

INTRODUCTION TO MULTIGRID METHODS

LONG CHEN

We give a short introduction to multigrid methods for solving the linear algebraic equation that comes from the discretization of the Poisson equation in one dimension.

Multigrid is among the most efficient iterative solvers. It can solve a linear system of size N with optimal complexity $\mathcal{O}(N)$. To see the large saving of an $\mathcal{O}(N)$ method compared with an $\mathcal{O}(N^2)$ one, consider the following estimate. Suppose $N = 10^6$ and a standard PC can sum 10^6 numbers in one second. Then an $\mathcal{O}(N)$ algorithm will finish in a few seconds, while an $\mathcal{O}(N^2)$ algorithm will take nearly 11.6 days.

1. FREQUENCY AND ALIASING

We consider the discretization of the Poisson equation with homogeneous Dirichlet boundary conditions in one dimension:

$$(1) \quad -u'' = f, \quad x \in (0, 1), \quad u(0) = u(1) = 0.$$

For a positive integer N , choose a uniform grid, denoted by \mathcal{T}_h , of the interval $[0, 1]$ as follows:

$$0 = x_0 < x_1 < \dots < x_N < x_{N+1} = 1, \quad x_i = ih, \quad i = 0 : N + 1,$$

where $h = 1/(N + 1)$ is the length of each subinterval. For finite difference methods, we consider the vector $\mathbf{v} = (v_1, \dots, v_N)^T$, where $v_i \approx v(x_i)$ for $i = 1 : N$. The boundary nodes ($i = 0, N + 1$) are excluded due to the homogeneous Dirichlet boundary condition. We can connect the function values at the grid points to form a piecewise linear function $v = \sum_{i=1}^N v_i \phi_i$, where ϕ_i is the hat basis function at x_i for $i = 1 : N$.

The algebraic system for a scaled finite element or finite difference discretization is

$$(2) \quad \mathbf{A}\mathbf{u} = \mathbf{b},$$

where

$$\mathbf{A} = \text{diag}(-1, 2, -1), \quad \mathbf{b} = (b_i), \quad b_i = h^2 f(x_i).$$

Due to the special structure of the matrix \mathbf{A} , we can write its eigenvalues and eigenvectors explicitly. Recall that the eigenpairs of the continuous problem

$$-u'' = \lambda u \quad \text{in } (0, 1), \quad u(0) = u(1) = 0$$

are

$$\lambda_k = k^2, \quad \xi^k(x) = \sin(k\pi x), \quad k = 1, 2, 3, \dots$$

At the continuous level, $\lambda_{\min} = 1$ and $\lambda_k \rightarrow \infty$ as $k \rightarrow \infty$. At the discrete level, the $N \times N$ matrix \mathbf{A} has only N eigenvalues, and therefore eigenpairs with $k > N$ cannot be represented. After the h^2 scaling, we have $\lambda_{\min}(\mathbf{A}) \rightarrow 0$ and $\lambda_{\max}(\mathbf{A}) \leq 4$.

To better describe the eigenvalues of \mathbf{A} , we introduce a uniform grid in the angle interval $(0, \pi)$:

$$0 = \theta_0 < \theta_1 < \dots < \theta_N < \theta_{N+1} = \pi, \quad \theta_k = k\pi h, \quad k = 0 : N + 1.$$

The discrete eigenvectors are obtained by applying the nodal interpolation I_h to the continuous eigenfunctions on the grid $\{x_i\}$, and the corresponding discrete eigenvalues are functions of the angles $\{\theta_k\}$. For a continuous function u , we define $I_h u(x_i) = u(x_i)$ for $i = 0, \dots, N + 1$, so $I_h u$ is the discrete function that matches u at the grid points.




FIGURE 1. Eigenvalues of matrix \mathbf{A} with $N = 7$.

Proposition 1.1. *If $\mathbf{A} = \text{diag}(b, a, b)$ is an $N \times N$ tridiagonal matrix, then its eigenvalues are, for $k = 1, \dots, N$,*

$$\lambda_k = a + 2b \cos \theta_k,$$

and the corresponding eigenvectors are

$$\xi^k = I_h \sin(k\pi x).$$

Proof. This can be verified directly by using the trigonometric identity

$$\sin((i-1)\theta) + \sin((i+1)\theta) = 2 \cos(\theta) \sin(i\theta).$$

□

We use the index $i = 0 : N + 1$ for the uniform discretization of $\Omega = (0, 1)$ with mesh size h , so that $x_i = ih$. Similarly, we use the index $k = 0 : N + 1$ for the uniform discretization of the angle interval $[0, \pi]$ with step size πh , that is, $\theta_k = k\pi h$. For each k , the eigenvector ξ^k can be viewed as a discrete function whose nodal values are given by

$$\xi^k(x_i) = \sin(k\pi x_i) = \sin(k\pi h i) = \sin(i\theta_k), \quad i = 0, \dots, N + 1.$$

