

Gap Detection with Electromagnetic Terahertz Signals

H. T. Banks ¹, Nathan L. Gibson ¹, William P. Winfree ²

August, 2004

Abstract

We apply an inverse problem formulation to determine characteristics of a defect from a perturbed electromagnetic interrogating signal. A defect (gap) inside of a dielectric material causes a disruption, via reflections and refractions at the material interfaces, of the windowed interrogating signal. We model the electromagnetic waves inside the material with Maxwell's equations. This leads to a non-standard, nonlinear optimization problem for the dimensions and location of the defect. Using simulations as forward solves, we employ a Newton-based, iterative optimization scheme to a novel modified least-squares objective function. Numerical results are given in tables and plots, standard errors are calculated, and computational issues are addressed.

1 Introduction

The problem we consider is that of detecting a gap inside of a dielectric material using high frequency electromagnetic interrogation. The idea is to observe the reflected and/or transmitted signals and use the data to solve an inverse problem to determine certain characteristics of the gap, e.g., location and/or width. Possible applications of this procedure include quality assurance in fabrication of critical dielectric materials, or damage detection in aging materials for safety concerns. Further applications of electromagnetic interrogation, as well as additional problem formulations and solutions, can be found in [2].

The particular motivation for this research is the detection of defects in the insulating foam on the fuel tanks of the space shuttles in order to help eliminate the separation (delamination) of foam during shuttle ascent. To this end we address the problem of detecting a single gap formed between a dielectric medium and a supra-conducting backing representing the foam and the metallic tank, respectively. However, first we develop our methodology on the slightly simpler problem of a gap formed in the interior of the foam (void), where for simplicity, we ignore the reflections from the back boundary (i.e., we impose absorbing boundary conditions instead). We also allow for the possibility in this case that the foam has been removed for testing, and therefore we are able to place sensors both on the front and back sides of the foam.

To be applicable to real world problems we must eventually be able to solve these inverse problems with length scales on the order of $20cm$ for the thickness of the foam, $.2mm$ for the width of the gap, and a wavelength of about $3mm$. This wavelength corresponds to a frequency of $100GHz$, which is the lower end of the terahertz frequency range ($.1 \sim 10 \times 10^{12}Hz$). The rationale for using this choice of frequency is that higher frequencies are significantly attenuated in the materials which we are interested in interrogating. Lower frequencies (larger wavelengths) have less resolution in detecting small gaps, and are less capable of sharply distinguishing between air and foam which may have a high air content.

First we simplify the problem by considering a linearly polarized, pulsed interrogating signal which reduces the problem to one spatial dimension. We then define an inverse problem for determining the gap's dimensions. We assume that we have data from sensors, located in front of and/or behind the material, that record the electromagnetic signal after it is reflected from (or passes through) the material interfaces. We compute simulated signals with approximations to the gap's characteristics and apply an optimization routine to a modified least squares error between this simulated signal and the given data. In our computations we use Maxwell's equations and a Debye polarization equation to model the signal, and solve these equations using a Finite Element method in space and Finite Difference methods in time. Thus the optimization routine finds those gap characteristics which generate a simulated signal that most closely matches the given data. In this sense we determined an estimate to the "true" gap characteristics.

¹Center for Research and Scientific Computation, Box 8205, NC State University, Raleigh, NC 27606-8205

²NASA Langley Research Center, MS 231, Hampton, VA 23681-2199

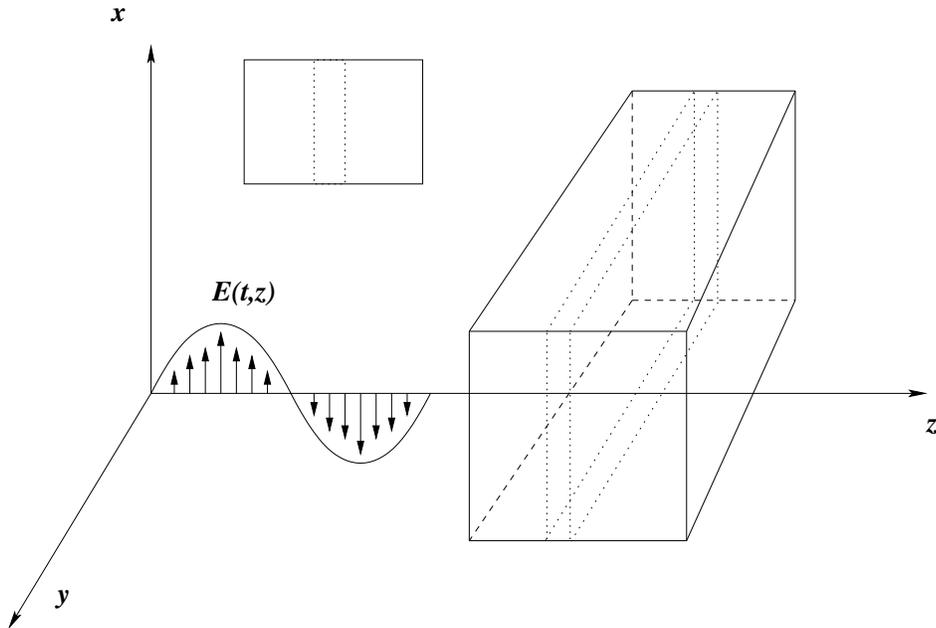


Figure 1: *Problem 1*: Dielectric slab with a material gap in the interior. Possible sensors in front and behind.

In Section 2 we define the equations that we have chosen in order to model the electromagnetic waves inside the material. We further distinguish between two problem types that we will address (denoted as *Problem 1* and *Problem 2*). Section 3 contains the details of our numerical methods for the simulations. We introduce the inverse problem formulation for *Problem 1* in Section 4, and later improve upon it in Section 4.2. Numerical results of the inverse problem are displayed in Section 4.3.

In Section 5 we begin addressing *Problem 2*. Similarities and differences between the computational issues between the two problems are pointed out. A more sophisticated optimization method is described in Section 5.3 and associated numerical results are given in Section 5.5. In Sections 5.5.1 and 5.5.2 we explore the effects of adding random noise to the data, both relative and constant variance. In the latter, we compute standard error estimates.

2 Problem Description

We interrogate an (infinitely long) slab of homogeneous nonmagnetic material by a polarized, windowed signal (see [2] for details) in the THz frequency range (see Figure 1). We assume a wave normally incident on a slab which is located in $\Omega = [z_1, z_4]$ with faces parallel to the x - y plane (see Figure 2). Note that we employ the “method of mappings” (see [2]) in our computations, therefore we may assume $0 < z_1 < z_4 < 1$. We denote the vacuum outside of the material by Ω_0 . The electric field is polarized to have oscillations in the x - z plane only. Restricting the problem to one dimension, we can write the electric and magnetic fields, \vec{E} and \vec{H} respectively, as follows

$$\vec{E}(t, \vec{x}) = \hat{i}E(t, z)$$

$$\vec{H}(t, \vec{x}) = \hat{j}H(t, z),$$

so that we are only concerned with the scalar values $E(t, z)$ and $H(t, z)$.

Maxwell’s equations then become:

$$\frac{\partial E}{\partial z} = -\mu_0 \frac{\partial H}{\partial t} \tag{1a}$$

$$-\frac{\partial H}{\partial z} = \frac{\partial D}{\partial t} + \sigma E + J_s, \tag{1b}$$

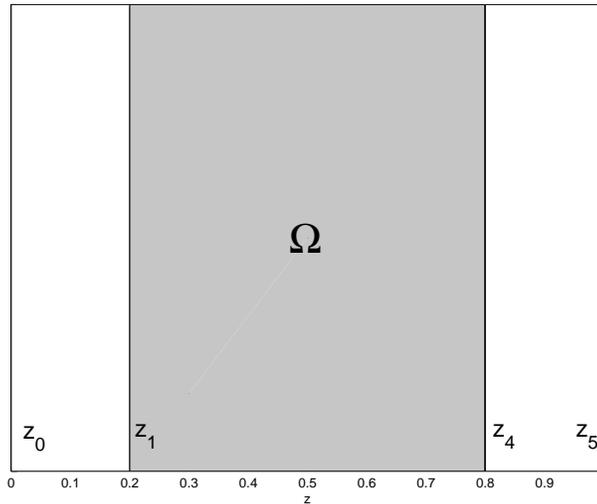


Figure 2: The domain of the material slab: $\Omega = [z_1, z_4]$.

where $D(t, z)$ is the electric flux density, μ_0 is the magnetic permeability of free space, σ is the conductivity, and J_s is a source current density (determined by our interrogating signal). We take the partial derivative of Equation (1a) with respect to z , and the partial of Equation (1b) with respect to t . Equating the $\frac{\partial^2 H}{\partial z \partial t}$ terms in each, and thus eliminating the magnetic field H , we have:

$$E'' = \mu_0 \left(\ddot{D} + \sigma \dot{E} + \dot{J}_s \right),$$

(where $'$ denotes z derivatives and $\dot{}$ denotes time derivatives).

Note that we have neglected magnetic effects and we have let the total current density be $J = J_c + J_s$, where $J_c = \sigma E$ is the conduction current density given by Ohm's law in the material.

For our source current, J_s , we want to simulate a windowed pulse, i.e., a pulse that is allowed to oscillate for one full period and then is truncated. Further, we want the pulse to originate only at $z = 0$, simulating an infinite antenna at this location. Thus we define

$$J_s(t, z) = \delta(z) \sin(\omega t) I_{[0, t_f]}(t),$$

where ω is the frequency of the pulse, $t_f = 2\pi/\omega$ is fixed, $I_{[0, t_f]}(t)$ represents an indicator function which is 1 when $0 \leq t \leq t_f$ and zero otherwise, and $\delta(z)$ is the Dirac delta distribution.

Remark 1 *Computationally, having a windowed signal introduces discontinuities in the first derivatives which are not only problematic in the numerical simulations (producing spurious oscillations), but are also essentially non-physical. Therefore in our implementation we actually multiply the sine function by an exponential function (see [3] for details) rather than the traditional indicator function. However, for notational consistency we will continue to denote this function as $I_{[0, t_f]}(t)$.*

The electric flux density inside the material, given by $D = \epsilon_0 \epsilon_\infty E + P$, is dependent on the polarization, P . Note that ϵ_0 is the permittivity of free space and ϵ_∞ is the relative permittivity in the limit of high frequencies. For computational testing we assume for this presentation that the media is Debye and thus we use the following polarization model inside Ω :

$$\tau \dot{P} + P = \epsilon_0 (\epsilon_s - \epsilon_\infty) E,$$

where ϵ_s is a static relative permittivity and τ is a relaxation time. We also assume $P(0, z) = 0$. Note that in the vacuum outside of Ω , $P = 0$.

In order to represent D in the entire domain, we use the indicator function I_Ω which is 1 inside Ω and zero otherwise. Thus

$$D = \epsilon_0 E + \epsilon_0(\epsilon_\infty - 1)I_\Omega E + I_\Omega P.$$

In order to have a finite computational domain, we impose absorbing boundary conditions at $z = 0$ and $z = 1$, which are modeled as

$$\begin{aligned} \left[\dot{E} - cE' \right]_{z=0} &= 0 \\ \left[\dot{E} + cE' \right]_{z=1} &= 0. \end{aligned}$$

With these boundary conditions, any boundary incident signal passes out of the computational domain, and does not return, i.e., we force it to be absorbed by the boundary. Also we assume zero initial conditions, i.e.,

$$\begin{aligned} E(0, z) &= 0 \\ \dot{E}(0, z) &= 0. \end{aligned} \tag{2}$$

Thus our entire system can be written

$$\begin{aligned} \mu_0 \epsilon_0 (1 + (\epsilon_\infty - 1)I_\Omega) \ddot{E} + \mu_0 I_\Omega \ddot{P} + \mu_0 \sigma I_\Omega \dot{E} - E'' &= -\mu_0 \dot{J}_s \quad \text{in } \Omega \cup \Omega_0 \\ \tau \dot{P} + P &= \epsilon_0 (\epsilon_s - \epsilon_\infty) E \quad \text{in } \Omega \\ \left[\dot{E} - cE' \right]_{z=0} &= 0 \\ \left[\dot{E} + cE' \right]_{z=1} &= 0 \end{aligned} \tag{3}$$

with (2) and

$$J_s(t, z) = \delta(z) \sin(\omega t) I_{[0, t_f]}(t). \tag{4}$$

Classical solutions to (3) should not be expected due to the windowed interrogating signal and the discontinuous dielectric parameters across interfaces. For this reason, and also to enable the application of the Finite Element Method, we prefer to convert (3) to weak form using spaces $H = L_2(0, 1)$ and $V = H^1(0, 1)$. Substituting $\epsilon_r = (1 + (\epsilon_\infty - 1)I_\Omega)$ and $\epsilon_d = \epsilon_s - \epsilon_\infty$ results in the following weak system

$$\begin{aligned} \langle \mu_0 \epsilon_0 \epsilon_r \ddot{E}, \phi \rangle + \langle \mu_0 I_\Omega \ddot{P}, \phi \rangle + \langle \mu_0 \sigma I_\Omega \dot{E}, \phi \rangle + \langle E', \phi' \rangle - \frac{1}{c} \dot{E}(t, 1) \phi(1) + \frac{1}{c} \dot{E}(t, 0) \phi(0) &= -\langle \mu_0 \dot{J}_s, \phi \rangle \\ \langle \tau \dot{P}, \phi \rangle_\Omega + \langle P, \phi \rangle_\Omega &= \langle \epsilon_0 \epsilon_d E, \phi \rangle_\Omega \end{aligned} \tag{5}$$

with (2) and (4). Note that $\langle \cdot, \cdot \rangle$ is modified from the traditional L_2 inner product due to the aforementioned use of the ‘‘method of mappings’’ (see [2]). Existence and uniqueness of systems of this type are treated in [2].

