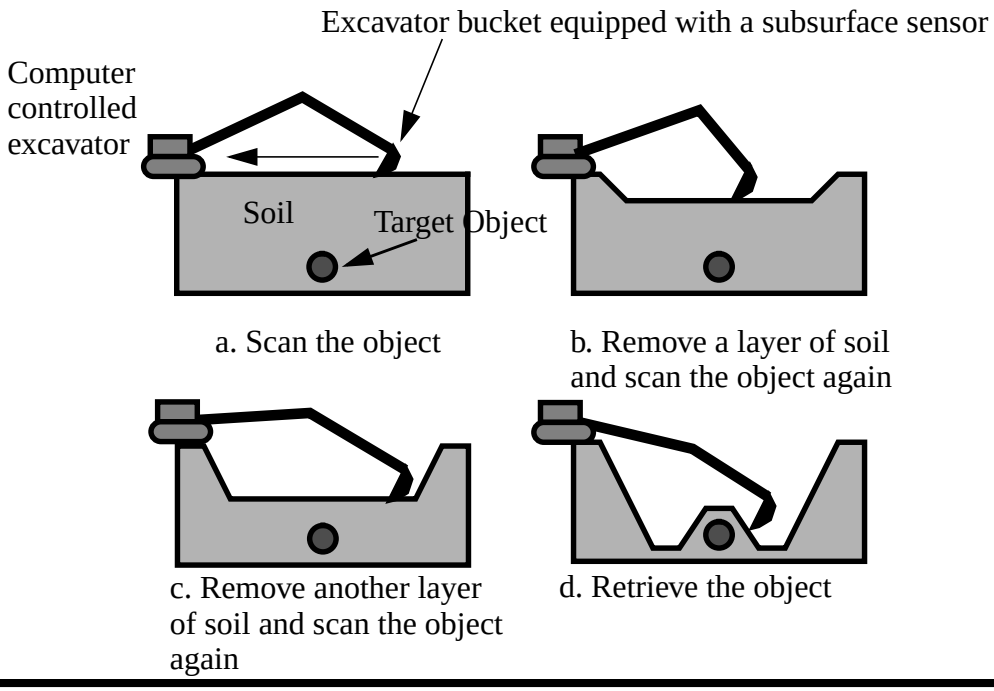

Figure 2: The scenario for retrieving buried object using sense and dig cycle

The actions during the sense and dig cycles can be seen in Figure 3. The main assumption of this approach is that the errors in the subsurface map decrease as we get closer to the buried objects. The errors can be caused by a wrong GPR propagation velocity estimate and noise from spurious reflections. Intuitively we can say that as the amount of soil between the antenna and the object decreases, there are fewer uncertainties in the GPR output. Therefore we should be able to get more accurate information as we get closer to the object.

This approach is in contrast with existing approaches which try to obtain an accurate and high resolution subsurface map using a single scan. These existing approaches often fail because the soil is not homogenous, the penetration depth of the GPR signal is shallow and the difficulty in interpreting GPR signals that are reflected from deeply buried objects. The biggest problem with just doing a single subsurface scan in the beginning of the retrieving process is in obtaining an accurate position and orientation of the buried object. Since the buried objects may be located at a significant distance from the surface, there are a lot of uncertainty in the medium between the surface of the soil and the buried object. This uncertainties cause error in the position and orientation estimate of the buried objects. By doing multiple subsurface scan each time a layer of soil above the object is removed, we can continually improve the position and orientation estimate. In addition, we can compute a more accurate parameters of the soil characteristic as we dig deeper to the soil.

Figure 4 shows the architecture of our integrated robotic subsurface mapper and buried object retriever. There are 4 main subsystems.First, we have the elevation map generator, which scans the ground surface to generate an elevation map. The subsurface mapper uses the elevation map to generate the path for the scanning motion of the sensor. The path is executed by the robotic excavator which is equipped with a subsurface sensor at its end effector. The same robotic excavator is also used for excavating the soil.

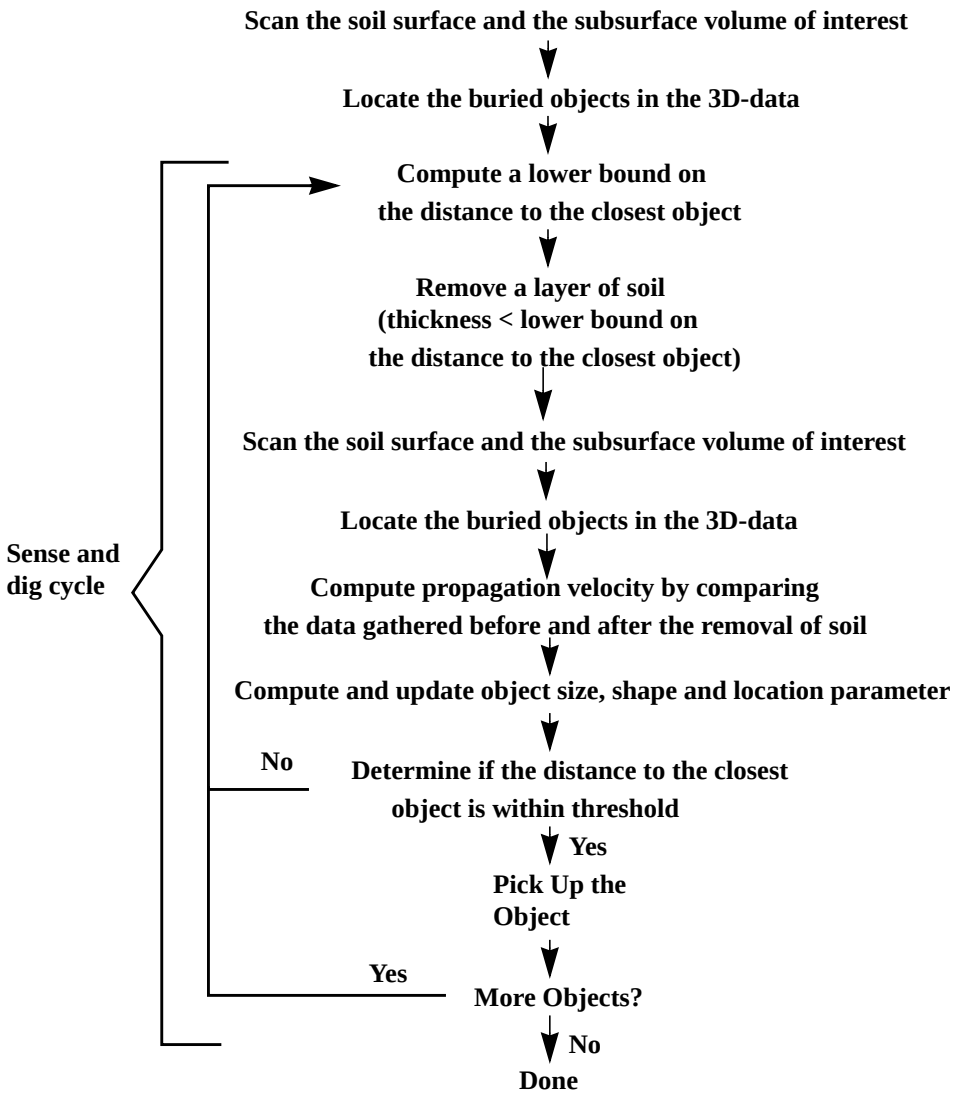

Figure 3: Processing steps within the sense and dig cycle

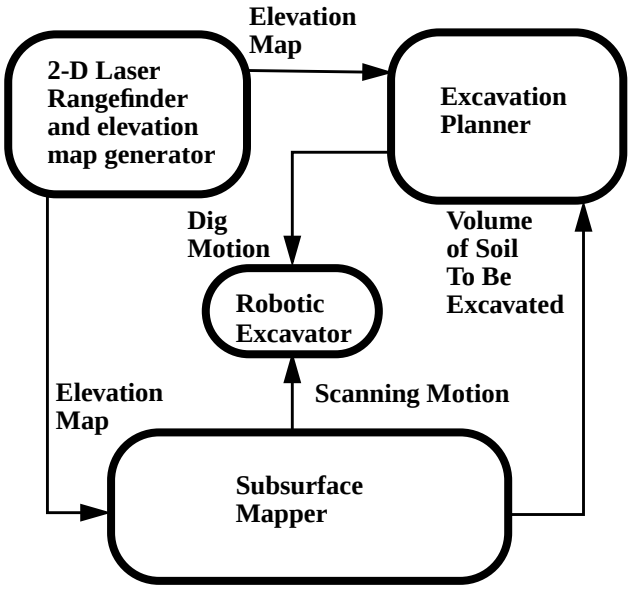

Figure 4: System Architectures

1.4. Rationales

Although subsurface mapping can also be done using manual methods, there are several important rationales for using an autonomous or semi autonomous system to build subsurface map. They can be categorized into several different categories:

1.4.1. Improved safety

By having an autonomous system, we can remove the human operators from the operation site, thus reducing the possible danger to the operators. This is especially true for mapping sites which contain potentially explosive, radioactive or toxic materials. Although the safety problem can also be alleviated using teleoperation, the latency and the bandwidth limitation for low level communication between the teleoperated machine and the operator limit the type of work that can be done. Autonomous and semi autonomous systems offer much more flexibility because the communication between the machine and the operator can happen at several different levels, each of which can be tailored to the task.

Safety is also improved by reducing the possibility of human error in interpreting the subsurface sensing output and in registering the objects' location in the subsurface map with its

actual location in the world. This is possible by using the same mechanism for mapping and excavating, which will eliminate most of the registration error.

1.4.2. Increased productivity

A fully autonomous system could, in principle, operate continuously day or night. We can also have multiple systems operating in parallel to speed up the operation. Due to the absence of human in the operation area, fewer safety precautions need to be taken, which should also increase the efficiency of the retrieval task. All of these factors contribute to the increase productivity in term of man hours required for the work.

1.4.3. Cost saving

Many of the applications of this work require mapping and retrieving buried objects in a wide area, which could easily reach several square miles. Due to the large scale of the problem, any increase in productivity should result in significant saving in both time and money. We will also save quite a lot of time and money since the automated system can be operated by operators with less expertise and skill. This is possible because the difficult process of data interpretation and low level machine control are done autonomously by the computer. Autonomous system usually incurs a large one time cost, which is also called the non recurring engineering cost. Once it is working, it can be duplicated at a reduced cost. On the other hand, a manual system needs experts to operate, which means that each new additional system requires training new experts.

1.4.4. New capability

An integrated mapper and excavator will be able to do precise operations that is not possible with manually operated equipments. Due to the precise information about the object's location and orientation gathered by the mapper, the excavator will be able to excavate soil very close to the buried object without actually touching the object. Our new improved subsurface data processing techniques also generate the object's location and orientation in 3-D, compared to existing techniques which mostly generates 2-D information.

1.5. Applications of the Robotic Subsurface Mapper

This work can be applied to many tasks that require subsurface sensing and/or retrieval of buried object. The following are some example applications in several distinct categories:

1.5.1. Subsurface Mapping

1.5.1.1. Mapping of subsurface utility structures

For this application, the robotic mapper builds the map of subsurface structures such as gas pipes. The subsurface data can be obtained by scanning in a regular grid or by tracking certain subsurface features, for example by tracking the buried gas pipe individually. Currently this is done by metal detector or by manual ground penetrating radar (GPR) operation. Metal detector does not give depth and it only works for metallic pipes. Manual operation of GPR has its own shortcomings, such as the need for expert operator and the difficulty in getting accurate registration between the location of the pipes in the GPR data and their actual locations in the world. It is also hard for even an expert to detect some features in the GPR data.

1.5.1.2. .Detection and mapping of unexploded ordnance and mines

A robotic subsurface mapper would be very useful in detecting and locating landmines. A robotic subsurface mapper can be deployed in advance of troops to identify a safe route. Currently landmine detection and localization are done manually using hand-held metal detectors or mechanical probes. The manual operation is very dangerous and is done at a very slow pace. Using a robotic landmine mapper, the operation can be made faster by automating the manual data collection and interpretation task. In addition, we are not risking any human life in trying to detect and locate the landmines.

1.5.2. Retrieval of Buried Object

1.5.2.1. Retrieval of hazardous waste containers or unexploded ordnance

In this application, the robot needs to map the buried objects, compute their shape and orientation, and generate a plan to remove them. In essence, this application is a continuance of the detection and mapping of unexploded ordnances or mines. In this application the robot does not stop when the subsurface objects are detected and located, but it proceeds to determine their shape and orientation. It uses the additional information to generate a plan to extricate or neutralize the unexploded ordnance or landmines. Automated scanning and interpretation are perfect for this application because of the reduced possible error in registering the location of the object in the GPR data and its location in the real world. The automated scanning can also collect a very high resolution 3-D data which should increase the accuracy of the subsurface map.

1.5.3. Collision prevention in excavation

1.5.3.1. Maintenance or repair of subsurface structure

In maintaining subsurface structures such as electrical lines, phone lines, or gas pipes, construction crews often need to excavate the soil around the structure. In the process of doing so, they sometimes hit the structure or other structures that are on their way. For example: a construction crew from a gas company might have an accurate map of the gas pipes, but during the excavation process, the crew might hit and break an electrical line. To prevent this from happening, the excavator needs to know that the next volume of soil to be excavated is devoid of any buried objects. So this problem is actually a little bit simpler than the buried object retrieval problem, since in this application the robotic subsurface mapper only needs to confirm that a certain volume of soil is devoid of any buried object.

Chapter 2. Related Work

2.1. Subsurface Mapping

The use of subsurface sensor as a sensing modality has received very little attention in robotics compared to other sensing modalities such as video images, range images or sonar. Therefore, it is not surprising to find that the proposed robotic subsurface mapper would be one of the first robotic systems to use a subsurface sensor as one of its sensing modalities. In this case, the use of the subsurface sensor enables the robot to see through certain solid medium, such as soil.

While very little work has been done in automated gathering and interpretation of subsurface data, there have been quite a lot of work in manual subsurface data gathering and interpretation. In the beginning, subsurface sensing is mainly used for geological explorations and landmine detections. These are done primarily using sound waves echo recorders or metal detectors. Many aspects of these two applications are at opposing extremes. Geological exploration equipment uses sound waves to scan a very large area, which could easily reach several square miles. The output of the scanning operation is large and usually used to map the macroscopic geological features. On the other hand, landmine detection using a metal detector operates on a much smaller scale. It is usually a point sensor that could detect a metal object underneath it. The sensor size is usually not more than 1 feet in diameter and the output of the sensor is usually only a single value denoting the strength of the signal

from the metal detection circuit. Subsequently, magnetometer is also used to detect and locate buried objects. It works by measuring the disturbance created by the buried objects on the earth magnetic field. Most recently, Ground Penetrating Radar (GPR) is also used for the detection and localization of buried objects. Lord et al. [Lord 84] did a good overview of these various subsurface sensing techniques and their characteristics.

Among all the above sensors, GPR might be the most versatile one. As a testament to its versatility, the variety of its uses has increased significantly during recent years. GPR has been used in numerous diverse fields such as archaeology [Imai 87], geology [Davis 89], non-destructive testing [Beck 94][Davis 94] and engineering [Ulricksen 82]. Some of the specific tasks include mapping soil stratigraphy [Davis 89], probing underground caves [Deng 94][Vaish 94], detecting landmines [Ozdemir 92], testing roads and runways [Beck 94][Davis 94][Saaraketo 94], mapping pipes and drums [Lord 84][Osumi 85][Gustafson 93] and locating persons buried under snow [Yamaguchi 91]. Peters and Young [Peters 86] give good examples of the diverse applications of GPR, while Ulricksen gave a good overview of the application of GPR in Civil Engineering [Ulricksen 82].

Most of the GPR data gathering and processing is currently done by manually scanning the area of interest with a handheld antenna or antenna towed by a human-operated motorized vehicle [Ulricksen 82][Bergstrom 93]. After the data have been obtained then experts examine the data to find the buried object. These two operations are very time consuming and prone to many human errors. After the buried object are located in the map, the real location of the object in the world must be determined. Using this information, a human operated excavating device can remove the soil above the objects and retrieve the buried objects. (end)

An important part in solving the detection and mapping of subsurface objects using GPR is the understanding on how a GPR pulse travels through the soil and is reflected by buried objects. This involves modeling the GPR system, signal propagation and reflection [Kim 92]. This work is important because through a better understanding of how a GPR system works is critical in making a better GPR system. Since GPR is a form of radar, many processing techniques in radar can also be used for GPR. Therefore it is important to understand the works that have been done in the field of radar signal processing. This is especially important because conventional radar signal processing is a mature field compare to GPR signal processing. A good review of modern radar and its processing methods can be found in the work by Eaves and Fitch [Eaves 87][Fitch 88].

As for the processing of GPR data, researchers have also experimented with multiple techniques to improve GPR data. Due to the similarities between GPR and seismic sensing technique, there have been some efforts to apply seismic processing methods to GPR, a good example is the work done by Chang [Chang 89] and Fisher [Fisher 92]. One disadvantage of

the seismic processing technique is the massive amount of required computational power. To alleviate this problem, Fiebrich [Fiebrich 87] and Fricke [Fricke 88] have implemented a seismic processing method called "Reverse Time Migration" on massively parallel supercomputer. Other optimization methods are also explored to reduce the needed resources in applying the seismic processing technique [Harris 92]. Due to this massive computational requirement and other differences between GPR and seismic [Daniels 93], seismic processing methods are ill-suited for realtime subsurface mapping.

Another processing technique that has been studied is the inverse scattering technique [Moghaddam 92][Oh 92]. This method requires a lot of computational power because the number of variables that need to be solved during the inverse scattering calculation is very large. This method also often requires a borehole antenna configuration which involves drilling into the ground. Due to these two problems the practical uses for this method are limited.

There are also some works dealing with the extraction of buried object's parameter from GPR data. Gustafson used a semi-automated method to do velocity depth inversion for cylindrical objects [Gustafson 93]. He needed to select the direction of the scanning profile manually. The reflections used in the velocity-depth inversion calculation are obtained by simple thresholding. Thus, even reflections that come from non-cylindrical objects are included in the computation if their strength exceeds the threshold. Roberts also used a semi automated method for velocity-depth inversion to calculate the EM wave propagation velocity in the soil [Roberts 93]. He assumes the reflection profile comes from a point reflector. The main disadvantage is that the direction of the scanning profile needs to be manually selected by the operator. Another work in localization of buried object is the work by Stolte which uses 2-D migration to highlight the location of buried pipes in the GPR data [Stolte 94]. In our research, we go one step further and uses high resolution 3-D data to find buried objects. The use of 3-D data in our research eliminates the need to know the orientation of the buried objects in advance.

These existing works have only addresses some of the issues that we need to solve, such as the poor lateral resolution of the GPR and the huge amount of data that need to be processed. As far as we know, no work has been done in the automated detection and mapping of buried objects using high resolution 3-D GPR data.

In our research, we use some of the existing methods, such as migration, in combination with new techniques adapted from computer vision field. To the best of our knowledge we are the first to use 3-D object recognition techniques to automate the GPR data interpretation process. It is also important to note that all our work is performed with high resolution 3-D volume data instead of 2-D data, although the methods can be modified so they can be applied to 2-D data set as well.

Recently, other researchers have begun to realize the potential for a high resolution 3-D GPR imaging [Daniels 93][Ulricksen 94]. In fact, Ulricksen has built a high resolution 3-D GPR scanning mechanism which is similar in many ways to our system [Ulricksen 94]. He uses a scanning mechanism with multiple antenna configurations to obtain a high resolution 3-D GPR data but he did not automate the process of buried objects detection and mapping.

There have also been some efforts in automating subsurface mapping using other type of sensors. One such system is the portable pipe mapper developed at Carnegie Mellon University by Lay et al. [Lay 87]. The portable pipe mapper uses electromagnetic induction technique, which works well for mapping ferrous pipes, but can not be used to map other kinds of buried objects.

Although they are not directly related to subsurface mapping, ultrasound and CT scanner also generate 3-D data which can be used to image the inside of a solid object. Many researchers have developed methods for automatically processing the ultrasound and CT scan data [Thomas 91][Sonka 95]. Similar with our GPR processing methods, these processing methods extract the object's parameters from the 3-D data. In this case, the object can be a blood vessel or an organ inside a human body. Unfortunately, the characteristics of 3-D data from ultrasound or CT scan are completely different with the characteristics of 3-D GPR data. Ultrasound or CT scan data are usually gathered using multiple transducers around an object. They usually work by measuring how the signal is propagated through the object. So they do not use the reflection of the signal to image the inside of the object. On the other hand, GPR data are usually gathered using one or two transducers. In addition, these transducers can only be positioned above the surface, which limits their viewing direction considerably. GPR also rely on the reflected signal to detect buried objects in the soil. Due to these differences, the ultrasound and CT scan processing methods are not directly applicable to GPR data.

2.2. Automated Excavation and Buried Object Retrieval

In the field of automatic excavation, there have been some works dealing with excavation planning, control and soil modeling. Apte discussed a representation for spatial constraints in material removal and its application to automatic mining and lawn mowing [Apte 89]. More recent work by Singh examines task planning for robotic excavation [Singh 92]. It looks for a set of digging movements that will efficiently excavate a given volume of soil using optimization methods. In our research we concentrate just on the detection and mapping of buried objects. During our experiment of buried object retrieval using our robotics subsurface mapper, we use the system developed by Singh [Singh 92] in controlling the robot for automated excavation.

There also have been some work in automatic pipe excavation by Whittaker et al. [Whittaker][Swetz 86]. In this case, no subsurface sensor is employed; the retrieval is done by carefully removing thin layers of soil until the pipe is detected by sonars. This is not suited for sensitive and explosive objects since the objects might be struck before they are detected by the sonars.

Chapter 3. Ground Penetrating Radar

3.1. Selection of Subsurface Sensors

There are many sensors that can be used to sense different types of buried objects. These sensors range from a simple device such as a metal detector to a very complicated chemical detection device such as a Thermal Neutron Activation (TNA) sensor. These sensors use different sensing modalities which are suitable for different tasks. A list of subsurface sensors along with their sensing modalities, advantages and disadvantages is listed in Table 9.

From these various sensors, we pick GPR because it can sense metallic and nonmetallic buried objects. It can also obtain geometric information about the buried objects, as well as its depth and horizontal location. GPR is also operable in propagation medium other than soil, such as fresh water [Mellet 93], which enables us to find buried objects at the bottom of lakes or rivers. All of these capabilities make GPR a much more versatile sensor compared to other subsurface sensors.

3.1.1. Method of Operation

Ground Penetrating Radar (GPR) is a special radar unit that is designed to transmit electromagnetic pulses down to the ground instead to the air. Therefore, the main difference with a conventional radar is the medium of propagation. In the case of a conventional radar, the

Sensor	Advantages	Disadvantages	Sensing Modality
Ground Penetrating Radar (GPR)	• High resolution 3-D volume data output (horizontal location and depth) • Very flexible transducers configuration • Able to detect metallic and nonmetallic objects	• Moderate cost • Processing intensive • Penetration depth depends on the characteristic of soil	Distance to metallic and non-metallic objects along with reflected signal intensity
Magnetometer	• Low cost • Simple processing • High sensitivity	• Detect only ferrous object • Mainly 2-D data (the horizontal position of the buried object)	Magnetic field
Metal Detector	• Low Cost • Easy to operate • Mature technology	• Detect only metallic object • Mainly 2-D data (the horizontal position of the buried object)	Electromagnetic interference caused by metallic objects
Seismic Sensors (Sound wave geophones and pulser)	• Large area coverage • Mature technology	• Low resolution due to the long wavelength of the signal that is used for mapping	Distance to major geological features along with reflected signal intensity
Chemical Sniffer such as Thermal Neutron Activation unit	• Very reliable chemical discrimination	• No depth information is available	Chemical discrimination

Table 9. Subsurface sensor technologies

medium is the air, while in the case of GPR, the medium is the soil. Since soil is much more heterogeneous and has much higher attenuation rate than air, the data generated by GPR is of much poorer quality than the ones generated by conventional radar. Due to the poorer quality data, processing the data becomes much more difficult. It also accounts for the need to have a different approach in processing GPR and conventional radar data.

A block diagram of a typical impulse GPR unit is shown in Figure 5. There are four main parts: the pulse generator, transmitting antenna, receiving antenna, and the receiver. The pulse generator produces a short pulse which is emitted by the transmitting antenna. The emitted electromagnetic pulse travels through the propagation medium until it encounters a

discontinuity in the electrical properties of the medium. The discontinuity can be caused by an object with different electrical characteristics with the medium. When the pulse hits the object, part of its energy is reflected back to the receiving antenna. The reflected electromagnetic pulse is converted by the receiving antenna to an electrical pulse. The electrical pulse is then recorded by the receiver. Since it is possible to have multiple objects in front of the antennas, the receiver produces a continuous wave signal which consists of superimposed reflected pulses. Figure 5 shows a separate transmitting and receiving antenna, which is called a bistatic configuration. We could also use a single physical antenna for transmitting and receiving by connecting the antenna through a multiplexer to the pulse generator and the receiver. The single antenna configuration is called a monostatic configuration.

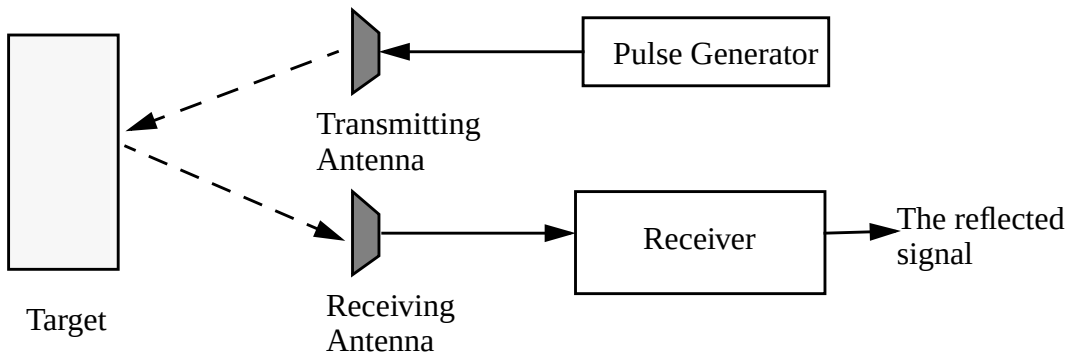

Figure 5. A simplified block diagram of an impulse GPR system

Figure 6 shows the operation of GPR in detecting different layers of soil or buried objects. The GPR antenna is placed very close to the ground in order to achieve maximum energy coupling. When the transmitted pulse hits an interface between two mediums with two different dielectric constants, part of the energy is reflected and part of it is refracted.

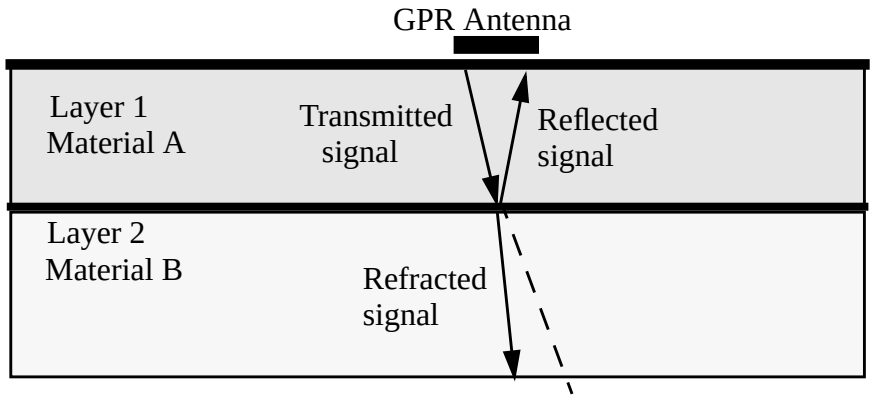

Figure 6: GPR signal reflection and refraction at the interface of two different layer of soil

The strength of the reflection and the refraction is determined by the electrical characteristics of the two materials. The interface can be caused by different soil layers or by buried objects which have different dielectric constants with the soil.

For an impulse GPR, to detect and determine range to the buried object, we must detect the reflected pulse in the output of the receiver. The distance is determined from the time difference between the start of the transmitted pulse and the start of the reflected pulse. The time difference multiplied by the propagation velocity of the signal in the soil results in the distance to the object. For large objects that return strong reflections, simple thresholding works to detect the reflected pulse in the receiver output. For weak reflections, other methods must be used, which will be discussed further in later chapters.

Although most of the current commercial GPR units transmit electromagnetic pulses, there are also other types of GPR which transmit continuous wave (CW) signal or frequency modulated (FM) signal. In CW radar, the range is obtained from the phase difference between the transmitted and reflected signals. In FM radar, the range is obtained from the beat frequency produced when the reflected signal is mixed with the transmitted signal.

Figure 7 shows an example of an impulse GPR signal consisting of the transmitted pulse emitted from the antenna and the reflected signal. At the beginning of the signal, we can see the transmitted pulse due to coupling between the transmitter and the receiver.

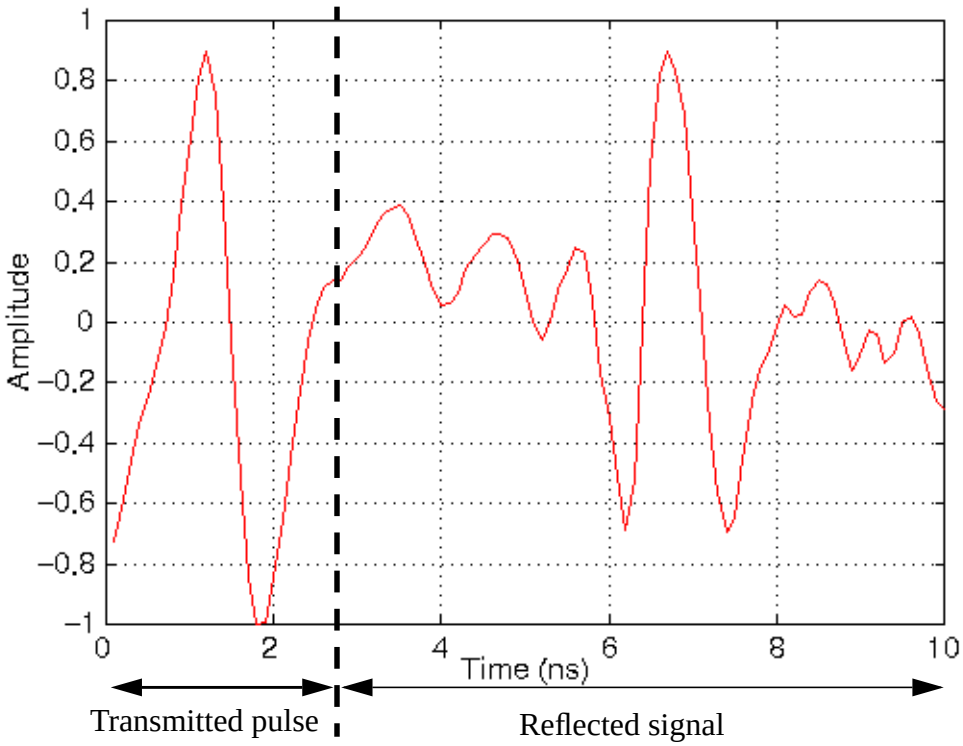

Figure 7: An example of a GPR signal consisting of the transmitted signal and the reflected signal

We sample this signal to generate its digital representation. Two parameters in sampling the signal are amplitude and time resolution. Currently, we sample the amplitude of the signal every 0.02 ns using 12 bits of data, which means that we can differentiate 4096 different amplitude levels. The 0.02 ns sampling interval translates into a spacing of 0.3 cm assuming a vacuum medium (15 cm/ns round trip propagation velocity). In sand, the same sampling interval translates into a 0.15cm, assuming a round trip propagation velocity of 7.5 cm/ns.

Figure 8 shows a series of GPR scans of a metallic plate at various distances from the antenna. As the distance between the antenna and the plate increases, the propagation times increase and the reflected pulses occur at later times in the reflected signal. If we multiply these propagation times with the propagation velocity, we will obtain the distances to the plate. Therefore, it is important to be able to determine the start of the reflected pulse to compute the correct propagation time. It is also important to have an accurate estimate of the propagation velocity.

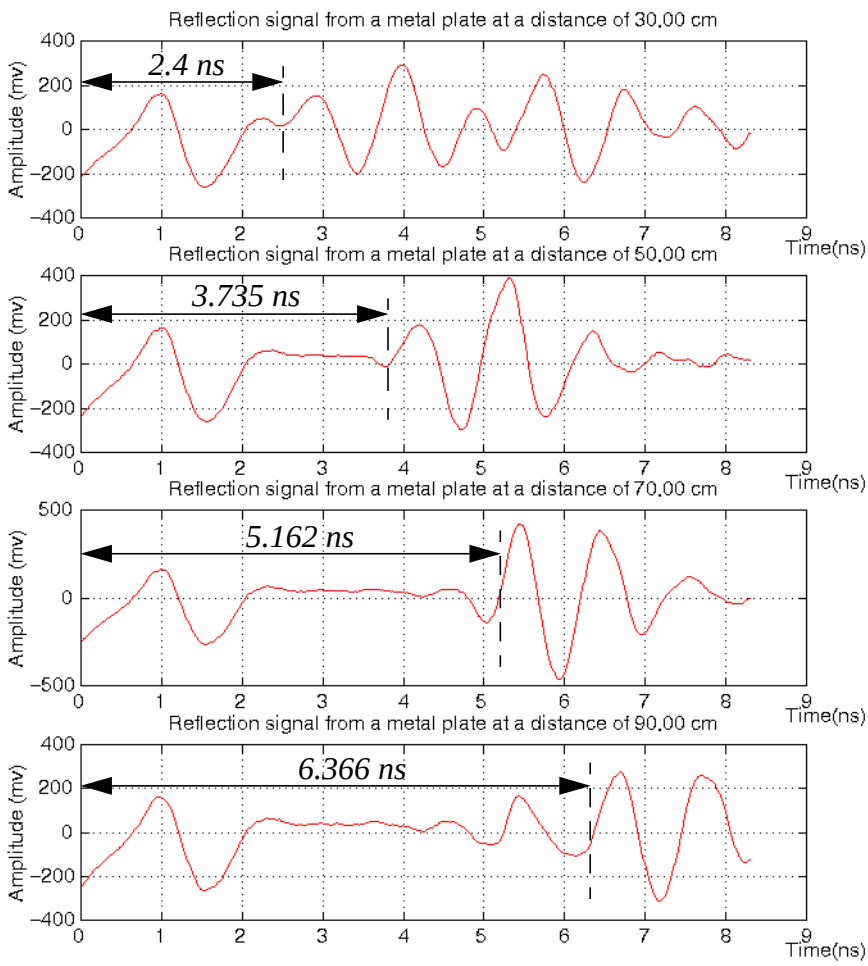

Figure 8. A series of GPR scans obtained by placing the antenna in front of a metal plate at a distance of 30cm, 50cm, 70cm and 90cm.

Another important choice to be made in determining the configuration of a GPR system is whether to use a monostatic or bistatic configuration for the antenna. A monostatic configuration uses a single antenna for both transmitting and receiving the electromagnetic pulse. On the other hand, a bistatic configuration uses one antenna for transmitting the pulse and another antenna for receiving the reflected pulse. A monostatic configuration is simpler to operate but a bistatic configuration does offer a number of advantages. In a bistatic configuration, noise from the soil heterogeneity near the surface can be minimized. This can be seen in Figure 9. The first interface is outside the intersection of the transmitting and receiving

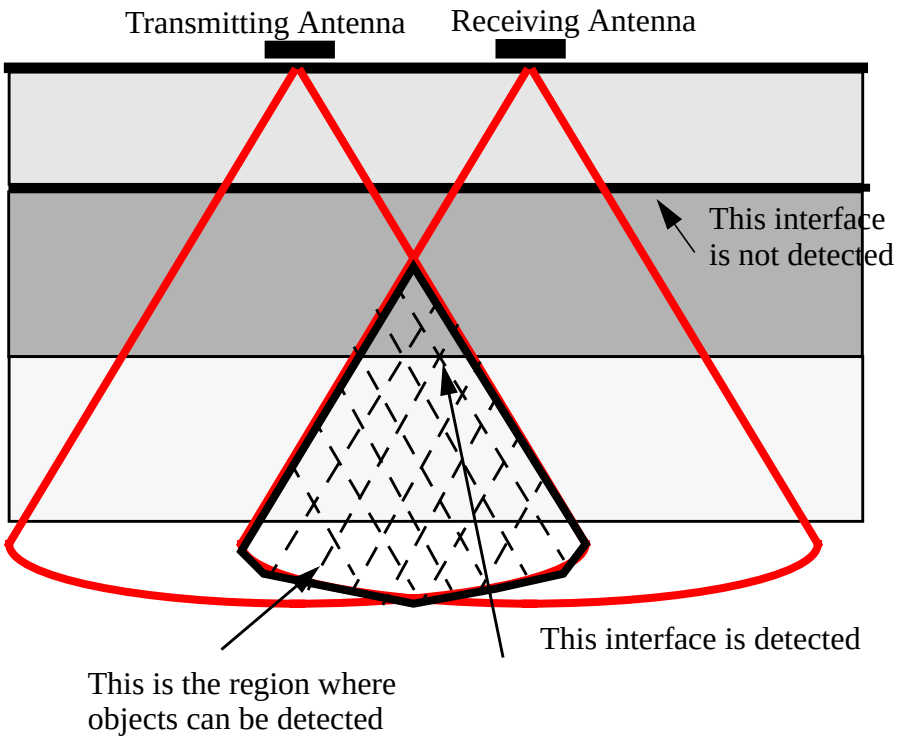

Figure 9: Sensitivity region of bistatic configuration

antennas' beamwidth and therefore, it is not detected. A section of the second interface can be seen by both the transmitting and receiving antenna, so it can be detected. So this configuration could prevent spurious reflections from disturbances near the surface of the soil. On the other hand, if the buried objects that we are interested in are located near the surface of the soil, a monostatic configuration is more suitable. Another advantage of bistatic configuration is that by using several scans taken with different antenna separations, it is possible to measure the propagation velocity of the electromagnetic signal in the soil. This is done by a method called "Common Mid Point" (CMP) method. In monostatic configuration, the propagation velocity must be obtained using other methods, such as using a soil's dielectric measurement device. The bistatic configuration can also be expanded to more than two antennas

where we have m different transmitting antennas and k different receiving antennas. Using this configuration, the CMP method can be done without moving the antennas at all because we can place different pairs of antennas at different spacings.

Bistatic configuration can also be used to collect polarimetric data. In order to do this, we need two antennas with orthogonal polarizations. If the first antenna has horizontal polarization and the second one has vertical polarization, we can collect data using three different polarizations. First, we can transmit and receive using the first antenna, resulting in horizontal polarization. Second, we can transmit and receive using the second antenna, resulting in vertical polarization. Finally, we can transmit using one antenna and receive using the other one, resulting in cross polarization. The polarimetric data can be used to differentiate objects of different shape, such as a pipe and a plate, because the reflections from those objects vary under different polarizations.

Regardless of the antenna configuration, the operating frequency of the GPR must be chosen carefully, since it influences the depth of penetration, antenna size, antenna beamwidth, and depth resolution. Low frequency signals can penetrate deeper because their attenuation rate is lower than high frequency signals. The size and the beamwidth of the antenna are also larger for lower frequency signals. In addition, low frequency signals result in lower depth resolution than high frequency signals. So a low frequency GPR is suitable for detecting objects that are buried deep in the soil. On the other hand high frequency signals can not penetrate as deeply into the soil, but offer a higher depth resolution, smaller antennas, and a thinner beamwidth. These characteristics make a high frequency GPR suitable for obtaining high resolution data of objects located close to the surface of the soil.

It is also possible for a GPR system to employ multiple frequencies. Such a system would use a low frequency signal to scan objects buried at great depths, to maximize the penetration depth, and a high frequency signal to scan shallowly buried objects, to maximize the resolution. This configuration is particularly suitable for buried objects retrieval because when the objects are buried deep in the soil, we need to be able to detect them but we do not need their precise locations. When the objects are closer to the surface then we need their precise locations because we need to approach the objects carefully in order to prevent collision. In our testbed, we only use a high frequency GPR since we are only dealing with shallowly buried objects.

