GPU performance analysis of a nodal discontinuous Galerkin method for acoustic and elastic models

A. Modave^{*1}, A. St-Cyr², and T. Warburton¹

¹Virginia Polytechnic Institute and State University, Blacksburg, Virginia, USA

²Shell Global Solutions International B.V., Rijswijk, The Netherlands

February 2, 2016

Abstract

Finite element schemes based on discontinuous Galerkin methods possess features amenable to massively parallel computing accelerated with general purpose graphics processing units (GPUs). However, the computational performance of such schemes strongly depends on their implementation. In the past, several implementation strategies have been proposed. They are based [exclusively on specialized compute kernels](#) tuned for each operation, or they can leverage BLAS libraries that provide optimized routines for basic linear algebra operations. In this paper, we present and analyze up-to-date performance results for different implementations, tested in a unified framework on a single NVIDIA GTX980 GPU. We show that specialized kernels written with a one-node-per-thread strategy are competitive for polynomial bases up to the fifth and seventh degrees for acoustic and elastic models, respectively. For higher degrees, a strategy that makes use of the NVIDIA cuBLAS library provides better results, able to reach a net arithmetic throughput 35.7% of the theoretical peak value.

Keywords: finite element, discontinuous Galerkin, seismic waves, time domain, GPU, BLAS, profiling

1 Introduction

High-performance compute resources are used intensively in modern computational seismology for applications including earthquake simulation and seismic imaging. State-of-the-art compute clusters consist of massively parallel many-core central processing units (CPUs) with graphics processing units (GPUs) or [co-processors](#) to provide a performance boost. Computing on GPUs has been made accessible to the scientific community thanks to programming frameworks that expand languages already widely used (*e.g.* C, C++ or Fortran). However, to fully leverage the capabilities of GPUs, applications must be developed taking into account specifics of each architecture. It is therefore critical to develop computational algorithms that can be efficiently implemented on these architectures.

In computational seismology, several implementations have been proposed for numerical methods suitable for GPUs [coprocessors](#). Examples include tuned implementations for finite-difference schemes [1, 31, 39, 46] or finite-element schemes [18, 25, 26, 34, 35]. Nowadays, the finite-difference schemes are the most widely used, and a large literature is available (see *e.g.* Virieux *et al.* [41] for a review). However, the stencil based reconstruction used in finite difference methods is not ideally suited to resolving wave propagation in realistic physical heterogeneous media typically suffering from loss of accuracy [40]. By contrast, finite-element methods based on unstructured meshes are better suited to handle such material interfaces. Several variants have been investigated, such as spectral finite-element methods [23, 24], continuous mass-lumped

^{*}Corresponding author: modave@vt.edu

finite-element methods [8, 9] or discontinuous Galerkin (DG) methods [10–12, 14, 27, 30]. In contrast to the standard schemes, the mentioned methods do not require the solution of large sparse linear systems of equations, and it is possible to use explicit time-stepping schemes. The DG approach, in particular, provides a framework that is both flexible for multi-scale modeling and appears to be well suited for GPU accelerated computing. First, it can easily handle local time-stepping strategies [2, 13, 32, 45], and hybrid discretizations can be deployed by mixing different discretization orders and several kinds of elements [5, 13, 14, 21]. Then, the weak element-to-element coupling and the dense algebraic operations required per element are suitable for parallel multi-threading computations with GPUs [22].

Discontinuous Galerkin schemes can be implemented in different ways for GPU computing. Klöckner *et al.* [22] proposed an implementation for first-order wave systems discretized using a nodal DG scheme. This implementation has been successfully adapted to several physical contexts [16, 17, 34], and we have recently presented an application for reverse-time migration with multi-rate time-stepping and GPU clusters [33]. Alternatively, Fuhry *et al.* [15] have proposed an implementation for two-dimensional problems with a modal DG scheme. All these implementations partition the computational work into tailored GPU kernels, which are optimized separately in order to improve performance. A key difference between the approaches of Klöckner *et al.* [22] and Fuhry *et al.* [15] is the programming strategy of kernels: each thread performs computations corresponding to one node and to one element, respectively. On the other hand, Witherden *et al.* [49, 50] recently developed an implementation based on a nodal DG method for applications of fluid dynamics. This implementation, named PyFR, makes use of external BLAS libraries that provide linear algebra routines optimized for hardware devices. Such a strategy has been successfully used for CPU computing [7, 19, 28, 43].

In this paper, we investigate and evaluate three strategies to implement time-domain DG schemes for GPU computing. All these strategies have been compared for three-dimensional acoustic and elastic cases with a unique computational framework. They were programmed in C++ using CUDA 7.5 through the abstract framework OCCA [29], and tested on a single Nvidia GTX980. We have tested one-element-per-thread and one-node-per-thread strategies, as well as a strategy that makes use of an external BLAS library. In this work, we have used the generalized single precision matrix-multiplication (SGEMM) routine of NVIDIA’s cuBLAS library [36]. We show that, similarly to CPU implementations [43], the best GPU implementation depends on the polynomial degree. The one-node-per-thread tailored kernels provided the best runtime for small degree, while the implementation with SGEMM is better for higher degrees.

This paper is organized as follows. In section 2, we present the mathematical models for acoustic and elastic wave propagation, and we describe the nodal discontinuous Galerkin method and the time-stepping scheme. In section 3, key aspects of GPU hardware and GPU programming are summarized, and all the implementations are described. Section 4 is dedicated to performance results and comparisons. We discuss the optimization of the SGEMM routine for the DG operations, and we analyse its performance using the roofline model. The implementations are then compared, and all the kernels are systematically profiled.

2 Discontinuous Galerkin schemes

We consider acoustic and isotropic elastic wave models discretized with a nodal discontinuous Galerkin method in space and on a third-order Adam-Bashforth method in time. We assume the physical coefficients to be constant over each mesh cell and discontinuous at interfaces. The variational forms of the models and the schemes are presented in sections 2.1 and 2.2, respectively.

2.1 Physical models and variational forms

Acoustic waves are governed by the pressure-velocity system, which reads

$$\begin{aligned}\frac{\partial p}{\partial t} + \rho c^2 \nabla \cdot \mathbf{v} &= 0, \\ \rho \frac{\partial \mathbf{v}}{\partial t} + \nabla p &= 0,\end{aligned}$$

with the pressure $p(\mathbf{x}, t)$, the velocity $\mathbf{v}(\mathbf{x}, t)$, the density $\rho(\mathbf{x})$ and the phase velocity $c(\mathbf{x})$. For each mesh cell D_k , we consider the variational form

$$\begin{aligned} \int_{D_k} \frac{\partial p}{\partial t} \psi \, d\mathbf{x} &= - \int_{D_k} \rho c^2 \nabla \cdot \mathbf{v} \psi \, d\mathbf{x} - \int_{\partial D_k} p_p \psi \, d\mathbf{x} \\ \int_{D_k} \frac{\partial \mathbf{v}}{\partial t} \psi \, d\mathbf{x} &= - \int_{D_k} \frac{1}{\rho} \nabla p \psi \, d\mathbf{x} - \int_{\partial D_k} p_{\mathbf{v}} \mathbf{n} \psi \, d\mathbf{x} \end{aligned}$$

where $\psi(\mathbf{x})$ is a test function, $p_p(\mathbf{x})$ and $p_{\mathbf{v}}(\mathbf{x})$ are penalty terms. The penalty terms corresponding to upwind fluxes provided by a one-dimensional Riemann solver are given by (see e.g. [20, 33])

$$\begin{aligned} p_p &= -(\rho c^2)^- k_c \alpha_c, \\ p_{\mathbf{v}} &= c^- k_c \alpha_c, \end{aligned}$$

with $\alpha_c = \llbracket p \rrbracket - (\rho c)^+ \llbracket v_n \rrbracket$, $k_c = 1/\{\rho c\}$ and the semi-jumps defined as $\llbracket p \rrbracket = (p^+ - p^-)/2$ and $\llbracket v_n \rrbracket = (\mathbf{v}^+ - \mathbf{v}^-) \cdot \mathbf{n}/2$. The brackets $\{\cdot\}$ denotes the mean value at the interface.

