

Underground Structures and Gravity Gradiometry

15 January 2002

Prepared by

M. W. MAIER
Computer and Systems Engineering Subdivision
Reconnaissance Systems Division

Prepared for

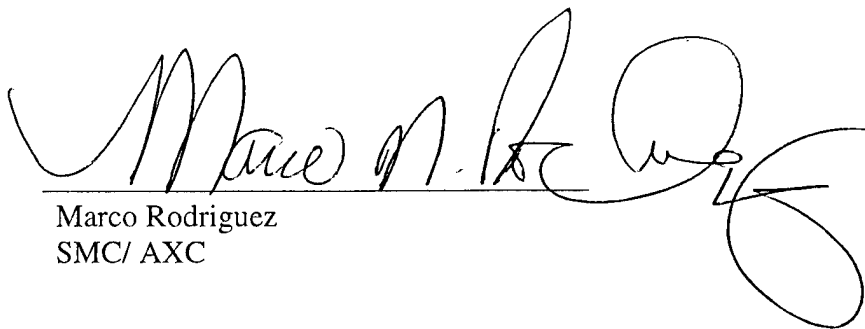
SPACE AND MISSILE SYSTEMS CENTER
AIR FORCE SPACE COMMAND
2430 E. El Segundo Boulevard
Los Angeles Air Force Base, CA 90245

Engineering and Technology Group

This report was submitted by The Aerospace Corporation, El Segundo, CA 90245-4691, under Contract No. F04701-00-C-0009 with the Space and Missile Systems Center, 2430 E. El Segundo Blvd., Los Angeles Air Force Base, CA 90245. It was reviewed and approved for The Aerospace Corporation by Nicola A. Nelson, Principal Director, Computer and Systems Engineering Subdivision.

This report has been reviewed by the Public Affairs Office (PAS) and is releasable to the National Technical Information Service (NTIS). At NTIS, it will be available to the general public, including foreign nationals.

This technical report has been reviewed and is approved for publication. Publication of this report does not constitute Air Force approval of the report's findings or conclusions. It is published only for the exchange and stimulation of ideas.



Marco Rodriguez
SMC/AXC

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY) 15-01-2002		2. REPORT TYPE		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE Underground Structures and Gravity Gradiometry				5a. CONTRACT NUMBER F04701-00-C-0009	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) M. W. Maier				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) The Aerospace Corporation Reconnaissance Systems Division El Segundo, CA 90245-4691				8. PERFORMING ORGANIZATION REPORT NUMBER TR-2001(3000)-1	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Space and Missile Systems Center Air Force Space Command 2450 E. El Segundo Blvd. Los Angeles Air Force Base, CA 90245				10. SPONSOR/MONITOR'S ACRONYM(S) SMC	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) SMC-TR-02-14	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This report evaluates the use of gravity gradiometry as a tool to characterize underground facilities. The gradient of the Earth's gravitational field can be used to make inferences about the Earth's mass density field. Mass deficits, such as those produced when an underground region is excavated, produce a characteristic change in the Earth's gravitational field. In principle, gradiometric measurements can make inferences about mass deficits that are covered by hundreds of meters of rock. In practice, the technique is much more limited. The report studies the potential for characterizing underground facilities from gravity data using matched filters and a variety of inversion techniques. All of the techniques indicate that the fundamental spatial resolution is directly related to the distance from the point of interest to the gradient sensor. Thus, resolution of feature sizes of man-made facilities requires quite close access by the gradient sensors. Detection applications are also limited by the unknown size of spatial gravity "clutter" due to the inhomogeneous Earth. Useful applications may be possible where high-sensitivity instruments can be used in proximity to the area of interest. Measurements from high-altitude aircraft or space may be strongly synergistic with short-distance measurements.					
15. SUBJECT TERMS Gravity gradiometry, Underground facilities					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 60	19a. NAME OF RESPONSIBLE PERSON Mark Maier
a. REPORT UNCLASSIFIED	b. ABSTRACT UNCLASSIFIED	c. THIS PAGE UNCLASSIFIED			19b. TELEPHONE NUMBER (include area code) (703)633-5350

Acknowledgements

Mr. Mike Martino developed the analysis code that generated some of the matched filter plots. Dr. Pete Palmadesso provided many very useful comments that improved the report.

Contents

1.	Executive Summary	1
2.	Introduction	5
2.1	Gravity Gradiometry Concepts.....	5
2.2	Usage Models.....	5
2.2.1	Target Model	6
2.2.2	Collection Model.....	6
2.2.3	Change Detection	6
2.2.4	Direct Detection.....	7
2.2.5	Characterization.....	7
2.2.6	Inversion to Images	7
2.3	Previous Work.....	8
2.3.1	Stanford/Yale/AFRL proposals	8
2.3.2	Jet Propulsion Laboratory Directors Innovation Initiative	8
2.3.3	Other Work.....	10
3.	Physics and Mathematics Preliminaries	11
3.1	Gradiometric Field	11
3.1.1	Tensor Plots	12
3.2	Noise Models.....	12
3.2.1	Gradiometer Sensitivity.....	12
3.2.2	Natural Background	15
3.3	Size of the Effects	16
3.4	Detection Models	17
3.5	Inverse Gravimetry and Gradiometry.....	17
3.6	Matched Filter Concepts.....	19
4.	Detection with Known Background	21

4.1	Simple Detection.....	22
4.2	Structured Discrimination.....	23
5.	Detection with Unknown Background.....	27
6.	Detection with Complex Signatures.....	29
7.	Inversion and Imaging with Matched Filters	31
7.1	Reference Geometry Model.....	31
7.2	Simple Matched Filter Processing Model.....	31
7.3	Volume Spread Function Equations.....	33
7.4	Matched Filter Point Spread Function	34
7.4.1	For Different Tensor Components.....	34
7.4.2	As a Function of Altitude.....	38
7.4.3	As a Function of Measurement Area.....	40
7.5	Signal-to-Noise and Clutter-to-Noise.....	42
7.6	Matched Filter Inversion Results	43
7.6.1	Basic Results	44
7.6.2	Spatial Resolution of Features	45
7.6.3	Monopole Signal Screening.....	47
7.6.4	Density Variation Impacts	47
8.	Constrained Inversion Algorithms	49
8.1	Basic Least-Squares Results.....	50
8.2	Spatial Resolution of Features.....	53
8.3	Monopole Signal Screening.....	53
8.4	Density Variation Impacts.....	53
9.	Summary of Feasibility.....	57
9.1	Inversion versus Detection	57
9.2	Inversion of Ground-Based Measurements.....	58
9.3	Inversion from Airborne Measurements	59
9.4	Inversion from Spaceborne Measurements.....	60
	References	61

Figures

1. Shape of the gravity gradient tensor T_{xy} for a point mass.....	12
2. Shape of the gravity gradient tensor T_{xx}	13
3. Shape of the gravity gradient tensor T_{yy}	13
4. Shape of the gravity gradient tensor T_{xz}	14
5. Shape of the gravity gradient tensor T_{yz}	14
6. Ratio of gradient terms as a function of sensor altitude.....	17
7. Many physically distinct objects have the same external gravitational field.....	18
8. Gradient deviation for a reference tunnel seen from a 200-km orbit.....	22
9. Gradient deviation for a reference tunnel seen from an aircraft at 1-km altitude.....	22
10. Probability of error curve for two-tunnel discrimination as a function of sensor height.....	24
11. Probability of discrimination error with the SNR raised to 1000.....	25
12. Classification performance with high SNR and increased measurement density.....	25
13. T_{xx} for a line mass measured at a distance of one-twentieth its length.....	29
14. T_{xx} field for a line mass measured at 50 times its length.....	30
15. Geometry used for evaluating gradiometric processing.....	32
16. Five-component point-spread function.....	34
17. Point-spread function for T_{xx} alone.....	35
18. Point-spread function for T_{yy} alone.....	35
19. Point-spread function for T_{xy} alone.....	36
20. Point-spread function for T_{xz} alone.....	36
21. Point-spread function for T_{yz} alone.....	37
22. X-axis cut through the point-spread functions.....	37

23. Z-axis cut through the point-spread functions.	38
24. The point-spread function shrinks directly with sensor altitude.	39
25. Comparison of cuts through the point-spread function at various sensor altitudes.	39
26. Product of T_{xx} and a displaced copy for x from 0 to 5 and displacements from 0 to 2	41
27. T_{xx} product plot in $10 \cdot \log_{10}$ of the absolute value of the product.	41
28. Product of the T_{xz} components along the x-axis from 0 to 5 for displacements of 0 to 2.	42
29. T_{xz} plot repeated but scaled to $10 \cdot \log_{10}$ of the absolute value to better show small values.	42
30. Matched filter inversion at a height of one unit.	44
31. Matched filter inversion of the two-tunnel case for an altitude equal to the spacing.	45
32. Zero-depth cut through the matched filter inversion.	46
33. Matched filter inversion of two tunnels with sensor altitude of 5 times the tunnel spacing.	46
34. Matched filter inversion of a screened tunnel with sensor height of 1.	47
35. Matched filter inversion of a screened tunnel from a height of five.	48
36. Matched filter inversion with 5% random density variation.	48
37. Original mass distribution in inversion experiment.	50
38. Estimated mass distribution with a sensor height of 0.25.	51
39. Estimated mass distribution with a sensor height of one.	51
40. Estimated mass density with a sensor height of two.	52
41. Estimated mass distribution with a sensor height of five.	52
42. Least-squares inversion for two tunnels with a sensor height of one.	53
43. Least-squares inversion of the two tunnel case for a sensor height of five.	54
44. Least-squares inversion of a screen tunnel from a height of one.	54
45. Least-squares inversion of the screen tunnel from a height of five.	55
46. Least-squares inversion for 5% randomized density points from a height of one.	55

Table

1. Summary of Gravity Gradiometry Feasibility.....	2
--	---

1. Executive Summary

This report summarizes studies on the use of gravity gradiometry as an intelligence/surveillance/reconnaissance (ISR) tool, specifically to characterize underground facilities. Measurements of the Earth's gravitational field using gravity gradiometry are used to make inferences about the Earth's mass density field. In particular, gradiometric data is used to find mass deficits when attempting to characterize underground facilities. This technique has great promise because inferences can be made about mass deficits that are covered by hundred meters of rock. Since no electromagnetic radiation is involved, camouflage cannot attenuate the signals of interest. This makes gravity a powerful approach to countering certain types of denial and deception.

Unfortunately, gravity gradiometry using only satellites is unlikely to be practical for military ISR applications. At orbital distances, inversion to the feature sizes of interest is almost certainly impossible, regardless of the sensitivity of the gradiometer. On the other hand, the technique appears to hold considerable promise for short stand-off applications, and probably in airborne applications as well. In both cases, the geometry and physics appear to allow inversion to the feature sizes of interest, and the necessary sensitivities are probably achievable, albeit after significant investment. The algorithms needed to invert the data have yet to be developed, but they appear (in principle) to be possible. A collection program that includes both measurements from space and at ranges of 1 km or less may have very positive synergies.

Feasibility in ground and aircraft applications should not be taken as demonstrated by this report because a number of incompletely modeled effects are present. For success there needs to be adequate sensitivity, a clear target signal that can be separated (spatially and/or temporally) from the naturally occurring background variation, and a method to invert the measurements into the desired inferences. For the surface/air applications, the most fundamental issues are the natural temporal variation in the gradiometric field at the time scales of interest and the background spatial spectral content of the gradiometric field. The issue of variation levels in the gradiometric field at and below the milli-Eotvos range is problematic and yet unresolved. There is, even in principle, no way to resolve these background clutter issues without building instruments of adequate sensitivity and collecting data.

The difficulties in accomplishing a useful capability from space make this a poor investment area for research unless all collection modes are considered. The existing studies (and this study) have only identified change detection as potentially feasible. The fundamental problems are the difficulty of inversion and the lack of knowledge about gradiometric clutter. While it is certainly very difficult to make an instrument with enough sensitivity to carry out the proposed functions from space, it may be possible. Past studies (and this study) show the difficulty in overcoming the inversion resolution limit of the sensor altitude. An inversion algorithm would need to exceed the natural resolution limit by at least 4 orders of magnitude. No algorithms that can exceed the limit by even 1 order of magnitude are known, and there are no paths of investigation that show clear promise for such an enormous improvement. Moreover, present levels of uncertainty in surface topography are typically the same

order of magnitude as the size of volume pixels that are needed for military ISR problems. This means the background clutter level is almost certainly much larger than the signals of interest. While the multi-dimensional nature of gradiometric signals might lead to especially effective clutter discrimination techniques, it is unlikely in the space-based measurement case. Solving the inverse problem will require algorithms with entirely different assumptions and capability from those proposed in the literature to date. In spite of the difficulties in accomplishing a useful capability from space make this a poor investment area for research, it may be worth investing in limited exploratory research to determine whether such algorithms can be constructed. This area also becomes much more favorable for research if ground or air applications are included.

While feasibility in ground and air applications is not demonstrated, it is reasonably likely, particularly for ground-based measurements. In these applications, a broader investment program can be justified. In particular, it would be worthwhile to make some significant ground-based collections at known underground facilities to test processing algorithms. It may also be reasonable to investigate instruments suitable for aircraft use and with sensitivities adequate for characterizing facilities at aircraft ranges. Sensitivities in the milli-Eotvos range should be sufficient for ground-based studies. The overall feasibility analysis is summarized in Table 1.

Change detection is where underground structures are detected by observing changes to the gradiometric field over long periods. This mode is much less sensitive to clutter since the time-invariant clutter can be estimated from the collections. All three collection scenarios are potentially possible for change detection. Feasibility in the space collection cases are dependent on facility size. Large structures will produce a signal large enough to be detected by proposed gradiometers. Moderate size structures are at the sensitivity edge of proposed gradiometers. Both space collection cases are marked with an asterisk because they are sensitive to the facility constructors' excavation tailing disposal policy. If the tailings are dropped within a few tens of kilometers of the facility, the size of the gradient signature is greatly reduced, and changes may no longer be detectable. If the tailings are transported 200 km or more away, then the signature is not attenuated. When the tailings are disposed of close to the facility, then a dipole signature remains, but its size is reduced, and it may be nearly non-existent if the tailings are spread around evenly.

