# Computational Electronics 

# Numerical Analysis 

Prepared by:
Dragica Vasileska
Associate Professor
Arizona State University

## Numerical Solution of Algebraic Equations

The solution of linear systems of algebraic equations is an important subject of linear algebra [1], and the computational considerations needed for computer implementation are usually treated in some detail in introductory numerical methods courses. This section simply represents a quick review or overview of the subject - it is not intended as a complete treatise on this topic. Students with little or no background in this area are referred to one of many good numerical methods texts that treat the subject in more detail. The numerical solution of large systems of algebraic equations is a direct consequence of the Finite Difference method for solving ordinary differential equations (ODEs) or partial differential equations (PDEs). Recall that the goal in these techniques is to break the continuous differential equation into a coupled set of algebraic difference equations for each finite volume or node in the system. When one has only a single independent variable (the ODE case), this process can easily lead to several hundred simultaneous equations that need to be solved. For multiple independent variables (the PDE case), systems with hundreds of thousands of equations are common. Thus, in general, we need to be able to solve large systems of linear equations of the form $A x=b$ as part of the solution algorithm for general Finite Difference methods.

There are two general schemes for solving linear systems: Direct Elimination Methods and Iterative Methods. All the direct methods are, in some sense, based on the standard Gauss Elimination technique, which systematically applies row operations to transform the original system of equations into a form that is easier to solve. In particular, this section overviews an algorithm for implementation of the basic Gauss Elimination scheme and it also highlights the LU Decomposition method which, although functionally equivalent to the Gauss Elimination method, does provide some additional flexibility for computer implementation. Thus, the LU decomposition method is often the preferred direct solution method for low to medium sized systems (usually less than 200-300 equations).

For large systems, iterative methods (instead of direct elimination methods) are almost always used. This switch is required from accuracy considerations (related to round-off errors), from memory limitations for physical storage of the equation constants, from considerations for treating nonlinear problems, and from overall efficiency concerns. There are several specific iterative schemes that are in common use, but most methods build upon the base Gauss Seidel
method, usually with some acceleration scheme to help convergence. Thus, our focus in this brief overview is on the basic Gauss Seidel scheme and on the use of Successive Over Relaxation (SOR) to help accelerate convergence. We also give a brief introduction to the incomplete-lower-upper (ILU) decomposition method and a short tutorial to the multigrid method for solving 2D and 3D problems.

## Direct Methods

## 1. Gauss Elimination Method

The Gauss Elimination Method forms the basis for all elimination techniques. The basic idea is to modify the original equations, using legal row operations, to give a simpler form for actual solution. The basic algorithm can be broken into two stages:

1. Forward Elimination (put equations in upper triangular form)
2. Back Substitution (solve for unknown solution vector)

To see how this works, consider the following system of equations:

$$
\begin{align*}
& a_{11} x_{1}+a_{12} x_{2}+\ldots a_{1 N} x_{N}=b_{1} \\
& a_{21} x_{1}+a_{22} x_{2}+\ldots a_{2 N} x_{N}=b_{2}  \tag{1}\\
& \vdots \\
& a_{N 1} x_{1}+a_{N 2} x_{2}+\ldots a_{N N} x_{N}=b_{N}
\end{align*}
$$

Now, with reference to this system of $N$ equations and $N$ unknowns, the Forward Elimination Step (with partial pivoting) becomes:
Step 0: $\quad$ Create an augmented matrix, $\tilde{A}=[A b]$
Step 1: Determine the coefficient in the $i^{\text {th }}$ column with the largest absolute value and interchange rows such that this element is the pivot element $(i=1,2,3$, to $N-1)$
Step 2: $\quad$ Normalize the pivot equation (i.e. divide by the $i, i$ element)
Step $3 \quad$ Multiply normalized eqn. $i$ by the $j, i$ element of eqn. $j$
Step $4 \quad$ Subtract the resultant equation in Step 3 from eqn. $j$
repeat Steps 3 and 4 for $j=i+1$ to $N$
go to Step 1 for next $i=i+1$ to $N-1$
and the Back Substitution Step is given by:
Step $5 \quad x_{N}=b_{N}^{\prime} / a_{N N}^{\prime}$

Step $6 \quad x_{i}=\left(b_{i}-\sum_{j=i+1}^{N} a_{i j} x_{j}\right) / a_{i i}^{d}$
repeat for $i=N-1, N-2$, to 1
Note: The primes here indicate that the coefficients at this stage are different from the original coefficients.

