Numerical Simulation of Linear and Nonlinear Mie Scattering Using the Finite Difference Time Domain Method

Sean Meenehan
Harvey Mudd College
May 2nd, 2008

1 Linear Mie Scattering

Mie scattering refers to the scattering of light from spheres with a radius comparable to the wavelength of the light. In the linear case, simple Mie scattering is well understood and widely utilized for many applications, including the sizing of small particles, studies of atmospheric optical effects, and describing the nature of interstellar dust by examining the propagation of starlight through dust clouds[2].

However, classical Mie theory, though it yields an exact series solution for the scattered field, makes some assumptions which are not representative of real scattering situations. In particular, the theory is derived under the assumption of a single wave sitting isolated in free space, illuminated by a single frequency plane wave. In a real laboratory situation, the particles in question are typically delivered in some sort of aerosol system (if they are liquid) or sit in an array on some dielectric substrate (if they are solid). The presence of additional dielectric materials in the vicinity of the sphere can potentially alter the scattering pattern in ways that ordinary Mie theory cannot account for.

Moreover, our incident field is typically a laser. Far from being infinite, continuous plane waves, lasers are usually Gaussian beams, meaning that the intensity of the field falls off like a Gaussian as one moves radially outward from the central axis of the beam,

and are often pulsed, meaning they have some temporal envelope and are composed of multiple frequencies.

In addition, the particles themselves may not be spherical, but rather may possess an ellipsoidal or stranger shape. Alternatively, the particles in question may have a much more complex structure than the homogenous isotropy normally assumed in Mie scattering.

If we are to utilize Mie scattering in the real world, we need a way of assessing what non-idealities in the observed fields may be caused by these different factors. Complex geometries and fields like those above resist analytical solutions, and must be solved for numerically.

2 Nonlinear Mie Scattering

Also of considerable interest is Mie scattering in the nonlinear optical regime. Nonlinear optics refers to optical and electromagnetic effects generated by source terms which are related to the electric field in a nonlinear manner (usually proprtional to some power of the electric field strength greater than 1). Contrast this with the usual, linear approximation, where the source term for electromagnetic effects, the polarization, is related to the electric field linearly throught the relation $\vec{P} = \vec{\chi}^{(1)}\vec{E}$, where $\vec{\chi}^{(1)}$ is the linear susceptibility of the medium in question.

One of the most common nonlinear optical effects is known as second harmonic generation, where light at frequency ω can, if of sufficient intensity, generate light at frequency 2ω when propagating through or scattering off of a dielectric. This harmonic light is generated by a nonlinear polarization source term of the form $P_i^{(2)} = \chi_{ijk}^{(2)} E_j E_k$, where $\chi_{ijk}^{(2)}$ is the nonlinear susceptibility tensor. This tensor typically has very small values, and so this nonlinear effect is only evident when the intensity of the light (which is proportional to E^2) is quite large, such as in a high-field, pulsed laser beam[6].

A common example of second harmonic generation is the everyday green laser pointer. It is more power efficient to lase at lower frequencies, and so if we take an infrared laser and send it through a highly nonlinear material, we can produce substantial amounts of coherent green light for less of a power expenditure than had we lased in the green part of the visible spectrum directly. Almost every green laser pointer you find is actually an infrared diode sent through a nonlinear crystal before being output (open one up if you don't believe me!).

Second harmonic generation is particularly interesting in the Mie regime due to what are known as whispering gallery modes. Named after the whispering gallery at St. Paul's Cathedral in London, where a circular chamber can amplify even faint whispers to an audible level, a whispering gallery mode is a highly resonant optical wave created inside spheres of a precise size. Basically, an electric field incident on a dielectric sphere can totally internally reflect inside the sphere and travel around the inner surface. If the circumference of the sphere is equal to an integer number of wavelengths, the field will remain in phase with itself as it returns to its starting point, and will have a long lifetime (well, nanoseconds, which is relatively long on the scale of optical processes) inside the sphere due to a lack of destructive interference. Essentially, the sphere acts as an extremely effective spherical resonating cavity.

This can lead to high field strengths, and the effect is particularly evident in the Mie regime. In the case of second harmonic, since the intensity of second harmonic is quadratically related to the linear intensity, we expect to see very strong contributions to second harmonic fields from Mie size particles.