Detection of previously built structures cannot be fully evaluated here because it is limited by spatial clutter. The gradiometric field will vary from many sources (e.g., topography, batholiths, geophysical activity, etc.). These natural, spatial variations may mimic the signature of underground facilities or may be easily separable. We do not know which since there has been no analysis of data at the

Table 1. Summary of Gravity Gradiometry Feasibility

Application	Ground collection	Airborne collection	Space collection
Change detection on large structures ($>10^6$ m ³)	Yes	Yes	Yes*
Change detection on moderate structures (10^4 m ³)	Yes	Yes	Maybe*
Detection of previously built structures	Yes, but unresolved issues	Maybe, but unresolved issues	Unlikely, unresolved issues
Characterization or inversion to the 10 m ³ scale	Probably, with advanced sensors	Unlikely, but possible	No

required sensitivity levels. It is possible that a study of the NASA GRACE mission data could partially resolve this issue. The GRACE instrument is not designed for the sensitivity that would be needed for a space-based military ISR mission, but the GRACE data could provide some indication of clutter levels. The needed analysis would be to try to fit a density model to the GRACE data and minimize the residuals between the model and the data, taking into account all known topographic and geophysical information. If the residual can be reduced to a power level roughly equivalent to the noise floor of the instrument, then we can upper bound the clutter level that would be experienced in a direct detection application. If, after extensive modeling, there is a residual that is similar in spectral content to underground facilities, but larger by several orders of magnitude, then we know that detection of pre-existing underground facilities from gradiometric data is probably impossible.

The characterization row is for the problem of actually inverting the gradiometric data into a density image. This is the preferred approach, from an ISR perspective, and is the approach most extensively studied here. Inversion performance is limited because the natural resolution of a gradiometric system is the altitude of the sensor. As a result, space gradiometers image on a scale of 100's of kilometers, airborne sensor on a scale of kilometers, and ground-based sensor on scales of 10's to 100's of meters. There appears to be no way to improve the resolution by more than an order of magnitude or so, rendering space applications in the row impossible and airborne application problematic.

The discussion is broken into an introduction, mathematical preliminaries, the main body of the analysis, and conclusions. Section 2 introduces the overall problem and modes of use of a gradiometric sensing system. Section 3 gives the relevant mathematics and physics background. Section 4 analyzes detection with known background. Section 5 analyzes detection with unknown background. Section 6 discusses detection with complex signatures or structured signatures. Section 7 addresses the main problem, that of direct inversion of gradiometer data to a density map. Section 8 looks at alternative inversion methods that may be able to improve on the matched filter results. Section 9 looks across all of the results and draws conclusions for the feasibility of ISR applications in each of the scenarios. Finally, Section 10 makes recommendations for future research.

2. Introduction

The potential use of gravity gradiometry is a complex technical problem, and one that involves several scientific and engineering disciplines. As a result, discussions about the possible utility of gradiometry, and various technical approaches, have been hindered by the lack of a common vocabulary. This introduction describes, in a self-contained way, a number of key concepts needed for understanding proposals in this area. The first section covers the basic concepts of gradiometry, how gradients can be measured, and how missions might be imagined. The second covers usage models, i.e., potential ISR applications described in a generic way to avoid specific requirements that cannot be discussed in the open literature. The third section covers previous work considered in this investigation.

2.1 Gravity Gradiometry Concepts

As is well known, a gravitational field generated by that body surrounds all massive bodies. The gravitational field is a vector field. At each point in space there is a three-dimensional vector representing the gravitational force at that point. The gradiometric field is the gradient of the gravitational field. The gradiometric field is a tensor field. Since the gravitational field is made up of three vectors, the gradient at each point has nine components. The nine components are the x , y , and z derivatives of the x , y , and z components of the gravitational vector. Because of symmetry there are only five independent components of the gradient tensor.

The goal of gravity gradiometry in ISR is to use measurements of the gravitational field (either directly or of its gradient) to make inferences about hidden objects. The tremendous advantage of gravity detection is that it is nearly impossible to screen. An object buried under 100 m of rock is essentially as visible as if displayed on the Earth's surface, in the sense of its gravity field being observable.

Scalar gravity measurements, vector gravity measurements, and gradient measurements might all be of military ISR interest. The discussion here focuses on gradient measurements because gradient measurements are by far the most suitable for space application, indeed for collection by any moving platform. Because acceleration and gravity are fundamentally indistinguishable, all gravity sensors have the problem of removing accelerations from their gravity measurements. For ground fixed observations, this is much less of a problem, although seismic microtremors and other terrestrial vibrations can be a problem even here. In space applications the basic gravity vector is unobservable because of the free-fall environment of space. However, gravity gradients, that is, the change in gravitational acceleration in space, are immune to vibration and are observable in free-fall.

2.2 Usage Models

The performance analysis of gravity methods depends on having an understood usage model. The usage model describes the targets to be characterized, how the gravity field data will be collected and

what inferences the analyst will seek to draw from its observation. We discuss six aspects of usage: The target model, the collection model, change detection, direct detection, characterization of underground facilities, and field inversion to images.

2.2.1 Target Model

There is an enormous range of underground facilities of potential ISR interest. They range from simple infiltration tunnels to large complexes housing industrial facilities. For the purposes of this report we will use a simple tunnel assumed to be 10 m in diameter and 100 m long. The exact depth and orientation of the tunnel is less important, and is considered in the calculations. The gradiometric signature of this tunnel can be scaled to other cases, and it is representative of a number of real ISR problems.

2.2.2 Collection Model

In all cases, we assume we can measure the gradiometric field in some region above the Earth. The size of the region, the altitude, the tensor components measured, and number of times we measure are all parameters to explore. In the simplest cases, we measure the gravity gradient along a curve of constant radius from the center of the Earth and seek to make inferences about the structures below. In more complex cases, we will measure all tensor components at many points throughout a three-dimensional region above the Earth.

2.2.3 Change Detection

The change detection case seeks to use gravity measurements to detect new underground construction. We assume we have a gradiometer repeatedly passing above some area of the Earth. Over time, we build up a baseline model for the gravity gradients. Each time we re-measure the gradients we compare to the baseline. If the deviation (perhaps integrated over many passes) varies from the baseline in a characteristic way, we say we have detected a new facility.

The simplest way of characterizing the performance of this case is by the minimum size of a new underground void that would pass the noise threshold of the system. A more sophisticated measure is the probability of detection and false alarm realized by a sensor for a given new underground void size. Performance in this case is fundamentally limited by the sensitivity level (that is the noise floor) of the gradiometer and the level of naturally occurring gradiometric noise (due to weather, geologic processes, astronomical variation, etc.).

We will also examine the special case of detection with known background, that of highly structured discrimination. Unstructured discrimination, that is, trying to infer aspects of the internal structure of the facility without any assumptions, is treated in the section on characterization and inversion to image. The structured case is where we want to use gravity data to discriminate between two precisely stated hypothesis about the facility. The case we will use is very simple, distinguishing between a single tunnel and two tunnels with the same total internal volume. We can use the two tunnel discrimination case to show the sensitivity of discriminatory power to sensor height.

2.2.4 Direct Detection

Direct detection is similar to change detection, except the idea is to detect all underground facilities without a baseline absent of those facilities. In this case, we collect the gradiometric field over some region and then search the collected data for some distinctive signature of an underground void. The appropriate performance measure is the probability of detection and false alarm in the face of sensor noise and natural gradiometric clutter.

2.2.5 Characterization

In characterization we know that an underground void exists but seek more detailed information about hidden characteristics. For example, we might observe two tunnel adits and want to estimate the total excavated volume. Or we might want to estimate the internal configuration (e.g., a simple tunnel, a tunnel and galleries, large spaces). Or we might try and characterize the contents of an excavation (e.g., empty or filled with machinery). The performance measures for this case depend on what we specifically are trying to estimate. If we are trying to estimate the interior volume, we should use the error of volumetric estimate in the presence of disturbances, including both sensor noise and natural geologic variation.

The reference tunnel size, 10-m diameter by 100 m length, gives a scale for characterization. If we are trying to judge the internal layout or extent, we will need to discriminate differences on the scale of a 10-m cubed block or rock. If we are interested in simple detection, we will need to be able to detect a mass deficit of approximately $10,000 \text{ m}^3$ of rock.

2.2.6 Inversion to Images

The real desire in this concept is to produce volumetric images (image in volume pixels or "voxels") of underground structures. In this case, we collect the gradiometric field over some region and then estimate the mass-density field that produced the measured gradiometric field. This is inverse gradiometry. Inverse gravimetry and gradiometry have been studied in the geophysics community and applied to some satellite measurements. Past applications have produced "images" with resolutions of many kilometers. These are sufficient for geodetic research, such as looking for undersea mountains, but are unsuitable for military ISR applications. Our goal is to find methods of gradiometric measurement and inversion that supports volumetric imaging on the scale of meters.

An important concept here is the distinction between local and global imaging. In the local case, we take measurements in some region and then do an inverse computation of the mass-density field only in the area of interest. In the global case, we have to estimate mass-density field parameters for the entire Earth in order to extract the mass-density field for the area of interest. Since the area of interest is much smaller than the entire Earth, it is obviously desirable to do local rather than global computations.

For inversion to images the natural performance measures are resolution, signal-to-noise ratio in the images, and the signal-to-clutter level (or related volumetric point-spread function parameters). Resolution in the imaging case refers to resolution in the mass density field as opposed to resolution

in the gradiometric field. The two are presumably related, but the relationship is a complex one that comes from the inversion mathematics.

2.3 Previous Work

Gradiometry and inverse gradiometry are known subjects in the geoscience community. Considerable work, including experiments, has been done by NASA, universities, and private companies. There is a modest body of published literature dating back to at least the 1970's. Most work has focused on building better gravity maps rather than inverting the maps to density fields. Work on inverting to density fields has concentrated on finding large features, such as underground mountains, rather than smaller scale characterization of features.

2.3.1 Stanford/Yale/AFRL proposals

The Air Force Research Laboratory (AFRL) has been collaborating with the National Imagery and Mapping Agency (NIMA) on gravity gradiometry and inversion to density maps. This grows out of long-term work by NIMA on building gravity maps for inertial navigation. There has been considerable recent interest by AFRL and NIMA in work being conducted at Yale and Stanford on ultra-sensitive gradiometers based on quantum effects. Most of the research proposals have been on instrument design and physics studies, but some consideration has also been given to detection and processing. The canonical example used by these groups is detection of new facilities. This use-case is taken up below as detection against a known background (since it is envisioned as change detection). In most respects, this is the easiest case, and the sensitivities involved are such that one can argue that the technique is feasible, even from space. However, this report argues below that the unknowns and operational limitations of this technique, especially in space applications, are the worst of all the use-cases, and it is unlikely to be practical in any circumstances, even though its sensitivity requirements are least.

2.3.2 Jet Propulsion Laboratory Directors Innovation Initiative

The NRO has funded several Directors Innovation Initiative (DII) projects in gravity gradiometry. The work done by the Jet Propulsion Laboratory (JPL) under the DII program is particularly germane. The work done by other DII projects on gradiometry is not discussed here.

The JPL projects have generated some interesting results in both instrument design and processing. At the time of this writing, one of the final reports from JPL was available. The report is divided into several parts; some are devoted to instrument issues, while others are devoted to processing issues. The part germane to this study is the part addressing processing by J. Clauser. Clauser's excellent work reaches conclusions essentially in agreement with the technical conclusions reached here. Clauser's work does not speak to the issues of research recommendations, and one should make no interpretation of such cross-support.

Clauser, this report, and other work conducted on the subject all have substantial agreement on the issue of spatial resolution and its interpretation. These studies all show that mass features with spatial

separations on the order of the sensor altitude are resolvable from each, but smaller features are not.¹ These studies all show that the fundamental resolution of a gravity gradient system is its altitude. As this report will show, there are some techniques that can modestly sharpen that resolution, but almost certainly not by orders of magnitude.

A good understanding of Clauser's results requires some discussion of spatial resolution. The fundamental meaning is the ability to discriminate one feature from another when they are separated by some distance or differ in size. For example, can one distinguish two tunnels running in parallel some distance apart from a single tunnel volume equal to the combined volume of the two separate tunnels? In this simple example, the answer is no if the tunnel separation is less than the sensor altitude. So, a spaceborne gradiometer is generally unable to distinguish features of less than a hundred kilometers or so in size from other features.

The issue of spatial resolution is different from the ability to detect, especially the ability to do change detection. This is discussed at length below and is carefully discussed in Clauser. Just as an optical system can detect bright points of light even when their angular extent is much smaller than the diffraction limit, a gradiometer system can (in principle) detect mass points smaller than its spatial resolution. As in an optical system, the ability to detect is not the same as the ability to image. There is an important difference between the optical case and the gravity case. An optical system can (when conditions are right) do either non-coherent or coherent resolution improvement. Non-coherent resolution improvement increases the resolution by taking several intensity images and processing or centroiding the objects of interest. Coherent resolution improvement uses phase interference among different observations to improve the fundamental resolution. As Clauser points out, there is no analog of coherent resolution improvement or interferometry in gradiometry. There is no phase information to interfere. Any resolution improvement must be done non-coherently on the intensity information. In optical systems, such improvements are usually modest (much less than a factor of 10). In the gradiometry case, we will see that intensity processing can improve resolution, but there are similar limits. For spaceborne sensors, this effectively precludes their use in characterizing facilities because the required improvement is many orders of magnitude over the sensor altitude. As will be shown in a later section, the ability to discriminate is non-existent for features much smaller than the sensor altitude, even in highly structured, artificial situations. For ground-based and possibly airborne sensors, the required degree of improvement is possibly achievable, as it is several orders of magnitude less.

The analysis by Clauser of noise sources is illuminating and appears to be the only careful study of this type at the precision levels that would be needed. An important conclusion is that full tensor measurement from space at the required accuracy levels is almost certainly unfeasible. Fortunately, as Clauser argues and computation below shows, there is little lost by using only a single tensor component. The inversion characteristics of the various tensor components are similar, at least to a factor of 2 or so. Clauser discusses, in the "Amazing-GRACE" mission outline, how a mission might be constructed that could measure one tensor component in the range of useful sensitivities.

¹ See, for example, Clauser at I.2, page 36. In this section, Clauser dismisses all applications except change detection.

Another issue introduced by Clauser is the possibility of looking at either monopole or dipole components of the field. This is important for any proposed spaceborne application since the monopole components are easily erased (either deliberately or by routine) in most underground facility constructions. The problem is that the excavated Earth is normally deposited somewhere within 200 km of the construction site (usually much closer). When the excavation is deposited relatively close to the construction site the inability of a spaceborne sensor to distinguish mass points closer than its altitude effectively erases the facility signature. The mass deficit of the excavation is cancelled by the additional mass of the excavation material being dropped. Smaller movement still creates a mass dipole, which may be detectable, but is much smaller than if the mass were removed. This issue is examined computationally in Sections 7 and 8 in selected cases.

Clauser proposes that one can look at dipole components of the field as a way around this problem. Dipole components are still present when the displacement of material is small compared to the spatial resolution of the system; however, they are small compared to the monopole component that would otherwise be present. This leads directly to the first use-case, change detection. This applies to the problem of finding a newly constructed facility against a known background. Clauser suggests that this is relevant to finding new, covertly constructed facilities or to detecting covert underground nuclear tests. However, it is important to point out that it does not assist with the other use-cases, detection without a known background or characterization. The application to nuclear testing also needs to be compared to other detection methods and the device size needed to produce a detectably vertical mass dipole. The gravity technique requires that there be a large vertical mass displacement to be detectable. It is unclear whether such a displacement stably occurs even in normal underground nuclear tests much less in covert tests where the movement of ground (which produces a seismic signature) has been carefully suppressed.

