Engineering Analysis with Boundary Elements 66 (2016) 168–182

Contents lists available at ScienceDirect

Engineering Analysis with Boundary Elements

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

Acceleration of isogeometric boundary element analysis through a black-box fast multipole method

R.N. Simpson *, Z. Liu

School of Engineering, University of Glasgow, Glasgow G12 8QQ, UK

A R T I C L E I N F O

Article history:
Received 12 August 2015
Received in revised form
2 March 2016
Accepted 4 March 2016
Available online 14 March 2016

Keywords:
Isogeometric analysis
T-splines
Boundary element method
Black-box fast multipole method

A B S T R A C T

This work outlines the use of a black-box fast multipole method to accelerate the far-field computations in an isogeometric boundary element method. The present approach makes use of T-splines to discretise both the geometry and analysis fields allowing a direct integration of CAD and analysis technologies. A black-box fast multipole method of $O(N)$ complexity is adopted that minimises refactoring of existing boundary element codes and facilitates the use of different kernels. This paper outlines an algorithm for implementing the open-source black-box fast multipole method BBFMM3D[1] within an existing isogeometric boundary element solver, but the approach is general in nature and can be applied to any boundary element surface discretisation. The $O(N)$ behaviour of the approach is validated and compared against a standard direct solver. Finally, the ability to model large models of arbitrary geometric complexity directly from CAD models is demonstrated for potential problems.

1. Introduction

In the majority of modern industrial engineering and design workflows Computer Aided Design (CAD) and analysis software play a crucial role in reducing the overall design lifecycle. The iterative nature of design requires tight integration of CAD and analysis software, but modern workflows are inhibited by cumbersome fixing and defeaturing algorithms that must be used in the transition from CAD models to analysis models. The disparity between CAD and analysis is one of the biggest challenges facing engineering design which has inspired research into new discretisation approaches that unify or greatly ease the transition from CAD to analysis and vice versa.

One of the most active research areas that aims to address the disparity between CAD and analysis is the field of isogeometric analysis (IGA) [28] that uses spline-based discretisations generated by CAD software as a basis for analysis thus providing a framework that unifies CAD and analysis. Since the seminal paper of [28], the concept has expanded rapidly into several applications including acoustics [47,38], vibrations [15], elasticity [1,42], electromagnetics [49,10] and fluid flow [3,17,26]. Early work on IGA has focussed on the use of Non-Uniform Rational B-Splines (NURBS) [39] due to their popularity within modern commercial CAD software, but limitations stemming from their tensor-product nature have prompted research into alternative CAD discretisations including subdivision surfaces [13,14], PHT splines [35,50], LR B-splines [29], T-splines [2] and THCCS [52]. From a commercial perspective, the two technologies which have made the largest impact include subdivision surfaces and T-splines. Subdivision surfaces are ubiquitous within the computer animation industry but at present, they have yet to penetrate the CAD software market. T-splines offer a promising route to overcome the tensor product nature of NURBS while also providing backwards compatibility with existing NURBS technology. From an analysis perspective, T-splines have opened up interesting routes for integrated design and analysis technologies through properties such as water-tight geometries and local refinement algorithms. T-splines were first used in an analysis context in [2] and subsequently analysis-suitable T-splines [32,9] were proposed that satisfy important analysis properties while retaining flexible geometry and modification algorithms. Further research includes efficient evaluation of T-spline basis functions through Bézier extraction [43].

A popular approach in CAD is to represent geometry in terms of a surface or Boundary-Representation (B-Rep) through appropriate geometry discretisations such as connected NURBS patches, T-spline or subdivision surfaces. Such surface discretisations are insufficient for volumetric analysis methods such as the finite element method but provide the necessary data structures for analysis methods based on surfaces such as shell and boundary integral formulations. The limitations of boundary integral approaches are well-known, but assuming the use of such an approach is valid, they are found to be a particularly attractive

* Corresponding author.
E-mail address: robert.simpson.2@glasgow.ac.uk (R.N. Simpson).

[1] https://github.com/ruoxi-wang/BBFMM3D

http://dx.doi.org/10.1016/j.enganabound.2016.03.004
0955-7997/

approach for integrated design and analysis. By adopting a common discretisation for both geometry and analysis, isogeometric boundary element methods completely circumvent meshing procedures and eliminate geometry error promoting design software that truly integrates CAD and analysis. The idea has been explored in the context of several applications including elastostatics [46,45,51], shape optimisation [7,19,31] acoustics [47,37,38] and underground excavations [6].

A well-known feature of the BEM approach is the debilitating O(N^2) asymptotic behaviour for matrix assembly and O(N^3) behaviour of direct solvers that eventually dominates for large problems. For practical engineering problems this manifests itself as large runtimes and heavy memory demands that often completely prohibit the use of direct solvers. Instead, matrix compression techniques which reduce the overall solver complexity to O($N \log N$) or O(N) must be used. At present, the most popular techniques include: the Fast Multipole Method (FMM) [21,11,12,36] and Hierarchical (H-) matrices [24,22,23] which make use of low-rank compression methods such as Adaptive Cross Approximation (ACA) [5,30]. These techniques are all based on the same fundamental concept of approximating the smooth nature of the kernel for far-field interactions through efficient hierarchical data structures that allow for fast matrix–vector computations within an iterative solver. More recent research has focussed on the development of fast direct solvers (e.g. [20,8]) that have shown advantageous properties over iterative techniques and offer a promising direction for future BEM solvers.