It is interesting to note that the eigenvectors are independent of a and b . In the special case $a = 2$ and $b = -1$, we obtain

$$(3) \quad \lambda_k(\mathbf{A}) = 2(1 - \cos \theta_k) = 4 \sin^2 \frac{\theta_k}{2}.$$

Notice that $\lambda_1 = \mathcal{O}(h^2)$ and $\lambda_N = \mathcal{O}(1)$ (see Fig. 1), and therefore $\kappa(\mathbf{A}) = \mathcal{O}(h^{-2})$, i.e., the matrix \mathbf{A} is ill-conditioned. For the finite difference method, the corresponding matrix is \mathbf{A}/h^2 , and for the finite element method it is \mathbf{A}/h . Such scalings multiply all eigenvalues by the same positive factor and thus do not change the condition number.

Exercise 1.2. Apply Proposition 1.1 to the mass matrix $M = (m_{ij})$ with $m_{ij} = \int_0^1 \phi_i \phi_j \, dx$ and conclude that M is well-conditioned. □

Exercise 1.3. Determine the eigenvalues and eigenvectors for the Neumann problem $-u'' = \lambda u$ in $(0, 1)$ with $u'(0) = u'(1) = 0$. □

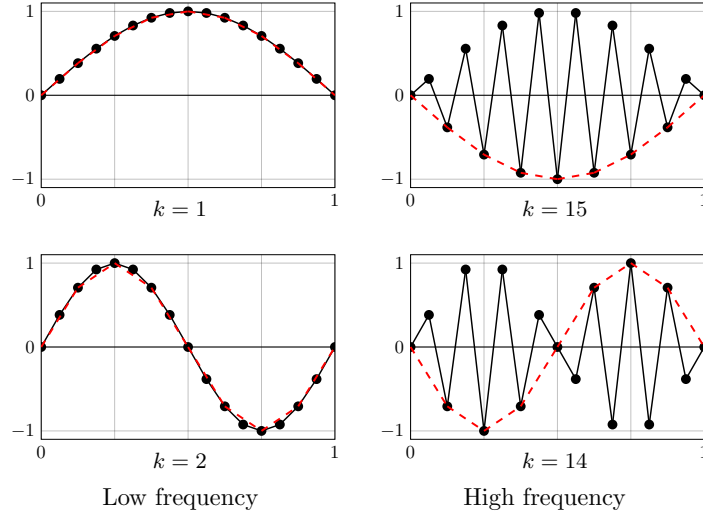


FIGURE 2. Selected eigenvectors of A for $h = 1/16$. Left: low-frequency modes that can be preserved on a coarse grid \mathcal{T}_{2h} . Right: high-frequency modes that will alias with low-frequency ones on the coarse grid \mathcal{T}_{2h} .

The integer k in the function $\sin(k\pi x)$ is called the *frequency*. For a uniform grid of $[0, 1]$ with mesh size $h = 1/(N + 1)$, the range of k is $1, \dots, N$. At the continuous level, there are infinitely many frequencies associated with all eigenfunctions, but frequencies higher than N cannot be represented on the grid \mathcal{T}_h . This limitation is known as *aliasing*.

To illustrate the aliasing phenomenon, consider a coarse grid with mesh size $2h$. When restricted to the coarse grid \mathcal{T}_{2h} , certain frequencies in the fine grid \mathcal{T}_h become indistinguishable from others. This can be verified by connecting coarse-grid values in Fig. 1. Specifically, for $k = 1 : N$, $k' = N + 1 - k$, so that $\theta_k + \theta_{k'} = \pi$, then

$$I_{2h}(\xi^k) = I_{2h}(-\xi^{k'}).$$

Indeed, this can be verified as follows:

$$\begin{aligned} I_{2h} \sin(k\pi x) &= (\sin(k\pi i 2h))_{i=1}^{N_c} = (\sin(i 2\theta_k))_{i=1}^{N_c}, \\ I_{2h} \sin(k'\pi x) &= (\sin(k'\pi i 2h))_{i=1}^{N_c} = (\sin(i 2\theta_{k'}))_{i=1}^{N_c} = -(\sin(i 2\theta_k))_{i=1}^{N_c}, \end{aligned}$$

using $\theta_k + \theta_{k'} = \pi$. Note that the eigenvalues on the coarse grid still range within $(0, 4)$, but the highest resolvable frequency of the eigenvectors is restricted by the mesh size.

