A STUDY OF A RECENT, MOMENT-METHOD ALGORITHM THAT IS ACCURATE TO VERY LOW FREQUENCIES

M. Burton and S. Kashyap
Electronic Warfare Division
Defence Research Establishment Ottawa
Ottawa, Ontario, Canada K1A 0Z4

Abstract—**We give an alternative description of a recently published moment-method algorithm, which uses divergence-free and rotation-free basis functions to maintain accuracy down to very low frequencies. The basic algorithm is restricted to simply-connected and non-self-intersecting surfaces. But this restriction has little practical impact—we show how multiply-connected surfaces, self-intersecting surfaces, and one-sided surfaces can easily be converted to the required topology without changing the solution. We examine a claim that the impedance matrix is diagonally dominant, which implies a guaranteed-to-converge Jacobi type of iterative solution of the matrix equation. Finally, we show how to control catastrophic-cancellation errors that occasionally appear in the voltage vector.**

INTRODUCTION

In moment-method algorithms finite computer memory and CPU time impose a high-frequency limit at about the first or second resonance of the body being tested. Typically, $W/\lambda_{min} \approx 1$, where W is some representative linear dimension of the body. A popular algorithm for finding the current on arbitrarily shaped conductors was described by Rao, Wilton and Glisson over ten years ago [1]. With this algorithm, a finite computer resource of another kind—word length—sets a *low*-frequency limit. Typically, $W/\lambda_{max} \approx 10^{-3}$, which limits the usefulness of the algorithm as a tool for studying electrostatic or magnetostatic problems.

Wilton, Lim, and Rao have recently described a new algorithm that has a much smaller low-frequency limit [2,3]. It is very effective and has an appealing structure. We think it deserves wide exposure and, therefore, we present here our own description of it and of the problem that it cures.

The basic algorithm is restricted to surfaces that are simply-connected and not self-intersecting. But this restriction has little practical impact—we show that multiply-connected surfaces, self-intersecting surfaces, and even one-sided surfaces can easily be converted into the required topological form without altering the solutions obtained for them.

The new algorithm also has a feature that may make it superior over the entire moment-method bandwidth. Its impedance matrix is closer to being diagonally dominant than the one from the earlier, Rao-Wilton-Glisson algorithm. This makes the matrix equation more likely to yield to an iterative method of solution. We discuss this possibility.

Finite word-length can occasionally corrupt some of the voltage vector of this new algorithm. We present a simple and effective fix based on the Faraday law.

THE CAUSES OF THE LOWER LIMIT

In the Rao-Wilton-Glisson algorithm a set of "rooftop" functions approximates the **J** field on a triangulated surface of the body [1]. There is a rooftop function for each interior edge of the triangulation. The domain of a rooftop is restricted to the pair of faces that share the interior edge. See Figure 1. The definition for the rooftop anchored to the i^{th} interior edge is

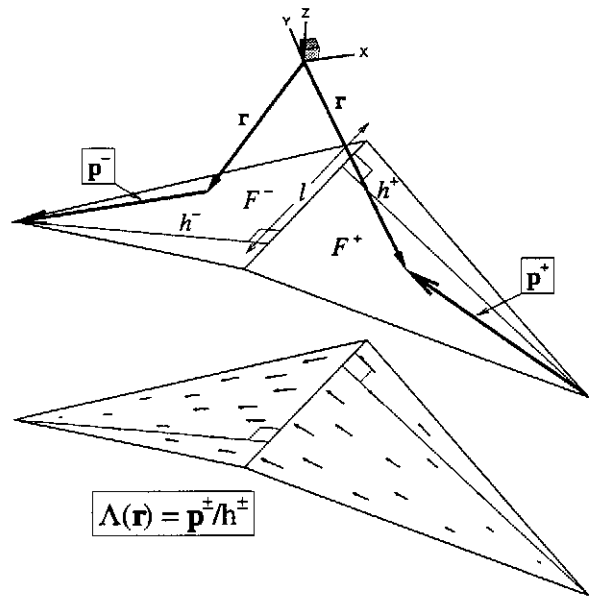

Figure 1 Parameters for defining a rooftop function.

$$\Lambda_i(\mathbf{r}) = \begin{cases} \dfrac{\mathbf{p}_i^{\pm}}{h_i^{\pm}}, & \mathbf{r} \in F_i^{\pm} \\ 0, & \textit{otherwise} \end{cases} \qquad (1)$$

where $h_i^{\pm}$ is the height of face $F_i^{\pm}$, when measured from the anchor edge to the free vertex; and $\mathbf{p}^{\pm}$ is $(\pm 1)\times$(the vector from the free vertex to $\mathbf{r}$). The area of face $F_i^{\pm}$ is $a_i^{\pm}$. The length of the anchor edge is l_i.

The rooftop function has four nice features: (1) it automatically satisfies Kirchhoff's current law at the anchor edge; (2) there is a charge density of $\rho_i^{\pm} = \pm l_i/(-j\omega a_i^{\pm})$ on each face, yet there is no *net* charge deposited—no need to test for conservation of charge on the body; (3) it is simple enough to permit efficient, robust, and accurate computation of the potential integrals [4]; (4) it imposes no restriction on the topology of surface—the surface can be open, closed, simply-connected, disjoint, multiply-connected, one-sided, two-sided, or self-intersecting.

There is a testing integral associated with each rooftop function. Its domain is an *open* path, beginning at the centroid of F^+ and following a streamline of $\Lambda(\mathbf{r})$ across the anchor edge to the centroid of F^-. Figure 2 shows the set of paths from a typical triangulation of a square plate. For each path there is a corresponding equation in the moment-method matrix equation $\boldsymbol{Zx} = \boldsymbol{b}$ stating that the tested scattered field $\int \mathbf{E}_{scat} \cdot d\mathbf{l}$ is equal to the negative of the tested incident field $-\int \mathbf{E}_{inc} \cdot d\mathbf{l}$.

But the scattered field is computed in two parts: $\mathbf{E}_{scat} = -\nabla V(\rho) - j\omega\mathbf{A}(\mathbf{J})$. The first part, $\nabla V(\rho)$, is inversely proportional to the frequency, because $\rho_i^{\pm} = \pm l_i/(-j\omega a_i^{\pm})$. The second part, $j\omega\mathbf{A}$, is directly proportional to the frequency. So, as the frequency goes lower, the $\mathbf{E}_{scat}$ field becomes more dominated by ∇V and testing integrals of $\mathbf{E}_{scat}$ become more path-independent (because ∇V is a conservative vector field). Ultimately, an integral over any given path in Figure 2 becomes identical to one over any concatenation of other paths having the same endpoints. The corresponding effect on $\boldsymbol{Z}$ is that any row becomes identical to a linear combination of other rows. In other words, as the frequency goes lower, $\boldsymbol{Z}$ becomes more ill-conditioned. Ultimately, it becomes singular.