In this formulation we have initially assumed a single slab of a dielectric contained in $\Omega = [z_1, z_4]$. Thus $I_\Omega = 1$ if $z_1 < z < z_4$, and zero otherwise. We now introduce a gap consisting of a vacuum in the interior of the material as depicted in Figure 3. If the gap is located in (z_2, z_3) then we redefine $\Omega = \{z | z_1 \leq z \leq z_2 \text{ or } z_3 \leq z \leq z_4\}$. We will refer to this formulation as *Problem 1* (recall Figure 1). Note that it is not necessary to enforce additional conditions at the interfaces as they are natural interface conditions and are implied within the weak form of the system. Later we will discuss a second problem formulation, *Problem 2*, where the gap is between the dielectric slab and a metallic (supra-conducting) backing, as shown in Figure 4. This will require slightly different boundary conditions (reflecting instead of absorbing at $z = 1$, where the metal backing begins), but otherwise the numerical solution methods and analysis are the same.

3 Numerical Solution

3.1 Finite Elements

We apply a Finite Element method using standard linear one dimensional basis elements to discretize the model in space. Let N be the number of intervals in the discretization of z , and $h = 1/N$, then the Finite

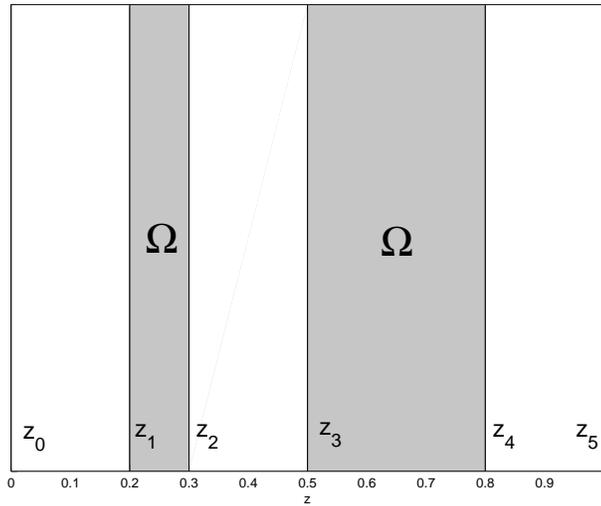


Figure 3: The domain of the material slab with an interior gap between z_2 and z_3 : $\Omega = \{z | z_1 \leq z \leq z_2 \text{ or } z_3 \leq z \leq z_4\}$.

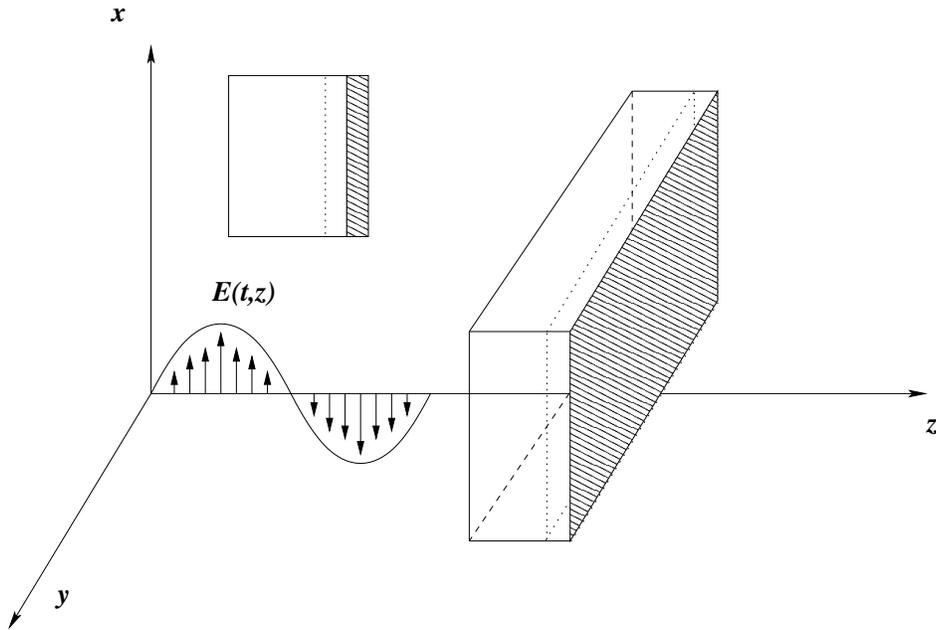


Figure 4: *Problem 2*: Dielectric slab and metallic backing with a gap between. Possible sensors only in front.

Element discretization has an order of accuracy of $O(h^2)$. For implementation we scale time by $\tilde{t} = ct$ and the polarization by $\tilde{P} = P/\epsilon_0$ for convenience. The resulting system of ordinary differential equations after the spatial discretization is the semi-discrete form

$$\epsilon_r M \ddot{e} + M^\Omega \ddot{p} + (\eta_0 \sigma M^\Omega + D + B) \dot{e} + K e = \eta_0 J \tag{6a}$$

$$M^\Omega \dot{p} + \lambda M^\Omega p = \epsilon_d \lambda M^\Omega e, \tag{6b}$$

where $\lambda = \frac{1}{c\tau}$, and $\eta_0 = \sqrt{\mu_0/\epsilon_0}$. Also e and p are vectors representing the approximate values of E and P , respectively, at the nodes $\tilde{z}_i = ih$. The mass matrix M has entries

$$M_{ij} = \langle \phi_i, \phi_j \rangle := \int_0^1 \phi_i \phi_j dz,$$

where $\{\phi_i\}_{i=1}^N$ are the basis functions (M^Ω is the mass matrix integrated only over Ω), while the stability matrix K has entries

$$K_{ij} = \langle \phi'_i, \phi'_j \rangle := \int_0^1 \phi'_i \phi'_j dz.$$

The matrices D and B result from the boundary conditions where

$$\begin{aligned} D_{1,1} &= 1 \\ B_{N+1,N+1} &= 1 \end{aligned}$$

and all other entries are zero. Finally, J is defined as

$$J_i = -\langle \phi_i, \dot{J}_s \rangle := -\int_0^1 \dot{J}_s \phi_i dz.$$

Note that by differentiating (6b) we can substitute into (6a) and obtain an equation only dependent explicitly on P (two substitutions are required to eliminate \ddot{P} and \dot{P}):

$$\epsilon_r M \ddot{e} + (\eta_0 \sigma M^\Omega + D + B + \epsilon_d \lambda M^\Omega) \dot{e} + (K - \epsilon_d \lambda^2 M^\Omega) e + \lambda^2 M^\Omega p = \eta_0 J.$$

Using shorthand we can write our entire coupled system as

$$M_1 \ddot{e} + M_2 \dot{e} + M_3 e + \lambda^2 \bar{p} = \eta_0 J \quad (7a)$$

$$\dot{\bar{p}} + \lambda \bar{p} = \epsilon_d \lambda M^\Omega e, \quad (7b)$$

where $\bar{p} = M^\Omega p$. It is important to mention that each matrix is tridiagonal due to the choice of the linear finite elements.

3.2 Finite Differences

In order to solve the semi-discrete form of our equations we consider two distinct finite difference methods. In the first method we convert the coupled second order system of equations into one larger first order system and simply apply a theta method (unless otherwise stated, we use $\theta = \frac{1}{2}$). In the second method we solve first for the polarization with a forward differencing scheme using the initial conditions and then use that to update a second order central difference scheme for the magnitude of the electromagnetic field. We then continue this process iteratively, alternating between solving for P and for E .

Both methods are second order in time and space for appropriately smooth data (and with $\Delta t = O(h)$). We have compared the errors and the run times of the two methods for several smooth test problems and have determined that the second method is as accurate, but twice as fast, as the first in all cases primarily due to the fact that the linear system is of smaller dimension in the second method. Also, the first method incidentally solves for \dot{e} in addition to e and p , which is superfluous.

In our second method we use a second order central difference scheme to solve (7a). Thus we must first find an approximation to $E(t_1, z)$ where $t_i = i\Delta t$. Note that approximating E with its Taylor expansion around $t_0 = 0$ and applying the initial conditions and ODE, one obtains

$$E(t_1, z) \approx -\frac{\Delta t^2}{2} \mu_0 \dot{J}_s(0, z).$$

Our approach is to first solve for \bar{p} using a θ -method, and then use that approximation to solve for e at the next time step. Thus, our finite difference approximation for (7b) is

$$\bar{p}_{n+1} = \bar{p}_n + \frac{\lambda \Delta t}{1 + \lambda \Delta t \theta} (\epsilon_d M^\Omega e_{n+\theta} - \bar{p}_n) \quad (8)$$

where $[e_n]_j = E(t_n, \tilde{z}_j)$, $[\bar{p}_n]_j = M^\Omega P(t_n, \tilde{z}_j)$, $\tilde{z}_j = jh$, and $e_{n+\theta} = \theta e_n + (1 - \theta)e_{n+1}$ is a weighted average of e_n and e_{n+1} for relaxation to improve the stability of the method. Once we have \bar{p}_{n+1} we can solve for e_{n+2} . Applying a second order central difference method with averaging to (7a) gives

$$A_1 e_{n+2} = A_2 e_{n+1} + A_3 e_n + \Delta t^2 \eta_0 J_{n+1} - \lambda^2 \Delta t^2 \bar{p}_{n+1}. \quad (9)$$

Note that in this case A_1 is tridiagonal and the matrix is the same for each time step, so we may store the Crout LU factorization and use back substitution to solve the system at each time step. For tridiagonal matrices the factorization and the back substitution are both order $O(N)$. See [3] for computational issues encountered in implementation, including the use of the method of maps to allow for gap sizes (δ) smaller than the mesh size (h).

3.3 Numerical Simulations

The following figure depicts the numerical solution of the amplitude of the electric field at various times (Figure 5). We considered a Debye medium with the following parameters:

$$\begin{aligned} \epsilon_s &= 78.2, \\ \epsilon_\infty &= 5.5, \\ \sigma &= 1 \times 10^{-5}, \\ \tau &= 3.16 \times 10^{-8}, \\ f &= 2GHz, \end{aligned}$$

and a material gap located at $[z_3, z_4] = [.6, .61]$.

We have used 2 GHz simply so that our computational domain of $z \in [0, 1]$ would not have to be scaled for this demonstration. Also, in practice, one would not compute a domain so much larger than the material, just as in an experiment the sensors should be as close to the material as possible to reduce noise. We did so here merely so that the full wavelength of the signal would be visible. See [3] for figures showing the signal recorded at receivers located at $z = 0$ and $z = 1$.

4 Problem 1

We now apply an optimization routine to the least squares error between a simulated signal and the given data to try to determine the gap characteristics. In particular we will be trying to find the depth, $d := z_2 - z_1$, and the width, $\delta := z_3 - z_2$, which will produce a simulated signal most closely similar (in the least squares sense) to the data. Existence and continuous dependence of inverse problems of this type are addressed in [2].

4.1 Inverse Problem

All of the following are solved with respect to a reference problem ($R1$) with these parameter values (see also Figure 3):

$$\begin{aligned} z_0 &= 0, z_1 = .2, z_2 = .3, z_3 = .5, z_4 = .8, z_5 = 1.0 \\ f &= 4 \text{ GHz}, t_f \text{ is one period} \\ \tau &= 8.1 \times 10^{-12}, \sigma = 1 \times 10^{-5}, \epsilon_s = 80.1, \epsilon_\infty = 5.5 \text{ in the material} \\ N &= 1024, N_t = 12926 \end{aligned}$$

The sample rate for the data is one sample per $sr = .05ns$

The corresponding values of (d, δ) are $(.1, .2)$. With this choice of parameters, the forward solve solution at $z = 0$ clearly shows distinct reflections from the z_1 , z_2 , and z_3 interfaces. This clear distinction will aid in our approximating the initial guesses, thus making this a relatively easy sample problem.

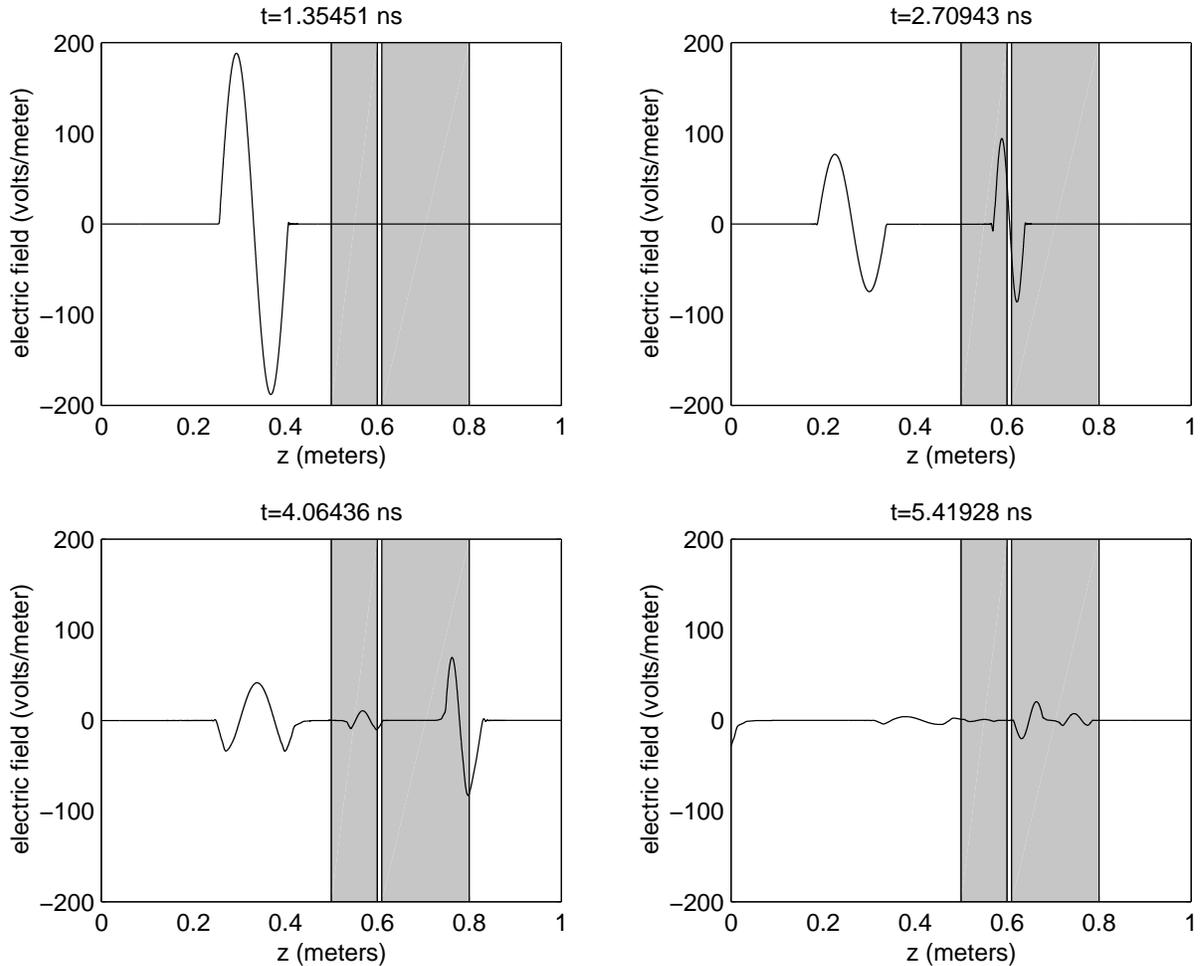


Figure 5: Computed solutions at different times of a windowed electromagnetic pulse incident on a Debye medium with a gap.

