Journal of Computational and Applied Mathematics 283 (2015) 1–16

Contents lists available at ScienceDirect

Journal of Computational and Applied Mathematics

journal homepage: www.elsevier.com/locate/cam

A multi-level preconditioned Krylov method for the efficient solution of algebraic tomographic reconstruction problems

Siegfried Cools [a,*], Pieter Ghysels [b], Wim van Aarle [c], Jan Sijbers [c], Wim Vanroose [a]

[a] Applied Mathematics Group, University of Antwerp, Middelheimlaan 1, 2020 Antwerp, Belgium

[b] Future Technologies Group, Computational Research Division, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA

[c] iMinds-Vision Lab, University of Antwerp, Universiteitsplein 1, 2610 Wilrijk, Belgium

ARTICLE INFO

Article history:
Received 20 March 2014
Received in revised form 13 November 2014

Keywords:
Tomography
Algebraic reconstruction
Krylov methods
Preconditioning
Multigrid
Wavelets

ABSTRACT

Classical iterative methods for tomographic reconstruction include the class of Algebraic Reconstruction Techniques (ART). Convergence of these stationary linear iterative methods is however notably slow. In this paper we propose the use of Krylov solvers for tomographic linear inversion problems. These advanced iterative methods feature fast convergence at the expense of a higher computational cost per iteration, causing them to be generally uncompetitive without the inclusion of a suitable preconditioner. Combining elements from standard multigrid (MG) solvers and the theory of wavelets, a novel wavelet-based multi-level (WMG) preconditioner is introduced, which is shown to significantly speed-up Krylov convergence. The performance of the WMG-preconditioned Krylov method is analyzed through a spectral analysis, and the approach is compared to existing methods like the classical Simultaneous Iterative Reconstruction Technique (SIRT) and unpreconditioned Krylov methods on a 2D tomographic benchmark problem. Numerical experiments are promising, showing the method to be competitive with the classical Algebraic Reconstruction Techniques in terms of convergence speed and overall performance (CPU time) as well as precision of the reconstruction.

1. Introduction

Computed Tomography (CT) is a powerful imaging technique that allows non-destructive visualization of the interior of physical objects. Besides its common use in medical applications [1], tomography is also widely applicable in fields such as biomedical research, materials science, and metrology. In all applications, a certain imaging source (e.g. an X-ray source) and an imaging detector (e.g. X-ray detector) are used to acquire two-dimensional projection images of the object from different directions. A three-dimensional virtual reconstruction can then be computed using one of the many reconstruction techniques that can be found in the literature. In practice, the most commonly used analytical methods for CT are Filtered Backprojection (FBP) and its cone-beam variant Feldkamp–Davis–Kress (FDK). These methods make use of various analytical properties of the projection geometries to compute the reconstructed object at a low computational cost. A major drawback of analytical methods is their inflexibility to different experimental setups and their inability to include reconstruction constraints which can be used to exploit possible prior information about the object.

* Corresponding author.
E-mail address: siegfried.cools@uantwerp.be (S. Cools).

http://dx.doi.org/10.1016/j.cam.2014.12.044

Iterative Algebraic Reconstruction Techniques (ART) form an interesting alternative to the aforementioned analytical methods. Here, the reconstruction problem is described as the solving of a system of linear equations. The Simultaneous Iterative Reconstruction Technique (SIRT) is a straightforward method that has been extensively studied in the literature, see [2] and the references therein. Another general class of algebraic solution methods are the Krylov solvers such as CGLS, GMRES, and BiCGStab, an overview of which can be found in [3]. Alternatively, one can resort to more powerful techniques that apply additional constraints to the reconstruction, which can lead to improved accuracy, especially when fewer projection images are available (i.e. scans with a lower radiation dose). Total variation minimization approaches such as FISTA [4], for example, assume that the variation between neighboring pixels is low inside a homogeneous object. Discrete tomography approaches such as DART [5] improve the reconstruction quality by limiting the number of gray level values that can be present in the reconstructed image.