For isotropic media, elastic waves can be simulated with the velocity-stress system, which reads

$$\begin{aligned} \frac{\partial \boldsymbol{\sigma}}{\partial t} &= \lambda(\nabla \cdot \mathbf{v}) \mathbf{I} + 2\mu \operatorname{sym}(\nabla \mathbf{v}), \\ \rho \frac{\partial \mathbf{v}}{\partial t} &= \nabla \cdot \boldsymbol{\sigma}, \end{aligned}$$

with the stress tensor $\boldsymbol{\sigma}(\mathbf{x}, t)$ and the Lamé parameters $\lambda(\mathbf{x})$ and $\mu(\mathbf{x})$. This system supports pressure waves and shear waves, which the phase velocities are

$$\begin{aligned} c_p &= \sqrt{(\lambda + 2\mu)/\rho}, \\ c_s &= \sqrt{\mu/\rho}, \end{aligned}$$

respectively. We consider the variational form

$$\begin{aligned} \int_{D_k} \frac{\partial \boldsymbol{\sigma}}{\partial t} \psi \, d\mathbf{x} &= \int_{D_k} \lambda(\nabla \cdot \mathbf{v}) \mathbf{I} + 2\mu \operatorname{sym}(\nabla \mathbf{v}) \psi \, d\mathbf{x} + \int_{\partial D_k} \mathbf{P}_{\boldsymbol{\sigma}} \psi \, d\mathbf{x}, \\ \int_{D_k} \frac{\partial \mathbf{v}}{\partial t} \psi \, d\mathbf{x} &= \int_{D_k} \frac{1}{\rho} (\nabla \cdot \boldsymbol{\sigma}) \psi \, d\mathbf{x} + \int_{\partial D_k} \mathbf{p}_{\mathbf{v}} \psi \, d\mathbf{x} \end{aligned}$$

where the upwind fluxes provided by an exact Riemann solver are given by [47]

$$\begin{aligned} \mathbf{P}_{\boldsymbol{\sigma}} &= k_p \alpha_p \lambda^- \mathbf{I} + 2(k_p \alpha_p - k_s \alpha_s) \mu^- \mathbf{N} + 2k_s \mu^- \operatorname{sym}(\mathbf{n} \otimes \boldsymbol{\alpha}_s), \\ \mathbf{p}_{\mathbf{v}} &= (k_p \alpha_p c_p^- - k_s \alpha_s c_s^-) \mathbf{n} + k_s c_s^- \boldsymbol{\alpha}_s, \end{aligned}$$

with

$$\begin{aligned} \alpha_p &= \llbracket \sigma_{nn} \rrbracket + (\rho c_p)^+ \llbracket v_n \rrbracket, & k_p &= 1/\{\rho c_p\}, & \llbracket \sigma_{nn} \rrbracket &= \mathbf{n} \cdot (\boldsymbol{\sigma}^+ - \boldsymbol{\sigma}^-) \cdot \mathbf{n}/2, \\ \alpha_s &= \llbracket \sigma_{nn} \rrbracket + (\rho c_s)^+ \llbracket v_n \rrbracket, & k_s &= 1/\{\rho c_s\}, & \llbracket \boldsymbol{\sigma} \mathbf{n} \rrbracket &= (\boldsymbol{\sigma}^+ - \boldsymbol{\sigma}^-) \cdot \mathbf{n}/2, \\ \boldsymbol{\alpha}_s &= \llbracket \boldsymbol{\sigma} \mathbf{n} \rrbracket + (\rho c_s)^+ \llbracket \mathbf{v} \rrbracket, & \mathbf{N} &= \mathbf{n} \otimes \mathbf{n}, & \llbracket \mathbf{v} \rrbracket &= (\mathbf{v}^+ - \mathbf{v}^-)/2. \end{aligned}$$

2.2 Numerical schemes

The approximate fields are built on a spatial mesh of the computational domain $\Omega \subset \mathbb{R}^3$, made of K non-overlapping tetrahedral cells, $\Omega = \bigcup_k D_k$, where D_k is the k^{th} cell. For the nodal discontinuous Galerkin method, all the scalar fields and the Cartesian components of vector and tensor fields are approximated by piecewise polynomial functions. The discrete unknowns correspond to the values of fields at nodes distributed

over the surface and interior of an element [19, 20]. Over each cell D_k , the approximate fields can then be represented by

$$\mathbf{q}_k(\mathbf{x}, t) = \sum_{n=1}^{N_p} \mathbf{q}_{k,n}(t) \ell_{k,n}(\mathbf{x}), \quad \forall \mathbf{x} \in D_k, \quad (1)$$

where N_p is the number of nodes per element, $\mathbf{q}_{k,n}(t)$ is the values of fields at node n of element k , and $\ell_{k,n}(\mathbf{x})$ is the associated multivariate Lagrange polynomial function. In this work, the position of nodes are chosen using the warp-and-blend technique [44]. The number of nodes for each element is given by $N_p = (N+1)(N+2)(N+3)/6$, where N is the maximal order of polynomial functions.

The semi-discrete equations are obtained by substituting the approximate representation of fields into the weak form, and using the Lagrange polynomials as test functions. For each element k , a system of equations can then be expressed as

$$\frac{d\mathbf{q}_k}{dt} = \mathbf{r}_k, \quad (2)$$

where the right-hand side vector is defined as

$$\mathbf{r}_k = \sum_{i=1}^3 \sum_{j=1}^3 g_{k,i,j}^{\text{vol}} \mathbf{D}_j \mathbf{f}_{j,k} + \sum_{f=1}^{N_{\text{faces}}} g_{k,f}^{\text{sur}} \mathbf{L}_f \mathbf{p}_{k,f}, \quad (3)$$

where N_{faces} is the number of faces, $\mathbf{f}_{j,k}$ correspond to the physical flux in the x_j -direction for all the nodes of element k , and the vector $\mathbf{p}_{k,f}$ contains the penalty terms for all the nodes belonging to face f . The first term of the right-hand side vector in equation (3) corresponds to the volume integrals of the variational form, while the second term corresponds to the surface integrals. The matrices $\mathbf{D}_j$ and $\mathbf{L}_f$ are respectively the differentiation matrices and the lifting matrices of the reference element. The geometric factors $g_{k,i,j}^{\text{vol}}$ and $g_{k,f}^{\text{sur}}$ depend on the shape of each element. The definitions of these matrices and factors, and a complete derivation of the semi-discrete equations are given in [33].

The global semi-discrete scheme is simply built by combining the local systems of all elements. We use the classical third-order Adams-Bashforth formula for time discretization. At each iteration, the discrete unknowns of all the elements are updated according to

$$\mathbf{q}_k^{m+1} = \mathbf{q}_k^m + \Delta t \sum_{s=1}^3 a_s \mathbf{r}_k^{m-s+1}, \quad (4)$$

with $a_1 = 23/12$, $a_2 = -16/12$ and $a_3 = 5/12$, and where m is the time index. This formula only depends on the right-hand side vector $\mathbf{r}_k$ at the three time levels t^m, t^{m-1}, t^{m-2} . For each step, this vector is thus computed with the values of unknowns local to the element, as well as those of the neighboring element at each interface.

3 Implementations for GPU computing

We consider three different GPU-based implementations of DG schemes. Each has been programmed in a C++ code using CUDA 7.5 through the abstract framework OCCA [29], and tested with a single NVIDIA GTX980. In section 3.1, we summarize key characteristics of NVIDIA's GPU architecture and CUDA programming. Specifications corresponding to the NVIDIA GTX980 are mentioned in brackets [...]. We present two implementations based [exclusively on specialized kernels](#) in section 3.2. In section 3.3, we propose an alternative strategy that makes use of the basic linear algebra subroutine SGEMM, from NVIDIA's cuBLAS library, optimized for dense algebraic operations on GPU.

3.1 GPU architecture and GPU programming

Heterogeneous computing systems consist of central processing units (CPUs) complemented with sidecar graphics processing units (GPUs), each with its own memory space. The programming framework CUDA [6, 37] provides extensions to the C programming language to allocate/deallocate storage in the global memory of GPU, to manage memory transfers between the CPU memory and the GPU memory, and to invoke *kernels* — the codes that run on the GPU. In our implementations, memory allocations and CPU-GPU data transfers are performed only at the beginning of the computation, and during the time stepping only when the solution must be exported for checkpointing or halo exchanges. The main challenge consists in optimizing the execution of tasks to be performed by the GPU.

The execution is based on the *single instruction multiple threads* (SIMT) model, where many *threads* are executed concurrently by the GPU. In the programming framework CUDA, threads are grouped into *thread blocks* [max. 1024 threads per thread block], which compose the *grid*. At the hardware level, each *thread block* is assigned to one *streaming multiprocessor* (SM) as the kernel executes [16 SMs on chip with 128 cores per SM], and successive threads of each thread block are grouped into *warps* [32 threads per warp]. Ideally, all the threads of a warp must perform the same instructions to avoid *warp divergence*. Warp schedulers manage the execution of warps, eventually context switching between warps to avoid the SM from stalling [max. 64 active warps per SM] due to high device memory load and store latencies. The percentage of active warps to the maximum theoretical number of active warps is called *occupancy*. It is generally advised that the number of threads per thread block should be chosen to be a multiple of 32 to maximize the occupancy [38], though, strategies with low occupancy can also be efficient [42].

