

Acceleration of Element-by-Element Kernel in Unstructured Implicit Low-order Finite-element Earthquake Simulation using OpenACC on Pascal GPUs

Kohei Fujita

Advanced Institute for Computational Science,
RIKEN
7-1-26 Minatojima-minami-machi, Chuo,
Kobe, Hyogo 650-0047, Japan
Email: kohei.fujita@riken.jp

Takuma Yamaguchi

Department of Civil Engineering,
The University of Tokyo
1-1-1 Yayoi, Bunkyo,
Tokyo, 113-0032, Japan
Email: yamaguchi@eri.u-tokyo.ac.jp

Tsuyoshi Ichimura

Earthquake Research Institute and
Department of Civil Engineering,
The University of Tokyo
1-1-1 Yayoi, Bunkyo,
Tokyo, 113-0032, Japan
Email: ichimura@eri.u-tokyo.ac.jp

Muneo Hori

Earthquake Research Institute and Department of Civil Engineering,
The University of Tokyo
Advanced Institute for Computational Science, RIKEN
1-1-1 Yayoi, Bunkyo, Tokyo, 113-0032, Japan
Email: hori@eri.u-tokyo.ac.jp

Lalith Maddegadara

Earthquake Research Institute and
Department of Civil Engineering,
The University of Tokyo
1-1-1 Yayoi, Bunkyo, Tokyo, 113-0032, Japan
Email: lalith@eri.u-tokyo.ac.jp

Abstract—The element-by-element computation used in matrix-vector multiplications is the key kernel for attaining high-performance in unstructured implicit low-order finite-element earthquake simulations. We accelerate this CPU-based element-by-element kernel by developing suitable algorithms for GPUs and porting to a GPU-CPU heterogeneous compute environment by OpenACC. Other parts of the earthquake simulation code are ported by directly inserting OpenACC directives into the CPU code. This porting approach enables high performance with relatively low development costs. When comparing eight K computer nodes and eight NVIDIA Pascal P100 GPUs, we achieve 23.1 times speedup for the element-by-element kernel, which leads to 16.7 times speedup for the 3×3 block Jacobi preconditioned conjugate gradient finite-element solver. We show the effectiveness of the proposed method through many-case crust-deformation simulations on a GPU cluster.

I. INTRODUCTION

To improve the estimation reliability of empirical methods, we have been developing an integrated earthquake simulation based on physics-based numerical analyses of each phase of an earthquake disaster. Since such earthquake simulations require unstructured, low-order, implicit finite-element analyses for appropriate modeling of nonlinear phenomena in complex-shaped domains, we have been speeding up this method to enable realistic earthquake simulations (e.g., Gordon Bell Prize Finalists in 2014 [1] and 2015 [2]). This method is regarded as a prospective next-generation earthquake simulation by companies and governmental sectors, and many joint research

projects are on-going. Our development has targeted the K computer [3], which is a CPU-based massively parallel super-computer; however, companies commonly have different types of computing systems, and there is demand for simulations that can run on a wide range of computing equipment. In addition, we may be able to further speed up the analysis by using the rapidly developing GPU accelerators. OpenACC [4], which enables offloading of loops and regions of a code to various accelerators by compiler directives, is suitable for porting pre-developed CPU applications to heterogeneous compute environments without altering the performance of the original CPU code. Thus, in this study, we port the fast solver developed for the K computer, which was published in SC15 proceedings [2], to GPU-CPU heterogeneous environment by OpenACC to make the simulation method available to a wider range of compute environments. We expect that this will also improve the runtime speed.