2.3.3 Other Work

There is a modest body of published literature on gradiometric instruments and inverse gradiometry. Of particular note is that there are private companies using gradiometry in mineral prospecting applications. One company, Bell Geospace (www.bellgeo.com), advertises a commercial service collecting and processing gradiometric tensor data. However, their methods have two important limitations. First, they only collect from ships. Their website indicates they are working on airborne collection but do not yet do it. And second, they only use the gradiometric data to interpret seismic data. That means they do not directly invert the gradiometric data into density maps. Instead they use it to constrain and refine reflection seismology data. There are ISR cases where this might be applicable (cooperative ground-based collection), but most spaceborne collection scenarios are not among them.

3. Physics and Mathematics Preliminaries

A number of preliminaries are necessary to understand the performance estimates. Here we cover models for the gradiometric field, noise in sensors and the natural environment, detection models, inverse problems, and matched filters.

3.1 Gradiometric Field

For this study, we use classical, Newtonian gravity. The gravitational potential of a point mass is given by:

$$U = \frac{Gm}{r}$$

Where U is the potential energy, r is the distance from mass to the field point of interest, m is the mass of the mass, and G is the gravitational constant. The gradiometry tensor can be obtained by differentiation of U twice with respect to the Cartesian product of the three axes, x , y , and z . The five independent components of interest (with the product Gm set to one) are:

$$T_{xx}(\mathbf{r}) = \frac{\|\mathbf{r}\|^2 - 3x^2}{\|\mathbf{r}\|^5}$$

$$T_{yy}(\mathbf{r}) = \frac{\|\mathbf{r}\|^2 - 3y^2}{\|\mathbf{r}\|^5}$$

$$T_{xy}(\mathbf{r}) = \frac{3xy}{\|\mathbf{r}\|^5}$$

$$T_{xz}(\mathbf{r}) = \frac{3xz}{\|\mathbf{r}\|^5}$$

$$T_{yz}(\mathbf{r}) = \frac{3yz}{\|\mathbf{r}\|^5}$$

Here we use mixed coordinates with r for the radial distance and x , y , and z the Cartesian displacements from the mass to the point of interest. In practice, to make computations, one must appropriately reconcile coordinate systems. This is done in the computations to be presented later, but is

unimportant here. For most of the computations, we will set the product Gm equal to one for convenience. In those cases where quantitative values are important, we reintroduce it. The units of gravity gradient are inverse seconds squared. In the gradiometry community it is common to use units of Eotvos. One Eotvos is 10^{-9} inverse seconds squared.

The effects of interest in this study are sometimes extremely small. That implies that a description in Newtonian gravity might not be sufficient for some applications. It is possible that general relativistic effects are significant at the gradient scales of interest. This subject has not been investigated here.

3.1.1 Tensor Plots

The individual tensors have characteristics shaped when visualized in the x-y plane. Some reach peaks at closest approach to a point mass, others have a zero at closest approach. Plots of the tensors are given in Figure 1 through Figure 5.

3.2 Noise Models

As soon as we talk about sensitivity, we have to talk about noise. To say that a gradiometer has a sensitivity of some number of Eotvos means that the instrumental noise must be somewhat below that number of Eotvos. Of course, measurements instruments are not the only source of noise. The gradiometric field of the Earth is not constant. At the scale of most human interest, it is constant, but at the scales of interest here it may vary quite a bit.

3.2.1 Gradiometer Sensitivity

Gradiometer developers typically quote sensitivity for their instruments in some fractions of an Eotvos. The developers do not explicitly describe a noise model for the instrument, but a lower limit on sensitivity implies that there is a noise floor only just below the quoted sensitivity level. In the




Figure 1. Shape of the gravity gradient tensor T_{xy} for a point mass. The x and y axes are in kilometers. The z axis is dimensionless.




Figure 2. Shape of the gravity gradient tensor T_{xx} . The x and y axes are in kilometers. The z-axis is dimensionless.




Figure 3. Shape of the gravity gradient tensor T_{yy} . The x and y axes are in kilometers. The z-axis is dimensionless.




Figure 4. Shape of the gravity gradient tensor T_{xz} . The x and y axes are in kilometers. The z-axis is dimensionless.




Figure 5. Shape of the gravity gradient tensor T_{yz} . The x and y axes are in kilometers. The z-axis is dimensionless.

simplest case, we can assume that each gradient measurement is corrupted by independent Gaussian noise with standard deviation somewhat below the instrument sensitivity level. How much below depends on our specification for detectability in the presence of noise. The exact definition is unimportant in the analysis we are doing here.

3.2.2 Natural Background

Since the Earth is a solid body, it would seem that the gradiometric field of the Earth must be constant. But, at the scales of interest here, it is clear that it is not. For this analysis, we need to distinguish between temporally constant, spatially varying natural gradiometric noise, and temporally varying gradiometric noise. If we move a gradiometer along a line of constant radius from the center of the Earth, we will see a great deal of variation. If the Earth were a perfect sphere, we would see no variation, except for the variation due to the rotation of the coordinate system. However, the Earth is not a sphere, and its surface is quite rough. As we move along a constant radius path we will see all the variation due to surface and subsurface topography.

If we were able to observe a single point for an extended time we would also see temporal variation in the gradient. Some sources for this variation include astronomical (movement of the moon and Earth relative to the sun), geophysical (movement of geological features in micro-earthquakes and mantle convection), weather and climatic (movement of the water table, density changes in the soil, Earth movement by erosion), biological (growth of trees), and human (vehicle traffic, construction). A significant risk area for this concept is that we have no models for this kind of gradiometric variation. It is quite possible that it is large compared to the signals of interest, but we have no means at present to observe it and find out. Our lack of a model for this noise source is a fundamental restriction on detection against unknown backgrounds.

There is another potential source of environmental noise that is not based on temporal variation. The actual gradiometric field will vary spatially with some unknown frequency content. The spatial variation of the field will have at least some high-frequency components from things as simple as rocks embedded in sand. These features nominally form part of the stationary gradiometric field and would not be considered noise. However, our sampling process is spatial. We do not sample the gradiometric field at every point in space, we only sample on a discrete grid with spacing determined by operational and instrumental issues. If the gradiometric field contains energy at spatial frequencies above the sampling limit for the grid, that energy will be aliased just as are undersampled components of any signal in conventional signal processing situations. Since many of the operational scenarios are looking for effects that are extremely small, both absolutely and relatively compared to the background field, even very small amounts of aliased spatial noise could be critically important. Again, we have no data on which to evaluate the size of this effect.

Some simple arguments lead one to expect that the "spatial noise" effect will be much more important in close-in measurement situation than from space. The characteristic spatial extent of the field resulting from an object is roughly the distance from the object to the sensor. In close-in (ground-based) scenarios, the grid spacing will often be larger than the distance from the sensor to massive objects. In this case, there is likely to be considerable aliasing of spatial variation. In space applications, the minimum distance from an object to the sensor will be several hundred kilometers. Grid spacings in space can be made much smaller than several hundred kilometers, eliminating the aliasing, although this may be difficult for particular sensor cases. For example, some sensor proposals use two satellites spaced about the altitude apart may have particular problem with this case since the two satellites may see parts of the gradiometric field with fluctuation due to the nearest objects.

3.3 Size of the Effects

For gradiometry to be possible we must be able to detect the distortion to the background gradiometric field produced by the cavity of an underground facility. One effect limiting detectability is noise, which was described previously and will be analyzed in more depth in later sections. A second limit is dynamic range. A gradiometric instrument must be able to detect the distortion of the cavity against the background of the Earth's field. Ideally, the instrument should have a linear response for gradiometric fields from that of the cavity all the way to the Earth's background. This can be estimated quantitatively from the ratio of the fields. If the Earth is approximated as a spherical body, its gradiometric field is just that of a point mass with the mass of the Earth, m_e . Let the coordinate system have its origin at the instrument location and the z-axis run through the center of the Earth. Then the xx-component of the gradiometric field due to the Earth is:

$$T_{e_{xx}} = Gm_e \frac{(x^2 + y^2 + (z + R_e + \text{alt})^2 - 3x^2)}{(x^2 + y^2 + (z + R_e + \text{alt})^2)^{5/2}},$$

where R_e is the Earth's radius.

The change to the gradiometric field due to the facility is the same as the field produced by a facility size and shape mass of rock at the facility location surrounded by a vacuum. As an approximation, to get a ratio of effects, let the facility be a point mass with mass equal to the reference volume (from above) times an average density of rock, ρ . Relative to the coordinate system it will be located on the z-axis at the altitude of the sensor plus the depth of the facility.

$$T_{f_{xx}} = GV_f \rho \frac{(x^2 + y^2 + (z + \text{alt})^2 - 3x^2)}{(x^2 + y^2 + (z + \text{alt})^2)^{5/2}},$$

For an estimate, let x, y, and z equal zero, and evaluate the ratio between the two terms.

$$\frac{T_{e_{xx}}}{T_{f_{xx}}} = \left(\frac{m_e}{V_f \rho} \right) \left(\frac{\text{alt}^3}{(R_e + \text{alt})^3} \right)$$

Using the reference values (the reference tunnel), and letting the altitude range from 100 m to 1000 km, yields the plot of Figure 6.

So, a gradiometer needs to detect changes of one part in 10^{10} to one part in 10^{15} to see a reference tunnel. The dynamic range is correspondingly larger if one is trying to discriminate individual voxels instead of the entire tunnel element. One part in 10^{12} level accuracy is certainly challenging, but does not appear to be beyond the capabilities of emerging instruments. A more significant problem may be instrumental linearity over this range, which is very large compared to most instrumental applications.

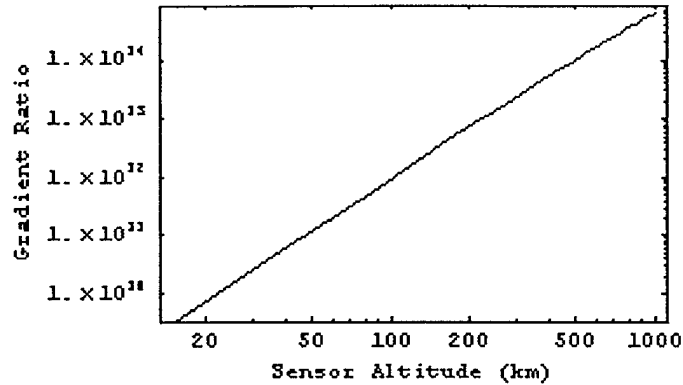


Figure 6. Ratio of gradient terms as a function of sensor altitude.

Inversions and detection against unknown background algorithms may be sensitive to instrument linearity. They are sensitive to linearity in other remote sensing applications. This issue is not investigated in this report, although it may be significant in applications of interest. It should be part of further research efforts on algorithm design.

3.4 Detection Models

What does it mean to say that we can detect a particular event by observing a signal? Mathematical detection theory provides a precise answer to this question. We assume that we collect a signal, s , and put the signal through a processing algorithm, $P(s)$. $P(s)$ yields a binary declaration, either the event of interest is present or it is not. The signal s is stochastic and either has statistics S_0 or S_1 . Statistics S_0 occur when the event of interest is not actually present. Statistics S_1 occur when the event of interest is actually present. This yields two figures of merit. The first, the probability of detection, is the probability that $P(s)$ will declare the event present when s has statistics S_1 . The second, the probability of false alarm, is the probability that $P(s)$ will declare the event present when s has the statistics S_0 . Optimal detection theory is concerned with how to design $P(s)$ to produce the maximum possible probability of detection for a given probability of false alarm. Obviously, a degenerate P can produce a probability of detection of one by always declaring the event, in which case the probability of false alarm is one as well.

3.5 Inverse Gravimetry and Gradiometry

Inverse gradiometry (or gravimetry) is the estimation of a mass-density field from a set of gradiometric (or gravimetric) measurements corresponding to that field. People often make the analogy to computed tomography, although the analogy is weak. In tomography, we measure the X-ray transmission function of an object along a circle circumscribing the object. We can then estimate the density field of the object from the measured transmission function. The problem of estimating a field of interest from measurements derived from that field is known in mathematics as "inverse problems." Inverse problems are very widely studied in many subfields.

The mathematical structure of the inverse gradiometry problem is straightforward. We have a density function, ρ , and a bounding region, V (the Earth/atmosphere surface). The function ρ and the volumetric region V have the constraints:

$$\begin{aligned}\rho(x, y, z) &\geq 0, \{x, y, z\} \in V \\ \rho(x, y, z) &= 0, \{x, y, z\} \notin V\end{aligned}$$

The mass density function ρ creates a gravitational potential function U throughout space. Here we assume Newtonian gravity for U . In real situations, it is possible that relativistic corrections would be necessary to achieve the required level of precision.

$$U(x, y, z) = \iiint_V \frac{G\rho(x_1, y_1, z_1)}{\sqrt{(x-x_1)^2 + (y-y_1)^2 + (z-z_1)^2}}$$

The gradiometric field is found by taking double derivatives of the potential function. Each gradiometric tensor component is given by

$$T_{ab} = \frac{\partial^2}{\partial a \partial b} U(x, y, z).$$

The complete equations were given previously, and plotted for each tensor component.

The inverse gradiometric problem can now be precisely stated. Suppose we measure one or more of the gradiometric tensors, T_{ab} , in a region M (not all of space, and usually on a non-uniform grid of points). How do we estimate the function ρ from those measurements? In general, this problem is ill-posed and does not have a unique solution. A simple thought experiment shows why. Suppose the body of interest is a perfect sphere of uniform density. The sphere can be shrunk to any radius (while maintaining the same mass), and it will have no effect at all on the gravitational field in space outside the radius of the original body. This is illustrated in Figure 7. Thus, no inverse algorithm can uniquely recover the radius of the perfect sphere since it is unobservable in the collected data.

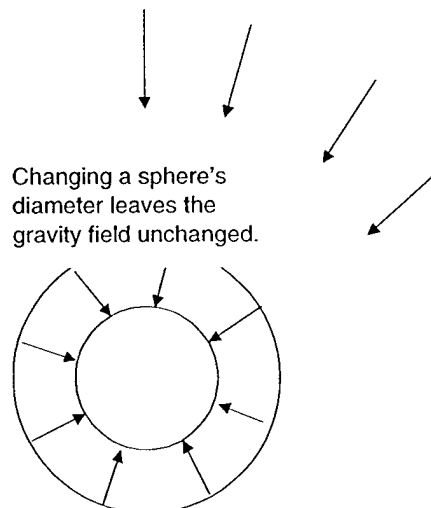


Figure 7. Many physically distinct objects have the same external gravitational field.