The GPU has a deep non-uniform memory architecture with a hierarchy of memory spaces, each with different purposes and characteristics that indicate their best use cases. *Global memory* [4GB] is the largest on device memory space and is accessible by all the threads, in all warps, in all thread-blocks, in all kernels. The bandwidth of this memory [max. 224GB/s] is generally a bottleneck. Common strategies to maximize usage of global memory include hiding latency due to data transfer by overlapping memory load and store requests with computations, and making coalesced memory transfers by requesting contiguous blocks of data that are page aligned to maximize the utilization of the memory bus. Similarly to CPUs, GPUs have also a system of lower-latency *caches* (L2 cache [2MB] and L1 cache [48KB per SM]), which allow reuse of data previously fetched from the global memory. These caches cannot be programmed at the level of CUDA, and their use is determined by the streaming multiprocessor cache management system.

Shared memory (smem) [96KB per SM, max. 48KB per thread block] is a relatively low latency memory that can be accessed by threads of a thread block, providing a scratch pad for thread-block level data sharing and collaboration. Shared memory can be considered as a programmable cache, where memory storage is allocated at the thread block execution, and released when it is finished. Shared memory accesses are served by 32 independent memory managers. Shared memory is organized into banks that are interleaved in the logical indexing of the memory addresses. Efficient data transactions are served by the shared memory managers if all threads in a warp that access shared memory do so by accessing the 32 banks. Finally, data that are local to each thread are stored in 32-bit *registers* [64K per SM, max. 255 per thread] that provide very low latency accesses. When the number of registers is insufficient to contain all the register variables for all the threads in a thread-block the *spilled register* data are stored the local memory space, that is mapped to global memory and cached in L1 and L2 increasing the latency. The storage allocated for spilling registers is called *local memory* (lmem). NVIDIA's GPUs also have constant memory and texture memory, but they are not used in this work.

3.2 Implementations with specialized kernels

The operations to perform in the DG implementations can be partitioned into several kernels, which provides flexibility to optimize each task considering its own characteristics. For each time-step update, the right-hand side vector (3) must be computed at the current time, and the values of fields must be updated with the update scheme (4). Following previous works [15–17, 33], we consider implementations with three main kernels: the *volume kernel* and the *surface kernel* compute respectively the volume terms and the surface

Name	Symbol	Granularity <i>one-element-per-thread</i>	Granularity <i>one-node-per-thread</i>
Values of fields at nodes	Q	$N_{\text{fields}} \cdot N_p \cdot K$	$K \cdot N_{\text{fields}} \cdot N_p$
Values of traces at face nodes	Qf	$N_{\text{traces}} \cdot N_{\text{faces}} \cdot N_{fp} \cdot K$	$K \cdot N_{\text{traces}} \cdot N_{\text{faces}} \cdot N_{fp}$
Right-hand side terms	$\mathbf{R}_{\{0,-1,-2\}}$	$N_{\text{fields}} \cdot N_p \cdot K$	$K \cdot N_{\text{fields}} \cdot N_p$
Differentiation matrices	$\mathbf{D}_i$	N_p^2	N_p^2
Merged lift matrix	L	$N_{\text{faces}} \cdot N_{fp} \cdot N_p$	$N_{\text{faces}} \cdot N_{fp} \cdot N_p$

Table 1: List of main arrays stored in the global memory of the device. Dimensions are given from the coarsest to the finest granularity of storage. N_p , N_{fp} , N_{faces} and K are respectively the number of nodes per element, the number of nodes per face, the number of faces per element, and the number of elements in the mesh. N_{fields} and N_{traces} are the number of fields and traces, respectively.

terms of the right-hand side vector, while the *update kernel* performs the time stepping. A similar partition of operations into tailored kernels has also been successfully tested for computing with coprocessors [18]. We have tested implementations based on two different strategies: in all the kernels, each thread deals with the computations corresponding to one element (*one-element-per-thread* strategy) or one node (*one-node-per-thread* strategy). These implementations are based on the work of Fuhry *et al.* [15] and Klöckner *et al.* [22]. In the following we describe key aspects of memory storage and kernel programming with these strategies.

For both considered implementations, all the required data are stored in global memory: the array **Q** with the values of fields at all the nodes, the array **Qf** with the values of traces (p and $\mathbf{n} \cdot \mathbf{u}$ for the acoustic case, and $\mathbf{u}$ and $\boldsymbol{\sigma} \mathbf{n}$ for the elastic case) at all the face nodes, the arrays $\mathbf{R}_0$, $\mathbf{R}_{-1}$ and $\mathbf{R}_{-2}$ with the values of the right-hand-side vector corresponding to the current and two previous time steps, the arrays $\mathbf{D}_i$ and **L** with the differentiation and lift matrices, and specific arrays for the geometric factors and physical coefficients. The granularity of storage depends on the implementation strategy. With the one-element-per-thread strategy, successive threads have to fetch/store values corresponding to successive elements. It then is advantageous to store all the arrays with the element index k as the finest level of granularity, which enables coalescing data transfers with global memory [15]. By contrast, with the one-node-per-thread strategy, successive threads deal with successive nodes. A granularity with the node index n at the finest level then is the best choice [33]. Dimensions and granularity of the main arrays are given in Table 1.

We describe the tasks that must be performed by each kernel. For the acoustic case, the right-hand side vector (3) can be rewritten as

$$\mathbf{r}_k^p = -\rho_k c_k^2 \sum_{j=1}^3 \mathbf{D}_j \left[\sum_{i=1}^3 g_{k,i,j}^{\text{vol}} \mathbf{q}_k^{v_i} \right] - \mathbf{L} \mathbf{p}_k^p, \quad (5)$$

$$\mathbf{r}_k^{v_i} = -\frac{1}{\rho_k} \sum_{j=1}^3 g_{k,i,j}^{\text{vol}} \left[\mathbf{D}_j \mathbf{q}_k^p \right] - \mathbf{L} \mathbf{p}_k^{v_i}, \quad \text{for } i = 1, 2, 3, \quad (6)$$

where the lift matrices have been merged into **L**, and the penalty vectors $\mathbf{p}_k^{\{p,v_i\}}$ contain the penalty terms multiplied with the geometric factors for all the face nodes of the element. The superscripts p and v_i denote vectors with values corresponding to the pressure p and the Cartesian component of velocity v_i . Using similar notations, the right-hand side vector of the elastic case can be rewritten as

$$\begin{aligned} \mathbf{r}_k^{v_i} &= \frac{1}{\rho_k} \sum_{n=1}^3 \mathbf{D}_n \left[\sum_{m=1}^3 g_{k,m,n}^{\text{vol}} \mathbf{q}_k^{\sigma_{im}} \right] + \mathbf{L} \mathbf{p}_k^{v_i}, \quad \text{for } i = 1, 2, 3, \\ \mathbf{r}_k^{\sigma_{ij}} &= \delta_{ij} \lambda_k \sum_{m=1}^3 \sum_{n=1}^3 g_{k,m,n}^{\text{vol}} \left[\mathbf{D}_n \mathbf{q}_k^{v_m} \right] + \mu_k \sum_{m=1}^3 \left(g_{k,i,m}^{\text{vol}} \left[\mathbf{D}_m \mathbf{q}_k^{v_i} \right] + g_{k,m,i}^{\text{vol}} \left[\mathbf{D}_i \mathbf{q}_k^{v_m} \right] \right) + \mathbf{L} \mathbf{p}_k^{\sigma_{ij}}, \quad \text{for } i, j = 1, 2, 3. \end{aligned}$$

The volume kernel performs element-wise tasks corresponding to the volume terms, which are the first terms of equations (5)-(6) in the acoustic case. For each element, the values of fields are fetched from the array

Name	Memory transfer <i>one-element-per-thread</i>	Memory transfer <i>one-node-per-thread</i>
Fields and traces	Fetch into registers	Fetch into registers
Temporary array	Store in registers	Store in shared memory
Geometric and physical factors	Prefetch into registers	Prefetch into shared memory
Elemental matrices	Prefetch into shared memory	Fetch into registers

Table 2: Memory transfers performed by the volume and surface kernels. Fetched data are transferred when required, while prefetched data are loaded at the beginning of each kernel execution. The temporary array stores the vectors to be used in the matrix-vector products, which are the penalty vectors for the surface kernel.