## 2. The LU Decomposition Method

The Gauss Elimination Method has the disadvantage that all right-hand sides (i.e. all the $b$ vectors of interest for a given problem) must be known in advance for the elimination step to proceed. The LU Decomposition Method outlined here has the property that the matrix modification (or decomposition) step can be performed independent of the right hand side vector. This feature is quite useful in practice - therefore, the LU Decomposition Method is usually the Direct Scheme of choice in most applications.

To develop the basic method, let us break the coefficient matrix into a product of two matrices,

$$
\begin{equation*}
A=L U \tag{2}
\end{equation*}
$$

where $L$ is a lower triangular matrix and $U$ is an upper triangular matrix. Now, the original system of equations $A x=b$, becomes

$$
\begin{equation*}
L U x=b \tag{3}
\end{equation*}
$$

This expression can be broken into two problems,

$$
\begin{equation*}
L y=b, \quad U x=y \tag{4}
\end{equation*}
$$

The rationale behind this approach is that the two systems given in Eq. (4) are both easy to solve; one by forward substitution and the other by back substitution. In particular, because $L$ is a lower triangular matrix, the expression $L y=b$ can be solved with a simple forward substitution step. Similarly, since $U$ has upper triangular form, $U x=y$ can be evaluated with a simple back substitution algorithm.

Thus the key to this method is the ability to find two matrices $L$ and $U$ that satisfy Eq. (4). Doing this is referred to as the Decomposition Step and there are a variety of algorithms available. Three specific approaches are as follows:

- Doolittle Decomposition:

$$
\left[\begin{array}{cccc}
1 & 0 & 0 & 0  \tag{5}\\
l_{21} & 1 & 0 & 0 \\
l_{31} & l_{32} & 1 & 0 \\
l_{41} & l_{42} & l_{43} & 1
\end{array}\right]\left[\begin{array}{cccc}
u_{11} & u_{12} & u_{13} & u_{14} \\
0 & u_{22} & u_{23} & u_{24} \\
0 & 0 & u_{33} & u_{34} \\
0 & 0 & 0 & u_{44}
\end{array}\right]=\left[\begin{array}{llll}
a_{11} & a_{12} & a_{13} & a_{14} \\
a_{21} & a_{22} & a_{23} & a_{24} \\
a_{31} & a_{32} & a_{33} & a_{34} \\
a_{41} & a_{42} & a_{43} & a_{44}
\end{array}\right]
$$

Because of the specific structure of the matrices, a systematic set of formulae for the components of $L$ and $U$ results.

- Crout Decomposition:

$$
\left[\begin{array}{cccc}
l_{11} & 0 & 0 & 0  \tag{6}\\
l_{21} & l_{22} & 0 & 0 \\
l_{31} & l_{32} & l_{33} & 0 \\
l_{41} & l_{42} & l_{43} & l_{44}
\end{array}\right]\left[\begin{array}{cccc}
1 & u_{12} & u_{13} & u_{14} \\
0 & 1 & u_{23} & u_{24} \\
0 & 0 & 1 & u_{34} \\
0 & 0 & 0 & 1
\end{array}\right]=\left[\begin{array}{llll}
a_{11} & a_{12} & a_{13} & a_{14} \\
a_{21} & a_{22} & a_{23} & a_{24} \\
a_{31} & a_{32} & a_{33} & a_{34} \\
a_{41} & a_{42} & a_{43} & a_{44}
\end{array}\right]
$$

The evaluation of the components of $L$ and $U$ is done in a similar fashion as above.

