Available online at www.sciencedirect.com
ScienceDirect
Comput. Methods Appl. Mech. Engrg. 290 (2015) 98–126
Computer methods in applied mechanics and engineering
www.elsevier.com/locate/cma

Numerical simulations of elastic wave propagation using graphical processing units—Comparative study of high-performance computing capabilities

P. Packo[a,*], T. Bielak[a], A.B. Spencer[b], T. Uhl[a], W.J. Staszewski[a], K. Worden[b], T. Barszcz[a], P. Russek[c], K. Wiatr[c]

[a] *Department of Robotics and Mechatronics, AGH University of Science and Technology, Al. Mickiewicza 30, 30-059 Krakow, Poland*
[b] *Department of Mechanical Engineering, University of Sheffield, Mappin Street, Sheffield S1 3JD, UK*
[c] *Academic Computer Centre Cyfronet AGH, AGH University of Science and Technology, ul. Nawojki 11, 30-950 Krakow, Poland*

Received 28 January 2014; received in revised form 5 March 2015; accepted 10 March 2015
Available online 19 March 2015

Abstract

High-performance computing is important in many areas of engineering. Increasing capabilities of modern workstations open new possibilities in scientific computing. This paper demonstrates how graphical processing units can be used efficiently for large models of elastic wave propagation in complex media. The method is based on the local interaction simulation approach and a parallel algorithm architecture. The focus of the work presented is on numerical implementation and covers aspects related to software modular architecture, computer memory organisation and optimisation. A domain decomposition approach allowing for calculations using multiple-GPU configurations is proposed, implemented and examined. The performance of the proposed simulation framework is tested for numerical models of different sizes, various computing precisions and hardware platforms. The results obtained are discussed in terms of graphical processing unit limitations. Obtained results indicate significant speed-up factors comparing to calculations using central processing units or different modelling approaches.

Keywords: CUDA; Parallel processing; Wave propagation; LISA

1. Introduction

Various types of elastic waves can propagate in solids. Waves travelling in an unbounded body of materials are known as bulk waves. These waves have a finite number of relatively simple longitudinal and shear components. In contrast, wave propagation in bounded media leads to guided waves that can exhibit an infinite number of complex dispersive modes. Analysis of elastic waves and their propagation has been investigated for many years in science

* Corresponding author. Tel.: +48 126173677, +48 604 526 235.
E-mail addresses: pawel.packo@agh.edu.pl (P. Packo), tuhl@agh.edu.pl (T. Uhl), w.j.staszewski@agh.edu.pl (W.J. Staszewski).

http://dx.doi.org/10.1016/j.cma.2015.03.002
0045-7825/

and engineering. These waves are governed by the fundamental elastodynamic equations based on linear (or non-linear) stress–strain relations and Newton's second law. It is well known that analytic solutions for this equation are only known for a range of canonical problems. Semi-analytic approaches include methods based on the theory of diffraction [1]. Therefore, various methods have been developed for modelling and numerical simulation of elastic wave propagation. This includes methods based on finite differences (FD) [2], velocity–stress FD [3], finite elements (FE) [4,5], spectral elements [6,7], the elastodynamic finite integration technique (EFIT) [8], local interaction simulation approach (LISA) [9–20] and cellular automata [21]. An excellent review of many different available approaches is given in [22].

The high frequencies and complex geometries involved in many applications often require large models of wave propagation. These models are challenging and computationally expensive; however parallel computation algorithms can be used to solve the problem [23,24]. This approach can be applied using either matrix manipulations or direct domain decomposition techniques. The former relies heavily on various properties of assembled equations and their coefficients (e.g. sparseness or diagonality) to speed up computations. However, the increase in calculation speed strongly depends on mathematical properties of matrices and model configuration (e.g. connectivity of elements). The latter method decomposes models into a number of sub-models that are calculated independently. Each sub-model is augmented by a set of boundary conditions that represent interactions with removed parts of original models. As a result original models are modified. Nevertheless, the speed-up factor obtained with the use of this approach is linear with respect to the number of domains, up to the hardware limitations. Hard disk space and speed is always the important limitation in hardware.

The aforementioned parallelisation techniques require specific hardware to carry out the computations. Typical computing units consist of a number of processors that can process data in parallel. Standard Central Processing Units (CPUs) typically include from eight to twelve processors. Although each core is relatively powerful, the total number of such units is relatively low. Usually when core numbers are increased costly hardware investments are required. Moreover, a significant computing power of these cores remains unused due to model subdivision into a number of relatively small sub-models. Although, data are exchanged between processors to advance computations, data transfers are limited by different factors, such as for example memory or bus speed. Therefore even large networks of computing units (e.g. computer clusters) are often exploited properly. Recent years have shown very fast evolution and development of Graphical Processing Units (GPUs). As a result, new architectures – available in low-cost graphical cards such as Computer Unified Device Architectures (CUDAs) – can be used very efficiently for numerical simulations [25]. GPUs consist of a large number of processors with extremely fast on-board memories. Therefore any computation problem can be decomposed into a number of threads and processed efficiently. However, it is important to note that these computing cores are not as powerful as standard CPUs and thus appropriate computation algorithms are required. The CUDA-based technology has been used to develop extremely efficient and accurate modelling tools in medicine, finance, seismology and computer aided design, as shown in [26–31].

The CUDA technology was introduced in 2006 by *Nvidia* in GPUs for graphics applications. Current *Nvidia* GPUs – widely available in graphic cards of high-end desktop PCs and laptops – have several computing multiprocessors, consisting of hundreds of cores, and capable of running thousands of light-weight threads in parallel. This technology, when combined with on-board, high-bandwidth memory, exceeds the performance and computing power of typical desktop CPUs. According to recent data the computational power of modern graphical cards is at least eight to ten times larger than the computation capability of CPUs. It is important to note that both approaches, i.e. CPU- and GPU-based, are used in practice for parallelisation of matrix operations. However, approaches employing large numbers of less powerful units have been found superior when compared with small numbers of more powerful units.

Efficient parallelisation algorithms are usually based on approaches that require less mathematical operations. Therefore explicit rather than implicit time integration techniques are preferred for wave propagation simulations. Also, it is important that analysed problems do not lead to increased computational complexities in single treads, for example due to intrinsic integral-related properties of the governing equations. Consequently, single treads should perform simple calculations involving small numbers of variables and operations. Therefore, when parallel computing algorithms are considered for numerical simulations of wave propagation, the LISA is one of the possible options. The method is perfectly suited for this task due to relative simplicity and proven record of various applications. The LISA has been proposed for wave propagation modelling in physics [14] and used for damage detection in structural health monitoring [9–13,18–20]. More recently, the new LISA-based platform for wave propagation modelling has been proposed in [20] and used for Lamb wave propagation modelling in structural damage detection applications.