$\mathbf{Q}$, and the main algebraic operations are linear combinations and six matrix-vector products with squared matrices of size $N_p \times N_p$. The surface kernel computes the penalty terms and performs the matrix-vector products for the second terms of equations (5)-(6). **In both volume and surface kernels, direct matrix-vector products are implemented. Let us mention that alternative DG methods with sparse elemental matrices can leverage sparse matrix-vector products (see e.g. [3, 4]).** The implementation of the surface kernel is trickier because computing the penalty terms requires the values of traces at face nodes corresponding to both sides of the interface (*i.e.* values for the current element and its neighbors), which leads to erratic memory accesses when fetching neighbor values. We use a mapping array that, for each face node, gives the location of the first trace corresponding to the neighbor element in the array $\mathbf{Qf}$. For each element, the surface kernel then loads both local and neighbor values of traces from the array $\mathbf{Qf}$, computes the penalty terms and multiplies them with geometric factors. The obtained values are then used in the merged matrix-vector product with a rectangular matrix of size $N_p \times N_{\text{faces}}N_{fp}$ and a vector of size $N_{\text{faces}}N_{fp}$. Both volume and surface kernels accumulate the final right-hand side terms in the current right-hand side array $\mathbf{R}_0$. The update kernel loads all the right-hand side arrays, updates the fields values in the array $\mathbf{Q}$ by performing the time-stepping, and updates the trace values in the array $\mathbf{fQ}$.

For the one-element-per-thread strategy, all kernels run with 256 threads per thread block. That number of threads has been suggested by Fuhry *et al.* [15], and is also optimal for our implementation. For both volume and surface kernels, the elemental matrices (differentiation and lift) are stored in global memory and prefetched into shared memory by all the threads in collaborative way (successive threads fetch successive values). Let us mention that, in the implementation of Fuhry *et al.* [15], the elemental matrices are stored in constant memory, and each value is fetched during the matrix-vector products. This strategy was however less efficient in our case. The geometric factors, physical parameters and fields/traces, which are local to each element, are stored in registers. The final values of the right-hand side terms are stored in temporary arrays before being transferred to local memory using coalesced data transfers. While this implementation uses a lot of registers, leading to spilling registers for high polynomial degree, it provides the best runtime with this strategy. The kinds of memory transfer are summarized in Table 2.

With the one-node-per-thread strategy, threads collaborate to perform the computations required for each element. For all the kernels, each thread block deals with several elements, which allow to have a reasonable number of threads per thread block for small polynomial degree. Therefore, we have $K_{\text{blkV}}N_p$ threads per thread block for the volume kernel, $K_{\text{blkS}}\max(N_p, N_{\text{faces}}N_{fp})$ for the surface kernel, and $K_{\text{blkU}}N_p$ for the update kernel. The numbers of elements per thread block, K_{blkV} , K_{blkS} and K_{blkU} , must be tuned. This implementation make use of shared memory to store temporary arrays with values that are useful for different threads. For more details about this implementation, we refer to our previous paper [33].

3.3 Implementations with SGEMM routine

We investigate an alternative implementation which makes use of an extern BLAS library. Such library provides linear algebra routines that are tuned for specific hardware devices. In our computational procedure, the matrix-vector products of equations (5)-(6) can be merged into global matrix-matrix products, which are BLAS-3 operations. They will be performed here with the SGEMM routine of the cuBLAS library

developed by Nvidia [36]. We describe hereafter the procedure only for the acoustic case, which can be straightforwardly adapted to the elastic case.

In order to integrate global matrix-matrix products in the computational procedure, we derive an alternative partition of tasks. The computation of the right-hand side terms can be rewritten with global operations. By merging equations (5)-(6) for all the elements, we have

$$\mathbf{R}^p = \sum_i \mathbf{D}_i \mathcal{L}_i^p(\mathbf{Q}^v) + \mathbf{L}\mathbf{P}^p, \quad (7)$$

$$\mathbf{R}^v = \mathcal{L}^v(\mathbf{D}_1 \mathbf{Q}^p, \mathbf{D}_2 \mathbf{Q}^p, \mathbf{D}_3 \mathbf{Q}^p) + \mathbf{L}\mathbf{P}^v, \quad (8)$$

where $\mathbf{Q}^{\{p,v\}}$, $\mathbf{P}^{\{p,v\}}$ and $\mathbf{R}^{\{p,v\}}$ are matrices for which the k^{th} column corresponds respectively to the values of fields, the penalty terms and the right-hand side vectors of element k . The operators $\mathcal{L}_i^p$ and $\mathcal{L}^v$ abstract the linear combinations to perform for the volume terms. A rapid study of equations (7)-(8) leads to the new partition of tasks:

1. computation of the *gradient* matrix-matrix product for equation (8) [*SGEMM routine*]:

$$\begin{bmatrix} \mathbf{D}_1 \\ \mathbf{D}_2 \\ \mathbf{D}_3 \end{bmatrix} \mathbf{Q}^p \rightarrow \begin{bmatrix} \mathbf{Q}_1^{\text{tmp}} \\ \mathbf{Q}_2^{\text{tmp}} \\ \mathbf{Q}_3^{\text{tmp}} \end{bmatrix}, \quad (\text{or } \mathbf{D}_{\text{grad}} \mathbf{Q}^p \rightarrow \mathbf{Q}^{\text{tmp}}) \quad (9)$$

2. computation of the linear combinations for the volume terms [*volume kernel*]:

$$\begin{aligned} \mathcal{L}^v(\mathbf{Q}_1^{\text{tmp}}, \mathbf{Q}_2^{\text{tmp}}, \mathbf{Q}_3^{\text{tmp}}) &\rightarrow \mathbf{R}^v, \\ \mathcal{L}_i^p(\mathbf{Q}^v) &\rightarrow \mathbf{Q}_i^{\text{tmp}}, \quad \text{for } i = 1, 2, 3, \end{aligned}$$

3. computation of the *divergence* matrix-matrix product for equation (7) [*SGEMM routine*]:

$$\begin{bmatrix} \mathbf{D}_1 & \mathbf{D}_2 & \mathbf{D}_3 \end{bmatrix} \begin{bmatrix} \mathbf{Q}_1^{\text{tmp}} \\ \mathbf{Q}_2^{\text{tmp}} \\ \mathbf{Q}_3^{\text{tmp}} \end{bmatrix} \rightarrow \mathbf{R}^p, \quad (\text{or } \mathbf{D}_{\text{div}} \mathbf{Q}^{\text{tmp}} \rightarrow \mathbf{R}^p) \quad (10)$$

4. computation of penalty vectors and storage in arrays $\mathbf{P}^p$ and $\mathbf{P}^{v_i}$ [*surface kernel*],

5. computation of the *lift* matrix-matrix products and addition to the right-hand side arrays [$2 \times$ *SGEMM routine*]:

$$\begin{aligned} \mathbf{L}\mathbf{P}^p + \mathbf{R}^p &\rightarrow \mathbf{R}^p, \\ \mathbf{L}\mathbf{P}^v + \mathbf{R}^v &\rightarrow \mathbf{R}^v, \end{aligned}$$

6. update of fields and traces with the time stepping scheme [*update kernel*].

In this new procedure, the first three operations compute the volume terms, the next two compute the surface terms, and the latter performs the update. The procedure requires three kernels and four calls of the SGEMM routine. We describe hereafter the arrays that must be stored in global memory, the different settings for the SGEMM routine, and the kernels.

In comparison with the previous implementations, arrays $\mathbf{Q}_i^{\text{tmp}}$ and $\mathbf{P}^{\{p,v\}}$ have been introduced to store intermediate results, which was not necessary before since these intermediate results were immediately used to compute matrix-vector products. For the sake of clarity, the fields, traces and right-hand side terms are stored in two sets of arrays, one for the pressure ($\mathbf{Q}^p$, $\mathbf{Q}\mathbf{f}^p$ and $\mathbf{R}^p$) and the other for the velocity ($\mathbf{Q}^v$, $\mathbf{Q}\mathbf{f}^v$ and $\mathbf{R}^v$). To efficiently use the SGEMM routine, the arrays $\mathbf{Q}_i^{\text{tmp}}$ are merged into the array $\mathbf{Q}^{\text{tmp}}$, and the differentiation arrays $\mathbf{D}_i$ are stored into two merged arrays, the gradient array $\mathbf{D}_{\text{grad}}$ and the divergence array

$\mathbf{D}_{\text{div}}$. The granularity of storage follows the one-node-per-thread strategy of the previous section, with the node index n at the finest level. This is required to use the SGEMM routine.

Conforming to BLAS terminology, the SGEMM routine computes the general matrix multiplication $\mathbf{C} = \alpha\mathbf{AB} + \beta\mathbf{C}$, where $\mathbf{A}$, $\mathbf{B}$ and $\mathbf{C}$ are respectively $m \times k$, $k \times n$ and $m \times n$ matrices and α and β are scalars. The memory storage of arrays must be column-major for the matrix $\mathbf{C}$, and either column-major or row-major for the others. They can be padded in order to improve computational performances. The padding for DG operations is studied in section 4.1.