3.1.2. Physical principles

The following are several main physical principles that govern the propagation of the GPR signal from the antenna to the soil and back to the antenna:

- Beamwidth is a function of antenna shape and size, for a dipole antenna the beam-

width is determined by the following equation:

$$Beamwidth = \frac{2\lambda}{D} \tag{1}$$

Therefore, a larger antenna has a smaller beamwidth and vice versa.

- Coupling efficiency is a function of antenna height. When the antenna is very close to the ground, it is coupled with the ground and the energy is transferred efficiently between the antenna and the ground.
- Beamwidth also depends on the soil electrical characteristics. This is especially true when the antenna is located at some distance from the soil. At the air and ground interface, the beam will be refracted. Characteristics of the soil determine the angle of refraction. A large angle of refraction will make the beam thinner while a small angle of refraction will keep the beam width more or less the same.
- At an interface between two different materials, the GPR signal experiences refraction and reflection. The strength and angle of the refraction and reflection depend on the incoming angle of the signal and the characteristic of the material. The incoming and the outgoing beam angle follows the Snell equation:

$$N1\sin(\alpha) = N2\sin(\beta) \tag{2}$$

$$N1 = material\ 1 \text{ refraction index}$$
$$N2 = material\ 2 \text{ refraction index}$$
$$\alpha = \text{incoming beam angle}$$
$$\beta = \text{refracted beam angle}$$

- The reflection strength at each interface also depends on the incoming beam angle and the characteristic of the two bordering materials.

3.1.3. Factors influencing GPR operation

3.1.3.1. Soil type

The depth of signal penetration depends on the electrical characteristics of the soil, such as attenuation rate, permeability, conductivity and dielectric constant. The heterogeneity of the soil also influences the quality of the image. Various types of soil have different electrical properties. For example sand has a very low permittivity, resulting in a low attenuation rate for electromagnetic wave in GPR frequency. Clay on the other hand causes a very high

attenuation rate of the GPR signal. As a result, GPR signal can not penetrate too deep in the soil where clay concentration is high.

3.1.3.2. Soil Moisture Content

Soil moisture content influences the electrical characteristics of the soil. Usually it increases the attenuation rate and in some cases a layer of soil with a very high moisture content can cause a strong reflection of the GPR signal.

3.1.3.3. Equipment parameters

GPR signal must have parameters that are suited to the task. These parameters include frequency of the signal, power, sampling rate, sampling resolution, antenna configuration and signal to noise ratio of the system.

3.1.3.4. Antenna to ground coupling

The height of the antenna, the orientation of the antenna and the shape of the soil under the antenna all influence the coupling of the electromagnetic energy from the antenna to the soil. For the deepest penetration depth, the amount of energy that is coupled to the ground must be maximized. In order to get maximum energy coupling between the antenna the antenna must be pointing normal to the ground plane and located very close to the ground.

3.1.3.5. Antenna characteristics

Included in this category is antenna shape and other electrical characteristics. Depending on the shape we can have different radiation patterns which affect antenna gain. For example, a horn is usually used if the antenna must be placed at some distance from the ground, such as in the case of high speed highway inspection. Another example is the parabolic antenna that can detect objects buried at some distance from the antenna [Peters 93].

3.2. GPR Data Collection and Data Format

The basic GPR output is a 1-D time varying signal from the GPR receiver. The amplitude of the 1-D signal is the strength of the reflected signal at that point. To obtain 2-D and 3-D data, the antenna must be scanned above the area of interest in one or two directions. Usually we move the antenna in a raster scanning pattern. In a typical manual GPR survey, the scan density is often much higher in one direction than the other direction. Our work, on the other hand, relies on the fact that we can obtain a dense 3-D volume data by scanning the antenna in a raster scanning pattern using very fine spacings in both directions as shown in Figure 10. This is possible because our GPR antenna is mounted on a highly accurate robotic scanning mechanism. This 3-D data are stored as a 3-D array where the value of the voxel is the

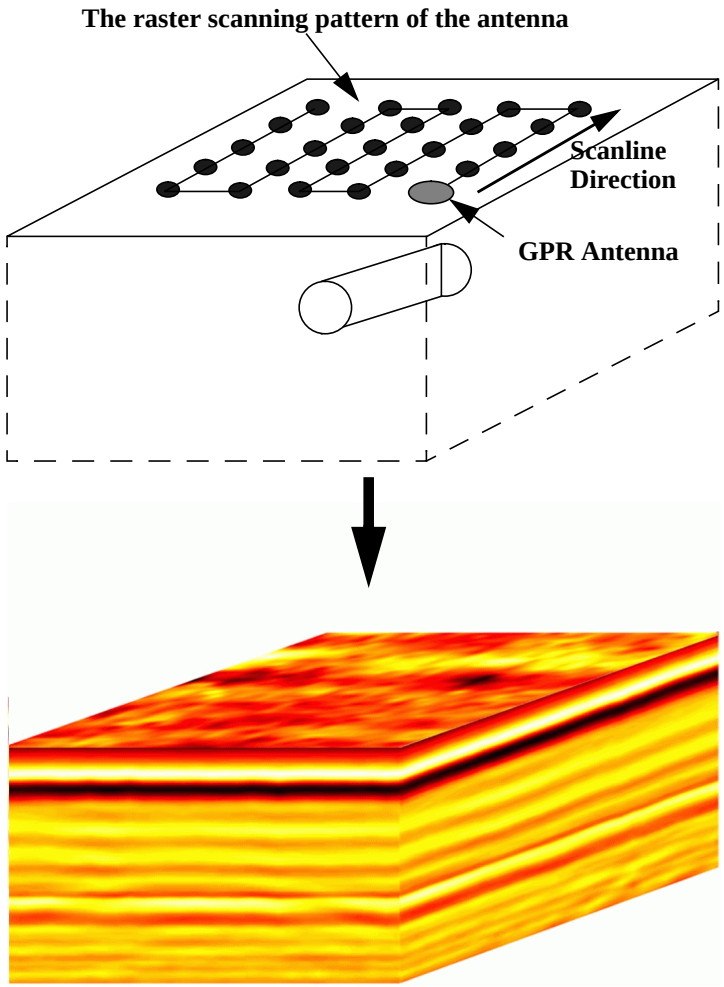

Figure 10: A 2-D scan above the buried object to obtain a 3-D volume data of the buried object

signal amplitude at that location. Each column of the 3-D grid is a single GPR scan as shown in Figure 7. The high density of the 3-D scans produce significant amount of 3-D volume data. If the antenna is scan using 2 cm grid, then we have 2500 scans for each square meter. If the number of samples in the depth direction is 500, it means we have 1.25 million values to consider for each square meter. Due to the large amount of data, the efficiency of the processing algorithms is crucial to their applicability to real world problems. This also underscores the need for automated processing of the data because it is extremely time consuming or impossible for a human to go through the massive amount of data that is generated by the GPR system.With automated and efficient processing of the data, it is feasible to process a high density 3-D data in real time. By real time, we mean that the processing units are able to process the data as fast as the GPR system can generate the scans.

3.3. Analysis of Different Antenna Array Configurations

In this section we will analyze the characteristics of different antenna array configurations., which is important to determine the best antenna array for a certain task. Since we are interested in finding shallowly buried objects as well as deeply ones, we only considers monostatic antenna arrays. Each array configuration has a different data acquisition rate, cost and complexity. To have a valid comparison, we assume that each antenna in the array is hooked up to a similar GPR system. For the maximum sample rate of each GPR system, we use the figure for a GSSI SIR-2 GPR system, which is about 200 scans per second.

3.3.1. Single antenna

In this configuration we have a single antenna, which must be mounted on a scanning mechanism in order to scan a 2-D area. The scanning mechanism could take a form of a mobile platform as shown in Figure 11, which also shows the scanning pattern of the antenna. Usu-

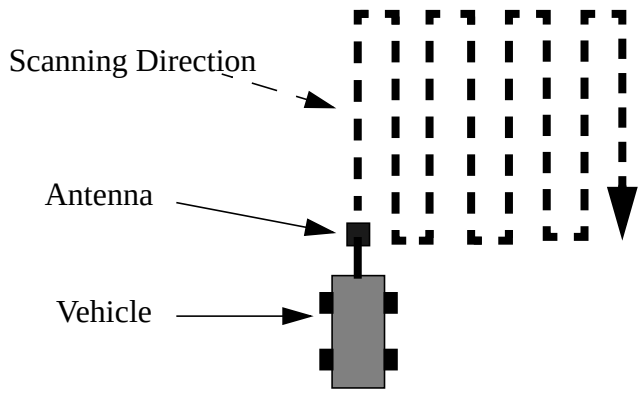

Figure 11: Scanning pattern for a single antenna configuration.

ally the speed of the mobile platform is the limiting factor in this configuration. The time that it takes to scan using this configuration is determined using the following equations:

$$\text{Scanning Time} = \frac{1}{\text{Sample Rate}} \times \frac{\text{Scanning Area}}{\text{Grid Spacing}} \tag{3}$$

$$\text{Sample Rate=Min(GPR Sample Rate,Vehicle Velocity/Grid Spacing)}$$

Notice that the bottleneck depends on the GPR equipment sample rate or vehicle velocity. This configuration is not suitable for scanning a very dense grid because a very accurate vehicle positioning system is needed. The maximum absolute error of such positioning sys-

tem should be in the order of the grid spacing. Anything more than that will decrease the accuracy of the scans significantly.

Instead of mounting the antenna on a vehicle as shown in the Figure 11, we can also mount the antenna at the end of a robotic manipulator or excavator. This is especially suited for scanning a limited area, which could be reach completely by the manipulator. A single axis or a two axis scanning mechanism can be used as shown in Figure 12.

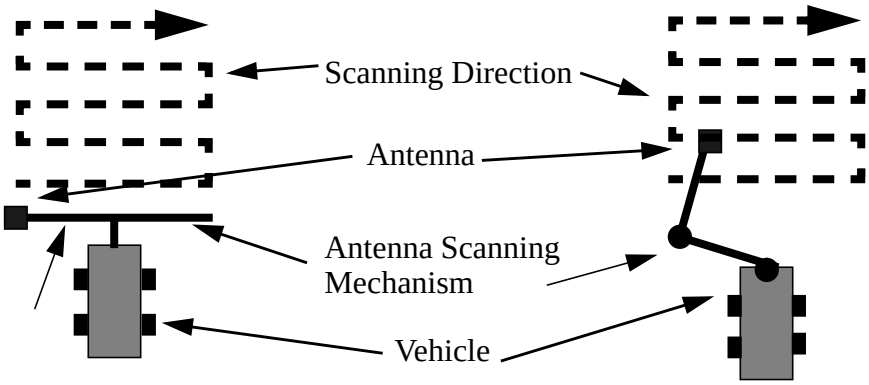

Figure 12: Scanning pattern for a single antenna with a one axis (left) or two axis (right) scanning

A specially designed vehicle with manipulator can be used, but a computer controlled excavator is also suited for the application. It is also possible to retrofit existing excavator as a scanning mechanism by equipping it with a precise positioning sensor.

With respect to the GPR equipment, this configuration has the minimum cost, we only need a single GPR system with a single antenna. It is also the best configuration for uneven surface since the height of the antenna can be adjusted for every single scan.

3.3.2. Linear Antenna Array

In order to minimize data acquisition time we can use a linear array of antenna as shown in Figure 13. By having multiple antennas we can decrease the amount of scanning that we need to do. The linear antenna array can be connected to a single or multiple GPR system using multiplexers. Even when they are connected to a single GPR system, this antenna configuration might still have a higher data acquisition rate than the single antenna system, because it eliminates the physical scanning motion in one axis. This configuration is well suited for scanning a continuous strip of limited width. Mine detection is the obvious application, but it is also suited for pipe mapping. Instead of mapping the pipes by scanning a large area, we can also map the pipes by tracking the pipes once they are detected. Once a pipe is detected, the linear antenna array enables the system to get the pipe reflection profile

instantaneously by orienting the linear array perpendicular to the pipe direction. One disadvantage of this configuration is the limitation on the scanning resolution. Since each individual antenna has a finite width, the scanning resolution along the antenna array is limited by the width of the antenna. However, the scanning resolution in the direction of the vehicle travel is only limited by the accuracy of the vehicle positioning system.

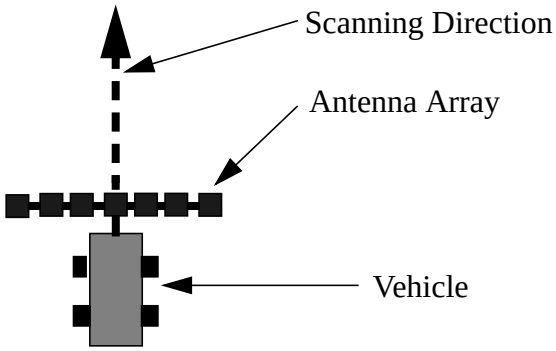

Figure 13: Scanning pattern for a linear antenna array configuration.

As in the single antenna configuration, an accurate positioning system is needed to merge different linear scans accurately. Any error in the positioning estimate will result in some error in the collected 3-D data.

Depending on the number of GPR system used, the cost of GPR equipment for this configuration can be moderate to expensive. We can use a single GPR system connected to all the antennas using a multiplexer. At the other extreme we can also have a GPR system for every antenna. Due to possible interference problem between closely located antenna, the latter option is not desirable even if cost is not an issue. To avoid interference, two closely located antennas should not be used at the same time, which means that even if two closely located antennas are connected to two different GPR systems, only one of them can be used at a single time. Therefore it is more cost effective to connect some of the antennas to the same GPR system using multiplexers. A multiplexing configuration is shown Figure 14, which also shows the firing order of the antennas. The figure shows nine antennas connected to

three GPR systems. At any moment, the active antennas are always separated by two inactive antennas.

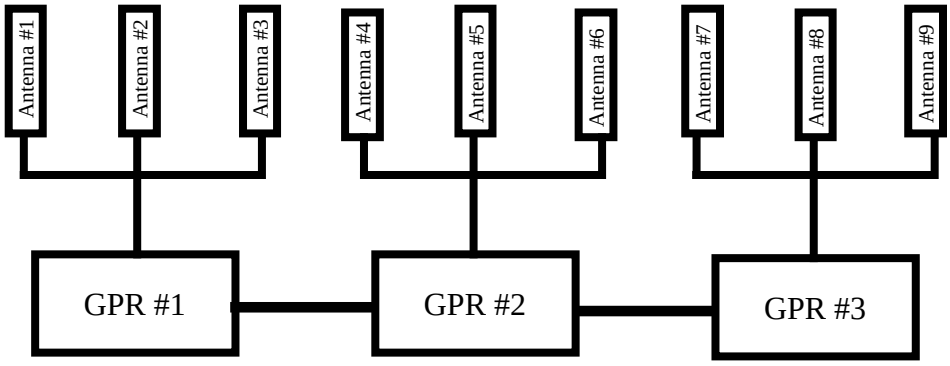

Antenna Firing order: 1 4 7 - 2 5 8 - 3 6 9 - 1 4 7.....

Figure 14: Multiplexed GPR antenna and their firing order.

3.3.3. Area Antenna Array

We can also have a 2-D array of GPR antennas as shown in Figure 15. If the size of the array is larger than the area that needs to be scanned, it is possible to obtain the 3-D volume data without any scanning action. The main disadvantage is the large number of antennas required for obtaining the data. For example, to locate a mine with a diameter of 20cm, we will conservatively need an array of antennas with a dimension at least 40cm by 40cm to cover the mine and its surrounding area. If the antenna spacing is 4cm, this means that we need 100 antennas. As in the case with the linear antenna array, we can have a single GPR system connected to all the antennas, or we can multiplex the antennas to several GPR system to increase the throughput. Like a linear antenna array, the size of the antenna limits the sampling resolution.

There is one advantage of using a 2-D array of antenna even when scanning motion is required. Similar to the linear array of antennas, the vehicle still supplies one axis of the scanning motion. But in this case, we do not need an accurate positioning mechanism as long as there is an overlap between the coverage of two consecutive scans as shown

in Figure 15. We can use the overlapped area as a positioning feedback, which enable us to merge the two sets of scans accurately. This will work if features in the overlapped data can be used as location references. In the absence of such features, we do not have sufficient

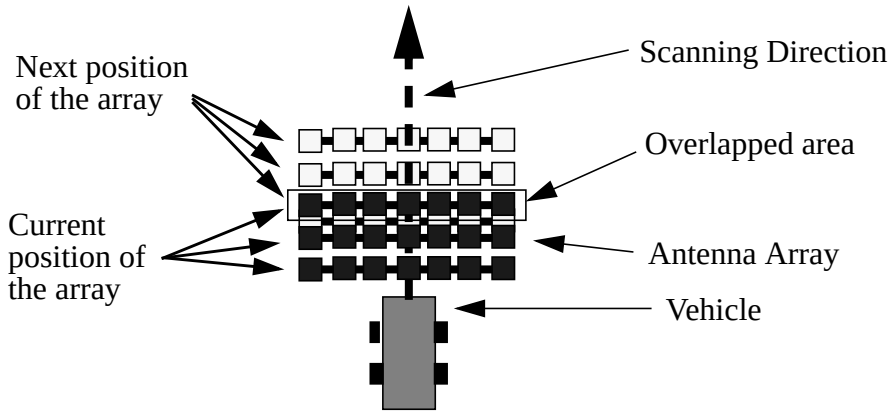

Figure 15: Scanning pattern for a 2-D array of antenna configuration.

information to merge the two data sets. This is not a problem in many cases where the buried objects are smaller than the array, since we do not need to merge consecutive scans to detect a buried object. In addition, a lack of feature often signifies that there is no buried object, and no merging needs to be done.

3.4. GPR Data Visualization

Once we collect a 3-D GPR volume data, we need to be able to visualize it before and after processing. The basic problem is to find a way to represent the 3-D GPR volume data in such a way that makes it easy for human to understand it. Visualization techniques range from a simple 2-D plotting technique to a sophisticated transparency based 3-D volume rendering. Each of these techniques has its own advantages and disadvantages. Typically each technique is well suited for one or a few specific tasks. It is important to choose the correct ones based on the objective of the visualization and the availability of the correct type of data. Some techniques are useful for visualizing both raw and processed data, while some only work with the latter.

We will show that the visualization problem becomes much easier to solve once we process the data. Our processing techniques transform the 3-D data into the parameters of the buried objects. We can use these parameters to display the 3-D representations of the objects, which are much easier to understand than the raw 3-D data.

3.5. 2-D Visualization

3.5.1. Vertical Cross Section

2-D visualization is the simplest way to visualize the 3-D GPR data. It works by first encoding the amplitude of a single GPR scan using intensity or color as shown in Figure 16.

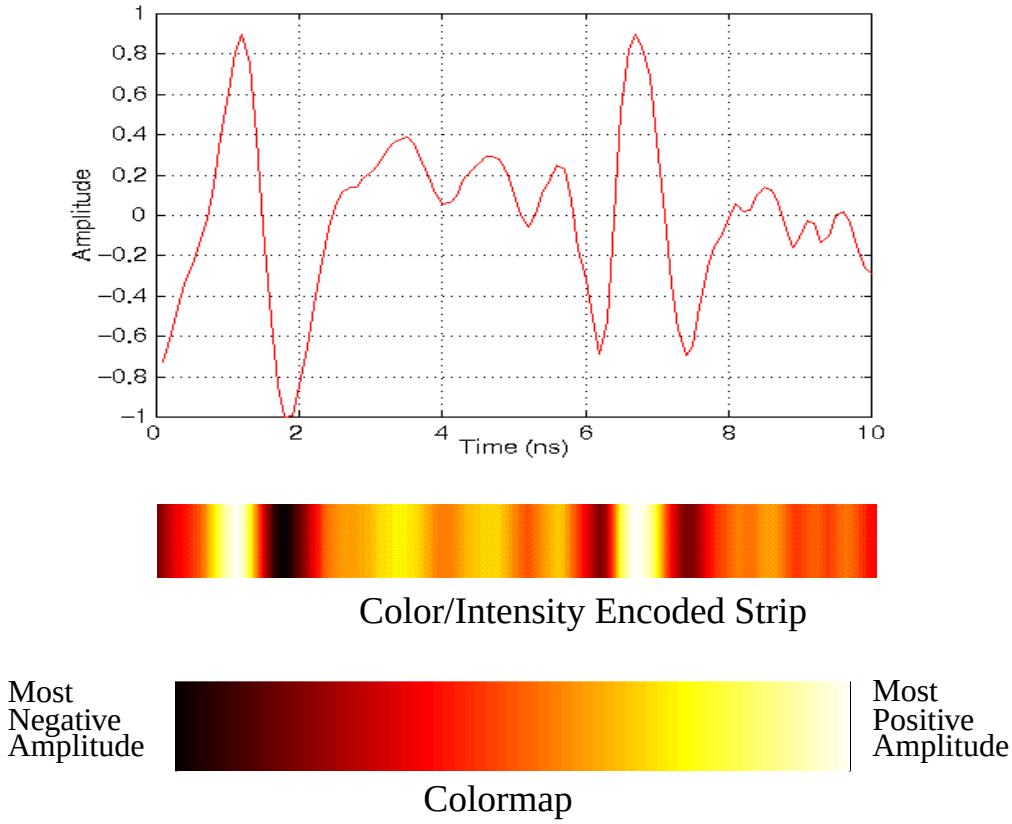

Figure 16: A single GPR scan shown as a graph and as an encoded strip of varying intensity. The intensity corresponds to the amplitude of the signal as shown in the colormap.

Once we encode every GPR scan as a color coded strip, we can put each strip next to each other to construct a 2-D image of a slice or cross section of the GPR volume data. A collection of slices completely represent the 3-D volume data. Figure 17 shows the 3-D volume data and how it can be decomposed into a series of 2-D slices.

Figure 18 shows one of the slices in detail. It can be clearly seen from that figure that there is a buried object in the middle of the scanline, although it is hard to tell what kind of object it is.

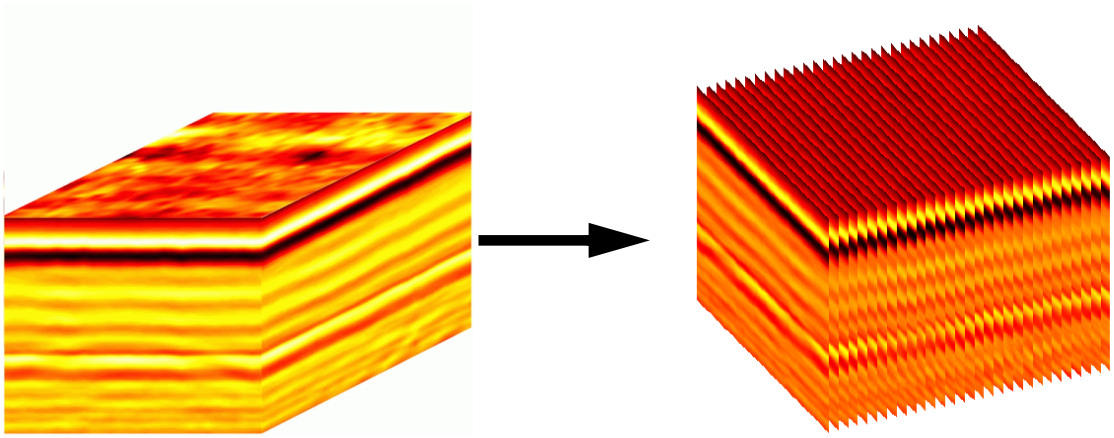

Figure 17. The 3-D GPR volume data decomposed into a series of 2-D slices.

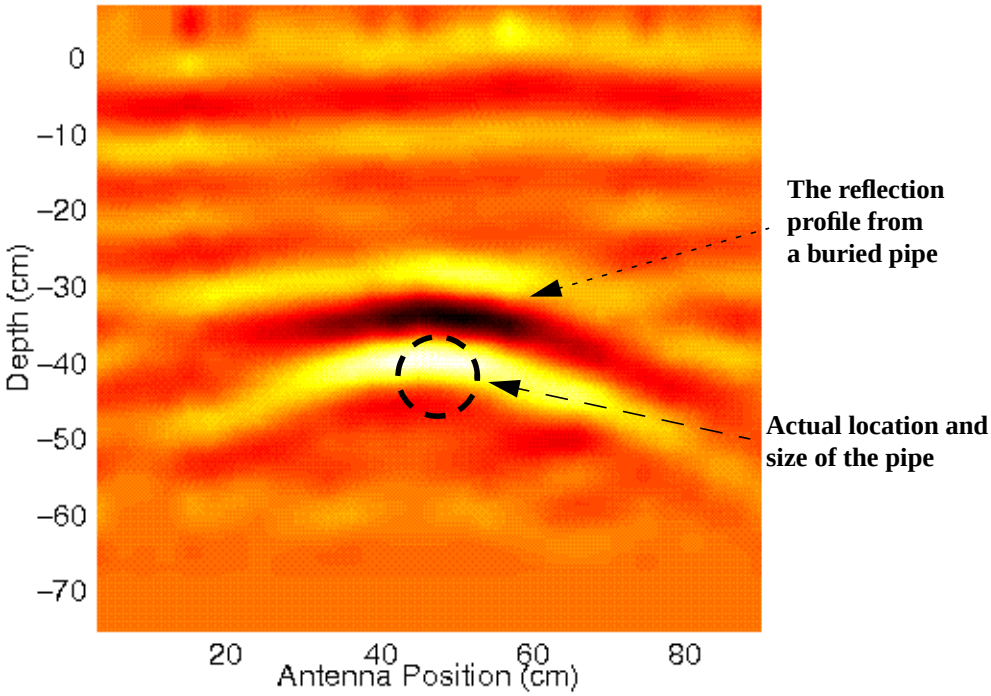

Figure 18: A vertical slice of GPR volume data constructed by putting each color-coded GPR scans next to each other. The reflection profile is produced by a buried pipe.

There are infinitely many directions to slice the 3-D volume data. Figure 19 shows slices obtained in the Y axis direction and Figure 20 shows slices obtained in the X axis direction. The most obvious is to construct the slices along the scanline, which is the direction in which the antenna moves the fastest. But there is really no reason not to slice the volume

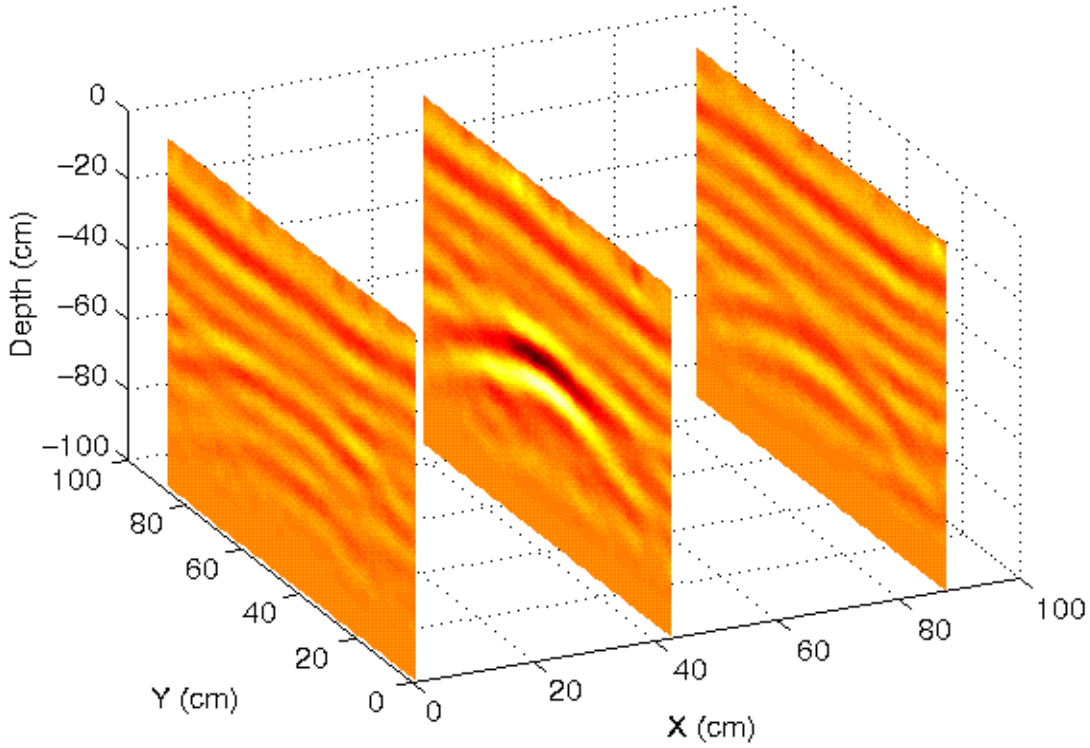

Figure 19: Several vertical slices of a 3-D GPR data in the X axis direction.

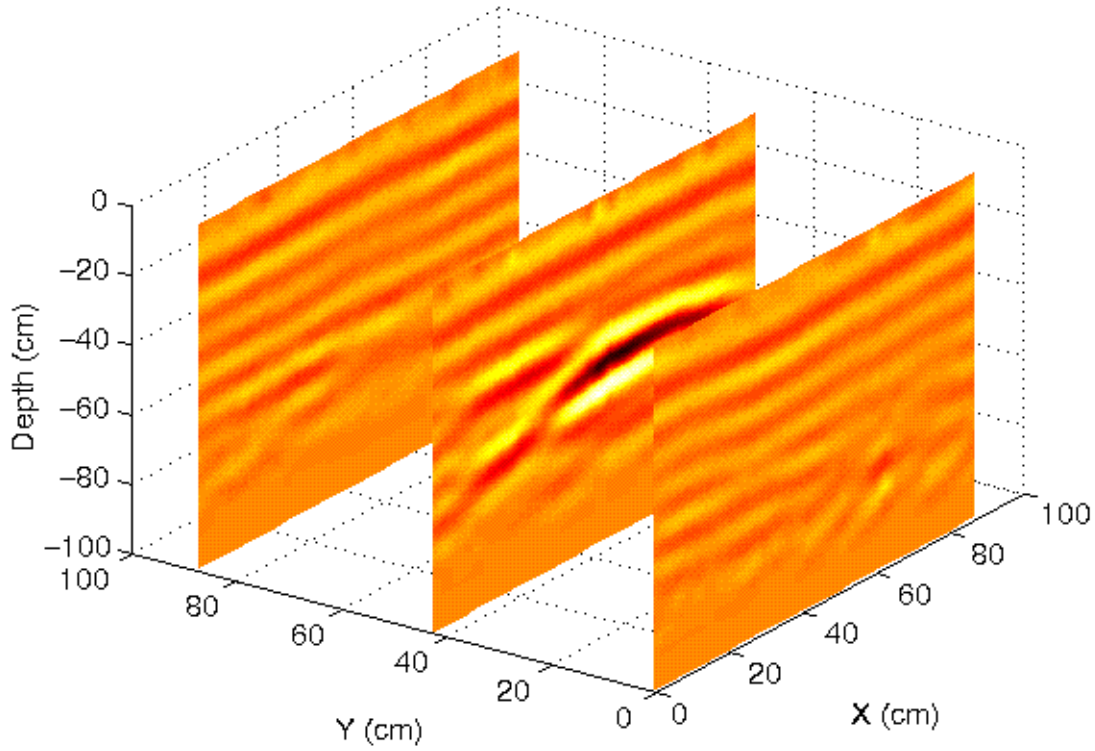

Figure 20: Several vertical slices of a 3-D GPR data in the Y axis direction.

data in different directions. The best 2-D slices of a 3-D data are the slices that can best show the buried objects that we are interested in. The problem is that we usually do not know ahead of time what is buried under the soil and what is its orientation. In order to illustrate the problems, let's display the slices in Figure 19 and Figure 20 as mosaics of vertical cross sections. These are displayed in Figure 21 and Figure 22. Looking at Figure 21, it

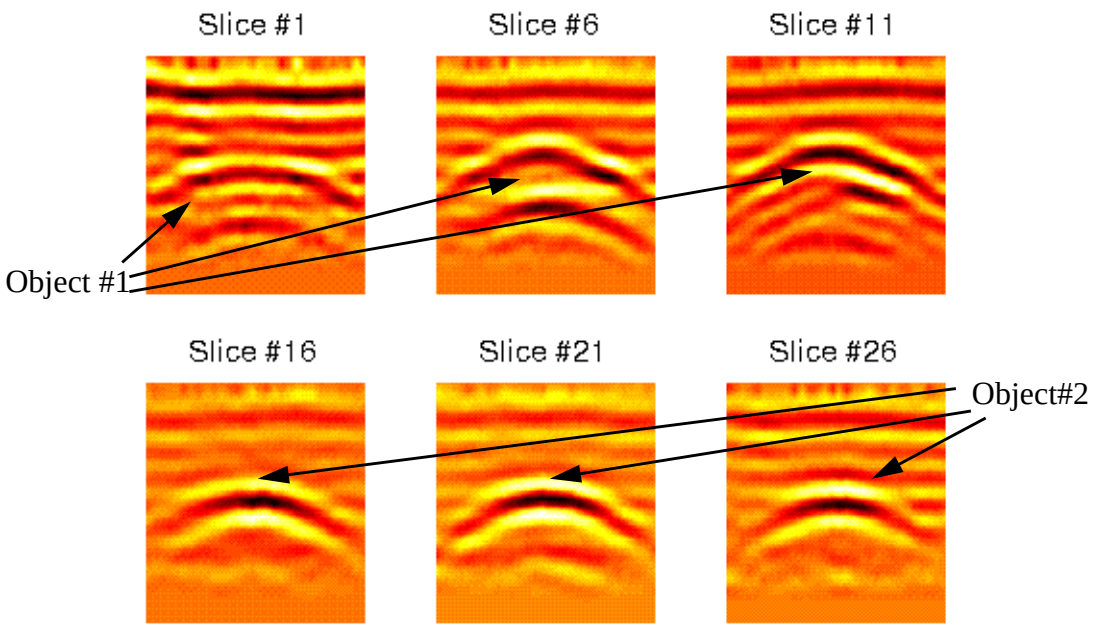

Figure 21: A series of slices in the X direction (same as Figure 19).

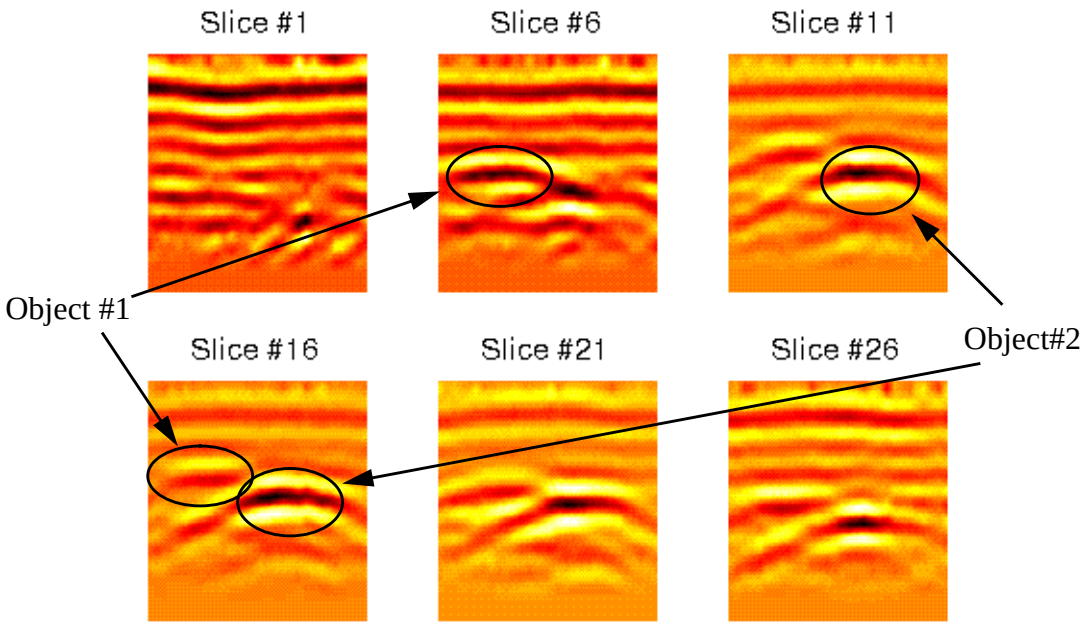

Figure 22: A series of slices in the Y direction (same as Figure 20).

seems that there is only one buried object, but there are actually two buried objects. Figure 22 on the other hand, shows both object quite clearly although one of them have a much stronger reflection than the other.

The main benefit of visualizing the vertical cross sections or slices is that we convert a 3-D visualization problem into a 2-D visualization problem. This can be done with very little computing power. The trade-off is the speed of inspection and the ability to locate 3-D structures that are embedded in the 3-D volume data. To inspect the 3-D data using 2-D slices means that we have to examine every slice of the data one by one. This is both time consuming and error prone. When an object's main feature is oriented in a certain direction, the operator might not be able to see the object when he is taking a look at the slice from other directions Since we are sensing 3-D objects, this can happen quite easily. Moreover the operator might misinterpret the image as shown in Figure 21 and Figure 22. This problem is worse for asymmetrical objects. For example, a pipe looks completely different depending on the direction that we see it. If it is seen from the side, it looks like a flat object, but if it is seen head on, then it looks like a circular object. The same exact problem also exists when we are visualizing the 2-D slices of a 3-D GPR data of other asymmetrical objects.

The problem with 2-D visualization of 3-D GPR data boils down the reduced amount of useful information that the operator can observe. In 2-D, many of the 3-D object's features degenerate into non-differentiable 2-D features. Despite of this shortcoming, 2-D visualization still plays an important roles in subsurface data visualization, because it only requires very minor processing of the data. In 3-D visualization on the other hand, we need to threshold the data or use transparency to be able to look at the raw data.

3.5.2. Horizontal Cross Section

Instead of slicing the 3-D data in the vertical direction, we could slice the data in the horizontal direction to generate horizontal cross sections. This technique is very useful if we are looking at small buried objects that are oriented horizontally. As long as the size of the object is small compared to the scan area, such an object can be located easily by looking at the horizontal slice. An example of a useful application for horizontal slice visualization is for mapping largely horizontal sets of pipe. Pipe directions and interconnection should be visible from visualizing the horizontal slice. Figure 23 shows a series of horizontal cross sections of a single horizontal pipes

On the other hand if the buried object does not have any horizontal feature such as a large horizontal plate, we will not be able to find it by visualizing the horizontal slices.

Imaging of antipersonnel landmine is another good application for horizontal cross section visualization. An antipersonnel mine is relatively small (about 10cm in diameter), so the disturbance can be readily seen in the image. We still need to differentiate the reflection from

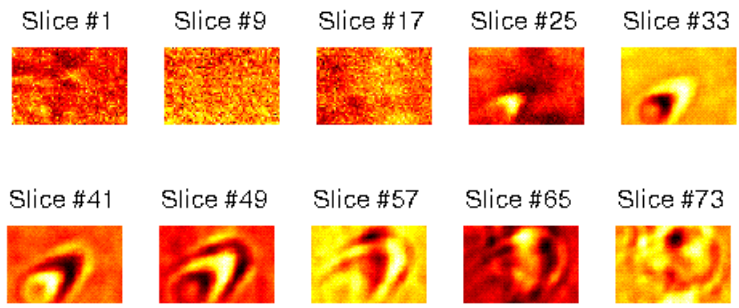

Figure 23: Horizontal cross sections of a pipe.

landmines with disturbances caused by other objects such as rocks. Since modern landmines are mostly made from plastic, the reflections from the rocks might even be stronger that the reflections from the landmines.

3.5.3. Projection of subsurface objects into the soil surface

Even in the best cases, sometimes it is hard to determine the horizontal location of an object from a 3-D image. In some cases this is a very critical and necessary information. For example, in landmine detection, the operator of the detection equipment needs to find out exactly where the mine is. Since landmines are usually buried shallowly, their horizontal locations are much more important than their depth. In such cases, it is possible to generate a projection of the buried object on the ground surface as shown in Figure 24. The projection itself can be color coded. We can color code objects near the surface red and objects that is buried deeper in the soil blue.