However, no exact solution exists for a complete treatment of second harmonic Mie scattering. Moreover, second harmonic generation generally involves high intensity, ultrafast lasers, which are difficult, if not impossible, to model without numerical methods. For the case of nonlinear light scattering, a numerical solution for the fields is almost certainly necessary.

3 Finite Difference Time Domain Method

The aim of my project, then, was to model complex linear Mie phenomena and hopefully to develop a numerical method for handling nonlinear scattering as well. The method I chose to attempt to model complex linear and nonlinear Mie phenomena was the finite difference time domain method. The theory behind this is somewhat involved, so we will first treat only the linear case.

3.1 Basic Concept

One of the most widely used methods for numerically determining electromagnetic fields is the finite difference time domain method. The basic idea starts with two of Maxwell's equations

$$\nabla \times \vec{E} = -\frac{d\vec{B}}{dt}, \tag{1}$$

$$\nabla \times \vec{H} = \frac{d\vec{D}}{dt} + \vec{J}, \tag{2}$$

where $\vec{E}$ and $\vec{B}$ are the electric and magnetic fields, $\vec{H}$ is the "auxiliary field" (no true consensus exist on a name for this field; some just call it $\vec{H}$), $\vec{D}$ is the electric displacement vector, and $\vec{J}$ the current. In a linear dielectric, these quantities are connected by the constituitive relations

$$\vec{D} = \vec{\epsilon}\,\vec{E}\,, \tag{3}$$
$$\vec{J} = \vec{\sigma}\,\vec{E}\,, \tag{4}$$
$$\vec{B} = \vec{\mu}\vec{H}\,, \tag{5}$$

where $\vec{\epsilon}$, $\vec{\sigma}$, and $\vec{\mu}$ are the electric permittivity, conductivity, and magnetic permeability tensors, respectively. These describe the response of charges in a dielectric material to applied electromagnetic fields.

Armed with equations (1) and (2), we proceed to discretize the space in and around our scatterer and express the spatial and temporal derivatives in Maxwell's equations using finite differences. We then split the total fields into the incident and scattered com-

ponents. The reason for this is simple. It is assumed that we can specify the incident field and its time derivatives exactly, usually through a closed form expression. Thus, we can separate it from the fields that we want, the scattered fields, and solve for these directly.

Once this is done, we perform some algebra to obtain expressions for the individual components of the fields at time step $n+1$ in terms of the incident field and the scattered field values at earlier time steps. This allows us to simply step forward in time, solving for the fields at each time step explicitly. In practice, we evaluate the electric and magnetic fields alternately at each time step in a method called "leapfrogging". That is, we solve for the electric fields everywhere at time $t = n\triangle t$, then proceed to use these values to solve for the magnetic fields at time $t = (n + \frac{1}{2})\ \triangle t$, which are in turn used to solve for the elctric fields at the next time step, etc.[1] The purpose of this is to allow us to use central difference approximations to the temporal derivatives, rather than resorting to less accurate backwards difference approximations. A similar spatial staggering of the fields is used to allow central differencing of the curl equations and to ensure that important boundary conditions on the fields are satisfied. To make this more concrete, we will show the derivation of one of the six field equations below.

3.2 Assumptions

Before we do that, however, we make a few simplifying assumptions. The first is the the material is isotropic. This means that the response of the material to one spatial component of the electric field is exactly the same as to the others. Mathematically, this means that the permittivity tensor in (3) becomes a constant. The reason for this is that off-diagonal elements in the permittivity tensor result in the mixing of field components. That is, all of a sudden our expression for, say, $E_x(t = (n + 1)\triangle t)$ would contain terms involving not only E_x at time $n\triangle t$, but also E_y and E_z. This coupling of the different field components makes simple iteration over the solution space for each field component at each time step impossible, since changes in one field component will affect the other field

components at the same location. Thus, off-diagonal permittivity elements require that we solve for all the field components at a certain time step simultaneously. While this could in principle be done with a matrix equation, the fact that we must solve for three different field components at each point in a 3D space makes this method cumbersome and tricky. Note that we can actually simulate anisotropic media, as long as the permittivity tensor remains diagonal (i.e. the dielectric responses in the x, y, and z directions may be different, but in each case only one component of the electric field stimulates a response in a given direction).

Another assumption is that $\mu \approx 1$, so that $\vec{B} = \vec{H}$, which simplifies (1) by reducing the number of constants we must keep track of. Most materials have a low magnetic permeability with a few rare exceptions (gadolinium and bismuth, for example).

