Solving the Generalized Poisson Equation Using the Finite-Difference Method (FDM)

James R. Nagel

September 30, 2009

1 Introduction

Numerical simulation is an extremely valuable tool for those who wish to understand complex electrostatic systems. Although there are several competing methods for achieving this goal, one of the simplest and more straightforward of these is called the *finite-difference method* (FDM). Basically, FDM works by sampling the voltage potential within some finite simulation domain and then approximating the derivative operation with a finite-difference. When applied to a time-independent partial differential equation, the net result is a linear system of equations that may be readily solved via matrix inversion.

Physically speaking, the ultimate governing equation for any electrostatic system is Gauss's law, which is expressed in point form as

$$\nabla \cdot \mathbf{D}(\mathbf{r}) = \rho(\mathbf{r}) \ . \tag{1}$$

In this context, $\mathbf{r} = x\hat{\mathbf{x}} + y\hat{\mathbf{y}} + z\hat{\mathbf{z}}$ is a position vector in space, ρ is the charge density function, and $\mathbf{D}$ is the electric flux density. Using the constitutive relation $\mathbf{D}(\mathbf{r}) = \epsilon(\mathbf{r})\mathbf{E}(\mathbf{r})$, Gauss's law may be rewritten in terms of the electric field intensity $\mathbf{E}$ as

$$\nabla \cdot \Big[\epsilon(\mathbf{r})\ \mathbf{E}(\mathbf{r})\Big] = \rho(\mathbf{r}) \ , \tag{2}$$

where $\epsilon(\mathbf{r})$ is the dielectric function. Gauss's law may be further rewritten in terms of voltage potential V by making the substitution $\mathbf{E}(\mathbf{r}) = -\nabla V(\mathbf{r})$:

$$\nabla \cdot \Big[\epsilon(\mathbf{r})\ \nabla V(\mathbf{r})\Big] = -\rho(\mathbf{r}) \ . \tag{3}$$

Although it is not commonly discussed in the literature, this is really nothing more than a generalized form of the *Poisson equation*, and is the expression we shall be most interested in throughout this paper. A far more familiar expression occurs if we next assume a uniform dielectric function with the form $\epsilon(\mathbf{r}) = \epsilon$. This gives us

$$\nabla^2 V(\mathbf{r}) = -\frac{\rho(\mathbf{r})}{\epsilon} \ , \tag{4}$$

which is the classical form for the Poisson equation as given in most textbooks.

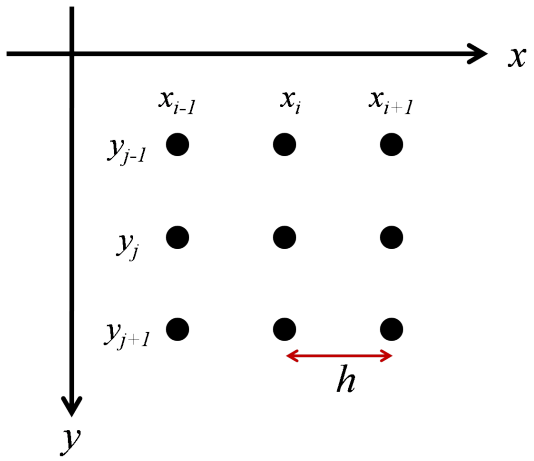

Figure 1: Mesh points for the FDM grid. For convenience with matrix indexing, the y-axis is inverted from the typical Euclidian convention.

Although the classical Poisson equation is much simpler to numerically solve, it also tends to be very limited in its practical utility. Realistically, the generalized Poisson equation is the true equation we will eventually need to solve if we ever expect to properly model complex physical systems. We shall therefore begin by using the classical Poisson equation as a demonstration case for how FDM works before expanding our algorithm to the generalized form. For brevity and simplicity, this paper will be strictly limited to two-dimensional systems, though a full three-dimensional solution follows a nearly identical derivation.