The kernels of this implementation, corresponding to steps 2, 4 and 6, perform only data transfers and algebraic operations that are local to nodes or face nodes. The volume and update kernels are implemented following the one-node-per-thread strategy, with K_{blkV} and K_{blkU} elements per thread block. In contrast with the previous implementations, the surface kernel iterates over all the interfaces of the mesh. Each thread computes the numerical fluxes corresponding to one face node for both sides of the interface, and each thread block deals with K_{blkS} interfaces.

4 Performance results

The different implementations are tested and profiled in a basic setting: the time stepping of an initial solution on a given mesh. We consider polynomial degrees ranging from 1 to 8 with, for each degree, a mesh corresponding to approx. 2 millions nodes. All the results therefore correspond to the same number of discrete unknowns (approx. 8 and 18 millions for the acoustic and elastic models, respectively). The computations are performed in single precision.

In section 4.1, we discuss the memory storage of the SGEMM routine for the DG operations required by the third implementation. We then profile the SGEMM routine, report the arithmetic throughput and the memory bandwidth, and combine them in the roofline model (section 4.2). In sections 4.3, we profile the DG kernels of the three implementations for the acoustic case, and we compare the global performances of the implementations for both acoustic and elastic cases.

4.1 Storage strategy for SGEMM routine

The performance of the SGEMM routine strongly depends on both the dimensions of matrices and their storage in memory. For the operations required by the DG implementation, better performances are obtained with a row-major storage for $\mathbf{A}$, while a column-major storage is used for $\mathbf{B}$ and $\mathbf{C}$. We have tested several padding strategies. Only a padding of both the leading storage dimension ldc and the matrix dimension m , with $ldc = m$, can significantly improve performances. Increasing ldc involves a larger memory storage for matrix $\mathbf{C}$, while increasing m also adds dummy floating-point operations. Though these tunings artificially increase the required computational resources, they lead to better runtimes for the SGEMM routine. The dimensions of matrices and storing arrays for the different SGEMM calls are summarized in Table 3.

The matrix $\mathbf{C}$ corresponds to different arrays in our DG implementation: $\mathbf{Q}^{\text{tmp}}$ for the gradient product, $\mathbf{R}^p$ for the divergence product, and $\mathbf{R}^p$ and $\mathbf{R}^v$ for the lift products. The value to pad is $3N_p$ for the temporary array $\mathbf{Q}^{\text{tmp}}$, and N_p for the right-hand side arrays $\mathbf{R}^p$ and $\mathbf{R}^v$. For the sake of consistency, the padding used for the right-hand side arrays is also used for the arrays $\mathbf{Q}^p$ and $\mathbf{Q}^v$, which have the same kind of memory storage. Since $\mathbf{Q}^p$ and $\mathbf{Q}^{\text{tmp}}$ correspond to the matrix $\mathbf{B}$ of operations gradient and divergence, respectively, the parameter ldb of these operations has to be padded (see Table 3). This additional padding does not change the performance, and is made only for compatibility.

The unpadded and optimized padded values for N_p and $3N_p$ are listed in Table 4. Padding N_p improves the runtime only over polynomial degree $N > 5$, while padding $3N_p$ improves runtime for all degrees. For nearly all the cases, the best padding corresponds to increasing N_p and $3N_p$ until the next power of 2. The single exception is for $3N_p$ at the seventh degree, where the next multiple of 32 is better. For small polynomial degrees ($N \leq 5$), we have observed speedups between 1.1 and 1.5 for the gradient operation. The paddings do not affect the other operations. For higher degrees, speedups are between 2.1 and 3.2 for

	m	n	k	lda	ldb	ldc
Gradient	$(3N_p)_{\text{pad}}$	K	N_p	N_p	$(N_p)_{\text{pad}}$	$(3N_p)_{\text{pad}}$
Divergence	$(N_p)_{\text{pad}}$	K	$3N_p$	$3N_p$	$(3N_p)_{\text{pad}}$	$(N_p)_{\text{pad}}$
Lift	$(N_p)_{\text{pad}}$	$N_{ftl}K$	$N_{\text{faces}}N_{fp}$	$N_{\text{faces}}N_{fp}$	$N_{\text{faces}}N_{fp}$	$(N_p)_{\text{pad}}$

Table 3: Dimensions of matrices for the four SGEMM operations $\mathbf{C} = \alpha\mathbf{AB} + \beta\mathbf{C}$. According to BLAS terminology, the dimensions of $\mathbf{A}$, $\mathbf{B}$ and $\mathbf{C}$ are respectively $m \times k$, $k \times n$ and $m \times n$, and α and β are scalars. Storage is row-major for $\mathbf{A}$ and column-major for both $\mathbf{B}$ and $\mathbf{C}$. lda , ldb and ldc are the leading dimensions of storing arrays. N_p , N_{fp} , N_{faces} and K are respectively the number of nodes per element, the number of nodes per face, the number of faces per element, and the number of elements in the mesh. For the lift operation, N_{ftl} is the number of fields to lift (1 for p and 3 for $\mathbf{v}$).

N	N_p	$(N_p)_{\text{pad}}$	$3N_p$	$(3N_p)_{\text{pad}}$
1	4	4	12	16
2	10	10	30	32
3	20	20	60	64
4	35	35	105	128
5	56	56	168	256
6	84	128	252	256
7	120	128	360	384
8	165	256	495	512

Table 4: Unpadded and optimum padded values for N_p and $3N_p$ when using the SGEMM routine for the DG operations. The padded values have been obtained by minimizing the runtimes for DG settings corresponding to approx. 8 millions discrete unknowns. Gray cells correspond to pads leading to increases of dimensions larger than 50%. No padding is used for N_p from the first to the fifth polynomial degree.

the gradient operation and between 1.4 and 1.8 for the others. In some cases (gray cells in Table 4), these paddings lead to an increase of the size of the system larger than 50%.

4.2 Profiling of SGEMM routine for DG operations

To evaluate hardware utilization, we compare both arithmetic throughput and memory bandwidth of kernels to the theoretical peak values provided by the constructor, which are respectively 4,616 GFLOP/s (single precision) and 224 GB/s for Nvidia’s GTX980. The number of floating-point operations required for the general matrix multiplication is

$$\#(\text{FLOP}) = 2mn(k + 1)$$

and the number of bytes transfered from and to the global memory is

$$\#(\text{Byte}) = 4(mk + nk + 2mn),$$

where the dimensions m , n and k are not padded. Dividing these values by the runtime of the routine provides *estimated arithmetic throughput* and *estimated memory bandwidth*, respectively. These metrics give information about the performances for the requested operations, but they assume that there are no dummy floating-point operations and that data is always useful and transferred only once. However, optimization strategies, such as padding, can artificially increase the work really performed by the device. To evaluate that work, we also consider the *effective arithmetic throughput* and the *effective memory bandwidth*, which are obtained using metrics provided by Nvidia’s profiler `nvprof`. The effective arithmetic throughput is computed using the number of floating-point operations given by `flop_count_sp`, and the effective memory bandwidth is obtained by summing `dram_read_throughput` and `dram_write_throughput`.

Figure 1: Arithmetic throughput (a) and memory bandwidth (b) for the SGEMM routine for DG operations with several polynomial degrees and, for each degree, $N_p K \approx 2$ millions. Continuous lines correspond to estimated performances (*i.e.* considering only requested operations), while discontinuous lines are for effective performances (*i.e.* considering all the operations achieved by the device).

Figure 1 shows the estimated and effective performances of cuBLAS SGEMM routine for the DG operations presented in section 3.3. Since the performance of the lift operation is similar for pressure and velocity fields, we show the results only for the pressure field. For all the operations, the estimated arithmetic throughput is far smaller than the effective value for low degrees, especially for divergence and lift operations, while m is not padded. Large differences are also observed at high degrees when the padding significantly increases m : for the gradient operation with $N = 5$ and for both divergence and lift operations with $N = 6$ and 8. In particular, the estimated throughput significantly decreases from $N = 7$ to $N = 8$ for both divergence and lift operations, while the effective throughput slightly increases. This difference of behavior is clearly related to the large number of dummy operations introduced by the large padding required for $N = 8$, while the padded and unpadded dimensions are very close for $N = 7$. Similar differences between estimated and effective memory bandwidths are also observed for cases with large padding.

In the remainder, we consider only estimated arithmetic throughput and memory bandwidth. Though these values underestimate the real work of the device, they allow us to compare different implementations without counting potential dummy operations and useless data transfers.

For all operations, the estimated arithmetic throughput is very small for low degrees (from few percents to 20% of the peak performance between $N = 1$ and 4) and grows until approximately 60% for $N = 7$. The largest throughput, 65%, is obtained with the gradient operation for $N = 8$. The estimated memory bandwidth increases for lower degrees and decreasing for higher degrees. The bandwidth reaches 54–57% of the peak performance (for $N = 4$ and $N = 6$) for the gradient operation, while the other operations exhibit slightly worse performances.