The real inverse gradiometric problem is somewhat more complicated. The real problem has many additional constraints. The size and shape of the Earth are known (although not exactly) and irregular. The known information about size and shape can be utilized by an inverse algorithm to reduce or eliminate ambiguity. There is potentially also information on surface and subsurface geology that further constrains the mass-density function.

3.6 Matched Filter Concepts

In many detection theory situations (though not all), it turns out that optimal performance can be achieved by thresholding the measured signal after passing it through a matched filter. A matched filter can be thought of in two ways. The simplest version is a filter that duplicates the signal produced by a target of interest. Put another way, in a matched filter, we correlate the measured data with the data that would have been produced had a target of interest been present with no noise or other clutter. The more complex version duplicates the signal of interest while canceling the interference background. The simple version maximizes the signal-to-noise ratio with the assumption of a white, Gaussian noise background. The more complex version maximizes the signal-to-interference ratio, taking into account all that is known about the statistics of the interference.

In the discussion to follow, we use the simplest possible target for the matched filter, a matched filter corresponding to a point mass at a given location and assume white, Gaussian noise. This point mass produces a gradiometric field, which we correlate with the actual, measured gradiometric field. The correlation is a measure of how much mass is near the given location of the matched filter mass. Where the correlation is large, we estimate a high mass density; where the correlation is low, we estimate a low mass density. For the case of inverse volumetric imaging, the specific equations are given in Section 7.2.

While this very simple matched filtering is not always optimal, it provides a useful benchmark on performance. Since one can always do at least as well as matched filtering, a result that shows that matched filter achieves a desired performance level is a demonstration of feasibility. If we need performance much better than the matched filter provides, we can estimate feasibility and difficulty from the amount of improvement over simple matched filtering.

An important question is how much performance might be improved if we could compute a true signal-to-interference-ratio maximizing filter. This is particularly germane to the case of detection with an unknown background, where it is discussed at greater length.

4. Detection with Known Background

The simplest detection case is where the background is known. That is, we are looking for new underground facilities when gravitational data on the site before the facility was constructed is available. Of course, this is an optimistic case. However, it has been proposed and used for some performance estimates and so should be considered here in more depth. A more realistic version of this case is where a spaceborne collector generates the baseline gravity data itself, and we then look for changes in the collected baseline.

To make the analysis simple, we will consider a one-dimensional, one tensor component version. Let a satellite orbit the Earth on a path with constant radius from the center of the Earth. Real orbits are like this, but the issues in how orbiting in a non-uniform field affects measurement and inversion are too detailed for this stage. While orbiting, the sensor periodically measures one component of the gradiometric field.

To make a performance model, we need to define a few functions and terms:

X is the linear distance along the orbital path

$T(x)$ is the true, non-varying, baseline gradient tensor value along the orbital path

$Ne(x)$ is the environmental gradient noise along x (e.g., weather, astronomical, geologic, etc.)

$To(x)$ is the gradient tensor due to the moved mass of the underground facility, as induced by the actual movement of mass to build the facility.

$Ns(x)$ is random noise in the gradient sensor from all sources

Then, the sensor actually produces $T(x) + Ne(x) + Ns(x) - To(x)$, and has $T(x)$ available for reference. Our challenge is to determine whether $To(x)$ is present or absent. It is important to understand the uncertainties and the randomness in order to apply detection theory. In this case, we are assuming that $T(x)$ is known and not random (obviously, this is somewhat unrealistic). In the real situation, we would have to estimate $T(x)$ from many measurements (and possibly side information). The error in estimating $T(x)$ would become another noise term in the analysis. The sensor noise term, $Ns(x)$, can probably be easily characterized. We can probably assume that it is additive, white, Gaussian noise, at least white on the spacing of our sensor's samples. Depending on the sensor design, it might or might not actually be Gaussian, and might not be white. Presumably, we know its statistics since the sensor is under our control. The magnitude of this noise term determines the ultimate sensitivity of the instrument, so we can readily quantify it from the sensitivity estimates of the various investigators. The environmental noise term, $Ne(x)$, is probably the least understood. It is clearly not zero, but we have little data to estimate how it might effect gradient measurements on the very tiny scales on which we want to work. Some sources are known and (in principle) removable, like Earth tides. In

other cases (e.g., micro-earthquakes, water table change, etc.), we do not know much about their potential impact.

In the studies so far, $T_0(x)$ has been assumed to have the shape of a point mass, with magnitude determined by the excavated volume of the facility. For the spaceborne case these assumptions are only partially correct. Assuming the shape to be that of a point mass is appropriate since the height of the sensor (200 km or more) above the facility will erase all shape due to internal structure. However, giving the facility a mass equivalent to the excavated volume is appropriate only if the excavated material has been transported a distance larger than the sensor height (typically more than 200 km). When the mass has not been so transported, then the net change to the background will be less than indicated by the volume of the facility. The mass deficit of the facility will be “screened” by the increased mass of the excavation materials deposited close by. As Clauser describes this, the excavation materials destroy the expected monopole field. In Sections 7.6.3 and 8.3, we show some quantitative results for inversion in this case demonstrating the effect.

4.1 Simple Detection

Before taking up these complexities, we’ll compute the simplest case. In the simplest case, we look only at the peak deviation of the facility-induced gradient from the background. For a satellite sensor in 200 km orbit and a low-flying aircraft at 1 km altitude, the deviation, in Eotvos, for a reference tunnel taken along one axis is shown in Figures 8 and 9.




Figure 8. Gradient deviation for a reference tunnel seen from a 200-km orbit.




Figure 9. Gradient deviation for a reference tunnel seen from an aircraft at 1-km altitude.

In the spaceborne application, we have to reliably detect deviations in the range of 10^{-9} Eotvos, as illustrated in Figure 8. From the air, the deviations are much larger, on the order of 100 mEotvos.

Without knowing how large the uncompensated random background is we can't make a quantitative estimate of detection and false alarm probabilities for a given scenario. However, we can set some bounds based on the geometry. There is little point in making detection decisions on a spacing finer than the width of the signal one wants to detect. One may sample the data more frequently, and then apply a matched filter, but the fundamental grid for detection decisions is roughly the width of the signal. For a spaceborne application take 200 km as the detection decision spacing. The surface area of a 200-km orbit is about 5.4×10^{14} m². This is about 13,600 patches 200 km on a side. Of course, most of the Earth is of no interest for searching for underground change characteristic of human construction. Probably only a few percent is important, meaning that we are interested in change in a few hundred patches. Thus, we will want a false alarm probability around 10^{-2} so detections are typically real (rather than biased toward most hits being false alarms). To get this false alarm probability, the detection threshold will have to be about 2.3 standard deviations above the noise level (assuming Gaussian noise in the aggregate from all sources). This implies that to reliably detect the reference tunnel from space, the net of all uncompensated interference sources (temporal variation in the gradient field, sensor noise, uncompensated astronomical effects, geodetic variation, etc.) will need to be brought below 10^{-9} Eotvos.

Some instrument builders think it may be possible to build an instrument that achieves this, but we currently have no knowledge of whether or not the environmental effects can be compensated for to this level. At present, there is simply no data to support or reject the possibility. Existing measured datasets are about 9 orders of magnitude too coarse to determine this.

4.2 Structured Discrimination

The simple statistical model for the signal and noise can be used to understand how sensor height affects the ability to discriminate. Consider the following, highly idealized, facility characterization problem. Suppose we know that the facility is composed either of a single point mass or two point masses (with the same combined mass as the single point mass) separated by a distance of one. There is no other background of mass to make the problem any more complicated. This is an extremely idealized version of the problem of discriminating two tunnels from a single tunnel of the same total volume. To measure the ability to discriminate, we compute the optimal discriminator by the method of maximum likelihood.

We assume that everything about the situation is known (geometry, masses, separation, etc.) and the only unknown is whether or not the object of interest is two masses or one. Let the noise model be additive, white, Gaussian. The rms value of the noise will be set to a fixed fraction of the peak gradient for the single mass. We measure the T_{xx} tensor along a line extending twice the sensor height on both sides of the masses of interest. Our performance metric is the probability of error for the maximum likelihood test on one tunnel versus two. Given these highly idealized assumptions, we can analytically compute the probability of error using standard methods in detection theory.

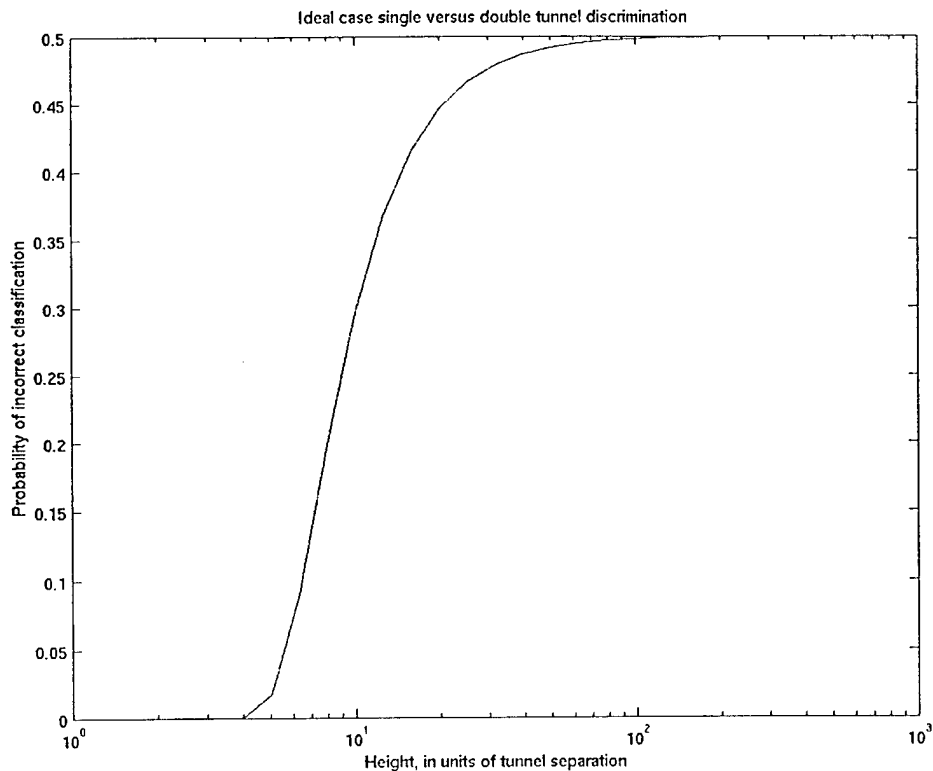


Figure 10. Probability of error curve for two-tunnel discrimination as a function of sensor height.

For 50 points along the measurement line, and a noise floor one twentieth of the peak gradient, the probability of error as a function of sensor height is shown in Figure 10.

If we drop the noise floor to one thousandth of the peak gradient value, the curve improves to that shown in Figure 11.

It is important to note that this is a very high signal-to-noise-ratio. This is the SNR calculated for the whole mass of interest at the sensor distance. For the reference tunnel at 200 km, this implies a noise floor of 0.002 nEotvos, which is almost certainly not achievable. In spite of this very low noise floor, the ideal discriminator cannot distinguish (with high probability) two tunnels from one when the distance is more than 50 times the separation. For a sensor height of 200 km this implies that facility components with spacings less than 4 km could not be distinguished. For the more reasonable SNR the limit is more like 40 km.

If the number of measurement points is increased, the ability to discriminate also increases, but only slowly. For example, if the number of measurement points is increased by a factor of 10 (from 50 to 500) in the measurement region used for Figure 11, the maximum altitude at which the separation can be distinguished increases by only about 20%. This is shown in Figure 12. In practice, this is likely to be extremely idealized. The measurement density cannot be increased on single satellite passes because it is limited by the physics of the device. As more passes are used to increase the measurement density, one also introduces more environmental noise terms.




Figure 11. Probability of discrimination error with the SNR raised to 1000.

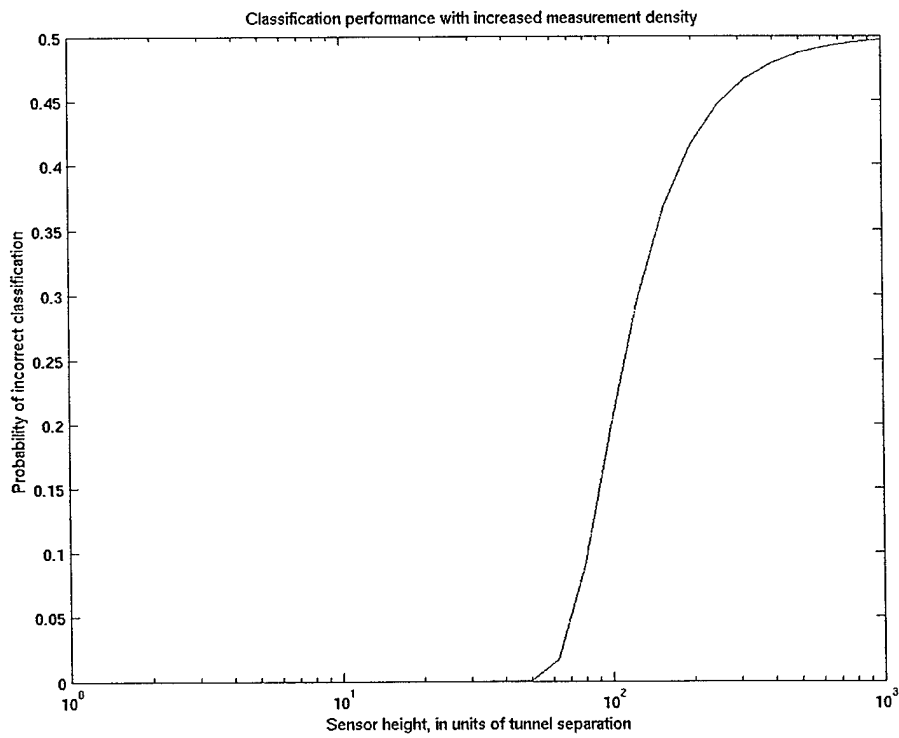


Figure 12. Classification performance with high SNR and increased measurement density.

5. Detection with Unknown Background

The previous section analyzed detection when the background is known. The reality of the ISR problem is that the background is not known. Known background detection is only applicable to change detection. In reality we don't have ground truth on where underground facilities are or are not. Being able to detect a new underground facility may be of interest, but a very large number of facilities will exist before we begin any program of high-precision gradiometric measurement. We also don't have a way of precisely computing what the natural gradiometric field should be. We have to use measured data to determine the natural background. When we use measured data to estimate the background, it is possible that a facility already exists there, and that we are including it in the background. Moreover, much of the concern over underground facilities is related to characterizing them, not to detecting them. A variety of observables can show that a facility is present, but cannot nearly as easily reveal how large the facility is, its configuration, or for what it is being used. So the more realistic problem of detecting a pre-existing facility is to detect an underground facility when the natural background is unknown.