- Cholesky Factorization:

For symmetric, positive definite matrices, where $A=A^{T}$ and $x^{T} A x>0$ for all $\mathrm{x} \neq 0$ then,

$$
\begin{equation*}
U=L^{T} \quad \text { and } \quad A=L L^{T} \tag{7}
\end{equation*}
$$

and a simple set of expressions for the elements of $L$ can be obtained (as above). Once the elements of $L$ and $U$ are available (usually stored in a single $N \times N$ matrix), Matlab's standard equation solver (using the backslash notation, $x=A l b$ ), uses several variants of the basic LU Decomposition method depending on the form of the original coefficient matrix (see the Matlab help files for details).

## 3. LU Decomposition in 1D

The LU decomposition method is very trivial for 1D problems where the discretization of the ODE or the PDE leads to a three point stencil and a tridiagonal matrix $A$. It is easy to show, that the system of equations that we need to solve, and for the purpose of clarity, we denote by $A x=f$, in matrix form reads

$$
\left[\begin{array}{cccccc}
a_{1} & c_{1} & 0 & 0 & \ldots & 0  \tag{8}\\
b_{2} & a_{2} & c_{2} & 0 & \ldots & 0 \\
0 & b_{3} & a_{3} & c_{3} & \ldots & 0 \\
\vdots & & & & & \\
\vdots & & & & & \\
0 & 0 & 0 & \ldots & b_{n} & a_{n}
\end{array}\right]\left[\begin{array}{c}
x_{1} \\
x_{2} \\
x_{3} \\
\vdots \\
\vdots \\
x_{n}
\end{array}\right]=\left[\begin{array}{c}
f_{1} \\
f_{2} \\
f_{3} \\
\vdots \\
\vdots \\
f_{n}
\end{array}\right] .
$$

The solution of this problem can be represented as a two-step procedure that is explained below.
Step 1: Decompose the coefficient matrix $A$ into a product of lower and upper triangular matrices:

$$
A=L U=\left[\begin{array}{cccccc}
1 & 0 & 0 & 0 & \ldots & 0  \tag{9}\\
\beta_{2} & 1 & 0 & 0 & \ldots & 0 \\
0 & \beta_{3} & 1 & 0 & \ldots & 0 \\
\vdots & & & & & \\
\vdots & & & & & \\
0 & 0 & 0 & \ldots & \beta_{n} & 1
\end{array}\right]\left[\begin{array}{cccccc}
\alpha_{1} & c_{1} & 0 & 0 & \ldots & 0 \\
0 & \alpha_{2} & c_{2} & 0 & \ldots & 0 \\
0 & 0 & \alpha_{3} & c_{3} & \ldots & 0 \\
\vdots & & & & & \\
\vdots & & & & & \\
0 & 0 & 0 & \ldots & 0 & \alpha_{n}
\end{array}\right]
$$

From the equality of the two matrices, we have:

$$
\begin{equation*}
\alpha_{1}=a_{1} ; \quad \beta_{k}=b_{k} / \alpha_{k-1} ; \quad \alpha_{k}=a_{k}-\beta_{k} c_{k-1} ; \quad k=2,3, \ldots, n \tag{10}
\end{equation*}
$$

Step 2: Solve the system of equations $L U x=f$, by first solving $L g=f$ using forward substitution, and then solving $U x=g$ using backward substitution. Then, the solution of $L g=f$ is represented as:

$$
\begin{equation*}
g_{1}=f_{1} ; \quad g_{k}=f_{k}-\beta_{k} g_{k-1} ; \quad k=2,3, \ldots, n, \tag{11}
\end{equation*}
$$

And the solution of $U x=g$ as:

$$
\begin{equation*}
x_{n}=g_{n} / \alpha_{n} ; \quad x_{k}=\left[g_{k}-c_{k} x_{k+1}\right] / \alpha_{k} ; \quad k=n-1, n-2, \ldots, 2,1 \tag{12}
\end{equation*}
$$