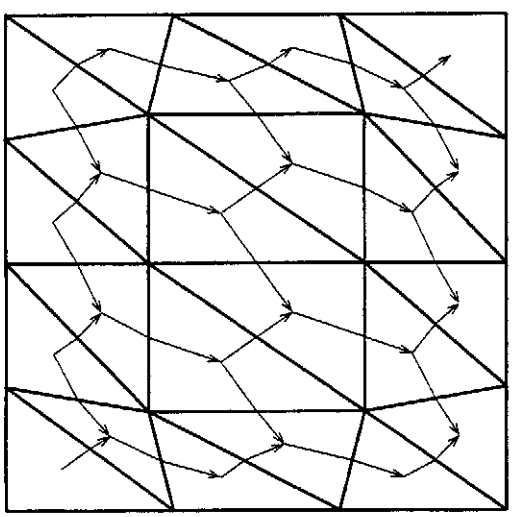

Figure 2 Integration paths for testing a rooftop model of a square plate.

Increasing condition number is one cause of the poor accuracy at low frequencies. Another cause is the loss of information about $j\omega\mathbf{A}$, due to insufficient word length.

As the frequency goes lower, the two quantities ∇V and $j\omega\mathbf{A}$ become increasingly different in magnitude, yet *their sum is stored in the same word.* The recoverable precision of the smaller quantity, $j\omega\mathbf{A}$, is inversely proportional to the square of the frequency.

But $j\omega\mathbf{A}$ is no less important than ∇V, even at extremely low frequencies. Although all of the paths in Figure 2 are open there are linear combinations of them that form *closed* paths. This implies that there are linear of combinations of equations in $\boldsymbol{Zx} = \boldsymbol{b}$ stating that the *closed*-path integral $\oint \mathbf{E}_{scat} \cdot d\mathbf{l}$ is equal to the *closed*-path integral $-\oint \mathbf{E}_{inc} \cdot d\mathbf{l}$. But ∇V has no curl (because it is a gradient of a scalar). By Stokes theorem it contributes nothing to a closed-path integral. The only contributor is $j\omega\mathbf{A}$.

The algorithm ignores this. While numerically computing the integrals $-\oint \mathbf{E}_{inc} \cdot d\mathbf{l}$, it needlessly computes the *theoretically vanishing* integrals $\oint \nabla V \cdot d\mathbf{l}$. The influence of $j\omega\mathbf{A}$ is lost in the resulting, unnoticed, catastrophic cancellations.

An ill-conditioned $\boldsymbol{Z}$ and loss of precision in $j\omega\mathbf{A}$ are the two causes of poor low-frequency accuracy. They, and the cure, were first described in [2] and [3].

HOW THE LOW-FREQUENCY ERRORS ARE REMOVED

To prevent catastrophic cancellations Wilton, Lim, and Rao *explicitly* create each closed path, which allows them to replace the numerical integration of $\oint \nabla V \cdot d\mathbf{l}$ with the exact value of zero. They finesse each closed-path integral into existence by designing a basis function with that integral as its testing integral. The new basis function is a superposition of original rooftop functions so that the new testing integral is a superposition of original testing integrals.

In this way, most of the original algorithm, and its corresponding computer code, stay the same.

But how to find *all possible* closed paths? The answer is to make a set of generator loops, that is, a set of closed paths from which all other closed paths can be assembled. Then, avoiding (numerical computation of) $\oint \nabla V \cdot d\mathbf{l}$ integrals on every loop in the set will avoid them on all possible loops. Making a set of generators is easy. A set of rooftop functions sharing a common node will, if they have the proper reference directions, have a testing integral whose path is a generator. To see this, pick an interior node in Figure 2 and reverse some of the nearby integration paths so that they all have the same sense of rotation around the node. This is a "loop function".

Loop functions do an excellent job of preserving the information contained in the $j\omega\mathbf{A}$ field. But they are an incomplete cure for low-frequency errors. There are two reasons. First, loop functions do not address the condition-number issue. Second, the charges deposited on each face of a loop function tend to cancel. In fact a body with equilateral triangulation would produce *complete* cancellation. (Recall that $\rho_i^{\pm} = \pm l_i / (-j\omega a_i^{\pm})$.)

A second basis function needs to be invented; one whose testing path cannot become part of a concatenation of paths in a path-independent integral, and thereby increase the condition number. And the second function should be *guaranteed* to contribute to the charge—the charges from contributing rooftop functions should *accumulate* rather than cancel. The set of rooftop functions sharing a common face will do both of these things, if their reference directions are properly chosen. As an illustration, pick a face in Figure 2 and reverse some of its integration paths so that they all have the same sense, outward or inward. The charges on the chosen face will accumulate because they all have the same sign. This is a "star function". Notice that the testing paths of star functions cannot be concatenated to form a non-reversing path around a loop.

Star functions anchored to interior faces are made from three rooftop functions. Those anchored to faces on boundaries are made from fewer rooftop functions.

Loop functions describe most of the rotation $\oint \mathbf{J} \cdot d\mathbf{l}$ of the $\mathbf{J}$ field. Star functions describe most of the divergence $\nabla \cdot \mathbf{J}$ of the $\mathbf{J}$ field. The duties of rotation and divergence usually would be shared by both functions, since the triangulation is rarely equilateral. Wilton, Lim, and Rao remove this awkward feature by adjusting the intensity of the contributing rooftop functions to compensate for unequal edge lengths. They define the loop function anchored to the i^{th} interior node as follows:

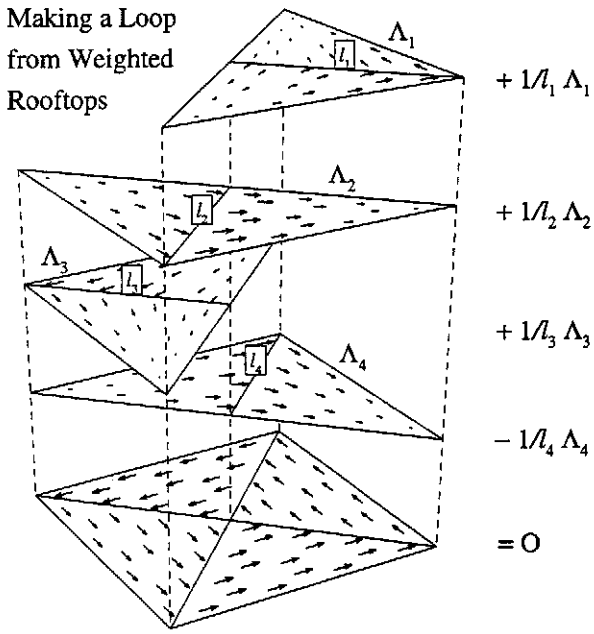

Figure 3 Making a loop function from weighted rooftop functions.