If the background is estimated from repeated measurement, then it can be treated as in the previous case, except with an additional noise term due to estimation error. We might consider using the case when there is some other source of information indicating that a facility will be built, but has not yet been started. Alternatively, it might apply in detecting changes, including battle damage assessment. But, when the background is truly unknown, as in wide area search, we must detect some signature characteristic of a facility of interest while rejecting false detections due to background that may have a similar signature.

In detection theory the essential step is to compute the conditional probability densities $p(m | \text{facility present})$ and $p(m | \text{facility not present})$, where m is the measured data. In conventional cases, like radar or infrared detection, we have Gaussian or other statistical models for the measured data in the presence and absence of targets. In this case, we don't have a statistical model for the data, with or without a facility being present. Lacking that statistical model, we cannot form the conditional densities and do not know what to threshold.

One approach is to assume nothing about the background. This is essentially what we do in the section on inversion with target-only matched filters. In formulating the matched filter for a point mass at a given location, we assume nothing about the background of masses at other locations or about temporal or spatial noise. We simply correlate to the point mass and compare it to the correlation with point masses at other locations. The relative sizes of the correlation is used as an estimate of mass density. If we were using this scheme for detection, we would have to estimate noise and clutter by looking at spatial variation in estimated density. We would declare an underground facility to be present when we found a group of volume pixels with estimated densities much lower than their surroundings, and physical locations under the known ground profile. This is conceptually similar to constant-false-alarm-rate (CFAR) systems in radar that use estimation of the actual received data to estimate clutter and noise levels, and thus set detection thresholds.

The performance of a system doing detection against an unknown background is entirely dependent on the characteristics of the clutter. A better approach would be to form a signal-to-interference-ratio maximizing filter that uses all that is known about the background clutter statistics. If the background statistics were favorable then it is possible that the filter would have excellent performance, even when the clutter power level was much higher than the signal level. For gradiometry, we have no data at the scales and sensitivities of interest; hence, we cannot make any quantitative models of performance. We could speculate on background statistics and correlation lengths (both spatial and temporal), but there is no data to allow one speculation to be preferred over another. Therefore, it is not useful to further explore modeling the clutter at this time, although it is of considerable practical importance.

Because of the importance of this issue, we need to look for arguments that would suggest whether the background is likely to be favorable for removal or not. The situations where removal usually works well are where we have multi-dimensional signals with a lot of dimensions. A simple example is frequency discrimination. If a signal is on one frequency and the interference is spectrally concentrated on other frequencies (and we have enough sampling time), then a filter can efficiently null the interference and retain the signal. The same principle can apply in other signal spaces where the dimensionality is obtained in frequency or space.

Unfortunately, space applications of gravity gradiometry are unlikely to match this model. In the space case, as will be shown, virtually all underground facilities will only yield on one independent spatial sample. With an ideal detector, we might get six tensor components, but are more likely to get only one (given the analysis of Clauser). Thus, at best, we have only a few signal dimensions on a target to work with in trying to find a discriminating filter. In contrast, for an airborne or close-access case, we might have 10's to 100's of independent measurements across a target and six components on each one. This provides a rich signal space in which to look for signal decompositions in which the signal projects into subspaces with little clutter energy.

6. Detection with Complex Signatures

Using target signatures with more information can refine the detection process, specifically, more information about how measurements with valid targets differ from measurements with only natural background. Most underground facilities are built from tunnels, which are long thin hollow spaces. A long thin hollow (or, symmetrically, a long thin mass) has a fairly distinctive gradiometric field. It is possible that one could match filter to the field of a long thin object instead of a point mass and improve performance. Some simple calculations show this may have merit in close-in measurement situations, but is unlikely to be useful in space applications.

The gradiometric field due to a tunnel can be approximated by the field of a line mass. A line mass is an assemblage of point masses, so we can compute the field of a line mass by integrating the field due to a point mass along a finite line. The equation for the TL_{ab} (a and b are any combination of x, y, and z) tensor field of a line mass of length L making an angle θ to the x-axis is given by:

$$TL_{ab}(x,y,z) = \int_{-L/2}^{L/2} T_{ab}(x-t*\cos(\theta), y-t*\sin(\theta), z) dt$$

If we plot the resulting field, we see it has a clear structure. What are particularly striking are the sharp peaks near the ends of the line mass, for the close-in measurement cases. Figure 13 shows one component for very close-by collection.




Figure 13. T_{xx} for a line mass measured at a distance of one-twentieth its length.

Unfortunately, this distinctive structure disappears when the observation distance is large compared the size of the tunnel. For the same tunnel, now observed at 50 times its length, the T_{xx} -tensor field now looks like Figure 14.

So, it is apparent that looking for complex signatures correlated to the structure of tunnels may have merit in close-in, or even airborne measurement applications, but will not improve the situation for space measurement.

Clauser suggests that we may be able to use dipole components of the field to make detections where there is no net mass deficit within a horizontal region roughly equal to the sensor height. However, he shows that this signature is small, and detecting it would be very challenging even under the most favorable circumstances.




Figure 14. T_{xx} field for a line mass measured at 50 times its length.

7. Inversion and Imaging with Matched Filters

The previous discussion focused on detecting underground facilities using gravity gradients. However, the tomographic analogy suggests imaging them, as in a volumetric image. The ISR value of inversion to a volumetric image is greater since it allows a facility to be characterized. For example, in principle it allows one to determine size, internal layout, and the presence or absence of machines, and to assess damage after an attack. To get to a volumetric image we need to attack the inversion problem directly. This section discusses one simple method for solving the inverse gradiometry problem, the method of matched filtering.

Matched filtering is not optimal for this problem. Theoretically, we don't know how far from optimal, or whether or not the concept of optimal processing has any meaning in this case. Nevertheless, the matched filter method is illuminating of the whole problem. First, if matched filtering proves to produce adequate volumetric images, then we have provided an existence proof that a useful inversion algorithm exists. Second, the performance of matched filter algorithms is frequently a kind of benchmark for the performance of other, more sophisticated algorithms. We can measure the performance of the matched filter algorithms and estimate how much of an improvement will be required by the more sophisticated algorithms to yield desired performance. If the margin of improvement is very large, we will know that success is unlikely.

7.1 Reference Geometry Model

The geometry model is the same as before, only in this section we make it explicit. The gradient sensor measures the gravity gradient in some region above the Earth. The coordinate system is rectangular and centered at the center of Earth. An arbitrary vector in the coordinate system is \mathbf{r} . A vector \mathbf{s} represents a vector to a point at which the gravity gradient is sampled. The vector \mathbf{r}_0 designates a reference point inside the area of the facility that we wish to image. The vector \mathbf{d} is a displacement from \mathbf{r}_0 . Our concern is how masses at points $\mathbf{r}_0 + \mathbf{d}$ affect the estimate of the mass density at \mathbf{r}_0 . This is conceptually the same as how we use point-spread functions in optical imaging or ambiguity functions in synthetic aperture radar to evaluate imaging performance. The geometry is illustrated in Figure 15.

7.2 Simple Matched Filter Processing Model

The processing model can be concisely stated. We start with the set S of points in the measurement region where we sample the gradient tensor.

$$S = \{\mathbf{s}_i : \text{each a vector in the measurement region}\}$$




Figure 15. Geometry used for evaluating gradiometric processing.

The measurement set M is the set of tensor measurements taken at each point in S . In general, M is a set of vectors or even matrix quantities. For this problem, each element of M is a vector with length up to five.

$$M = \{m_i : \text{The tensor measurement at } s_i\}$$

A reference function G is the set of tensor measurements at each point in S that would be made if we had a point mass at an arbitrary point r . The function T refers to whichever tensor component functions one had decided to use. The functions for point masses were given in Section 3.1.

$$G(r) = \{T(s_i - r) \forall s_i \in S\}$$

Put another way, G is the measurement set we would expect if we placed a unit point mass at r . Now the processing model for estimating the mass density field is trivially simple. To make an unnormalized estimate of the mass density at a point r , take the dot product of the reference function for r with the actual measurement set.

$$\rho(r) = G(r) \bullet M$$

This estimate is unnormalized. It can be used to find the relative density between two points, but does not produce a normalized actual density for a given point. This simple matched filter also assumes equal weighting on all five tensor components, which may not be desirable.

This method has an important desirable property. Using matched filters, we need only compute the mass-density field estimates for the points of interest. The estimate for the density at one point is not affected by the estimate for the density at another point. To be sure, the estimate at one point is affected by the actual density at another point, but the estimates of density at different points do not affect each other. This is important because it means we need to compute only the points we want. To contrast, a globally dependent algorithm would need to compute the density estimate everywhere to get the mass density at any individual point. This would mean that to compute a volumetric image of a facility on a 10-m cubed grid, we would have to compute the mass-density field of the entire Earth on a 10-m cubed grid, clearly undesirable and probably infeasible.

A truer matched filter needs both a signal model and an interference model. Since that formulation is not used here, its details are not given. In outline, we set up an additive signal model with the signal (known) added to an interference vector (whose statistics are known). Then we compute a set of weights that maximizes the ratio of the summed signal parts to the summed interference parts. This is normally given by a product of the signal vector and a set of weights that whitens the interference. Since the exact interference statistics are not usually known, but must be estimated from the received data, a practical algorithm must account for the joint estimation. There are many methods in the literature that do this.

The question of interest is whether or not such an approach would yield good performance for a detector looking for pre-existing underground facilities. Unfortunately, we can't analyze that question because we have no data on the real interference statistics. Any assumption (Gaussianity, correlation length, spectral content, etc.) will only be a guess. An analysis of this issue will have to wait for the availability of data. There have been some ground-based gradiometric surveys at roughly the 1-Eotvos resolution level. Those surveys suggest there is considerable spatial variation (clutter). Since there is evident clutter at the 1-Eotvos level it is almost certain that the clutter only becomes worse as sensitivity levels go milli-Eotvos or micro-Eotvos. The datasets collected do not extend over large enough geographic areas to suggest whether or not the clutter has subspace properties that distinguish it from underground facilities.

7.3 Volume Spread Function Equations

We can compute the volume spread function by picking a reference point (relative to the measurement region where a facility of interest would be), computing a reference function for it, using that as the measurement set, and then recomputing reference functions for points near the reference point. Mathematically, we write this by starting with a reference function at the reference point \mathbf{r}_0 .

$$G_0 = G(\mathbf{r}_0)$$

The volume spread function $P(\mathbf{d})$ is then calculated by taking dot products of G_0 with the reference function for other nearby points. The volume spread function is a function of the displacement vector \mathbf{d} from the reference point \mathbf{r}_0 .

$$P(\mathbf{d}) = G(\mathbf{r}_0 + \mathbf{d}) \cdot G_0$$

If desired, the volume spread function can be normalized by dividing by the dot product of G_0 with itself. This guarantees that the function reaches a peak of 1.0 at the point \mathbf{r}_0 .

7.4 Matched Filter Point Spread Function

The important performance question for gradiometric inversion by matched filters is how well it performs as a function of operational parameters. In particular, we want to know how different tensor components perform, how the function is affected by the minimum altitude of the measurement region, and how the function is affected by the size and shape of the measurement region.

7.4.1 For Different Tensor Components

We start by looking at how the function is affected by each of the five tensor components. First, for reference, we have the volume spread function for all five tensor components equally weighted in Figure 16. The altitude is 100 km; the measurement region is 1500 km on a spherical patch centered above the reference point.