## Iterative Methods

For large systems of equations, an iterative solution scheme for the unknown vector can always be written in the form

$$
\begin{equation*}
x^{p+1}=B x^{P}+c \tag{13}
\end{equation*}
$$

where $B$ is the iteration matrix, $c$ is a constant vector and $p$ is an iteration counter. Convergence of this scheme is guaranteed if the largest eigenvalue of the iteration matrix is less that unity, where $\rho=$ spectral radius $=|\lambda|_{\max }$. Therefore, if $\rho<1$ the iterative scheme will converge. If
$\rho \ll 1$, the iterative scheme converges very rapidly. If $\rho \approx 1$ but less than unity, the scheme will be slowly converging. The iteration algorithm will diverge if the spectral radius is greater than unity. Convergence is tested during the iterative process by computing the largest relative change from one iteration to the next, and comparing the absolute value of this result with some desired tolerance. If the maximum relative change is less than the desired accuracy, then the process is terminated. If this condition is not satisfied, then another iteration is performed.

## 1. The Gauss Seidel Method

Let us take the original system of equations given by $A x=b$ and convert it into the classical Gauss Seidel iterative scheme. To do this, let us break the original matrix into three specific components, or

$$
\begin{equation*}
A=L+D+U \tag{14}
\end{equation*}
$$

where the three matrices on the right hand side, in sequence, are strictly lower triangular, diagonal, and strictly upper triangular matrices. Now, substituting this into the original expression gives

$$
\begin{equation*}
(L+D) x+U x=b \tag{15}
\end{equation*}
$$

or

$$
\begin{equation*}
(L+D) x=b-U x \tag{16}
\end{equation*}
$$

If we premultiply by $(L+D)^{-1}$ and notice that the solution vector appears on both sides of the equation, we can write the equation in an iterative form as

$$
\begin{equation*}
x^{p+1}=-(L+D)^{-1} U x^{p}+(L+D)^{-1} b \tag{17}
\end{equation*}
$$

Clearly this is in the standard form for iterative solutions as defined in Eq. (13), where the iteration matrix is given by

$$
\begin{equation*}
B=-(L+D)^{-1} U \tag{18}
\end{equation*}
$$

and the constant vector is written as

$$
\begin{equation*}
c=(L+D)^{-1} b \tag{19}
\end{equation*}
$$

This form of the iteration strategy is useful for the study of the convergence properties of model problems. It is, however, not particularly useful as a program algorithm for code implementation. For actual implementation on the computer, one writes these equations differently, never having to formally take the inverse as indicated above. In practice, Eq. (17) is written in iterative form as

$$
\begin{equation*}
D x^{p+1}=b-L x^{p+1}-U x^{p} \tag{20}
\end{equation*}
$$

or

$$
\begin{equation*}
x^{p+1}=D^{-1}\left(b-L x^{p+1}-U x^{p}\right) \tag{21}
\end{equation*}
$$

This specific form is somewhat odd at first glance, since $x^{p+1}$ appears on both sides of the equation. This is justified because of the special form of the strictly lower triangular matrix, $L$. This can be seen more clearly if the matrix equations are written using discrete notation. In discrete form Eq. (21) can be expanded as

$$
\begin{equation*}
x_{i}^{p+1}=\frac{1}{a_{i i}}\left(b_{i}-\sum_{j=1}^{i-1} a_{i j} x_{j}^{p+1}-\sum_{j=i+1}^{N} a_{i j} x_{j}^{p}\right) \tag{22}
\end{equation*}
$$

where the diagonal elements of $D^{-1}$ are simply $1 / a_{i i}$ and the limits associated with the summations account for the special structure of the $L$ and $U$ matrices.

## 2. The Successive Over-Relaxation (SOR) Method

To improve the rate of convergence, one might consider using a weighted average of the results of the two most recent estimates to obtain the next best guess of the solution. If the solution is converging, this might help extrapolate to the real solution more quickly. This idea is the basis of the SOR method. In particular, let $\alpha$ be some weight factor with a value between 0 and 2. Now, let's compute the next value of $x^{p+1}$ to use in the Gauss-Seidel method as a linear combination of the current value, $x^{p+1}$, and the previous solution, $x^{p}$, as follows:

$$
\begin{equation*}
\left.x^{p+1}\right|_{\text {new }}=\alpha x^{p+1}+(1-\alpha) x^{p} \quad \text { with } \quad 0<\alpha<2 \tag{23}
\end{equation*}
$$