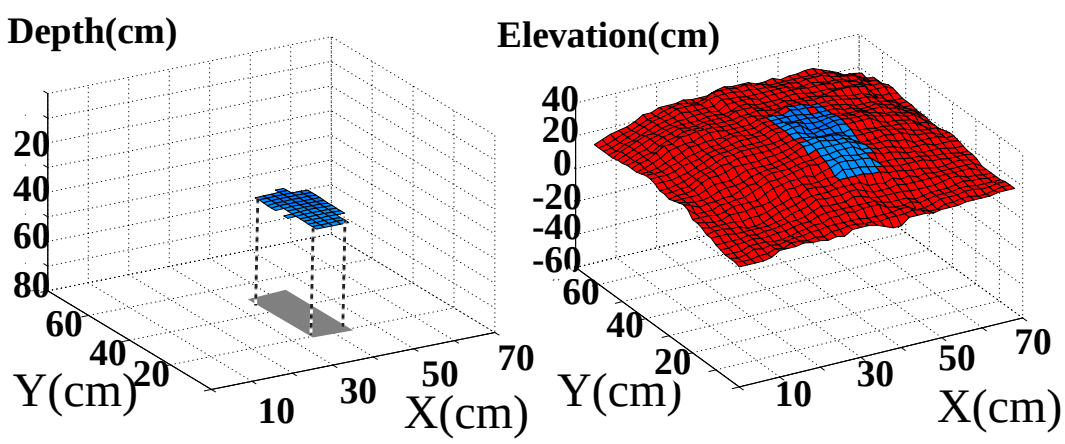

Figure 24:The detected top surface of the cylinder and its projection to the elevation map

3.6. 3-D Visualization

To preserve most of the information embedded in GPR data, we need to able to visualize the data in 3-D. This is harder than visualizing 2-D cross sections, but now we do not have to choose the best slicing direction for each set of 3-D volume data. 3-D visualization techniques can be divided into two main categories. One is visualizing the raw 3-D volume data and the other is visualizing the processed 3-D volume data.

3.6.1. Raw Data Visualization (Volume Data Rendering)

3.6.1.1. Transparency

By mapping the amplitude at each voxel to a different transparency level we can visualize a 3-D volume data. This works well if there are only a few voxels that have a high amplitude or low level of transparency. Voxels with low transparency value occlude all the voxels behind them. Therefore if there are too many of them, it is hard to visualize the 3-D data due to occlusion. The mapping from amplitude to transparency can be varied dynamically to emphasize voxels with certain amplitude. Figure 25 shows two views of a 3-D visualization of GPR data using transparency. The images are generated using the package BOB from University of Minnesota.

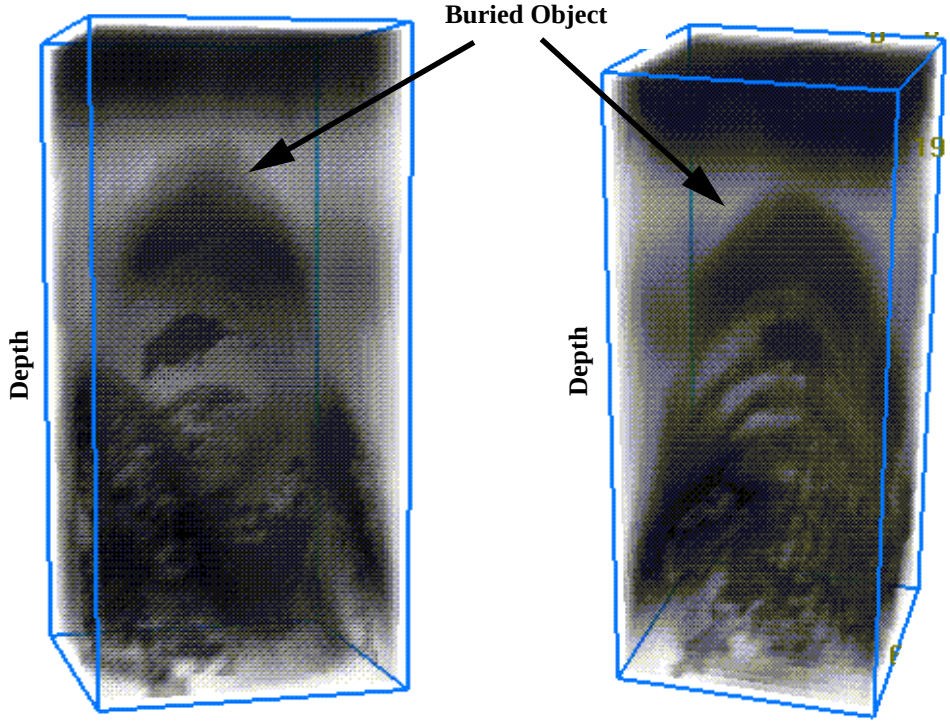

Figure 25: Two views of a buried object using transparency based 3-D volume data visualization

3.6.1.2. Thresholding

Another method of visualizing raw 3-D volume data is to threshold the voxels value. Only voxels that have a value bigger than certain threshold are displayed. They are displayed as a solid cube while the voxels that have values lower than the threshold are not displayed at all. Different threshold values should be tried for each 3-D data set to obtain the best result. Sometimes it is impossible to obtain a threshold value that can show all the buried objects. This will happen when we have one object which returns a very strong reflection and another object that returns a very weak reflection. If we set the threshold very low in this case, the strong reflection from the first object will obscure the second object. If we set the threshold too high, we can not see the second object because its reflection energy is below the threshold.

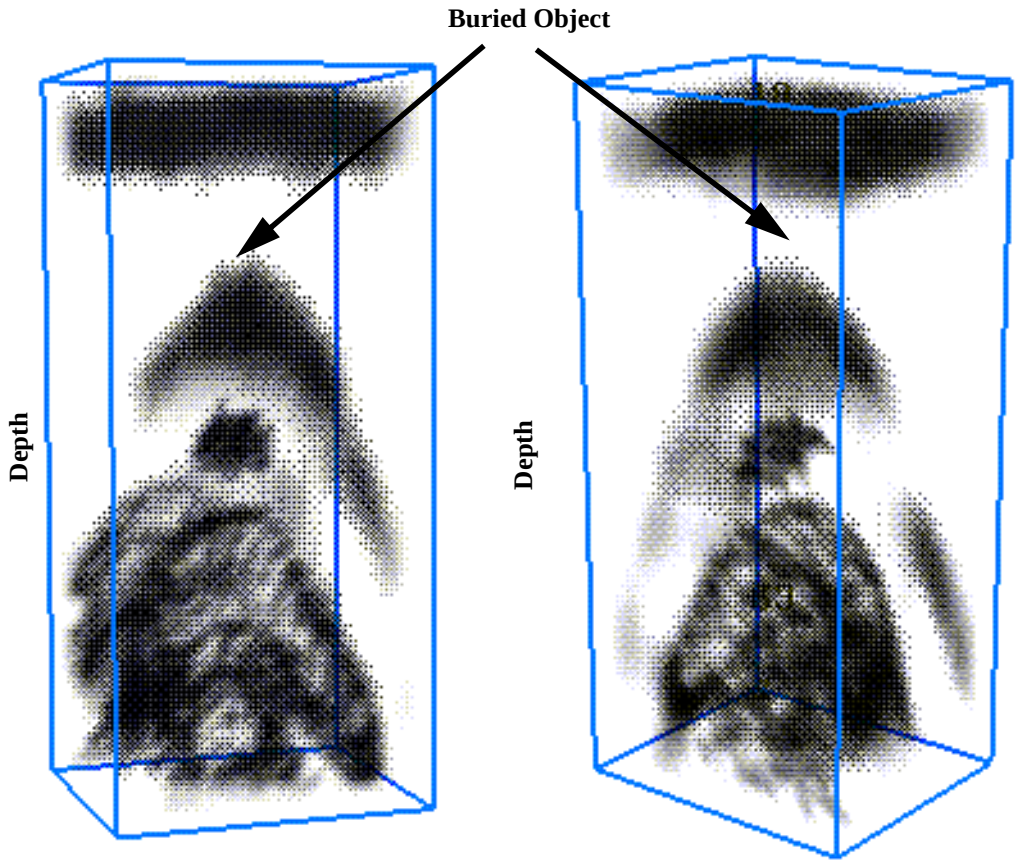

Figure 26: Two views of a buried object using thresholded 3-D volume data visualization

3.6.2. Processed data visualization

3.6.2.1. Object Surface

It is much easier to visualize 3-D GPR data once the objects embedded in the volume data have been located. Then it is possible to reduce the 3-D data into a series of objects and display the object surfaces. An example of an object surface is shown in Figure 27.

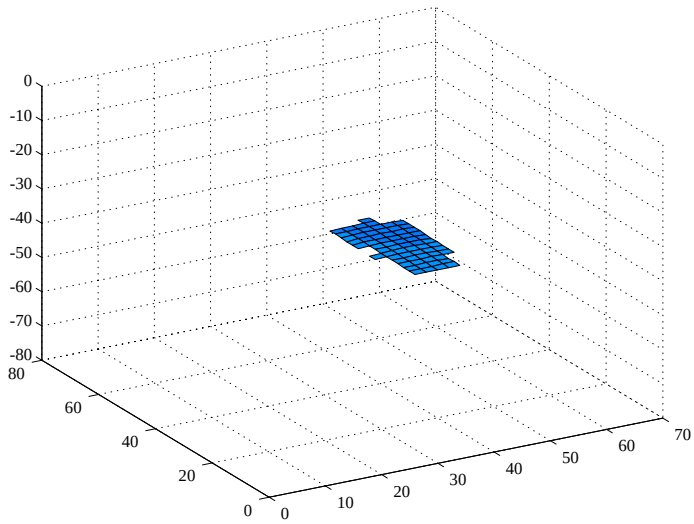

Figure 27: The top surface of a buried object.

3.6.2.2. Fitted Object Model

Once the object is located, it is also possible to fit its shape to a suitable model. We can then display this object's model instead of the actual reflection profile of the object as shown in Figure 28. This is the easiest display to understand because all the important parameters are already implicitly represented on the display. Such parameters can be radius and length of a pipe, along with its orientation and location. This is one of the most important benefit of having an algorithm that can automatically extract buried objects from 3-D GPR volume data. Instead of having the human operator struggling to reconstruct the 3-D data in his or her mind, we can just represent the object model to him or her in a very easy to understand 3-D graphics. Figure 25 and Figure 28 illustrate the difference in looking at the raw 3-D data and the processed data. It is also possible to show this 3-D graphics with images from a live video using a head mounted display so the operator has the impression of looking through the soil. This would be especially beneficial for tasks such as landmines removal.

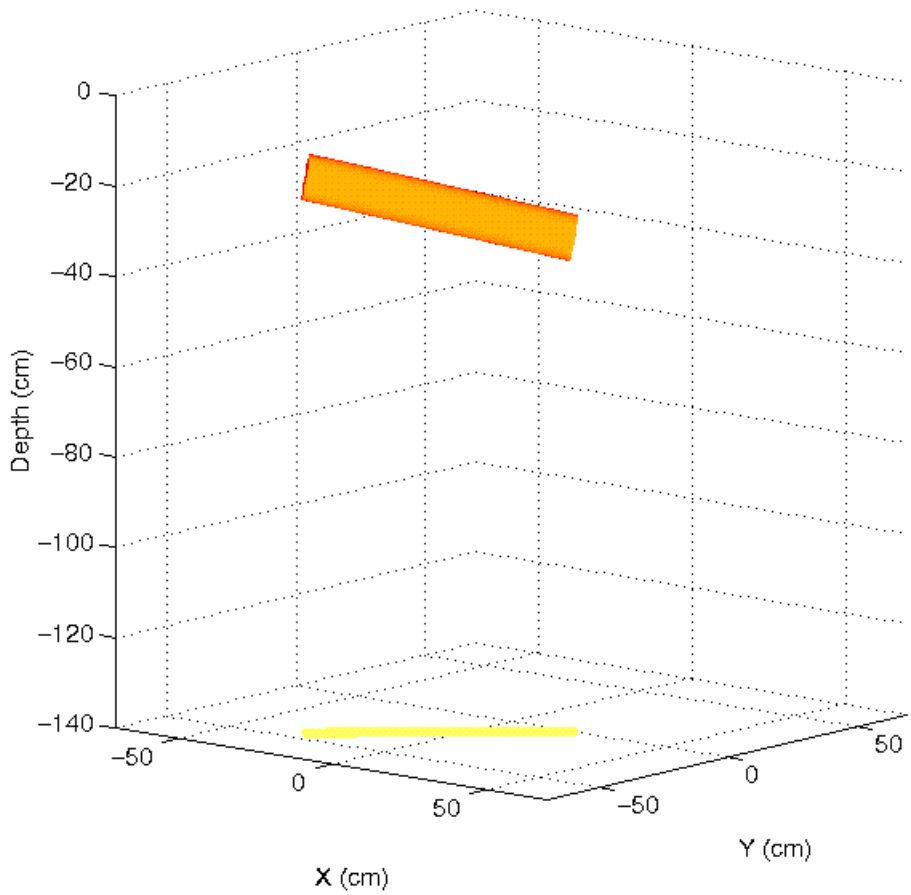

Figure 28: Buried object model (a cylinder).

3.7. GPR Data Processing and Interpretation

As we have seen in the previous sections, it is much easier to visualize the processed rather than the raw 3-D data. The 3-D raw data also leaves a lot of room for subjective interpretations, which can varied significantly from one expert to another. In this section we will explain in detail the deficiencies of manual GPR data interpretation and we will also examine the issues in automatic interpretation of GPR data.

3.7.1. The problems with manual GPR data interpretation

GPR has been around for a while, but most of the collected data are interpreted manually by human experts. As the amount of data grows and new uses of the data appear, the manual

interpretation process can become the bottleneck in the subsurface sensing problem. Below are some of the problems encountered when the GPR output is interpreted manually:

- Vast amount of data. Using GPR to scan a wide area results in a vast amount of data, which could easily reach gigabytes of data. This is especially true if 3-D high resolution data are needed. In order to inspect the data manually, the operator needs to inspect each 2-D slice or cross-section of the data. The spacing of the slice needs to be in the order of the smallest buried object that needs to be detected. To detect a 50cmx50cm object buried at maximum depth of 5m, we need to visualize 40000 slices, each slice represents a vertical slices of 50mx5m. If an expert needs 30 seconds to examine each slice and record the object location in each slice, then the whole task would take over 300 man-hours.
- Understanding the GPR data by visualizing the raw 3-D images is not trivial. It is extremely hard for human being to visualize 3-D volume data without any preprocessing of the data. As a result usually 3-D GPR data are interpreted on a slice by slice basis.
- Visualization of 3-D data using a series of 2-D slices does not consider the out of plane scattering effect. Out of plane scattering happens when the source of a visible feature in a 2-D slice actually comes from neighboring slices. In another word, when we are visualizing a 2-D slice, some of the objects that are visible in that slice are physically located somewhere else. This is possible because the antenna has a wide beamwidth and is able to detect objects that are not located directly under it.
- Constructing the shape of a buried object by visualizing a series of 2-D slices is very difficult.
- Data management and record-keeping are also time consuming and prone to error if done manually.

Due to the above problems, there is a strong need for automated data collection and interpretation of 3-D GPR data. It will not only increase the throughput of a subsurface mapping operation using GPR, but it will reduce error as well.

3.7.2. Automated Interpretation of GPR data

Although automated interpretation of GPR data alleviates some of the problems associated with manual interpretation of GPR data, there are still many issues to be solved for its practical use. Here are some of the issues that need to be solved in order to have automated interpretation of GPR data that is both useful and practical:

- Detection of the reflected signals from small objects as well as big objects. This

means that the detection process can not be solely based upon the strength of the reflected signal. So simple thresholding will not work.

- Transformation of the 3-D volume data into parameterized buried object. This involves segmentation of the volume data into surfaces and fitting object models to the surfaces. This is especially difficult because GPR data contain a lot of spurious echoes and reverberations from the buried objects and other disturbances in the soil. So it is necessary to differentiate the reflections from the object from their reverberations and other spurious echos.
- Due to the noise and the heterogeneity of the soil, it is hard to fit an exact model to the reflection profile of the reflecting signal from the buried object. The model fitting process must be able to tolerate this.
- As in the manual interpretation process, the vast amount of data also raises an issue for the automated interpretation of GPR data. In this case, the algorithm must be really efficient in order to minimize processing time.

3.8. Example of GPR Data

In this section we will show some cross sections of raw 3-D GPR data. We will use the examples to illustrate some of the effects that occur in GPR imaging. We have to consider these effects in our GPR processing algorithms so we can extract the accurate buried object parameters from the 3-D data.

3.8.1. Blurring Effect

Wide beamwidth of a GPR antenna often causes objects in GPR data to appear blurry and larger than their actual size. The effect is similar in concept to taking a picture with out of focus lens. To illustrate this effect, we scan a buried pipe using GPR and show a vertical cross section of the data in Figure 29. The dashed circle denotes the location and size of the actual pipe, and the dashed curve denotes the reflection profile that is created when we scan the pipe using GPR. The figure shows that the reflection profile is much wider than the pipe. It also shows the cause of this blurring effect. Since the antenna has such a wide beamwidth, the barrel is detected even when the antenna is not directly above it. As a result the resulting reflection profile is much wider than the object's actual size. The width of the reflection profile is related to the depth of the object and beamwidth of the antenna. Let us assume that the barrel is buried 1 meter deep and the beamwidth is 90 degrees. We can then compute that the reflection profile would be at least 2 meters wide ($2x(\tan(45^o)x1meters)$). This is just a theoretical estimate, the actual width of the reflection profile depends on many more parameters such as the signal attenuation rate, object's size and shape.

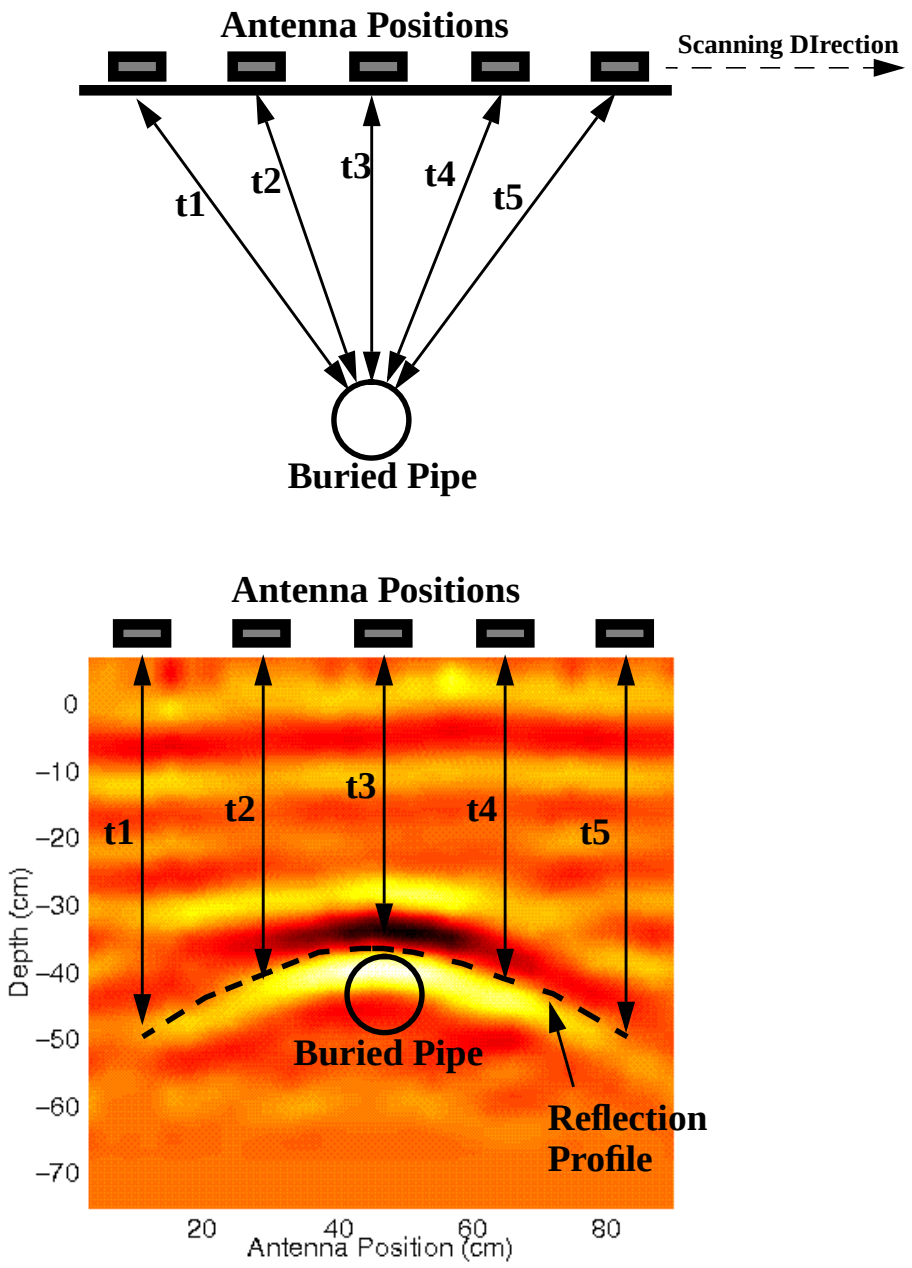

Figure 29. A cross section of 3-D GPR data resulting from scanning a buried pipe.

In chapter 4, we will discuss algorithms that can remove this blurring effect. The algorithms are based on migration, which is a processing method that sharpen GPR data and correct other GPR imaging effects.

3.8.2. Shifting and Tilting effect

This phenomenon is also caused by the wide beamwidth of the antenna which caused reflection profile of a flat object to be shifted to a different location and tilted at a different angle than its actual angle. An example of this effect can be seen in Figure 30. It shows a cross section of a GPR scan of a plate. The dashed line denotes the actual location and orientation of the plate which are different with the location and orientation of the reflection profile.

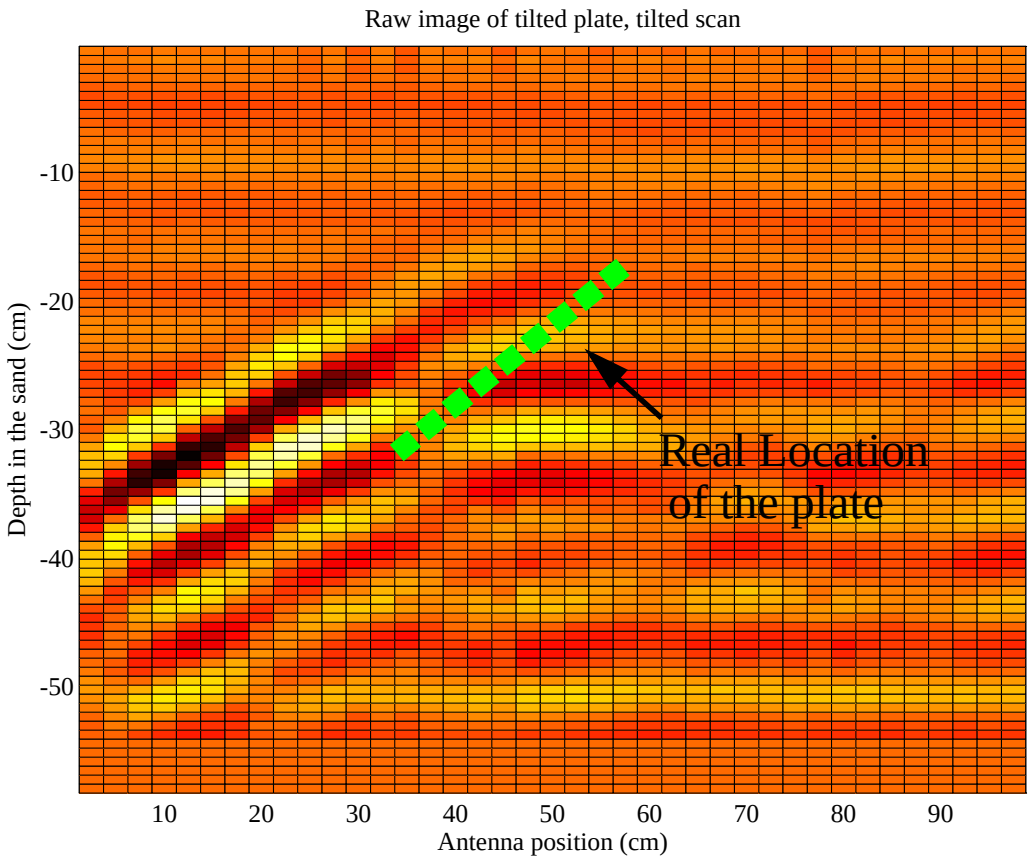

Figure 30: A cross section of a buried plate showing the shifting and tilting effect

This effect occurs because the reflections do not come directly from under the antenna. as can be seen in Figure 29. As a result the position of the object is shifted and the angle of the reflection profile will be less than the actual angle of the object.

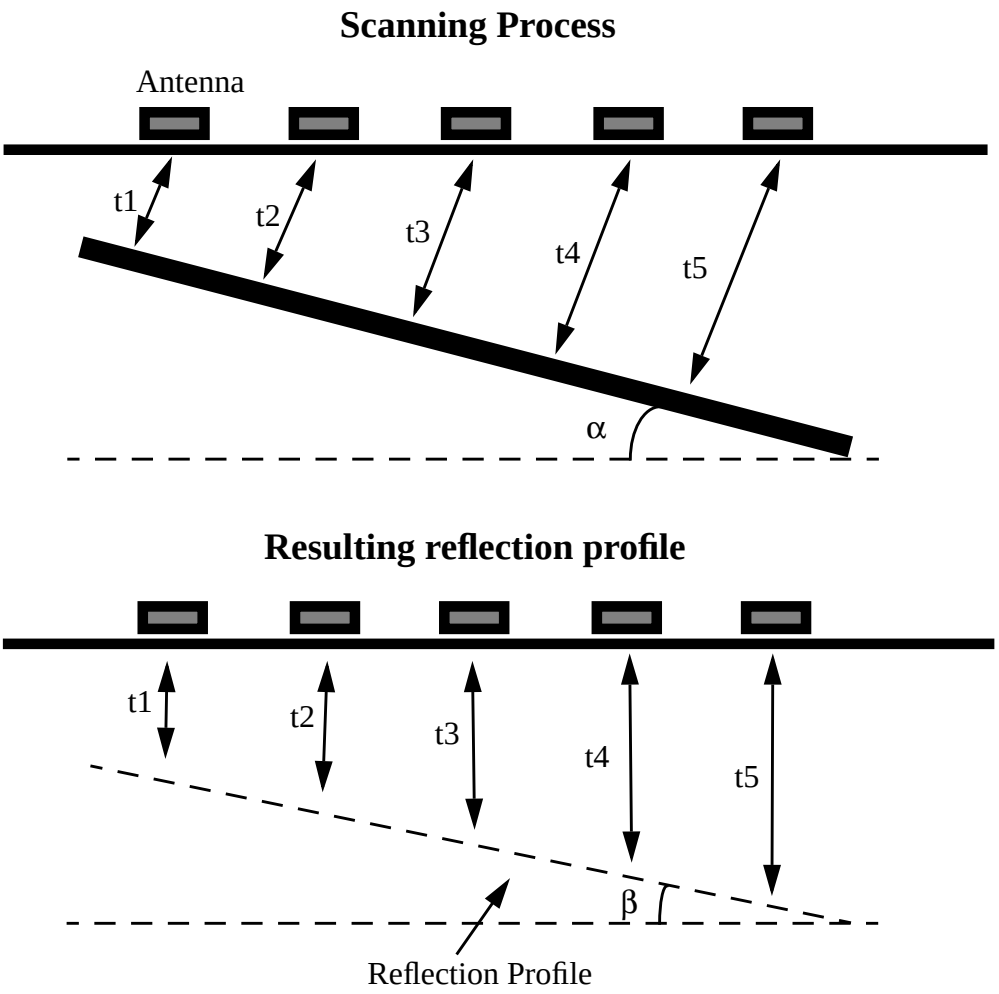

$$\alpha = \operatorname{asin}(\tan(\beta))$$

α = the true angle of the surface

β = apparent angle of the surface in the GPR image

Figure 31: A line scan above a tilted plate and its resulting reflection profile.

Chapter 4. Volume Based Processing

4.1. Overview of Volume Based GPR Data Processing

In this chapter, we will explain our development and implementation of volume based GPR data processing for mapping buried objects. We have two different algorithms which directly find the buried object in the 3-D GPR volume data The first algorithm is a 3-D migration based on coherent summation. Although the concept is not new, we extend the migration process to include detection and localization of buried objects. We also implement this migration method utilizing parallel processing techniques, resulting in a much faster execution time. The second method is also based on migration, but it uses a new method which migrates the voxels in the 3-D volume data much more efficiently. We call it migration using reflector pose estimation. Both of these methods operate directly on 3-D GPR volume data. The main assumption for both methods is the uniform propagation velocity in the soil above the buried objects. If this assumption is not true, then both methods will produce inaccurate, but still meaningful subsurface maps.

Before we discuss the algorithms, we will explain a preprocessing step which is used to remove the background noise. This preprocessing step is applied to all GPR data before further processing.

4.2. Background Noise Removal

Although most of the energy is transmitted downward by a GPR antenna, the antenna also radiates some energy upward. Some of this energy is reflected by the robot end effector which holds the antenna, and causes spurious reflections. Other sources of background noise include coupling between the transmitter and receiver and reflections from the antenna mounting frame. Since these spurious reflections and static noise can be much stronger than a reflection from a small buried object, we need to remove them from the GPR data.

To remove the background noise, we aim the GPR antenna to an empty space and take a scan. This scan is then subtracted from all the GPR scans. This background noise reduction

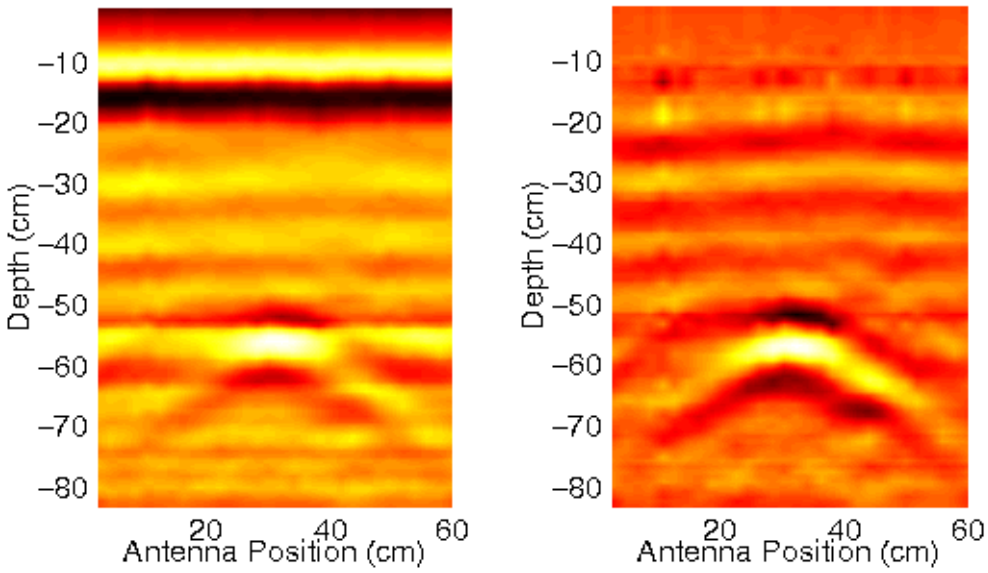

Figure 32: A vertical cross section of 3-D GPR data before (left) and after (right) the background noise is subtracted.

only affects static noise which does not depend on the position of the antenna. Figure 32 shows a vertical slice of GPR scan before and after the background noise is subtracted. The coupling between the transmitter and the receiver is almost completely eliminated and the reflection profile from the buried object becomes much more visible after the removal of the background noise. We always apply this filtering technique before any further processing is done.

4.3. 3-D Migration Using Coherent Summation

The purpose of migration is to focus the energy of the reflections in order to obtain a sharper image. This is necessary because the beamwidth of the antenna is quite wide. For the 1GHz antenna that we used in our experiment, the 3dB beamwidth is about 80 degrees. By focusing the widely spread energy we correct some of the artifacts created by the wide beamwidth of the antenna. As a result, closely spaced objects can be clearly seen and the objects are moved to their true location. Migration also corrects the objects' orientation. The main advantage of the migration process is that it does not make any assumption of the object's shape and size. The main disadvantage is that it requires the propagation velocity of the GPR signal in the soil.

4.3.1. Coherent Summation Migration

The idea of migration using coherent summation is to combine the reflection energy from neighboring scans to obtain a focusing effect. This process is also called digital "beamforming", because we focus the energy of the antenna by postprocessing the digital representation of the GPR output. Another name for this method is synthetic antenna array processing. All these names refer to a method that is used to combine radar scans gathered at different antenna positions in a coherent way.

The principle of this processing method is pretty simple. It is illustrated by Figure 33. We obtain 3 scans from antenna position A1, A2 and A3. In each of these scans, we obtain a reflection from the buried object. From the time delay of the reflections, we can compute

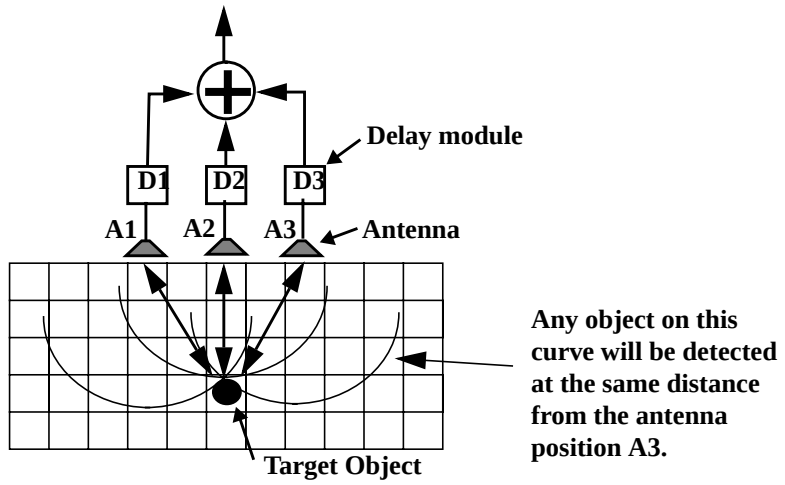

Figure 33: Focusing of GPR beam by coherent summation.

the distance of the object from each antenna position. This is the only constraint that we can obtain from the individual scan. The object can be located anywhere within the antenna's beamwidth or field of view as long as it satisfies the distance constraint. The idea of coherent summation algorithm is that by taking into account the distances of the object to multiple antenna locations, we can pinpoint the location of the object within the antenna's field of view. Mathematically, we can view the operation as taking the intersection of the three curves shown in Figure 33. Each curve is a possible object location that satisfies the distance constraint, so their intersection is the most probable place for the location of the object.

In coherent summation algorithm, the output signal at each antenna position is coherently summed with the signals obtained from the neighboring location. At the true locations of the reflectors the coherent summation should produce constructive interferences while at other locations the summation should produce destructive interferences. To achieve coherent summation, appropriate delays are inserted before the signals from different antenna positions are summed together, as shown in Figure 33. Another way to view this migration process is illustrated in Figure 34. This figure illustrates combining the out signal from five different position to pinpoint the object location. Each individual scan is distributed to every possible reflectors' location. At the true locations of the reflectors, constructive interferences will occur and at other locations, destructive interferences will occur. So migration is very similar to hough transform in computer vision. Basically each scan votes on every possible reflectors' positions. The locations that have many consistent votes are classified as the true locations of the reflectors. Figure 34 only illustrates the migration process in 2-D. In 3-D, the return signal energy is distributed to an equidistant surfaces not curves.

In our implementation, we actually compute the migrated value of a voxel by summing all the raw value of the voxels along its reflection profile as illustrated in Figure 35. If there is a buried object at the voxel location, we will have constructive interference. The reflection profile is different at different depth. In order to do this as efficiently as possible, we construct a migration table. This is a lookup table where we can find the reflection profile for any depth. Each depth position has an entry that lists all locations in relative coordinate which need to be summed to determine the value of a voxel at that depth. This is shown in Figure 36. Using this table, we only need to compute the reflection profile once for every depth value.

4.3.2. Postprocessing of the migrated data

Once the data are migrated, it is necessary to threshold the migrated data in order to determine if a voxel is occupied or not. There are two methods to threshold the migrated data. One is based on the intensity of the voxel and the other one is based on the consistency of the votes during the migration process. The intensity based thresholding is simpler, but picking the intensity threshold is difficult. Another serious problem with intensity based thresh-

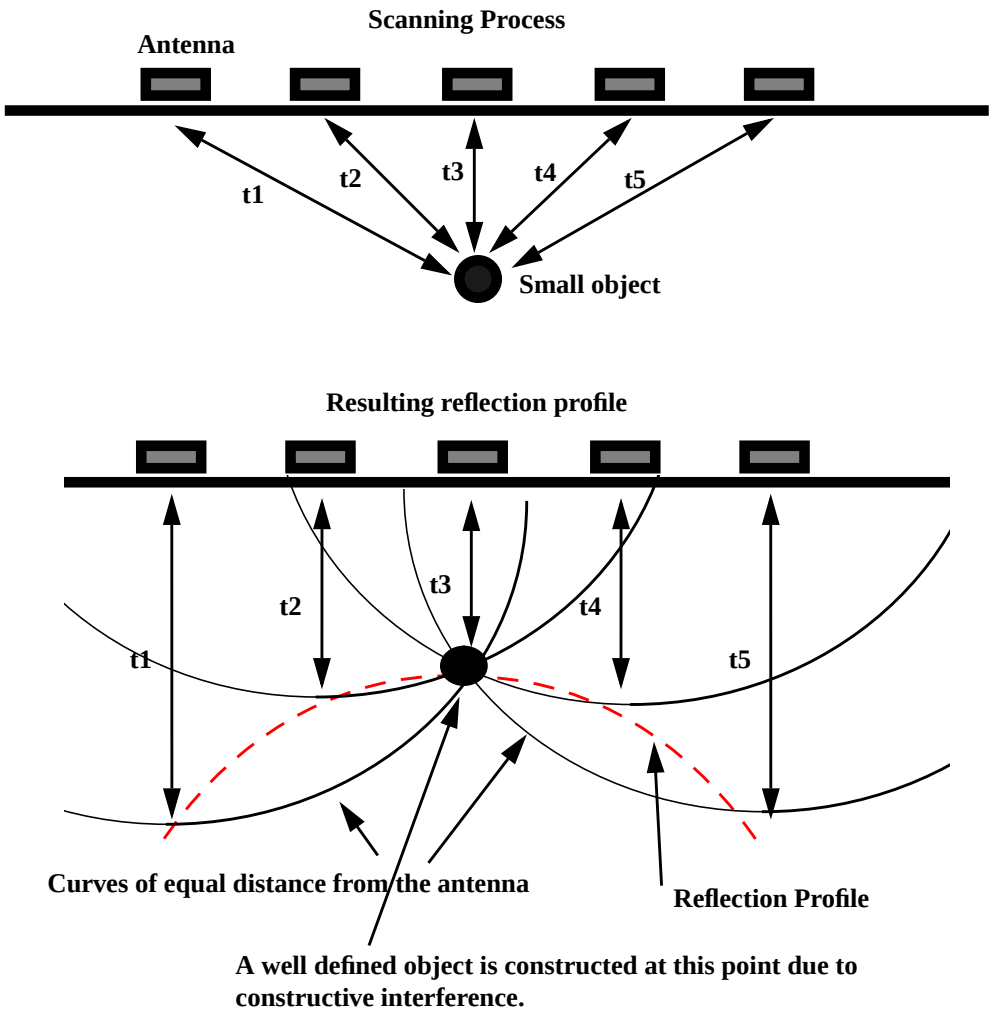

Figure 34: A line scan above a small object and its resulting reflection profile.

olding is the possibility that a large object will mask out reflections from smaller objects because its reflection intensity is much higher than the others.