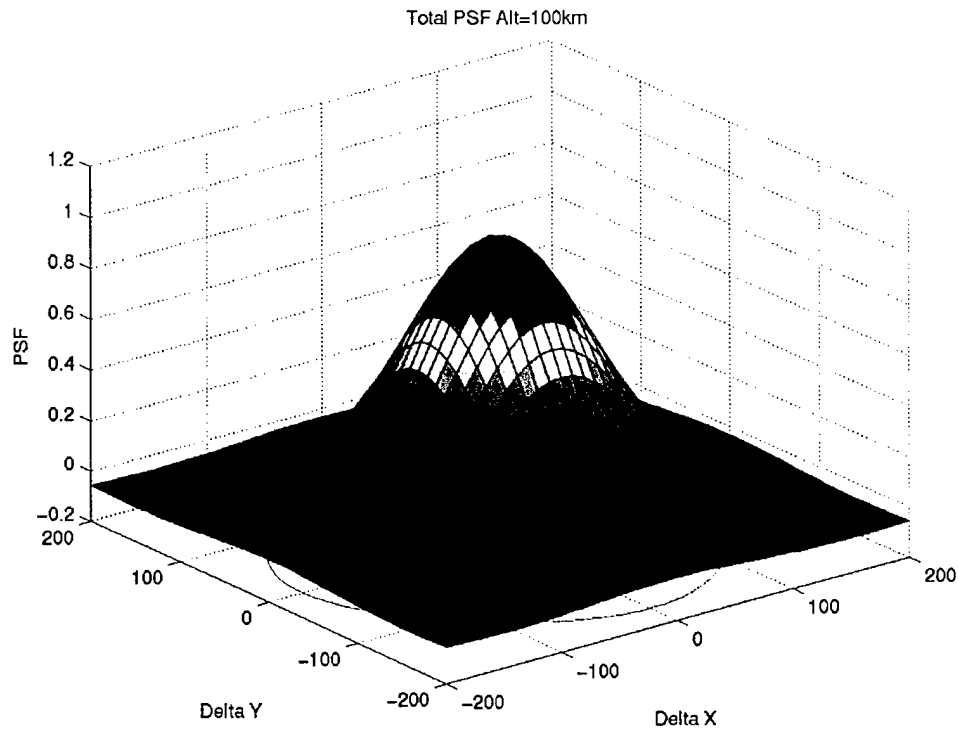
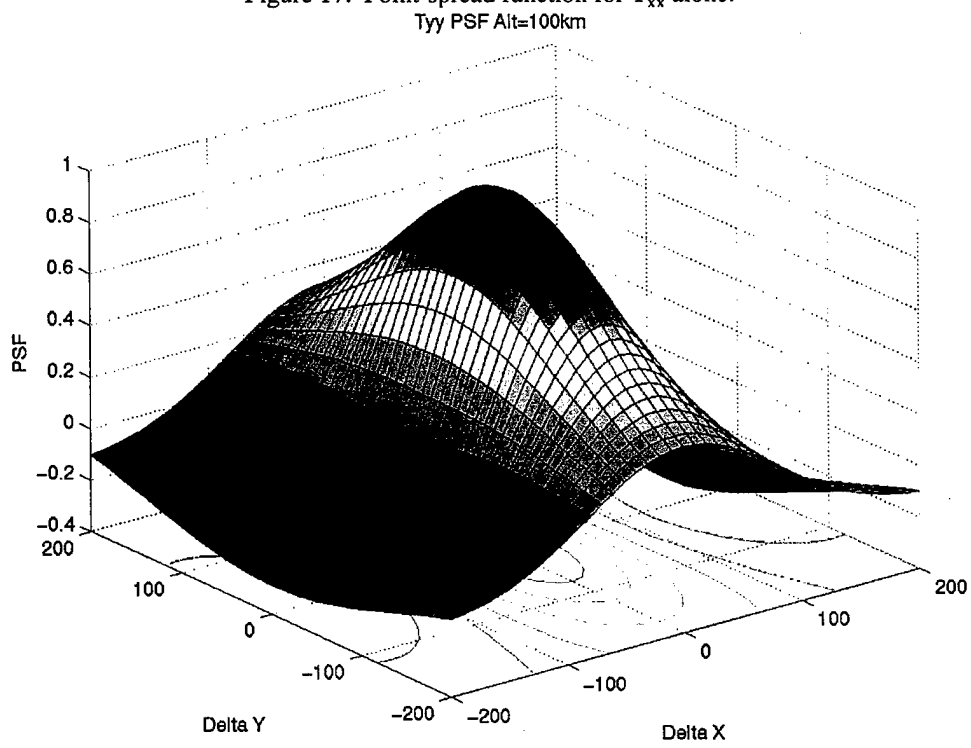
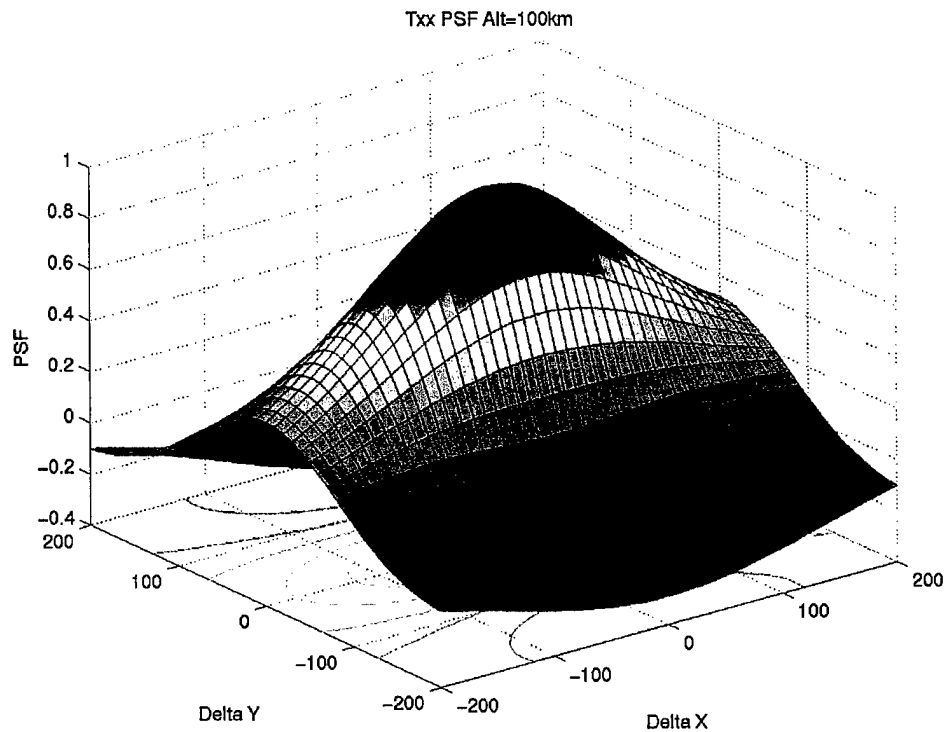


Figure 16. Five-component point-spread function.

Now we show the total function for each of the tensor components individually in Figures 17–21.






Figure 19. Point-spread function for T_{xy} alone.

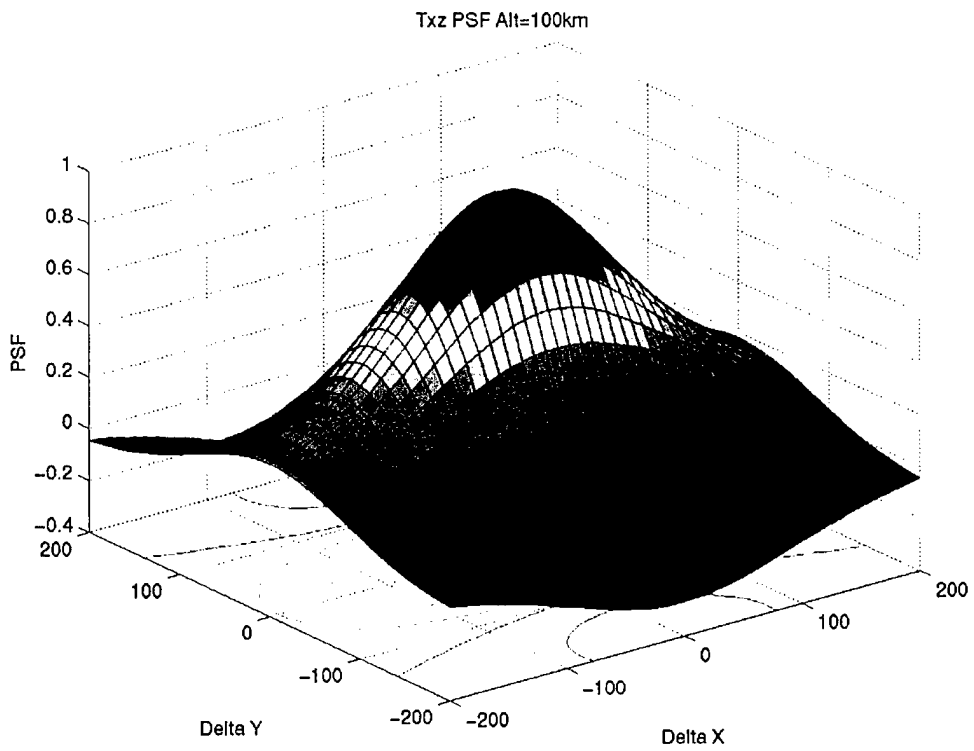


Figure 20. Point-spread function for T_{xz} alone.




Figure 21. Point-spread function for T_{yz} alone.

These can be more easily compared quantitatively by looking at a $y = 0$ slice through each all plotted on the same graph. The five components are compared in Figure 22.




Figure 22. X-axis cut through the point-spread functions.

Although it appears that T_{xx} and T_{xz} are narrowest, the previous figures in both x and y show that their shapes are non-symmetric. All five components have roughly similar performance, but not in the same dimensions. Equal weighting, as in Figure 16, may not be optimal for particular resolution goals, but is reasonable if all three dimensions are equally important in resolution.

Since this is really a volume spread function it also has z-direction performance. The function also has fall-off with depth. The shape of the function in the z-direction is illustrated in Figure 23, though in this case for a measurement altitude of 10 km rather than 100 km as in the previous figures. As the next section will show, this is easily scaled for measurement region altitude.

7.4.2 As a Function of Altitude

An essential question in deciding whether gradiometry is suitable for space application is how its performance scales with the altitude of the measurement region. There are two important effects. First, there is the impact on resolution. Second, there is the impact on sensitivity and noise floor. We deal first with the impact on resolution. If we lower the altitude of the measurement region to 10 km from 100 km, we get Figure 24.

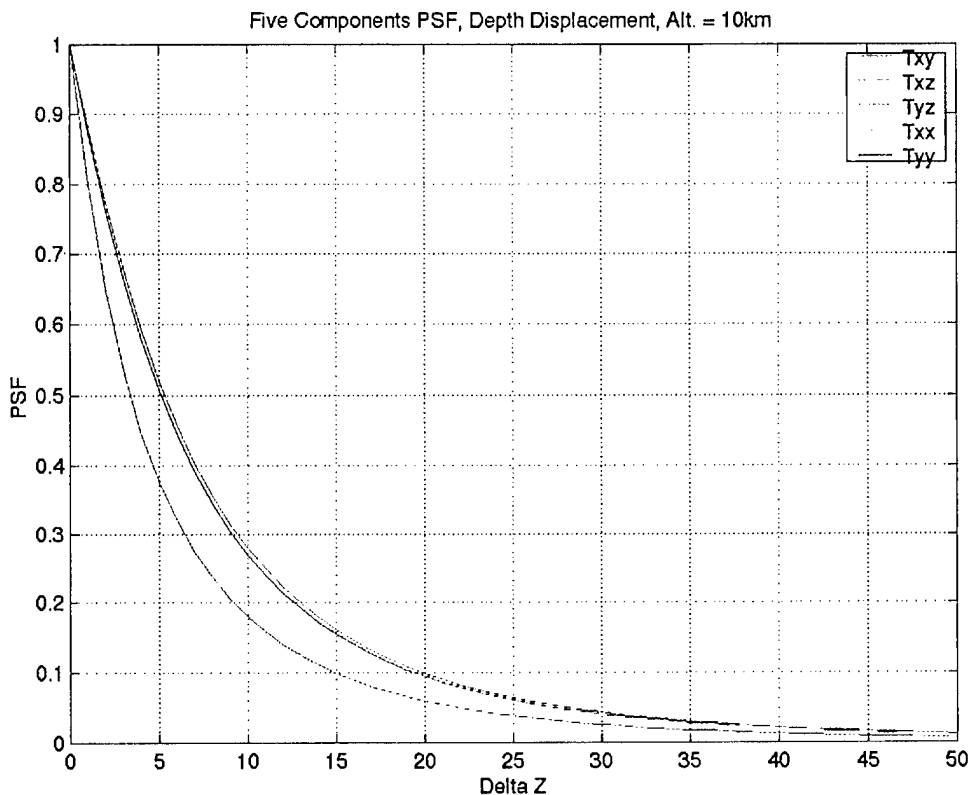


Figure 23. Z-axis cut through the point-spread functions.

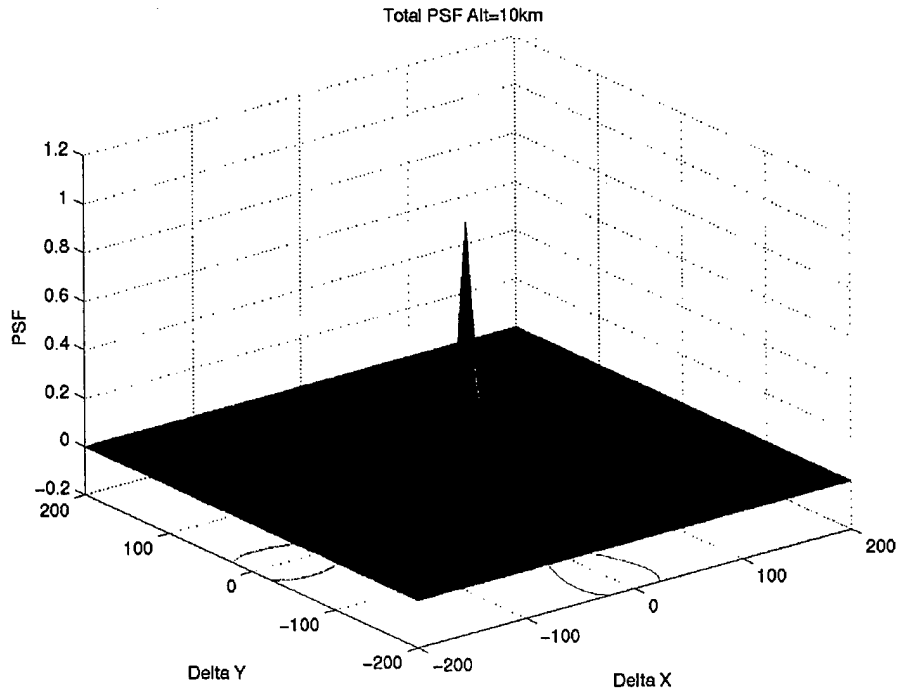


Figure 24. The point-spread function shrinks directly with sensor altitude.

Obviously, as the altitude reduces, the width of the function becomes narrower. Quantitatively, the relationship is shown in Figure 25.

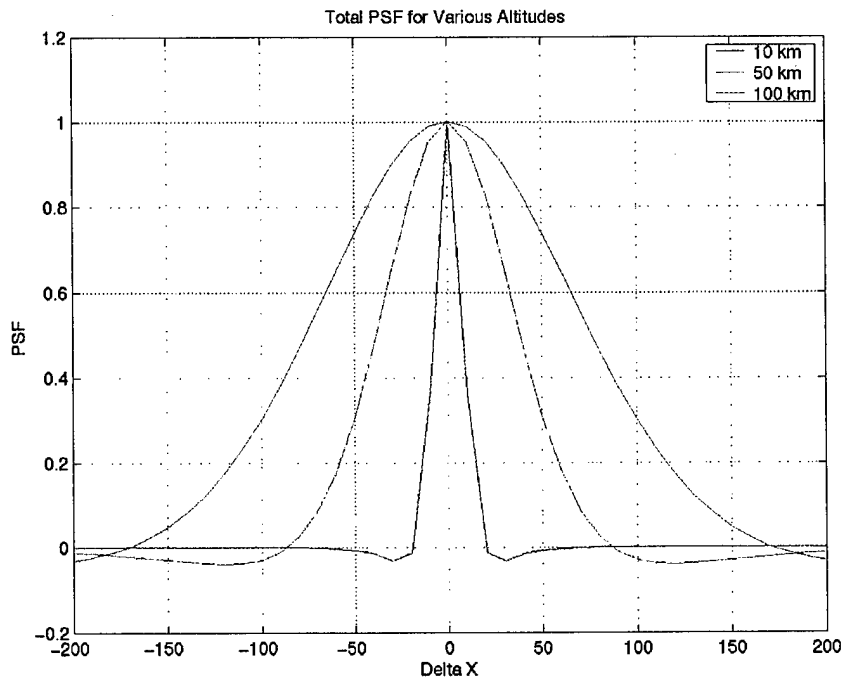


Figure 25. Comparison of cuts through the point-spread function at various sensor altitudes.

The important relationship in Figure 25 is that the width of the spread function is roughly equal to the altitude. Several investigators have reached this same conclusion. Some think the width may be narrowed moderately with better choice of weights or measurement region, but the order-of-magnitude relationship stands.

This order-of-magnitude relationship sizes the problem of moving to more sophisticated techniques. If our resolution goal is 10 m from a 300 km altitude, then we need to drop the matched filter resolution by almost 5 orders of magnitude. This is likely to be very difficult.