4.1.1 Initial Guesses

Assuming the physical parameters are given (either known or from a previous estimation), we want to determine the depth of any gap (d), and the width of that gap (δ), using reflection and/or transmission signals. First we attempt to get very close approximations to d and δ using information about the travel times of the data signal, then we use these values as initial guesses in the optimization routine. See [3] for a complete discussion of the methods used in determining initial estimates.

4.1.2 Optimization of Least Squares Error

With initial estimates to d and δ established, we define our inverse problem to be: find $q = \{d, \delta\} \in Q_{ad}$ such that the following least squares error between the simulation and the observed data is minimized:

$$J(q) = \frac{1}{2MS} \sum_{j=1}^M \sum_{i=1}^S (E(t_i, z_j^O; q) - \hat{E}_{ij})^2. \quad (10)$$

Here the \hat{E}_{ij} are measurements of the electric field taken at M specific locations (e.g., $z_1^O = 0$ and/or $z_2^O = 1$) and S distinct times (e.g., every $sr = 0.06ns$). The $E(t_i, z_j^O; q)$ are solutions of the simulations evaluated at

the same locations and times corresponding to the data, \hat{E}_{ij} , and using parameter values q . The set Q_{ad} is the feasible set of q values determined such that d and δ are realistic (e.g. positive).

We apply an inexact Gauss-Newton iterative method to the optimization problem. That is, we re-write the objective function as

$$J(q) = \frac{1}{2MS} R^T R$$

where $R_k = (E(t_i, z_j^O; q) - \hat{E}_{ij})$ for $k = i + (j - 1)M$ is the residual. To update our approximation to q we make the Inexact Newton update step $q_+ = q_c + s_c$ where

$$\begin{aligned} s_c &= - (R'(q_c)^T R'(q_c))^{-1} \nabla J(q_c) \\ &= - (R'(q_c)^T R'(q_c))^{-1} R'(q_c)^T R(q_c) \end{aligned}$$

is the step, q_c is the current approximation, and q_+ is the resulting approximation. This is an inexact method because we have disregarded the S Hessians of $(E(t_i, z_j^O; q) - \hat{E}_{ij})$, which is generally acceptable for small residual problems [7].

In this simple case we have a 2×2 matrix inverse, so we can compute it explicitly. Each iteration requires one function evaluation and a forward difference gradient, which is two additional function evaluations (since we have two parameters). Each function evaluation is equivalent to a simulation. Therefore we want as few iterations as possible.

4.1.3 Convergence

Initial testing (with $z = 0$ data only) shows convergence to 8 decimal places of each parameter in 6 iterations of Gauss-Newton with initial guesses having at most 5% relative error in δ and 2% in d . The algorithm does not converge to the correct solution if the initial guess for d has 5% relative error or if the initial guess for δ has 10% relative error.

One reason that the algorithm fails to converge is that this objective function is poorly behaved. In Figure 6 we show a plot of the objective function with respect to δ . The two very large peaks in J on

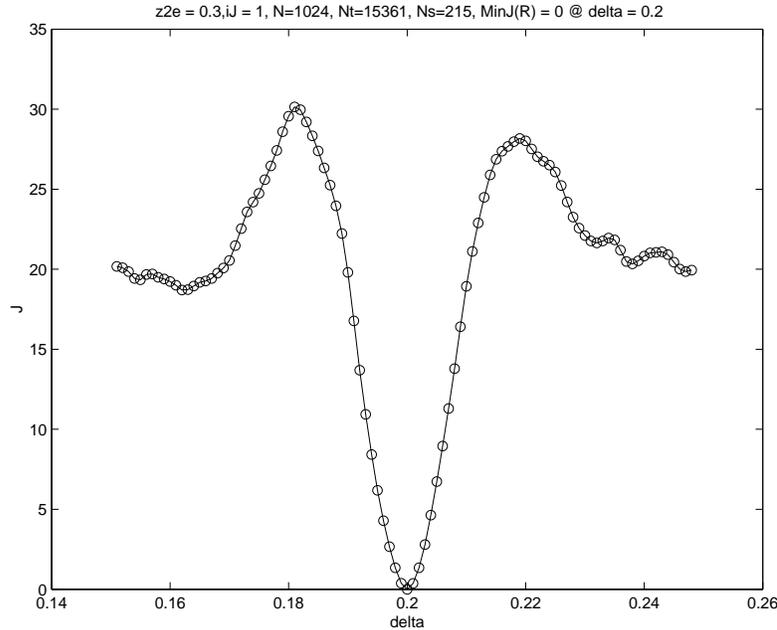


Figure 6: Nonlinear Least Squares objective function versus δ for a small range of δ values (data at $z = 0$ only)

either side of the exact minimizer are due to the simulated solution going in and out of phase with the exact solution. For this example, the simulated solution is most out of phase with the exact solution at $\delta = .181$ and $\delta = .219$ (i.e., approximately $\delta^* \pm \frac{\lambda}{4}$), which correspond to the first and second peak in J respectively. The same phenomenon occurs in the d direction, for the same reasons. See [3] for a thorough demonstration of this behavior.

Very few optimization routines can provide convergence without initial conditions between the two peaks in J . The effective convergence region for this objective function applied to this problem (with or without observations at $z = 1$) is within about 8% of the actual value of δ when d is exact and within about 7.5% of the actual value of d when δ is exact.

Note also that the convergence region is very dependent on the frequency of the interrogating signal; for higher frequencies, the region is even smaller. This is because the distance between the two peaks in J is linearly dependent on λ , the wave length of the interrogating signal. This has profound implications for our desire to interrogate with signals in the Terahertz range. There are also peaks in J with respect to d as well, for the same reasons.

Further, this is considering only a one parameter minimization problem with the other parameter held fixed at the exact solution. Convergence is much worse for the actual two parameter problem. In fact, in the full surface plot of J , a diagonal “trench” occurs approximately along the line

$$d = -\frac{1}{2.3}(\delta - \delta^*) + d^* \quad (11)$$

(where $*$ denotes the exact solution). This phenomena is further explained, and depicted in various figures, in [3].

4.2 An Improved Objective Function

As demonstrated above, the usual Nonlinear Least Squares objective function when plotted with respect to either d or δ has two large peaks in J on either side of the exact minimizer. The reason for these peaks in J is that the simulated solution goes in and out of phase with the data as d or δ change. When they are precisely out of phase, there is a very large absolute error, which when squared, causes the objective function to have large peaks. Therefore, one solution is to not consider the absolute error, but instead the error of the absolute values, i.e., the following objective function:

$$J_2(q) = \frac{1}{2MS} \sum_{j=1}^M \sum_{i=1}^S \left(|E(t_i, z_j^O; q)| - |\hat{E}_{ij}| \right)^2. \quad (12)$$

This non-standard mechanism will prevent the fact that the signals go in and out of phase with each other from having an impact on the objective function, since positive magnitudes cannot cancel each other out. Therefore it gives a more accurate measure of the difference between two signals. Note that the orientation of the interrogating signal (e.g., peak first) precludes the possibility of a solution E from having the same magnitude but opposite sign as \hat{E} . Further, note that $J_2(q)$ is not differentiable on a set of measure zero; this is very unlikely to affect the finite difference computations of the gradients, and did not present problems in our numerical testing. We plot $J_2(q)$ versus δ in Figure 7. See [3] for other plots including surface plots as J_2 varies over d and δ .

Note that while we have effectively eliminated the peaks on either side of the exact solutions, in essence we have merely converted them to local minima! But, since the minima occur for the same reasons the peaks in J had been occurring, they occur at the same values of δ . Note that we can see from plotting the signals that they were exactly out of phase when they were shifted by $\frac{\lambda}{4}$, where λ is the wavelength of the interrogating signal. Therefore δ is off by $\frac{\lambda}{4}$. Since we determine the frequency of the interrogating signal, this is a known quantity, and we can predict where these local minima will occur a priori!

Most optimization routines will continue until they find a local minimum, and since the two false minima described above are at least close to “predictable” locations, we can easily test on either side of any detected minimum to determine if it is in fact global. If J_2 is less at a fixed distance in the δ direction (e.g., $\frac{\lambda}{4}$) on either side of a detected minimum, i.e., at either *test location* of the local minimum, we restart Gauss-Newton

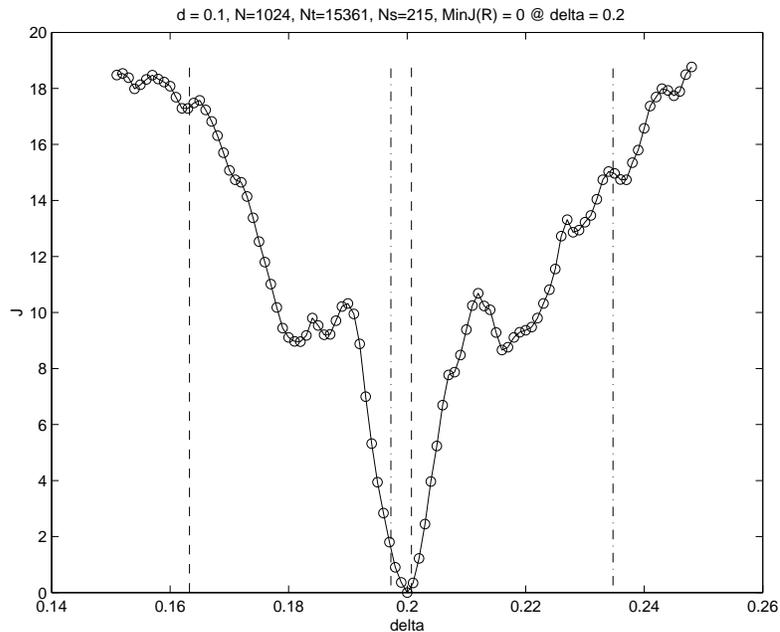


Figure 7: Our modified Nonlinear Least Squares objective function (J_2) versus δ for a small range of δ values. The dotted lines represent the delta values that will be tested if a local minimum is found

at the new best guess. Thus if either of the two local minima described above is found, we will eventually have global convergence if one of their *test locations* are sufficiently close to the global minimizer.

To graphically demonstrate this approach, we have added dotted and dash-dotted lines to the the graph of $J_2(\delta)$ in Figure 7 to represent the test values of the first and second local minima respectively. One can clearly see that if either local minimum is detected, the test value either to its right if it is the first one, or to its left if it is the second, will give a smaller J_2 and should eventually lead to the global minimizer being detected. Using this method we have in principle increased our convergence region to about 25% of δ when d is exact. The same approach works for the d direction, increasing its convergence region from about 7.5% to about 15% when δ is exact.

Other possible modifications to the Least Squares objective function having similar effects include squaring the signal instead of taking the absolute value (thus preserving smoothness everywhere), or just halving the reference signal so that we only have a positive amplitude to begin with. Each of these options will still have the local minima problems described above, as well as their own unique disadvantages.

4.3 Testing J_2

In order to determine the limitations of an optimization routine to minimize our objective function J_2 in a more practical setting we examine J_2 versus q when error is present. In particular we try both adding random noise to the data signal, as well as testing bad initial guesses for δ and d . It should be noted that in the tests reported on below we assume that data at $z = 1$ is not available and used only observations at $z = 0$.

4.3.1 Sensitivity to Initial Guesses

For J_2 described in (12), the objective function is more sensitive to d than to δ , therefore it is imperative that our initial guess for d is as good as possible. To give an idea of what may happen if our d estimate were not within the 15% our testing has determined is necessary, we examined plots of the objective function versus δ for three values of d , which are 3%, 15%, and 30% off respectively (these are displayed in [3]). With errors greater than 15% an erroneous global minimum appears for small δ values. This occurs because

the first reflection of the data is not matched by the simulation, but the second reflection matches it if δ is small enough (see [3] for details and sample plots). It turns out that the distance between the erroneous global minimum and the correct minimum is exactly $\delta^* = 0.2$, which is what would be expected. However, we cannot apply the same idea as before where we add or subtract a fixed amount to test for other local minima, since for one, the “more optimal” of the two is farther from the “true” solution, and also, we would have to know δ in order to add or subtract it (but δ is what we are trying to estimate!).

4.3.2 Random Observation Noise

In order to test the feasibility of this procedure as an estimation method, we have produced synthetic data for our observations \hat{E}_i . In an actual experiment, one must assume that the measurements are not exact. To simulate this we have added random noise to the original signal. The absolute value of the noise is relative to the size of the signal. If E_i is the data sampled, then we define $\hat{E}_i = E_i(1 + \nu\eta_i)$, where η_i are independent normally distributed random variables with mean zero and variance one. The coefficient ν determines the relative magnitude of the noise as a percentage of the magnitude of E_i , in particular, $\nu = 0.05$ corresponds to 10% noise and $\nu = 0.025$ to 5% noise.

Plots of the resulting objective functions for various values of ν ranging from 2% to 40% are shown in [3]. Summarizing these results, we note that the structure of the curves is not significantly affected, nor is the location of the global minimum. However the magnitude of the minimum of the objective function is increased, making Inexact Newton methods slightly less reliable due to the larger residual. Still, our results show that the correct minima were consistently found and within a reasonable amount of time. Select examples are summarized in Table 1. Corresponding initial estimates ranged from, in the $\nu = 0$ case, $(d_0, \delta_0) = (0.093689, 0.20986)$ to, in the $\nu = .2$ case, $(d_0, \delta_0) = (0.109668, 0.172385)$.