While iterative methods for tomography have become widely accepted in the scientific community, practical applications have not yet adopted these techniques [6], mostly due to the variable computational cost and storage requirements of the iterative process (contrary to the fixed costs of analytical methods based on FFT-type algorithms). The development of efficient new iterative solvers is therefore crucial. This efficiency can be accomplished in two ways. Firstly, the computation time of each iteration can be reduced by optimally exploiting parallelism of the projection and backprojection operators with the use of modern hardware accelerated computer architectures such as NVIDIA GPU's [7] or the Intel Xeon Phi [8]. Secondly, a solver with a fast convergence rate, requiring only a limited number of iterations should be used. Additionally, the convergence rate of the ideal solver should not depend on the problem size.

In this work, an approach that fits into the second category will be introduced for non-constrained iterative reconstruction. By analyzing the spectral properties of the standard SIRT method, it will be shown that the convergence of classical algebraic reconstruction techniques (stationary iterative schemes) is notably slow. As it appears, the alleged smoothing property does not hold in the case of tomographic reconstruction problems. Krylov methods prove to be more efficient, yet are generally more expensive in terms of memory and computation cost. Therefore, when using Krylov methods, it is mandatory to define an efficient preconditioner, which allows faster convergence. This approach is very common in a wide range of PDE-type problems, yet is still fairly new for tomographic reconstruction. Related work in the setting of tomographic reconstruction includes the research on multilevel image reconstruction by McCormick et al. [9,10], and more recently the work done on multigrid methods for tomographic reconstruction by Webb et al. [11] and Rüde et al. [12].

Originally introduced as a theoretical tool by Fedorenko in 1964 [13] and later adopted as a solution method by Brandt in 1977 [14], multigrid (MG) solvers are commonly used as efficient and low-cost Krylov preconditioners for high-dimensional problems in the PDE literature, see e.g. [15,16]. One of the key concepts of the multigrid scheme is the representation of the original fine grid reconstruction problem on a coarser scale resolution, where the problem is computationally cheaper to solve. However, we show that the standard multigrid approach [17–20] does not act as an efficient preconditioner for algebraic tomographic systems. Indeed, the ineffectiveness of the smoother in eliminating the oscillatory modes causes the key complementary action of smoother and coarse grid correction to fail, resulting in an inefficient multigrid scheme for algebraic tomographic reconstruction problems.

In this work a new wavelet-based multigrid (WMG) preconditioner is introduced, which is more suited for tomographic reconstruction. The proposed method combines elements from standard multigrid with the theory of wavelets, and shows some similarities to the work on wavelet-based multiresolution tomographic reconstruction in [21,22]. Additionally, the main advantage of the proposed method, i.e. projection of the large fine-scale system onto smaller, easy-to-solve subproblems, resembles key features of the Hierarchical Basis Multigrid Method (HBMM) [23,24]. It is shown through an eigenvalue analysis that WMG-preconditioning significantly increases Krylov convergence speed, which is confirmed by various numerical experiments. Additionally, we show that the WMG-preconditioned Krylov solver allows for an accuracy which is generally unobtainable by classical SIRT reconstruction. The numerical results presented in this work show promise, validating the proposed WMG scheme as an efficient Krylov preconditioning technique for algebraic tomographic reconstruction.

The paper is structured as follows. In Section 2 the classical SIRT and MG-Krylov solvers for iterative tomographic reconstruction are reviewed and analyzed. Section 3 introduces a novel preconditioning approach to account for the defects of the MG preconditioner, which greatly improves convergence speed of the BiCGStab Krylov solver. In Section 4, a series of experimental simulations is presented to validate our contribution. Ultimately, Section 5 concludes this work with an overview of the main results in this paper and a discussion on possible future research options.

2. Notation and key concepts of tomographic reconstruction

2.1. Algebraic tomographic reconstruction

Consider a data vector $b \in \mathbb{R}^M$, with $M = m \times n$, where m is the number of projection angles and n is the number of beams. We assume that the number of pixels in every spatial dimension equals n, such that the data is reconstructed on a 2D $n \times n$ grid. We denote the total number of pixels in the image by $N = n \times n$. Algebraic reconstruction methods consider tomographic reconstruction as the problem of solving the linear system of equations

$$Wx = b, \tag{1}$$