2 The Five-Point Star

The first step in applying FDM is to define a *mesh*, which is simply a uniform grid of spatial points at which the voltage function will be sampled. Letting h be the distance between each sample, the points that lie on the mesh may be defined by

$$x_i = ih \ , \quad \text{and} \tag{5}$$

$$y_j = jh \ , \tag{6}$$

where i and j are integers. In practice, i and j will eventually be used as indices for a matrix of voltage samples, so it helps to use the convention depicted in Figure 1 with the inverted y-axis. We may then define the voltage samples at each grid point using the short-hand notation

$$V(i,j) = V(x_i, y_j) \ . \tag{7}$$

Along with being more compact, this notation serves as an analogy to the matrix indexing that will be required when one actually codes the FDM algorithm into a programming language like Matlab. In a similar fashion, we may also define the charge density samples along the same mesh by using the $\rho(i,j)$ notation.

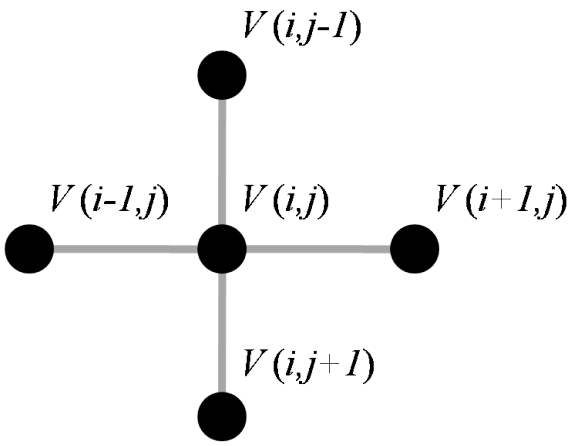

Figure 2: Basic stencil for the 5-point star.

The next step is to expand the Poisson equation by explicitly showing the partial derivatives in space:

$$\frac{\partial^2 V(i,j)}{\partial x^2} + \frac{\partial^2 V(i,j)}{\partial y^2} = -\frac{\rho(i,j)}{\epsilon} \ . \tag{8}$$

The reason for doing this is so that we may approximate the derivative operators through the use of finite-differences. The easiest way to do this is through the three-point approximation for the second-derivative, which is given as

$$\frac{\partial^2}{\partial x^2} V(i,j) \approx \frac{V(i-1,j) - 2V(i,j) + V(i+1,j)}{h^2} \ . \tag{9}$$

with a similar expression for the y-component. Plugging back into Equation (8) then gives us

$$V(i-1,j) + V(i+1,j) + V(i,j-1) + V(i,j+1) - 4V(i,j) = -\frac{h^2}{\epsilon}\rho(i,j) \ . \tag{10}$$

Finally, if we solve for $V(i,j)$, we find

$$V(i,j) = \frac{1}{4}\left[V(i-1,j) + V(i+1,j) + V(i,j-1) + V(i,j+1) + \frac{\rho(i,j)h^2}{\epsilon}\right] \ . \tag{11}$$

What this expression tells us is that every voltage sample $V(i,j)$ is dependent only on $\rho(i,j)$ and voltage at the four nearest neighbors. A graphical depiction of this is called a *computational molecule*, and is shown in Figure 2. Because of its unique geometry, this five-point stencil is often referred to as the *five point star*.

Because each voltage sample $V(i,j)$ is linearly dependent on its four nearest neighbors, the solution over all (i,j) may be represented as a simple matrix-vector equation. This is readily achieved if we first vectorize the voltage samples by some convention, such as

$$\mathbf{x} = \begin{bmatrix} V(1,1) & V(2,1) & V(3,1) & \cdots & V(2,1) & V(2,2) & V(2,3) & \cdots \end{bmatrix}^T \ . \tag{12}$$