ν	d	δ	J	Iterations	CPU Time (s)
0	0.1	0.2	1.32319E-10	7	160
0.01	0.099994	0.199969	0.00792792	8	186
0.05	0.099974	0.199835	0.199489	13	291
0.2	0.099928	0.199204	3.04619	20	435

Table 1: Number of Iterations and CPU Time for Gauss-Newton given various relative magnitudes of random error

5 Problem 2

We next apply the most useful techniques obtained from investigations of Problem 1 to a new formulation of the interrogation problem. In Problem 2 we consider a dielectric slab and a metallic backing (conductor) with a possible gap between the two (see Figures 4 and 8). Applications of this specific formulation included detecting delamination of insulation from metallic containers, e.g., insulating foam on a space shuttle fuel tank. In order for this numerical approach to be useful in this particular application we must be able to resolve a gap of width $.2mm$ inside of a slab with a thickness of at least $20cm$ using a frequency of $100GHz$.

We will again assume the same physical parameters for our dielectric and consider the gap as a vacuum. The variables d and δ are still the depth and the width of the gap respectively. One major difference is that in this problem we are only able to detect the electromagnetic signal in front of the material. Also, since the metallic backing reflects much of the signal, we have considerably more overlapping of the reflections to worry about. These properties contribute to the fact that this formulation leads to a much more difficult inverse problem. For this reason we will be using more sophisticated optimization routines including a Levenberg-Marquardt parameter and Implicit Filtering. We will also need to develop different approximation methods for our initial guesses.

The implementation of this problem has several minor differences from the previous one. First, we now only need to represent two interfaces \tilde{z}_1 and \tilde{z}_2 , with \tilde{z}_0 and \tilde{z}_3 being the front and back computational boundaries, respectively. Thus now we define the depth of the gap as $d := \tilde{z}_2 - \tilde{z}_1$ and the width as

$\delta := \tilde{z}_3 - \tilde{z}_2$. Also, as previously mentioned, the conductive metal backing reflects the signal, and hence we must change our absorbing boundary conditions at $z = 1$ (for a finite computational domain), to an actual fixed, Dirichlet boundary condition ($E = 0$). We must modify our finite element matrices accordingly, as well. Otherwise, the numerical method for simulation is the same as it was for Problem 1, namely standard finite element methods for spatial derivatives, and an alternating implicit/explicit centered difference time stepping scheme. Sample solutions are plotted in Figure 9.

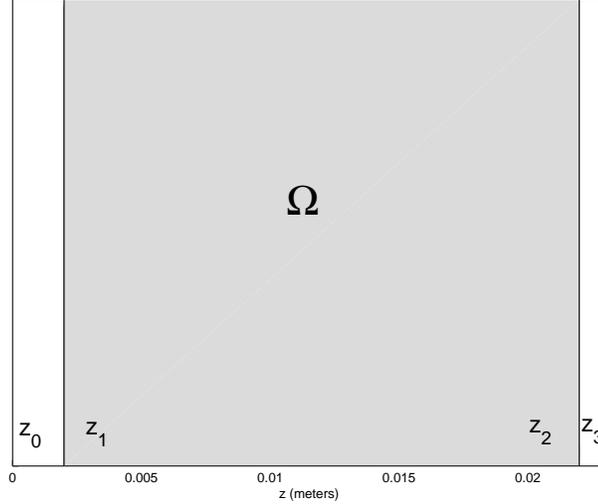


Figure 8: The domain of the material slab with a gap between the medium and a metallic conductive backing: $\Omega = \{z | z_1 \leq z \leq z_2\}$.

We again define our inverse problem to be: find $q := \{d, \delta\} \in Q_{ad}$ such that an objective function representing the error between the simulation and the observed data is minimized:

$$\min_{q \in Q_{ad}} J(q).$$

Here the measurements of the electric field, \hat{E}_i , are taken only at $z = 0$, but still at S distinct times (e.g., every $0.06ps$). The solutions of the simulations, $E(t_i, 0; q)$, are evaluated at the same location and times corresponding to the given data, and using parameter values q . In lieu of actual data from experiments, we again create our observed data by using the simulator, however, the only information that is given to the minimizer is the data observed at $z = 0$, which we will denote by \hat{E} .

The system that we use to model the propagation of the electric field, and thus simulate in order to solve our inverse problem, is as follows, and includes the above mentioned Dirichlet condition at $z = 1$:

$$\begin{aligned} \mu_0 \epsilon_0 (1 + (\epsilon_\infty - 1) I_\Omega) \ddot{E} + \mu_0 I_\Omega \ddot{P} + \mu_0 \sigma \dot{E} - E'' &= -\mu_0 \dot{J}_s && \text{in } \Omega \cup \Omega_0 \\ \tau \dot{P} + P &= \tau \epsilon_0 (\epsilon_s - \epsilon_\infty) E && \text{in } \Omega \\ [\dot{E} - cE']_{z=0} &= 0 \\ [E]_{z=1} &= 0 \\ E(0, z) &= 0 \\ \dot{E}(0, z) &= 0 \end{aligned}$$

with

$$J_s(t, z) = \delta(z) \sin(\omega t) I_{[0, t_f]}(t).$$

See Section 2 for a complete description.

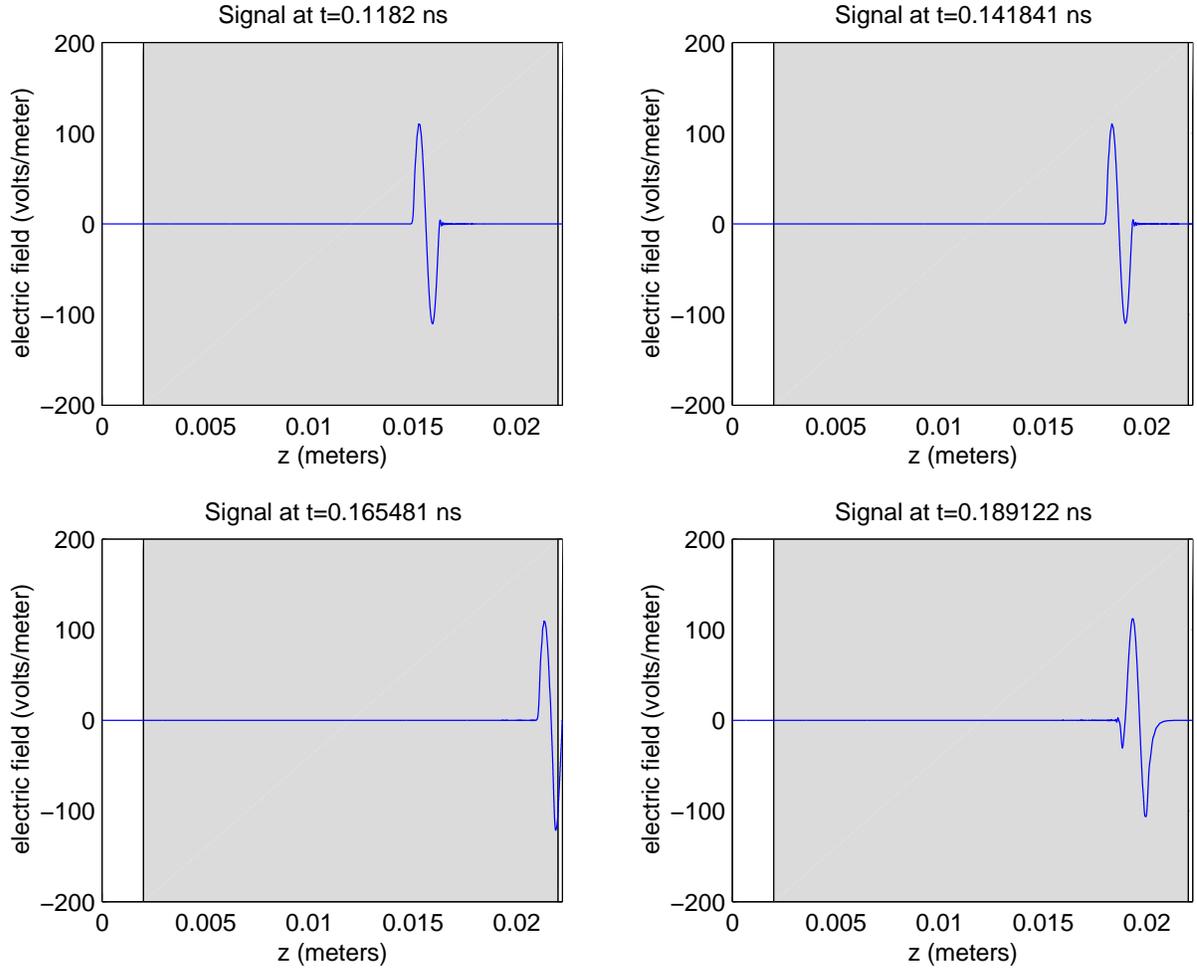


Figure 9: Computed solutions at different times of a windowed electromagnetic pulse incident on a Debye medium with a gap between the medium and a metallic conductive backing. The width of the slab is $d = .02m$ and the width of the gap is $\delta = .0002$ (barely visible at the far right of the gray region).

5.1 Objective Function

As in the previous problem, we encounter difficulties when attempting to use the standard Least Squares objective function to compute the error between the simulated signal and the observed data. The constructive interference of peaks and troughs produces peaks in J in the objective function on all sides of the global minimum which make it nearly impossible to find the solution in the middle. The peaks in J are clearly apparent in Figure 10. In contrast, Figure 11 shows a surface plot of our modified least squares objective function

$$J_2(q) = \frac{1}{2S} \sum_{i=1}^S \left| |E(t_i, 0; q)| - |\hat{E}_i| \right|^2.$$

It is clear, as before, that the initial guess is crucial to the success of any optimization routine. Notice that although J_2 does not exhibit the familiar peaks in J of J_1 , it does however still have many local minima, which are just as difficult to avoid in a minimization routine.

The local minima in J_2 for this problem occur approximately every $\frac{\lambda}{4}$ along the line

$$d = -\frac{1}{\sqrt{\epsilon_\infty}}\delta + b.$$

This happens for the same reason as in Problem 1 (see [3] for details and illustrations). Because we cannot eliminate these local minima, we must appeal to the procedure that worked in the previous problem, namely testing “check points”. Since we know where these local minima are occurring with respect to the global minimum, if our minimization routine finds what it suspects to be a local minima, say (d_1, δ_1) , we simply check $(d_1 \pm \alpha \frac{\lambda}{4}, \delta_1 \mp \alpha \sqrt{\epsilon_\infty} \frac{\lambda}{4})$, where $\alpha = 1/\sqrt{1 + \epsilon_\infty}$. If we find a lower objective function value, we restart our optimization routine at that “check point”.

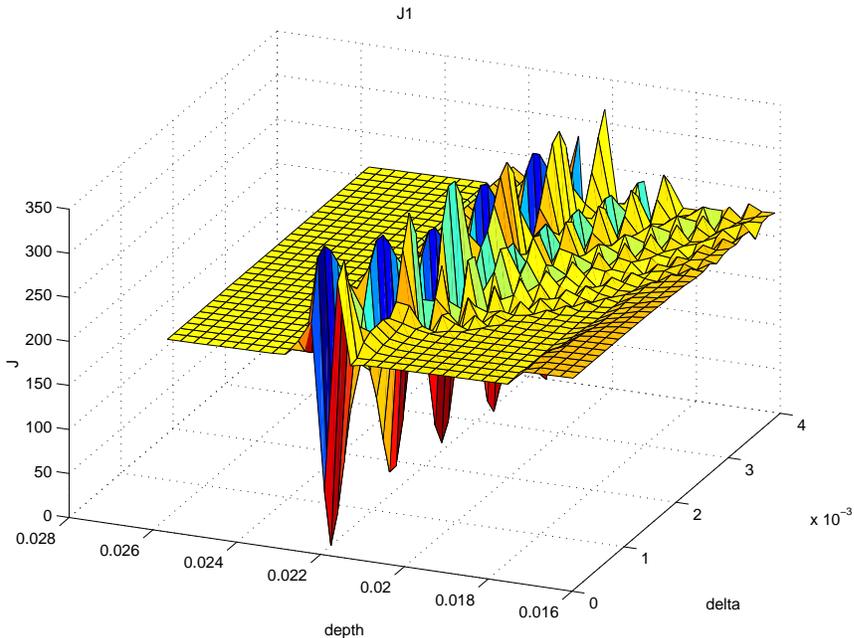


Figure 10: Close up surface plot of Least Squares objective function demonstrating peaks in J , and exhibiting many local minima.

5.2 Initial Guesses

In spite of our faith in the “check point” method, we still desire to find the best initial guesses for our optimization routine as possible so that we may hopefully find the global minimum without restarting. As before, we use the travel time of the first trough to approximate the location of the first interface. However, in this formulation we can take advantage of some of the characteristics of the signals. For example, the first reflection off of the gap is always trough-first, and the second (as well as each subsequent reflection) is always peak-first. For this reason, if we want to locate the first trough we can simply find the largest peak (belonging to the second reflection) and back track. It is a very simple matter to find a maximum or minimum of a vector of values. After the location of the largest peak is found, we back track to find the minimum in front of it, namely that belonging to the first reflection off of the gap. Then using the procedure described in Section 4.1, we approximate the root immediately in front of this trough. That gives us the travel time for the first reflection off of the gap, which in turn gives us the depth d of the gap.

Finding δ is, unfortunately, not nearly as straightforward. There are two main possibilities, and therefore, two differing approaches to approximating δ , depending on the nature of the reflected signal. We consider the two cases:

- (i) The leading trough of the first reflection and the second reflection are disjoint (i.e., $\delta > \frac{\lambda}{8}$). In this case we can find the locations of the peak and trough and use the travel time between the two to approximate δ . We denote this approximation by δ_1 . (Note that the observed peak is not necessarily the same as the original peak unless $\delta > \frac{3\lambda}{8}$, but it is still a good approximation). See Figure 12.