The general behavior of both throughput and bandwidth can be explained using the roofline model. Because of the physical limitations of the device, the performance of any operation is theoretically bounded either by the throughput or by the bandwidth, but rarely by both together. The roofline model provides a tool to evaluate the dominant bound, and how far we are from that bound [48]. The arithmetic throughput is plotted as a function of the operational intensity, which is defined as

$$\text{operational intensity} = \frac{\#(\text{FLOP})}{\#(\text{Byte})} = \frac{2mn(k+1)}{4(mk+nk+2mn)}.$$

This metric determines the theoretical achievable throughput, which is

$$\text{achievable throughput} = \min \{ \text{peak bandwidth} \times \text{operational intensity}, \text{peak throughput} \}$$

Figure 2: Roofline analysis (a) of the SGEMM routine for DG operations, and percentage to achievable performance (b) considering the specific device. For every curves, each bullet corresponds to one polynomial degree. The leftmost bullet always corresponds to the first degree.

The operation is bandwidth or throughput bounded if the operational intensity is small or large, respectively. Figure 2a shows the curve of the achievable throughput (called the *roofline*) together with the obtained throughput for the different DG operations. The performance of the SGEMM routine for the DG operations are theoretically bounded by the bandwidth until the sixth or seventh degree, and by the arithmetic throughput beyond these polynomial degrees. This explains the low throughput observed for small degrees, as well as the stagnation of the throughput and the decaying memory bandwidth observed for the highest degrees.

Arithmetic throughput and memory bandwidth are representative of the implementation efficiency only when the kernel is throughput or bandwidth bounded, respectively. To quantify the effectiveness of the SGEMM implementation in a unique way, we consider the percentage of the obtained throughput to its achievable value. When the kernel is bandwidth bound, this percentage actually also corresponds to the percentage of the obtained bandwidth to the peak bandwidth. This metric, that we call *percentage to achievable performance*, is shown on figure 2b. We observe that the performance of the implementation increases with the polynomial degree, but only for the cases **without** large padding already discussed above.

4.3 Comparison and analysis of DG implementations

We compare and analyse the performance of the three DG implementations. The global runtime per node per time step is plotted on figure 3 for the acoustic case. We also show the detail kernel-by-kernel for each implementation. For the one-element-per-thread strategy, results are shown only up to the fifth polynomial degree. Larger degrees were not feasible due to memory limitations. This is discussed later.

For all implementations, the global runtime per node decreases with polynomial degree until some minimum value is achieved. After this, the runtime increases with polynomial order. The minimum is reached at $N = 3$ for both one-element-per-thread and one-node-per-thread strategies, and $N = 7$ for the strategy with the SGEMM routine. Until the fifth degree, the one-node-per-thread strategy provides the best performance. The runtime is larger by a factor two or more with the one-element-per-thread strategy. For higher degrees, the strategy with SGEMM is more efficient. Let us mention that these results are not necessarily correlated to the most suitable polynomial degree for applications. A complete performance analysis would include discussions of accuracy and time stepping. For instance, higher polynomial degrees allow for a better representation of high frequency modes, but they also reduce the allowed time step. Such an analysis is out of the scope of this paper.

In order to explain these results and to identify some bottlenecks in each implementation, we present the roofline and performance analyses of all the kernels in figure 4. Some statistics provided by Nvidia's compiler

Figure 3: Global runtimes per node per time step with the three DG implementations in the acoustic case (a), and runtimes kernel-by-kernel for each implementation (b)-(c)-(d). For each polynomial degree, a mesh corresponding to approx. 2 millions of nodes has been used.

`nvcc` (allocated registers, local memory and shared memory) and Nvidia’s profiler `nvprof` (occupancy with `achieved_occupancy` and hit rate for global loads with `global_hit_rate`) are listed in table 5, together with the optimum numbers of elements per thread block K_{blk} . All the kernels have been tuned and the K_{blk} ’s optimized to minimize the runtime as much as possible.

The roofline analyses show that all the kernels are bandwidth bound, except those that perform matrix-vector products with a large polynomial degree. All the update kernels and both volume and surface kernels of the strategy with SGEMM have small constant operational intensities. These kernels mainly perform streaming operations, and require less than one floating-point operation per transferred byte. The volume and surface kernels of both one-element-per-thread and one-node-per-thread strategies perform matrix-vector products, with which the number of required floating-point operations increases with the polynomial degree. These kernels are in the bandwidth bound regime until $N = 4$ (volume kernel) or $N = 5$ (surface kernel), and in the throughput bound regime beyond.

We first analyse the performance of the one-element-per-thread kernels. The runtime of the volume kernel is very low in comparison to the runtime corresponding to the same operation with the other strategies (blue areas on figure 3). However, the update kernel does not perform as well as the other versions, and the surface kernel gives very bad results. A major bottleneck of those kernels is the storage required in local and shared memory. This storage increases with the polynomial degree N . The maximum number of registers

Figure 4: Roofline and performance analyses of kernels for the different strategies in the acoustic case. For every curves, each bullet corresponds to one polynomial degree. The leftmost bullet always corresponds to the first degree.

Volume kernel (*one-element-per-thread*)

N	reg.	lmem	smem	occ.	gbl.h.
1	56	0	0.19	0.45	53%
2	64	0	1.17	0.35	51%
3	128	0	4.69	0.24	51%
4	186	0	14.35	0.12	51%
5	255	0.02	36.75	0.12	51%

Surface kernel (*one-element-per-thread*)

N	reg.	lmem	smem	occ.	gbl.h.
1	116	0	0.19	0.24	27%
2	149	0	0.94	0.12	24%
3	225	0	3.12	0.12	24%
4	255	0.07	8.20	0.12	24%
5	255	0.41	18.37	0.12	25%

Update kernel (*one-element-per-thread*)

N	reg.	lmem	smem	occ.	gbl.h.
1	72	0.07	0.05	0.36	52%
2	214	0.16	0.09	0.12	50%
3	255	0.38	0.16	0.12	50%
4	40	0.55	0.23	0.74	50%
5	40	0.88	0.33	0.71	49%

Volume kernel (*one-node-per-thread*)

N	K_{blk}	reg.	smem	occ.	gbl.h.
1	16	32	1.69	0.93	76%
2	16	35	3.19	0.88	91%
3	11	37	3.91	0.61	95%
4	5	32	2.95	0.90	97%
5	9	37	8.26	0.73	95%
6	3	37	4.07	0.74	94%
7	4	37	7.67	0.67	91%
8	3	32	7.86	0.98	87%

Surface kernel (*one-node-per-thread*)

N	K_{blk}	reg.	smem	occ.	gbl.h.
1	11	32	2.15	0.67	62%
2	2	32	0.58	0.95	75%
3	2	32	0.83	0.92	80%
4	2	31	1.14	0.93	84%
5	1	31	0.76	0.94	76%
6	9	32	8.79	0.89	85%
7	7	32	8.59	0.88	83%
8	5	31	7.54	0.85	82%

Volume kernel (*strategy with SGEMM*)

N	K_{blk}	reg.	smem	occ.	gbl.h.
1	8	27	0.34	0.49	41%
2	2	27	0.09	0.48	48%
3	1	27	0.04	0.48	40%
4	1	27	0.04	0.91	44%
5	1	27	0.04	0.91	50%
6	1	27	0.04	0.90	46%
7	1	27	0.04	0.92	50%
8	1	27	0.04	0.85	50%

Surface kernel (*strategy with SGEMM*)

N	K_{blk}	reg.	smem	occ.	gbl.h.
1	8	26	0.47	0.50	59%
2	4	26	0.33	0.50	53%
3	2	26	0.23	0.49	55%
4	1	26	0.15	0.49	52%
5	1	26	0.20	0.49	63%
6	1	26	0.25	0.49	59%
7	1	26	0.32	0.94	58%
8	1	26	0.39	0.93	62%

Update kernel (*one-node-per-thread and strategy with SGEMM*)

N	K_{blk}	reg.	smem	occ.	gbl.h.
1	16	30	1.00	0.87	53%
2	16	32	2.5	0.88	48%
3	10	32	3.12	0.77	53%
4	2	32	1.09	0.93	51%
5	1	32	0.88	0.91	58%
6	1	32	1.31	0.92	55%
7	1	32	1.88	0.88	56%
8	1	32	2.58	0.87	56%

N	Polynomial degree
K_{blk}	Number of thread per thread block
reg.	Number of register per thread
lmem	Local memory allocated per thread [KB]
smem	Shared memory allocated per thread block [KB]
occ.	Occupancy
gbl.h.	Hit rate for global loads

Table 5: Profiled statistics of each kernel for the different strategies in the acoustic case. The number of threads per thread block is always 256 for the one-element-per-thread strategy. There is no local memory allocated (lmem) with the one-node-per-thread strategy and the strategy with SGEMM.