The second method of thresholding, which is based on the consistency of the votes during migration process is more suitable for the object detection and localization. As explained above during the migration process is done using coherent summation. Instead of just keeping track of the sum at each voxel, we also keep track of the number of positive or negative votes.

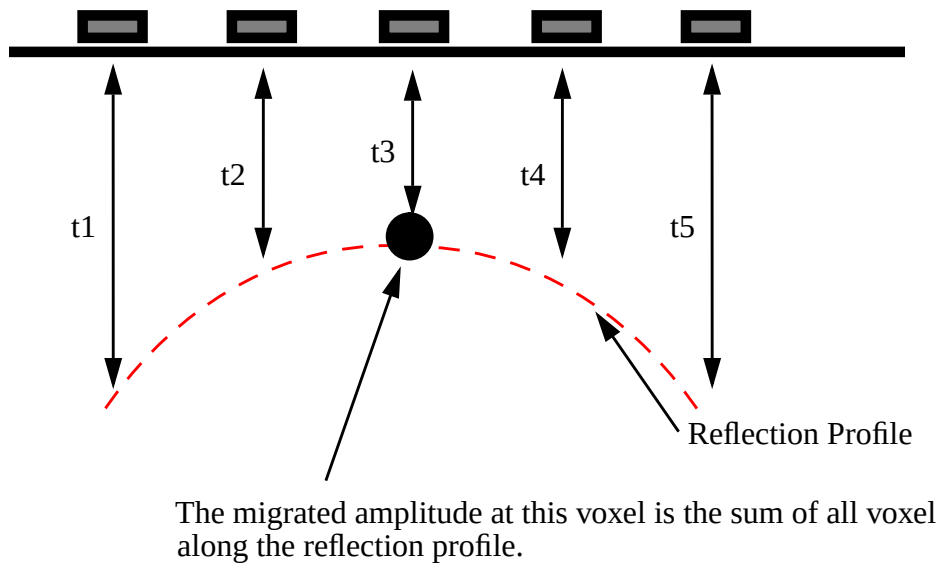

Figure 35: Computing the migrated value of a voxel by summing the reflection energy along its reflection profile.

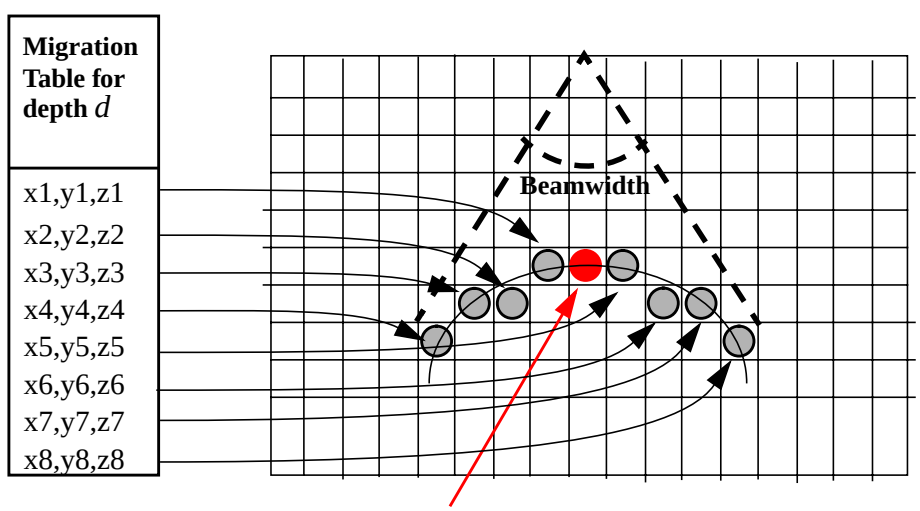

Figure 36: Use of migration table to compute a value of a voxel.

At the end of the migration process, we calculate the consistency of that voxel using the following equations if the voxel value is positive:

$$c = \frac{p}{p+n} \qquad a > 0.0 \tag{4}$$

or the following equation, if the voxel value is negative:

$$c = \frac{n}{p+n} \qquad a < 0.0 \tag{5}$$

where c is the consistency ratio, p is the number of positive votes for that voxel, n is the number of negative votes and a is the amplitude for that voxel. We determine that a voxel is occupied with a buried object if it has a value of c higher than a threshold. A value between 50% and 75% for the threshold usually produces good result. Using this approach, a buried object that returns weak reflections can be detected as well as the object that strongly reflects GPR signal. After the thresholding, we are left with a 3-D data consisting of occupied and empty voxels.

In order to locate and compute the parameters of the buried objects, we group connected voxels into regions. We perform thinning and growing operation on these regions to remove spurious objects. For each group of voxels, we can compute the object's location from its centroid, orientation and size from the shape of the grouped voxels.

4.3.3. Parallel Processing

Migration is very computation intensive due to the amount of processing involved and the size of the 3D data set. A 3-D GPR scan covering a volume of 1m by 1m wide and 1m depth with resolutions of 2cm in the horizontal directions and 0.5 cm in the depth directions consists of 500000 voxels. The migration process needs to be done for each of this voxels. In order to make the computation time as short as possible, many optimization methods are used. First, the data are reduced by decimating the data in the depth direction. The decimation process must take into account the wavelength of the GPR signal and the smallest size objects that we want to detect. It is important not to decimate the data too much because it will cause aliasing in the resulting data. Even if aliasing does not occur, the accuracy of the depth information will deteriorate due to too much decimation. Usually we decimate the data so it has a depth resolution of 0.5 cm or 1.0 cm.

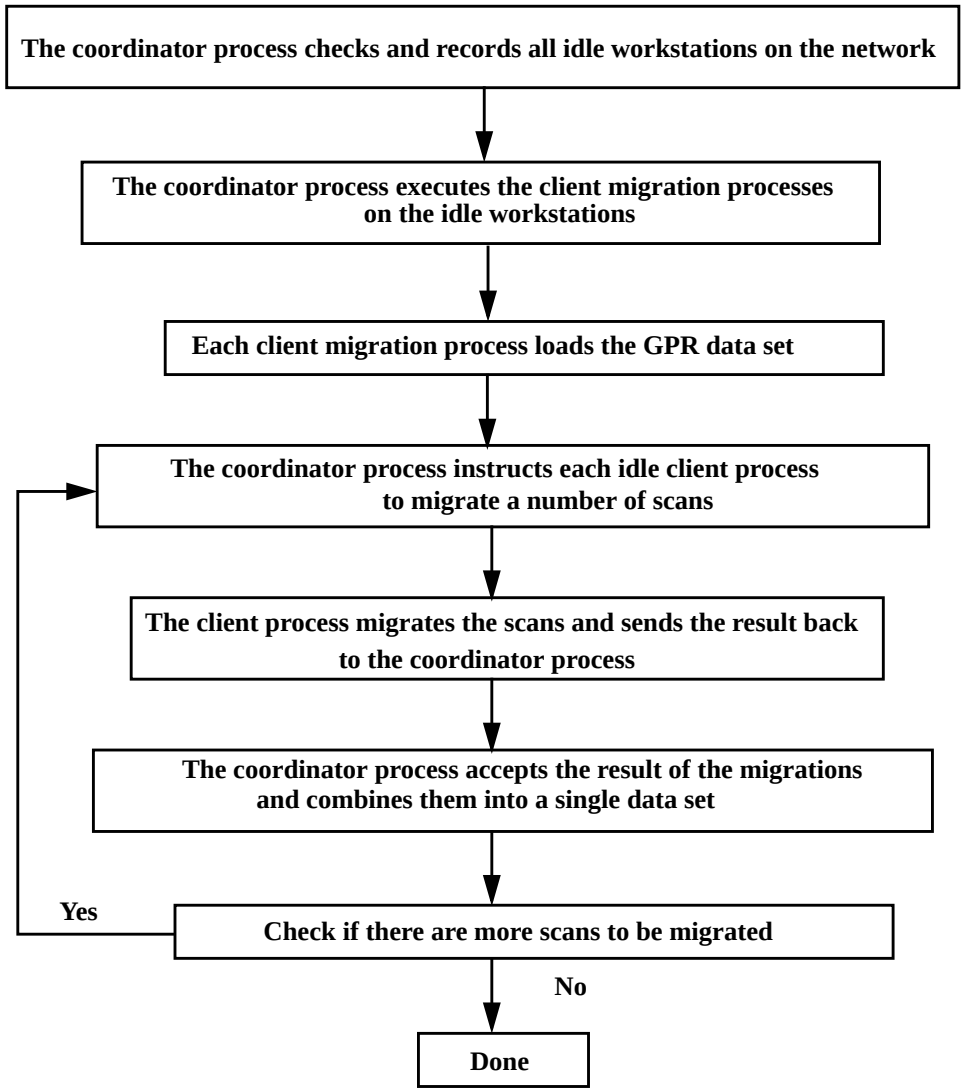

Figure 37: Algorithm for distributed migration

Even with all the optimization, it still took about 30 minute to migrate about 180000 points. This is completely unacceptable because it does not facilitate a real time operation of the system. Since the migration process is inherently a parallel problem, we develop a distributed version of the above algorithm. The distributed migration program uses idle workstations on the network. The programs on different workstations communicate through the network to exchange data. There are basically two different programs for executing the distributed migration. First is the coordinator program which starts the client programs, distribute the data and combine the individual results. Second is the client program which actually migrates the individual scans, based on instructions from the coordinator program.

Algorithm for the distributed migration program is shown in Figure 37. In the current implementation, each client program loads the whole data set, although it might only need a sub-

set of them. This can be a problem if the whole data set does not fit in the workstation memory. In this case the client program can be easily modified so it only loads a subset of the data that is needed to migrate a particular scan.

The features of the distributed migration algorithm are:

- Much faster execution time (depends on the number of workstation used).
- Automatic load balancing. Since the tasks are distributed as small pieces at a time, every time a client process finishes migrating some scans, it is automatically instructed to migrate a new set of scans. So we can use a mix of fast and slow workstations without any detrimental effect.
- Ability to process a massive data set. With a minor modification, the distributed migration program can be used to migrate a very large data set which will not fit in the memory of a single workstation.

The gain in the execution time is almost linear with the number of workstation used. Of course there is some overhead for setting up the processes, loading the data sets etc., but this overhead is negligible compared to the saving in the execution time. The following is an example of the executions time in migrating a 30x30x250 (225000 voxels) data set:

# of workstation	Execution Time
1	24 minutes
5	6 minutes
10	4 minutes

Table 10. The execution time for the migration program running on different number of workstations.

The workstations are mostly Sun Sparcstations 1+ with one or two Sun Sparcstations 2. It is also possible to modify the algorithm to use a cluster of digital signal processors or DSP instead of a cluster of workstation. Migration is especially well suited for DSP because migration needs a lot of multiply-add operations. Many DSP are able to compute a multiply-add operation in a single clock cycle, which should result in a very fast migration processing.

4.3.4. Result

4.3.4.1. Sharpening of the GPR images

The wide beamwidth of the antenna causes GPR data to appear blurred. After the migration process, some of the blurring effect is eliminated. This is especially useful for separating two nearby objects. Figure 38 shows a vertical cross section of two buried objects. After

migration, the two objects are clearly separated and are much more focused. The figure clearly shows that migration transforms the reflection profiles in the raw data into a shape that is a better representation of the objects's actual shape.

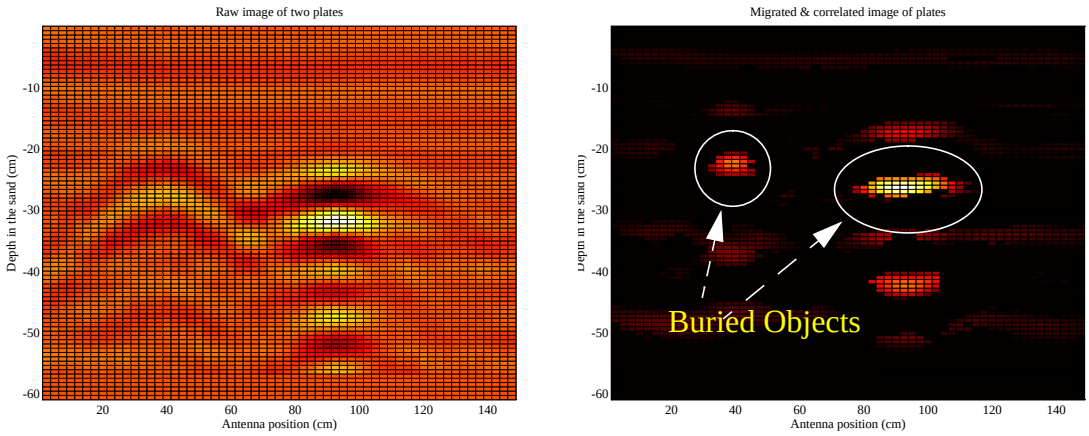

Figure 38: A vertical slice of 3-D GPR data before and after migration. The left buried objects are a small metallic cylinder (radius = 3cm) and the right one is a small metal plate (10cmx10cm).

4.3.4.2. Correcting the tilting effect and putting all the time events at the right depth

Since the beamwidth of the antenna is quite large, reflections detected by the GPR antenna do not necessarily come from the direction below the antenna. This is especially true of the reflections from a non-horizontal flat object such as a diagonal metallic plate. Using migration, we can correct this artifact and move the reflection from below the antenna location to the correct location Figure 39 shows a result of migrating GPR data set of a tilted buried plate. After the migration, the plate is placed at the correct depth and location. The orientation of the plate is also corrected.

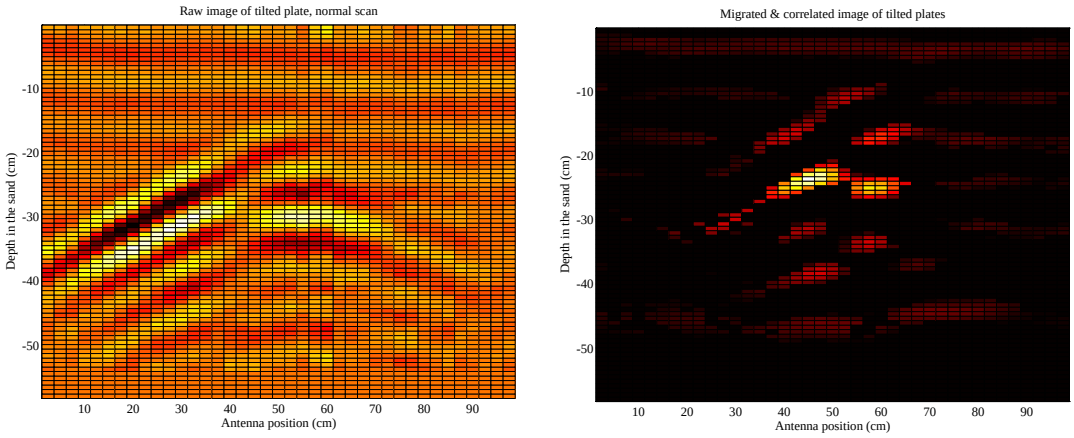

Figure 39: A vertical slice of 3-D GPR data before and after migration. The buried object is a metallic plate buried diagonally. Notice how the reflection is moved to the correct place.

4.3.4.3. Localization of buried object

The result of migration can be processed further to extract the location of the buried object. We group the occupied voxels into connected region. Then for each region we can compute its location and the orientation. An example of this is shown in Figure 40. There are two buried objects and both of them are automatically found by thresholding and grouping the connected voxels after migration processing.

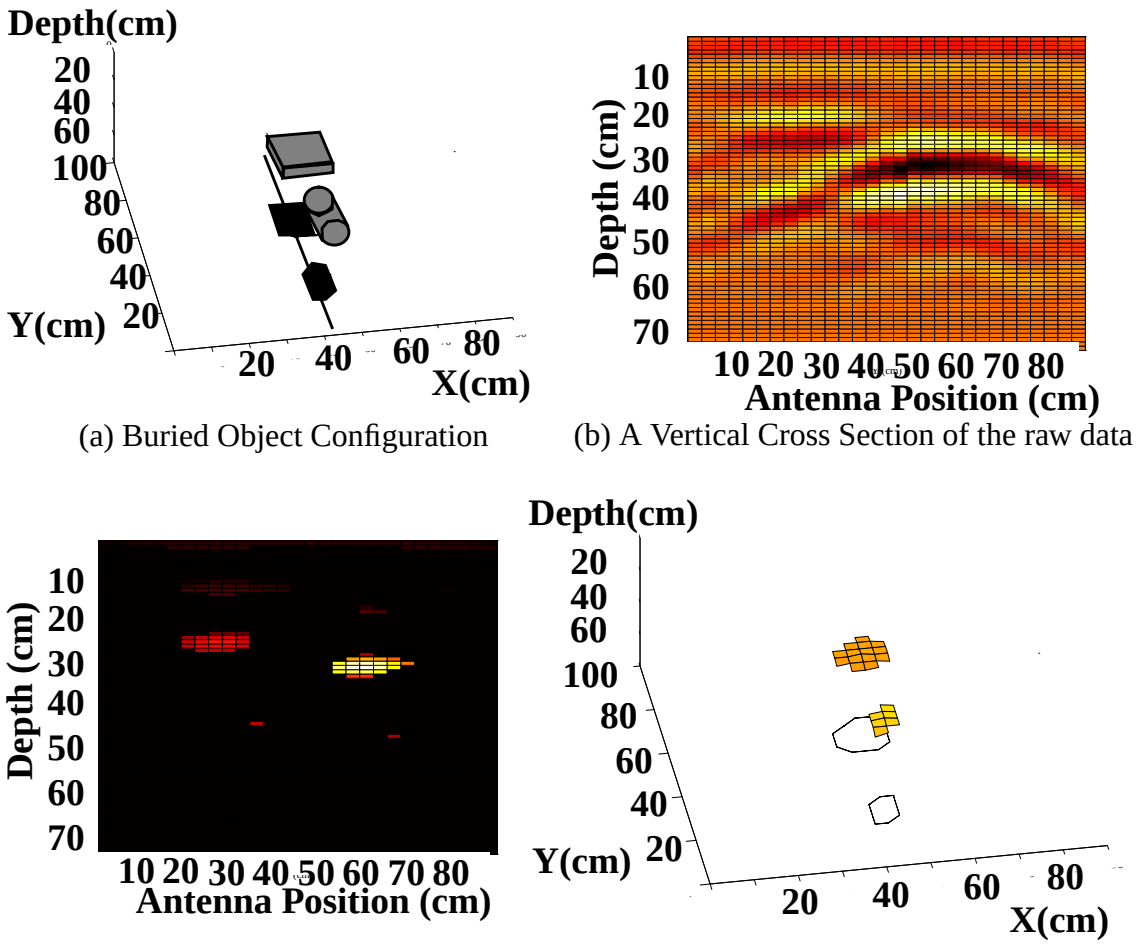

(a) Buried Object Configuration

(b) A Vertical Cross Section of the raw data

(c) A Vertical Cross Section of the processed data

(d) Buried Object Map

Figure 40. Localization of two buried objects using distributed migration.

4.4. 3-D Migration using Reflector Pose Estimation

Upon further investigation, we found that the above migration process is not really efficient. Each possible reflections in the reflected signals must be distributed along all possible incoming directions. Since most reflections come from a single surface, this is a waste of computation time.

If we assume that the reflection is coming from a single flat surface, the above migration process can be made much more efficient. The method is illustrated in Figure 41. Let assume we take a GPR scan at Antenna Position #1 and after processing, be able to compute the distance to a surface of a buried object. In the figure the distance is denoted by *range1*. Then we move the antenna to position#2 and take another scan. In the second scan the distance to the object is *range2*. From the individual scan, the only information we have is the distance to the reflector. The reflector can be located anywhere at a distance of *range1* within the antenna beamwidth. The key observation is the following. If we make the assumption that the reflector is a flat surface, we can use the relationship of *range1* and *range2* to constrain the location of the reflector. From the two range values *range1* and *range2* to the reflector, we can compute the direction of the reflection using equation(6).

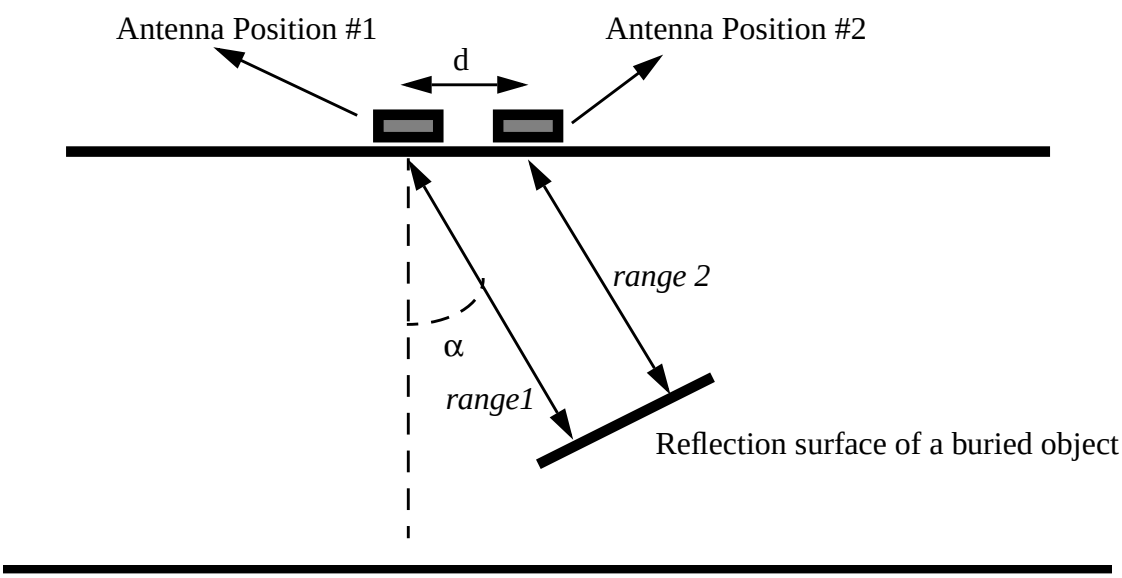

Figure 41. Computing the direction of the reflecting surface from two distance measurements

$$\alpha = \operatorname{asin}\left(\frac{range1 - range2}{d}\right) \tag{6}$$

In a GPR image, the angle of the reflection can be computed directly from the angle of the reflector in the GPR data. Figure 42 shows the geometry of the problem and the equation to compute the true angle of the surface from the apparent angle of the surface in GPR data. Once the true angle is computed, the reflection surface can be moved or migrated to its true location as shown in Figure 43. Figure 44 shows an example of the result of reflector pose

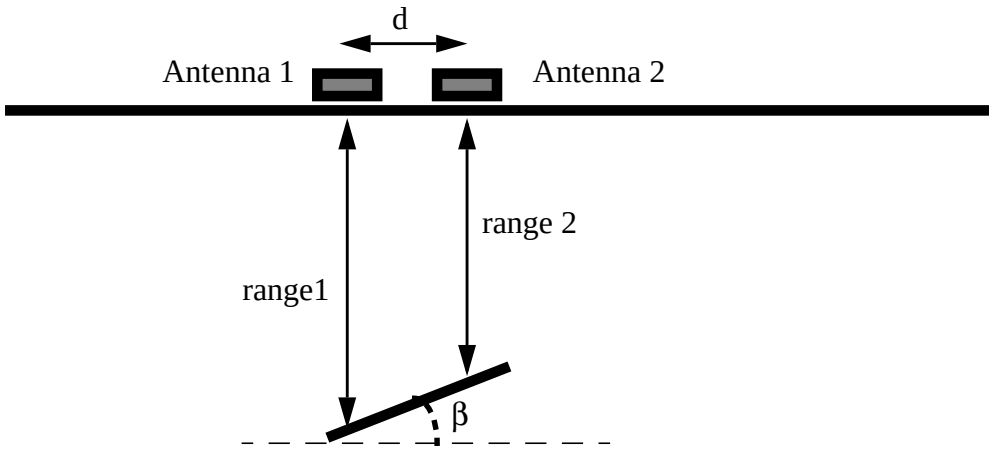

$$\alpha = \operatorname{asin}(\tan(\beta))$$

α = true angle of the surface

β = apparent angle of the surface in the GPR image

Figure 42. The resulting apparent angle in the GPR image as a result of the two scans of the plate in Figure 41.

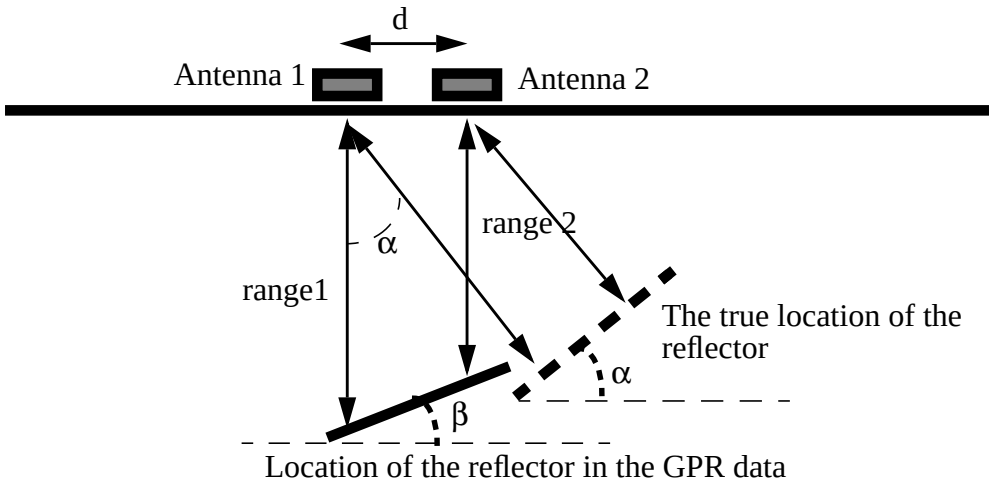

Figure 43. The true location of the reflector can be computed once the true angle is computed from the orientation of the reflector in the GPR data.

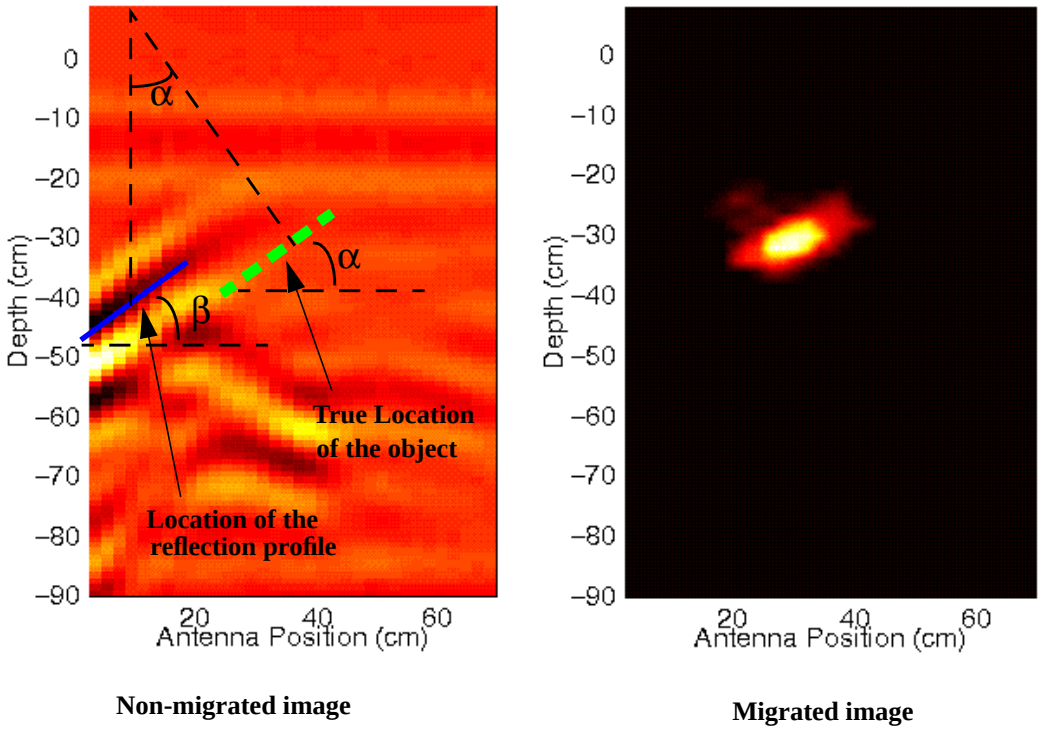

Figure 44. An example of migration using reflector pose estimation.

estimation on real GPR data. It shows a cross section of a buried plate. In the non-migrated image, the plate is located at the left side of the image, although its actual location is at the center of the image (denoted by the dashed line). In the migrated image, the object is moved to its correct location. In 3-D the same process can be done, except instead of a single orientation angle, we need to compute the real surface normal from the apparent surface normal. So instead of having a single angle denoting the orientation of the object, in 3-D we have two angles to define the surface normal.

Figure 41 to Figure 44 show how to compute the actual location of reflector under the assumption that the reflector is flat. We can also compute the actual location of the reflector under the assumption that the reflections are coming from a point reflector, as shown in Figure 45.

In this case, the direction of the reflection α is computed as follow:

$$\alpha = \operatorname{asin}\left(\frac{(range1^2 - range2^2) + d^2}{2.0 \times d \times range1}\right) \tag{7}$$

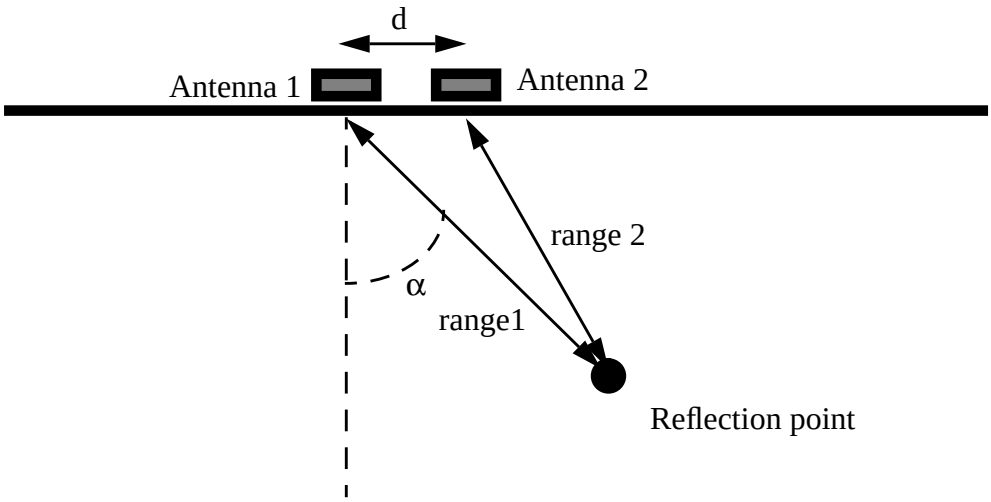

Figure 45. Computing the direction of the small reflecting object from two distance measurements

Figure 46 shows that the reflection angle computation for a point reflector yields a very similar result regardless whether the angle is computed using the point reflector or the flat reflector assumption. Because the error is very small and negligible, we are able to compute the angle of reflection using either assumptions. In our experiment, we assume that the reflection is coming from a flat reflector, although in reality, it might come from a point reflector.

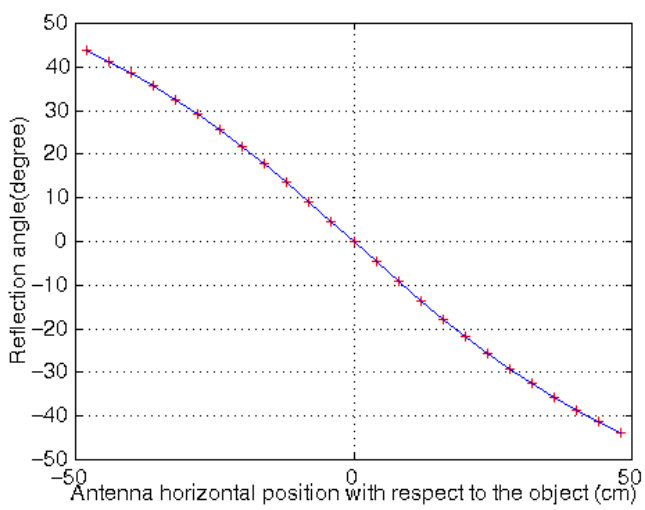

Figure 46. Computation of incoming reflection angle assuming that the reflector is a flat reflector (solid line) and a point reflector(+ sign).

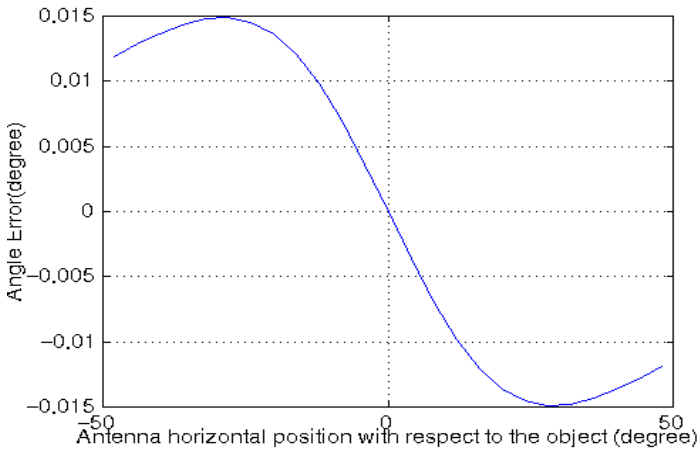

Figure 47.The error in the reflection angle computation using the flat reflector method for a point reflector.

One important thing to notice is that in order to be able to compute the real surface orientation, we need the distance to the reflector. As in coherent summation migration, an accurate propagation velocity estimate is needed to convert the propagation time to distance. An error in the velocity estimate will result in imperfect migration, since the reflector will be migrated to incorrect location. If the velocity estimate is lower than the actual propagation velocity, then the reflector will be moved too little. If the velocity estimate is higher than the actual propagation velocity, then the reflector will be moved too much.

The actual 3-D migration involves computing the 3-D gradient of each voxel. This 3-D gradient is used as the apparent surface normal. We start with a 3-D volume data. These data are stored in an array, let's call it $v(x,y,z)$. x and y is the horizontal location of the scan while z is the time index. $v(x,y,z)$ is the amplitude of the signal at horizontal location x and y and time index z. Now to migrate the data, we need to compute the real location of each of the voxels. This is done by computing the orientation of the 3-D gradient at that voxel. In order to do this, we need to find the voxels that has the closest value to this voxel in the neighbouring scans. Since the value of a single voxel is noisy, we need to find similar signal patches in neighboring scans or columns instead. For the similarity value between two small patches from the migrated voxel's scan and a neighbouring scan, we use the sum of absolute differences. So for each voxel, we need to find the location of a similar patch in the neighbouring scans. Figure 48 illustrate this operation. It shows three similar patches in neighboring scans that can be used to compute the 3-D gradient of a voxel.

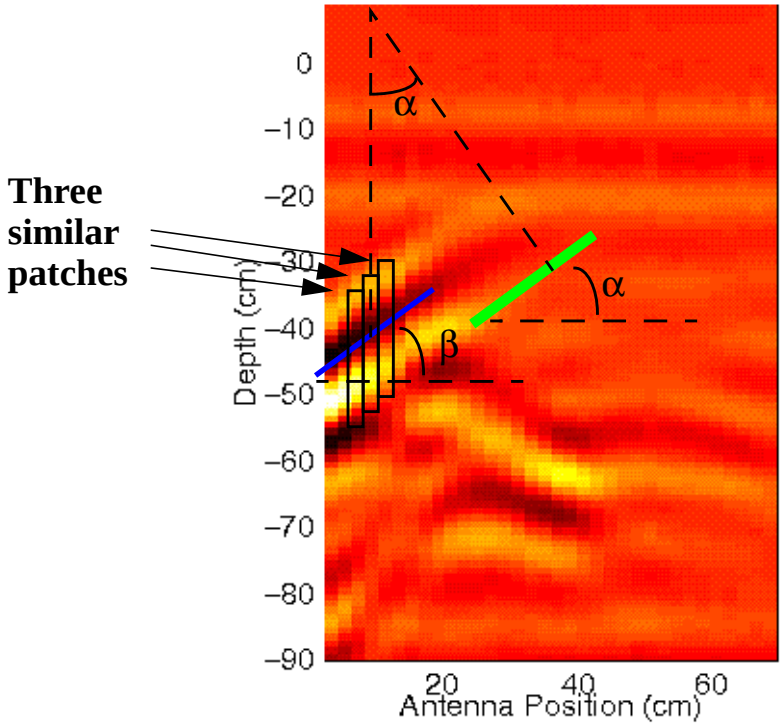

Figure 48. Computing the 3-D gradient of a voxel by locating similar patches at neighboring scans.

This is done by finding the offset l_{min} that minimizes the following equation:

$$d_{i,j} = \sum_{k=-\text{max_offset}}^{\text{max_offset}} abs(v(x, y, z+k) - v(x+i, y+j, z+l_{min}+k)) \qquad (8)$$

$$-max_offset < l_{min} < max_offset$$

max_offset is the maximum permissible offset value. This is determined by the following equation:

$$max_offset = \frac{horizontal\ distance\ between\ the\ two\ scans}{depth\ resolution} \qquad (9)$$

If there is uncertainty in the propagation velocity, we use a lower estimate of the propagation velocity to compute *depth resolution*. This will result in a conservatively large estimate of the *max_offset*. We compute the value l_{min} for every neighboring columns of the migrated voxel, resulting in 8 values. Each of this l_{min} value is converted to distance value:

$$z_{min} = l_{min} \times \text{depth resolution} \tag{10}$$

At this point we have 8 depth values for the 8 neighboring columns of the migrated voxel. They can be displayed as a 2-D array as shown in the following figure:

$Z_{min}(-1,1)$	$Z_{min}(0,1)$	$Z_{min}(1,1)$
$Z_{min}(-1,0)$	*Depth of the migrated voxel*	$Z_{min}(1,0)$
$Z_{min}(-1,-1)$	$Z_{min}(0,-1)$	$Z_{min}(1,-1)$

Figure 49. A 2-D array of depth values of the migrated voxel and its neighbouring voxels. The neighbouring voxels are the voxels which have the same amplitude as the migrated voxel.

From this 8 depth values, we can fit a plane using a least square fitting technique. After the fitting process, we obtain the plane equation:

$$ax + by + cz + d = 0 \tag{11}$$

From this equation, we can compute the normalized surface normal n, where $n_x = a$, $n_y = b$, $n_z = c$,. The surface normal is then converted into two angles, the elevation and the azimuth angle. The elevation angle is computed using:

$$elevation = \operatorname{acos}\left(\frac{sqrt(n_x^2 + n_y^2)}{n_z}\right) \tag{12}$$

If *elevation* is larger than a very small threshold then the azimuth angle is computed using;

$$azimuth = \operatorname{atan2}(-n_y, -n_x) \tag{13}$$

From these two angles, we can compute the real location of the voxel. First we need to compute the real elevation angle:

$$real_elevation = \operatorname{asin}(\tan(elevation)) \tag{14}$$

Using the *real_elevation* and *azimuth* value, we can compute the new location for the migrated voxel. If the migrated voxel is originally located at *x, y, depth, t*hen the real location is computed using the following equations:

$$x_{\text{real}} = x + \cos(azimuth) \times \sin(elevation) \times depth \tag{15}$$

$$y_{\text{real}} = y + \sin(azimuth) \times \sin(elevation) \times depth \tag{16}$$

$$z_{\text{real}} = \cos(elevation) \times depth \tag{17}$$

It is obvious from the above equations that we need an accurate estimate for the propagation velocity in order to migrate the voxels accurately. This is the main weakness of the migration method, although in most cases the propagation velocity can be obtained pretty easily using a dielectric measurement tool or other methods. In addition, the two migration methods also require that the soil above the buried objects to have a uniform average propagation velocity. If this is not true then the result of the migration will be inaccurate.