$$\mathbf{O}_i(\mathbf{r}) = \sum_{n_i \in loop_i} \frac{\sigma_{n_i}}{l_{n_i}} \mathbf{\Lambda}_{n_i}(\mathbf{r}) \qquad (2)$$

where $\mathbf{\Lambda}_{n_i}$ is the n^{th} member of the set of rooftop functions whose common node is the anchor node of $\mathbf{O}_i$; $\sigma = \pm 1$; and l_{n_i} is the length of the n^{th} edge attached to the anchor node. On each face there are two contributions to the charge density; $1/l_{n_i}$ gives them identical magnitudes; σ_{n_i} makes them cancel. In doing so, σ_{n_i} also forces each rooftop current to have the same direction as the loop orientation. The $1/l$ weighting causes the $\mathbf{J}$ streamlines to form closed paths around the anchor node (rotational flow). See Figure 3.

Wilton, Lim, and Rao define the star function anchored to the i^{th} face as follows:

$$\bigstar_i(\mathbf{r}) = \sum_{n_i \in face_i} \frac{\nu_{n_i}}{l_{n_i}} \mathbf{\Lambda}_{n_i}(\mathbf{r}) \qquad (3)$$

where $\mathbf{\Lambda}_{n_i}$ is the n^{th} of the set of rooftop functions whose common face is the anchor face of $\bigstar_i$; $\nu = \pm 1$; and l_{n_i} is the length of the n^{th} edge of the anchor face. On the anchor face there are (usually) three contributions to the charge; $1/l_{n_i}$ gives them equal magnitudes; ν_{n_i} makes them accumulate. In doing so, ν_{n_i} also forces all rooftop currents to flow out of the face. The $1/l$ weighting causes the $\mathbf{J}$ streamlines to emanate from the *centroid* of the anchor

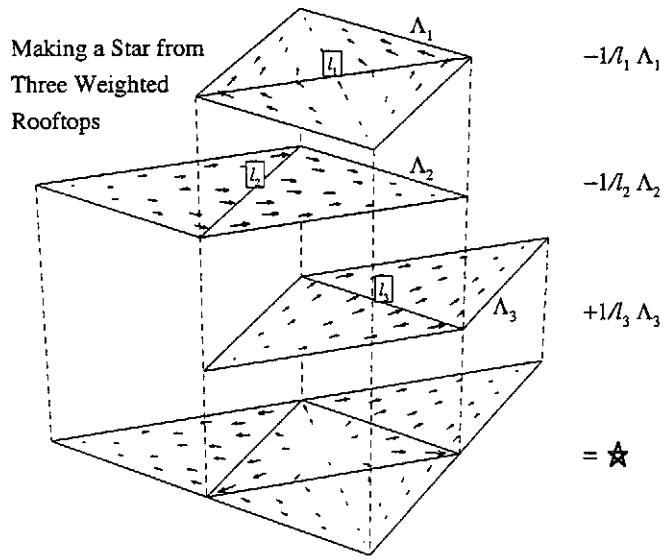

Figure 4 Making a star function from weighted rooftop functions.

face (but only when there are *three* contributing rooftop functions). See Figure 4.

HELMHOLTZIAN COMPLEMENTARITY

When defined this way, loop and star functions have pure Helmholtzian complementarity [5]. Loop functions describe only the rotation of the **J** field; star functions describe only the divergence of the **J** field. This feature is unaffected by the triangulation of the surface. A loop function's contributions to a star function's irrotational flow come in self-cancelling pairs. Hence, no superposition of loop functions can produce an irrotational flow of current. Similarly, a star function's contributions to a loop function's rotational flow also come in self-cancelling pairs. Hence, no superposition of star functions can produce a rotational flow of current.

There is no coupling via $\int \nabla V \cdot d\mathbf{l}$ between the two kinds of functions; star-to-loop coupling is zero because the path is a loop; loop-to-star coupling is zero because the integrand is zero. The only vehicle for mutual influence between loops and stars is the integral $j\omega \int \mathbf{A} \cdot d\mathbf{l}$. Hence, the two sets of functions become mutually independent as the frequency goes to zero.

Loops and stars form a complete description of the **J** field on a simply-connected surface, as does the original set of rooftop functions from which they were assembled. Hence, a moment-method equation $\boldsymbol{Zx} = \boldsymbol{b}$ will have the same number of unknowns whether it is based on a loop-and-star model or on a rooftop model. But only if the surface is simply-connected. This is the source of a lingering difficulty with the definition of the loop function. (See MULTIPLY-CONNECTED SURFACES below.)

PERFECTLY CONDUCTING SPHERE IN A MAGNETOSTATIC FIELD

Figure 5 shows a perfectly conducting sphere of 1 m radius immersed in a static magnetic field of 1 A/m. The exact, analytical solution for this problem is $\mathbf{J}_\phi = -1.5\, H_{inc} \sin\theta\, \mathbf{u}_\phi$ [6].

To do a numerical simulation we illuminated the sphere with a 377 V/m plane wave at 30 kHz. We then sampled the current that crossed the dashed line. The results show that the loss of information in $j\omega\mathbf{A}$ and the large condition number of $\boldsymbol{Z}$ have made the rooftop model useless, and that the loop-and-star model suffers from neither of these ailments.

