Commun. Comput. Phys.
doi: 10.4208/cicp.OA-2016-0168

Vol. **22**, No. 4, pp. 997-1014
October 2017

GPU-Accelerated LOBPCG Method with Inexact Null-Space Filtering for Solving Generalized Eigenvalue Problems in Computational Electromagnetics Analysis with Higher-Order FEM

A. Dziekonski*, M. Rewienski, P. Sypek, A. Lamecki and M. Mrozowski

Department of Microwave and Antenna Engineering, Faculty of Electronics, Telecommunications and Informatics, Gdansk University of Technology, Gdansk 80-23, Poland;
CUDA Research Center for Computational Electromagnetics atGdansk University of Technology (2012-2016).

Received 29 September 2016; Accepted (in revised version) 7 March 2017

Abstract. This paper presents a GPU-accelerated implementation of the Locally Optimal Block Preconditioned Conjugate Gradient (LOBPCG) method with an inexact nullspace filtering approach to find eigenvalues in electromagnetics analysis with higher-order FEM. The performance of the proposed approach is verified using the Kepler (Tesla K40c) graphics accelerator, and is compared to the performance of the implementation based on functions from the Intel MKL on the Intel Xeon (E5-2680 v3, 12 threads) central processing unit (CPU) executed in parallel mode. Compared to the CPU reference implementation based on the Intel MKL functions, the proposed GPU-based LOBPCG method with inexact nullspace filtering allowed us to achieve up to 2.9-fold acceleration.

AMS subject classifications: 65N25, 68W10, 65Y05, 65F08, 65F10, 74S05

Key words: LOBPCG, inexact nullspace filtering, multilevel preconditioning, FEM, GPU, parallel computing.

1 Introduction

Eigenvalue problems derived from Maxwell's equations are an important class of problems in electromagnetic research, as the associated eigenvalues and eigenvectors provide key characteristics of the examined systems—including resonant frequencies for

*Corresponding author. *Email addresses:* adziek@eti.pg.gda.pl (A. Dziekonski), mrewiens@eti.pg.gda.pl (M. Rewienski), psypek@eti.pg.gda.pl (P. Sypek), adam.lamecki@ieee.org (A. Lamecki), m.mrozowski@ieee.org (M. Mrozowski)

resonator cavities that are essential for material characterization, particle accelerators or coupled resonator filter design, propagation coefficients for guided waves, poles or transmission zeros for filters, and so on. Let us consider a closed region Ω, forming a cavity in three dimensions. If the cavity is loaded with a medium whose properties are represented by permeability μ and permittivity ϵ, then Maxwell's equations can be transformed to the following Helmholtz wave partial differential equation (PDE):

$$\nabla\times(\mu^{-1}\nabla\times\mathbf{E})=\omega^2\epsilon\mathbf{E},\quad \nabla\cdot(\epsilon\mathbf{E})=0\quad \text{in } \Omega, \tag{1.1}$$

where ω is the angular frequency and $\mathbf{E}$ is the electric field. If we impose boundary conditions on $\partial\Omega$ (e.g., $\mathbf{E}\times\mathbf{n}=0$, where $\mathbf{n}$ is the vector normal to $\partial\Omega$), then solving Eq. (1.1) yields the discrete spectrum of the double curl operator on the left-hand side of Eq. (1.1). In practice, only a few nonzero low-order eigenvalues (which means small positive eigenvalues) and their associated eigenvectors are of interest. To find these, the finite element method (FEM) with higher-order basis functions [1] may be applied, resulting in a generalized matrix eigenvalue problem of the form:

$$\mathbb{K}x=\left(\frac{\omega}{c}\right)^2\mathbb{M}x, \tag{1.2}$$

where $\frac{\omega}{c}$ is a wavenumber, ω is the angular frequency, c is the speed of light, $\mathbb{K}$ is a large sparse symmetric and positive semidefinite matrix, $\mathbb{M}$ is a large sparse symmetric and positive definite matrix [2], and their sparsity decreases as the order of the FEM basis functions increases.

Symmetric generalized eigenvalue problems are also encountered in many other areas of physical modeling, including structural simulations, hydrodynamics, and solid-state physics [3, 4]. For such large-scale eigenproblems, Krylov subspace-based methods—such as the Lanczos, Arnoldi algorithms (e.g., Implicitly Restarted Arnoldi Method (IRAM) [5]), the Locally Optimal Block Preconditioned Conjugate Gradient (LOBPCG) method [6], and the Jacobi-Davidson-type techniques such as JDQZ [7])—have been developed. It is worth noting that Krylov subspace-based methods are designed for finding a few extreme eigenvalues.

As the size of the matrix and the number of nonzero matrix elements grow, the time required to solve such sparse generalized eigenproblems increases. To mitigate this problem, parallel computing techniques have been proposed [8–10]. In most cases massively parallel implementations are intended for clusters, but such computer systems are costly and not readily available for researchers and engineers. One of the most promising trends in parallel computing, allowing the acceleration of numerical code via massive parallelism on relatively low-cost hardware, is the utilization of graphic processing units (GPUs). GPUs support thousands of cores and different levels of fast memory which, if used appropriately, have been shown to be effective in accelerating sparse matrix processing, including iterative methods for sparse linear algebra exploiting the concept of a Krylov subspace. For these methods, the speedups achieved depend on the matrix sparsity pattern, the matrix compression scheme, and the way in which a crucial operation