Journal of Computational Physics 321 (2016) 369–413

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

Scheduled Relaxation Jacobi method: Improvements and applications

J.E. Adsuara [a], I. Cordero-Carrión [b], P. Cerdá-Durán [a], M.A. Aloy [a,*]

[a] Departamento de Astronomía y Astrofísica, Universidad de Valencia, E-46100, Burjassot, Spain
[b] Departamento de Matemática Aplicada, Universidad de Valencia, E-46100, Burjassot, Spain

ARTICLE INFO

Article history:
Received 6 November 2015
Received in revised form 6 April 2016
Accepted 24 May 2016
Available online 27 May 2016

Keywords:
Iterative method
Jacobi method
Finite differences method
Elliptic equations

ABSTRACT

Elliptic partial differential equations (ePDEs) appear in a wide variety of areas of mathematics, physics and engineering. Typically, ePDEs must be solved numerically, which sets an ever growing demand for efficient and highly parallel algorithms to tackle their computational solution. The Scheduled Relaxation Jacobi (SRJ) is a promising class of methods, atypical for combining simplicity and efficiency, that has been recently introduced for solving linear Poisson-like ePDEs. The SRJ methodology relies on computing the appropriate parameters of a multilevel approach with the goal of minimizing the number of iterations needed to cut down the residuals below specified tolerances. The efficiency in the reduction of the residual increases with the number of levels employed in the algorithm. Applying the original methodology to compute the algorithm parameters with more than 5 levels notably hinders obtaining optimal SRJ schemes, as the mixed (non-linear) algebraic-differential system of equations from which they result becomes notably stiff. Here we present a new methodology for obtaining the parameters of SRJ schemes that overcomes the limitations of the original algorithm and provide parameters for SRJ schemes with up to 15 levels and resolutions of up to 2^{15} points per dimension, allowing for acceleration factors larger than several hundreds with respect to the Jacobi method for typical resolutions and, in some high resolution cases, close to 1000. Most of the success in finding SRJ optimal schemes with more than 10 levels is based on an analytic reduction of the complexity of the previously mentioned system of equations. Furthermore, we extend the original algorithm to apply it to certain systems of non-linear ePDEs.

1. Introduction

Partial differential equations (PDEs) are the appropriate mathematical language for modeling many phenomena [1]. In particular, we are interested in elliptic PDEs (ePDEs), that arise when we face the solution of equilibrium problems, in which the time evolution of the system is either neglected or irrelevant. Poisson and Laplace equations are prototype second order ePDEs, with and without source terms respectively.

Though the aforementioned Poisson and Laplace equations possess analytic solutions in a limited number of simple cases, we usually need a numerical solution when more general problems are considered. One of the standard approaches

* Corresponding author.
E-mail addresses: jose.adsuara@uv.es (J.E. Adsuara), isabel.cordero@uv.es (I. Cordero-Carrión), pablo.cerda@uv.es (P. Cerdá-Durán), miguel.a.aloy@uv.es (M.A. Aloy).

http://dx.doi.org/10.1016/j.jcp.2016.05.053

for solving these equations numerically is using finite differences methods. In this approach, both functions and operators are discretized on a numerical mesh, leading to a system of linear algebraic equations, which can be solved with direct or iterative methods. One of the simplest and most studied iterative schemes is the so called Jacobi method [2,3], whose main drawback is its poor convergence rate.

In oder to improve the efficiency of the Jacobi method, many alternatives have been considered. A popular possibility is the use of preconditioners [4–6] applied to linear systems, that make the associated Jacobi and Gauss–Seidel methods converge asymptotically faster than the unpreconditioned ones. Indeed, the method we improve on here, can be adapted as a preconditioner for other methods (e.g., the conjugate gradient method). Very widespread is the use of multigrid methods [e.g., 7] that, in many cases, provide the solution with $\mathcal{O}(N)$ operations, or that can even be employed as preconditioners. Relaxation algorithms [originally introduced in 3], improve the performance of the Jacobi method by considering modifications of the Gauss–Seidel algorithm that include a weight, for instance, successive overrelaxation (SOR) methods [8].