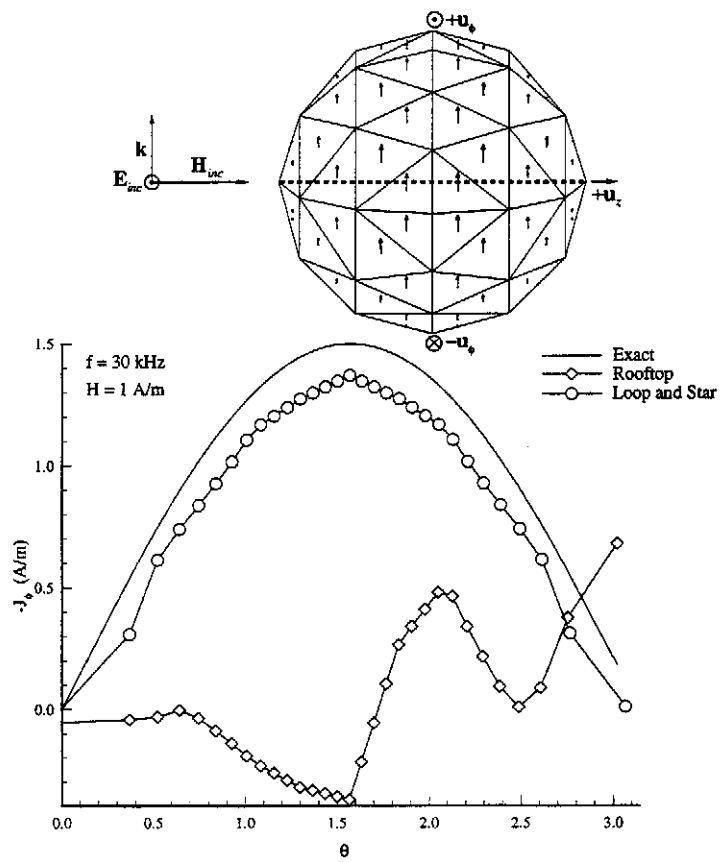

Figure 5 Comparison of surface currents computed by rooftop and loop-and-star models of a conducting sphere.

CONDITION NUMBER

Figure 6 shows the LINPACK estimate [7] of the condition number from a rooftop model and from a loop-and-star model of the sphere in Figure 5. The condition number from the rooftop model is inversely proportional to the square of the frequency, down to about 100 kHz. Below 100 kHz the curve becomes more complicated. We suspect that the roundoff error in the entries in $\boldsymbol{Z}$ actually *prevent* the straight $1/f^2$ trend towards total dependence below 100 kHz.

The condition number from the loop-and-star model is frequency independent down to at least 3 Hz, showing how effectively the condition number has been controlled.

The LINPACK estimate of the condition number is most accurate when all entries in $\boldsymbol{Z}$ have the same absolute error. Since the error in each z_{ij} is roughly proportional to its magnitude, we meet this requirement by scaling the rows and columns of $\boldsymbol{Z}$ so that all entries on the diagonal have unit magnitude. (Both rows *and* columns were scaled to preserve the near symmetry of the matrix.)

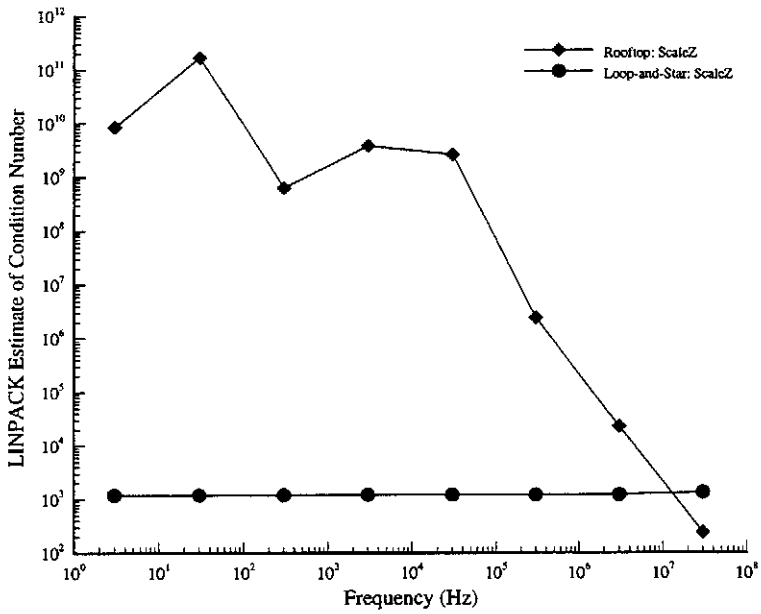

Figure 6 Comparison of condition numbers from the rooftop and loop-and-star models.

SIMPLY-CONNECTED OPEN AND CLOSED SURFACES

When the surface is simply-connected there is no difficulty with the loop-and-star model. The loop function definition (2) forms a complete set of generator loops. There is a loop function at each interior node if the surface is open; one less if the surface is closed (to avoid making a redundant generator).

The number of star functions is whatever is needed to make the total number of unknowns equal to the number of unknowns from a rooftop model.

DISJOINT BODIES

Both the original Rao-Wilton-Glisson rooftop model and the new loop-and-star model compute the mutual influence between sources on *arbitrarily located* faces. It does not matter that those faces might reside on disjoint surfaces. (However, each surface must be simply-connected if the loop-and-star model is to be used).

MULTIPLY-CONNECTED SURFACES

When the surface is multiply-connected the loop function definition (2) yields an incomplete set of generator loops. As examples, it misses a loop for every aperture, and on a torus it misses two loops.

The most obvious solution would be to broaden the definition in some manner to include the missing loops. But each missing loop is just one of a large set of topologically equivalent loops; any one of them can be used. It might not be possible to design software that reliably and efficiently finds the required sets of topologically equivalent loops and then picks only one of them.

A workable alternative would be to supply the missing loops by hand after inspecting a three-dimensional view of the triangulated surface. Each loop could then be supplied to the program in the form of the set of faces that are traversed. The software would need to be capable of assembling the faces into a loop function, which would then be processed the same as the original loop functions.

Another alternative also involves some manual labor. It is based on the fact that loop functions and star functions are assembled from rooftop functions. Suppose there was some way to change the topology to simply-connected without, in any way, changing the set of rooftops. The solution would then be identical to that produced by either of the previous alternatives, because the *same* set of rooftops would be used to assemble the loops and stars.

Exactly that happens when the surface is "cut" along a selected sequence of edges in such a way that the

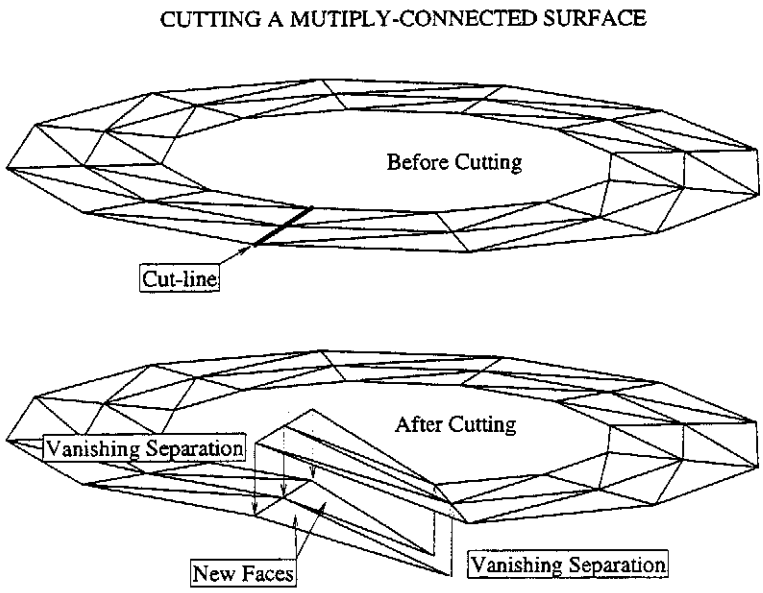

Figure 7 Cutting a multiply-connected surface into a simply-connected surface.

overlapping rooftop functions on each face become topologically, *but not physically*, delaminated. See Figure 7. The surface becomes simply-connected *yet the size, shape, and position of its rooftops are not changed.*

Three new edges, and one new face, are added for each edge in the cut. But the number of *interior* edges is not changed. There are new nodes but no new node coordinates.