The final assumption is that the medium is dispersionless. That is, we assume that the reponse of the material to an applied field is inpedendent of the frequency, and thus all frequencies propagate throught the medium at the same speed. This is competely untrue in general, but as long as we restrict our examinations to small bands of frequencies it is a reasonable approximation. Methods exist to handle frequency dependent material parameters, but for the time being we will restrict ourselves to this simpler approximation.

3.3 Finite Differencing: an example

Let us begin with (2), and immediately take note of the fact that, since the incident fields must satisfy the free space equations (where $\epsilon \to \epsilon_0$, the vacuum permittivity, and $\vec{J}$ disappears altogether), we have,

$$\nabla \times \vec{H}_{\text{inc}} = \epsilon_0 \frac{d\,\vec{E}_{\text{inc}}}{dt}. \tag{6}$$

Combining (2) and (6), and moving the temporal derivative of the incident field on to the right side of equation (2). Then, considering only the x component, (2) becomes

$$\frac{dH_{y,\text{scatt}}}{dz} - \frac{dH_{z,\text{scatt}}}{dy} = \epsilon_{xx}\frac{dE_{x,\text{scatt}}}{dt} + (\epsilon_{xx} - \epsilon_0)\frac{dE_{x,\text{inc}}}{dt} + \sigma_{xx}(E_{x,\text{inc}} + E_{x,\text{scatt}}). \tag{7}$$

After approximation of the scattered field derivatives the equation beomes

$$\frac{H^{n+\frac{1}{2}}_{y,\text{scatt}}(i,j,k) - H^{n+\frac{1}{2}}_{y,\text{scatt}}(i,j,k-1)}{\triangle z} - \frac{H^{n+\frac{1}{2}}_{z,\text{scatt}}(i,j,k) - H^{n+\frac{1}{2}}_{z,\text{scatt}}(i,j-1,k)}{\triangle y} =$$

$$\epsilon_{x\,x}\frac{E^{n}_{x,\text{scatt}}(i,j,k) - E^{n-1}_{x,\text{scatt}}(i,j,k)}{\triangle t} + (\epsilon_{x\,x} - \epsilon_0)\frac{d\,E^{n}_{x,\text{inc}}(i,j,k)}{dt} + \sigma_{x\,x}(E^{n}_{x,\text{inc}} + E^{n}_{x,\text{scatt}}), \tag{8}$$

where the superscripts indicate the time step at which the fields are evaluated, and the (i, j, k) additions indicate the spatial grid indices at which the fields are evaluated. Equation (8) can be solved for $E^{n}_{x,\text{scatt}}$ to get the formula for updating the x-component of the scattered field at each point in space using only the most recent magnetic fields and the previous electric field values.

Earlier the claim was made that the finite differences used were central differences, but here they appear at first glance to be backwards differences. However, with regards to the time derivative note that we are in effect taking a central difference about time $t = (n + \frac{1}{2})\triangle t$, with the effective time step $\triangle t' = \frac{1}{2}\triangle t$. Thus, we are taking a centered difference in time of the *electric* fields about the time step of the previously calculated *magnetic* fields. This is valid because $\vec{E}$ and $\vec{H}$ are not just any old vector fields, but are rather intimately connected by Maxwell's equations. In a very real sense they are both part of the same physical electromagnetic field. This also sheds light on why we alternate between the evaluation of the two fields with each time step. Were we to evaluate both $\vec{E}$ and $\vec{H}$ at each point in time, we would be forced to find $\vec{E}^{\,n}$ using backwards differences, or to keep both fields out to two time steps back, doubling our memory requirement.

The claim that the spatial differences of $\vec{H}$ present in (8) are central differences is justified by considering the geometry that we use to evaluate our fields. Though we index the fields by discretized points, we are not actually evaluating the fields at each point (i, j, k).