The next step is to express the linear relationship between voltage samples into a matrix $\mathbf{A}$. This effectively converts the entire problem into a matrix-vector equation with the form

$$\mathbf{A}\mathbf{x} = \mathbf{b} \ , \tag{13}$$

where $\mathbf{b}$ simply contains all the information about charge density and Dirichlet boundary conditions. The numerical solution to the system is finally found by simply inverting the matrix $\mathbf{A}$ to find

$$\mathbf{x} = \mathbf{A}^{-1}\mathbf{b} \ . \tag{14}$$

3 Successive Over-Relaxation

For relatively small simulation domains, the direct matrix inversion of Equation (14) works perfectly well for obtaining a solution. However, it is important to realize that the size of $\mathbf{A}$ grows directly with the square of the simulation domain. For example, given a simulation domain of 100×100 voltage samples, the matrix $\mathbf{A}$ will need to be $10,000 \times 10,000$ elements. Because direct matrix inversion is such an intense operation, it is easy to see how even small simulations can quickly require excessive computational resources.

To reduce the computational cost required by direct matrix inversion, it helps to realize that $\mathbf{A}$ is a *sparse* matrix, meaning the vast majority of elements in $\mathbf{A}$ are all zeros. This is a direct consequence of Equation (11), which shows that each voltage element is only dependent on four other samples. As a result, each row in $\mathbf{A}$ has, at most, only five nonzero entries (or even one nonzero entry if the voltage sample is a fixed Dirichlet boundary). This allows us arrive at a solution through the use of iterative methods that take advantage of this property, and one of the more simple algorithms is called *successive over-relaxation* (SOR).

The first step when utilizing SOR is to define a *residual* $R(i,j)$ as the degree to which each voltage sample $V(i,j)$ does not satisfy Equation (11):

$$R(i,j) = V(i-1,j) + V(i+1,j) + V(i,j-1) + V(i,j+1) - 4V(i,j) + \frac{\rho(i,j)h^2}{\epsilon} \tag{15}$$

The next step is to loop over every sample in $V(i,j)$ and add a correction factor defined by the residual $R(i,j)$. This process is then repeated over many iterations until the residual falls below some acceptable error value. For the kth iteration in the loop, we therefore have

$$V^{k+1}(i,j) = V^k(i,j) + R^k(i,j) \ . \tag{16}$$

This method is referred to as *successive relaxation*, and is guaranteed to eventually converge on the correct solution. However, it is also possible to speed up the convergence process by multiplying R with a *relaxation factor* ω, such that

$$V^{k+1}(i,j) = V^k(i,j) + \omega R^k(i,j) \ . \tag{17}$$

This is the method of successive *over*-relaxation, and will also converge as long as we enforce the condition that $0 < \omega < 2$. The only difficult part is choosing an ideal value for ω which, unfortunately, can only be found through trial-and-error. However, rapid converge may be typically obtained with a value on the order of $\omega = 1.9$.

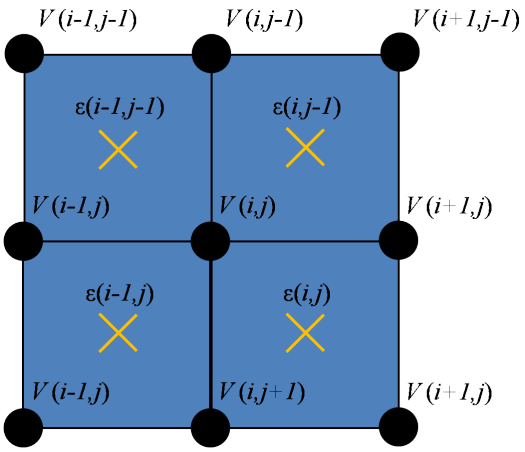

Figure 3: Finite-difference mesh for the generalized Poisson equation.