We use the cutting option because it changes only the *triangulation* of the surface; the *algorithm* remains in its basic form. Cutting involves manual work, which takes time. But creating the triangulated surface itself usually involves manual work, and sometimes a lot of it. This is at least as time consuming as the cutting.

number of unknowns in the matrix equation would be *greater* than the number of interior edges in the triangulation. This method would leave the triangulation intact and would change the code.

The original Rao-Wilton-Glisson rooftop algorithm could easily be encoded to use this method but the loop-and-star algorithm might present some trouble. It might be hard to write code that reliably assembles loop-and-star functions in the vicinity of a generalized edge.

For us, a better method is to cut the self-intersecting surface into surfaces that are not self-intersecting. See Figure 8. The cut surface is then disjoint and can be treated by *the basic forms of* either the rooftop model or the loop-and-star model. (See DISJOINT SURFACES above.) As with multiply-connected surfaces, cutting a self-intersecting surface does not change

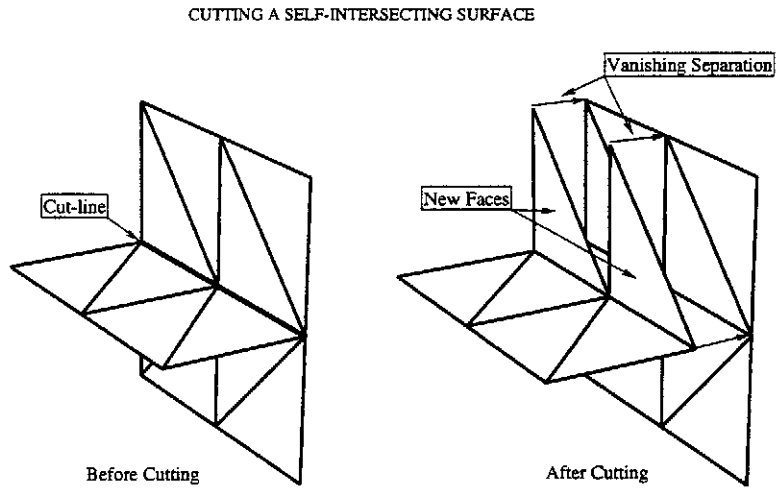

Figure 8 Cutting a self-intersecting surface into a disjoint surface.

SELF-INTERSECTING SURFACES

When a self-intersecting body is triangulated, some edges will have more than two faces attached to them. There are (at least) two methods that could handle this situation, both of which satisfy Kirchhoff's current law at each edge.

In one method, each edge at a self-intersection could be treated as a generalized edge: an edge with f faces attached to it would be assigned $f-1$ rooftop functions. No special care would be needed when choosing which pair of faces to assign to which rooftop function—rooftop currents could even pass *through* each other at a generalized edge. The in any way the fundamental set of rooftops, so the solution does not change either.

ONE-SIDED SURFACES

A one-sided (non-orientable) surface, like a Moebius strip, can always be cut into a two-sided surface in the manner already shown in Figure 7. Once cut, the surface becomes simply-connected, allowing the basic loop-and-star model to be used. As always, cutting changes the triangulation and leaves the encoded algorithm intact.

Although the encoded version of the original rooftop algorithm of Rao, Wilton, and Glisson is not the main topic of this paper, its popularity and the present heading make it appropriate to now discuss some pertinent details.

The original Rao-Wilton-Glisson rooftop code cannot be applied to a one-sided surface; nor was it intended to be. Its authors chose to give the user another way to check the integrity of the triangulation: they designed the code to compute the volume whenever a closed surface was encountered. So a consistent normal had to be assigned to every face. For simplicity this was done for every type of surface, even open ones. It then seemed convenient to use the normal to help set a reference direction for each rooftop function. Each rooftop current was made to cross its anchor edge in the same direction as $\hat{\mathbf{e}} \times \hat{\mathbf{n}}$, where $\hat{\mathbf{e}}$ is the (arbitrarily defined) direction of the anchor edge.

But a basis function's reference direction is arbitrary in any algorithm. (In the Rao-Wilton-Glisson algorithm it depends on an arbitrary choice of the direction of $\hat{\mathbf{e}}$.) There is no need to involve the surface normal in an arbitrary decision. By removing it from the decision process, the rooftop model of Rao, Wilton, and Glisson immediately can be applied to one-sided surfaces (and it can still compute the volume of closed surfaces).

Finally, we note that the subroutine that sets the surface normals (and, hence, the reference directions) is quite complex and contains a subtle bug, which takes effect only on some triangulations.

DIAGONAL DOMINANCE AND ITERATIVE SOLUTIONS OF $\boldsymbol{Zx} = \boldsymbol{b}$

It turns out that a matrix from a loop-and-star model is closer to being diagonally dominant than one from a rooftop model. $\boldsymbol{Z}$ is said to be diagonally dominant when the "diagonal-dominance" ratio (DDR)

$$\frac{|z_{ii}|}{\sum_{j:j\neq i} |z_{ij}|} \qquad (4)$$

is greater than 1 for all $i=1,\cdots,N$ [12].

The explanation for the greater DDR lies in the shape of the paths of the testing integrals: in loop-and-star modeling the paths are made from directed line-segments whose orientations tend to cancel; in rooftop modeling the orientations tend to accumulate. This makes the off-diagonal integrals smaller, relative to the diagonal integral, than those in a rooftop model.