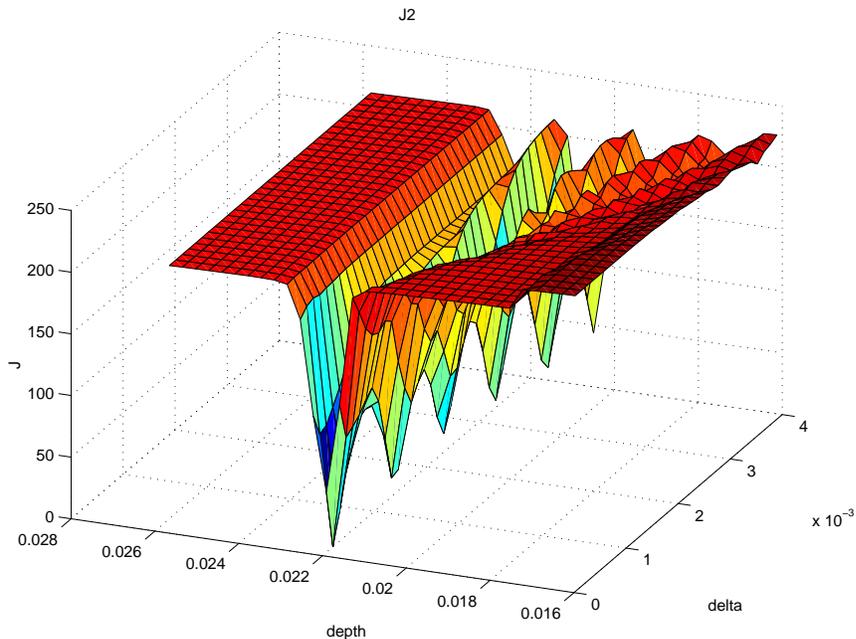


Figure 11: Close up surface plot of modified Least Squares objective function demonstrating lack of peaks in J , but exhibiting many local minima.

- (ii) The second reflection partially truncates the trough of the first (i.e., $\delta < \frac{\lambda}{8}$). As a rough approximation, we can assume that the location of the actual minimum (trough) is where the two signals begin to interfere with each other (the observable minimum). See Figure 13. We denote this approximation by δ_4 .

A more accurate method is to use triangles to approximate the two reflections. By knowing the location of the maximum and minimum (peak and trough, respectively), and also the beginning of the first signal (from Section 4.1) and the rough approximation to the beginning of the second signal using δ_4 , we can estimate the slopes of the two triangles with finite differences. Also note that since the two signals are added, the observed root between the peak and trough in the combined signal is actually an equilibrium point between the two signals. By setting equal to each other the two linear approximations for each of the two signals, evaluated at the equilibrium point, we can solve for the distance between the starting point of each signal, and thus for δ_3 . See Figure 14. Specifically, let (p_1, q_1) be the location of the trough of the combined signal and (p_2, q_2) be the location of the peak. Let r_1 be the location of the root in front of the trough, and r_2 be the root between the trough and peak. Estimate the slope of the first signal, $m_1 < 0$, using (p_1, q_1) and r_1 . Now if we let $y = r_2 - r_1$ and say x is the actual distance between r_1 and the beginning of the second signal, then setting the linear approximations equal in magnitude, but opposite in sign, at r_2 yields

$$-m_1 y = m_2 (y - x).$$

Now we can estimate the slope of the second signal, $m_2 > 0$, using (p_2, q_2) and $(r_2, -m_1 y)$. Also, we can re-write the above equation as

$$x = \left(\frac{-m_1 + m_2}{m_2} \right) y.$$

To find δ_3 we simply divide x by 2 and the (scaled) speed of light in the material, i.e., $\sqrt{\epsilon_\infty}$.

Since each of the two situations above is dependent on the parameter it is approximating, we must also determine which of the above methods is most appropriate to use. Thus we use the most precise of the available methods to determine the situation, i.e., δ_4 , instead of δ_3 since in general δ_3 underestimates δ so we

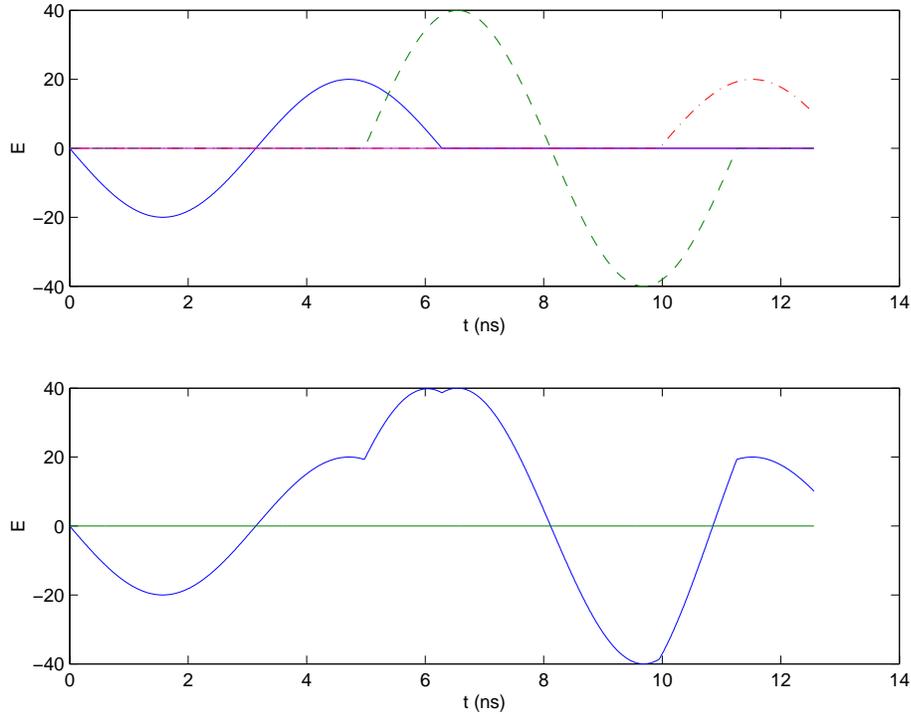


Figure 12: The top plot represents several signals which may be observed in a simulation of Problem 2. The bottom plot shows the sum of the top signals. The peak of the second signal is just beginning to be obscured by the first when δ becomes less than $\frac{3\lambda}{8}$. Thus the observable maximum is still a good approximation of the peak of the second signal, and a trough to peak distance can be used to estimate δ .

do not want to use it as a criterion for determining whether δ is small. (Note that when δ is indeed small, δ_3 is more accurate than δ_4 .) The estimate for δ_4 tends to be an overestimate, and is only valid if $\delta < \frac{\lambda}{8}$. Unfortunately, δ_1 also tends to be an overestimate, so we prefer to only trust it entirely if it is larger than $\frac{\lambda}{4}$. If neither δ_1 nor δ_3 is a sufficient approximation we choose to use the average of the two, and call it δ_2 .

Therefore our algorithm for approximating δ is as follows:

- (a) If $\delta_4 < \frac{\lambda}{8}$ then use δ_3
- (b) else if $\delta_1 > \frac{\lambda}{4}$ then use δ_1
- (c) else use δ_2 (average between δ_1 and δ_3).

We tested our approximating methods on exact depth (d) values of: .02, .04, .08, .1, and .2 m , and values of width (δ): .0001, .0002, .0004, .0006, and .0008 m . Since $\frac{\lambda}{8}$ is the transition point between the two situations, it is understandable why δ close to this value is the most difficult to accurately resolve. We chose this range of δ 's because our choice of frequency gives $\frac{\lambda}{8} = 3.7475 \times 10^{-4}m$. See the tables in [3] for the initial estimates of d and δ .

The approximations improve slightly as the number of finite elements is increased, and appeared to converge to fixed values. This suggests that numerical error (and instability) can affect the estimates. For each case there is a significant amount of visible numerical error in the simulations below a certain number of elements, therefore in approximating δ we chose to use the number of elements just above the threshold.

While the initial estimates were relatively inaccurate, some δ approximations being almost 100% off from the true solution value, in the numerical tests we performed, all the initial estimates were sufficiently close to the true, global minima as to not cause the optimization routine to result in a false, local minima. While our “check point” method is available if needed, it is much more efficient to have an accurate initial estimate

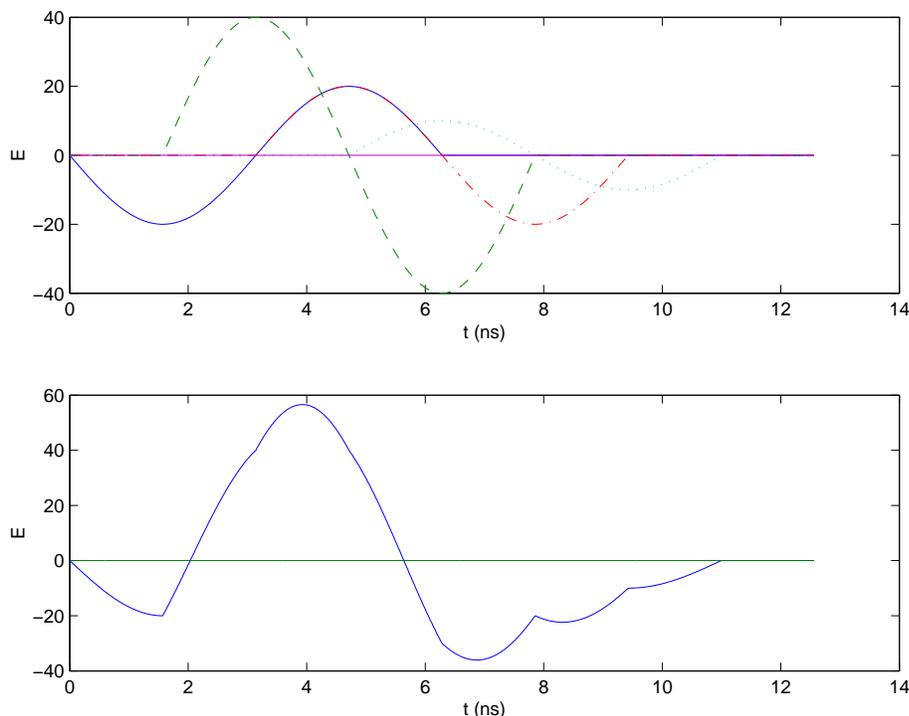


Figure 13: The top plot represents several signals which may be observed in a simulation of Problem 2. The bottom plot shows the sum of the top signals. The trough of the first signal is partially truncated by the second signal. In this case the observed minimum is still a good approximation to where the second signal begins. For smaller δ , a linear approximation must be used.

than to restart after optimizing from a bad one. Still, the report [3] describes several very real examples where the “check point” method would be a necessary last resort.

5.3 Optimization Method

Now that we have approximated our initial guesses, we need to minimize the objective function in order to solve the inverse problem. In Problem 1, Gauss-Newton was sufficient to find the global minimum for most cases. In this formulation, however, we will apply more sophisticated methods, reverting to Gauss-Newton whenever possible since its convergence rate is best.

The first modification we make to Gauss-Newton is to add a Levenberg-Marquardt parameter, ν_c (see [7]). The Inexact Newton step becomes

$$s_c = - (R'(q_c)^T R'(q_c) + \nu_c I)^{-1} R'(q_c)^T R(q_c).$$

The parameter adds regularization by making the model Hessian positive definite. The method uses a quadratic model Hessian, and also has a built-in line search with a sufficient decrease condition. The line search is based on the predicted decrease computed from the quadratic model. If the actual improvement of the objective function, J , is close to the amount predicted by the model Hessian after a step is taken, then the method decreases the Levenberg-Marquardt parameter, ν_c , effectively increasing the relative size of the next step, which hopefully accelerates the convergence. As ν_c is decreased to 0 the method becomes Damped Gauss-Newton (meaning Gauss-Newton with a line search). If, however, the actual improvement of J after a step is not sufficient (or is even negative), ν_c is increased, effectively scaling back the Newton step, and we retest. If there are too many reductions then we declare a “line search failure” meaning that too small a step is required to decrease the objective function.

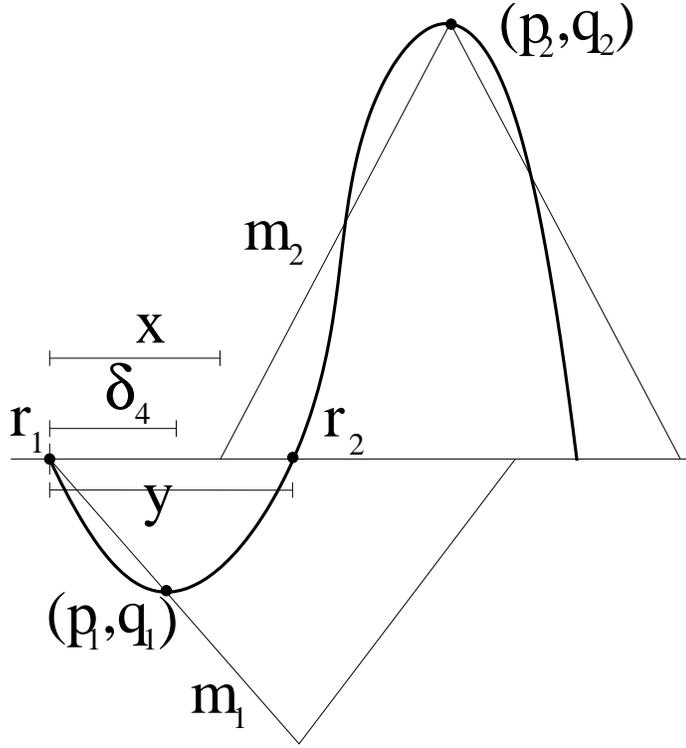


Figure 14: This schematic shows the roots, extrema, distances, and slopes used in the computation of δ_3 .