4.4.1. Postprocessing

As in the coherent summation migration, the result of the efficient migration can be thresholded to further improve the image quality and to locate the buried objects. We use a similar method as discussed in Section 4.3.2. We threshold the data based on the number of consistent votes that each voxel get, instead of the amplitude of the voxel. Since this method does not take the strength of the reflection into account, it does not discriminate weak reflections from small objects.Using this thresholding criteria, only the voxels that consistently get votes with the same sign are visible.

The threshold value must be picked carefully so as to avoid getting rid of useful information. The desirable rate of false positive or false negative must be taken into account. For most buried object retrieval application, it is desirable to minimize the false negative rate. It is preferable to think falsely that there is a buried object rather than to think falsely that there is no buried object. As a result it is better to set the threshold value to the lowest value possible.

Once the volume data are thresholded, we groups the voxel into connected regions. For each regions, we use principal component analysis to compute the object's parameters.

4.4.2. Result

In this section, we will show several 3-D GPR data before and after migration using reflector pose estimation. For each data sets, we will show two series of vertical cross sections taken in the direction of X and Y axis. We will also show a horizontal cross section of the 3-D data. For each of the cross sections, we display the data before and after migration. For each buried object, we will show a table of the computed parameters along with their actual values. The parameters that we compute are the location of the horizontal centroid position, length, width, azimuth, elevation angle and the depth of the reflection surface at the centroid location. The azimuth value has meaning only when the object is asymmetrical. It does not have any useful meaning if the object is more or less symmetrical. All the others parameters should be self explanatory except for the last one, the depth of the reflecting surface at the centroid. This is the depth of the object's surface at the horizontal location of the centroid. This is the most meaningful depth information, since GPR measures the distance to the reflector's surface, not to the centroid of the object. We will show the results of migration using reflector pose estimation for a single horizontal metallic pipe, two metallic objects, a non-horizontal metallic pipe and a plastic anti-tank mine. The figures for all results are shown from page 82 to page 93.

4.4.2.1. Migration of a horizontal metallic pipe

Figure 50 shows the vertical cross section (X axis) of the buried pipe before migration. The reflection profile created by the pipe can be clearly seen. After 3-D migration, the reflection profile is focused and a clearer image of the pipe appears, as shown in Figure 51. Figure 52 and Figure 53 show the same buried object but in the direction of Y axis. The buried pipe looks like a flat reflector in this direction. Figure 54 and Figure 55 show the horizontal cross sections and Figure 56 shows two 3-D views of the migrated data. The computed parameters of the buried pipe are shown in Table 11.

Parameter	Real Value	Computed Value	Error
Centroid horizontal position (cm)	(0.0,-3.5)	(-4.21,-0.36)	5.71
Width (cm)	12.0	15.31	3.31
Length (cm)	30	24.14	-5.86
Azimuth (degrees)	90.0	91.44	1.44
Elevation (degrees)	0.0	0.944	0.94
Depth of the reflecting surface at the centroid (cm)	25	30.39	5.39

Table 11. Parameter for the metallic pipe in Figure 56.

4.4.2.2. Migration of a metallic coffee cup and a small metallic plate

In this experiment we buried two small metallic object separated by about 12cm. The objects are a coffee cup and the a metal plate. Figure 57 and Figure 58 shows the vertical cross sections of the data in the X axis before and after migration. It is not clear from these cross sections that there are two buried objects. If we display the cross section in the Y axis as shown in Figure 59 and Figure 60, we can clearly see the two separate objects. They can also be clearly seen in the horizontal cross sections in Figure 62 and in 3-D views in Figure 63. The computed the parameters of the two objects are shown in Table 12 for the coffee cup and in Table 13 for the small plate.

Parameter	Real Value	Computed Value	Error
Centroid horizontal position (cm)	(0.0,-18.0)	(-1.33,-20.67)	2.98
Widths (cm)	10.0	11.29	1.29
Length (cm)	15	18.29	3.29

Table 12. Parameter for the metallic coffee cup

Parameter	Real Value	Computed Value	Error
Azimuth (degrees)	90.0	97.26	7.26
Elevation (degrees)	0.0	3.61	3.61
Depth of the reflecting surface at the centroid (cm)	24.0	25.78	1.78

Table 12. Parameter for the metallic coffee cup

Parameter	Real Value	Computed Value	Error
Centroid horizontal position (cm)	(0.0,7.0)	(-1.5,10.53)	3.84‘
Width (cm)	11.0	21.75	10.75
Length (cm)	11.0	24.22	13.22
Elevation (degrees)	0.0	5.21	5.21
Depth of the reflecting surface at the centroid (cm)	25.0	33.26	8.26

Table 13. Parameter for the metallic plate.

4.4.2.3. Migration of a non-horizontal metallic pipe

So far we have presented the result of 3-D migration of horizontal objects. In order to see if we can extract the orientation of the objects as well as the location, we buried a pipe with an azimuth angle of 45 degrees and elevation angle of -16 degrees. Figure 64 to Figure 67 shows the vertical cross sections before and after 3-D migration. It is very hard to determine the orientation of the pipe from these cross sections, even after the data are migrated. On the other hands the horizontal cross section, as shown in Figure 68 and Figure 69, show the horizontal orientation of the pipe quite clearly. The vertical and the horizontal orientation of the pipe can be clearly seen in the two 3-D views in Figure 70. The computed parameters are shown in Table 14. The biggest error is in the length of the pipe. If we lower the threshold value, then the computed length will closer to the actual length. The orientation of the object is computed quite accurately, especially the azimuth angle.

Parameter	Real Value	Computed Value	Error
Centroid horizontal position (cm)	(-9.0,-2.5)	(-11.24,-14.24)	11.95
Width (cm)	6.0	6.90	0.90
Length (cm)	60.0	22.00	-38.0
Azimuth (degrees)	36.25	36.0	-0.25

Table 14. Parameter for the metallic pipe in Figure 96.

Parameter	Real Value	Computed Value	Error
Elevation (degrees)	-16.0	-11.57	4.03
Depth of the reflecting surface at the centroid (cm)	21.0	20.40	-0.6

Table 14. Parameter for the metallic pipe in Figure 96.

4.4.2.4. .Migration of a plastic anti-tank mine

All the results that we have presented so far involves imaging metallic objects. Since the ability to detect and locate non-metallic objects is one of GPR advantages, we buried a plastic anti tank mine. Figure 71 through Figure 76 shows the cross sections of the data before and after migration in various directions. The mine can be clearly seen in all the cross sections. Two 3-D views of the mine are shown in Figure 77. The computed parameters are shown in Table 15.

Parameter	Real Value	Computed Value	Error
Centroid horizontal position (cm)	(0.0,0.0)	(-2.07,2.78)	3.47
Width (cm)	30.0	9.77	-20.23
Length (cm)	30.0	15.48	-14.52
Elevation (degrees)	0.0	-0.33	-0.33
Depth of the reflecting surface at the centroid (cm)	17.0	19.58	2.58

Table 15. Parameter for the plastic anti-tank mine.

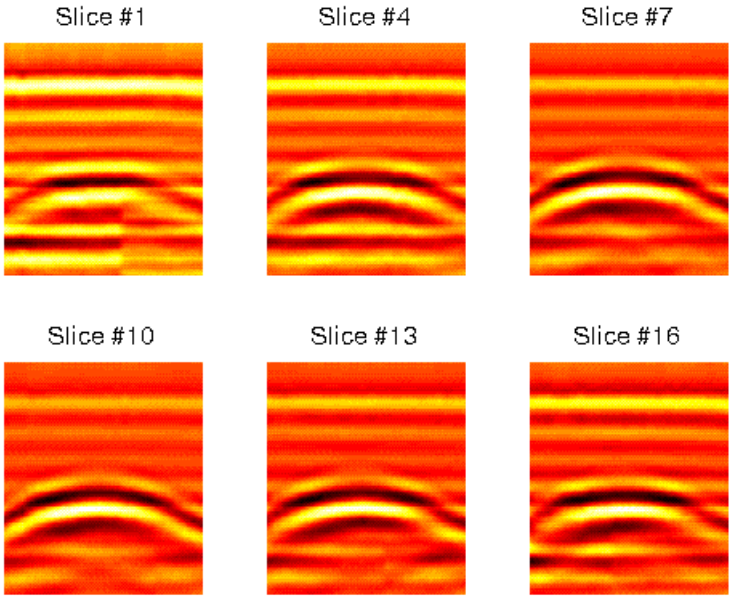

Figure 50: Vertical cross sections (X axis) before 3-D migration. The object is a metallic pipe (diameter=12cm, length=30cm).

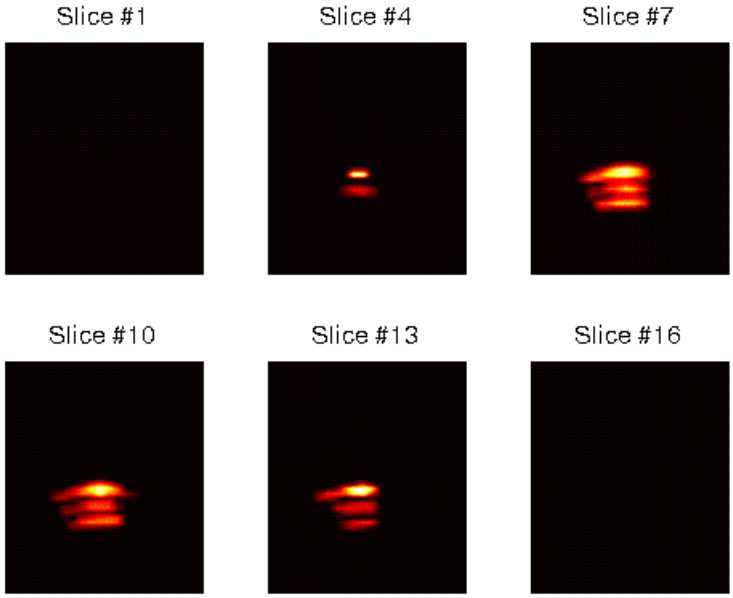

Figure 51: Vertical cross sections (X axis) after 3-D migration. The object is a metallic pipe (diameter=12cm, length=30cm).

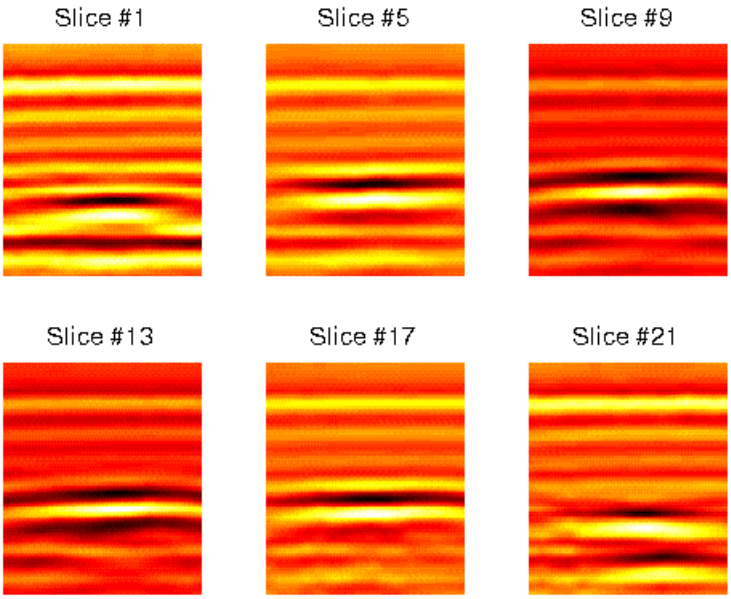

Figure 52: Vertical cross sections (Y axis) before 3-D migration. The object is a metallic pipe (diameter=12cm, length=30cm).

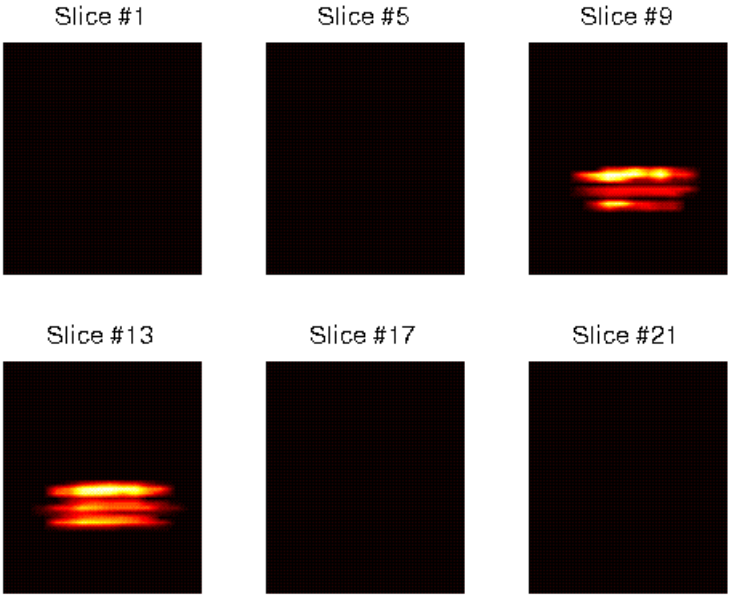

Figure 53: Vertical cross sections (Y axis) after 3-D migration. The object is a metallic pipe (diameter=12cm, length=30cm).

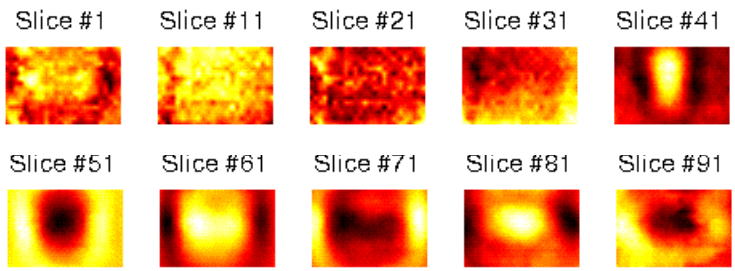

Figure 54: Horizontal cross sections before 3-D migration. The object is a metallic pipe (diameter=12cm, length=30cm).

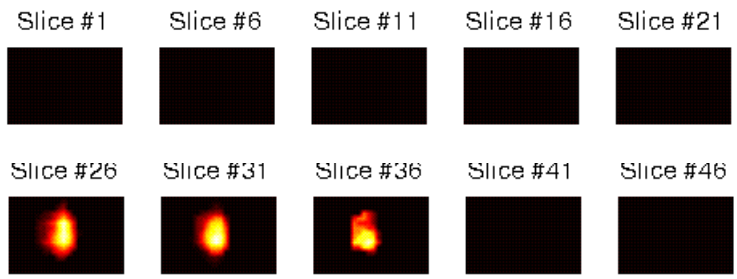

Figure 55: Horizontal cross sections after 3-D migration. The object is a metallic pipe (diameter=12cm, length=30cm).

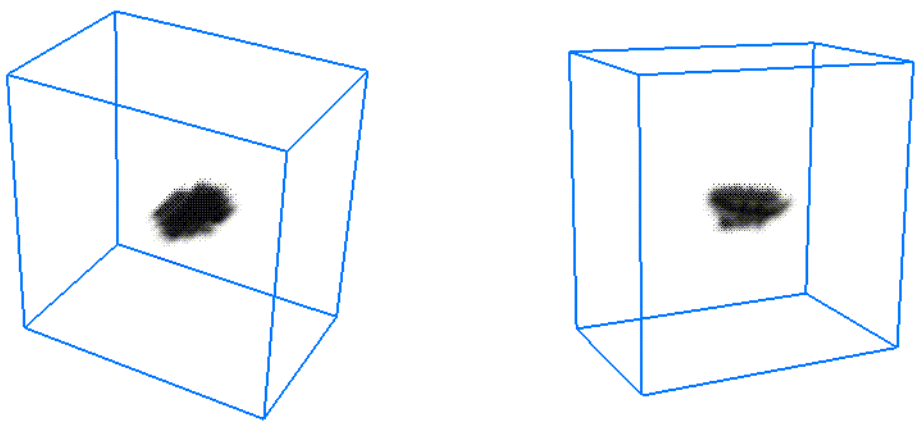

Figure 56: Two views of the occupied voxels showing the buried pipe.

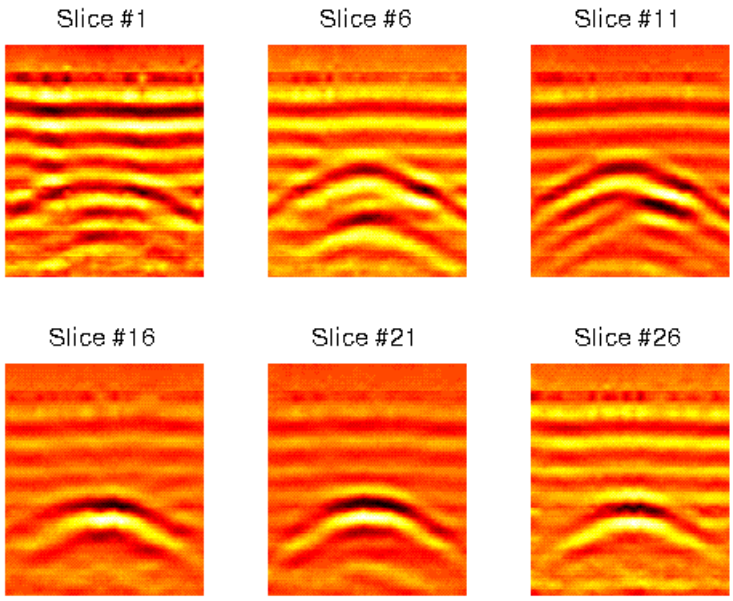

Figure 57: Vertical cross sections (X axis) before 3-D migration. The first object is a metallic coffee cup (diameter= 10cm, length = 15cm) and the second object is a small metallic plate (11cm by 11cm).

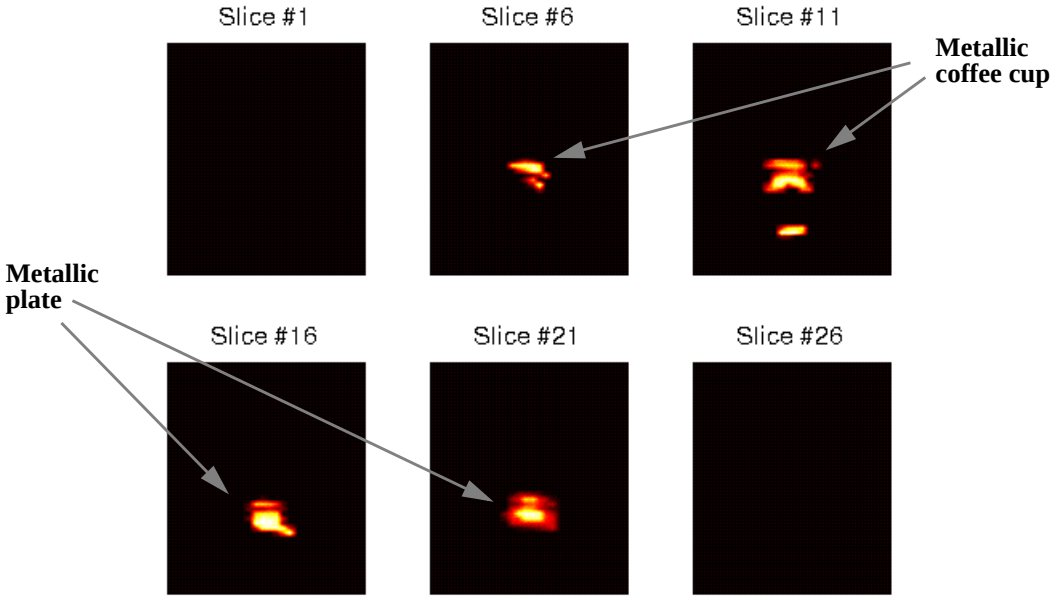

Figure 58: Vertical cross sections (X axis) after 3-D migration. The first object is a metallic coffee cup (diameter= 10cm, length = 15cm) and the second object is a small metallic plate (11cm by 11cm).

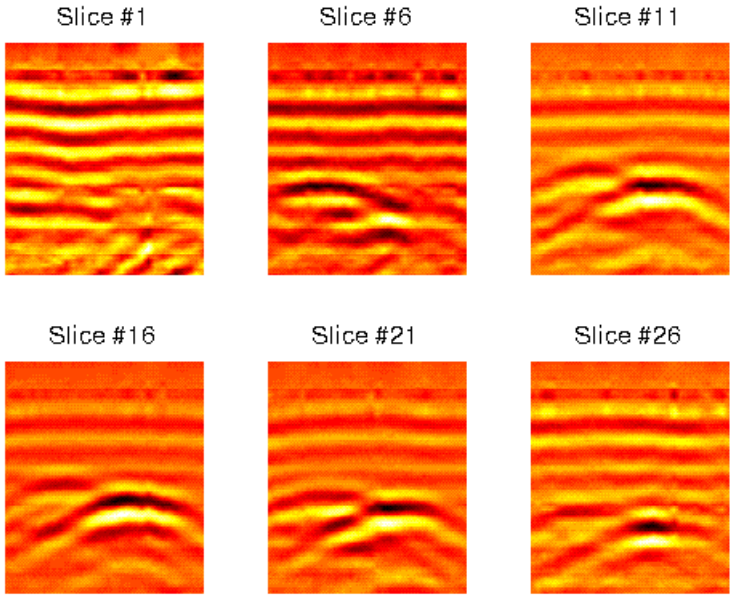

Figure 59: Vertical cross sections (Y axis) before 3-D migration. The first object is a metallic coffee cup (diameter= 10cm, length = 15cm) and the second object is a small metallic plate (11cm by 11cm).

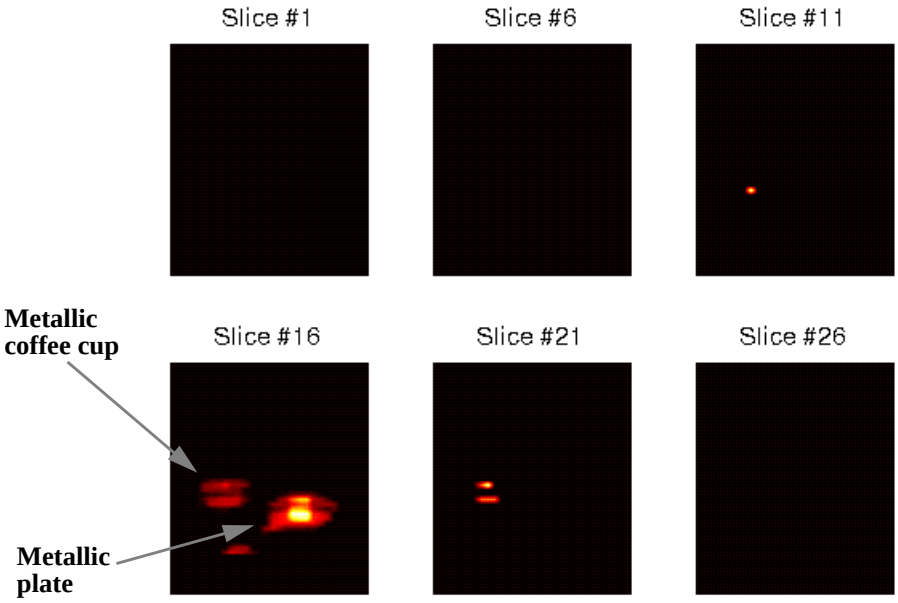

Figure 60: Vertical cross sections (Y axis) after 3-D migration. The first object is a metallic coffee cup (diameter= 10cm, length = 15cm) and the second object is a small metallic plate (11cm by 11 cm).

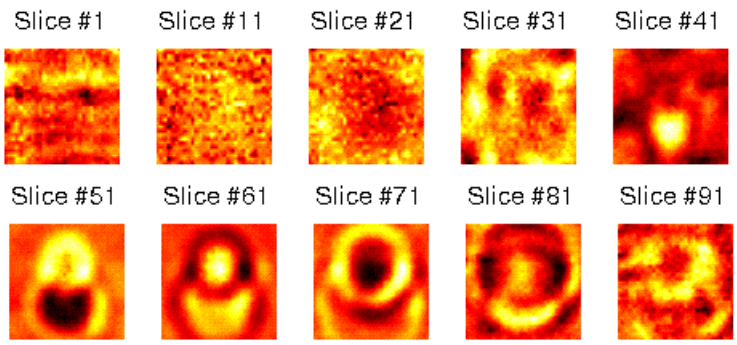

Figure 61: Horizontal cross sections before 3-D migration. The first object is a metallic coffee cup (diameter= 10cm, length = 15cm) and the second object is a small metallic plate (11cm by 11 cm).

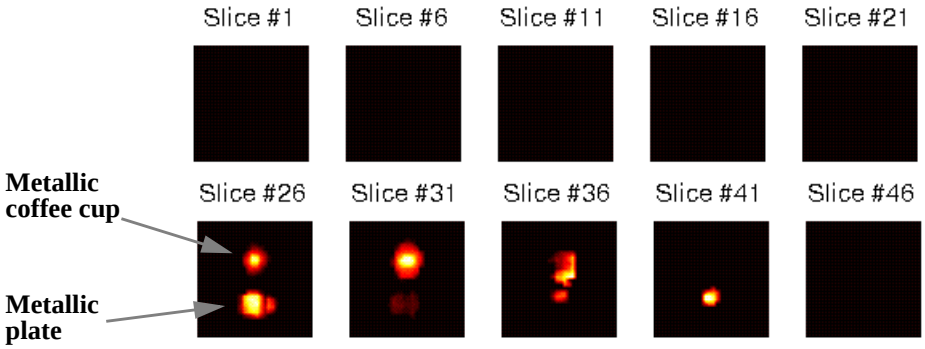

Figure 62: Horizontal cross sections after 3-D migration. The first object is a metallic coffee cup (diameter= 10cm, length = 15cm) and the second object is a small metallic plate (11cm by 1cm).

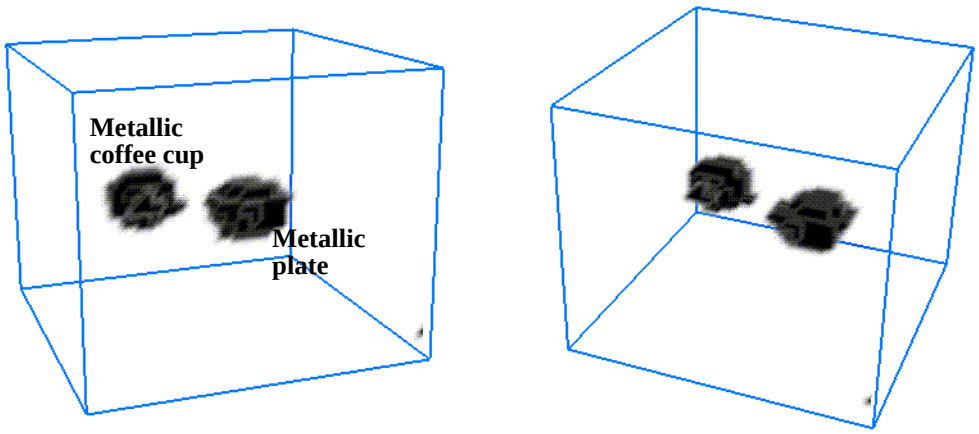

Figure 63: Two views of the occupied voxels after 3-D migration showing the buried metallic coffee cup and plate.

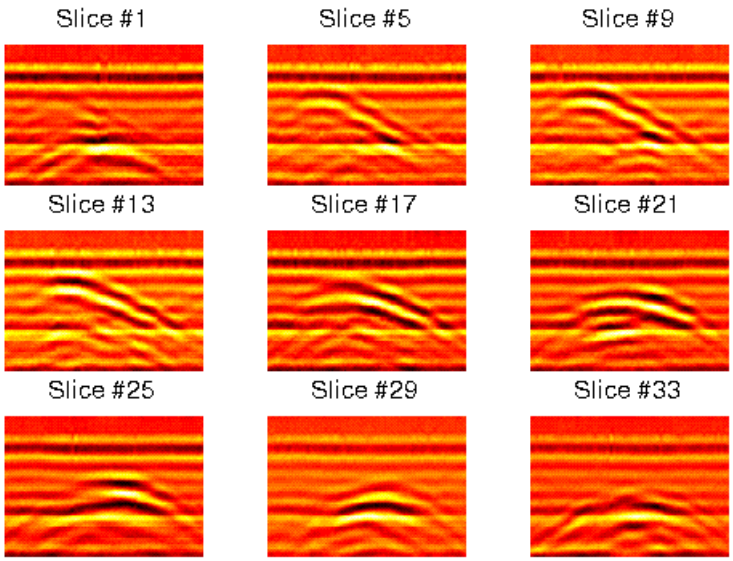

Figure 64: Vertical cross sections (X axis) before 3-D migration. The object is a diagonally oriented metallic pipe (diameter=7cm, length=60cm).

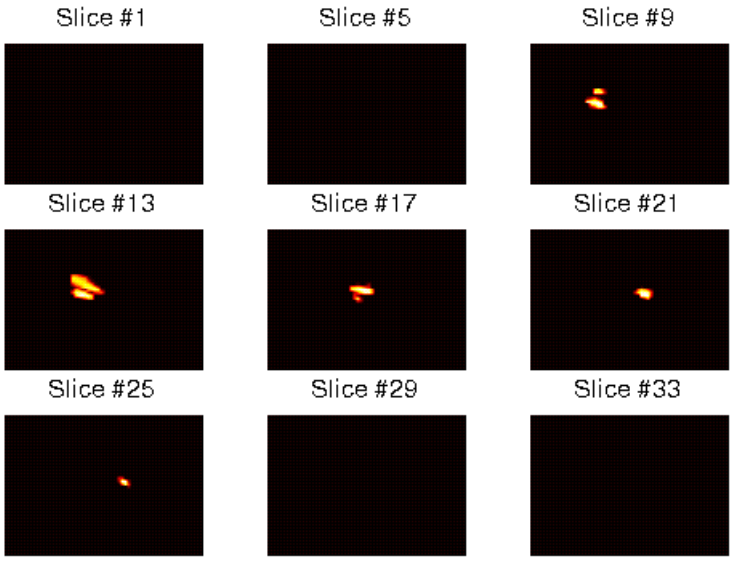

Figure 65: Vertical cross sections (X axis) after 3-D migration. The object is a diagonally oriented metallic pipe (diameter=7cm, length=60cm).

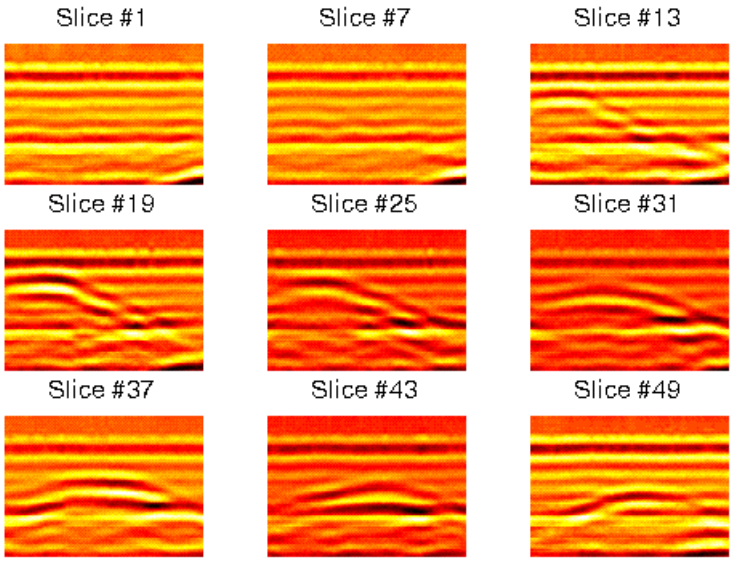

Figure 66: Vertical cross sections (Y axis) before 3-D migration. The object is a diagonally oriented metallic pipe (diameter=7cm, length=60cm).

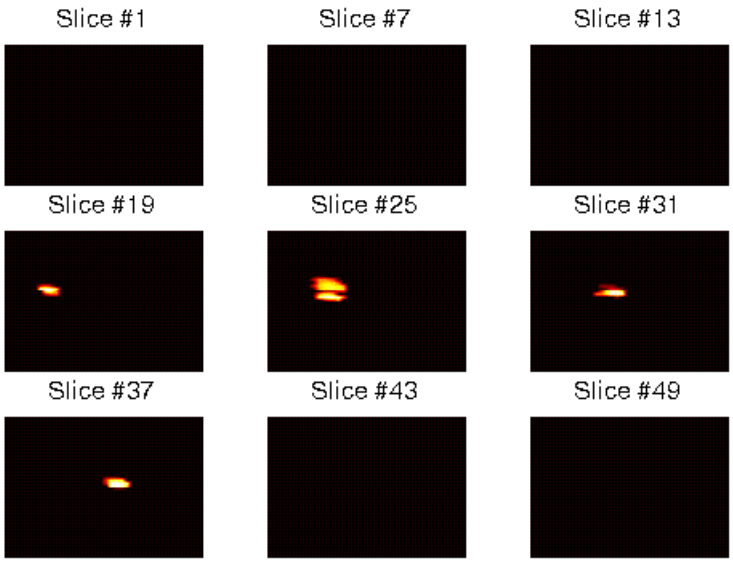

Figure 67: Vertical cross sections (Y axis) after 3-D migration. The object is a diagonally oriented metallic pipe (diameter=7cm, length=60cm).

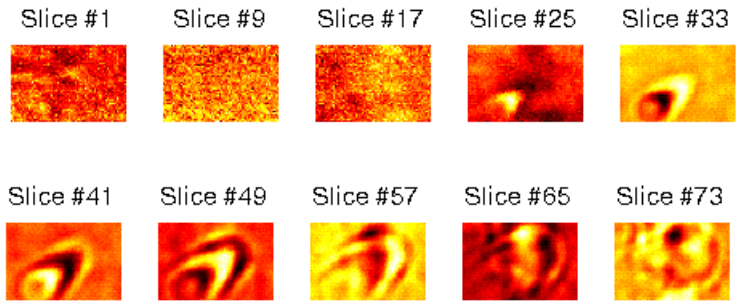

Figure 68: Horizontal cross sections before 3-D migration. The object is a diagonally oriented metallic pipe (diameter=7cm, length=60cm).

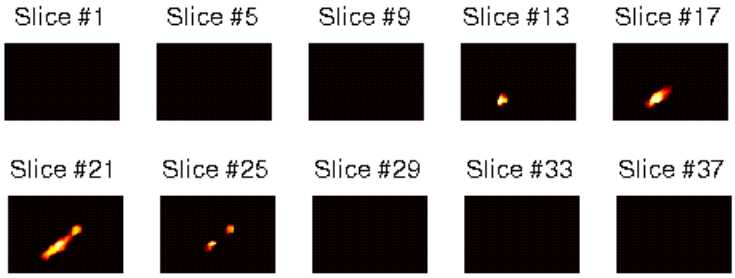

Figure 69: Horizontal cross sections after 3-D migration. The object is a diagonally oriented metallic pipe (diameter=7cm, length=60cm).

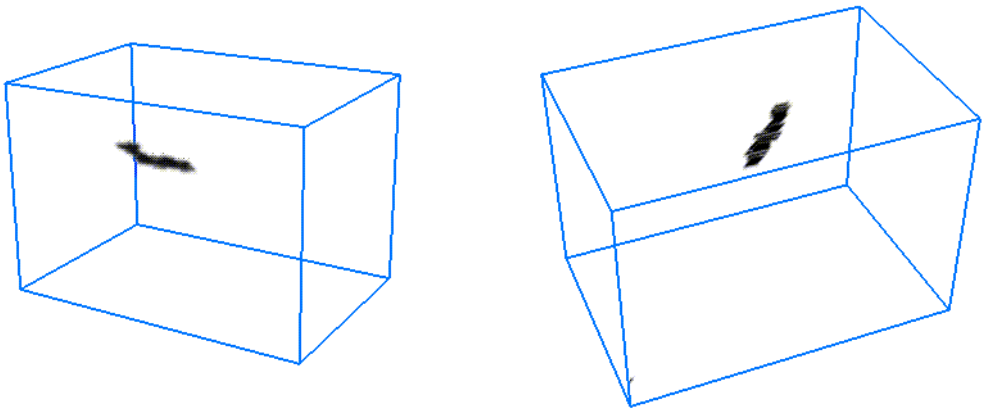

Figure 70: Two views of the occupied voxels after 3-D migration showing the diagonal pipe.

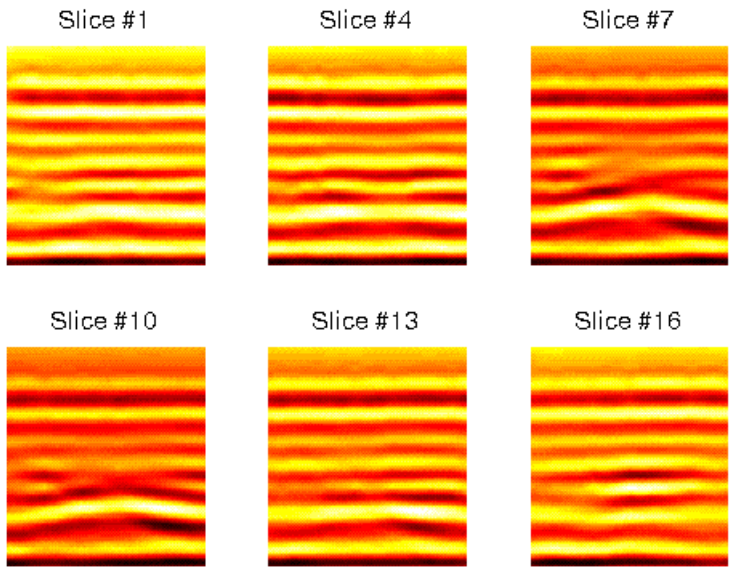

Figure 71: Vertical cross sections (X axis) before 3-D migration. The object is an anti tank plastic mine (diameter=30cm, thickness=10cm).

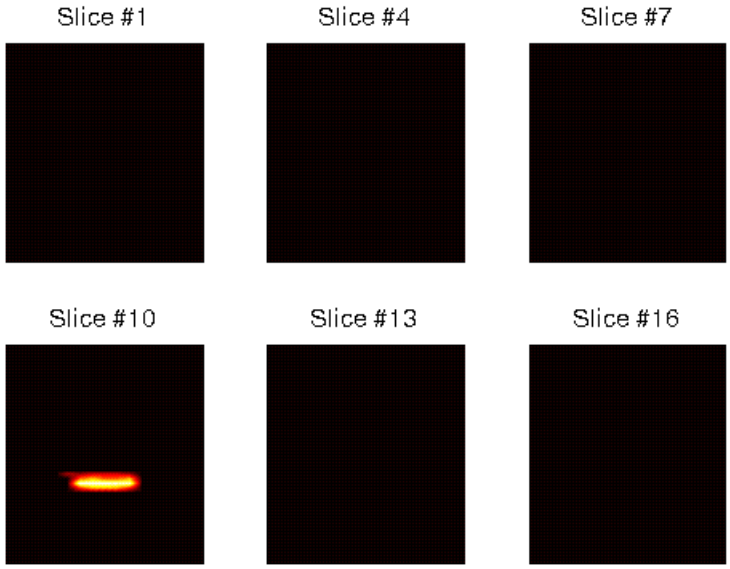

Figure 72: Vertical cross sections (X axis) after 3-D migration. The object is an anti tank plastic mine (diameter=30cm, thickness=10cm).