4 Generalized Poisson Equation

Let us now return to the generalized Poisson equation:

$$\nabla \cdot \Big[\epsilon(\mathbf{r})\ \nabla V(\mathbf{r})\Big] = -\rho(\mathbf{r})\ . \tag{18}$$

Although we could easily begin by directly applying numerical derivatives to this expression, a more accurate approximation may be reached by first defining the permittivities along the *staggered* grid shown in Figure 3. Mathematically, this may be written as

$$\epsilon(i,j) = \epsilon(x_i + h/2, y_j + h/2)\ , \tag{19}$$

with $V(i,j)$ and $\rho(i,j)$ defined along the same grid points as before. What this allows us to do is use the central-difference method on the first-order derivatives that will come out from this expression. It also places the voltage samples along the boundaries of the dielectric permittivities. Another change that turns out to be very convenient is to convert the Poisson equation into its integral form. If we define Ω_{ij} as the square region around a single voltage sample $V(i,j)$, we may take the volume integral to find

$$\int\limits_{\Omega_{ij}} \nabla \cdot \Big[\epsilon(\mathbf{r})\ \nabla V(\mathbf{r})\Big] d\Omega = -\int\limits_{\Omega_{ij}} \rho(\mathbf{r}) d\Omega\ . \tag{20}$$

An example of this volume element is depicted in Figure 4.

Looking first at the right-hand side of Equation (20), we note that the integral over the charge density is simply the total charge enclosed inside the volume. We may therefore make the replacement

$$-\int\limits_{\Omega_{ij}} \rho(\mathbf{r}) d\Omega = -Q(i,j)\ . \tag{21}$$

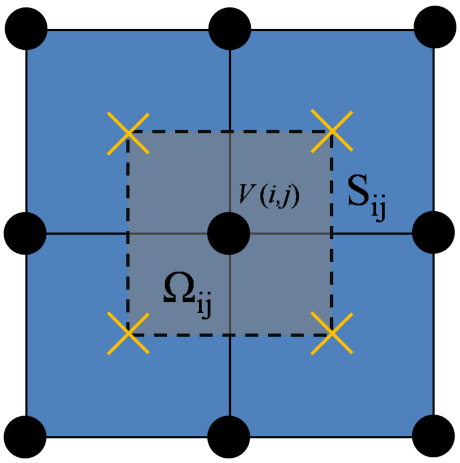

Figure 4: Volume element around voltage sample $V(i,j)$. The region Ω_{ij} represents the region enclosed by the outer surface S_{ij}.

The left-hand side of Equation (20) may also be simplified by applying the divergence theorem. This converts the volume integral over Ω_{ij} into a surface integral around its outer border, giving

$$\int\limits_{\Omega_{ij}} \nabla \cdot \Big[\epsilon(\mathbf{r})\ \nabla V(\mathbf{r})\Big] d\Omega = \oint\limits_{S_{ij}} \Big[\epsilon(\mathbf{r})\ \nabla V(\mathbf{r})\Big] \cdot d\mathbf{S} = Q(i,j)\ , \tag{22}$$

where S_{ij} is the enclosing surface and $d\mathbf{S}$ is the differential unit normal vector. It is this form of the Poisson equation that we will now apply our finite-difference approximations to.

With the desired expression in hand, the next step is to expand out the gradient operator to find

$$\oint\limits_{S_{ij}} \Big[\epsilon(\mathbf{r})\ \nabla V(\mathbf{r})\Big] \cdot d\mathbf{S} = \oint\limits_{S_{ij}} \left[\epsilon(\mathbf{r}) \left(\frac{\partial}{\partial x} V(\mathbf{r})\hat{\mathbf{x}} + \frac{\partial}{\partial y} V(\mathbf{r})\hat{\mathbf{y}}\right)\right] \cdot d\mathbf{S} \tag{23}$$

In three dimensions, the surface S_{ij} would normally be a cube, but in our two-dimensional example is simply a square. We may use this simplified geometry to rewrite the total surface integral as a series of sub-integrals around each side of the square. For brevity, we shall simply write these as integrals over the four surfaces $S_1 \cdots S_4$:

$$\oint\limits_{S_{ij}} \left[\epsilon(\mathbf{r}) \left(\frac{\partial}{\partial x} V(\mathbf{r})\hat{\mathbf{x}} + \frac{\partial}{\partial y} V(\mathbf{r})\hat{\mathbf{y}}\right)\right] \cdot d\mathbf{S} = \int\limits_{S_1} + \int\limits_{S_2} + \int\limits_{S_3} + \int\limits_{S_4} \ . \tag{24}$$

This geometry is depicted in Figure 5, which highlights the contour of integration over all four sides, taken in the counter-clockwise direction.