Rather, in each grid cell within our space the components of $\vec{E}$ and $\vec{H}$ are staggered in an arrangement known as Yee cells[3]. In this arrangement, the electric field components are located on the edges of each grid cell, while the magnetic field components are located in the middle of each face. That is, for a grid point $(i,\ j,\ k)$, the physical locations of the field components are as follows

$$E_x(i,j,k) \to (i+\tfrac{1}{2},j,k) \qquad E_y(i,j,k) \to (i,j+\tfrac{1}{2},k) \qquad E_z(i,j,k) \to (i,j,k+\tfrac{1}{2})$$
$$H_x(i,j,k) \to (i,j+\tfrac{1}{2},k) \quad H_y(i,j,k) \to (i+\tfrac{1}{2},j,k+\tfrac{1}{2}) \quad H_z(i,j,k) \to (i+\tfrac{1}{2},j+\tfrac{1}{2},k).$$

Notice that, using the following field locations that the finite differences of H_y and H_z in (8) can be considered central differences about the point $(i+\frac{1}{2},j,k)$, which is the point at which $E_x(i,\ j,\ k)$, the field component we are evaluating, resides. This spatial interleaving of the electric and magnetic fields allows us to make our spatial derivatives central differences as well.

The other advantage of the Yee cell arrangment is that, since the components of the magnetic field are located on face of the grid cells to which they are normal (e.g H_y is located in the middle of a face aligned in the $x\,z$ plane), and the components of the electric field are located on edges to which they are tangential (e.g. E_x is located on the edge parallel to the x-axis), these field locations allow us to ensure the well known boundary conditions requiring continuity of the normal components of $\vec{H}$ and the tangential components of $\vec{E}$ at the boundary surface of the dielectric (at least to the extent that it is accurately represented by cubic cells)[3].

3.4 Stability and Accuracy Issues

The primary factor in determining how accurate our solution will be is our grid size. Namely, we make some errors in approximating non-cubic shapes, such as a sphere, using cubic cells (FDTD can be formulated in terms of spherical coordinates, but its much more

cumbersome). The location of the boundaries necessarily get smeared out, as we assign dielectric properties on a cell-by-cell basis, thus overlooking any fine structure to the boundary on the scale of less than the size of a grid cell.

In addition, choosing the grid cell size determines how many samples per wavelength we are getting of the electric and magnetic fields. A general rule of thumb is that we want the grid spacing less than or eual to $\frac{\lambda}{10}$, where λ is the smallest wavelength of interest in our problem[1]. This ensures adequate sampling of the field to prevent aliasing effects, while not requiring too much computational time to iterate over the solution space.

The stability condition for our time step is simply the three-dimensional formulation of the well-known Courant stability condition, requiring

$$\triangle t \leq \frac{1}{c\sqrt{\left(\frac{1}{\triangle x}\right)^2 + \left(\frac{1}{\triangle y}\right)^2 + \left(\frac{1}{\triangle z}\right)^2}}, \tag{9}$$

where c is the speed of light in a vacuum.

3.5 Absorbing boundaries

The boundary of our solution space presents a problem. In reality, the scattered fields would propagate out to infinity. This is the definition of radiation, in fact: light which carries energy out to infinity. However, because we cannot solve in an infinitely large solution space, we necessarily impose an artificial boundary around our scatterer. If we are not careful this will result in reflections of the scattered fields back into the space, which will severely impact the accuracy of our solution.

One method to circumvent this is what is known as an "absorbing boundary condition", first developed by Mur[4]. Without going into too much detail, the general idea is to find the value of a field component at some point x on the boundary by interpolating between the previous time step's field value at x and the current and previous field values just inside the solution space near x, say at some point y.

Essentially, by looking at the field values just inside the boundary at y, both now and previously in time, and by using the known time and distance intervals between the points, we can approximate how the fields that *were* at y have now propagated out towards x. By using this in conjunction with an approximation for how the fields that were at x have propagated out towards infinity, we can find a new value for the fields at x which properly simulates the outgoing fields. This can be done to higher order (i.e. looking back multiple steps in time and looking multiple steps inward into the solution space), but in practice first or second order Mur boundary conditions are usually sufficient for most applications. One of the first-order Mur conditions, that for computing E_z on the $x=0$ face of the boundary, is presented here for concreteness

$$E_z^n\left(0, j, k+\frac{1}{2}\right) = E_z^{n-1}\left(1, j, k+\frac{1}{2}\right) + \frac{c\triangle t - \triangle x}{c\triangle t + \triangle x}\left(E_z^n\left(1, j, k+\frac{1}{2}\right) - E_z^{n-1}\left(0, j, k+\frac{1}{2}\right)\right). \quad (10)$$

The outward travelling fields will become more and more like pure transverse waves (i.e the electric and magnetic field vectors are perpendicular to the direction of propagation), as they move away from the scatterer, and their progress towards the boundary will be more accurately modeled. Thus, we ideally want many grid cells between the boundary of the scatterer and the solution space boundary. In practice 10 cells on all sides is generally sufficient for simple scattering from spheres[1].