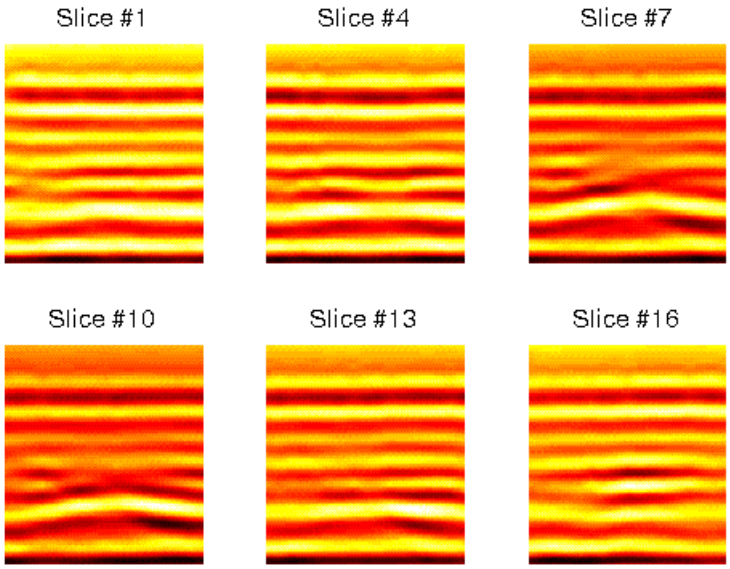

Figure 73: Vertical cross sections (Y axis) before 3-D migration. The object is an anti tank plastic mine (diameter=30cm, thickness=10cm).

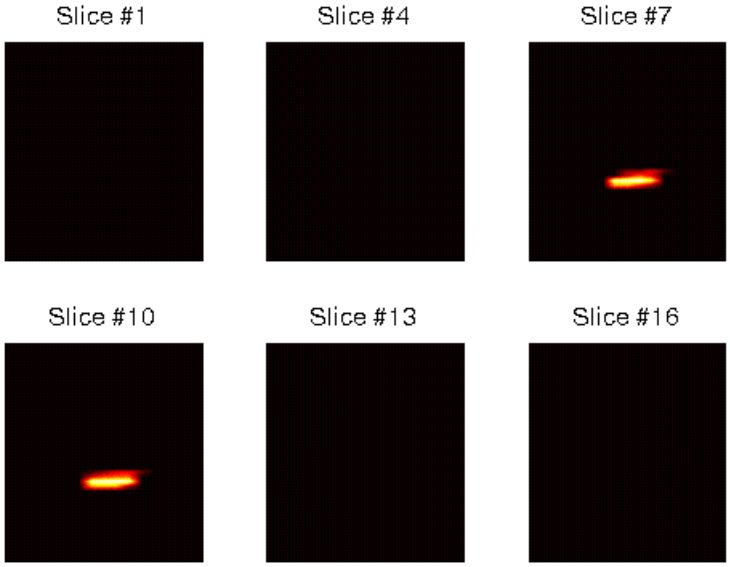

Figure 74: Vertical cross sections (Y axis) before 3-D migration. The object is an anti tank plastic mine (diameter=30cm, thickness=10cm).

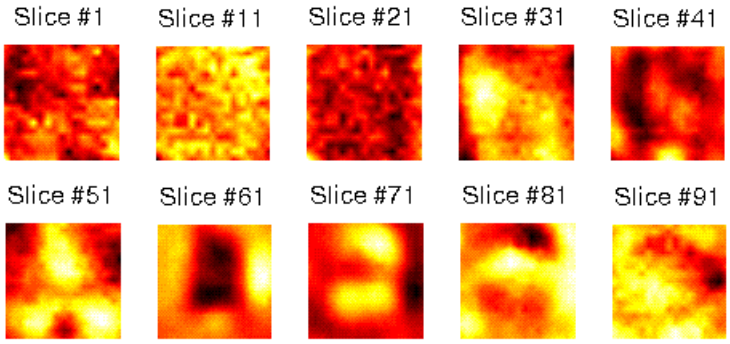

Figure 75: Horizontal cross sections before 3-D migration. The object is an anti tank plastic mine (diameter=30cm, thickness=10cm).

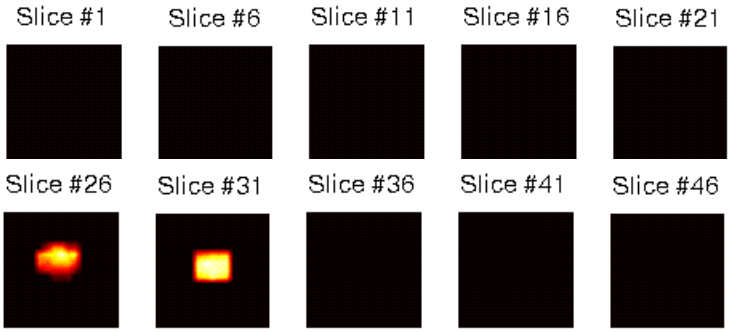

Figure 76: Horizontal cross sections before 3-D migration. The object is an anti tank plastic mine (diameter=30cm, thickness=10cm).

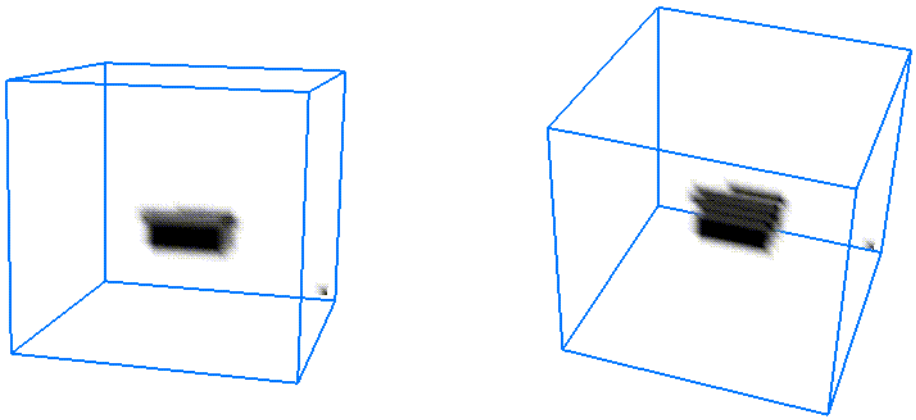

Figure 77. Two views of the occupied voxels after 3-D migration showing the buried anti tank plastic mine.

Chapter 5. Surface Based Processing

5.1. Overview of Surface Based GPR Data Processing

In volume based processing, we directly process the 3-D volume data to find buried objects embedded in it. The result is also a 3-D volume data, where buried objects are represented by the occupied voxels. In surface based GPR data processing, we first reduce the 3-D volume data into a series of 2.5-D surfaces using 3-D segmentation technique. The surfaces denote possible buried objects or reflectors. For each of the surfaces, we compute its parameter, such as curvature, principal axes and surface area. By examining the parameters, we can determine whether a surface belongs to a buried object or not. This sequence of processing is shown in Figure 78. The output of the processing sequence is a list of possible buried objects, along with their parameters. These objects still need to be migrated to their true location and orientation. The concept is similar with the migration in volume based processing algorithms, except the migration process is done in the object's parameter space, not the data space. So instead of migrating millions of voxels, we only need to migrate the objects' location and orientation. Another advantage of the surface based processing method are the extra information it can extract from the GPR data, such as object's type, pose and size. It is also able to process and find buried objects in the absence of an accurate propagation velocity estimate, because the detection process is based on the parameters of the surface shape and these parameters are not sensitive to error in the propagation velocity estimate.

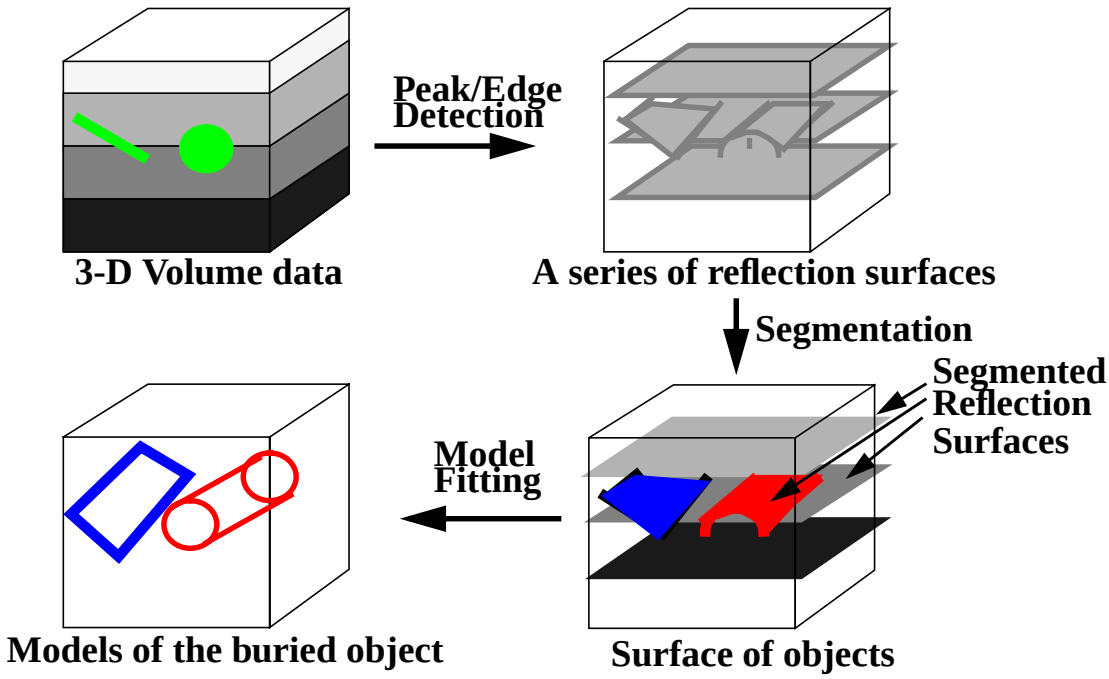

Figure 78. The sequence of processing required to convert buried objects embedded in the 3-D volume data into parametric object models

5.2. Preprocessing

Before we begin the segmentation process, we need to preprocess the raw data to a form that is suitable for the geometrical 3-D segmentation.

The first preprocessing are similar to the migration process, which is background noise subtraction. In addition to this, we need to use three additional preprocessing steps to reduce the 3-D volume data into a series of reflection points which could be segmented into 2.5-D surfaces. The 3 preprocessing steps are correlation of the received signal with the transmitted pulse, peak detection and range conversion, and reverberation elimination.

5.2.1. Correlation

In the returned signal, the beginning of the reflection is denoted by the zero crossing of the reflected pulse, instead of its peak. If we use the location of the peak to signify the beginning of the reflection, it will have an error of half the pulse length. On the other hand, the strength of the reflection is denoted by the amplitude of the peak. In order to get both the correct reflection's location and strength, we correlate the returned signal with the transmitted pulse. The location of the peak in the resulting signal signifies the beginning of the reflection and

the amplitude of the peak corresponds to the strength of the reflection. Figure 79 shows a vertical cross section of a 3-D GPR data before and after the correlation step.

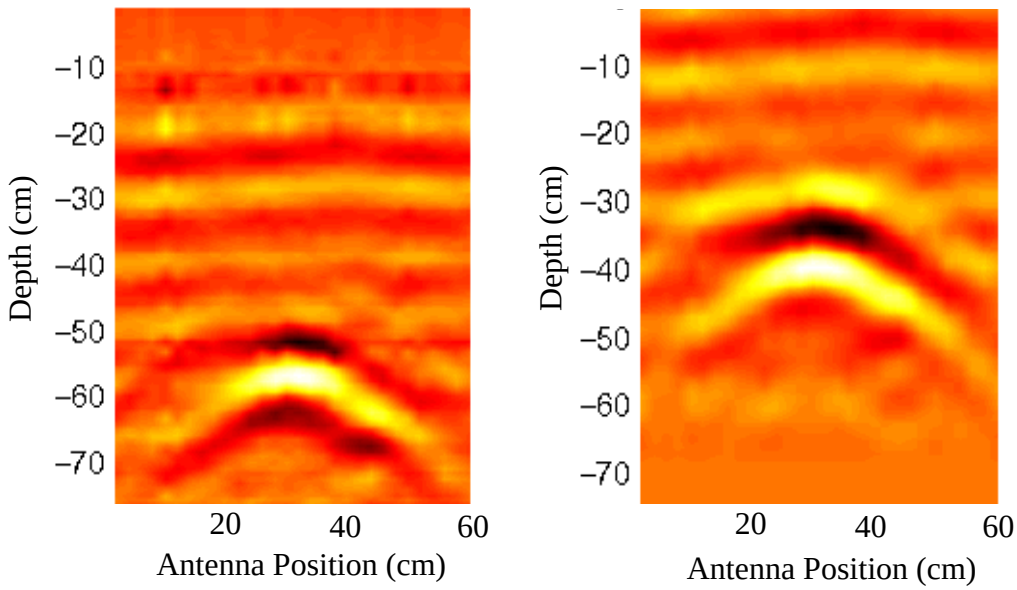

Figure 79. A vertical cross section of a GPR scan before (left) and after (right) correlation with the transmitted pulse.

In order for the correlation to be meaningful, the reflected pulse must retain a similar shape to the original transmitted pulse. This is only true if the pulse does not experience too much dispersion and other distortions as it travels through the soil. Therefore we only correlate the reflected pulse with the transmitted pulse when the reflection comes from objects that are located not too deep in the soil. If the buried objects are located very deep in the soil, the reflected pulse from the objects will experience a lot of distortion. In this case correlation is not meaningful nor needed, because at great depth, an error of half a pulse length becomes insignificant compared to the depth of the object. In other word, the relative error is very small.

As a result of this processing step, the location and amplitude of each peak in the individual GPR scan signifies the beginning and the strength of possible reflections from a buried object.

5.2.2. Peak detection and range conversion

Peak detection reduces the 3-D GPR volume data into a series of 2.5-D surfaces which is much more compact, but still contains most of the geometrical features of the 3-D data. It is important to remember that each column in the 3-D data is an individual GPR scan, which is

a discretized time varying signal. After the correlation process, peaks in this signal signifies possible reflections from buried objects. Since the signal is smoothly changing, most of the information is contained in the time offsets and amplitudes of the peaks in the signal. The peak detection process detects these peaks and reduces the signal to a series of time offsets and signed amplitudes of its positive and negative peaks.

Figure 80 shows a single GPR scan consisting of 500 sample points before and after the peak detection process. After peak detection, the data are reduced to about 12 pairs of numbers, which are the time offsets and signed amplitudes of the peaks. As we can see the amount of reduction in the data size is very significant.

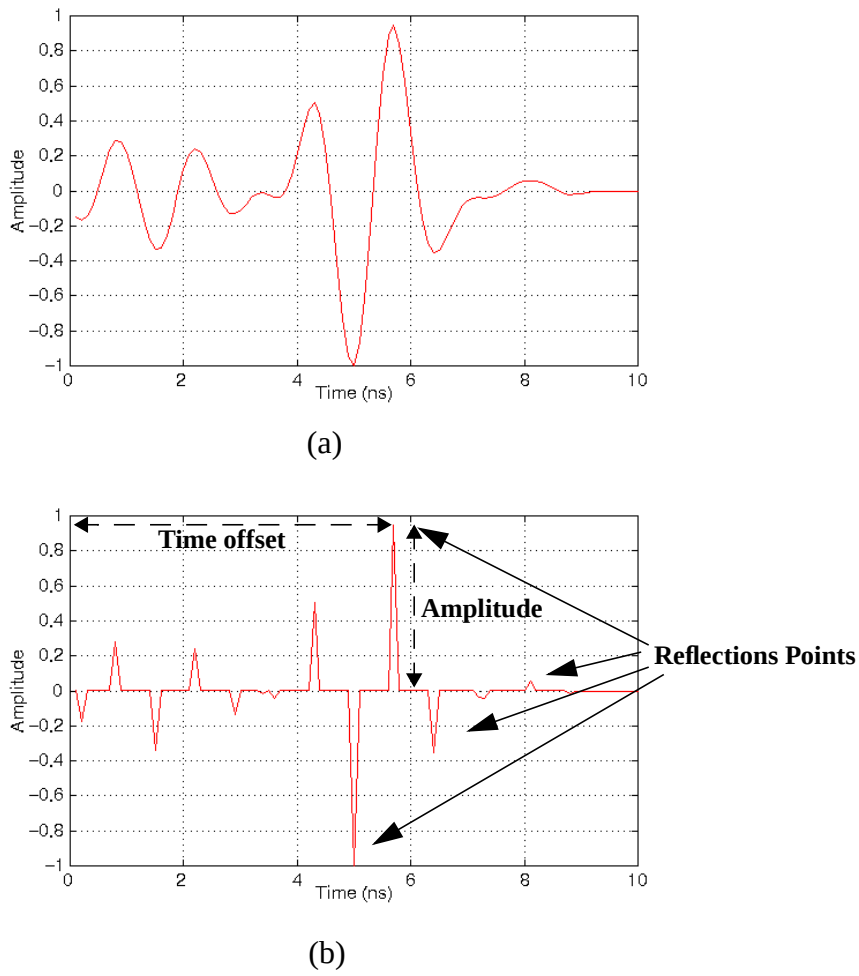

Figure 80. A single GPR scan before (a) and after peak detection (b)

After the peak detection process, we convert the time offset of each peak to distance by multiplying it with the estimated round trip propagation velocity of the GPR signal in the soil. As a result, the individual scan is represented by a list of locations and amplitudes of the peaks.

The result of these processing in 2-D can be seen in Figure 81. It shows a vertical cross section of a 3-D GPR data of a buried cylinder before and after peak detection. Before the peak detection, the vertical cross section is a 2-D image, but after the peak detection process, the vertical cross section is a series of 2-D curves. The figure shows that peak detection preserves the shape of the pipe's reflection profile.

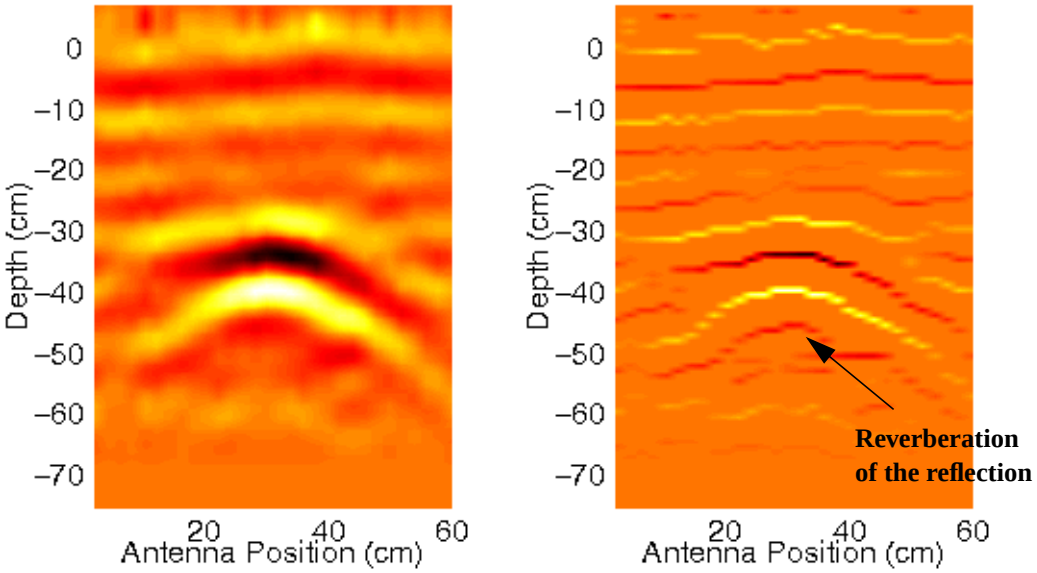

Figure 81. A vertical slice of GPR scan before (left) and after (right)peak detection

For the 3-D volume data, the peak detection and range conversion process reduce the volume data into a series of 2.5-D surfaces. The 2.5-D surfaces can be represented by a 2-D image where each pixel contains a list of the locations and amplitudes of the peaks in the reflected signal that is obtained at an antenna location. We can also think of each pixel as a list of distances to possible reflectors and their reflection strengths.

It is important to notice that each pixel do not only contains a list of the peaks' location, but also the peaks' signed amplitude. This is really important for the 3-D segmentation process, where we only group peaks which have similar sign.

5.2.3. Reverberation Elimination

One of the main problems in with GPR data processing is the large amount of noise in the data. This can be caused by random noise, static noise (unwanted background reflections) and reverberations or ringings. We have remove most of the background noise using back-

ground noise subtraction technique, which is described in section 4.2.. Now we will describe a method that we used to remove reverberations or ringings

In Figure 81, we see that even a single object can creates multiple reflections or reverberations. In order to reduce the reverberations, we further filter the output of the peak detection step. We compare each peak with the previous peak of the same sign. If a peak's magnitude is less than the previous peak's, then that peak is eliminated. We only compare a peak with a previous peak of the same sign, so a negative peak is compared in magnitude to the previous negative peak and a positive peak is compared in magnitude to the previous positive peak. This is shown for a single GPR scan in Figure 82. In this case, the number of peaks is reduced from 14 to 8.

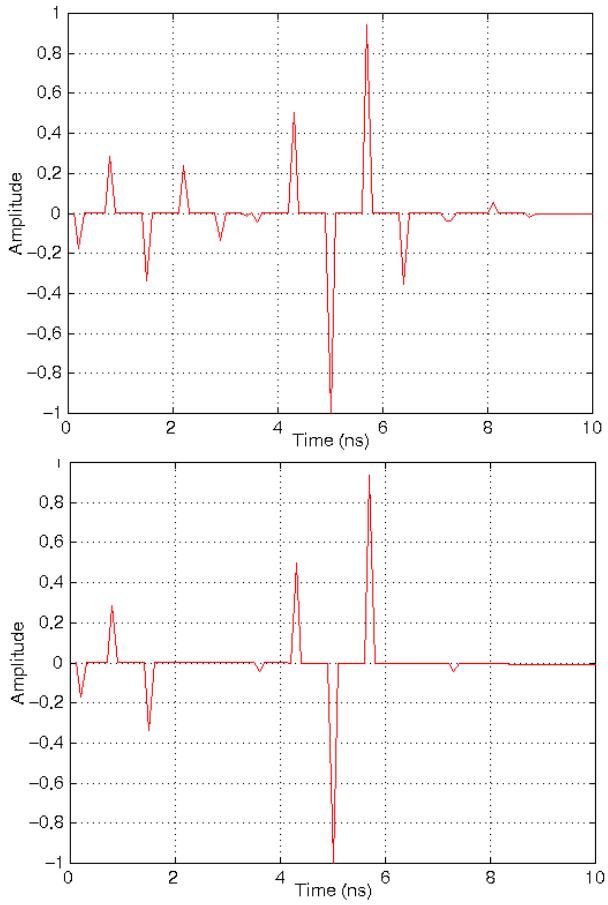

Figure 82. A graph of the peaks of a single GPR scan before (top) and after (bottom) reverberation elimination.

The effect can be seen clearer in Figure 83 where a vertical slice or cross section is shown before and after the reverberation elimination. The elimination process works so well because it does not depend on the absolute amplitude of the peaks. It only detects a relative

decrease in the magnitude. It is similar in concept to the edge detection process in computer vision.

There is a problem with the reverberation elimination when the objects are located very close to the surface, because the reflection from the air-ground interface is usually very strong, so it can mask the reflection from the buried object. This will happen when the reflected signal from the soil interface is only separated by less then a wavelength with the reflection from a buried object. The wavelength itself depends on the frequency of the antenna and the propagation velocity of this signal in the soil. For a 1 GHz antenna and a propagation velocity of 8 cm/ns, half the wavelength is 4 cm.

One solution for this problem is to place the antenna very close to the interface or if possible touching the surface, thereby eliminating the air-ground interface's reflections. This is not always possible when the ground surface is uneven or rough. In this case it will be necessary to modify the reverberation elimination process so it takes into accounts the reflections from the interface and ignores them during the reverberation elimination process.

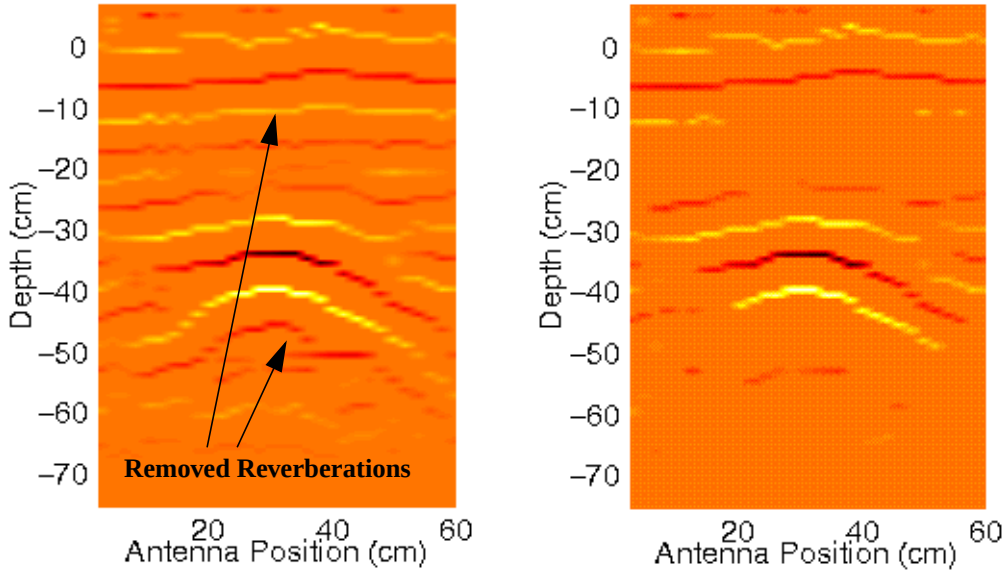

Figure 83. A vertical slice of GPR scan before (left) and after (right) reverberation

5.3. 3-D Segmentation

After peak detection and reverberation elimination, the GPR 3-D volume data are reduced to a 2-D image where each pixel contains a list of distances to reflectors and their reflection strengths. The distances and the reflection strengths are obtained from the locations and signed amplitudes of the peaks in the individual GPR scan. Another way to look at the data is to think of it as a 2-D image, where each pixel contains multiple range and intensity values. This is in contrast to a traditional 2-D image, where each pixel contains only a single range or intensity value.

A traditional 2-D image is usually produced by a sensor which only generate a single value for each pixel. If the sensor is a video camera, each pixel contains an intensity or color value. If the sensor is a scanning laser rangefinder, each pixel contains a range to an object. In the case of GPR, we must move the antenna in a raster scan pattern to generate an image. Due to the wide beamwidth of the antenna and the ability of GPR signal to penetrate solid objects, multiple objects can be detected at each scan position. One way to represent these data is by using 3-D volume data representation where each vertical column is a single scan. After the peak detection process, we can also represent the same data as a 2-D image where each pixel contains a list of peaks that represent distances to the reflectors and their reflection strengths.

The reflectors can be the air-ground interface, soil's inhomogeneity or buried objects. In order to be able to detect, locate and categorize the reflectors, we need to group these reflectors into several 2.5-D surfaces. Each surface corresponds to a reflector's surface that causes the reflections. We can compute the characteristics of each surface, then based on them, we can determine whether the surface belongs to a buried object. If it does, then we compute the parameters of the objects.

The process of grouping the reflectors into different surfaces is called 3-D segmentation. The method that we use is similar in concept with the range image segmentation algorithm developed by Faugeras and Hebert [Faugeras 86]. Our algorithm is more complicated because in our case each pixel in the image can have multiple range and intensity values.

We begin the 3-D segmentation process with a series of location and signed amplitude of the peaks in the GPR scans. Each of these peaks constitutes a reflection point (look at Figure 80). Combined with the antenna positions, the reflection points contain all the necessary geometric information to reconstruct the 3-D GPR data. Let us call the array which contains the ranges to these reflection points $R(i,j,k)$, where the i and j is the index to the x and y location of the antenna where the reflection is detected, and k denotes the kth reflection obtained at that antenna location. Associated with the array $R(i,j,k)$ is $A(i,j,k)$ which contains the signed amplitude information of the reflection points. The x and y coordinates of each

R(i,j,k) are stored in the array *X(i,jk)* and *Y(i,j,k)*. The *z* coordinate is the depth value which is -*R(i,j,k)*.

The detailed steps of the segmentation is shown in Figure 84. In the beginning, each reflection point is a region with only a single point, so we start with as many regions as points. We then build a mesh connecting these regions. Each region is connected to another region if any point in one region is connected to another point in other region. There is a constraint

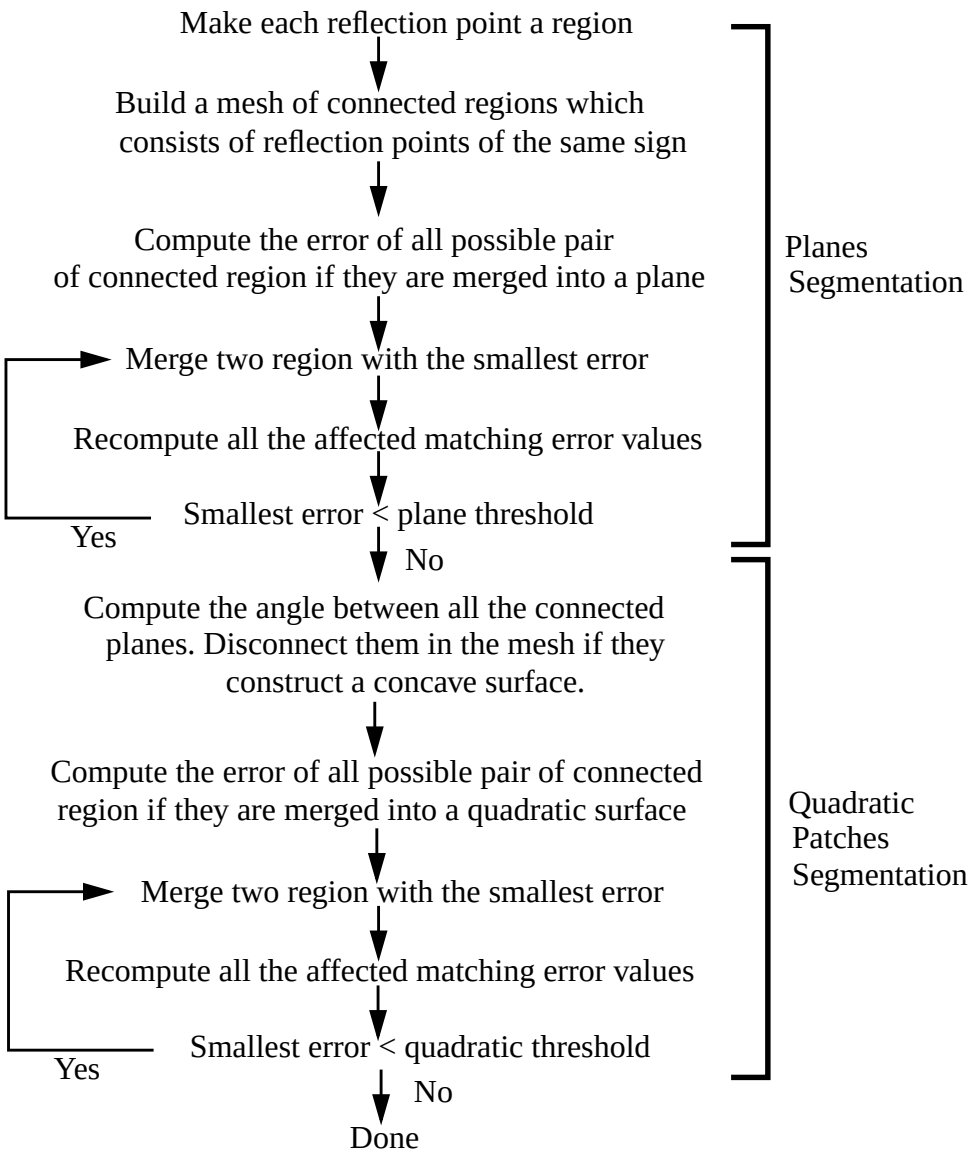

Figure 84. The steps in the segmentation of 3- D GPR data.

that the two connected points must have *R(i,j,k)* values that are within a certain distance threshold and whose *A(i,j,k)* values have the same sign. The distance threshold is set to the maximum difference between two adjacent range reading assuming that the reflections come from the same point reflector. Two neighboring range readings $R(i_1,j_1,k_1)$ and $R(i_2,j_2,k_2)$ will have the largest difference when the reflections are coming from a point reflector that is located at the edge of the beamwidth. This is illustrated in Figure 85.

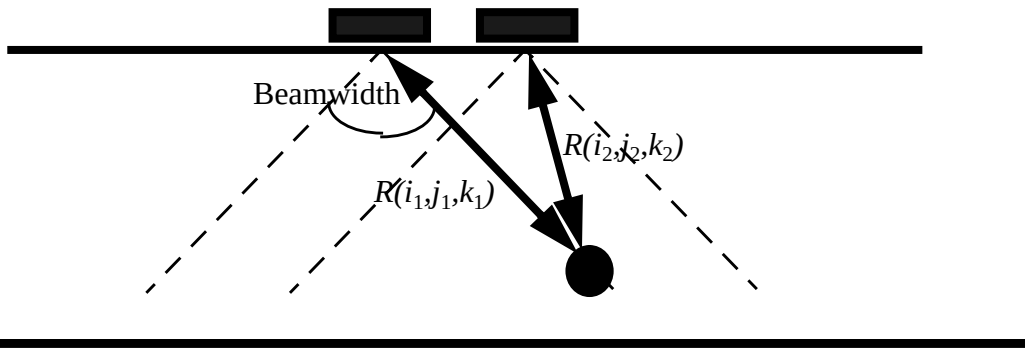

Figure 85. Two range readings to the same buried object that is located at the edge of the antenna beamwidth.

The difference in $R(i_1,j_1,k_1)$ and $R(i_2,j_2,k_2)$ also grows larger as the object is located farther away from the two antenna. The maximum difference happens when the reflector is located infinitely far away from the two antennas. The maximum distant is:

$$Distant_{max} = \lim_{R(i,\,j,\,k) \to \infty} abs(R(i_1, j_1, k_1) - R(i_2, j_2, k_2)) \tag{18}$$

$$\text{where: } abs(i_1 - i_2) + abs(j_1 - j_2) = 1$$

$$Distant_{max} = dh \cdot \cos\left(\frac{beamwidth}{2.0}\right) \tag{19}$$

We set the distant threshold to $Distant_{max}$. If two adjacent reflections are separated by a distance which is larger than $Distant_{max}$, the reflections must come from two different objects.

Once we build the mesh which connects neighboring regions, we can begin merging them into planes. This is the plane segmentation phase. At each iteration, for every possible pairs of connected regions, we compute the matching error. The pair of regions which have the

smallest error are merged together. To compute the matching error, we fit the points in both regions to a plane using least square error fitting technique and compute the maximum absolute deviation of the points from the plane.

After combining a pair of regions, we compute the new match error values between the new region and all its connected neighbouring regions. Then we repeat the whole process until the smallest match error is larger than a threshold.

At the beginning of the merging process, each region contains only a single point, so it is possible to merge two regions which have a total of two point. In this case and in the case where the combined region has less than 3 non-coaxial points, we fit the points to a line instead of a plane to avoid ill defined computation. We also have another constraint during the merging process. The planes can not have an elevation angle larger than half of the beamwidth because if a reflector is tilted at a larger angle, it will not be detected.

When the smallest match error is larger than a preset threshold, the plane segmentation phase ends and the quadratic patch segmentation phase begins. At this point, the reflection points have been grouped into planes. In the next phase, the planes are combines into quadratic patches. We chose quadratic patches because it models the reflection profile from a flat, cylindrical and spherical reflector pretty accurately. We begin this phase by examining all the connected planes in the mesh. If a pair of planes can be combined to make a concave surface, we disconnect the pair in the mesh so that the planes can not be directly merged together. This is to prevent the construction of concave quadratic patches. Once this is done, we begin to compute the match errors of combining every pair of connected regions into quadratic patches. As in the planes segmentation phase, we use a least square fitting technique to fit the points in the two regions into a quadratic patch. The match error measurement is the maximum absolute deviation of the points from the quadratic surface. We chose the pair of regions that have the smallest match error and combined them into a quadratic patch. We then recompute all the affected match errors between the new region and its neighboring regions. We repeat this process until the smallest match error is larger than a preset threshold.

At the end of the quadratic patch segmentation phase, we have grouped the reflection points into a collection of quadratic patches. Each of these patch should represent a single surface of a buried object.

5.4. Surface Type Categorization

The segmentation process produces a series of quadratic surfaces. In order to determine whether any of these quadratic patches belongs to a buried object, we need examine the type of the quadratic surfaces. The surfaces can be flat, cylindrical or spherical. If a quadratic sur-

face has a flat shape, then it is obvious that it is produced by a flat reflectors. If it does not have a flat shape, the quadratic surface might be caused by a flat, cylindrical or spherical reflectors. One might ask how a flat reflector can produces a non-flat surface. The reason for this is that the edges of a flat reflector also reflects some energy and it produces a reflection surface that is similar to the one produced by a cylindrical or spherical reflector. This effect is illustrated in Figure 86.

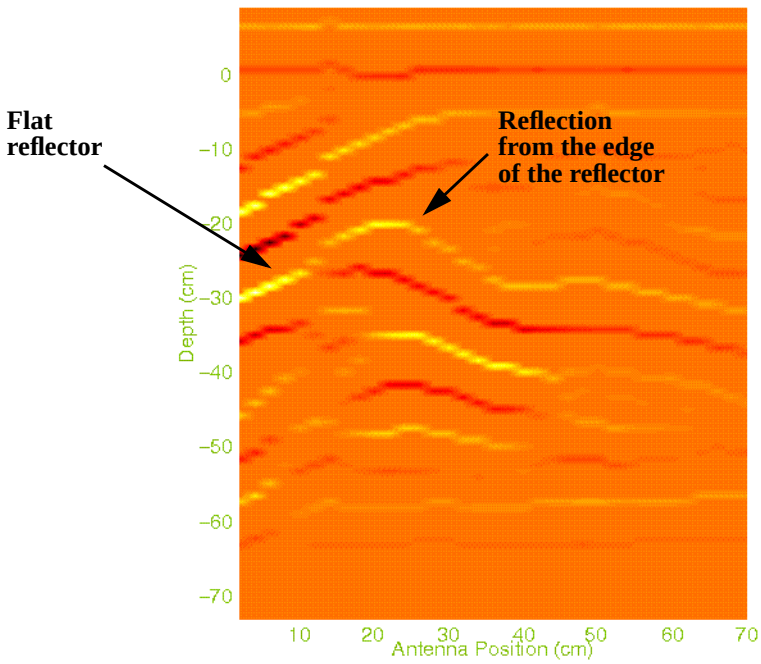

Figure 86. The reflection profile from the edge of a flat reflector.

To be able to differentiate the reflector's type despite the above effect, we use the orientation histogram of each quadratic patch. We compute the orientation histogram by computing the orientation of each point on the quadratic patches. We compute the orientation of each point by fitting a plane at that point using a collection of neighbouring points. For each point on the quadratic patch we fit a plane to its local neighborhood and compute the following plane equation:

$$z = ax + by + c \tag{20}$$

Since the beamwidth of the antenna is about 45 degrees, a and b should be less than 1, so least square fitting technique should work fine. We then discretize the value of a and b. We

add the amplitude of the reflection point to the orientation histogram value at *a* and *b*. So the orientation histogram is weighted by the amplitude of the reflection point.

From the resulting orientation histogram, the type of the reflector can be robustly determined. For a flat reflector, the strongest reflections comes from the surface of the reflector itself, not from the edge, and all the points on the surface of the reflector have more or less the same surface normal. So the resulting orientation histogram has a very small area with most of the orientations pointed in the direction of the surface normal of the flat reflector. This can be seen in Figure 87.

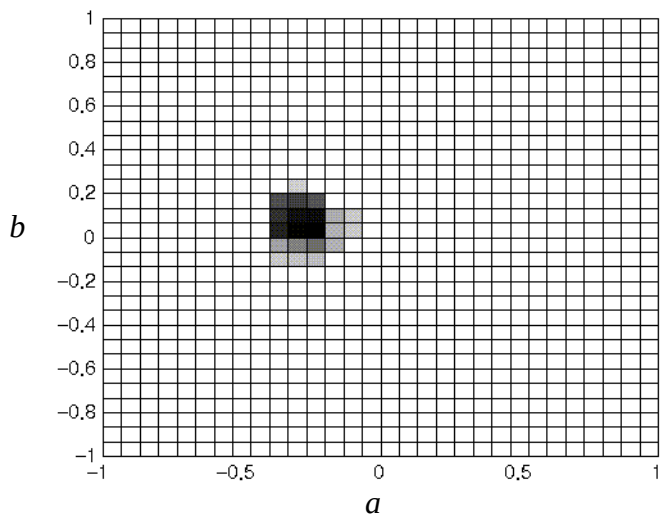

Figure 87: Orientation histogram of a flat reflector

For a spherical reflector, the orientations in the histogram are more evenly distributed since each reflection is obtained from a surface with a different normal vector. An example of an orientation histogram for a spherical reflector is shown in Figure 88. For a cylindrical reflector, the orientations are spread out in only one direction as shown in Figure 89.The direction of the elongation gives us the azimuth angle of the cylindrical reflector.