7.4.3 As a Function of Measurement Area

The previous cases have shown that resolution is apparently equivalent to the measurement altitude. This result has appeared in other related work and should not be surprising. It is inherent in the structure of the gradiometric equations that fall off in linear distance in proportion to altitude. An important aspect not explored so far is the impact of the region over which measurements are taken. The previous cases have been for a measurement region that is a single-layer spherical patch. The size of the patch could be increased, and it could be thickened to encompass measurement in some depth along the z-axis. We explored this issue by expanding and contracting the size of the measurement patch and by running correlations with several layers of measurement arrayed in depth.

The results were not encouraging. It appears that changing the size of the measurement region has only a weak effect on the inverted resolution, at least for the matched filter case. This is somewhat non-intuitive since many signal processing problems show a direct correlation between the measure of the set on which the function of interest is measured and the inverted resolution. However, in the gradiometry case, there are reasons to believe that the effect is weak or non-existent, at least over a broad range of measurement region sizes. It is not as clear whether the effect is present for non-linear inversion techniques.

To understand the lack of impact of measurement area, we can look at how the product of a tensor component and its displaced copy vary in space. The correlation process integrates the reference copy of the tensor component with a displaced copy. In conventional correlation processing (e.g., radio frequency interferometry) this product is oscillatory and decaying. It is important to integrate over a large area because the integral often cancels to high precision when the integration area is large enough. This is a consequence of the combined oscillatory and decaying structure of these functions. However, the product of gravity gradient components does not oscillate, though it does decay. The result is that integrating the product over a region a few times the altitude captures virtually all of the effect. This is illustrated in Figures 26 through 29. In Figure 26, we show the product of the T_{xx} component and a displaced copy (along the positive x-axis only) for displacements from 0 to 2 times the altitude. As seen, the function decays very rapidly in both x-axis distance and in displacement. There is virtually no residual value in extending the measurement area beyond twice the altitude. For further clarification, we plot $10 \cdot \log_{10}$ of the absolute value in Figure 27.




Figure 26. Product of T_{xx} and a displaced copy for x from 0 to 5 and displacements from 0 to 2. The z -axis is in dimensionless units.




Figure 27. T_{xx} product plot in $10 \cdot \log_{10}$ of the absolute value of the product.

This behavior is not an artifact of choosing the T_{xx} component. Figures 28 and 29 show essentially the same results for T_{xz} .




Figure 28. Product of the T_{xz} components along the x-axis from 0 to 5 for displacements of 0 to 2.




Figure 29. T_{xz} plot repeated but scaled to $10 \cdot \log_{10}$ of the absolute value to better show small values.

7.5 Signal-to-Noise and Clutter-to-Noise

Two figures of merit for an imaging algorithm are the signal-to-noise ratio (SNR) and clutter-to-noise ratio (CNR) or signal-to-clutter ratio (SCR). The SNR is a measure of what fraction of the luminance of a given element (voxel in this case) is due to noise instead of the actual density at that point. The

CNR is a measure of how much of the energy in a voxel is due to the actual density in the corresponding region versus mass elsewhere. Both ratios are relatively easily calculated for this case, although their actual values depend on instrument characteristics.

For the SNR, the simplest assumption is that the noise is independent and identically distributed from sample to sample, and is additive. Then the signal and noise model is:

$$\mathbf{M} = \mathbf{G}(\mathbf{r}) + \mathbf{N},$$

and the SNR is given by:

$$\text{SNR} = \frac{\|\mathbf{G}(\mathbf{r}) \cdot \mathbf{G}(\mathbf{r})\|^2}{\text{E}[\|\mathbf{G}(\mathbf{r}) \cdot \mathbf{N}\|^2]} = \frac{\mathbf{G}(\mathbf{r}) \cdot \mathbf{G}(\mathbf{r})}{\sigma^2}.$$

So, there may be a considerable integration gain if there are many points and the SNR at the maximum point is not too low. The noise variance can divide each term in the numerator dot product to make it the sum of squares of point-by-point SNR. If the peak value is large and there are many independent points, the integrated SNR will be much larger than the point-by-point peak.

Clutter, in this context, is the contribution to a filter due to mass at locations other than the mass point of interest. An ideal matched filter response is approximately a delta function. It peaks very strongly at the point to which it is matched, and has zero response to mass at other points. We can compute clutter impacts from the point-spread function. The part of the point-spread function that we want to be near one is the portion around the peak that covers the resolution we want. So, for example, if we wanted 100-m resolution, we would define the signal part as all filter output due to mass within 50 m of the center of the filter. Filter response to masses farther away than 50 m are clutter because they interfere with using the filter's output as an estimate of the mass at the point of interest.

In gradiometry, the signal-to-clutter performance is likely to be very poor. Looking at Figures 16 and 22, we see that the matched filter response is large compared to one-tenth out to several times the altitude away from the peak. Since signal levels are fairly uniform because the mass field is fairly uniform, these areas will strongly contribute to the total signal in the filter. In all cases, we would expect performance to be clutter limited rather than noise limited.

7.6 Matched Filter Inversion Results

We can see how matched filter inversion works with some simple computational experiments. The geometry for these cases is greatly simplified. The ground is a two-dimensional sheet with mass points on a unit spacing. The width is 100 and the depth is 4 (four layers of mass points). We measure the T_{xx} component of the gravity gradient along a line of constant altitude above the surface, then invert using the previous equations to make an un-normalized estimate of the density. To simulate an underground facility, the mass point field has some points set to zero. In the simplest case, a single

point in the second layer, in the middle of the field, is set to zero. In other cases, we set additional points to zero, or add mass to other points.

7.6.1 Basic Results

The first cases show how the matched filter inversion of a field with a single hole is affected by sensor height. Figure 37 shows the actual mass distribution.

The sensor altitude in Figure 30 is “one,” which means that it is the same as the spacing of the mass points. For a spaceborne sensor system, this is equivalent to looking for 200-km-diameter holes. Thus, this is more analogous to what we would actually see in ground-based collection. To see the effect of altitude, we simply raise the height at which the gradient is computed. The effect is more interesting for separated features and screened features, which is taken up in subsequent plots.

The plot, Figure 30, is notable for two features. First, the presence of the hole’s signature is clear, although its depth is not resolved at all. The dip in the matched filter occurs at zero depth (as if the tunnel were a valley). As will be shown later, least-squares inversion can remedy this problem. Second, the overall inversion is fairly clean, although it does not quantitatively determine the background density. Again, with least-squares estimation, this is changed, although not entirely for the better.




Figure 30. Matched filter inversion at a height of one unit.

7.6.2 Spatial Resolution of Features

Now, to illustrate spatial resolution issues, let there be two tunnels (points of zero mass) separated by one unit. Figure 31 shows the matched filter inversion for this case with sensor altitude of 1.

With two tunnels, the gradient dip is clearly visible, as it was for one tunnel. It appears that the dip really has two minima, as it should if the two tunnels are being resolved. To see this more clearly, we look at the zero depth cut in Figure 32. This shows the two tunnels quite clearly as two minima.

One should be aware of how idealized this is. The sensor altitude is equal to the tunnel spacing. While this is reasonable for ground-based collection, it is unrealistic for spaceborne collection. The sensor altitude for spaceborne collection will be at least 200 km, and tunnel complexes need to be resolved with spacings of 10's of meters, not hundreds of kilometers. Second, this is using essentially ideal gradient measurements. There is no noise model, and real sensor noise will certainly distort this picture.

As the altitude increases, the ability to resolve the two tunnels should disappear. In fact, it does, as shown in Figure 33. In Figure 33, the sensor altitude has been raised to five (that is, five times the feature separation). In this case, the two tunnels appear as a single dip in the gradient.

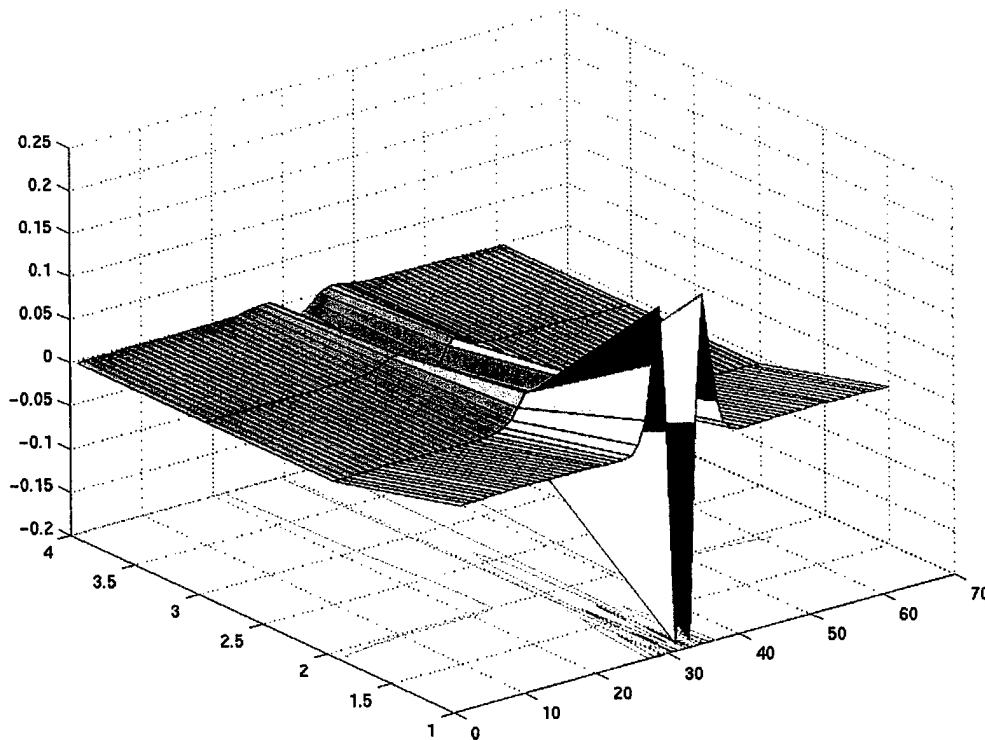


Figure 31. Matched filter inversion of the two-tunnel case for an altitude equal to the spacing.




Figure 32. Zero-depth cut through the matched filter inversion.




Figure 33. Matched filter inversion of two tunnels with sensor altitude of 5 times the tunnel spacing.

7.6.3 Monopole Signal Screening

Another issue of concern is the screen of the monopole signal when mass is not moved a sufficient distance. This issue is studied by adding the missing mass of the tunnel back in proximity to the tunnel. In this case, the mass was added back (in one third parts) to the mass points immediately above and to the sides of the tunnel. Figure 34 shows the resulting inversion for an altitude of 1. Figure 35 repeats the computation for a sensor height of five. What is important to note is that the nature of the signature has considerably changed. What had been a clear dip (due to the mass deficit) has become something more complex. The presence of the screened tunnel has affected the gradient field, but the change is not the change expected by the simple model where the mass is moved. The matched filter technique (because of the limited spatial resolution of the geometry) can "see" that the gradient field is changing, but cannot resolve things in depth. As will be shown, the least-squares technique can do somewhat better, although it too is limited by the natural spatial resolution of the sensor height.

7.6.4 Density Variation Impacts

We know that the real density field of the Earth is not uniform, it has variation on many scales. To evaluate the potential impact of this, we show some simple results in which the background density is randomly varied from point to point, in this case by 5%. Figure 36 shows the matched filter inversion at a height of one for a single tunnel. As is clear, the quality of the inversion is seriously damaged, and the signature of the tunnel disappears into the noise. This is essentially the signal-to-clutter problem discussed in the previous section. The clutter variation swamps the signal variation because of the poor sidelobe behavior of the gradiometric matched filter.




Figure 34. Matched filter inversion of a screened tunnel with sensor height of 1.




Figure 35. Matched filter inversion of a screened tunnel from a height of five.




Figure 36. Matched filter inversion with 5% random density variation.

8. Constrained Inversion Algorithms

The inversion technique investigated in the foregoing was based on matched filters or correlations. The matched filter technique does not take advantage of a great deal of additional information that could assist the inverse problem. The matched filter technique also does not exploit the signal-to-noise-ratio (if high) to attempt to sharpen resolution. Algorithms that do take these into account we term constrained inversion algorithms.

Most algorithms of this type have similar structure. They model the mass density of the Earth as the weighted sum of functions. An obvious example is modeling the Earth as the weighted sum of three-dimensional square waves, each corresponding to one volume pixel. A weighted sum over spherical harmonics is commonly used in estimation of the gravity field. Our problem is a little different since we are trying to estimate the mass density field from the gradiometric field instead of estimating the gravity field of a restricted set of measurements of that field. The expansion should be constructed so that all of its coefficients can be positive, and so that a non-zero coefficient does not violate any observed constraint. For example, there should be no coefficients that would put mass in places it is known not to exist, such as areas known from a topographical model of the Earth to be filled only with air. This automatically introduces the constraint that rock-like densities should be present only for points interior to the Earth (of course, we have to allow interior points to have near zero density as well or we'll never find a tunnel).

Given this expansion that will have to be truncated to some degree compatible with the mass density resolution we desire, then we try to find a set of weights that best matches the observed data (using some appropriate distance metric). In general, the problem is highly non-linear and may not have a unique solution. In restricted cases, such as were, the objective metric is quadratic, and the weights are always positive, one can use special algorithms (quadratic programming, in the example case).

Some quick numbers show how difficult this will be unless there is a very clever approach. If we cut the Earth into 10-m cubes, there will be about 10^{18} of them. Solving an optimization problem with 10^{18} variables is likely to be outside the computationally feasible range since the largest linear programs solved that I can find reference to have less than a billion variables (10^9). The complexity of all optimization algorithms is non-linear in the number of variables, so our problem will be much worse than a billion times more difficult.

It is possible that a happy choice of expansion could mitigate some of the problems. The expansion may also lead to computational problems even less tractable than the brute size. For example, the order of a spherical harmonic expansion that would be needed is orders of magnitude larger than we can stably compute today [Clouser]. Ideally, an expansion should put most effects measured into a few coefficients and minimize the number of coefficients that need to be computed to yield the desired resolution. There has been recent research on spherical wavelets that might apply to this problem. A spherical wavelet expansion would have low-scale terms that encompass the bulk mass

of the Earth and high-scale terms that are non-zero only in the geographic range of interest. Most high-scale terms could be zeroed because they will be too distant from the area of interest to have an effect. The size hierarchy of the wavelets may be usable to allow a rigorous choice of which coefficients to eliminate. Since the effect of a mass falls off with the cube of distance from the area of interest, and its effect rises with its non-zero volume, there is a straightforward way of computing the largest possible impact of a coefficient and then choosing to zero it or not.

