Journal of Computational Physics 274 (2014) 695–708

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

Acceleration of the Jacobi iterative method by factors exceeding 100 using scheduled relaxation

Xiyang I.A. Yang, Rajat Mittal *

Department of Mechanical Engineering, Johns Hopkins University, MD 21218, United States

ARTICLE INFO

Article history:
Received 12 November 2013
Received in revised form 5 June 2014
Accepted 8 June 2014
Available online 27 June 2014

Keywords:
Iterative method
Jacobi method
Elliptic equations

ABSTRACT

We present a methodology that accelerates the classical Jacobi iterative method by factors exceeding 100 when applied to the finite-difference approximation of elliptic equations on large grids. The method is based on a schedule of over- and under-relaxations that preserves the essential simplicity of the Jacobi method. Mathematical conditions that maximize the convergence rate are derived and optimal schemes identified. The convergence rate predicted from the analysis is validated via numerical experiments. The substantial acceleration of the Jacobi method enabled by the current method has the potential to significantly accelerate large-scale simulations in computational mechanics, as well as other areas where elliptic equations are prominent.

1. Introduction

Elliptic equations appear routinely in computational fluid and solid mechanics as well as in heat transfer, electrostatics and wave propagation. Discretization of elliptic partial differential equations using finite-difference or other methods leads to a system of linear algebraic equations of the form $Au = b$, where u is the variable, b the source term, and A a banded matrix that represents the coupling between the variables. In the context of computational fluid mechanics, which is of particular interest to us, the Poisson equation for pressure appears in the majority of incompressible Navier–Stokes solvers [1,2], and is by far, the most computationally intensive component of such simulations. Thus, any effective methods that can accelerate the numerical solution of such equations would have a significant impact on computational mechanics and numerical methods.

The early history of iterative methods for matrix equations goes back to Jacobi [3] and Gauss [4], and the first application of such methods to a finite-difference approximation of an elliptic equation was by Richardson [5]. The method of Richardson, which can be expressed as $u^{n+1} = u^n - \omega_n(Au^n - b)$, where n and ω are the iteration index and the relaxation factor respectively, was a significant advance since it introduced the concept of convergence acceleration through successive relaxation. Richardson further noted that ω could be chosen to successively eliminate individual components of the residual. However, this required knowledge of the full eigenvalue spectrum of A, which was impractical. Given this, Richardson's recipe for choosing ω was to distribute the "nodes" (or zeros) of the amplification factor evenly within the range of eigenvalues of A. This was expected to drive down the overall amplification factor for the iterative scheme, and the advantage of this approach was that it required knowledge of only the smallest and the largest eigenvalues of A.

* Corresponding author.
E-mail addresses: xyang44@jhu.edu (X.I.A. Yang), mittal@jhu.edu (R. Mittal).

http://dx.doi.org/10.1016/j.jcp.2014.06.010
0021-9991/

Richardson's method subsequently appeared in the seminal doctoral dissertation of Young [6]. Noting, however, that "*... it appears doubtful that a gain of a factor of greater than 5 in the rate of convergence can in general be realized unless one is extremely fortunate in the choice of the values of* ω", Young discarded this method in favor of successive over-relaxation of the so-called Liebmann Method [7], which was essentially the same as the Gauss–Seidel method. This seems to signal the end of any attempts to accelerate Jacobi or Jacobi-like methods for matrix equations resulting from discrete approximations of elliptic equations.

In this article we describe a new approach for accelerating the convergence of the Jacobi iterative method as applied to the finite-difference approximation of elliptic equations. Using this approach, gains in convergence-rate well in excess of a factor of 100 are demonstrated for problem sizes of practical relevance. The increase of processor count in parallel computers into the tens of thousands that is becoming possible with multi-core and GPU architectures [8], is leading to an ever-increasing premium on parallelizability and scalability of numerical algorithms. While sophisticated iterative methods such as multigrid (MG) are highly efficient on a single processor [9], it is extremely difficult to maintain the convergence properties of these methods in large-scale parallel implementations. The domain decomposition approaches associated with parallel implementations negatively impact the smoothing properties of the iterative solvers used in MG, and also limit the depth of coarsening in such methods; both of these can significantly deteriorate the convergence properties of MG methods. In addition to this, the ratio of computation to communication also decreases for the coarse grid corrections, and this further limits the scalability of these methods. Within this context, the iterative method described here, preserves the insensitivity of the Jacobi method to domain decomposition, while providing significant convergence acceleration.

