Journal of Computational Physics 304 (2016) 170–188

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

A fast block low-rank dense solver with applications to finite-element matrices

AmirHossein Aminfar [a,*,1], Sivaram Ambikasaran [b,2], Eric Darve [c,1]

[a] 496 Lomita Mall, Room 104, Stanford, CA, 94305, United States
[b] Warren Weaver Hall, Room-1105A, 251, Mercer Street, New York, NY 10012, United States
[c] 496 Lomita Mall, Room 209, Stanford, CA, 94305, United States

ARTICLE INFO

Article history:
Received 15 August 2014
Received in revised form 14 September 2015
Accepted 5 October 2015
Available online 3 November 2015

Keywords:
Fast direct solvers
Iterative solvers
Numerical linear algebra
Hierarchically off-diagonal low-rank matrices
Multifrontal elimination
Adaptive cross approximation

ABSTRACT

This article presents a fast solver for the dense “frontal” matrices that arise from the multifrontal sparse elimination process of 3D elliptic PDEs. The solver relies on the fact that these matrices can be efficiently represented as a hierarchically off-diagonal low-rank (HODLR) matrix. To construct the low-rank approximation of the off-diagonal blocks, we propose a new pseudo-skeleton scheme, the boundary distance low-rank approximation, that picks rows and columns based on the location of their corresponding vertices in the sparse matrix graph. We compare this new low-rank approximation method to the adaptive cross approximation (ACA) algorithm and show that it achieves better speedup specially for unstructured meshes. Using the HODLR direct solver as a preconditioner (with a low tolerance) to the GMRES iterative scheme, we can reach machine accuracy much faster than a conventional LU solver. Numerical benchmarks are provided for frontal matrices arising from 3D finite element problems corresponding to a wide range of applications.

1. Introduction

In many engineering applications, solving large finite element systems is of great significance. Consider the system

$$Ax = b$$

arising from the finite element discretization of an elliptic PDE, where $A \in \mathbb{R}^{N\times N}$ is a sparse matrix with a symmetric pattern. In many practical applications, the matrix A might be ill-conditioned and thus, challenging for iterative methods. On the other hand, conventional direct solver algorithms, while being robust in handling ill-conditioned matrices, are computationally expensive ($\mathcal{O}(N^{1.5})$ for 2D meshes and $\mathcal{O}(N^2)$ for 3D meshes). Since one of the main bottlenecks in the direct multifrontal solve process is the high computational cost of solving dense frontal matrices, we mainly focus on solving these matrices in this article. Our goal is to build an iterative solver, which utilizes a fast direct solver as a preconditioner for the dense frontal matrices. The direct solver in this scheme acts as a highly accurate pre-conditioner. This approach combines the advantages of the iterative and direct solve algorithms, i.e., it is fast, accurate and robust in handling ill-conditioned matrices.

* Corresponding author. Tel.: +1 650 644 7624.
E-mail address: aminfar@stanford.edu (A. Aminfar).
[1] Mechanical Engineering Department, Stanford University.
[2] Courant Institute of Mathematical Sciences, New York University.

http://dx.doi.org/10.1016/j.jcp.2015.10.012
0021-9991/

To be consistent with our previous work, we adopt the notation used in [3]. We should also mention that 'n' refers to the size of dense matrices and 'N' refers to the size of sparse matrices (e.g., number of degrees of freedom in a finite-element mesh).

In the next section, we review the previous literature on both dense structured solvers and sparse multifrontal solvers. We then introduce a hierarchical off-diagonal low-rank (from now on abbreviated as HODLR) direct solver in Section 4. In Section 5, we introduce the boundary distance low-rank (BDLR) algorithm as a robust low-rank approximation scheme for representing the off-diagonal blocks of the frontal matrices. Section 6 discusses the application of the iterative solver with a fast HODLR direct solver preconditioner to the sparse multifrontal solve process and demonstrates the solver for a variety of 3D meshes. We also show an application in combination with the FETI-DP method [21], which is a family of domain decomposition algorithms to accelerate finite-element analysis on parallel computers.