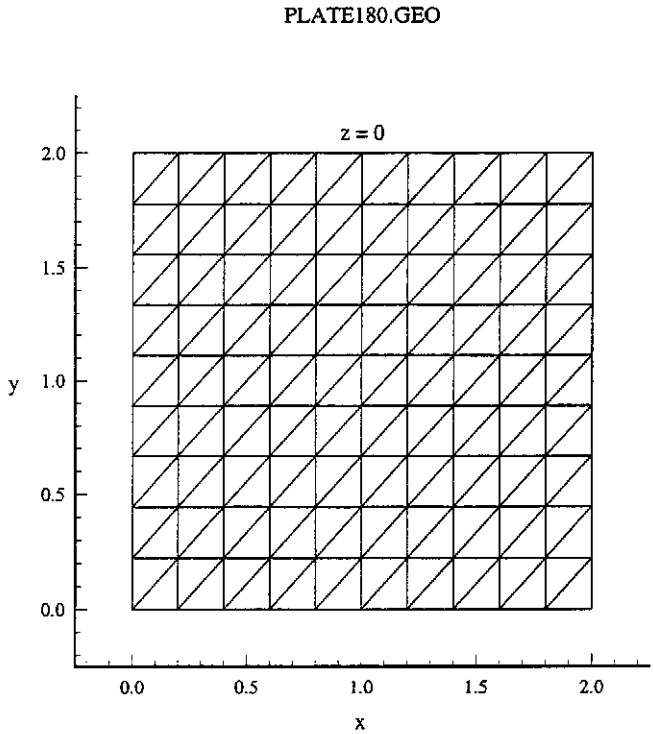

Figure 9 The square plate used to study diagonal dominance.

Figure 10 compares the DDR from a rooftop model to the DDR from a loop-and-star model. (The device-under-test is the square plate in Figure 9.) The rows that test the loops (rows 1 to 72) have DDRs consistently greater than 1. The DDRs of the rows that test stars are about half as big and are usually less than one. Other structures that we have studied also show this behavior. (Perhaps star functions could be redesigned to produce DDRs as high as those of loop functions.)

The DDR is insensitive to the frequency as long as the triangulation is dense enough to allow an adequate description of the sources on the surface. See Figure 11. Also, it is totally independent of the direction and polarization of the incident field, because $\boldsymbol{Z}$ itself is totally independent of these parameters.

Wilton, Lim, and Rao claim that the $\boldsymbol{Z}$ matrix is diagonally dominant [2,3]. We have found no geometries where this is so. We can only say that the $\boldsymbol{Z}$ from a loop-and-star model is closer to being diagonally dominant than one from a rooftop model.

In any case, the improved diagonal dominance may have a practical benefit at all frequencies (below the upper limit). The standard method for solving $\boldsymbol{Zx} = \boldsymbol{b}$ is by LU decomposition, for which the required CPU time is proportional to N^3. (N is the number of unknowns). The

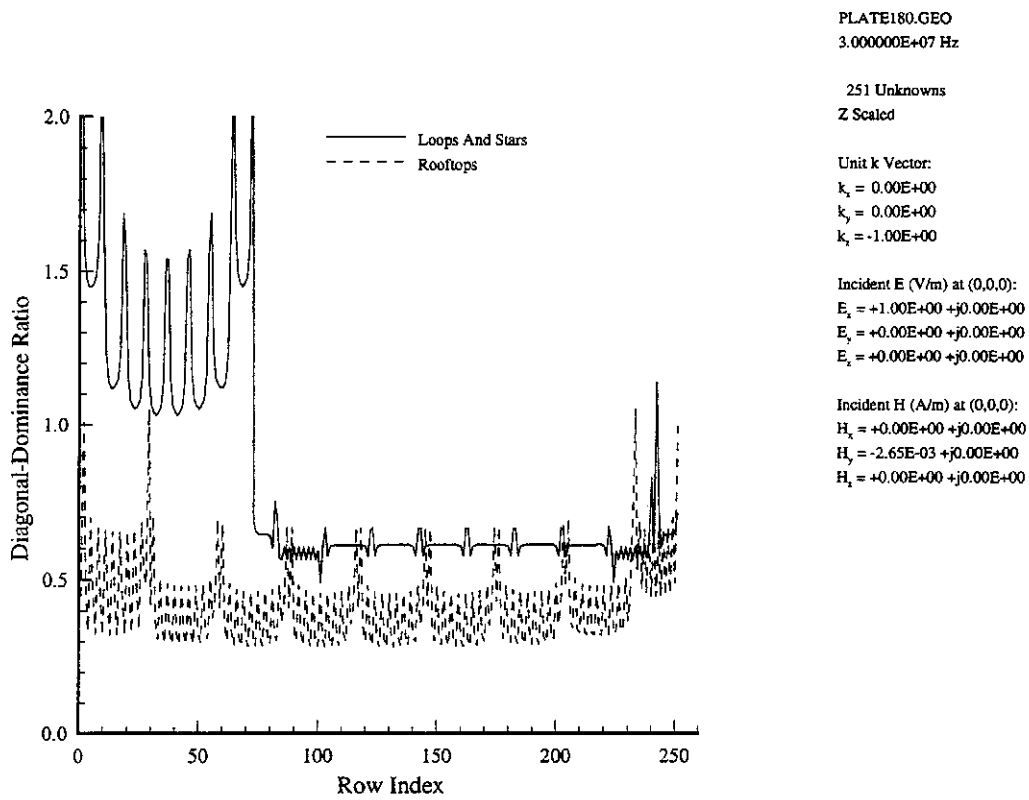

Figure 10 Comparison of diagonal-dominance ratios for a square plate.