Another class of methods that is extensively used for solving elliptic equations is conjugate gradient (CG) [10]. CG methods, however, require effective preconditioners in order to produce high convergence rates; in this context, the method proposed here could eventually be adapted as a preconditioner for CG methods. Thus, the method described here could be used as an alternate to or in conjunction with these methods, and as such, could have a significant impact in computational mechanics as well as other fields such as weather and climate modeling, astrophysics and electrostatics, where elliptic equations are prominent.

Finally, the slow convergence rate of the Jacobi iterative method and the inability to accelerate this method using relaxation techniques is, at this point, considered textbook material [10–12]. In most texts, a discussion of the Jacobi method and its slow convergence is followed immediately by a discussion of the Gauss–Seidel method as a faster and more practical method. In this context, the method described here demonstrates that it is in-fact, relatively easy to increase the convergence rate of the Jacobi method by factors exceeding those of the classical Gauss–Seidel method. It is therefore expected that the method presented here will have a fundamental impact on our view of these methods, and spur further analysis of the acceleration of these basic methods.

2. Jacobi with successive over-relaxation (SOR)

We employ a 2D Laplace equation in a rectangular domain of unit size as our model problem: $\partial^2 u/\partial x^2 + \partial^2 u/\partial y^2 = 0$. A 2nd-order central-difference discretization on a uniform grid followed by the application of the Jacobi iterative method with a relaxation parameter ω, leads to the following iterative scheme:

$$u_{i,j}^{n+1} = (1-\omega)u_{i,j}^{n} + \frac{\omega}{4}\left(u_{i,j-1}^{n} + u_{i,j+1}^{n} + u_{i-1,j}^{n} + u_{i+1,j}^{n}\right) \tag{1}$$

where n is the index of iteration. von Neumann analysis [11] of the above scheme results in the following amplification factor:

$$G_\omega(\kappa) = (1-\omega\kappa) \quad \text{with } \kappa(k_x,k_y) = \sin^2(k_x\Delta x/2) + \sin^2(k_y\Delta y/2) \tag{2}$$

where Δx and Δy are the grid spacings and k_x and k_y the wave-numbers in the corresponding directions. The largest value of κ is given by $\kappa_{\max} = 2$. The smallest non-zero value of κ depends on the largest independent sinusoidal wave that the system can admit. A Neumann (N) problem allows waves to be purely one-dimensional, i.e. elementary waves can have $k_x = 0$ or $k_y = 0$, whereas for Dirichlet (D) problems, one-dimensional waves are, in general, not admissible, and k_x, k_y must all be non-zero. Therefore the corresponding $\kappa_{\min}$ are given by:

$$\kappa_{\min}^{\mathrm{N}} = \sin^2\left(\frac{\pi/2}{\max(N_x,N_y)}\right); \qquad \kappa_{\min}^{\mathrm{D}} = \sin^2\left(\frac{\pi/2}{N_x}\right) + \sin^2\left(\frac{\pi/2}{N_y}\right) \tag{3}$$

The above expressions are true for a uniform mesh and Section 8 describes the extension of this approach to non-uniform meshes.

The convergence of the iterative scheme requires $|G| < 1$ for all wave numbers, and it is easy to see from Eq. (2) that over-relaxation of the Jacobi method violates this requirement. Furthermore, for a given grid, $\kappa_{\min}^{\mathrm{N}} < \kappa_{\min}^{\mathrm{D}}$; thus Neumann problems have a wider spectrum and are therefore more challenging than the corresponding Dirichlet problem. We therefore focus most of our analysis on the Neumann problem. We also note that while the above analysis is for 2D problems, corresponding 1D and 3D problems lead to exactly the same expressions for the amplification factors, and similar expressions for $\kappa_{\min}$, with a pre-factor different from unity.