Usually a method would exit after a line search failure, returning the best approximation so far. But we use this failure to call an adaptive mesh size routine, i.e., an Implicit Filtering technique. The idea is that the failure is likely due to the fact that the direction the finite difference gradient chose is probably not an actual “descent direction” in the global sense. In other words, the finite differencing is most likely differentiating noise. In the same manner that a smooth surface may look rough under a microscope, using too small of a differencing step amplifies effects from round-off error and other sources of numerical noise. Our technique is to increase the relative differencing step, \hat{h} , recompute the gradients, and then try the Levenberg-Marquardt method again. The relative differencing step, \hat{h} , is such that the gradient, $\nabla_{\hat{h}}$, of $J(q) = J([d, \delta])$ is computed with

$$\nabla_{\hat{h}} J([d, \delta]) = \begin{bmatrix} \frac{J((1+\hat{h})d, \delta) - J(d, \delta)}{\hat{h}d} \\ \frac{J(d, (1+\hat{h})\delta) - J(d, \delta)}{\hat{h}\delta} \end{bmatrix}$$

We apply a similar approach to modifying the differencing step \hat{h} as we do for changing ν_c in that after a successful step we decrease \hat{h} , but if we have another failure we increase \hat{h} even more. Since the convergence rates of gradient based methods are dependent on the size of \hat{h} (for example Gauss-Newton is $O(\hat{h}^2)$), we want \hat{h} to be as small as possible and still be effective, similarly with ν_c . We use a three tiered approach to changing \hat{h} . Initially we set $\hat{h} = 10^{-9}$. To increase \hat{h} we raise it to the $\frac{2}{3}$ power, to decrease we raise it to the $\frac{3}{2}$ power. Additionally we define 10^{-4} to be the maximum allowable differencing step value. Thus $\hat{h} \in \{10^{-9}, 10^{-6}, 10^{-4}\}$.

In general an optimization method exits with “success” if the norm of the current gradient is less than tol times the norm of the initial gradient. However, in our method we do not immediately trust the finite difference gradients, and instead call Implicit Filtering again when the gradients appear small. When we have verified small gradients on all three scales (the various values of the differencing step \hat{h} defined above), then we exit with “success”.

Remark 2 *In practice, a very good solution is found within a couple of Levenberg-Marquardt steps, and then an equal number of Implicit Filtering iterations verify, and sometimes enhance, this solution. In the interest of efficiency, and since this is a parameter identification problem, we exit early with “success” if our objective function is satisfactorily small (i.e., `tol` times the initial value), which can save on average about half of the possible iterations.*

Additionally we impose a restriction on the number of “pullbacks” on each linesearch, and on the number of iterations, effectively limiting the total number of function calls. If a small gradient has not been verified on all scales before exhausting the maximum number of iterations, we exit with “failure”.

5.4 Numerical Issues

For small N the difficult cases are those with large depth. This is because the computational domain is effectively increased when the depth is increased, making the mesh sizes larger and increasing the level of numerical error. The magnitude of δ does not seem to have a significant effect on the convergence of the method.

An obvious disadvantage to having a large N is that each simulation takes much longer. In general the total execution time is quadrupled when the number of elements is doubled. This is consistent with the fact that complexity of the most time consuming part of the simulation, the linear solves, is $O(N)$, and the number of time steps N_t is also $O(N)$. So when we double the number of finite elements we are also doubling the number of time steps. Therefore, we get an overall complexity of $O(N^2)$. Thus, as mentioned before, in our inverse problem we choose to use the number of elements just above the threshold of when numerical error is apparent.

We should also mention that in order to create data, in lieu of actual experimental data, we perform a simulation at a higher resolution believing it to be more accurate. Specifically, we double the number of finite elements. Since the time step, and therefore the effective sample rate if the time step is too large, are both dependent upon the mesh size (refer to [3]), the sample times of the simulated data do not necessarily correspond with the sample times of the simulations at the lower resolution. (In general we have twice as many samples from the higher resolution.) Thus in order to compute the modified least squares error between the two vectors, we perform a linear interpolation of the simulated data onto the sample times at the lower resolution. See Figure 15. Note that in the usual case where we simply have twice as many sample points from the higher resolution simulation, we are in effect discarding sample points rather than doing a true interpolation.

For comparison we compute the low resolution simulation using the values d^* and δ^* (note that this is not the same as taking the high resolution simulation and interpolating it onto the low resolution time steps, which we actually use as our observed data). In every case that we have tested, J , when computed with the d and δ values found from the optimization routine (d_{min} and δ_{min}), is less than or equal to J when computed with the original values (d^* and δ^*). This suggests that an actual global minimum of the objective function has been found, even though the final estimates of d and δ themselves are not necessarily equal to d^* and δ^* . Note in Figure 15 that the simulation using original values, (d^* , δ^*), is in fact closer to the original data, but the simulation using the minimizer values, (d_{min} , δ_{min}), is closer to the interpolated data (see for example the [.335, .3352] interval).

Although we could compute our optimization routine at the same resolution as the simulated data to get a better fit in our tests, this would not properly represent the real-life phenomenon of sampling data. Sampled data is inherently not a completely accurate representation of a physical observation. We believe that our interpolation approach gives a more realistic expectation of how our method would perform given actual experimental data. In order to further test the robustness of our inverse problem solution method we introduce random noise to the detected data in Section 5.5.

5.5 Numerical Results

Tables 2 and 3 show the final computed approximations for the depth of the slab (d_{min}) and the width of the gap behind it (δ_{min}). The relative differences from the original values used to generate the data (d^* and δ^*), are: for depth, on the order of .0001 and for δ , on the order of .01. However, this does not imply that the optimization routine was unable to find the optimal solution. Recall that since our data is generated

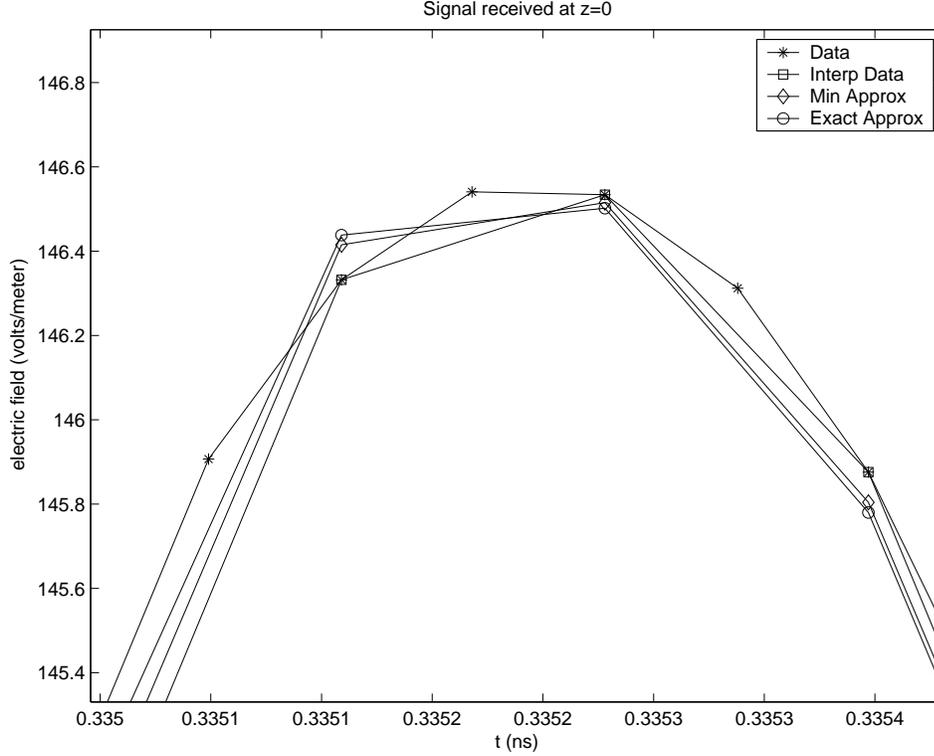


Figure 15: Plotted are the actual simulated data ($N = 2048$), the interpolation of the simulated data onto the low resolution sample times ($N = 1024$), the result of the minimization routine ($N = 1024$), and a low resolution ($N = 1024$) simulation using the exact values of d and δ .

with essentially a different simulator than our forward solves, the original values do not necessarily minimize the objective function. The objective function values give a better indicator of how well the optimization routine works since it shows the fit to the generated data. Table 4 shows the final objective function values. In each of these cases, the final objective function value (J_{min}) was less than $J^* := J(q^*)$. In fact, the ratios $J_r := J_{min}/J^*$ were on average .3008. We consider any $J_r < 1$ to represent a successful convergence.

Although δ values that are near $\frac{\lambda}{8} = 3.7475 \times 10^{-4}m$ are the most difficult for which to obtain initial approximations, we see that the objective function values in these cases are just as small (and the final estimates are just as close) as for other δ values.

The execution time, in seconds, as well as the number of function calls, are given in [3]. While the above tables establish that we were actually able to resolve the case of 20cm depth, there was a price we had to pay. The average execution times for each of different mesh sizes ($N = 1024, 2048, 4096, 8192$, and 16384) were 39, 248, 1452, 6229, and 35509 seconds, respectively. Each represents an increase in time over the previous mesh size by a factor of 6.4, 5.9, 4.28, and 5.7, respectively. This is consistent with the fact that the forward solves are order $O(h^2)$. However, the additional sample points for the larger N cases allowed for smaller initial objective function values which resulted in increasingly more iterations to satisfy the relative tolerance in our stopping criteria. This explains why we do not see ratios closer to the expected 4 for order $O(h^2)$ methods.

5.5.1 Relative Random Noise

We add random noise to the signal, as mentioned above, in order to more closely simulate the experimental process in data collection. As in Section 4.3.2, we start with relative noise where the absolute value of the

d		δ				
		.0001	.0002	.0004	.0006	.0008
.02	(N=1024)	0.0200053	0.0200022	0.0200006	0.0200005	0.0200002
.04	(N=2048)	0.0399948	0.0399974	0.0400005	0.0400005	0.0399999
.08	(N=4096)	0.0799973	0.0799987	0.0800006	0.0800006	0.0800003
.1	(N=8192)	0.0999945	0.0999974	0.1	0.1	0.0999999
.2	(N=16384)	0.200011	0.200005	0.2	0.2	0.200001

Table 2: The final estimates of d .

d		δ				
		.0001	.0002	.0004	.0006	.0008
.02	(N=1024)	9.40622e-05	0.000196754	0.000398642	0.000597275	0.00079707
.04	(N=2048)	0.000106435	0.000203916	0.000394204	0.000592156	0.000793622
.08	(N=4096)	0.000103585	0.000202273	0.000395791	0.000593861	0.000794401
.1	(N=8192)	0.000106593	0.000203876	0.000396203	0.000594976	0.000795985
.2	(N=16384)	8.7456e-05	0.000191808	0.00040297	0.000602902	0.00080129

Table 3: The final estimates of δ .

d		δ				
		.0001	.0002	.0004	.0006	.0008
.02	(N=1024)	0.00786171	0.00906699	0.0115657	0.0233783	0.0447687
.04	(N=2048)	0.021516	0.0343314	0.0514108	0.0700747	0.0927117
.08	(N=4096)	0.0116105	0.0145428	0.0201004	0.0272513	0.0344458
.1	(N=8192)	0.00304723	0.00547532	0.00779186	0.00931778	0.0118529
.2	(N=16384)	0.000609258	0.00133978	0.00146975	0.000962975	0.000766141

Table 4: The objective function value of the final estimates.

noise is proportional to the size of the signal. If E_i is the data sampled, then we define $\hat{E}_i = E_i(1 + \nu_r \eta_i)$, where η_i are independent normally distributed random variables with mean zero and variance one. Again, the coefficient ν_r determines the relative magnitude of the noise as a percentage of the magnitude of E_i , in particular, $\nu_r = 0.01$ corresponds to 2% noise. We tested relative magnitude levels of 2%, 10%, and 20% (corresponding to $\nu_r = .01, .05$, and $.1$ respectively). See [3] for tables of initial estimates. In nearly all the cases the estimate was close enough for the optimization method to converge ($J_r < 1$) to the expected minimum. The only exceptions were with $\nu = .1$ and $\delta = .0004$, which are understandably the most difficult cases.

The final approximations d_{min} and δ_{min} in the presence of noise are also given in [3]. Some approximations with high noise appear to be better approximations than some with little or no noise. For example, with $\delta^* = .0001$, $d^* = .04$, the $\nu_r = .1$ final approximations are an order of magnitude closer to the original values than the $\nu_r = 0$ final approximations. This is not to say that the noise helps the approximation method. Rather, it is for the same reason that, for example, as shown in Figure 15, the actual parameter values produced a signal farther away (in the Least Squares sense) from the generated data than a signal computed with the approximated parameter values. The resulting objective function values give a better indication of the accuracy of the approximation to the data. The final objective function values corresponding to $\nu_r = 0$ were two orders of magnitude smaller on average than those resulting from $\nu_r = .1$. Thus, it is clear that the data without noise is more accurately matched by its approximations than those with noise.

5.5.2 Standard Error Analysis

In an actual inverse problem using data collected by experiment, one desires to have confidence intervals on all parameter estimates. We will apply standard error techniques to an Ordinary Least Squares (OLS)

formulation of our problem to obtain confidence intervals on our estimates. In order to rewrite our objective function in an OLS formulation, we define $y(t; q) = |E(t, 0; q)|$ to be our estimate to $\hat{y} = |\hat{E}|$, which is the data we are trying to fit by determining $q = (d, \delta)$. Now it is clear that our objective function can be written in the standard OLS form

$$J(q) = \frac{1}{N_s} \sum_{i=1}^{N_s} |y(t_i; q) - \hat{y}_i|^2.$$

For simplicity of terminology, in this section alone, we will refer to $|\hat{E}_i|$ as the data and to $|E(t_i, 0; q)|$ as the simulations.

With the relative random noise described above we do not have constant variance, as is demonstrated in Figures 16 and 17. Here we have plotted the residual $r_i := |E(t_i, 0; \hat{q}_{OLS})| - |\hat{E}_i|$ against time, t_i , and also against $|E(t_i, 0; \hat{q}_{OLS})|$. As one would expect with noise that is relative in size to the signal value, we have a pattern in Figure 16 that follows the pattern of the original signal. Figure 17 demonstrates the fan shape associated with noise that is dependent upon the size of the signal, i.e., non constant variance.

Since constant variance is most conveniently assumed in standard error analysis, we further consider estimates obtained from an inverse problem applied to data with constant variance random noise added. In particular, the data we now consider is generated by

$$\hat{E}_j = E(t_j, 0; q^*) + \beta \nu_r \eta_j$$

where

$$\eta_j \sim \mathcal{N}(0, 1)$$

and the constant β is a scaling factor chosen simply so that the noise level, ν_r , will somewhat correspond to the parameter ν_r used in the previous section on relative noise. Specifically, $\beta = \max_i \hat{E}_i / 10$ ensures that J^* in the constant variance cases is on the same order of magnitude as those in the relative noise cases above for all choices of d and δ that we have considered.