Note that if $\alpha$ is unity, we simply get the standard Gauss Seidel method (or whatever base iterative scheme is in use). When $\alpha$ is greater that unity, the system is said to be over-relaxed, indicating that the latest value, $x^{p+1}$, is being weighted more heavily (weight for $x^{p}$ is negative). If, however, $\alpha$ is less than one, the system is under-relaxed, this time indicating that the previous solution, $x^{p}$, is more heavily weighted (positive weight values). The idea, of course, is to choose the relaxation parameter to improve convergence (reduce the spectral radius). This is most often done in a trial and error fashion for certain classes of problems (experience helps here). Some more advanced codes do try to estimate this quantity as part of the iterative calculation, although this is not particularly easy.

## 3. Incomplete LU decomposition for 2D and 3D problems

Within incomplete factorization schemes [2] for 2D problems, the matrix $A$ is decomposed into a product of lower $(L)$ and upper $(U)$ triangular matrices, each of which has four non-zero diagonals in the same locations as the ones of the original matrix $A$. The unknown elements of the $L$ and $U$ matrices are selected in such a way that the five diagonals common to both $A$ and $A^{\prime}=L U$ are identical and the four superfluous diagonals represent the matrix $N$, i.e., $A^{\prime}=A+N$. Thus, rather than solving the original system of equations $A x=b$, one solves the modified system $L U x=b+N x$, by solving successively the matrix equations $L V=b+N x$ and $V=U x$, where $V$ is an auxiliary vector. It is important to note that the four superfluous terms of $N$ affect the rate of convergence of the ILU method. Stone [3] suggested the introduction of partial cancellation, which minimizes the influence of these additional terms and accelerates the rate of convergence of the ILU method. By using a Taylor series expansion, the superfluous terms appearing in $A^{\prime}$ are partially balanced by subtracting approximately equal terms.

## 4. Multigrid Method

The multi-grid method represents an improvement over the SOR and ILU methods in terms of iterative techniques available for solving large systems of equations [4]. The basic principle behind the multi-grid method is to reduce different Fourier components of the error on grids with different mesh sizes. Most iterative techniques work by quickly eliminating the highfrequency Fourier components, while the low-frequency ones are left virtually unchanged. The result is a convergence rate that is initially fast, but slows down dramatically as the highfrequency components disappear. The multi-grid method utilizes several grids, each with consecutively coarser mesh sizes. Each of these grids acts to reduce a different Fourier component of the error, therefore increasing the rate of convergence with respect to single grid based methods, such as an SOR.

Practical multigrid methods were first introduced in the 1970s by Brandt [5]. These methods can solve elliptic PDEs discretized on $N$ grid points in $O(N)$ operations. The "rapid" direct elliptic solvers discussed in Ref. [6] solve special kinds of elliptic equations in $O(N \log N)$ operations. The numerical coefficients in these estimates are such that multigrid methods are comparable to the rapid methods in execution speed. Unlike the rapid methods, however, the multigrid methods can solve general elliptic equations with nonconstant coefficients with hardly
any loss in efficiency. Even nonlinear equations can be solved with comparable speed. Unfortunately there is not a single multigrid algorithm that solves all elliptic problems. Rather there is a multigrid technique that provides the framework for solving these problems. You have to adjust the various components of the algorithm within this framework to solve your specific problem. We can only give a brief introduction to the subject here. In this approach, the method obtains successive solutions on finer and finer grids. You can stop the solution either at a prespecified fineness, or you can monitor the truncation error due to the discretization, quitting only when it is tolerably small.