3.6 Far Zone Tranformation

In practice the solution space only extends a small physical distance beyond the scatterer (e.g. 10 grid cells for optical frequencies is on the order of μm), and so to translate the fields into what we might actually see requires tranforming the near fields into the corresponding fields in the far zone. We accomplish this by approximating a surface integral

of the tangential fields over the boundary of our solution space at each time step[5]. This can be transformed into the far zone fields by the formula

$$E_\theta^{\text{far zone}} = -\eta W_\theta - U_\phi, \tag{11}$$

$$E_\phi^{\text{far zone}} = -\eta W_\phi + U_\theta, \tag{12}$$

where $\eta = \frac{1}{\epsilon_0 c}$ is a constant known as the impedance of free space, and the vector fields $\vec{W}$ and $\vec{U}$ are defined as

$$\vec{W}(t) = \frac{1}{4\pi c}\frac{d}{dt}\left(\int_S dA\ \vec{J}(t + (\vec{r}\,' \cdot \hat{r})/c) \right), \tag{13}$$

$$\vec{U}(t) = \frac{1}{4\pi c}\frac{d}{dt}\left(\int_S dA\ \vec{M}(t + (\vec{r}\,' \cdot \hat{r})/c) \right). \tag{14}$$

The integration here is over a suitable surface which completely encloses the scattering objects of our solution space, and the fields inside the integrals are just the tangential components of the electromagnetic fields, $\vec{J} = \hat{n} \times \vec{H}$, $\vec{M} = -\hat{n} \times \vec{E}$, where $\hat{n}$ is the surface normal at the vector. The vector $\vec{r}\,'$ is the vector from a chosen reference point (usually the center of the space) to the point of integration in the surface, while $\hat{r}$ is the unit vector from this reference point to the far zone observation point for whihc we are calculating the fields.

To get the surface integral, we approximate the appropriate field as constant over each cell face, with a value equal to the average of the surrounding relevant fields, then multiply this spatial average by the area of the cell face. Note that since the magnetic fields tangential to a surface do not actually lie on the surface in a Yee cell geometry, we must get the average at the center of the face by using field values outside the integration surface. For this reason, we set the faces of our integration surface one cell inside the boundary of our solution space on all sides.

The time derivative is approximated by a finite difference method, and from these we can figure out at which time steps will a particular field component (evaluated at a partic-

ular time) will contribute to the surface integral. Any particular field component at any particular time step will contribute mutiple times to the surface integral, over a time interval of $2\triangle t$, due to its appearance in two finite difference equations.

For example, evaluating the contribution of E_x to U_z at time $(n+\frac{1}{2})\triangle t$ involves a finite difference proportional to $E_x^{n+1} - E_x^n$, while the contribution at time $(n-\frac{1}{2})\triangle t$ will be proprtional to $E_x^n - E_x^{n-1}$. Thus, E_x^n has multiple contributions to U_z at different times. In practice we handle this by keeping an array of all possible time steps of interest for $\vec{W}$ and $\vec{U}$. At each time step, after calculating the fields within the boundary, we add the contribution for each field element to whatever time steps of $\vec{W}$ and $\vec{U}$ it affects, keeping a running total for each time step of all the contributions.

3.7 Advantages of the FDTD Method

A primary advantage of this method is that it solves for the time dependent electromagnetic fields directly, rather than decomposing the incident field into different frequency components, solving for the radiated fields at each frequency, and then reconstructing the scattered fields from the individual scattered frequencies, as is frequently done in electrodynamical problems. Though the this technique can be powerful, there are many situations where the incident fields cannot be easily specified in terms of its constituent frequency components, such as a laser with a Gaussian intensity profile as mentioned above, or where this method of spectral decomposition can be cumbersome due to a large number of frequency components, such as in an ultrafast laser pulse.

Another advantage to this method is that it allows for the modeling of arbitrary geometries, the accuracy being limited only by one's resolution in terms of grid spacing. In the case of Mie scattering, this allows realtively straightforward modeling of multiple spheres in close proximity, spheres on a dielectric substrate, inhomogeneous spheres (spheres which have multiple interior regions with different dielectric properties), etc.