The variance of this data is

$$\sigma^2 = \mathcal{E}[\beta^2 \nu_r^2 \eta_j^2] = \beta^2 \nu_r^2 \mathcal{E}[\eta_j^2] = \beta^2 \nu_r^2$$

where \mathcal{E} denotes the expectation. Therefore, we do have constant variance. Note further the resulting lack of patterns in Figures 18 and 19. The suspicious looking phenomenon of many points on the line $E = 0$ is simply because in the original data E is very close to zero most of the time. Figure 20 demonstrates graphically the difference between relative noise and constant variance noise. The relative noise case is particularly difficult in our inverse problem since most of our initial estimates are based on accurately determining the peak locations, yet this is exactly where most of the relative noise is concentrated.

With constant variance, and further, assuming that each η_j is identically independently (normally) distributed, we have that (see [4]) in the limit as $N_s \rightarrow \infty$

$$\hat{q}_{OLS} \sim \mathcal{N}_2 \left(q_0, \sigma_0^2 [\mathcal{S}^T(q_0) \mathcal{S}(q_0)]^{-1} \right).$$

Here $\mathcal{S}(\hat{q}) = \frac{\partial |E|}{\partial q}(\hat{q})$ which is an $N_s \times 2$ matrix since $q = (d, \delta)$ and $|E|$ is evaluated at N_s sample times. Also, the scale parameter σ_0^2 is approximately given by

$$\sigma_0^2 = \frac{1}{N_s - 2} \sum_{i=1}^{N_s} \left(|E(t_i, 0; q_0)| - |\hat{E}_i| \right)^2.$$

In the above equations, q_0 denotes the theoretical ‘‘true’’ value of the parameter that best describes the system from which the data is taken. Note that in this case, this is not necessarily the same as q^* since the method used to generate the data is different from the forward solve simulator. Therefore q_0 is generally unknown even in examples with simulated data.

As demonstrated in the previous sections, our \hat{q}_{OLS} is often a better minimizer than even the original value of q^* , therefore we will approximate q_0 in the above equations by \hat{q}_{OLS} . In particular, if we denote the covariance matrix as $C_0 = \sigma_0^2 [\mathcal{S}^T(q_0) \mathcal{S}(q_0)]^{-1}$, then we will approximate C_0 by $C = \sigma_{OLS}^2 [\mathcal{S}^T(\hat{q}_{OLS}) \mathcal{S}(\hat{q}_{OLS})]^{-1}$, where

$$\sigma_{OLS}^2 = \frac{1}{N_s - 2} \sum_{i=1}^{N_s} \left(|E(t_i, 0; \hat{q}_{OLS})| - |\hat{E}_i| \right)^2.$$

We compute σ_{OLS}^2 by multiplying our J_{min} by an appropriate conversion factor, since they are defined in a similar manner. However, in order to compute the partial derivatives with respect to d and δ in \mathcal{S} we employ forward differencing, which requires an additional forward simulation for each q_j . For $\hat{q} = \hat{q}_{OLS}$ we have, for example

$$\mathcal{S}_{i1} = \frac{\partial |E|}{\partial q_1}(t_i, 0; \hat{q}) \approx \frac{|E(t_i, 0; [\hat{q}_1, \hat{q}_2])| - |E(t_i, 0; [(1 - h_d)\hat{q}_1, \hat{q}_2])|}{h_d \hat{q}_1}$$

and similarly for each \mathcal{S}_{i2} . In our computations we used the relative differencing factor of $h_d = 1 \times 10^{-4}$. One could also use a sensitivity equations approach (e.g., see [1] and the references therein), but since the variational equations are quite difficult to solve for this example, we choose instead to approximate the partials with respect to q directly with our simulations.

We also need to point out that while taking the absolute value of a function limits differentiability at a small number of points, the derivative does exist almost everywhere. The absolute value function does not change the magnitude of the derivative where it exists, which is what we need to compute the dot product of \mathcal{S} with itself. By using finite differences to estimate derivatives, we are essentially under-estimating at the discontinuities. Under-estimating a few points out of thousands is not going to significantly change our covariance matrix. (Alternatively, one could have defined the objective function by squaring the signals instead of taking absolute values to avoid this problem. In this research we were interested in comparing J_1 and J_2 in previous sections above and changing the scale of E by squaring it would have prevented this.)

With \mathcal{S} calculated, we can now evaluate $C = \sigma_{OLS}^2 [\mathcal{S}^T(\hat{q}_{OLS})\mathcal{S}(\hat{q}_{OLS})]^{-1}$. Then the standard error for $q_1 = d$ is estimated by $\sqrt{C_{11}}$ while the standard error for $q_2 = \delta$ is estimated by $\sqrt{C_{22}}$. See Tables 5 through 12 for confidence intervals relating to various d^* , δ^* and ν_r values. For example, in the case of $d^* = .02$, $\delta^* = .0002$ and with $\nu_r = .01$ our covariance matrix is

$$C = \begin{bmatrix} 2.37122 \times 10^{-15} & -4.43815 \times 10^{-15} \\ -4.43815 \times 10^{-15} & 9.1829 \times 10^{-15} \end{bmatrix}$$

which results in the confidence intervals $d \in (2.00004 \pm 4.86952 \times 10^{-6}) \times 10^{-2}$ and $\delta \in (1.9941 \pm 0.000958274) \times 10^{-4}$.

The width of these bounds are $\pm 0.000243471\%$ and $\pm 0.0480555\%$ of the approximation value respectively. For the $d^* = .02$ case, the average size of the confidence intervals for $\nu_r = .01, .05, .1$ respectively were $\pm .0002\%, \pm .001\%, \pm .002\%$ (averaged over various δ^* values ranging from .0001 to .0008). It is interesting that the widths of the confidence intervals nearly exactly double, on average, when the noise level doubles. For the $d^* = .04$ case the average size of the confidence intervals were $\pm .0001\%, \pm .0006\%, \pm .001\%$. Likewise, when the widths of the confidence intervals for $\delta^* = .0002$ are averaged over several various d^* values (.02, .04, .08) we get $\pm .05999\%, \pm .2883\%, \pm .5718\%$ for $\nu_r = .01, .05, .1$ respectively. For $\delta^* = .0004$ the averages are $\pm .03331\%, \pm .1575\%, \pm .3154\%$. In general, larger d^* and δ^* values have smaller (tighter) confidence intervals. This suggests that the approximations found in these cases are better than those estimating small parameters. While this is intuitive, it is not apparent looking at the estimates themselves or even the final objective function values (see, for example, Table 4).

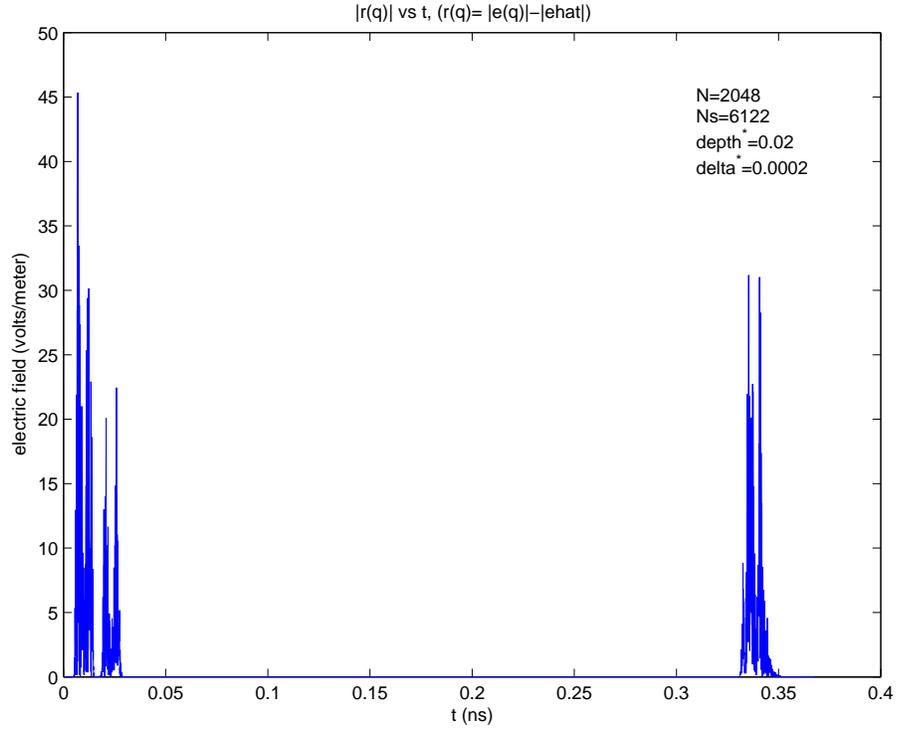


Figure 16: Plots of the absolute value of the residual $r_i = |E(t_i, 0; \hat{q}_{OLS})| - |\hat{E}_i|$ versus time t_i when the data contains relative random noise.

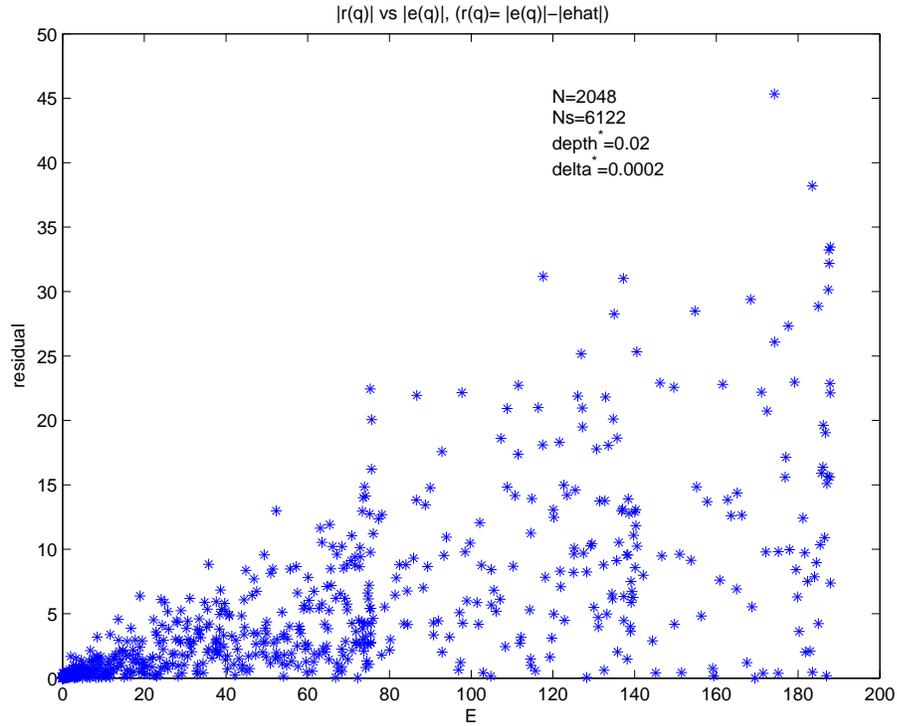


Figure 17: Plots of the absolute value of the residual $r_i = |E(t_i, 0; \hat{q}_{OLS})| - |\hat{E}_i|$ versus the absolute value of the electric field $E(t_i, 0; \hat{q}_{OLS})$ when the data contains relative random noise.

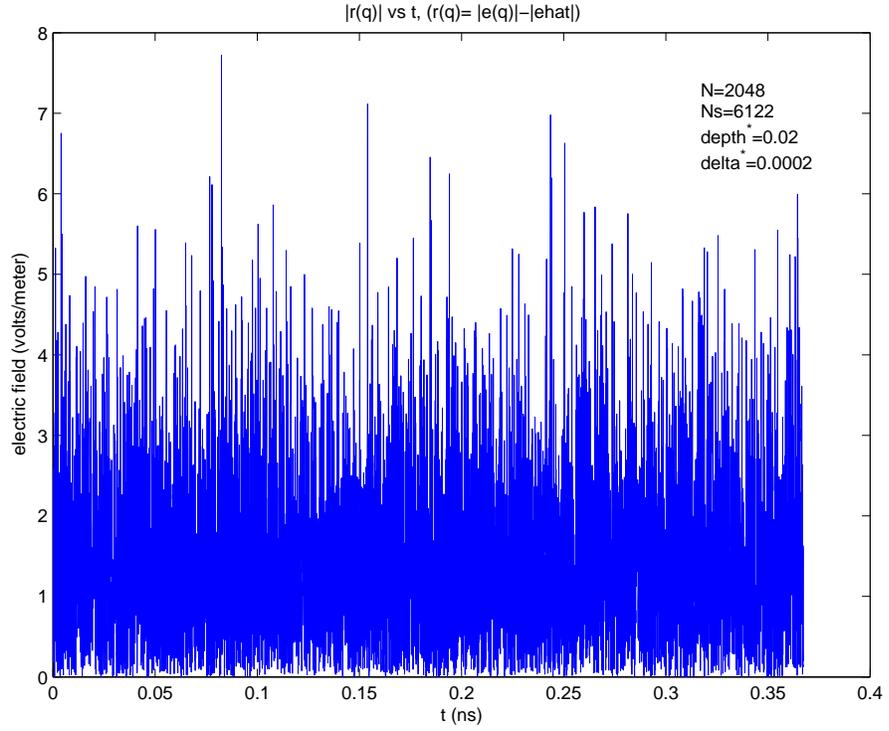


Figure 18: Plots of the absolute value of the residual $r_i = |E(t_i, 0; \hat{q}_{OLS})| - |\hat{E}_i|$ versus time t_i when the data contains constant variance random noise.