Mathematically, we determine the type of the reflector by computing the width and the length of the orientation distribution in the histogram. If the ratio of length to width is larger than a threshold, then the reflector is cylindrical. If the ratio is smaller than the threshold, we need to look at the distribution of the weighted votes in the orientation histogram. If the votes is concentrated in a small area around the centroid of the votes, then the type of the reflector is flat. If the votes are evenly spread out then the type of the reflector is spherical.

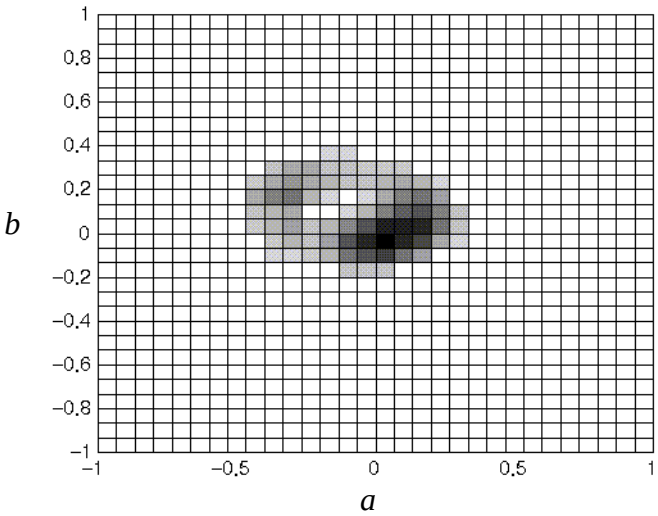

Figure 88: Orientation histogram of a spherical reflector

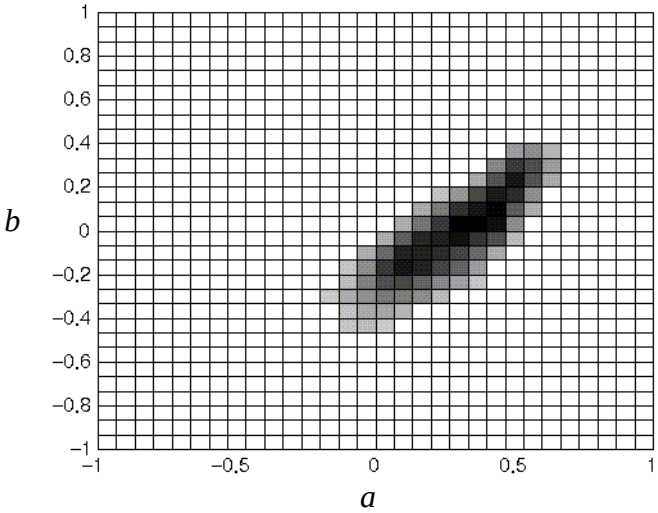

Figure 89: Orientation histogram of a cylindrical reflector

5.5. Parameter Estimation

Once the type of the reflector is determined, we need to compute the relevant object's parameters. The methods that we use to compute the parameters for each type of reflector are different. We will discuss the methods that we use to compute these parameters in the following sections.

5.5.1. Parameters Computation for a Flat Reflector

We compute the centroid of the reflector from the centroid of all the voxels on the reflection profile. The width, length, azimuth and elevation are computed using the information from principal component analysis. The length is computed as the extent of the surface in the direction of the eigenvector with the largest eigenvalue. The width is computed as the extent of the surface in the direction of the eigenvector with the second largest eigenvalue.Azimuth angle is the horizontal angle of the eigenvector with the largest eigenvalue, while elevation angle is computed as the vertical angle of the eigenvector with the largest eigenvalue.

5.5.2. Parameters Computation for a Cylindrical Reflector

In order to compute the parameters for a cylindrical reflector, we need to locate the cylinder axis first. We start by obtaining a rough estimate of the cylinder orientation from the orientation histogram. Then we generate the profile lines on the cylinder by sampling the surface along the direction perpendicular to the cylinder orientation as shown on Figure 90. The radius of the cylinder can be computed from these profile lines as shown in Figure 91 and Equation (21).

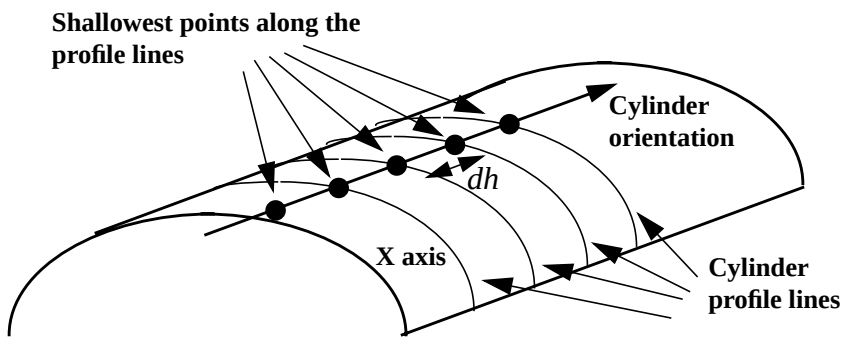

Figure 90: Resampling the cylindrical surface to locate the cylinder axis.

To locate the cylinder axis we locate the shallowest points along these profile lines. These points are fitted into a line. This line is then translated downward by a distance which is equal to the computed radius. The translated line is the cylinder axis. The azimuth and the elevation angles can be computed the orientation the cylinder axis.

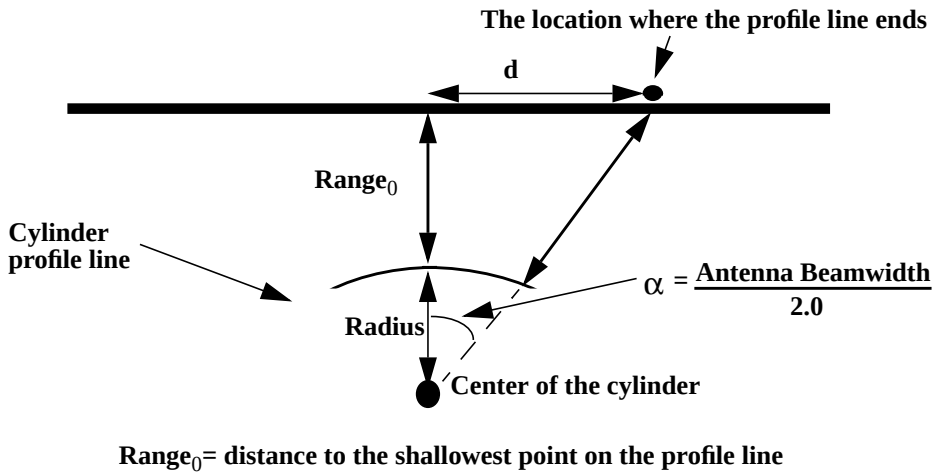

Figure 91: Computing the radius of a cylinder from the profile line

$$radius = \frac{d}{\tan(\alpha)} - range_0 \tag{21}$$

5.5.3. Parameters Computation for a Spherical Reflector

We use the peak of the quadratic patches to determine the location of the centroid. The radius is computed using the same method as in the case of cylindrical reflectors. Since the reflection surface resulting from a spherical reflector is more or less symmetrical, the direction of the profile lines that is used in the radius computation is not important.

5.6. Parametric Migration

Once the buried object's location, orientation and size are computed, these parameter must be migrated to their true value, for the same reason that migration is necessary to correct the various imaging effects in GPR data. The difference here is that rather than migrating the raw 3-D volume data, we only need to migrate the parameters that characterize the object location, orientation and size.This difference causes a very significant reduction in the number of necessary computation. The migration operation is an O(cN) algorithm where N is the number of voxels and c is a very large constant, while the parametric migration is only an O(N) algorithm where N is the number of parameters. While the number of voxels can eas-

ily exceeds several millions, the number of parameter for each object is at most ten or twenty.

The parametric migration process needs to be done for reflections that are reflected from surfaces that are not horizontal. These include non-horizontal flat surfaces and cylinders that are oriented diagonally. Spherical or point reflectors do not need to be migrated since the centroid of the detected object is already at the correct location. For a non-horizontal flat surface we migrate each vertex using Equation (15), Equation (16) and Equation (17). For a non-horizontal cylinder, we need to migrate the two vertices at the end of the cylinder axis using the same set of equation. The elevation angle for the non-horizontal planes and cylinders must also be migrated using Equation (14).

5.7. Propagation Velocity Computation

Once the object is detected and located, an automated excavator can remove a layer of soil that is known to be devoid of any buried object. After the soil is removed, we can detect and locate the same object again. From the difference in the depth of the object, we can compute a new estimate for the GPR signal propagation velocity in the soil. It is computed using this simple equation:

$$Velocity = \frac{Thickness\ of\ the\ removed\ soil}{T_0 - T_1} \tag{22}$$

where T_0 is the propagation time of the reflected signal from the shallowest point on the object before the removal of soil and T_1 is the propagation time after the removal of soil. We can repeatedly do this as the excavator remove additional layers of soil above the object. If the soil is heterogeneous, we can compute the trend in the propagation velocity versus the depth.

An example of propagation velocity computation is shown in Figure 92. A cylindrical object is detected under the soil and then a 6cm layer of soil is removed. Another scan is done and the object is detected again. Then we remove another 5cm layer of soil before doing another scan. Using Equation (22) we compute the propagation velocity. The result is shown in Table 16 .

The original depth of the object is 25cm. It can be seen from Table 16 that the velocity estimates are quite accurate, resulting in accurate depth estimates of the object. This is partly due to the orientation of the buried object. Since the buried object is a horizontal cylinder it is very easy to locate the same point on the surface of the object for the propagation time comparison.

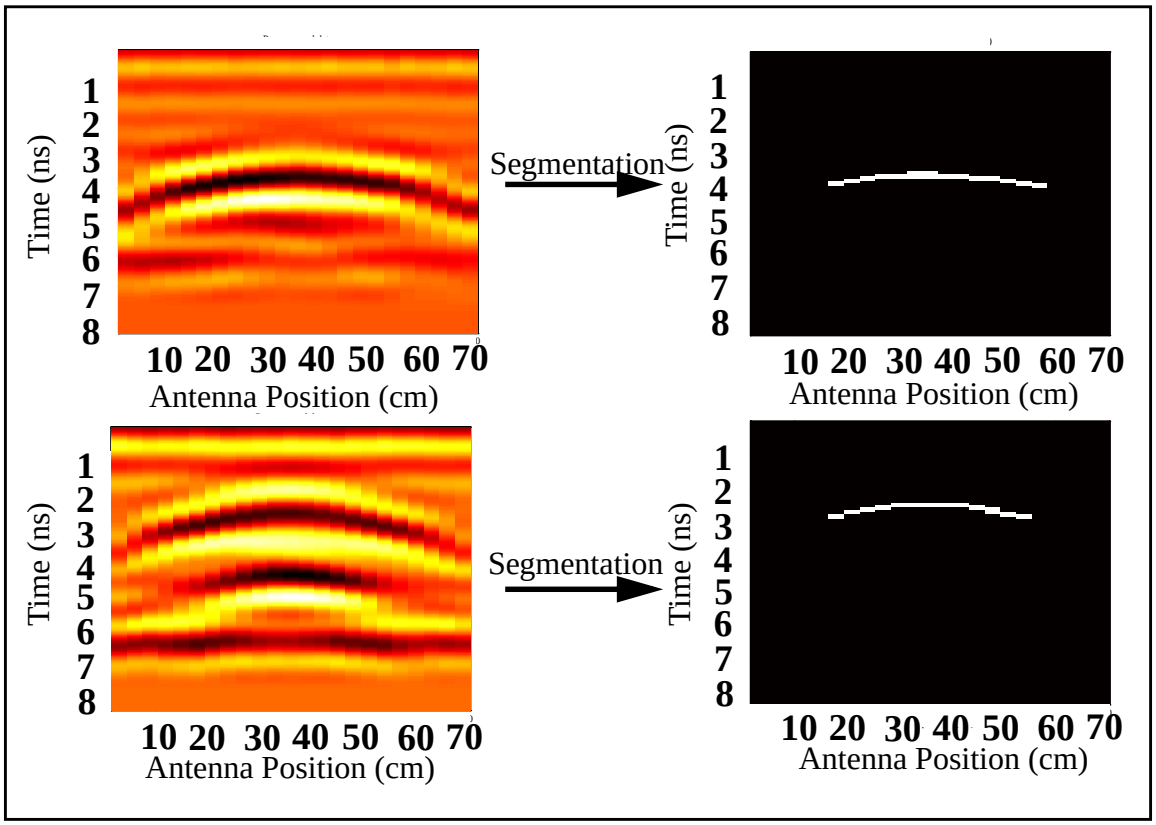

Figure 92. The segmentation result of a buried cylinder (top) and the same object (bottom) after a 11 cm layer of soil is removed.

	Computed velocity (cm/ns)	Computed depth (cm)	Real depth (cm)	Error (%)
After removal of 6 cm layer of soil	9.52	18.37	19.0	3.3
After removal of 11 cm layer of soil	9.62	13.49	14.0	3.6

Table 16. The computation of propagation velocity after removal of two layers of soil.

5.8. Limitation of the Surface Based Processing

The most important limitation of any geometrical features based mapping of buried object is the fact that they can not differentiate the material of the mapped object. So if there is a rock with the same dimension and shape as a landmine, it would be impossible to differentiate them just based on their geometrical feature. Since we are mainly interested in man made

buried objects, this limitation is not too critical. Most man made structures have a distinctive geometrical features that can be used to differentiate them from natural objects. For example a gas pipe might look the same as a buried tree trunk, but we can differentiate them based on their radius and length. We therefore believe that for many applications in buried object mapping and retrieval, this limitation is a minor disadvantage compare to the advantages of having automated subsurface mapping capability.

5.9. Analysis and Result

In this section we will show the result of the 3-D segmentation, parameter estimation, categorization and parameter migration. During the experiment we mainly use three types of buried object. The first one is a point or spherical reflector, such as a metallic ball or a very small object. The second one is a cylindrical reflector, such as a pipe or a barrel. The third one is a flat reflector, such as a plate or objects with flat surfaces. We tried using non-metallic material as well as the metallic one, to see whether our processing techniques can map non-metallic buried objects as well as metallic buried objects. We also buried an inert version of an actual non-metallic anti-tank mine and use our algorithm to find it in the GPR data.

In addition of varying the type of the objects used in the experiment, we also varied the number of objects buried. First, we will show the result of imaging a single buried object of various type, then we will show the result of imaging multiple buried objects. The latter is definitely more challenging because there is inter-reflection between the objects and one objects may obscure part of the other object.

For each set of result, we will show the computed object model and we will also show a table comparing the computed parameter to the actual parameter. It is important to note that although great care is taken to measure the actual location and orientation of the object, it has some amount of error. Our estimate is that the actual location measurements of the buried objects location are accurate to $\pm$ 1.0 cm, while the orientation measurements are accurate $\pm$ 2.0 degrees. For each detected object, the computed model of the object is shown along with the projection of the object into the horizontal plane. The horizontal projection makes it easier to determine the horizontal location of the object when looking at the 3-D image of the buried object model. In some cases, in addition to the actual buried object, some erroneous ghost objects are detected. These ghost objects are created by noise or by the reverberation of the reflections from actual buried objects. The reverberation elimination process eliminates most but not all the reverberations. In almost every case, the ghost object is located under the actual object. So we advise that any detected object that is located under another detected object might just be a reverberation of the topmost object.

The GPR subsurface data that is used in this section are obtained using a 1 GHz antenna. The antenna is scanned at a scanning interval of 2cm or 3cm and it is placed very close to the surface of the soil. The full scale of the GPR system is set at 10ns and the data are sampled at 600 sample/scan, which results in the range resolution of 0.0167ns multiplied by the propagation velocity.

The number of computed parameters varies with the type of reflectors. Some of the computed parameters, such as elevation angle is only valid for a flat or cylindrical reflector. Table 17 shows the computed parameter and their applicability for each type of reflector.

	Flat Reflector	Cylindrical Reflector	Spherical Reflector
Horizontal location of the object's centroid	Yes	Yes	Yes
Depth of the object's centroid	Yes	Yes	Yes
Radius	No	Yes	Yes
Length	Yes	Yes	No
Width	Yes	No	No
Elevation angle	Yes	Yes	No
Azimuth angle	Yes	Yes	No
Depth of the reflecting surface at the centroid	No	Yes	Yes

Table 17: The detected reflectors and their computed parameters.

The depth of the reflecting surface at the centroid is a more informative depth information compare to the depth of the centroid. This is the depth of the object's surface at the horizontal location of the centroid. The computed depth of the object's centroid depends on the computed radius. So if there is an error in the radius estimate, the error propagates to the depth of the centroid. So a more useful and reliable depth information is the depth of the reflecting surface at the centroid. The depth of the reflecting surface does not take into account the object's radius or thickness so any computed error in the radius does not affect the computed depth of the object's surface. Figure 93 illustrates the relationship between these computed parameters.

For a cylindrical reflector, the relationship between the centroid depth, radius and depth of the reflecting surface at the centroid is the following:

$$\textit{centroid depth} = \textit{depth of the reflecting surface} - \frac{radius}{\cos(elevation)} \tag{23}$$

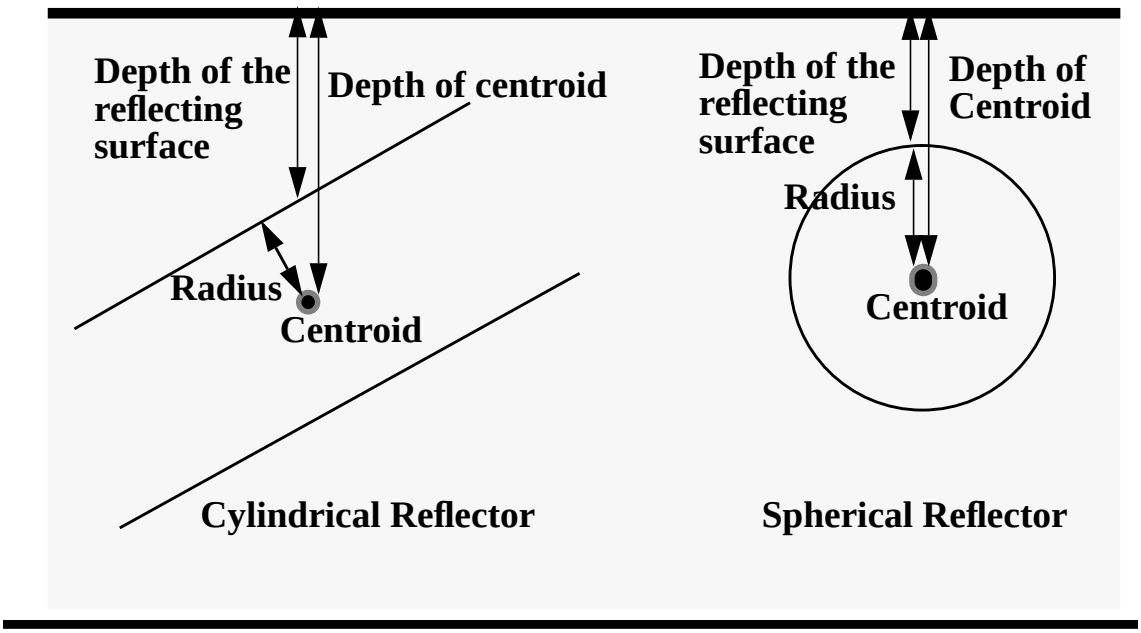

Figure 93. The relationship of the depth of the centroid, radius and the depth of the surface at the horizontal location of the centroid.

In the case of a spherical reflector which does not have any elevation angle estimate, the following equation applies:

$$\textit{centroid depth} = \textit{depth of the reflecting surface} - radius \tag{24}$$

As can be seen from the above equations, it is possible to have a large error in the computed radius and centroid depth, but only a small error in the computed depth of the object surface. The reason for this is the computed depth of the object surface mainly depends on the accuracy of the propagation velocity estimate, while the radius and the centroid depth is calculated using the information obtained from the shape of the reflection surface.

5.9.1. An aluminum foil covered tennis ball

We start by trying to map a single metallic object consisting of an aluminum covered tennis ball. The result of the mapping is shown in Figure 94, while the actual and computed parameters are shown in Table 18. The location of the ball and its radius are computed pretty accurately. In addition to the detected ball, the segmentation process also detects a ghost object. This false positive error is likely to be caused by reverberation from the reflections of the real objects. The reverberation elimination process eliminates many echos, but not all of them.

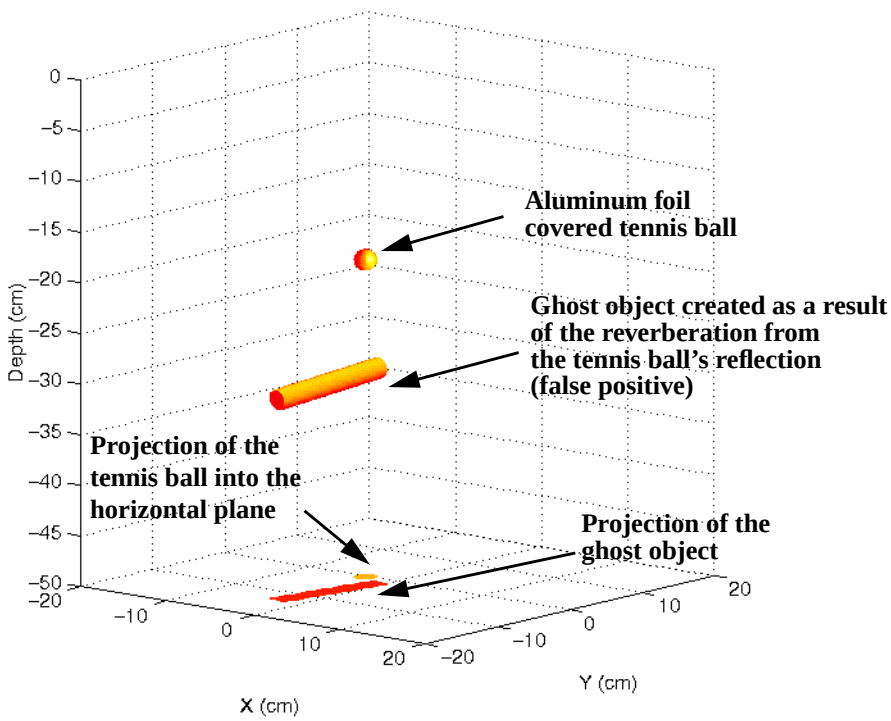

Figure 94. An aluminum foil covered tennis ball.

Parameter	Real Value	Computed Value	Error
Centroid horizontal position (cm)	(0.0,0.0)	(-1.34,0.45)	1.41
Centroid Depth (cm)	25	27.53	2.53
Radius (cm)	3.0	3.53	0.53
Depth of the reflecting surface at the centroid (cm)	22	24.0	2.0

Table 18. Parameter for the aluminum covered tennis ball in Figure 94.

5.9.2. A horizontal metallic pipe

In Figure 95 we show a result of mapping a metallic pipe (length= 30cm, radius= 6cm). The pipe is detected correctly and all the computed parameters are pretty close to the actual parameters.

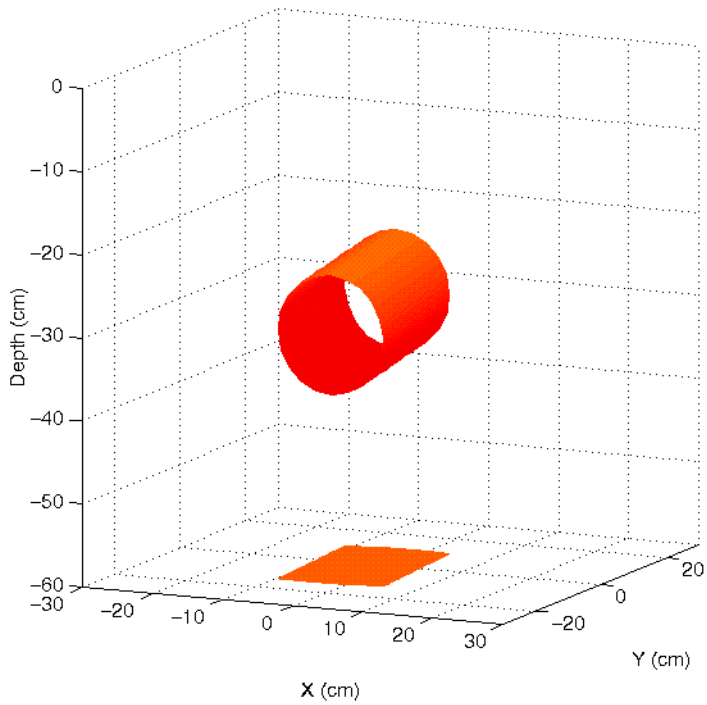

Figure 95. A metallic pipe (length= 30cm, radius= 6cm).

Parameter	Real Value	Computed Value	Error
Centroid horizontal position (cm)	(0.0,-3.5)	(-3.09,-1.59)	3.66
Centroid Depth (cm)	31.0	30.41	-0.59
Radius (cm)	6.0	8.28	2.28
Length (cm)	30	27.93	-2.07
Azimuth (degrees)	90.0	93.97	3.97
Elevation (degrees)	0.0	3.67	3.67
Depth of the reflecting surface at the centroid (cm)	25.0	22.13	-2.88

Table 19. Parameter for the metallic pipe in Figure 95.

5.9.3. A non-horizontal metallic pipe

Now that we have shown that we can map a metallic ball and a pipe, we would like to illustrate the capability of the mapping system to determine the orientation of a non horizontal pipe as well. We buried a pipe (length= 60cm, radius= 3cm) at an elevation angle of 16 degrees from the horizontal plane and azimuth angle of 36.25 degrees from the X axis. Figure 96 shows the 3-D model of the computed object and Table 20 shows the computed parameters. All the parameters are computed close to their actual values.

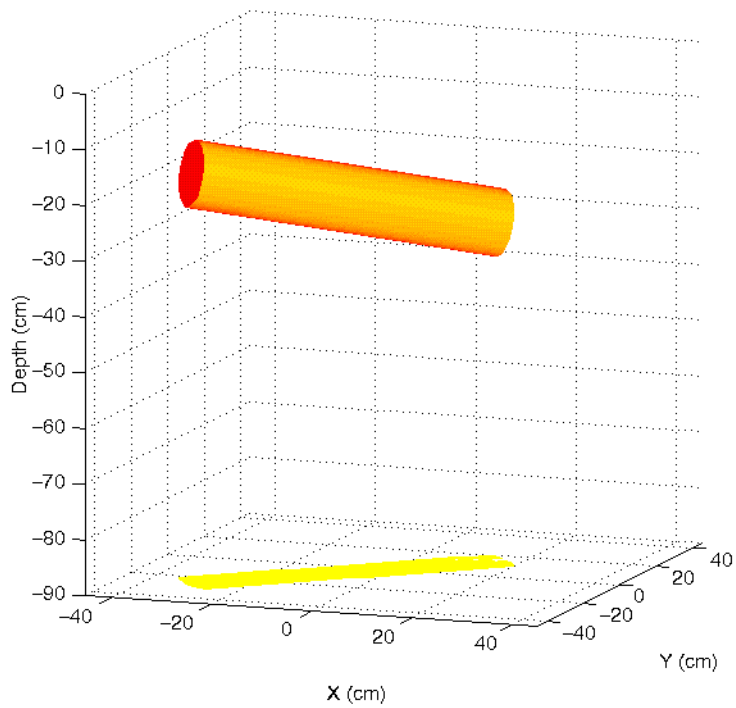

Figure 96. A metallic pipe (length=60cm, radius = 3cm)

Parameter	Real Value	Computed Value	Error
Centroid horizontal position (cm)	(-9.0,-2.5)	(-7.58,-5.36)	3.19
Centroid Depth (cm)	24.0	23.63	-0.37
Radius (cm)	3.0	6.75	3.75
Length (cm)	60.0	67.20	7.20
Azimuth (degrees)	36.25	41.31	5.06
Elevation (degrees)	-16.0	-13.47	2.52
Depth of the reflecting surface at the centroid (cm)	21.0	16.88	-4.12

Table 20. Parameter for the metallic pipe in Figure 96.

5.9.4. A horizontal pvc pipe

We have shown that we can map metallic objects pretty well. In Figure 97 we show the result of mapping a non-metallic object, a pvc pipe to be precise. Only a section of the pvc pipe is mapped, so we can't compare the computed length of the pipe and its actual length. As in the case of the metallic pipe, the computed parameters for the plastic pipe is pretty close to the actual values.

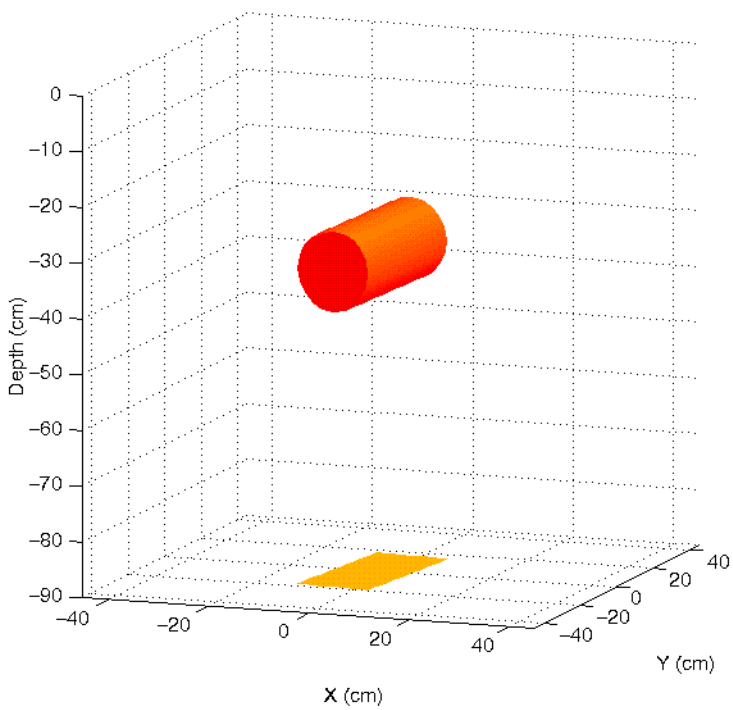

Figure 97. A pvc pipe (length= 120cm, radius= 7.5cm). Only the middle section of the pipe is mapped.

Parameter	Real Value	Computed Value	Error
Centroid horizontal position (cm)	(0.0,0.0)	(-3.24,-0.48)	3.28
Centroid Depth (cm)	27	24.53	-2.47
Radius (cm)	7.5	3.91	-3.59
Azimuth (degrees)	90.0	84.17	-5.83
Elevation (degrees)	0.0	1.40	1.40
Depth of the reflecting surface at the centroid (cm)	22	20.62	-1.38

Table 21. Parameter for the pvc pipe in Figure 97.

5.9.5. A metallic coffee cup and a metallic plate

So far we have shown mapping of a single buried objects. Since resolving multiple buried objects is one of the purpose of GPR data processing, we show in Figure 98 the result of mapping a metallic coffee cup and a metallic plate. Since the metallic plate is quite small (11cm by 11cm), it is detected as a point or spherical reflector. The positions and orientations of both objects are computed accurately. The computed dimensions for both objects are overestimated.

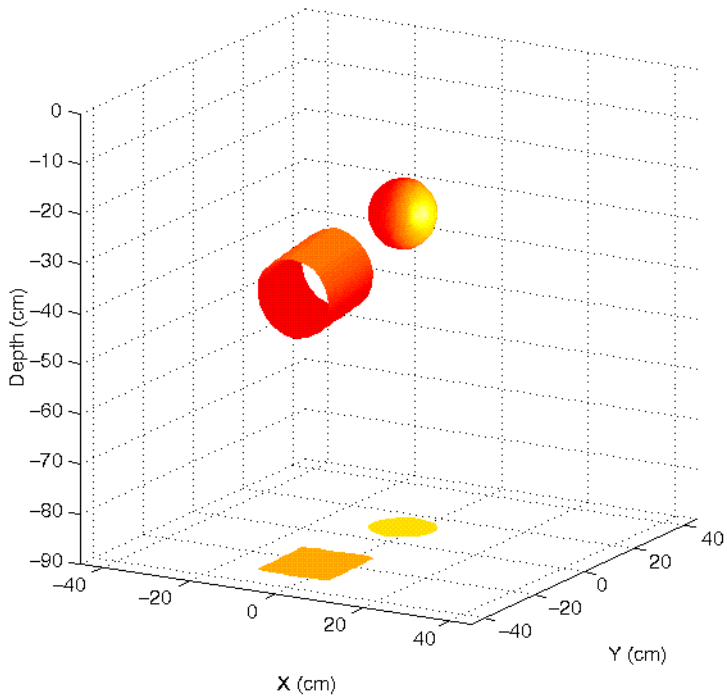

Figure 98. A metallic coffee cup and a small metallic plate.

Parameter	Real Value	Computed Value	Error
Centroid horizontal position (cm)	(0.0,-18.0)	(-4.53,-21.4)	5.66
Centroid Depth (cm)	28.0	35.50	7.50
Radius (cm)	5.0	9.62	4.62
Length (cm)	15	22.01	7.01
Azimuth (degrees)	90.0	91.28	1.28
Elevation (degrees)	0.0	5.8	5.8
Depth of the reflecting surface at the centroid (cm)	24.0	25.88	1.88

Table 22. Parameter for the metallic coffee cup

Parameter	Real Value	Computed Value	Error
Centroid horizontal position (cm)	(0.0,7.0)	(-2.99,10.77)	4.8‘
Centroid Depth (cm)	25.0	33.07	-2.71
Radius (cm)	5.5	8.32	2.82
Depth of the reflecting surface at the centroid (cm)	25.0	24.75	-0.25

Table 23. Parameter for the metallic plate. It is detected as a point/ spherical reflector.

5.9.6. A metallic plate and a metallic pipe

Figure 99 shows another result of mapping multiple objects. In this case they are a metallic plate and a metallic pipe. Both of the objects are detected correctly. The pipe’s radius and length have significant relative errors. We observe that the accuracy of the computed object parameters decrease when we have multiple objects located close to each other. We think that this is due to interferences between the reflections from the neighboring objects.

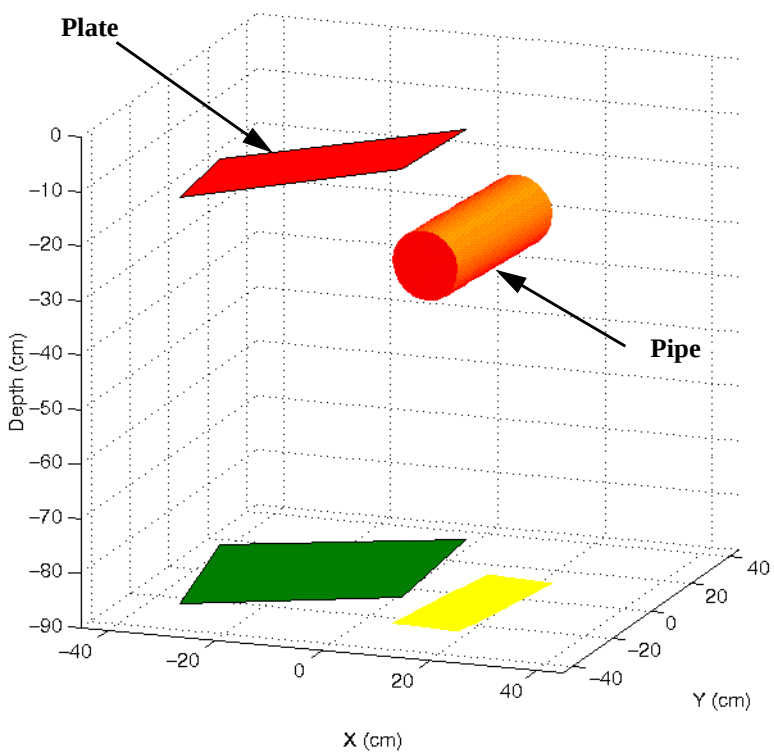

Figure 99. A metallic plate with a metallic cylinder

Parameter	Real Value	Computed Value	Error
Centroid horizontal position (cm)	(-17.0,0.0)	(-19.33,5.34)	5.83
Centroid Depth (cm)	21.0	-15.05	-5.95
Width (cm)	40.0	24.08	-15.92
Length (cm)	40.0	38.78	-1.22
Azimuth (degrees)	NA	154.72	NA
Elevation (degrees)	30.96	9.62	-21.34

Table 24. Parameter for the metallic plate.

Parameter	Real Value	Computed Value	Error
Centroid horizontal position (cm)	(18.0,0.0)	(13.67,-5.58)	7.06
Centroid Depth (cm)	18	22.91	4.91
Radius (cm)	0.5	6.44	5.9
Length (cm)	50.0	36.39	-13.61
Azimuth (degrees)	90.0	83.25	-6.75
Elevation (degrees)	0.0	2.24	2.24
Depth of the reflecting surface at the centroid (cm)	17.5	16.47	-1.03

Table 25. Parameter for the metallic cylinder.

5.9.7. Two metallic pipes in a cross configuration

Figure 100 shows two metallic pipes arrange in a cross configuration. Only the top pipe is detected correctly, the bottom pipe is detected as a small spherical reflector. We suspect the reflections from the top pipe distort the shape of the reflection profile of the bottom pipe.

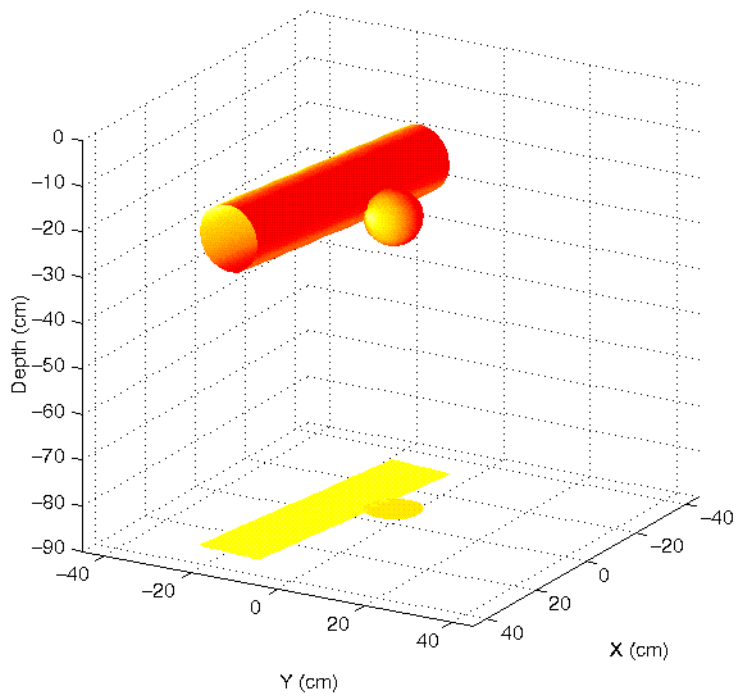

Figure 100. Two metallic cylinders (length= 50cm, radius= 0.5cm) in a cross configuration. Only

Parameter	Actual Value	Computed Value	Error
Centroid horizontal position (cm)	(0.0,-10.0)	(-1.06,-15.82)	5.92
Centroid Depth (cm)	17.0	22.68	5.68
Radius (cm)	0.5	8.06	6.83
Length (cm)	50	67.72	17.72
Azimuth (degrees)	0.0	-6.60	-6.60
Elevation (degrees)	0.0	1.01	1.01
Depth of the reflecting surface at the centroid (cm)	16.5	14.62	-1.88

Table 26. Parameter for the detected metallic cylinder.