## From One-Grid, through Two-Grid, to Multigrid

The key idea of the multigrid method can be understood by considering the simplest case of a two-grid method. Suppose we are trying to solve the linear elliptic problem

$$
\begin{equation*}
L u=f \tag{24}
\end{equation*}
$$

where $L$ is some linear elliptic operator and $f$ is the source term. When one discretizes Eq. (24) on a uniform grid with mesh size $h$, the resulting set of linear algebraic equations arises

$$
\begin{equation*}
L_{h} u_{h}=f_{h} \tag{25}
\end{equation*}
$$

Let $\tilde{u}_{h}$ denote some approximate solution to Eq. (25). We will use the symbol $u_{h}$ to denote the exact solution to the difference equations. Then the error in $\tilde{u}_{h}$ or the correction is

$$
\begin{equation*}
v_{h}=u_{h}-\tilde{u}_{h} \tag{26}
\end{equation*}
$$

The residual or defect is

$$
\begin{equation*}
d_{h}=L_{h} \tilde{u}_{h}-f_{h} \tag{27}
\end{equation*}
$$

Since $L_{h}$ is linear, the error satisfies

$$
\begin{equation*}
L_{h} v_{h}=-d_{h} \tag{28}
\end{equation*}
$$

At this point we need to make an approximation to $L_{h}$ in order to find $v_{h}$. The classical iteration methods, such as Jacobi or Gauss-Seidel, do this by finding, at each stage, an approximate solution of the equation

$$
\begin{equation*}
\hat{L}_{h} \hat{v}_{h}=-d_{h} \tag{29}
\end{equation*}
$$

where $\hat{L}_{h}$ is a "simpler" operator than $L_{h}$. For example, $\hat{L}_{h}$ is the diagonal part of $L_{h}$ for Jacobi iteration, or the lower triangle for Gauss-Seidel iteration. The next approximation is generated by

$$
\begin{equation*}
\tilde{u}_{h}^{\text {new }}=\tilde{u}_{h}+\hat{v}_{h} \tag{30}
\end{equation*}
$$

Now consider, as an alternative, a completely different type of approximation for $L_{h}$, one in which we "coarsify" rather than "simplify." That is, we form some appropriate approximation $L_{H}$ of $L_{h}$ on a coarser grid with mesh size $H$ (we will always take $H=2 h$, but other choices are possible). The residual equation is now approximated by

$$
\begin{equation*}
L_{H} v_{H}=-d_{H} \tag{31}
\end{equation*}
$$

Since $L_{H}$ has smaller dimension, this equation will be easier to solve than Eq. (28). To define the defect $d_{H}$ on the coarse grid, we need a restriction operator $R$ that restricts $d_{h}$ to the coarse grid:

$$
\begin{equation*}
d_{H}=R d_{h} . \tag{32}
\end{equation*}
$$

The restriction operator is also called the fine-to-coarse operator or the injection operator. Once we have a solution $\tilde{v}_{H}$ to Eq. (31), we need a prolongation operator $P$ that prolongates or interpolates the correction to the fine grid:

$$
\begin{equation*}
\tilde{v}_{H} P \tilde{v}_{H} \tag{33}
\end{equation*}
$$

The prolongation operator is also called the coarse-to-fine operator or the interpolation operator. Both Rand Pare chosen to be linear operators. Finally the approximation $\tilde{u}_{h}$ can be updated:

$$
\begin{equation*}
\tilde{u}_{h}^{\text {new }}=\tilde{u}_{h}+\tilde{v}_{h} \tag{34}
\end{equation*}
$$

One step of this coarse-grid correction scheme is thus:

- Compute the defect on the fine grid from Eq. (27).
- Restrict the defect by Eq. (32).
- Solve Eq. (31) exactly on the coarse grid for the correction.
- Interpolate the correction to the fine grid by Eq. (33).
- Compute the next approximation by Eq. (34).