In general, the grid \mathcal{T}_h cannot see frequencies higher than $1/h$. Therefore, the coarser grid \mathcal{T}_{2h} can only resolve frequencies less than $1/(2h)$. The frequencies that can be captured by the fine grid \mathcal{T}_h but not by the coarse grid \mathcal{T}_{2h} lie in the range

$$\frac{1}{2}(N + 1) \leq k \leq N,$$

and will be referred to as the *high-frequency* components (relative to \mathcal{T}_{2h}). We shall show that classical iterative methods quickly smooth out the high-frequency part of the error, while the low-frequency part decays much more slowly.

2. SMOOTHING PROPERTY OF THE RICHARDSON ITERATION

As an illustrative example, we apply the simplest iterative method, the Richardson iteration, to solve (2). Recall that one iteration of the Richardson method is

$$\mathbf{u}_{m+1} = \mathbf{u}_m + \alpha(\mathbf{b} - \mathbf{A}\mathbf{u}_m).$$

The corresponding error equation is

$$\mathbf{u} - \mathbf{u}_{m+1} = (I - \alpha\mathbf{A})(\mathbf{u} - \mathbf{u}_m),$$

and therefore,

$$\|\mathbf{u} - \mathbf{u}_{m+1}\|_{\mathbf{A}} \leq \rho(I - \alpha\mathbf{A}) \|\mathbf{u} - \mathbf{u}_m\|_{\mathbf{A}}.$$

When $\alpha \in (0, 2/\lambda_{\max}(\mathbf{A}))$, the iteration is convergent, and when $\alpha = 2/(\lambda_{\min}(\mathbf{A}) + \lambda_{\max}(\mathbf{A}))$, it achieves the optimal convergence rate

$$\rho = \frac{\kappa(\mathbf{A}) - 1}{\kappa(\mathbf{A}) + 1} \leq 1 - Ch^2.$$

This means that the total error, viewed as the sum of the squared contributions from all frequency components, decays at an overall rate of only $1 - Ch^2$. Here, we use m as the iteration index, since k is reserved for the frequency index.

We now apply a refined analysis to obtain different decay rates for different frequency components. First, we change the coordinate system. The set of eigenvectors of \mathbf{A} forms an orthogonal basis of \mathbb{R}^N with respect to both the Euclidean inner product (\cdot, \cdot) and the energy inner product $(\cdot, \cdot)_{\mathbf{A}}$. We expand the initial error $\mathbf{e}_0 = \mathbf{u} - \mathbf{u}_0$ as

$$\mathbf{e}_0 = \sum_{k=1}^N c_k \boldsymbol{\xi}^k.$$

After m Richardson iterations,

$$\mathbf{e}_m = \mathbf{u} - \mathbf{u}_m = (I - \alpha\mathbf{A})^m \mathbf{e}_0 = \sum_{k=1}^N c_{m,k} \boldsymbol{\xi}^k,$$

where

$$c_{m,k} = (1 - \alpha\lambda_k)^m c_k.$$

Since the eigenvectors are \mathbf{A} -orthogonal, we have

$$\|\mathbf{e}_m\|_{\mathbf{A}} = \left(\sum_{k=1}^N c_{m,k}^2 \|\boldsymbol{\xi}^k\|_{\mathbf{A}}^2 \right)^{1/2}, \quad \|\mathbf{e}_0\|_{\mathbf{A}} = \left(\sum_{k=1}^N c_k^2 \|\boldsymbol{\xi}^k\|_{\mathbf{A}}^2 \right)^{1/2}.$$

Therefore, the coefficient of the k -th frequency component decays at the rate

$$|c_{m,k}| \leq (1 - \alpha\lambda_k)^m |c_k|.$$

We choose $\alpha = 1/4$ to simplify the rate as

$$\rho_k = |1 - \alpha\lambda_k| = \frac{1}{2}(1 + \cos \theta_k).$$

From the graph of ρ_k (see Fig. 3(a)), it is easy to see that

$$\rho_1 \approx 1 - Ch^2,$$

while

$$\rho_N \leq Ch^2, \quad \rho_{(N+1)/2} = \frac{1}{2}.$$

This shows that *high-frequency components are damped very quickly*, a feature known as the smoothing property, whereas low-frequency components converge very slowly.