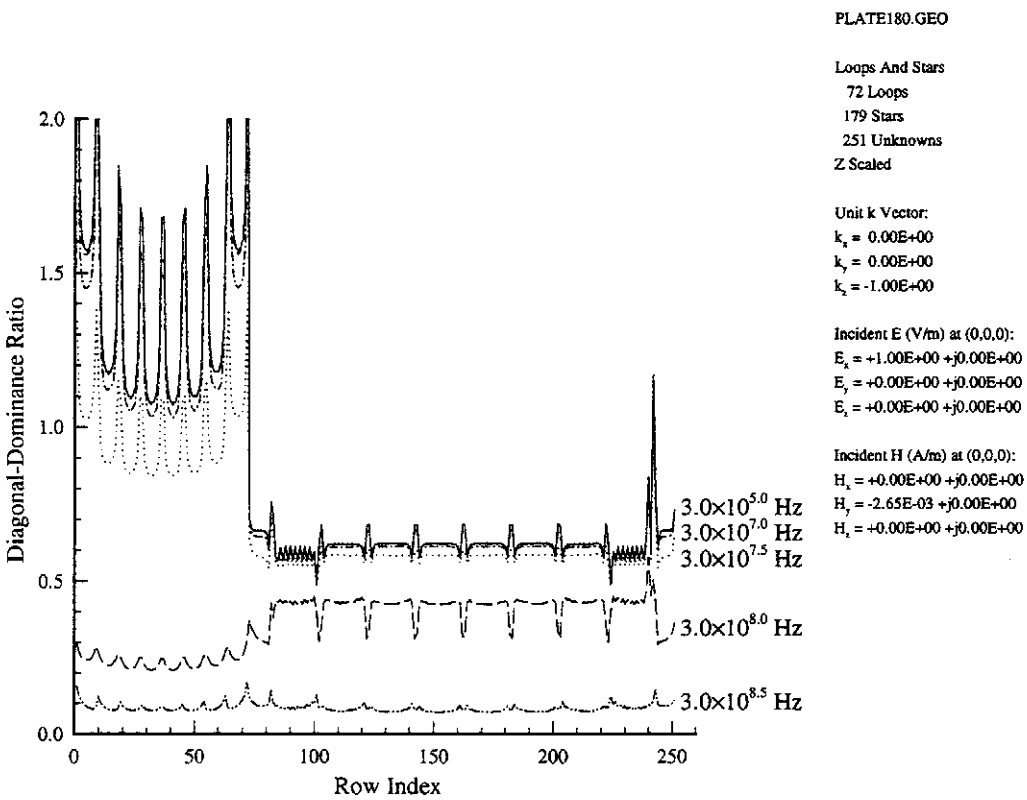

Figure 11 The effect of frequency on the diagonal-dominance ratio.

CPU time for an iterative solution, such as Jacobi iteration, is proportional to N^2, a great advantage when N is large. Solution by point Jacobi iteration is *guaranteed* to converge when $\boldsymbol{Z}$ is diagonally dominant [8,9,10,11]. A matrix equation in which $\boldsymbol{Z}$ is not diagonally dominant can *sometimes* permit a converging Jacobi iteration. Intuition suggests the closer $\boldsymbol{Z}$ is to diagonal dominance the more likely the iteration will converge and, if it does converge, the faster it will converge.

The number of iterations needed to get a solution depends on the initial guess. Hence, an iterative solution has an extra advantage when doing a sweep of a spectrum since each solution is a good initial guess for the next.

$\boldsymbol{Zx} = \boldsymbol{b}$ could also be solved by *block* Jacobi iteration. Loop-and-star analysis permits a natural way to divide $\boldsymbol{Z}$ into submatrices, if the loops are numbered from 1 to N_{loops} and the stars are numbered from $N_{loops}+1$ to N. The equation then becomes

$$\begin{bmatrix} Z^{oo} & Z^{*o} \\ Z^{o*} & Z^{**} \end{bmatrix} \begin{bmatrix} x^{o} \\ x^{*} \end{bmatrix} = \begin{bmatrix} b^{o} \\ b^{*} \end{bmatrix} \qquad (5)$$

which can be solved by block Jacobi iteration, if $\boldsymbol{Z}$ is well behaved. Good behaviour is guaranteed at 0 Hz because Z^{*o} and Z^{o*} are then filled with zeroes. (See HELMHOLTZIAN COMPLEMENTARITY above.) It may be that the Helmholtzian complementarity of loop functions and star functions will guarantee the convergence of this iteration at higher frequencies.

The potential speed-up is not as great as that for point Jacobi iteration because both Z^{oo} and Z^{**} need to be LU decomposed. The required time, relative to an LU decomposition of $\boldsymbol{Z}$, is $(1/3)^3 + (2/3)^3 = 1/3$ (since the numbers of loops and stars are approximately $N/3$ and $2N/3$, respectively).

CATASTROPHIC CANCELLATIONS IN MAKING b^o

Each entry in b^o in equation (5) is the integral $\oint \mathbf{E}_{inc} \cdot d\mathbf{l}$ around the closed testing path of a loop function. Sometimes the orientation of the surface and polarization of the illumination can combine so that the integral vanishes. This will happen, for example, when a uniform plane wave is normally incident on a flat plate. Since the integral is computed numerically, the zero is crudely approximated by catastrophic cancellations. To avoid this, we use Faraday's law to replace the closed-path integral $\oint \mathbf{E}_{inc} \cdot d\mathbf{l}$ with the surface integral $-j\omega \int_S \mathbf{B}_{inc} \cdot \hat{\mathbf{n}} da$, where $\boldsymbol{S}$ is some surface bounded by the path. A convenient choice for $\boldsymbol{S}$ is the part of the triangulated surface enclosed by the testing path. We do the numerical integration by sampling $\mathbf{B}_{inc}$ at the centroids of the four-sided "facelets" inside the testing path. See Figure 12. The centroid of each facelet is at $5/12\,\mathbf{r}_i + 7/12\,\mathbf{r}^c_{n_i}$. The area of the n^{th} facelet is one-third the area of the n^{th} face.

There are some path integrals for which the Faraday alternative is hard to implement. They are the testing integrals of those loop functions that are lost whenever the surface is multiply-connected. The paths of these integrals will not enclose any triangulated surfaces at all! In these cases the Faraday alternative would require the construction of special (non-conducting) surfaces on which to sample $\mathbf{B}_{inc}$. Also the basic loop-and-star model would have to be modified to handle multiply-connected surfaces. We avoid these problems by cutting the multiply-connected surface into a simply-connected one. (See MULTIPLY-CONNECTED SURFACES above.) Then no testing path will enclose an untriangulated surface.

It is conceivable that a numerical evaluation of the Faraday surface integral can also suffer from catastrophic cancellation. It could happen when the loop function is anchored at the tip of a very sharp cone or at a very sharp edge on the surface. But cases like these will not occur if the surface is triangulated so that there are no acute angles between adjacent faces. The no-acute-angle rule is implied by the standard rule for a good triangulation: it must be

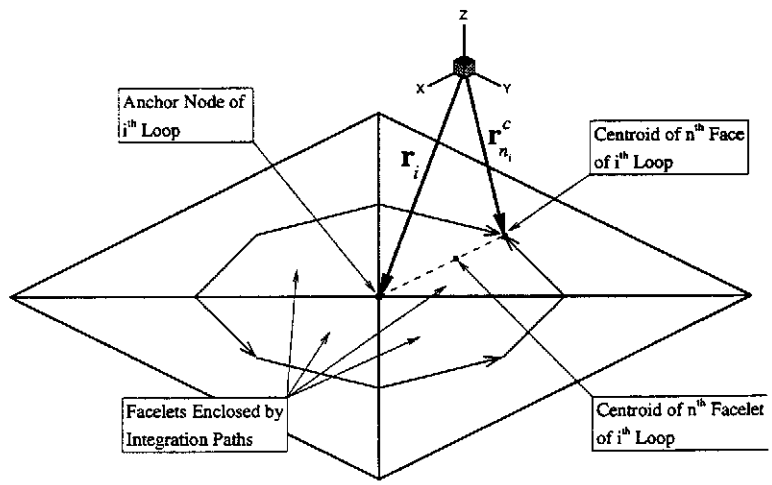

Figure 12 Computing the centroid of a facelet inside an integration loop.

dense enough to adequately describe the surface and the sources on it.