4 Results

The FDTD scheme presented above was coded up in C++. To evaluate the validity of the method and the correctness of the code, as simple test case was chosen: ordinary Mie scattering of a single frequency plane wave by a single dielectric sphere in vacuum. This provides an excellent benchmark, as we have a readily available exact theory against which to compare. The Mie theory calculations used to check the results were provided by another piece of code to implement the exact solution (the details of which are not relevant here; the interested reader is directed to Ref. 2), which has been thoroughly checked against other implementations of Mie theory and data from Prof. Tom Donnelly's laboratory.

The test case was an 800 nm plane wave incident upon a 1 μm sphere of water (dielectric constant of $\epsilon = 1.755\,\epsilon_0$ (the conductivity σ is assumed negligible). Observation points in the far field were chosen to be evenly spaced between $\theta = 0$ and $\theta = \pi$ at intervals of $\frac{\pi}{12}$. The exact solution using Mie theory is displayed in Fig. 1.

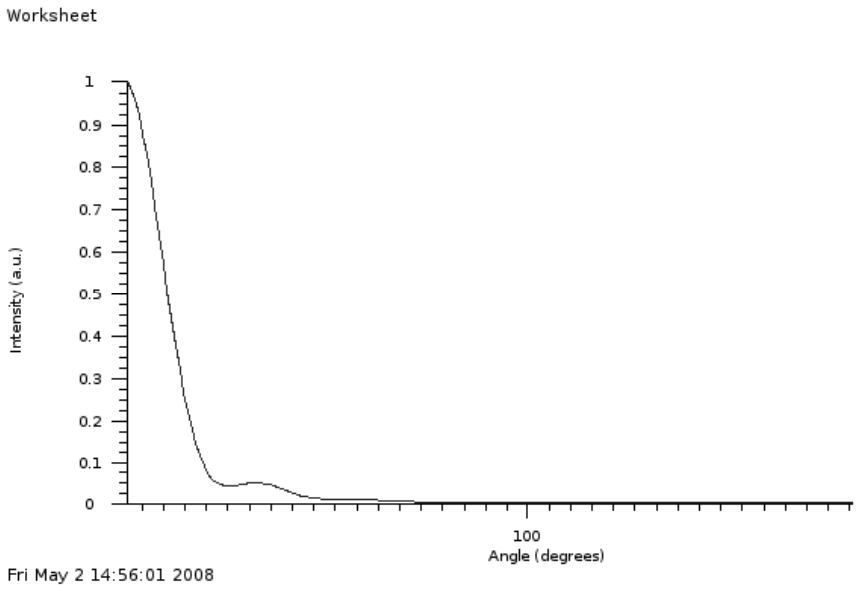

Figure 1. Plot of intensity vs. angle (normalized to 1 at peak intensity) for 800 nm incident plane wave, 1 μm radius spheres of water. Note the dominant forward scattering, the hallmark of Mie scattering.

Using the FDTD model proposed above, we have a grid spacing of 0.08 μm and a time step of about 0.15 fs (fs = 10^{-15} s) .

To check our model, we take the square of the numerically determined far zone fields (which is proportional to the intensity) after they have reached a steady state value (as they should when illuminated by a constant plane wave). Ideally, the values should then correspond exactly to the corresponding points on the plot in Fig. 1. The results are shown in Fig. 2