8.1 Basic Least-Squares Results

As discussed previously, several investigators, and this work, have concluded that the “natural” spatial resolution of a gradiometric imaging system is the sensor altitude. For spaceborne measurements, this leads to no ISR value, at least in the imaging use-case. Is there a realistic prospect that a nonlinear sharpening technique can image features smaller than the sensor distance? As an experiment, we created a simple case where we could perform the computation. In this case, we model the ground as a finite, two-dimensional rectangle. It contains unit masses on a regular rectangular grid. We record one component of the gradiometric field (T_{xx}) along a line some height above the surface. The point masses are one unit apart (no physical units, we scale everything else to the mass spacing). Now we do a least-square inversion of the gradiometric measurements back into discrete mass distribution. One unit mass point in the middle of the rectangular region at a depth of one unit has been zeroed to simulate a tunnel or other void. The first figure, Figure 37, shows the original mass distribution. The masses are all one, except for one voided point.

Now we simulate gradiometric measurements at a height of 0.25 (sensor height a small fraction of the mass spacing). This produces the estimated mass distribution in Figure 38.




Figure 37. Original mass distribution in inversion experiment.




Figure 38. Estimated mass distribution with a sensor height of 0.25.

Note that the computed mass distribution has the hole, clearly visible, in the right place. Also note that the inversion at greater depths, a depth of three to four, is very ragged.

Now we move the sensor out to heights of one, two, and five in Figure 39, Figure 40, and Figure 41, respectively.

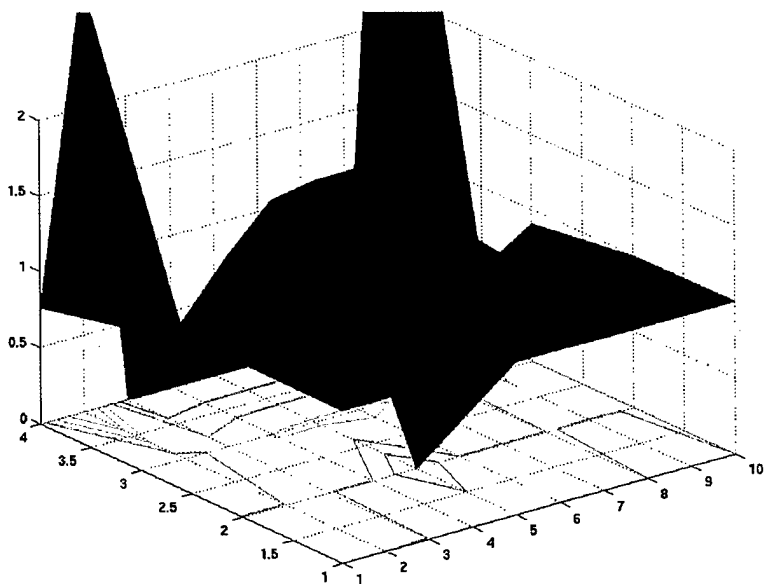


Figure 39. Estimated mass distribution with a sensor height of one.




Figure 40. Estimated mass density with a sensor height of two.




Figure 41. Estimated mass distribution with a sensor height of five.

8.2 Spatial Resolution of Features

Now, as was done in the matched filtering case, we look at how two features can be resolved from each other. The case is the same as before, two zero mass tunnels separated by a unit mass point. Figure 42 shows the resulting inversion with a sensor height of one, Figure 43 shows the least-squares inversion with a sensor height of five. When the sensor height is one, the two tunnels are well resolved. The least-squares technique can see both tunnels accurately, sees the unit density wall between them, and resolves their depth. These are all improvements over the matched filter inversion case. For the sensor at a height of five, however, we can no longer see two tunnels, and even the evidence for one tunnel is suppressed.

8.3 Monopole Signal Screening

The screening case examines the impact of not removing the excavated material when creating the tunnel. As before, the removed mass is added back to the mass elements above and nearby. Figure 44 shows the least-squares inversion of this case from a height of one. Figure 45 shows the inversion from a height of five.

8.4 Density Variation Impacts

The final case to examine is the impact of random density variation. As before, the unit masses are randomly varied with a standard deviation of 5%. In the matched filter case, a 5% density variation was sufficient to swamp the tunnel signature, even with a sensor height of one. In the least-square case, as shown in Figure 46, 5% variation does not swamp the clear indication of the tunnel. The tunnel shows up clearly as the dip in estimated density at the proper location.




Figure 42. Least-squares inversion for two tunnels with a sensor height of one.




Figure 43. Least-squares inversion of the two tunnel case for a sensor height of five.




Figure 44. Least-squares inversion of a screened tunnel from a height of one.




Figure 45. Least-squares inversion of the screen tunnel from a height of five.




Figure 46. Least-squares inversion for 5% randomized density points from a height of one.

9. Summary of Feasibility

The object of this analysis has been to evaluate the feasibility of using gravity gradiometry for ISR purposes, either detection or characterization of underground facilities. The weight of the analysis suggests the following conclusions:

- Inversion to density maps, while technically very challenging, is more useful for ISR purposes than detection schemes. Detection schemes are likely to be useful only in very limited circumstances, if at all.
- It is likely that ground-based gradiometric measurements can achieve useful volumetric resolutions (a few meters) given developmental instruments and advanced processing algorithms.
- Similar resolutions may be achievable from airborne measurements, but success is much less likely, and major developments are necessary in both instruments and processing approaches.
- Achieving useful volumetric resolution from space is extremely unlikely, and achieving reliable direct detection of underground facilities is even less likely. Achieving the required instrumental sensitivity is likely to be very difficult, though not impossible. There are no known approaches to achieving the degree of processing "sharpening" on the inverse problem needed for useful results.

We take each of these points in turn, referring back to the main body of the argument for the details.

9.1 Inversion versus Detection

Inversion to a density map appears to be more mathematically challenging than matching to the gradiometric signature of a void, but is probably simpler. First, inversion applies to all collection modes (ground, airborne, space). In each case, the goal is to characterize a known facility by determining the layout and density structure of the construction. Detection is appropriate to wide-area search for underground facilities. However, wide area search can only be conducted from high airborne altitudes or space. When the altitude is high enough for wide area searching (10's to 100's of kilometers) the detection schemes may fail due to the excavation spoil. In Section 7.4, we saw that the gradiometric response of a point mass had extended approximately equal to the altitude of sensor. When the sensor is 10's to 100's of kilometers away, the signature of interest will be spread over a similar area. It is unlikely that the excavation spoilings will be moved 10's to 100's of kilometers away from the facility area. If they are not, then the accumulation of spoilings will produce a region of increased mass and will distort the gradiometric field just as the voids of the facility do. The two distortions are opposite in direction (an increase from the accumulation of spoil, a decrease from the voids) and

almost completely overlap since their horizontal separation will normally be much less than the horizontal width of their gradiometric fields (at the sensor altitude). The two distortions are also about the same size, and one is the mass excess produced by the mass deficit of the other.

An inversion system that could sharpen the responses or invert to point masses much smaller than the spoil-facility separation could resolve this problem. But, the whole point of detection is to avoid inversion. If inversion at useful volumetric resolutions is possible, we can use it for both characterization and search (simply by scanning the inverted density map for characteristically low-density regions below the local Earth surface).

A second fundamental difficulty with the detection schemes is that the noise and clutter statistics are unknown, both spatial and temporal. Since we lack any knowledge of these, it is possible that they swamp the signature effects we are interested in. Of course, it is also possible that the clutter statistics and the target statistics for man-made facilities differ so substantially as to make the problem much easier. Unfortunately, this is unlikely to be the case for space-based measurement. The most obvious characteristic of man-made facilities, long linear voids, are too small to generate much of a signature from space. In the airborne case, especially with a full-tensor instrument, the case might be different. A measurement program could, in principle, resolve this clutter problem, but it cannot solve the problem of signature masking by spoil.

9.2 Inversion of Ground-Based Measurements

In this case, we are measuring the gravity gradients on some grid on the surface of the Earth. The altitude of the sensor above the facility is simply the depth of the facility. For our generic case, this means that we are looking for volumetric resolutions of a few meters from sensor measurements taken roughly 100 m away. The natural width of the matched filter inversion is about that distance, or 100 m (see Section 7.4). We want a resolution of a few meters, so we need to sharpen the response by a factor of 20–50 over the matched filter.

A cube of rock equal in volume to the desired volumetric resolution (1 m^3 to about 100 m^3) produces a gravity gradient in the range of 1 milliEotvos to 0.1 Eotvos at 100 m. This is modestly better than today's instruments, but should be much larger than the sensitivity of quantum-effects-based sensors that could be available in the next few years. This means that developmental instruments used in ground-based collection will generate high signal-to-noise ratios (20 dB) on units of rock equal to the desired volumetric resolution. This suggests that we can make use of aggressive non-linear algorithms for resolution sharpening. A rule of thumb in other signal processing applications is that resolution can be improved over the matched filter by roughly the square root of the SNR. For a 20 dB SNR on each cell, we should be able to realize a factor of 10 improvement resolution, which would bring the ground-based application within the desired bounds.

An additional positive factor in the ground-based case is that the inverse cubed fall-off in the gravity gradient should restrict the volume over which the inverse computation must be performed. If the facility of interest is 100 m below the sensor, then rocks 200 m below the sensor have their gradients attenuated by a factor of 8. At 1 km the attenuation is a factor of 1000. The volume of rock in a constant thickness shell rises with the square of the distance, so the net attenuation is just the distance.

Still, for a facility 100 m down, we might need to compute a cubic kilometer of Earth to adequately account for all variations (sidelobe-induced distortions). There are 10^9 one-meter cubes of rock in a cubic kilometer. If we were computing on a density map with a resolution of one meter, this implies solving a non-linear inversion problem of order one billion. This is not a small computational challenge, but is not clearly unfeasible either.

In summary, with developmental instruments used in ground-based collection over an underground facility, the collected data will be of high SNR, the region that might interfere with the inversion accuracy is tractable, and the scope of the inversion problem appears to be tractable. We cannot, today, carry out the required program of measurement and computation, but it appears to be a feasible problem. To actually carry out the objective, we would need to develop tractable gradiometric inversion algorithms and support or obtain gradiometers with sensitivities in the range of one milliEotvos. Some research and development is needed to produce such instruments, but the roadmap to doing so is fairly clear.

9.3 Inversion from Airborne Measurements

The case for airborne measurements covers the same points, but is orders of magnitude more difficult and introduces one new problem. First, we deal with the issues of scaling. In the airborne case, the collections will be made from an altitude of (nominally) 1 to 10 km. So, the distance is 10 to 100 times greater than in the ground-based case. The problems scale accordingly. The gradiometric field of a point mass scales directly with the distance, as does the size of a voxel inverted by a matched filter. To get to a volumetric resolution of a few meters, we now must sharpen by a factor of a few hundred to a few thousand. The depth of influence also rises linearly with the altitude, but the volume within that area goes up with the cube of the altitude. So, where in the ground-based case, we might have needed to compute the density over a cubic kilometer with 10^9 meter cells, we now have to compute over 10^3 to 10^6 cubic kilometers containing from 10^{12} to 10^{15} meter cells. Computationally, this is on the edge of, or beyond, computational feasibility.

To accomplish much sharpening requires a reasonable SNR on the volumetric elements in question. Over the airborne distance range, the gradients due to the rock in one desired resolution cell range from 10^{-4} to 10^{-7} Eotvos. Based on recent work in quantum-effects-based instruments, this is not out of the question, but it is very technically challenging. More difficult is achieving the required SNR. Since the algorithms for doing the required inversion do not exist, we cannot determine their required SNR. By the rule-of-thumb used previously, we want to sharpen by a factor of 10^3 , so we may need an SNR of as much as 10^6 . Being more optimistic, perhaps an SNR of only 10^3 will be required. Regardless, it implies that an instrument with a noise floor of from 10^{-7} to 10^{-13} Eotvos would be required. Again, this is not impossible, but this is pressing the optimistic estimates of what might be ultimately achieved.

The airborne case adds an additional problem, the dynamic environment of an aircraft. In the airborne case, the instrument cannot be fixed to the ground or allowed to float in space. Instead it has to be attached to a powered aircraft, which implies a high level of vibration and atmospheric dynamics. In principle, a gradiometer can cancel accelerations and detect only gravitational gradients. But, in the airborne case it will be required to cancel effects many orders of magnitude larger than the effects

of interest. This has to be regarded as a large risk, although we have no data today with which to evaluate it.

9.4 Inversion from Spaceborne Measurements

The spaceborne case is similar to the airborne case, just worse by another 1–2 orders of magnitude in distance (and thus a factor of a thousand to a million in other areas). With the sensor 300 km away from the ground, we now face the requirement to sharpen a factor of about 10^5 to get into the useful range. The depth of influence is now 3000 km, or a substantial fraction of the Earth. In the spaceborne case, it might be necessary to invert the density map of the entire planet in order to find and characterize the facilities of interest. The order of the inversion problem rises to near 10^{20} , and is almost certainly beyond any contemplated computer system.

The gradient signal of an interesting resolution cell is now less than 10^{-11} Eotvos, below contemplated sensitivities. The instrument noise floor would have to be orders of magnitude below even that to allow significant sharpening on the inversion. Those who are hopeful about space applications note that the gradiometric signature of the whole tunnel is several orders of magnitude larger, and so might be detectable. However, this runs into all of the problems of detection-only mentioned before (natural clutter, spatial aliasing, spoil cancellation of the signature, etc.). The problem of clutter is likely to become dramatically worse as the sensor altitude rises because the depth of influence rises to include much of the Earth. For ground-based measurements only the upper kilometer or so of the crust significantly influences the measurements collected. For space measurements, the depth of influence extends through the crust and deep into the mantle where entirely unknown geologic processes may be at work.

9.5 Recommendations

The task set was to evaluate the feasibility of detecting and characterizing underground facilities using gradiometric sensors in space. The clear conclusion is that space-based gradiometric sensors are very unlikely to provide useful military ISR against underground facilities. The only development that could change this conclusion is the availability of inverse gradiometric processing techniques far more effective and computationally efficient than any now known. No amount of sensitivity improvement in gradiometers can change the conclusion, although with a processing breakthrough it is possible that gradiometer sensitivity could become a make-or-break issue. Ground-based and airborne measurement looks much more feasible, with the ground-based looking likely and the airborne case as possible though highly challenging. As a result, from the perspective of space-based implementation, we recommend a program of research in the gradiometric inverse problem, but not in other areas. This research should be fairly fundamental since only major advances are likely to have any chance of making the spaceborne case feasible. Organizations interested in ground-based or airborne applications would be justified in investing in both inversion algorithms and instruments, although the inversion still appears to be the most difficult and the primary barrier to feasibility.

References

Clauser, J., "Surveillance and gravity-imaging of the Earth's surface, and satellite gravity gradiometry with atom interferometers," Appendix D of Final Report on Coherent, Atomic, Matter Wave Gravity Gradiometry:

A Quantum Breakthrough in Remote Sensing from Space, by Dowling, J, and collaborators. Published as an NRO/DII-2000 report.