FIGURE 3. Contraction rate of the Richardson iteration for different parameters

The Gauss–Seidel method, as a more effective iterative scheme, shows the same behavior. In practice, it serves as a better smoother than the Richardson method. However, the analysis of its smoothing property is considerably more difficult.

Not all convergent iterative methods have the smoothing property. For example, the Jacobi method, which is equivalent to $\alpha = 1/2$ in the Richardson method, does not have a smoothing effect; see Fig. 4(b). For the same example, the decay rate for $\alpha = 1/2$ is

$$\rho_k = |\cos \theta_k|, \quad \rho_1 = \rho_N \approx 1 - Ch^2.$$

As a smoother, the weighted Jacobi method αD^{-1} (with recommend weight $\alpha = 2/3$) is more effective and is therefore used more often in practice.




FIGURE 4. Smoothing effect of classic iterative methods.

3. MULTIGRID METHODS

The multigrid method is built on two basic observations:

- High-frequency errors are effectively damped by a smoother.
- Low-frequency errors can be well approximated on a coarse grid.

A key idea is that part of the low-frequency error on a fine grid becomes a high-frequency error on a coarser grid. On the coarse grid, we can again apply smoothing and scale separation. A recursive use of this process across multiple grid levels gives the classical multigrid method.

We present below the recursive subroutine of a multigrid method.

```

1 function e = MG(r,J,mu)
2 % r: residual; J: level; mu: smoothing steps
3 if J == 1 % coarsest level: exact solve
4     e = A{J}\r;
5     return
6 end
7 e = 0;
8 % Presmoothing
9 for i = 1:mu
10     e = e + R(r-A{J}*e);
11 end
12 % Restriction
13 rc = Res(r-A{J}*e);
14 % Coarse grid correction
15 ec = MG(rc,J-1,mu);
16 if W-cycle
17     ec = ec + MG(rc-A{J}*ec,J-1,mu); % W-cycle
18 end
19 % Prolongation
20 e = e + Pro(ec);
21 % Postsmoothing
22 for i = 1:mu
23     e = e + R'(r-A{J}*e);
24 end

```

The operator (or subroutine) R used in the pre-smoothing step can be chosen as the Richardson, weighted Jacobi, or Gauss–Seidel iteration. In the post-smoothing step, we use its transpose R' to make the overall iteration symmetric. Two additional subroutines, Res and Pro , which connect the fine and coarse levels, will be discussed in detail in *Programming of Multigrid Methods*.

The function call $e = MG(r, *, *)$ indicates that an MG cycle is used to solve the residual equation $Ae = r$ and serves as an iterator in the residual-correction form:

$$u_{k+1} = u_k + MG(f - Au_k).$$

Because the iteration is linear, we can also write a direct-update version of the multigrid method as $u = MG(u, f, J, mu)$; see *Programming of Multigrid Methods*.

The above MG scheme is the *V-cycle*. If we apply one additional coarse-grid correction using MG (lines 16–18), we obtain the *W-cycle*. The two typical multigrid cycles are illustrated in Figure 5.

With one more coarse-grid correction, the W -cycle becomes more robust, but the computational work increases accordingly. For a hierarchy of nested finite element spaces $\{\mathbb{V}_i\}_{i=0}^J$ with $\dim \mathbb{V}_i = N_i$ and refinement ratio $N_i = \gamma N_{i-1}$ for some $\gamma > 1$, the total work of a multigrid cycle is obtained by adding the work over all levels. A single V -cycle performs one descent and one ascent through the levels, so

$$T_V = c \sum_{i=0}^J N_i \leq \frac{c}{1 - \gamma^{-1}} N_J,$$

based on the geometric relation $N_i = \gamma^{i-J} N_J$. Thus the V -cycle has complexity $\mathcal{O}(N_J)$ for any $\gamma > 1$.

A W -cycle recursively visits each coarse level twice, which leads to

$$T_W(i) = cN_i + 2T_W(i-1),$$




FIGURE 5. Multigrid Cycles. Left: V-cycle. Right: W-cycle

and the closed form

$$T_W = cN_J \sum_{k=0}^J (2/\gamma)^k.$$

If $\gamma > 2$ (for example, $\gamma \approx 2^d$ when $d \geq 2$), the geometric series stays bounded, and $T_W = \mathcal{O}(N_J)$. When $\gamma = 2$, the series grows like $J \sim \log N_J$, giving

$$T_W = \mathcal{O}(N_J \log N_J).$$

Hence, for a W-cycle, it is preferable that the coarsening ratio $\gamma = N_i/N_{i-1}$ satisfies $\gamma \geq 2$.