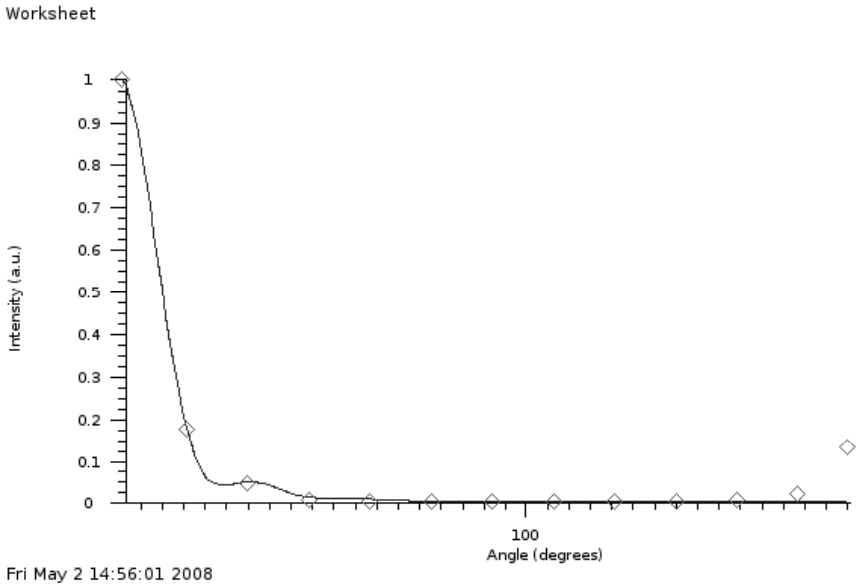

Figure 2. The diamond points represent the fields at the observation points of the model. Notice that we have excellent numerical agreement until we get to the far backwards angles, at which point the numerical solution deviates substantially from Mie theory.

Generally the numerical model perfectly replicates the predictions of Mie theory. However, there is substantial numerical disagreement in the backscatter. To try and discern whether this was an issue with the actual FDTD implementation in the sphere, or an error in the implementation of the far zone transform in the rear direction, I plotted the fields inside the sphere for the first 300 time steps (out of a total of 1000), shown in Fig. 3

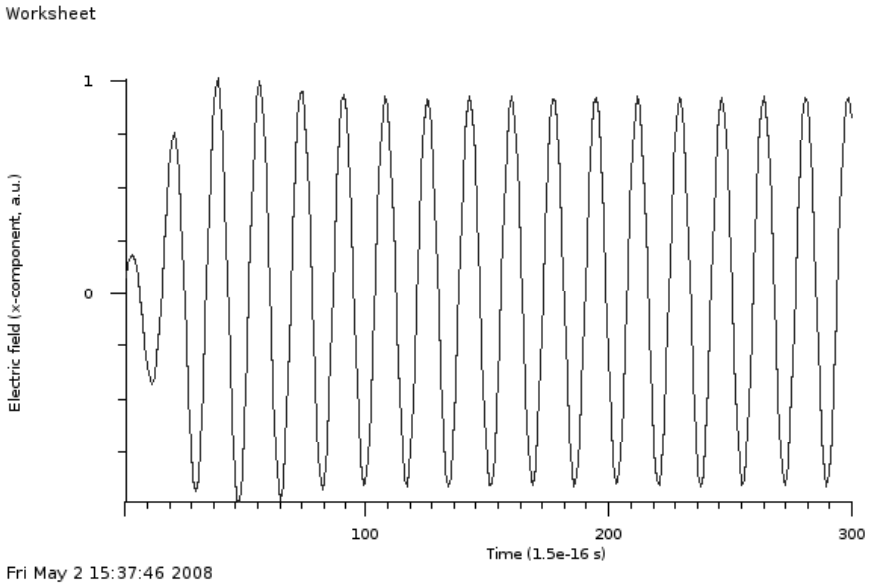

Figure 3. Electric field inside the sphere (x-component)

Looking at the above plot we see that over the course of roughly 80 time steps (8 minor ticks on the x-axis, or $\sim 12\,\text{fs}$) we have about 4 oscillations. 4 periods in about 12 fs equals a frequency of 3.33×10^{14} Hz, corresponding to about 900 nm. This is a rough estimate, of course, but it seems like the wavelength is mostly on target inside the sphere. Even if it weren't 900 nm is still well within the Mie regime for this size of sphere, so we shouldn't see backscatter. Thus, my guess is that something in the far-zone transformation code is off. Unfortunately, it took forever to get all the kinks worked out of the code (well, most of them, apparently), so I have no time to investigate this further. Sorry, Prof. Yong. I suck, big time...

References

1. Kunz, K., and Luebbers, R., *The Finite Difference Time Domain Method for Electromagnetics* (CRC Press, Boca Raton, FL, 1993)
2. Craig F Bohren, Donald R. Huffman, *Absorption and Scattering of Light by Small Particles* (John Wiley & Sons, New York, NY, 1983)
3. Yee, K.S., *IEEE Trans. Ant. Prop.*, **14**, 302 (1966)
4. Mur, G., *IEEE Trans. Electromag. Compat.*, **23**, 1073 (1981)
5. Luebbers, R., Kunz, K., Scneider, M., and Hunsberger, F., *IEEE Trans. Ant. Prop.*, **39**, 429 (1991)
6. Boyd, R.W., *Nonlinear Optics, 2nd. Ed.* (Academic Press, San Diego, 2003)