Figure 5: Global runtimes per node per time step with the three DG implementations in the elastic case (a), and runtimes per node kernel-by-kernel for each implementation (b)-(c)-(d). For each polynomial degree, a mesh corresponding to approx. 2 millions of nodes has been used.

per thread for the device is not sufficient to store the local data from $N = 5$ for the volume kernel, from $N = 4$ for the surface kernel and for all polynomial degrees for the update kernel. This causes register spilling where local data is stored global memory (eventually cached in L1 or L2), which increases latency. For $N \geq 6$, the shared memory required by the volume kernel is larger than the maximum allowed (48KB), and the strategy cannot be used. The poor performance of the surface kernel is also partly explained by non-coalescing memory transfers. With this kernel, each thread has to load data corresponding to one element and its neighbors. Because the memory accesses are not aligned with the thread indexing for the neighbor elements, the memory transfers cannot coalesce. The hit rate for global loads, 24%–27% (see table 5), is far smaller than for the other kernels. We have observed that this rate rises until approx. 50% if the memory transfers corresponding to the neighbor elements are removed.

With the one-node-per-thread strategy, the surface kernel is the most expensive kernel for small polynomial degrees, while the volume kernel takes most of the runtime for large degrees (see figure 3c). For small degrees, the performance of the surface kernel is less efficient than for the volume kernel (see figure 4d), and it requires more floating-point operations (not shown for the sake of brevity). As N increases, the percentage of achievable performance of both kernels decreases, and reaches a plateau around 23% (volume kernel) and 30% (surface kernel) in the throughput bounded region. Over $N = 5$, the volume kernel requires more floating-point operations than the surface kernel, which explains the larger runtime. The one-node-

Figure 6: Net arithmetic throughput of the acoustic (a)-(b) and elastic (c)-(d) implementations. The net arithmetic throughput is obtained by dividing the total number of FLOP required per time step by the runtime required for one time step update.

per-thread update kernel, which is also used for the strategy with SGEMM, exhibits a constant and high percentage of achievable performance (around 70%). The volume kernel of that strategy also has a similar performance (around 67%). For the surface kernel, the performance is bad for small degrees, but it increases until a plateau around 52% for $N \geq 6$.

The performance results then are consistent with the observed runtimes: both volume and surface kernels perform better for small degrees with the one-node-per-thread strategy, while both surface kernel and SGEMM routine perform better for large degree for the strategy with SGEMM. The profiled statistics of the kernels are quite good. There is no register spilling, and both the number of registers per thread and the shared memory per thread block are reasonable. Increasing the number of elements per thread block helps to improve the occupancy for small polynomial degrees. The occupancy and the hit rate for global loads cannot explain the variations of performance that are observed. We have obtained from NVIDIA’s profiler `nvprof` that most of the stalls occur because of data requests (metric `stall_data_request`) and execution dependencies (metric `stall_exec_dependency`), which are difficult to avoid.

The observations made for the acoustic case also apply to the elastic case, which requires more floating-point operations and memory transfers. We show the global runtime per node for the three elastic implementations and the detailed kernel-by-kernel runtimes on figure 5. The simulations have been done on the same meshes. The one-node-per-thread strategy remains the most efficient for small degrees. The strategy with SGEMM is as efficient as the one-node-per-thread strategy at $N = 7$, and is better for $N = 8$. The one-element-per-thread strategy is again very bad, with results far worse than before. Since the operations require a larger local memory, register spilling occurs for N smaller than with the acoustic case, which worsens the results.

The one-node-per-thread strategy is only 1.8 – 2.25 times slower than for the acoustic case, while the number of fields increases by a factor 2.25. This is partly explained by the operational intensity that is larger in the elastic case for both volume and surface kernels. The strategy with SGEMM is 2.14 – 2.49 times slower than in the acoustic case.

Finally, to have an idea of the global performance of each strategy as a whole, we plot the net arithmetic throughput of each implementation for both acoustic and elastic cases in figure 6. The net arithmetic throughput is obtained by dividing the total number of floating-point operations required for one full time step by the corresponding runtime required. The implementation with SGEMM reaches 31.4% of the peak throughput for the acoustic case, and 35.8% for the elastic case. A larger increase is observed with the one-node-per-thread strategy, from 21.7% to 33%. This made of this strategy an interesting candidate for DG schemes requiring denser computations.

5 Conclusion

We have presented and analyzed three GPU implementations for time-domain discontinuous Galerkin simulations of acoustic and elastic waves. Two of them are based on tailored kernels programmed by associating each thread to one element or one node, following strategies described by Klöckner *et al.* [22] and Fuhry *et al.* [15], respectively. We have also considered an alternative implementation with a specific partition of work which makes use of the SGEMM routine of an external BLAS library. All the implementations have been optimized and compared in the same computational framework with a three-dimensional benchmark.

The computational results highlight the strong performance of the one-node-per-thread kernels until $N = 5$ for acoustic and $N = 7$ for elastic waves. The strategy with SGEMM performs better for higher polynomial degrees, though it requires more data transfers and more data storage in global memory. The one-element-per-thread implementation does not work well for three-dimensional problems in our framework. We have analysed both the SGEMM routine and the acoustic DG kernels using the roofline model, and we have identified bottlenecks for each computational strategy.

In the perspective of large-scale applications, it is essential to combine the selected implementation with a multi-rate time stepping scheme and parallel computation with several GPUs. The one-node-per-thread implementation has already been adapted to multi-rate time stepping [16, 17] and multi-GPU computation [33]. The implementation with SGEMM can be adapted in very similar ways. In the future, we plan to investigate kernels for anisotropic wave propagation and artificial boundary conditions.

Acknowledgements

This work has been partly supported by NSF (award number DMS-1216674), Shell Global Solutions International B.V. and Shell Oil Company. Axel Modave is partially supported by an excellence grant from Wallonie-Bruxelles International (WBI). He is an Honorary Fellow of the Belgian American Educational Foundation (BAEF) and a Postdoctoral Researcher on leave with the F.R.S-FNRS. The authors acknowledge Ty McKercher and NVIDIA Corporation for providing GPU equipment. The authors gratefully thank Jesse Chan for his helpful comments on the manuscript.

References

- [1] R. Abdelkhalek, H. Calandra, O. Coulaud, G. Latu, and J. Roman. Fast seismic modeling and reverse time migration on a graphics processing unit cluster. *Concurrency and Computation: Practice and Experience*, 24(7):739–750, 2012.
- [2] C. Baldassari, H. Barucq, H. Calandra, and J. Diaz. Numerical performances of a hybrid local-time stepping strategy applied to the reverse time migration. *Geophysical Prospecting*, 59(5):907–919, 2011.
- [3] A. Breuer, A. Heinecke, S. Rettenberger, M. Bader, A.-A. Gabriel, and C. Pelties. Sustained petascale performance of seismic simulations with seissol on supermuc. In *Supercomputing*, pages 1–18. Springer, 2014.
- [4] J. Chan and T. Warburton. GPU-accelerated bernstein-bezier discontinuous galerkin methods for wave problems. 2015. <http://arxiv.org/abs/1512.06025>.
- [5] J. Chan, J. Wang, A. Modave, J.-F. Remacle, and T. Warburton. GPU-accelerated discontinuous Galerkin methods on hybrid meshes. 2015. <http://arxiv.org/abs/1507.02557>.
- [6] J. Cheng, M. Grossman, and T. McKercher. *Professional Cuda C Programming*. John Wiley & Sons, 2014.