2. Previous work

2.1. Fast direct solvers for dense hierarchical matrices

Hierarchical matrices are data sparse representation of a certain class of dense matrices. This representation relies on the fact that these matrices can be sub-divided into a hierarchy of smaller block matrices, and certain sub-blocks (based on the admissibility criterion) can be efficiently represented as a low-rank matrix. We refer the readers to [29,33,27,30,12,15,13] for more details. These matrices were introduced in the context of integral equations [29,33,62,43] arising out of elliptic partial differential equations and potential theory. Subsequently, it has also been observed that dense fill-ins in finite element matrices [60], radial basis function interpolation [3], kernel density estimation in machine learning, covariance structure in statistic models [16], Bayesian inversion [3,5,6], Kalman filtering [45], and Gaussian processes [4], can also be efficiently represented as data-sparse hierarchical matrices. Broadly speaking, these matrices can be grouped into two general categories based on the admissibility criterion: (i) Strong admissibility: sub-blocks that correspond to the interaction between well-separated clusters are low-rank; (ii) Weak admissibility: sub-blocks corresponding to non-overlapping interactions are low-rank. Ambikasaran [1] provides a detailed description of these different hierarchical structures.

We review some of the previously developed structured dense solvers for hierarchical matrices and discuss them in relation to our work. Grasedyck and Hackbusch [29,27] introduced the concept of $\mathcal{H}$-matrices, which are the most general class of hierarchical matrices with the strong admissibility criterion [29,27,30,32,31,33,34,9,10,12]. Contrary to the HODLR matrix structure, in which the off-diagonal blocks are low-rank, in $\mathcal{H}$-matrices, the off-diagonal blocks are further decomposed into low-rank and full-rank blocks. Thus, the rank can be kept small. In HODLR, we make a single low-rank approximation for the off-diagonal blocks and the rank is larger as a result. Hence, the HODLR structure makes for a much simpler representation and is often used because of its simplicity compared to the $\mathcal{H}$-matrix structure. Grasedyck and Hackbusch [27] suggest a recursive block low-rank factorization scheme for $\mathcal{H}$-matrices. This method is based on the idea that all the dense matrix algebra (matrix multiplication and matrix addition) can be replaced by $\mathcal{H}$-matrix algebra. As a result, the inverse of an $\mathcal{H}$-matrix can also be approximated as an $\mathcal{H}$-matrix itself. This results in a computational complexity of $\mathcal{O}(n\log^2(n))$ for an $\mathcal{H}$-matrix factorization.

We note that the approach in this paper is based on the Woodbury matrix identity. It is therefore different from the algorithm in Grasedyck and Hackbusch [27] for example. The latter is based on a block LU factorization, while the Woodbury identity reduces the global solve to block diagonal solves followed by a correction update.

The HODLR matrix structure is the most general off-diagonal low-rank structure with weak admissibility. Solvers for this matrix class have a computational cost of $\mathcal{O}(n\log^2 n)$. In an HODLR matrix, the off-diagonal low-rank bases do not have a nested structure across different levels [3]. The HSS matrix is an HODLR matrix but, in addition, has a nested off-diagonal low-rank structure. Solvers for the HSS matrices have an $\mathcal{O}(n)$ complexity [61,14].

Martinsson and Rokhlin [51] discuss an $\mathcal{O}(n)$ direct solver for boundary integral equations based on the HSS structure. Their method is based on the fact that for a matrix of rank r, there exists a well-conditioned column operation, which leaves r columns unchanged and sets the remaining columns to zero. Using this idea, they derive a two-sided compressed factorization of the inverse of the HSS matrix. Their generic algorithm requires $\mathcal{O}(n^2)$ operations to construct the inverse. However, they accelerate their algorithm to $\mathcal{O}(n\log^{\kappa}(n))$ when applied to two-dimensional contour integral equations.

Chandrasekaran et al. [15] present a fast $\mathcal{O}(n)$ direct solver for HSS matrices. In their article, they construct an implicit ULV^H factorization of an HSS matrix, where U and V are unitary matrices, L is a lower triangular matrix and ^{H} is the transpose conjugate operator. Their method is based on the Woodbury matrix identity and the fact that for a low-rank representation of the form UBV^H, where U and V are thin matrices with r columns, there exists a unitary transformation Q, in which only the last r rows of QU are nonzero. They use this observation to recursively solve the linear system of equations. Since this method requires constructing an HSS tree, the authors suggest an algorithm that uses the SVD or the rank revealing QR decomposition, recursively, to construct the HSS tree in $\mathcal{O}(n^2)$ time.

Gillman et al. [25] discuss an $\mathcal{O}(n)$ algorithm for directly solving integral equations in one-dimensional domains. The algorithm relies on applying the Sherman–Morrison–Woodbury formula (see for example [3]) recursively to an HSS tree structure to achieve $\mathcal{O}(r^2n)$ solve complexity, where r is the rank of the off diagonal blocks in the HSS matrix. They also describe an $\mathcal{O}(r^2n)$ algorithm for constructing an HSS representation of the matrix resulting from a Nyström discretization of a boundary integral equation.