Let us contrast the advantages and disadvantages of relaxation and the coarse-grid correction scheme. Consider the error $v_{h}$ expanded into a discrete Fourier series. Call the components in the lower half of the frequency spectrum the smooth components and the highfrequency components the nonsmooth components. We have seen that relaxation becomes very slowly convergent in the limit $h \rightarrow 0$, i.e., when there are a large number of mesh points. The reason turns out to be that the smooth components are only slightly reduced in amplitude on each iteration. However, many relaxation methods reduce the amplitude of the nonsmooth components by large factors on each iteration: They are good smoothing operators. For the twogrid iteration, on the other hand, components of the error with wavelengths $2 H$ are not even
representable on the coarse grid and so cannot be reduced to zero on this grid. But it is exactly these high-frequency components that can be reduced by relaxation on the fine grid! This leads us to combine the ideas of relaxation and coarse-grid correction:

- Pre-smoothing: Compute $\bar{u}_{h}$ by applying $v_{1} \geq 0$ steps of a relaxation method to $\tilde{u}_{h}$.
- Coarse-grid correction: As above, using $\bar{u}_{h}$ to give $\bar{u}_{h}^{\text {new }}$.
- Post-smoothing: Compute $\tilde{u}_{h}^{\text {new }}$ by applying $v_{2} \geq 0$ steps of the relaxation method to $\bar{u}_{h}^{\text {new }}$.

It is only a short step from the above two-grid method to a multigrid method. Instead of solving the coarse-grid defect Eq. (31) exactly, we can get an approximate solution of it by introducing an even coarser grid and using the two-grid iteration method. If the convergence factor of the two-grid method is small enough, we will need only a few steps of this iteration to get a good enough approximate solution. We denote the number of such iterations by $\gamma$. Obviously, we can apply this idea recursively down to some coarsest grid. There the solution is found easily, for example by direct matrix inversion or by iterating the relaxation scheme to convergence. One iteration of a multigrid method, from finest grid to coarser grids and back to finest grid again, is called a cycle. The exact structure of a cycle depends on the value of $\gamma$, the number of two-grid iterations at each intermediate stage. The case $\gamma=1$ is called a V-cycle, while $\gamma=2$ is called a W-cycle (see Fig. 1). These are the most important cases in practice. Note that once more than two grids are involved, the pre-smoothing steps after the first one on the finest grid need an initial approximation for the error $v$. This should be taken to be zero.

## Smoothing, Restriction, and Prolongation Operators

The most popular smoothing method, and the one you should try first, is Gauss-Seidel, since it usually leads to a good convergence rate. The exact form of the Gauss-Seidel method depends on the ordering chosen for the mesh points. For typical second-order elliptic equations like our model problem, it is usually best to use red-black ordering, making one pass through the mesh updating the "even" points (like the red squares of a checkerboard) and another pass updating the "odd" points (the black squares). When quantities are more strongly coupled along one dimension than another, one should relax a whole line along that dimension simultaneously. Line relaxation for nearest-neighbor coupling involves solving a tridiagonal system, and so is still efficient. Relaxing odd and even lines on successive passes is called zebra relaxation and is
usually preferred over simple line relaxation. Note that SOR should not be used as a smoothing operator. The over-relaxation destroys the high-frequency smoothing that is so crucial for the multigrid method.

A succint notation for the prolongation and restriction operators is to give their symbol. The symbol of $P$ is found by considering $v_{H}$ to be 1 at some mesh point $(x, y)$, zero elsewhere, and then asking for the values of $P v_{H}$. The most popular prolongation operator is simple bilinear interpolation. It gives nonzero values at the 9 points $(x, y),(x+h, y), \ldots,(x-h, y-h)$, where the values are $1, \frac{1}{2}, \ldots, \frac{1}{4}$.

$\gamma=1$


Figure 1 Structure of multigrid cycles. S denotes smoothing, while E denotes exact solution on the coarsest grid. Each descending line $\backslash$ denotes restriction ${ }^{\circledR}$ and each ascending line / denotes prolongation (P). The finest grid is at the top level of each diagram. For the V-cycles $(\gamma=1)$ the E step is replaced by one 2-grid iteration each time the number of grid levels increases by one. For the W-cycles $(\gamma=2)$ each E step gets replaced by two 2-grid iterations.