5.9.8. A plastic anti tank landmine

Here we show the result of mapping a plastic anti-tank landmine. The mine is shown beside a coke can in Figure 101.

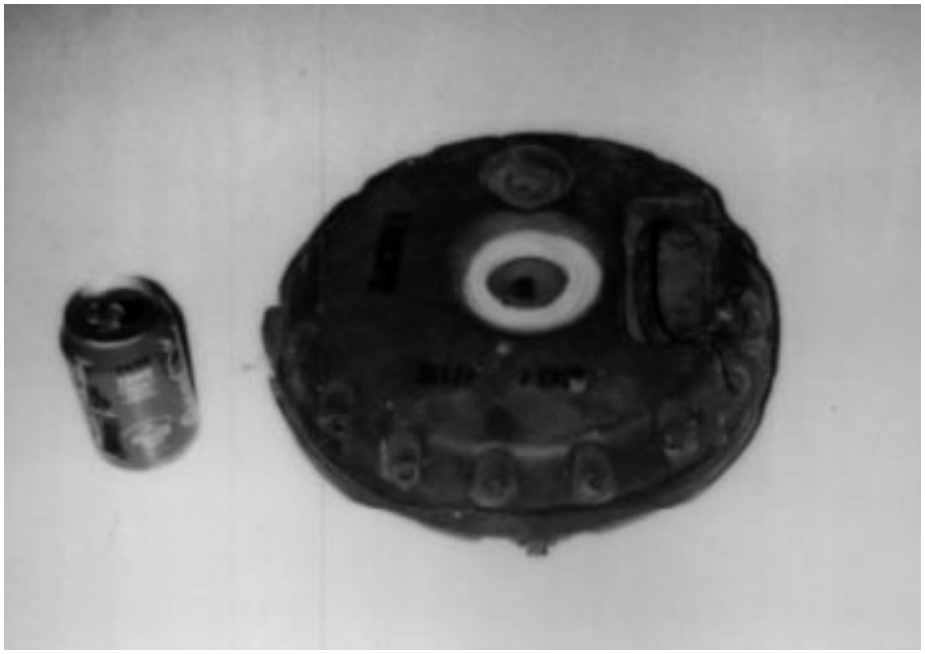

Figure 101: A plastic anti tank landmine.

The density in the center of the mine is higher that at the edge. This is why it is detected as a spherical reflector as shown in Figure 102. The horizontal position and the depth is computed accurately, but the computed size is off by a significant error. We suspect that this is due to the lower density of the material near the edge of the landmine. If the lower density material has a similar dielectric constant as the soil, the reflections from the edges of the landmine will be very weak and they might not be detected.

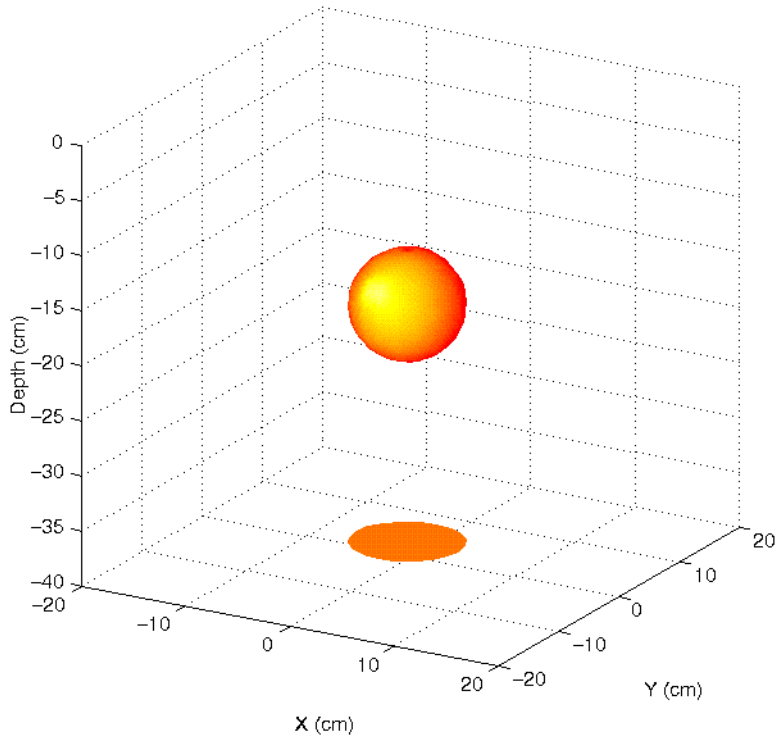

Figure 102. A plastic anti-tank mine (radius= 15cm, thickness=10 cm). It is detected as a spherical reflector.

Parameter	Real Value	Computed Value	Error
Centroid horizontal position (cm)	(0.0,0.0)	(-2.12,3.14)	3.79
Centroid Depth (cm)	22.0	-18.42	-0.87
Radius (cm)	15.0	4.92	-10.08
Depth of the reflecting surface at the centroid (cm)	17.0	13.5	-3.5

Table 27. Parameter for the plastic anti-tank mine.

Chapter 6. Subsurface Mapping Testbed

6.1. Testbed Description

We use a setup illustrated in Figure 103 to do experiment in autonomous mapping and retrieval of buried object. The testbed consists of a robotic manipulator, a GPR, a 2-D scanning laser rangefinder, and a large rectangular soil container for burying various objects. During the experiment, we fill the soil container with wet sand, dry sand or top soil. An industrial robotic manipulator (Cincinnati Milacron T3) is used both for scanning and excavation. When it is used for excavation, a small excavator bucket is attached to the toolplate and when it is used for scanning, an antenna mount is attached to the toolplate instead. We construct the antenna mount from lexan to minimize any interference that might be caused by metallic objects located nearby the GPR antenna. The structure of the mounting adapter is shown in Figure 104. The adapter also moves the antenna away from the metallic robotic arm so the back of the antenna is not located nearby any metallic object. The flexible lexan mounting adapter also provides some protection for the antenna. If the robot smashes the antenna to the ground, instead of crushing the antenna, the mounting adapter flexes to absorp the energy of the impact. This has been proven to be a very valuable feature.

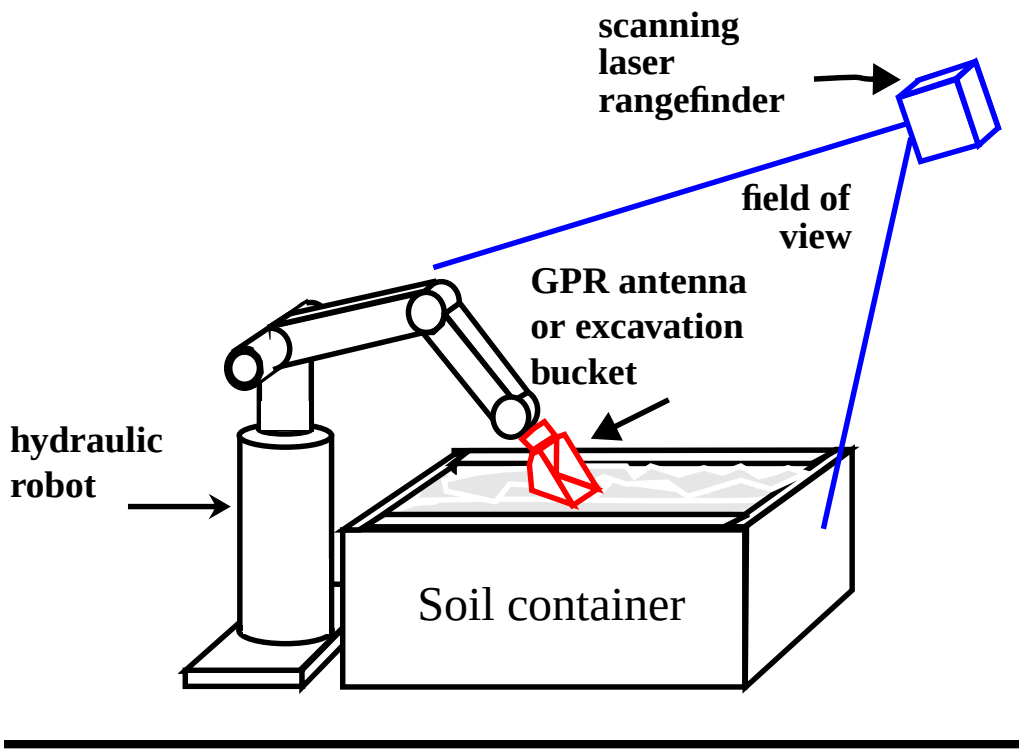

Figure 103. The autonomous subsurface mapper testbed

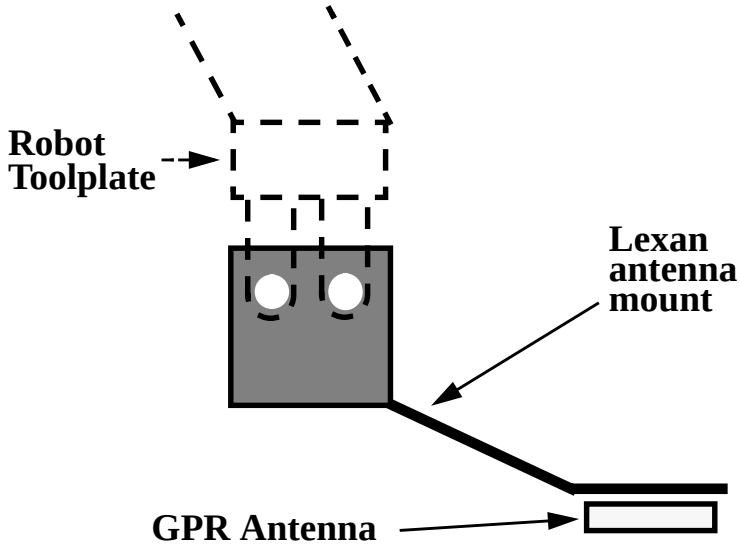

Figure 104. The lexan mounting adapter for the GPR antenna

6.2. Software Architecture

Figure 105 shows the architecture for the experimental buried object mapper and retriever. It consists of four main modules. They are the surface mapper, surface mapper, subsurface mapper, robot controller and excavation planner.

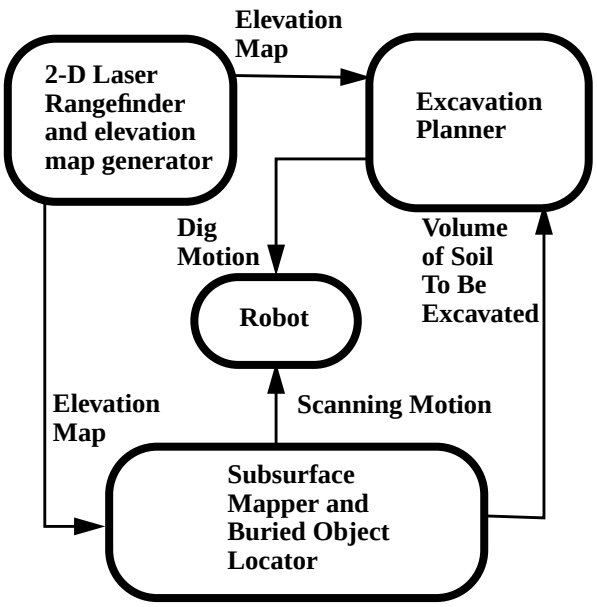

Figure 105: System Architectures

The surface mapper consists of a 2-D laser rangefinder and an elevation map generator. The 2-D laser rangefinder measures the height of the ground in a very fine grid. The elevation map generator then creates an elevation map based on these range data. This elevation map is used by the excavation planner and the subsurface mapper. The excavation planner uses the elevation map for planning the automated excavation. The subsurface mapper uses the elevation map to plan the scanning motion of the GPR antenna. By following the terrain contour, the antenna can be located very close to the ground. Once it collects the GPR data, the subsurface mapper processes the data to locate the buried objects. The buried objects are then shown to the operator. After the operator approves the result, the information about the objects' locations are sent to the excavation planner. The excavation planner in turn will generate the appropriate command to the robot to excavate the soil above the objects.

6.2.1. Surface mapper

The surface mapper module has one main task, which is scanning the soil surface and generating the elevation map. The elevation map is used for several purposes. First, it supplies the excavation planner with the necessary data to plan the digging motion. Second, it is used by the subsurface mapping module to position the antenna close to the surface during subsurface scanning. Last, it is used to create a graphic display of the buried object's projection on the soil surface. In the future it will also be used to scan the exposed portion of a partially buried object. The data will then be merged with the subsurface data obtained by scanning the buried part of the object. In order to do this we must be able to register the data from the surface mapper and the subsurface mapper. Since both the surface and the subsurface sensor are calibrated with respect to the robot position and orientation, it is easy to obtain the correct registration between the data obtained from both sensors.

We use a 2-D scanning laser rangefinder to create an elevation map of the soil surface. The 2-D Perceptron laser scanner that we use is theoretically able to produce a range image with range resolution of 1cm within the maximum range of 40m. In practice the raw range data from a single frame is only accurate to several cm, so we need to filter it to improve the accuracy of the elevation map.This is done using several filtering steps.

First, we remove the outliers by filtering the laser data using temporal median filter. We use the median values from a moving window of 5 range images. Then we filter the resulting range image using a temporal mean filter to remove the random noise. The temporal mean filter computes the means over twenty range images. These two filtering techniques enable us to achieve accuracy close to 2cm. An example of an elevation map generated using these steps is shown in Figure 106.

The availability of a highly accurate elevation map also enables us to improve the processing of GPR data when the soil surface is not flat. Figure 107 shows an example of using the elevation map to compensate for the shape of the soil surface during GPR data processing. If we do not take into account the shape of the soil surface, the location of the buried object will be computed at a wrong location.

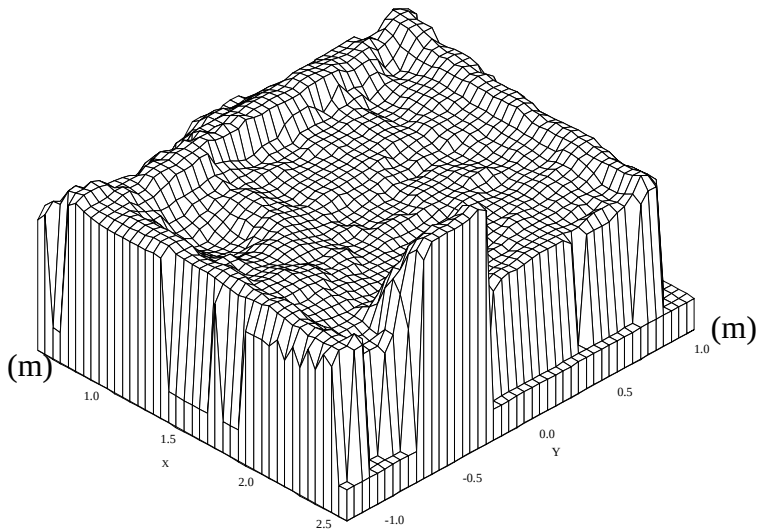

Figure 106: Elevation map of the testbed

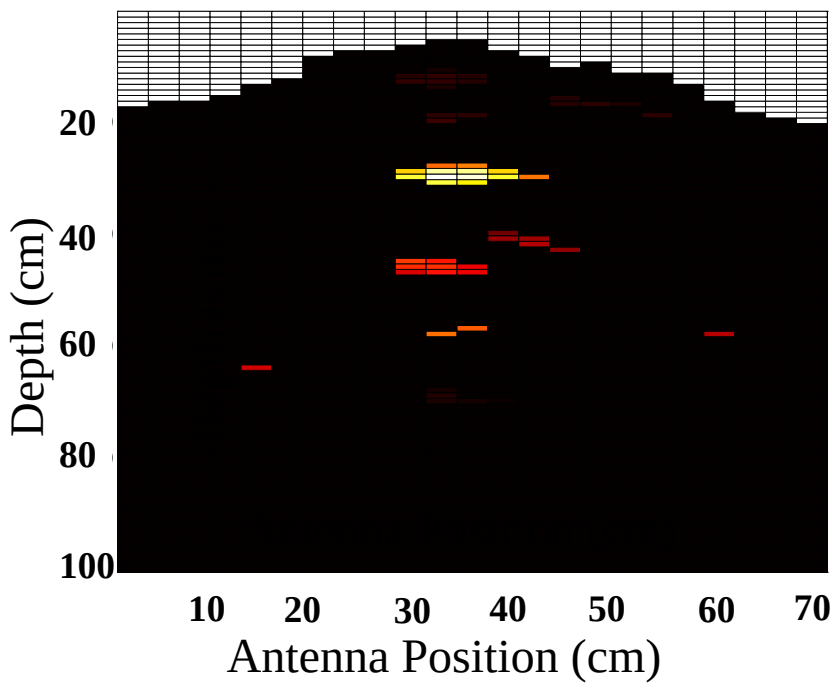

Figure 107. Coherent summation migration of a cylinder buried under a non-even surface incorporating the information from the elevation map.

6.2.2. Subsurface mapper

The subsurface mapper detects, locates and measures buried objects located under the scan area. It can also be used to ensure that there is no buried object in a section of soil that is going to be excavated. We use a GSSI SIR-3 GPR system equipped with a single monostatic 1 GHz antenna. The antenna has a 3db beamwidth of 80 degrees. We set the scan rate to 16 scan/second. Each scan is sampled at 12 bits resolution with a sampling rate of 600 points/ samples. Since our soil container is only about a meter deep, we set the maximum time delay of the GPR system to 10ns. In dry sand which has a propagation velocity of 9 cm/ns, this translate to a maximum distance of 90 cm.

The output of the GPR system is tagged with the antenna position. The position tagged data are then processed using our algorithms. The results are then displayed on a computer monitor, showing the location and shape of the buried objects.

The buried object's location is also used to compute which part of the soil needs to be excavated, in order for the robot to be able to reach the object. Once we know the thickness of soil that the excavator needs to excavate, the information is sent to the excavation planner.

6.2.3. Excavation Planner and Excavator

For the excavator, we use an industrial robot Cincinnati Milacron T3, which is capable of lifting 100 pounds at its end effector. We equip the end effector with a small excavator bucket. It is used to excavate the soil above the buried object and to retrieve the object once it is within reach. The excavation planner uses the information from the subsurface mapper and the elevation map to plan its move. During excavation, it will not excavate too close to the object because of the possibility that it might collide with the object. For more information on the autonomous excavation planner, consult the work by Singh [Singh 92].

We do not currently implement multiple "sense and dig" cycle into our planning software, so basically the planning is a single sense and dig cycle. The GPR antenna is mounted on the robot arm and it is scanned above the area of interest. After the scan is finished, the data are processed to locate the buried object. Once the object is located, the excavation planner develops a plan that will enable the robot to remove the soil above the object. When the object is exposed, the excavator will retrieve it.

6.3. An Example of Mapping and Retrieval of a Buried Object

Figure 108 shows a picture of our testbed. It also shows an example of autonomous buried object retrieval using subsurface sensing to locate the object and then using automated excavation to retrieve the object. The robot excavates the overburden before retrieving the small metallic object. The method used to locate the object is the coherent summation 3-D migra-

tion running on 10 Sun Sparcstation 1 and Sparcstation 2. The processing time is about 2 minutes. If we only use one machine, it will take almost 20 minutes of processing time.

Figure 108 also shows how the end effector of the robot can be equipped with the antenna or a small excavator bucket. At the present time, we need to manually mount the antenna or the bucket, although it is possible to automate such an operation.

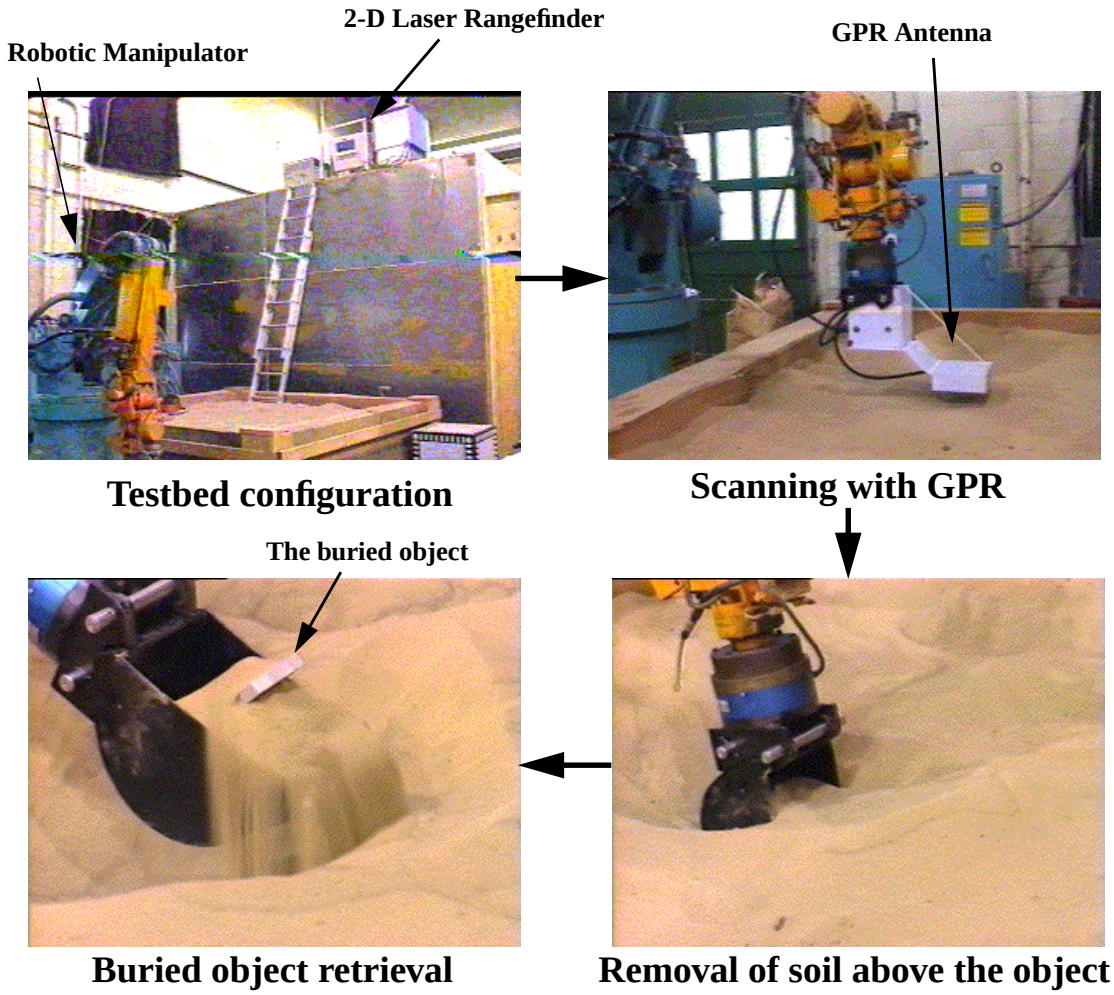

Figure 108. An example of buried object retrieval using subsurface sensing and automated

Chapter 7. Conclusions

7.1. Conclusion

In this thesis, we have examined the problem of subsurface mapping and developed several algorithms for automatically finding buried objects in high resolution 3-D GPR data. Using these algorithms, we have built a robotic system for autonomous subsurface mapping of buried objects. The algorithms that we developed have very fast processing times and can be integrated easily with autonomous excavation to make an autonomous buried object retrieval system. In fact, the ability to do multiple "Scan and Dig Cycle" during excavation actually enables us to obtain a better subsurface map by integrating the information from a series of scans.

The first two algorithms that we developed are called volume based processing, since it operates directly on the 3-D volume data. We improve existing coherent summation migration method by developing a 3-D version which also employs parallel processing technique. In order to extract the buried objects after migration, we use a novel thresholding method that is not dependent on the reflection strength. We also developed and implemented a second migration processing technique which is much more efficient and suitable for real time subsurface mapping application. Unlike the coherent summation technique, this migration method can process the data in small sections. It can begin processing the data even before the scanning process ends. So we can pipeline the data gathering and processing steps,

which results in shorter subsurface mapping time. We call this method "Reflector Pose Estimation".

In addition to volume based processing, we also developed and implemented another method for finding buried objects in high resolution 3-D GPR data. This approach is called surface based processing, because it first reduces the 3-D volume data to a set of 2.5-D surfaces. It then segments these surfaces into different possible objects' surfaces. For each segmented surface which matches one of our buried object's models, it computes the parameters of the buried object.

Volume and surface based processing have their own advantages and disadvantages. It is very important to understand when and where these processing methods should be used, depending on what objects need to be detected and the condition of the soil. Table 28 shows a comparison of the two processing methods.

Processing Methods	Output	Ability to handle soil heterogeneity	Computation time	Need for accurate velocity estimate	Capable of utilizing parallel processing
Volume Based Processing	Object's location, orientation and size	Poor	Single processor: slow/ moderate Multiple processors: fast	Yes	Yes
Surface Based Processing	Object's location, orientation, size and shape	Excellent	Moderate	No	Partially

Table 28. Comparison of two GPR data processing methods

It is important to note that the main disadvantage of the volume based processing method is not the slow to moderate computation time. Despite the amount of computation that is needed for migration, the process is inherently parallel and the processing time can be minimized using parallel processing as we have shown. With the advance in multiprocessor digital signal processing (DSP) equipment, it may be possible to do migration process in real time. The main disadvantage of the volume based processing is the requirement for an accurate propagation velocity estimate.

The surface based processing, on the other hand, is not completely amenable to parallel processing. Only parts of it can be executed in parallel. Its main advantage is its insensitivity to error in the propagation velocity estimate. It also produces additional informations, such as the object's shape.

Since each of the processing methods has its own unique advantages and disadvantages, it is necessary to understand fully the purpose of each subsurface mapping task in order to chose the correct processing method. If detection and localization of buried object are the main concerns, and the propagation velocity is known, then migration will be better suited to the task. A specialized migration processor, composed of several multiprocessor DSP boards, should be able to process the data in real time. This will increase the speed of the object detection and localization. On the other hand, if the shape of the buried object along with its accurate orientation and location are needed, or an accurate propagation velocity estimate is unavailable, then surface based processing of GPR data is more suitable.

The advantages and disadvantages of the two processing methods are almost completely complementary, so it is only natural to combine them by applying them both to a subsurface mapping task. For example, if we just begin to map an unknown area with unknown propagation velocity, we can locate the buried objects using surface based processing. Using the method described in Section 5.7., we then compute the propagation velocity. Once the propagation velocity is known, we can use a volume based processing technique to determine the location of the objects during the sense and dig cycle, since its processing time could be made much faster using parallel processing.

7.2. Contribution

The objective of our research is to solve the subsurface mapping problem by developing the necessary algorithms for autonomous mapping of buried object. We have shown through the development of several processing methods that we are able to map buried object autonomously. Although there are still many improvements that can be made, our experimental results prove that automated subsurface mapping is feasible. In addition, the results also show that such automated mapping techniques are accurate and practical. Since the processing methods extract the buried objects from the GPR data, 3-D visualization of the detected buried objects is possible and very easy to understand. We have also shown that such processing can be done in a minimal amount of time. Although all of our processing is done on general purpose workstations, the processing time is fast enough that currently the bottleneck is the data acquisition step.

In summary, this thesis's contributions are:

- Development and implementation of three new processing methods to find buried objects in 3-D GPR volume data. The first method is a new surface based processing method that uses 3-D segmentation to recognize and locate buried objects from 3-D GPR volume data. This method is radically different with all existing GPR processing methods. Instead of migration, it uses 3-D computer vision techniques to obtain the location, orientation, size and shape of buried objects in the 3-D data. The other two methods are based on the principle of 3-D migration. They are 3-D coherent summation migration and reflector pose

estimation. They are also unique because they do not just migrate the 3-D data, but they also compute the location, orientation and size of buried objects in the 3-D data.

- Demonstration of the first known robotic system to autonomously map and retrieve buried object.
- Development of the necessary methods for integrating subsurface mapping and buried object retrieval process. These include the development of a method to compute the propagation velocity of GPR signal in the soil from the difference in the data obtained before and after removal of soil.
- Laying the ground work for more advance robotics applications of subsurface sensing.

There is definitely a growing need for a fast and accurate subsurface mapping technology. We hope that our work will fill some of these needs. We also hope that our work will create more interest in the integration of subsurface mapping and robotics.

7.3. Future Direction

The logical next step is to try using this subsurface mapping system on real world problems. In order to do this, we need to automate an excavator to be used both for scanning and excavation. By using the system on real world applications, we will learn invaluable lessons on ways to improve the algorithms.

We also need to make our subsurface mapping techniques to work with 2-D as well as 3-D data sets. This is necessary because most of the existing GPR data sets are 2-D. Although we will not be able to extract 3-D parameters from the 2-D data, we believe that having an automated buried object mapper that can operate on 2-D data will increase the rate of acceptance of our techniques for automated processing of GPR data.

Another possible improvement to our processing techniques is the ability to process 3-D data collected on a non-regular grid. When the terrain is not flat, it is not always possible to have a regular raster scanning pattern and the density of the scans will vary in different area.

References

[Apte 89] Apte, S.M., “Object Oriented Implementation of a Domain Specific Planner for a Robotic Excavator”, Masters Thesis, Civil Engineering Department, Carnegie-Mellon University, April 89.

[Beck 94] Beck, A. and Ronen, A., "The Application of Ground Penetrating Radar in Israel: Selected Case Histories", in Proc. of Fifth International Conference on Ground Penetrating Radar, Vol 3, Kitchener, Ontario, Canada, June 12-16, 1994.

[Bergstrom 93] Bergstrom, K.A., Mitchell, T.H., and Kunk, J.R., "Interpretation of Ground Penetrating Radar Data at the Hanford", in Proc. of Second Government Workshop on GPR, Columbus, Ohio, October 26-28, 1993.

[Besl 88] Besl, P., Surfaces in Range Image Understanding, Springer-Verlag, New York, 1988.

[Chang 89] Chang, W.F., McMechan, G.A., “3D Acoustic Reverse-Time Migration”, Geophysical Prospecting 37, pp.243-256, 1989.

[Daniels 93] Daniels, J.J., et al., "3-D GPR", in Proc. of Second Government Workshop on GPR, Columbus, Ohio, October 26-28, 1993.

[Davis 89] Davis, J.L., Annan, A.P., “Ground-Penetrating Radar for High-Resolution Mapping of Soil and Rock Stratigraphy”, Geophysical Prospecting 37, pp. 531-551, 1989.

[Davis 94] Davis, J.L., et al., "Quantitative Measurement of Pavement Structures Using Radar", in Proc. of Fifth International Conference on Ground Penetrating Radar, Vol 3, Kitchener, Ontario, Canada, June 12-16, 1994.

[Deng 94] Deng, S., Zuo, Z. and Wang, H., "The Application of Ground Penetrating Radar to Detection of Shallow Faults and Caves", in Proc. of Fifth International Conference on Ground Penetrating Radar, Vol 3, Kitchener, Ontario, Canada, June 12-16, 1994.

[Eaves 87] Eaves and Reedy, Principle of Modern Radar, Van Nostrand Reinhold, 1987.

[Faugeras 1986] Faugeras, O.D., Hebert, M., "The Representation, Recognition and Locating of 3-D Objects", International Journal of Robotics Research, 5 (3), pp. 27-52, 1986.

[Fiebrich 87] Fiebrich, R.D., "Reverse Time Migration- A Seismic Processing Application on the Connection Machine", Thinking Machine Corp. TMC-87

[Fitch 88] Fitch, Synthetic Aperture Radar, Springer-Verlag, 1988.

[Fisher 92] Fisher, E., McMechan, G.A., Annan, A.P., Cosway, S.W, "Examples of Reverse-Time Migration of Single-Channel, Ground-Penetrating Radar Profiles", Geophysics, Vol. 57, No 4, April 1992, pp. 577-586.

[Fricke 88] Fricke, J.R., "Reverse-Time Migration in Parallel: A Tutorial", Thinking Machine Corp. TMC-88.

[Garmon 85] Garmon, J.P., "Ground Probing Radar Site Parameters Study", Georgia Tech Project A-3869, January 1985.

[Gustafson 93] Gustafson, D.E, and Cohen, M., "Computerized Analysis of Ground Penetrating Radar Images for Hazardous Waste Site Characterization", CSDL-P-3267, Charles Stark Draper Laboratory, April 1993.

[Harris 92] Harris, C.E., McMechan, G.A., "Using Downward Continuation to Reduce Memory Requirements in Reverse Time Migration", Geophysics, Vol. 57, No. 6, June 1992, pp.848-853.

[Herman 94] Herman, H., Stentz, A., "Subsurface Sensing for the Autonomous Retrieval of Buried Objects", in Proc. of Fifth International Conference on Ground Penetrating Radar, Vol 3, Kitchener, Ontario, Canada, June 12-16, 1994.

[Imai 87] Imai, T., Sakayama T., and Kanemori, T., "Use of Ground-probing Radar and Resistivity Surveys for Archaelogical Investigations", Geophysics, Vol 52, February 1987, pp. 137-150.

[Kaufmann 90] Kaufman, A., "3D Volume Visualization", in Advances in Computer Graphics VI, pp. 175-203, IEEE Computer Society Press, 1990.

[Kim 92] Kim, H., Ling, H.,”Electromagnetic Scattering from an Inhomogeneous Object by Ray Tracing”, IEEE Transactions on Antennas and Propagation, Vol 40, No 5, May 1992, pp. 517-525.

[Lay 87] Lay, N.K., Whittaker, W.L., and Motazed, B., "An Automated Mapping System for Ferrous Pipes", in Proc. of Artificial Intelligence in Engineering: Tools and Techniques, Cambridge, Massachusetts, August 1987, pp. 81-94,

[Lord 84] Lord, A.E.Jr.,Tyagi, S., and Koerner, R.M., “Use of NDT Methods to detect and locate buried containers verified by ground truth measurements”, Departments of Civil Engineering and Physics, Drexel University, May 1984.

[Mellet 93] Mellet, J.S., "Bathymetric Studies of Ponds and Lakes Using Ground Penetrating Radar", in Proc. of Second Government Workshop on GPR, Columbus, Ohio, October 26-28, 1993.

[Moghaddam 92] Moghaddam, M. and Chew, W.C., “Nonlinear Two-Dimensional Velocity Profile Inversion Using Time Domain Data”, IEEE Transactions on Geoscience and Remote Sensing, Vol 30, No 1, January 1992, pp. 147-156.

[Oh 92] Oh, Y., Sarabandi, K. and Ulaby, F.T., “An Empirical Model and Inversion Technique for Radar Scattering from Bare Soil Surface”. IEEE Transactions on Geoscience and Remote Sensing, Vol 30, No 2, March 1992, pp. 370-381.

[Olhoeft 94] Olhoeft, G.R., "Modeling Out-of-Plane Scattering Effects", in Proc. of Fifth International Conference on Ground Penetrating Radar, Vol 3, Kitchener, Ontario, Canada, June 12-16, 1994.

[Osumi 84] Osumi, N. and Ueno, K.,”Microwave Holographics Imaging Method with Improved Resolution”, IEEE Transactions on Antennas and Propagation, Vol AP-32, No 10, October 1984, pp. 1018-1026.

[Osumi 85] Osumi, N. and Ueno, K.,”Microwave Holographics Imaging of Underground Objects”, IEEE Transactions on Antennas and Propagation, Vol AP-33, No 2, February 1985, pp. 152-159.

[Ozdemir 92] Ozdemir, T., Roy, S.,Berkowitz, R.S.,”Imaging of Shallow Subsurface Objects: An Experimental Investigation”, IEEE Transactions on Geoscience and Remote Sensing, Vol 30, No3, May 1992, pp. 472-481.

[Peters 86] Peters, L. Jr., and Young, J.D., “Application of Subsurface Transient Radar”, Time Domain Measurements in Electromagnetic, Ed. By Edmund K. Miller, pp. 297-351, Van Nostrand Reinhold, New York, 1986.

{Peters 93] **Peters, L. Jr., "A Focused Reflector as a Ground Penetrating Radar Antenna", in Proc. of Second Government Workshop on GPR, Columbus, Ohio, October 26-28, 1993.**

[Roberts 93] **Roberts, R., and Daniels, J.J., "Analysis of the Effectiveness of Velocity-Depth Inversion Using Bistatic GPR Data Collected over Targets of Different Sizes, Shapes and Orientation", Second Government Workshop on GPR Proceeding, Columbus, Ohio, October 1993.**

[Romero 89] **Romero-Lois, H., Hendrickson, C., Oppenheim, I., "A Strategic Planner for Robot Excavation," in Proc. Sixth International Symposium on Automation and Robotics in Construction, June 1989.**

[Saaraketo 94] **Saarenketo, T., Nikkinen, T., and Lotvonen, S., "The Use of Ground Penetrating Radar for Monitoring Water Movement in Road Structures", in Proc. of Fifth International Conference on Ground Penetrating Radar, Vol 3, Kitchener, Ontario, Canada, June 12-16, 1994.**

[Singh 91] **Singh, S., "An Operation Space Approach to Robotic Excavation," in Proc. IEEE Symposium on Intelligent Control, Alexandria, August 1991.**

[Singh 92] **Singh, S., Simmons, R., "Task Planning for Robotic Excavation," In Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems, Raleigh, 1992.**

[Sonka 95] **Sonka, M., et al., "Three-Dimensional Automated Segmentation of Coronary Wall and Plaque from Intravascular Ultrasound Pullback Sequences", in Proc. of Computers in Cardiology 1995, Vienna, Austria, September 10-13 1995, pp. 637-640.**

[Stolte 94] **Stolte, C., Nick, K.P., "Eccentricity Migration: A Method to Improve the Imaging of Pipes in Radar Reflection Data", in Proc. of Fifth International Conference on Ground Penetrating Radar, Vol 3, Kitchener, Ontario, Canada, June 12-16, 1994.**

[Swetz 86] **Swetz, S., "A System Architecture for Robotic Pipe Excavation," Master's Thesis, Advisor: William L. Whittaker, Department of Civil Engineering, Carnegie Mellon University, Pittsburgh, PA, 1986.**

[Thomas 91] **Thomas, J.G., Peters, R.A., and Jeanty, P., "Automatic Segmentation of Ultrasound Images Using Morphological Operators", IEEE Transactions on Medical Imaging, vol.10, no. 2, June 1991, pp. 180-186.**

[Ulricksen 82] **Ulricksen, C.P.F., Application of Impulse Radar to Civil Engineering, Doctoral Thesis, Department of Engineering Geology, Lund University of Technology, Sweden, 1982.**

[Ulricksen 94] **Ulricksen, P., Stolte, C., "Computerized Scanner System SARA: For the Acquisition and Visulaization of High Resolution 3-D GPR Data", in Proc. of Fifth International Conference on Ground Penetrating Radar, Vol 3, Kitchener, Ontario, Canada, June 12-16, 1994.**

[Vaish 94] **Vaish, J.N. and Gupta, S.C., "Detection of Abondoned Mining Channel under a Residential Complex by Ground Penetrating Radar", in Proc. of Fifth International Conference on Ground Penetrating Radar, Vol 3, Kitchener, Ontario, Canada, June 12-16, 1994.**

[Volakis 83] **Volakis, J.L., Peters. JR., “Improved Identification of Underground Targets Using Video Pulse Radars by Elimination of Undesired Natural Resonances”, IEEE Transactions on Antennas and Propagation, Vol AP-31, NO. 2, March 1983, pp. 334-340.**

[Whittaker 85] **Whittaker, R., et al., “First Result in Automated Pipe Excavation”, Civil Engineering and Construction Robotics Laboratory, Carnegie-Mellon University, June 1985.**

[Yamaguchi 91] **Yamaguchi, Y., Maruyama, Y., Kawakami, A., Sengoku, M., Abe, T., “Detection of Objects Buried in Wet Snowpack by an FM-CW Radar”, IEEE Transactions on Geoscience and Remote Sensing, Vol 29, No 2, March 1991, pp. 201-208.**