From an implementation standpoint, preference is often given to ACA and H-matrix methods which perform matrix compression in a purely algebraic manner, in contrast to the majority of FMM implementations which require extensive changes to BEM software. However, there exist black-box FMM implementations that overcome these limitations [53,18,33] opening up efficient O(N) FMM algorithms to BEM software. The present paper is based on such techniques.

Previous work on accelerating isogeometric BEM computations includes FMM compression for 2D Laplace problems [48], H-matrices to accelerate 2D and 3D elasticity applications [34] and a comparative study of Wavelet, FMM and ACA compression defined over parametric surfaces [25]. All of these studies have made use of tensor product parameterisations in the form of NURBS or rational Bézier surfaces.

The present paper outlines an approach for accelerating BEM computations in the framework of isogeometric analysis by employing a black-box FMM and adopting T-splines to discretise both the surface geometry and analysis fields. A collocation approach is chosen in the present study, but the techniques are applicable also to Galerkin and Nyström methods. Through the use of a black box FMM algorithm, the changes required to any existing BEM code are kept to a minimum. The use of T-splines allows direct integration of computational geometry and analysis technology while overcoming the inherent limitations of tensor product surfaces. The combination of these technologies offers a significant step forward towards integrated design and analysis for industrial applications.

The paper is organised as follows: a brief overview of the black-box FMM algorithm is given highlighting common FMM terminology and its relation to traditional BEM notation; the boundary element discretisation procedure that allows a system of equations to be formed is stated; an overview of T-spline discretisation technology is described; the algorithm for computing the matrix–vector product through the black-box FMM for fast BEM solve times and reduced memory consumption is detailed and finally, numerical examples are given to verify the implementation and assess its asymptotic behaviour against a standard direct solver for potential problems. All algorithms and numerical examples in the present work are based on three-dimensional problems.

2. Fast multipole methods

Fast multipole methods were originally developed to overcome the intractable computational complexity of N-body problems when solved by direct means. Such problems can be expressed as

$$f(\mathbf{x}_i) = \sum_{j=1}^{N_s} K(\mathbf{x}_i, \mathbf{y}_j)\sigma_j, \quad i = 1, 2, \ldots, N_f \tag{1}$$

where $f(\mathbf{x}_i)$ is the desired force or field, $K(\mathbf{x}, \mathbf{y})$ is a problem specific kernel, $\{\sigma_j\}_{j=1}^{N_s}$ is a set of charges, $\{\mathbf{x}_i\}_{i=1}^{N_f}$ a set of field points and $\{\mathbf{y}_j\}_{j=1}^{N_s}$ a set of source points. In the case $N_s = N_f = N$ and (1) is applied directly, the computation time scales as O(N^2) which necessitates acceleration methods for large problems. Early work on the FMM applied to three-dimensional problems formulated methods that scale as O($N \log N$) with subsequent improvements in algorithms achieving scaling of O(N). Many variants of the FMM exist, but all are based on the same fundamental algorithm:

1. *Prescribed tolerance*: a tolerance ϵ is prescribed to determine the number of terms retained in far-field expansions.
2. *Subdivision of space*: a hierarchical subdivision of space is constructed consisting of m levels indexed by $k = 0, 1, \ldots, m$ (see Fig. 1 for an illustration of increasing levels of hierarchical subdivision). Octree subdivision is commonly used for three-dimensional problems. For each level of the octree, a set of cells $\{\mathcal{Y}_a^k\}_{a=1}^{n_{cell}^k}$ is defined. Source and field points are assigned to cells in every level.
3. *Upwards pass*: far-field expansions are computed for each cell at the lowest level m of the tree. Far field expansions for cells in level $m-1$ and higher are computed from expansions in lower levels through a Moment-to-Moment (M2M) translation operator.
4. *Downwards pass*: working down the tree, local expansions are formed for each cell $\mathcal{Y}_a^k$. These are calculated through a Moment-to-Local (M2L) operator for cells in the interaction list[2] of $\mathcal{Y}_a^k$ and a Local-to-Local (L2L) operator applied to the parent cell of $\mathcal{Y}_a^k$.
5. *Evaluation*: working at the lowest level of the tree, the FMM approximation of $f(\mathbf{x}_i)$ is computed by finding the cell at the lowest level of the tree which contains $\mathbf{x}_i$. The local expansion of this cell is used to compute the far-field approximation with the near-field computed directly by summing over all near-neighbours.

2.1. Black-box fast multipole method

In the case of the black-box algorithm of [18], far-field expansions are based on Chebyshev interpolation and M2L operators are constructed through reduced rank operators calculated by Singular Value Decomposition (SVD). A particularly beneficial feature of this approach is its ability to handle arbitrary kernels in contrast to conventional FMM implementations that often require significant code rewrites for alternative kernels. This justifies the use of such an approach in the present study.

To accelerate N-body computations using the black-box code of [18], the following specific inputs are required:

1. *Tolerance parameters*: consisting of the target precision ϵ used to compute SVD cutoff parameters and n_{ch}, the number of Chebyshev nodes used to interpolate in each coordinate direction.
2. *Hierarchical subdivision parameters*: comprising m, the number of levels in the tree hierarchy and L, the side-length of the smallest cube enclosing the domain.
3. *Kernel*, $K(x, y)$: prescribed either analytically or numerically.

[2] See [4] for a thorough definition of FMM terms including well-separated, near-neighbours and the interaction list of a cell.