As an example case, let us evaluate the integral over S_1, which we shall define as the right side of the square where $d\mathbf{s} = \hat{\mathbf{x}} dy$. For simplicity, it also helps to assume that $V(i,j)$ lies at the origin, though the result is equivalent at any location. We may therefore express the integral

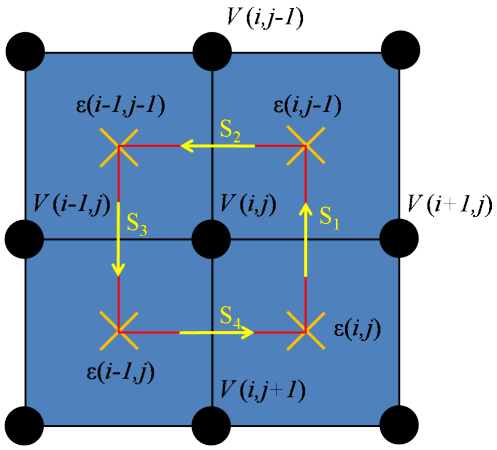

Figure 5: Schematic of the contour of integration along the staggered grid.

over S_1 as

$$\int_{S_1} = \int_{-h/2}^{h/2} \epsilon(x,y) \left(\frac{\partial}{\partial x} V(x,y)\hat{\mathbf{x}} + \frac{\partial}{\partial y} V(x,y)\hat{\mathbf{y}} \right) \cdot (\hat{\mathbf{x}}) dy = \int_{-h/2}^{h/2} \epsilon(x,y) \frac{\partial}{\partial x} V(x,y) dy \ . \tag{25}$$

Next, we note that the integral consists entirely along the border between the regions defined by $V(i,j)$ and $V(i+1,j)$. We may therefore approximate the partial derivative by using a central difference between the two samples, and assume that it is constant across the entire border. Calculating the integral across the two dielectric regions therefore gives

$$\begin{aligned} \int_{S_1} &\approx h \left[\frac{\epsilon(i,j) + \epsilon(i,j-1)}{2} \right] \left[\frac{V(i+1,j) - V(i,j)}{h} \right] \\ &= (1/2) \Big[\epsilon(i,j) + \epsilon(i,j-1) \Big] \Big[V(i+1,j) - V(i,j) \Big] \ . \end{aligned} \tag{26}$$

Carrying out this same operation over the other three sides thus gives

$$\int_{S_2} \approx (1/2) \Big[\epsilon(i,j-1) + \epsilon(i-1,j-1) \Big] \Big[V(i,j-1) - V(i,j) \Big] \tag{27}$$

$$\int_{S_3} \approx (1/2) \Big[\epsilon(i-1,j-1) + \epsilon(i-1,j) \Big] \Big[V(i-1,j) - V(i,j) \Big] \tag{28}$$

$$\int_{S_4} \approx (1/2) \Big[\epsilon(i-1,j) + \epsilon(i,j) \Big] \Big[V(i,j+1) - V(i,j) \Big] \ . \tag{29}$$

For notational compactness, we now define the following constants:

$$\begin{aligned}
a_0 &= \epsilon(i,j) + \epsilon(i-1,j) + \epsilon(i,j-1) + \epsilon(i-1,j-1) \\
a_1 &= (1/2)\Big[\epsilon(i,j) + \epsilon(i,j-1)\Big] \\
a_2 &= (1/2)\Big[\epsilon(i,j-1) + \epsilon(i-1,j-1)\Big] \\
a_3 &= (1/2)\Big[\epsilon(i-1,j) + \epsilon(i-1,j-1)\Big] \\
a_4 &= (1/2)\Big[\epsilon(i,j) + \epsilon(i-1,j)\Big] \ .
\end{aligned}$$