- [7] N. Chevaugeon, J.-F. Remacle, X. Gallez, P. Ploumhans, and S. Caro. Efficient discontinuous Galerkin methods for solving acoustic problems. In *11th AIAA/CEAS Aeroacoustics Conference*, page 2823, 2005.
- [8] M. J. S. Chin-Joe-Kong, W. A. Mulder, and M. Van Veldhuizen. Higher-order triangular and tetrahedral finite elements with mass lumping for solving the wave equation. *Journal of Engineering Mathematics*, 35(4):405–426, 1999.
- [9] G. Cohen, P. Joly, J. E. Roberts, and N. Tordjman. Higher order triangular finite elements with mass lumping for the wave equation. *SIAM Journal on Numerical Analysis*, 38(6):2047–2078, 2001.
- [10] S. S. Collis, C. C. Ober, and B. G. van Bloemen Waanders. Unstructured discontinuous Galerkin for seismic inversion. In *Proceedings of the 80th SEG Annual Meeting*, 2010.
- [11] J. de la Puente, M. Käser, M. Dumbser, and H. Igel. An arbitrary high-order discontinuous Galerkin method for elastic waves on unstructured meshes - IV. Anisotropy. *Geophysical Journal International*, 169(3):1210–1228, 2007.
- [12] M. Dumbser and M. Käser. An arbitrary high-order discontinuous Galerkin method for elastic waves on unstructured meshes – II. The three-dimensional isotropic case. *Geophysical Journal International*, 167(1):319–336, 2006.
- [13] M. Dumbser and M. Käser. A p-adaptive discontinuous Galerkin method with local time steps for computational seismology. In *High Performance Computing in Science and Engineering, Garching/Munich 2007*, pages 569–584. Springer, 2009.
- [14] V. Etienne, E. Chaljub, J. Virieux, and N. Glinsky. An hp-adaptive discontinuous Galerkin finite-element method for 3-D elastic wave modelling. *Geophysical Journal International*, 183(2):941–962, 2010.
- [15] M. Fuhry, A. Giuliani, and L. Krivodonova. Discontinuous Galerkin methods on graphics processing units for nonlinear hyperbolic conservation laws. *International Journal for Numerical Methods in Fluids*, 76(12):982–1003, 2014.
- [16] R. Gandham, D. Medina, and T. Warburton. GPU accelerated discontinuous Galerkin methods for shallow water equations. *Communications in Computational Physics*, 18(1):37–64, 2015.
- [17] N. Gödel, S. Schomann, T. Warburton, and M. Clemens. GPU accelerated Adams-Bashforth multirate discontinuous Galerkin FEM simulation of high-frequency electromagnetic fields. *Magnetics, IEEE Transactions on*, 46(8):2735–2738, 2010.
- [18] A. Heinecke, A. Breuer, S. Rettenberger, M. Bader, A.-A. Gabriel, C. Pelties, A. Bode, W. Barth, X.-K. Liao, K. Vaidyanathan, et al. Petascale high order dynamic rupture earthquake simulations on heterogeneous supercomputers. In *High Performance Computing, Networking, Storage and Analysis, SC14: International Conference for*, pages 3–14. IEEE, 2014.
- [19] J. S. Hesthaven and T. Warburton. Nodal high-order methods on unstructured grids: I. Time-domain solution of Maxwell’s equations. *Journal of Computational Physics*, 181(1):186–221, 2002.
- [20] J. S. Hesthaven and T. Warburton. *Nodal discontinuous Galerkin methods: algorithms, analysis, and applications*, volume 54. Springer, 2007.
- [21] R. M. Kirby, T. C. Warburton, I. Lomtev, and G. E. Karniadakis. A discontinuous Galerkin spectral/hp method on hybrid grids. *Applied numerical mathematics*, 33(1):393–405, 2000.
- [22] A. Klöckner, T. Warburton, J. Bridge, and J. S. Hesthaven. Nodal discontinuous Galerkin methods on graphics processors. *Journal of Computational Physics*, 228(21):7863–7882, 2009.

- [23] D. Komatitsch and J. Tromp. Introduction to the spectral element method for three-dimensional seismic wave propagation. *Geophysical journal international*, 139(3):806–822, 1999.
- [24] D. Komatitsch and J.-P. Vilotte. The spectral element method: an efficient tool to simulate the seismic response of 2D and 3D geological structures. *Bulletin of the seismological society of America*, 88(2):368–392, 1998.
- [25] D. Komatitsch, D. Michéa, and G. Erlebacher. Porting a high-order finite-element earthquake modeling application to NVIDIA graphics cards using CUDA. *Journal of Parallel and Distributed Computing*, 69(5):451–460, 2009.
- [26] D. Komatitsch, G. Erlebacher, D. Göddeke, and D. Michéa. High-order finite-element seismic wave propagation modeling with MPI on a large GPU cluster. *Journal of Computational Physics*, 229(20):7692–7714, 2010.
- [27] J. R. Krebs, S. S. Collis, N. J. Downey, C. C. Ober, J. R. Overfelt, T. M. Smith, B. G. van Bloemen-Waanders, and J. G. Young. Full wave inversion using a spectral-element discontinuous Galerkin method. In *Proceedings of the 76th EAGE Conference & Exhibition*, 2014.
- [28] E. Marchandise, J.-F. Remacle, and N. Chevaugeon. A quadrature-free discontinuous Galerkin method for the level set equation. *Journal of Computational Physics*, 212(1):338–357, 2006.
- [29] D. S. Medina, A. St.-Cyr, and T. Warburton. OCCA: A unified approach to multi-threading languages. 2014. <http://arxiv.org/abs/1403.0968>.
- [30] E. D. Mercerat and N. Glinsky. A nodal high-order discontinuous Galerkin method for elastic wave propagation in arbitrary heterogeneous media. *Geophysical Journal International*, 201(2):1101–1118, 2015.
- [31] D. Michéa and D. Komatitsch. Accelerating a three-dimensional finite-difference wave propagation code using GPU graphics cards. *Geophysical Journal International*, 182(1):389–402, 2010.
- [32] S. Minisini, E. Zhebel, A. Kononov, and W. A. Mulder. Local time stepping with the discontinuous Galerkin method for wave propagation in 3D heterogeneous media. *Geophysics*, 78(3):T67–T77, 2013.
- [33] A. Modave, A. St.-Cyr, W. Mulder, and T. Warburton. A nodal discontinuous Galerkin method for reverse-time migration on GPU clusters. *Geophysical Journal International*, 203(2):1419–1435, 2015.
- [34] A. Modave, A. St.-Cyr, T. Warburton, and W. A. Mulder. Accelerated discontinuous Galerkin time-domain simulations for seismic wave propagation. In *Proceedings of the 77th EAGE Conference & Exhibition*, 2015.
- [35] D. Mu, P. Chen, and L. Wang. Accelerating the discontinuous Galerkin method for seismic wave propagation simulations using multiple GPUs with CUDA and MPI. *Earthquake Science*, 26(6):377–393, 2013.
- [36] NVIDIA Corporation. cuBLAS Library: User Guide. <http://docs.nvidia.com/cuda/cublas/>, 2015.
- [37] NVIDIA Corporation. CUDA Toolkit Documentation v7.5. <http://docs.nvidia.com/cuda/>, 2015.
- [38] NVIDIA Corporation. CUDA C Best Practices Guide. <http://docs.nvidia.com/cuda/cuda-c-best-practices-guide/>, 2015.
- [39] F. Rubio, M. Hanzich, A. Farrés, J. De La Puente, and J. M. Cela. Finite-difference staggered grids in gpus for anisotropic elastic wave propagation simulation. *Computers & Geosciences*, 70:181–189, 2014.
- [40] W. W. Symes and T. Vdovina. Interface error analysis for numerical wave propagation. *Computational Geosciences*, 13(3):363–371, 2009.

- 575 [41] J. Virieux, H. Calandra, and R.-É. Plessix. A review of the spectral, pseudo-spectral, finite-difference
576 and finite-element modelling techniques for geophysical imaging. *Geophysical Prospecting*, 59(5):794–
577 813, 2011.
- 578 [42] V. Volkov. Better performance at lower occupancy. In *Proceedings of the GPU Technology Conference*,
579 *GTC*, volume 10. San Jose, CA, 2010.
- 580 [43] P. E. J. Vos, S. J. Sherwin, and R. M. Kirby. From h to p efficiently: Implementing finite and spectral/hp
581 element methods to achieve optimal performance for low-and high-order discretisations. *Journal of*
582 *Computational Physics*, 229(13):5161–5181, 2010.
- 583 [44] T. Warburton. An explicit construction of interpolation nodes on the simplex. *Journal of Engineering*
584 *Mathematics*, 56(3):247–262, 2006.
- 585 [45] T. Warburton. Accelerating the discontinuous Galerkin time-domain method. *Oberwolfach Report*, 36:
586 58–60, 2008.
- 587 [46] R. M. Weiss and J. Shragge. Solving 3D anisotropic elastic wave equations on parallel GPU devices.
588 *Geophysics*, 78(2):F7–F15, 2013.
- 589 [47] L. C. Wilcox, G. Stadler, C. Burstedde, and O. Ghattas. A high-order discontinuous Galerkin method
590 for wave propagation through coupled elastic–acoustic media. *Journal of Computational Physics*, 229
591 (24):9373–9396, 2010.
- 592 [48] S. Williams, A. Waterman, and D. Patterson. Roofline: an insightful visual performance model for
593 multicore architectures. *Communications of the ACM*, 52(4):65–76, 2009.
- 594 [49] F. Witherden, A. Farrington, and P. Vincent. PyFR: An open source framework for solving advection–
595 diffusion type problems on streaming architectures using the flux reconstruction approach. *Computer*
596 *Physics Communications*, 185:3028–3040, 2014.
- 597 [50] F. Witherden, B. Vermeire, and P. Vincent. Heterogeneous computing on mixed unstructured grids
598 with PyFR. *Computers & Fluids*, 120:173–186, 2015.