Along this line, [9, YM14 henceforth] has recently presented a significant acceleration (of the order of 100) over the Jacobi algorithm, employing the Scheduled Relaxation Jacobi (SRJ) method. The SRJ method is a generalization of the weighted Jacobi method which adds an overrelaxation factor to the classical Jacobi in a similar fashion to the SOR. This generalization includes a number P of different levels, in each of which, the overrelaxation (or underrelaxation) parameter or weight is tuned to achieve a significant reduction of the number of iterations, thus leading to a faster convergence rate. The optimal set of weights depends on the actual discretization of the problem at hand. Although the method greatly improves the convergence rate with respect to the original Jacobi, the schemes presented by YM14, optimal up to $P = 5$ and resolutions of up to 512 points per spatial dimension, are still not competitive with other methods used currently in the field (e.g., spectral methods [10], or multigrid methods as commented above). The main advantage of the SRJ method over other alternatives to solve numerically ePDEs is its simplicity and the straightforward parallelization, since SRJ methods preserve the insensitivity of the Jacobi method to domain decomposition (in contrast, e.g., to multigrid methods).

Following basically the same procedure as in YM14, [11, ACCA15 henceforth] has obtained optimal SRJ algorithms with up to $P = 10$ levels and multiple numerical resolutions. However, the limitations of the methodology of YM14 to compute optimal parameters for multilevel SRJ schemes prevent to develop algorithms with more than 10 levels. In this paper, we will show a new methodology to evaluate the parameters of optimal SRJ schemes with up to $P = 15$ levels and resolutions of up to 2^{15} points per spatial dimension of the problem, which in some cases may yield accelerations of order 10^3 with respect to the Jacobi method. For an straightforward use of the newly developed SRJ schemes, we provide the readers with a comprehensive set of tables for different SRJ schemes and different resolutions.

We begin the paper giving an overview of the SRJ method (Sect. 2) and describing the original methodology for obtaining optimal schemes together with the improvements on them already made in ACCA15 (Sect. 3). Then, we will present in Sect. 3.4, some analytical work which reduces the number of unknowns to solve for to $\mathcal{O}(P)$ (instead of $\mathcal{O}(P^2)$ as in YM14 and ACCA15). In Sect. 4 we show a comparison of the new method to compute optimal parameters for SRJ schemes with that of YM14.[1] Furthermore, we test the SRJ methods in a case study, namely a Poisson equation with Dirichlet boundary conditions (Sect. 5.1) that has analytic solution, and show that optimal SRJ parameters computed for resolutions close to that of the problem at hand can bring two orders of magnitude smaller number of iterations than the Jacobi method to solve such the problem. We have also assessed the performance of the new SRJ schemes with a large number of sublevels with respect to other standard methods to solve ePDEs (Sect. 5.2). In particular, we compare SRJ schemes with $P = 6$ and $P = 15$ to direct inversion methods and to spectral methods implemented in the `LAPACK` and `LORENE` packages, respectively. We outline the most prominent conclusions of our study and discuss the limitations of the current methodology in Sect. 6.

2. SRJ schemes

In this section we recap the most salient results obtained by YM14 and set the notation for the rest of the paper.

First of all, if we define $\omega_i J$ as a single step in a weighted Jacobi iteration using the weight ω_i $(i = 1, \ldots, P)$, then the SRJ method can be cast as a successive application of elementary relaxation steps of the form

$$\overbrace{\overbrace{\omega_1 J \ldots \omega_1 J}^{q_1}\underbrace{\omega_2 J \ldots \omega_2 J}_{q_2} \ldots \overbrace{\omega_P J \ldots \omega_P J}^{q_P}}^{M}\overbrace{\omega_1 J \ldots \omega_1 J}^{q_1} \ldots ,$$

where the largest weight, ω_1 is applied q_1 times, and each of the remaining and progressively smaller weights ω_i $(i = 1, \ldots, P)$ is applied q_i times, respectively. A single cycle of the scheme ends after M elementary steps, where $M := \sum_{i=1}^{P} q_i$. In order to reach a prescribed tolerance goal, we need to repeat a number times the basic M-cycle of the SRJ method. Both, a vector of weights and a vector with the number of times we use each weight, define each optimal scheme. We emphasize that, from the point of view of the implementation, the only difference with the traditional weighted Jacobi is that, instead of having a fixed weight, SRJ schemes of P-levels require the computation of P weights.

[1] ASCII files containing the optimal parameters for the SRJ algorithms shown in this paper with P values between 6 and 15 can be found at `http://www.uv.es/camap/SRJ.html`.