Finally, we put it all together to find

$$\oint_{S_{ij}} \approx -a_0 V(i,j) + a_1 V(i+1,j) + a_2 V(i,j-1) + a_3 V(i-1,j) + a_4 V(i,j+1) \ . \tag{30}$$

Including the right-hand term, we finally arrive at

$$\boxed{-a_0 V(i,j) + a_1 V(i+1,j) + a_2 V(i,j-1) + a_3 V(i-1,j) + a_4 V(i,j+1) = -Q(i,j)} \ . \tag{31}$$

Just like Equation (11), this expression represents a numerical stencil for the generalized Poisson equation. It is therefore a straightforward matter to generate a system of linear equations of the form $\mathbf{Ax} = \mathbf{b}$. However, just like before, $\mathbf{A}$ is a very large and sparse matrix, thereby making direct inversion an impractical option. We shall therefore follow the same procedure of successive over-relaxation by first solving for $V(i,j)$:

$$V(i,j) = \frac{1}{a_0}\Big[a_1 V(i+1,j) + a_2 V(i,j-1) + a_3 V(i-1,j) + a_4 V(i,j+1) + Q(i,j)\Big] \ . \tag{32}$$

We next define the residual as

$$R(i,j) = \frac{1}{a_0}\Big[a_1 V(i+1,j) + a_2 V(i,j-1) + a_3 V(i-1,j) + a_4 V(i,j+1) + Q(i,j)\Big] - V(i,j) \ . \tag{33}$$

And once again, the iteration formula is exactly as we found before:

$$V^{k+1}(i,j) = V^k(i,j) + \omega R^k(i,j) \ . \tag{34}$$

5 Electric Fields

In order to properly extract the electric field intensity out of the simulation, one must be careful to avoid any of the discontinuities that occur at the boundaries between dielectrics. This is readily accomplished as long as we remember that only the normal components of electric fields are actually discontinuous across boundaries. We therefore define a new staggered grid for the electric fields by using the convention depicted in Figure 6. In terms of the the central-difference approximation, this may be written as

$$E_x(i,j) = -\frac{V(i+1,j) - V(i,j)}{h} \ , \tag{35}$$

$$E_y(i,j) = -\frac{V(i,j+1) - V(i,j)}{h} \tag{36}$$

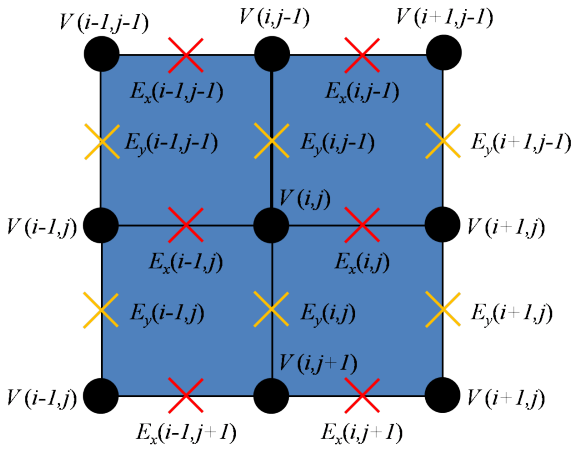

Figure 6: Grid stencil for obtaining the electric field samples. Note how only the tangential components of the E-field are defined at any of the boundaries, thus preventing any discontinuities to exist within the mesh.