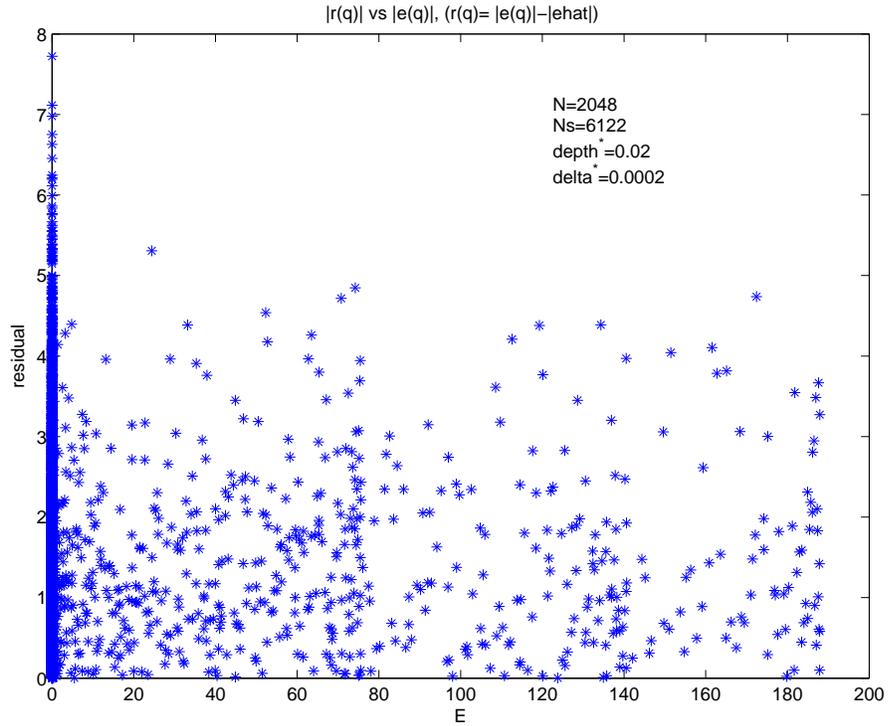


Figure 19: Plots of the absolute value of the residual $r_i = |E(t_i, 0; \hat{q}_{OLS})| - |\hat{E}_i|$ versus the absolute value of the electric field $E(t_i, 0; \hat{q}_{OLS})$ when the data contains constant variance random noise.

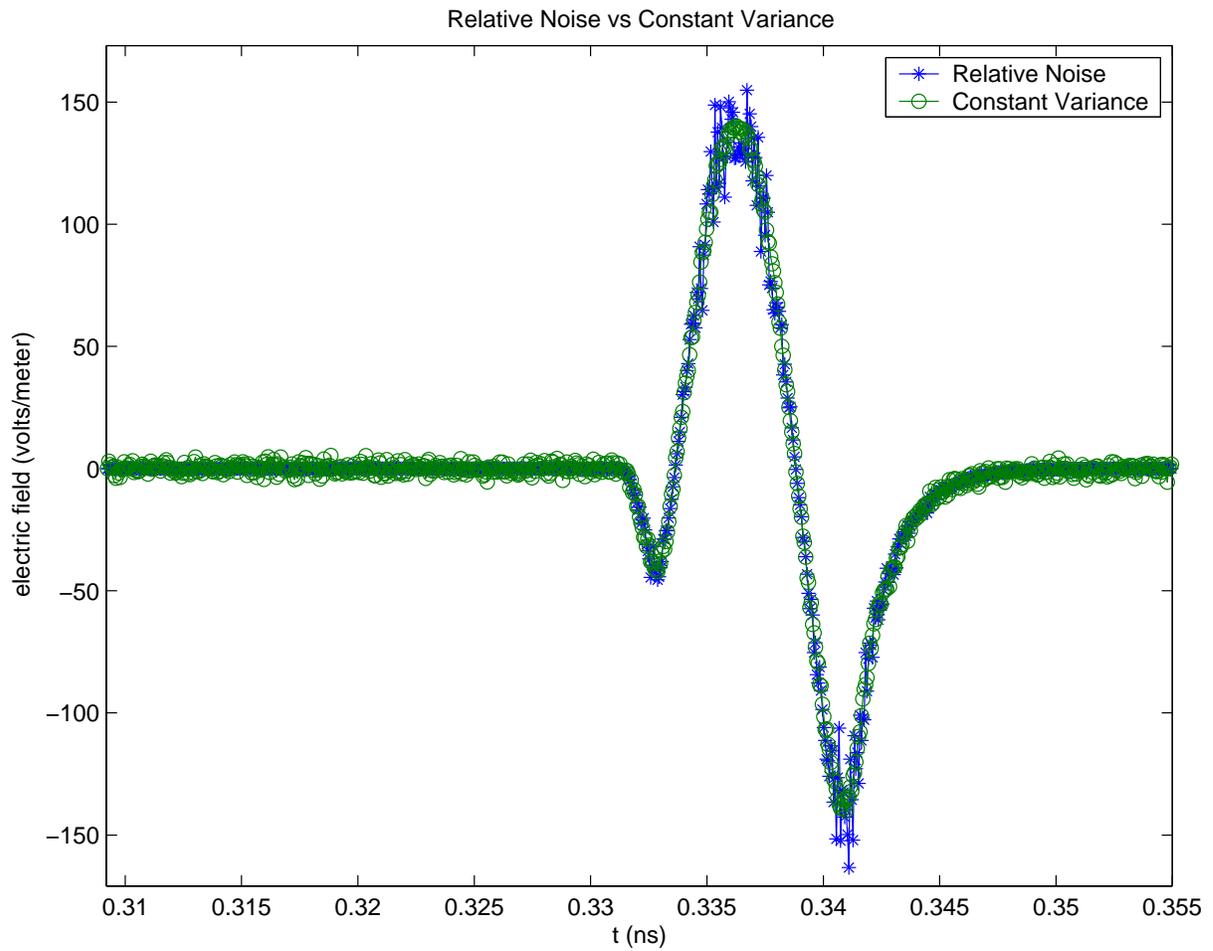


Figure 20: The difference between data with relative noise added and data with constant variance noise added is clearly evident when E is close to zero or very large.

δ	$d^* = .02 (N = 2048)$	$d^* = .04 (N = 4096)$
.0002	$(2.00005 \pm 9.30284 \times 10^{-7}) \times 10^{-2}$	$(4.00013 \pm 1.62162 \times 10^{-6}) \times 10^{-2}$
.0004	$(2.00001 \pm 6.50411 \times 10^{-7}) \times 10^{-2}$	$(4.00001 \pm 1.19064 \times 10^{-6}) \times 10^{-2}$
.0008	$(2.00001 \pm 4.91232 \times 10^{-7}) \times 10^{-2}$	$(4.00002 \pm 9.05240 \times 10^{-7}) \times 10^{-2}$

Table 5: Confidence intervals for the OLS estimate of d when the data is generated with no noise (i.e., $\nu_r = 0.0$).

δ	$d^* = .02 (N = 2048)$	$d^* = .04 (N = 4096)$
.0002	$(2.00004 \pm 4.86952 \times 10^{-6}) \times 10^{-2}$	$(4.00013 \pm 5.69385 \times 10^{-6}) \times 10^{-2}$
.0004	$(2.00001 \pm 3.50259 \times 10^{-6}) \times 10^{-2}$	$(4.00001 \pm 4.02428 \times 10^{-6}) \times 10^{-2}$
.0008	$(2.00001 \pm 2.87772 \times 10^{-6}) \times 10^{-2}$	$(4.00001 \pm 3.32933 \times 10^{-6}) \times 10^{-2}$

Table 6: Confidence intervals for the OLS estimate of d when the data is generated with noise level $\nu_r = .01$.

δ	$d^* = .02 (N = 2048)$	$d^* = .04 (N = 4096)$
.0002	$(2.00004 \pm 2.41541 \times 10^{-5}) \times 10^{-2}$	$(4.00014 \pm 2.76640 \times 10^{-5}) \times 10^{-2}$
.0004	$(2.00000 \pm 1.68896 \times 10^{-5}) \times 10^{-2}$	$(4.00001 \pm 1.90853 \times 10^{-5}) \times 10^{-2}$
.0008	$(2.00003 \pm 1.40398 \times 10^{-5}) \times 10^{-2}$	$(4.00000 \pm 1.60390 \times 10^{-5}) \times 10^{-2}$

Table 7: Confidence intervals for the OLS estimate of d when the data is generated with noise level $\nu_r = .05$.

δ	$d^* = .02 (N = 2048)$	$d^* = .04 (N = 4096)$
.0002	$(2.00000 \pm 4.72903 \times 10^{-5}) \times 10^{-2}$	$(4.00014 \pm 5.48283 \times 10^{-5}) \times 10^{-2}$
.0004	$(2.00003 \pm 3.39327 \times 10^{-5}) \times 10^{-2}$	$(4.00002 \pm 3.87474 \times 10^{-5}) \times 10^{-2}$
.0008	$(2.00003 \pm 2.79911 \times 10^{-5}) \times 10^{-2}$	$(4.00003 \pm 3.19526 \times 10^{-5}) \times 10^{-2}$

Table 8: Confidence intervals for the OLS estimate of d when the data is generated with noise level $\nu_r = .1$.

δ	$d^* = .02 (N = 2048)$	$d^* = .04 (N = 4096)$
.0002	$(1.99272 \pm 0.000182978) \times 10^{-4}$	$(1.98142 \pm 0.000317616) \times 10^{-4}$
.0004	$(4.00035 \pm 0.000201885) \times 10^{-4}$	$(4.00737 \pm 0.000369841) \times 10^{-4}$
.0008	$(7.99833 \pm 0.000136586) \times 10^{-4}$	$(8.00332 \pm 0.000251291) \times 10^{-4}$

Table 9: Confidence intervals for the OLS estimate of δ when the data is generated with no noise (i.e., $\nu_r = 0.0$).

δ	$d^* = .02 (N = 2048)$	$d^* = .04 (N = 4096)$
.0002	$(1.99410 \pm 0.000958274) \times 10^{-4}$	$(1.98029 \pm 0.00111475) \times 10^{-4}$
.0004	$(4.00170 \pm 0.00108740) \times 10^{-4}$	$(4.00667 \pm 0.0012499) \times 10^{-4}$
.0008	$(7.99882 \pm 0.000800042) \times 10^{-4}$	$(8.00486 \pm 0.000923838) \times 10^{-4}$

Table 10: Confidence intervals for the OLS estimate of δ when the data is generated with noise level $\nu_r = .01$.

δ	$d^* = .02 (N = 2048)$	$d^* = .04 (N = 4096)$
.0002	$(1.99606 \pm 0.00475672) \times 10^{-4}$	$(1.98106 \pm 0.00541764) \times 10^{-4}$
.0004	$(4.00190 \pm 0.00524360) \times 10^{-4}$	$(4.01214 \pm 0.00593246) \times 10^{-4}$
.0008	$(7.99045 \pm 0.00391181) \times 10^{-4}$	$(8.00947 \pm 0.00444525) \times 10^{-4}$

Table 11: Confidence intervals for the OLS estimate of δ when the data is generated with noise level $\nu_r = .05$.

δ	$d^* = .02 (N = 2048)$	$d^* = .04 (N = 4096)$
.0002	$(2.00017 \pm 0.00932701) \times 10^{-4}$	$(1.97674 \pm 0.0107203) \times 10^{-4}$
.0004	$(4.00070 \pm 0.0105331) \times 10^{-4}$	$(4.01229 \pm 0.0120445) \times 10^{-4}$
.0008	$(7.99698 \pm 0.00778563) \times 10^{-4}$	$(8.00361 \pm 0.00886925) \times 10^{-4}$

Table 12: Confidence intervals for the OLS estimate of δ when the data is generated with noise level $\nu_r = .1$.

6 Conclusion

In this presentation, we have explored a “proof of concept” formulation of an inverse problem to detect and characterize voids or gaps inside of, or behind, a dielectric medium. We have simplified the problem to one dimension and used Maxwell’s equations to model a pulsed, normally incident electromagnetic interrogating signal. We use Finite Element discretization in space, and Finite Differences in time, to simulate the electric field in the time domain. This is coupled with a Levenberg-Marquardt scheme in an optimization step with an innovative cost functional appropriate for reflected waves where phase differences can produce ill-posedness in the inverse problem when one uses the usual ordinary least squares criterion. We have successfully demonstrated that it is possible to resolve gap widths on the order of $.2mm$ between a dielectric slab of $20cm$ and a metal (perfectly conducting) surface using an interrogating signal with a $3mm$ wavelength.

Future work on this problem will likely involve more efficient computational methods since currently the inverse problem involving a $20cm$ slab takes 10 hours. Further, more sophisticated models for describing the polarization mechanisms in non-homogeneous materials must be developed. Finally, in order to take scattering and non-normally incident electromagnetic signals into account, multi-dimensional models will be necessary.

7 Acknowledgments

This research was supported (NLG) by the NASA Langley Graduate Student Researcher’s Program under grant NGT-1-02004 and in part (HTB) by the Air Force Office of Scientific Research under grant AFOSR F49620-01-1-0026. The research was also facilitated by stimulating discussions at the Statistical and Applied Mathematical Sciences Institute (SAMSI) which is funded by the National Science Foundation under grant DMS-0112069.

References

- [1] B. M. Adams, H. T. Banks, J. E. Banks and J. D. Stark *Population Dynamics Models in Plant-Insect Herbivore-Pesticide Interactions*, CRSC Technical Report CRSC-TR03-12, North Carolina State University, August, 2003.
- [2] H. T. Banks, M. W. Buksas, and T. Lin *Electromagnetic Material Interrogation Using Conductive Interfaces and Acoustic Wavefronts*, SIAM: Frontiers in Applied Mathematics, Philadelphia, 2000.
- [3] H. T. Banks, N. L. Gibson and W. P. Winfree *Electromagnetic Crack Detection Inverse Problems using Terahertz Interrogating Signals*, CRSC Technical Report CRSC-TR03-40, North Carolina State University, October, 2003.
- [4] M. Davidian and D. M. Giltinan *Nonlinear Models for Repeated Measurement Data*, Chapman and Hall: Monographs on Statistics and Applied Probability 62, New York, 1995.
- [5] R. S. Elliott *Electromagnetics: History, Theory, and Applications*, IEEE Press, New York, 1993.
- [6] C. Johnson *Numerical Solution of Partial Differential Equations by the Finite Element Method*, Cambridge University Press, New York, 1995.
- [7] C. T. Kelley *Iterative Methods for Optimization*, SIAM: Frontiers in Applied Mathematics, Philadelphia, 1999.