4. INTERPRETATION AS SUCCESSIVE SUBSPACE CORRECTION METHOD

We can interpret the multigrid method as a special case of the successive subspace correction method by choosing a multilevel decomposition of the space:

$$\mathbb{V}_1 \subset \mathbb{V}_2 \subset \cdots \subset \mathbb{V}_J = \mathbb{V},$$

together with a space decomposition

$$(4) \quad \mathbb{V} = \sum_{i=1}^J \mathbb{V}_i.$$

We introduce the following operators:

- $I_i : \mathbb{V}_i \hookrightarrow \mathbb{V}$, the natural inclusion;
- $Q_i : \mathbb{V} \mapsto \mathbb{V}_i$, the projection in the inner product (\cdot, \cdot) ;
- $A_i : \mathbb{V}_i \mapsto \mathbb{V}_i$, the restriction of A to the subspace $\mathbb{V}_i \times \mathbb{V}_i$;
- $R_i : \mathbb{V}_i \mapsto \mathbb{V}_i$, an approximation of A_i^{-1} , often called a *smoother*.

By definition,

$$(Q_i^\top v_i, v) = (v_i, Q_i v) = (v_i, v) = (I_i v_i, v), \quad \forall v_i \in \mathbb{V}_i, v \in \mathbb{V},$$

therefore Q_i^\top coincides with the natural inclusion I_i , or equivalently, $I_i^\top = Q_i$. Hence, the restriction of A to subspaces can be written as

$$A_i = I_i^\top A I_i = Q_i A Q_i^\top.$$

Given the current iterate u_k , one step of the successive (multiplicative) subspace correction (SSC) method is defined as follows: for $i = J, \dots, 1$, we approximately solve

$$A_i e_i = r_i, \quad \text{with } r_i = Q_i(f - A v^i),$$

and use the smoother $e_i = R_i r_i$.

Given the current iterate u_k , one step of the successive (multiplicative) subspace correction (SSC) method can be carried out without returning to the finest space at each subspace correction when the subspaces are nested.

Suppose the residual in \mathbb{V}_i is $r_i = Q_i(r - A e_{\text{old}})$, and the correction e_i is used to update $e_{\text{new}} = e_{\text{old}} + e_i$. We can then compute r_{i-1} recursively:

$$\begin{aligned} r_{i-1} &= Q_{i-1}(r - A e_{\text{new}}) \\ &= Q_{i-1} Q_i(r - A e_{\text{old}} - A e_i) \\ &= Q_{i-1}(r_i - Q_i A Q_i^\top e_i) \\ &= Q_{i-1}(r_i - A_i e_i). \end{aligned}$$

In the second line, we used the nestedness $\mathbb{V}_{i-1} \subset \mathbb{V}_i$, which gives $Q_{i-1} = Q_{i-1} Q_i$.

The correction can also be updated recursively:

$$e_i = e_i + I_{i-1} e_{i-1}, \quad i = 2, \dots, J.$$

Thus only the prolongation and restriction operators between consecutive levels are required. The work on each level is $\mathcal{O}(N_i)$, and the total cost is $\mathcal{O}(N)$ since $\sum_i N_i = \mathcal{O}(N)$.

From this viewpoint, applying the SSC method to a nested space decomposition naturally produces a V-cycle multigrid method. This interpretation not only clarifies the convergence mechanism but also leads to the following non-recursive implementation of the multigrid V-cycle.

```

1 function e = Vcycle(r, J)
2 ri = cell(J, 1); ei = cell(J, 1);
3 ri{J} = r;
4 for i = J:-1:2
5     ei{i} = R{i}*ri{i}; % pre-smoothing: one step
6     ri{i-1} = Res{i-1}*(ri{i}-Ai{i}*ei{i}); % update and restrict residual
7 end
8 ei{1} = Ai{1}\ri{1}; % exact solver in the coarsest level
9 for i = 2:J
10    ei{i} = ei{i} + Pro{i}*ei{i-1}; % prolongation and correction
11    ei{i} = ei{i} + R{i}'*(ri{i}-Ai{i}*ei{i}); % post-smoothing: one step
12 end
13 e = ei{J};

```

We refer to *Subspace Correction Method and Auxiliary Space Method* for a detailed explanation and convergence analysis of subspace correction methods, and to *Programming of Multigrid Methods* for implementation details.