The downside to this staggered grid approach is that E-field components are not sampled along the same points in space. One simple way to fix this is by choosing a new set of sample locations that lie in between the boundaries and then averaging over the nearest neighbors. Letting E'_x and E'_y represent the new set of grid samples, this is written as

$$E'_x(i,j) = (1/2)\Big[E_x(i,j+1) + E_x(i,j)\Big] \tag{37}$$

$$E'_y(i,j) = (1/2)\Big[E_y(i+1,j) + E_y(i,j)\Big]\ . \tag{38}$$

This new geometry is depicted in Figure 7, which shows that the new electric field samples are effectively placed at the same grid locations as the permittivities. Such an arrangement avoids any of the confusion that occurs at the boundaries between dielectric surfaces, since normal components are discontinuous. It also places all E-field samples along a grid that is interior to the voltage samples. For this reason, the number of rows and columns in the E-field and permittivity matrices will be one less than those of the voltage matrix.

6 Example: Parallel-Plate Capacitor

In this final section, we conclude with an example simulation of a parallel-plate capacitor using FDM with SOR. The simulation domain is 338×205 grid steps in size, and uses Dirichlet boundaries of $V = 0$ at all four edges. Using a relaxation factor of $\omega = 1.9$, this simulation required only 670 iterations before converging. Figure 8 shows the voltage mapping of the system. The top plate is at a potential of $+1.0$ V while the bottom plate is -1.0 V. Superimposed on the image is a quiver plot to represent the electric field vectors.

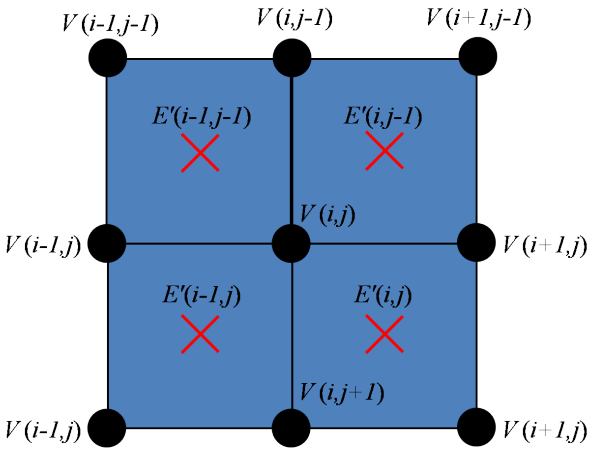

Figure 7: Modified E-field stencil. Each field component is defined by the average between the two nearest neighbors from the previous stencil. This places the E-fields along the same staggered grid as the permittivities.

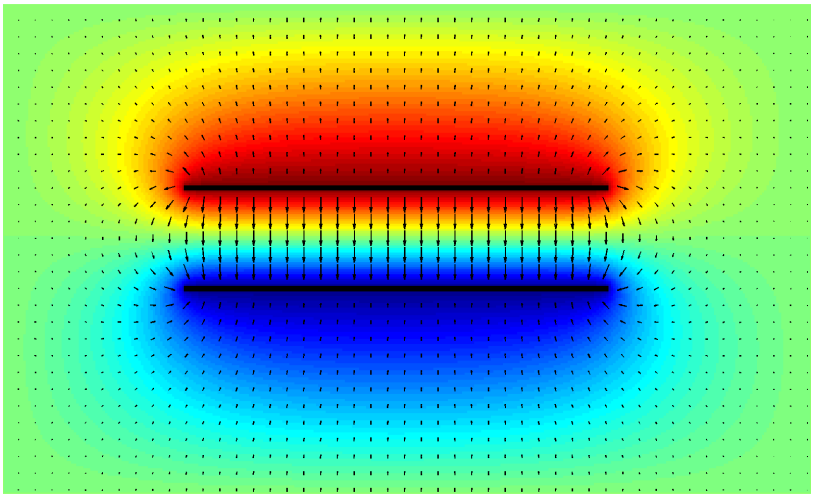

Figure 8: FDM simulation of a parallel-plate capacitor. The top plate is at a potential of +1.0 V while the bottom plate is at −1.0 V. The color mapping represents voltage potential throughout space while the quiver plots represent electric field vectors.