The CPU times needed to compute the b^O vector by either side of the Faraday equation are roughly the same and are minuscule when compared to the time needed to fill Z.

PROGRAMMING

Each entry z_{ij} in Z is the integration of $\mathbf{E}_{scat}$, due to the j^{th} basis function, over the testing path of the i^{th} basis function. The most straightforward way to fill Z is to compute an integral for each permutation of i and j encountered in a column-wise scan of Z. (Row-wise if the language is not FORTRAN.) But this is not the most time-efficient way. Each integration calls a potential-integral routine [4], which computes the potentials at a *single* field point due to a source on a *single* face. Because basis functions overlap and testing paths overlap, a scan of Z would actually call this routine many times with the *same combination of source face and field point.* It is therefore much faster to do the computation for a given combination *only once* and to then *add* the result to *all* the z_{ij} that are affected by that combination. This requires a scan of the geometry for all permutations of faces taken two at a time (since there is a field point at the centroid of each face.)

Notice that the scan is independent of the choice of model; the faces of loop functions, star functions, and rooftop functions all carry the same Rao-Wilton-Glisson source function (to within a weighting factor). Hence, it is natural to design the program to do either loop-and-star modeling or rooftop modeling. There is only a small amount of code that actually depends on the choice of basis function. It determines which z_{ij} (and which b_i) are to receive an accumulation. It also determines the weighting of each accumulation.

In our program this basis-function-specific code is executed *before* any entries in Z or b are computed. It simply fills two identically dimensioned arrays. One array links each edge of each face to the weighting factors, $\pm 1/l$, of the four possible functions (two loops and two stars) that cross that edge. The other array supplies the corresponding indices, $i = 1,N$, of the loops or stars. The inner dimension of each array is four because there can be up to four basis functions that can cross each interior edge.

The rest of the code operates without any knowledge of the type of basis function being used. It simply uses the two generic arrays of weighting factors and indices to fill Z and b, and then solves $Zx = b$. For example, if the program is to do rooftop modeling then there is only one weighting factor, ± 1, and one index per edge. The remaining three locations along the inner dimension are simply filled with zeroes so that the other accumulations amount to nothing.

The time to execute the basis-function-specific code is very insignificant when compared to the $O(N^3)$ time needed to solve $Zx = b$ and the $O(N^2)$ time needed to fill Z.

CONCLUSIONS

The loop-and-star algorithm of Wilton, Lim, and Rao effectively controls the catastrophic cancellation errors and condition number that corrupt the low-frequency results from the rooftop algorithm of Rao, Wilton, and Glisson. The high-frequency limit of the loop-and-star algorithm is as large as that of the rooftop algorithm, $W/\lambda \approx 1$; and the low-frequency limit is much smaller, at least $W/\lambda \approx 2\times10^{-8}$. Electrostatic and magnetostatic simulations can now be done with confidence.

The basic loop-and-star algorithm can treat only simply-connected surfaces. But this restriction has little practical impact because multiply-connected surfaces, self-intersecting surfaces, and one-sided surfaces can be readily "cut" into the required topology without altering the set of rooftop functions that they carry.

The impedance matrix from a loop-and-star model is not diagonally dominant. But it is more so than one from a rooftop model. In fact, the rows of Z that test the loop functions all have a diagonal-dominance ratio greater than one. It is only the rows that test the star functions that have a diagonal-dominance ratio less than one. It might be possible to redesign the star function so that all rows of Z are diagonally dominant. In any case, the matrix equation $Zx = b$ from a loop-and-star model is more likely to be solved by a converging Jacobi iteration.

Catastrophic cancellations can corrupt those entries in the b vector that participate in loop testing. This can be avoided by using the Faraday law to replace the closed-path integral $\oint \mathbf{E}_{inc} \cdot d\mathbf{l}$ with the surface integral $-j\omega \int_S \mathbf{B}_{inc} \cdot \hat{\mathbf{n}} da$, where S is some surface bounded by the testing path.

REFERENCES

[1] S. M. Rao, D. R. Wilton, and A. W. Glisson, "Electromagnetic Scattering by Surfaces of Arbitrary Shape," ***IEEE Trans. Antennas Propagat.***, vol. AP-30, no. 3, May 1982, pp. 409-418.

[2] D. R. Wilton, J. S. Lim, and S. M. Rao, "A Novel Technique to Calculate the Electromagnetic Scattering by Surfaces of Arbitrary Shape," ***URSI***

Radio Science Meeting, University of Michigan, June 1993, page 322.

[3] J. S. Lim, ***Electromagnetic Scattering from Arbitrarily Shaped Bodies at Very Low Frequency Range Using Triangular Patch Modelling***, Ph. D. dissertation, Auburn University, Auburn, AL 36849, March 1994.

[4] D. R. Wilton, S. M. Rao, A. W. Glisson, D. H. Schaubert, O. M. Al-Bundak, and C. M. Butler, "Potential Integrals for Uniform and Linear Source Distributions on Polygonal and Polyhedral Domains," ***IEEE Trans. Antennas Propagat.***, vol. AP-32, no. 3, March 1984, pp. 276-281.

[5] R. Plonsy and R. E. Collin, ***Principles and Applications of Electromagnetic Fields***, International Student Edition, McGraw-Hill Book Co. and Kogakusha Co., 1961, pp.29-38.

[6] J. Van Bladel, ***Electromagnetic Fields***, revised printing, Hemisphere Publishing Corporation, New York, 1985, pp. 275-279.

[7] J. J. Dongarra, C. B. Moler, J. R. Bunch, and G. W. Stewart, ***LINPACK Users' Guide***, SIAM, Philadelphia, 1979, pp. 1.11-1.22.

[8] L. A . Hageman and D. M. Young, ***Applied Iterative Methods***, Academic Press, 1981, pp. 23-26.

[9] G. E. Forsythe and C. B. Moler, ***Computer Solution of Linear Algebraic Equations***, Prentice-Hall, 1967, p. 130.

[10] W. H. Press, S.A. Teukolsky, W. T. Veterling, and B. P. Flannery, ***Numerical Recipes in FORTRAN, The Art of Scientific Computing***, Cambridge University Press, 1992, p. 856.

[11] K. Kalbasi and K. Demarest, "A Multilevel Formulation of the Method of Moments," ***IEEE Trans. Antennas Propagat.***, vol. AP-41, no. 5, May 1993, pp. 589-599.

[12] G. H. Golub, C. F. van Loan, ***Matrix Computations***, Second Edition, The Johns Hopkins University Press, 1989, p. 119.