Its symbol is therefore

$$
\left[\begin{array}{ccc}
\frac{1}{4} & \frac{1}{2} & \frac{1}{4}  \tag{35}\\
\frac{1}{2} & 1 & \frac{1}{2} \\
\frac{1}{4} & \frac{1}{2} & \frac{1}{4}
\end{array}\right]
$$

The symbol of $R$ is defined by considering $v_{h}$ to be defined everywhere on the fine grid, and then asking what is $R v_{h}$ at $(x, y)$ as a linear combination of these values. The simplest possible choice for $R$ is straight injection, which means simply filling each coarse-grid point with the value from the corresponding fine-grid point. Its symbol is "[1]." However, difficulties can arise in practice with this choice. It turns out that a safe choice for $R$ is to make it the adjoint operator to $P$. Then, take $P$ to be bilinear interpolation, and choose $u_{H}=1$ at $(x, y)$, zero elsewhere. Then take $P$ to be bilinear interpolation, and choose $u_{H}=1$ at $(x, y)$, zero elsewhere. Finally, the symbol of $R$ is

$$
\left[\begin{array}{ccc}
\frac{1}{16} & \frac{1}{8} & \frac{1}{16}  \tag{36}\\
\frac{1}{8} & \frac{1}{4} & \frac{1}{8} \\
\frac{1}{16} & \frac{1}{8} & \frac{1}{16}
\end{array}\right]
$$

Note the simple rule: The symbol of $R$ is $1 / 4$ the transpose of the matrix defining the symbol of $P$. This rule is general whenever $R=P^{\dagger}$ and $H=2 h$. The particular choice of $R$ in Eq. (36) is called full weighting. Another popular choice for $R$ is half weighting, "halfway" between full weighting and straight injection. Its symbol is

$$
\left[\begin{array}{ccc}
0 & \frac{1}{8} & 0  \tag{37}\\
\frac{1}{8} & \frac{1}{2} & \frac{1}{8} \\
0 & \frac{1}{8} & 0
\end{array}\right]
$$

A similar notation can be used to describe the difference operator $L_{h}$. For example, the standard differencing of the model problem, Eq. (29), is represented by the five-point difference star

$$
L_{h}=\frac{1}{h^{2}}\left[\begin{array}{ccc}
0 & 1 & 0  \tag{38}\\
1 & -4 & 1 \\
0 & 1 & 0
\end{array}\right]
$$

If you are confronted with a new problem and you are not sure what $P$ and $R$ choices are likely to work well, here is a safe rule: Suppose $m_{p}$ is the order of the interpolation $P$ (i.e., it interpolates polynomials of degree $m_{p}-1$ exactly). Suppose $m_{r}$ is the order of $R$, and that $R$ is the adjoint of some $P$ (not necessarily the $P$ you intend to use). Then if $m$ is the order of the differential
operator $L_{h}$, you should satisfy the inequality $m_{p}+m_{r}>m$. For example, bilinear interpolation and its adjoint, full weighting, for Poisson's equation satisfy $m_{p}+m_{r}=4>m=2$.

Of course the $P$ and $R$ operators should enforce the boundary conditions for your problem. The easiest way to do this is to rewrite the difference equation to have homogeneous boundary conditions by modifying the source term if necessary. Enforcing homogeneous boundary conditions simply requires the $P$ operator to produce zeros at the appropriate boundary points. The corresponding $R$ is then found by $R=P^{\dagger}$.

## References

1 G. Strang, Linear Algebra and Its Applications (Academic Press, New York, second edition, 1980)

2 G. V. Gadiyak and M. S. Obrecht, Simulation of Semiconductor Devices and Processes: Proceedings of the Second International Conference (1986) 147.

3 H. L. Stone, SIAM J. Numer. Anal., 5 (1968) 536.
4 W. Hackbush, Multi-Grid Methods and Applications (Springer-Verlag, Berlin, 1985).
5 A. Brandt, "Multi-level adaptive technique (MLAT): The Multi-grid Method", IBM Research Report RD-6026, IBM T. J. Watson Research Center, Yorktown Heights, Ney York, 1976.

6 R.E. Bank and D.J. Rose, "An $\mathrm{O}\left(\mathrm{N}^{*} \mathrm{~N}\right)$ method for solving constant coefficient boundary value problems in two dimensions," SIAM Journal of Numerical Analysis, Vol. 12, pp. 529-540 (1975).