Compared to conventional finite-element solvers used for large-scale analyses, the unstructured implicit finite-element solver reported in the SC15 paper is designed to reduce communication and memory access and increase the ratio of single-precision computation to double-precision computation. This algorithm is suitable for current CPU architectures, and the algorithm attained 18.6% of peak (1.97 PFLOPS) for the whole solver when running on the full K computer, which is very high for an implicit, unstructured, low-order finite-element analysis. This single-precision computation-rich algorithm is also suitable for GPU architectures; however, the

performance of some complex kernels can severely degrade due to changes in the design of CPU and GPU platforms. Indeed, the element-by-element computation [7], which is a matrix-free matrix-vector multiplication method, is complex due to the unstructured list access and data recurrence. The use of the algorithm designed for multi-cores, SIMD units, and the large cache used in this kernel may not achieve adequate performance when ported to GPUs. Since this kernel is the most computationally expensive part of the solver, acceleration of this kernel becomes the key for fast earthquake simulation on GPU-CPU heterogeneous environments. Thus, in this study, we first develop algorithms suitable for GPU platforms for the element-by-element computation, and later port this kernel by OpenACC. Other parts of the code that do not require changing the algorithm are ported directly by adding OpenACC directives to CPU code. This porting approach enables high performance with relatively low development costs. To demonstrate the performance of the developed element-by-element kernel, we measure performance on the newest NVIDIA Pascal GPUs [5]. We also present an application example of fault slip estimation of the 2011 Tohoku Earthquake by many-case crust-deformation simulations enabled using the ported earthquake simulation on a GPU cluster.

The remainder of this paper is organized as follows. Section 2 summarizes the solver developed for the K computer [2]. Section 3 presents the detailed algorithm suitable for a GPU and the OpenACC porting details. The performance relative to algorithmic choice and different implementations is also presented in this section. Section 4 shows an application example, and Section 5 summarizes related work. Section 6 summarizes the paper.

II. BASELINE FINITE-ELEMENT EARTHQUAKE SIMULATION ON THE K COMPUTER

A. Target Problem

Targeting estimation of the distribution of fault slip, we solve the crust-deformation problem for fault slip as

$$\frac{\partial \sigma_{ij}}{\partial x_j} = f_i,$$

with

$$\sigma_{ij} = c_{ijkl} \epsilon_{kl}, \quad \epsilon_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right).$$

Here, σ , x , f , ϵ , c , and u indicate stress, coordinates, outer force, strain, elastic tensor, and displacement, respectively. To estimate the slip state properly, we must model the stress and strain change between plates as well as the three-dimensional geometry and material properties connecting the fault and observation points. Thus, we use second-order tetrahedral elements for discretization of the problem, which enables accurate computation of stress and strain in complex-shaped domains. The target problem thus becomes

$$Ku = f. \quad (1)$$

Here, u , and f are displacement and nodal force vectors, respectively, and K is the stiffness matrix, which is sparse and

symmetric positive definite. Since nearly all of the computation cost is spent on solving this linear system of equations, acceleration of the solver is important.

B. Solver Overview

Solving Eq. (1) is a common target problem arising from low-order implicit finite-element methods used in solid continuum mechanics problems in science and engineering fields. In these low-order (e.g., second-order) implicit simulations, highly efficient high-order methods cannot be used. Thus, iterative solvers, such as conjugate gradient (CG) methods, using matrices stored explicitly in memory are generally used. Since these solvers involve significant random access and intensive memory bandwidth access, it is difficult to achieve adequate performance with current CPU and GPU systems. Previous work [2] accelerated this solver by reducing the data transfer size, communication size, and computation cost compared to standard solvers. The main idea here is the combination of an inexact preconditioned CG method, a multigrid method, and mixed-precision computation. Below, we explain the contents following Algorithm 1.

- Inexact preconditioned CG method [6]. An inexact preconditioned CG method is a preconditioned CG method that uses another solver to solve the preconditioning matrix equation rather than multiplying a fixed preconditioning matrix. As it is common to use another iterative solver to solve the preconditioning matrix equation, we refer to the original CG loop as the “outer loop” (Algorithm 1, lines 8-32) and the preconditioning solver as the “inner loop” (Algorithm 1, lines 10-19). As the inner loop does not need to be solved exactly, this makes room for improvements such as multigrid and mixed-precision arithmetic.
- Multigrid method. We reduce the computation cost and communication size of the inexact preconditioned CG method by using a multigrid solver for the inner loop. One limitation of multigrid solvers is the possibility of poor scalability in massively parallel compute environments due to the communication required to map between grids, or the difficulty in obtaining load balance for all grid levels. We circumvent this problem by using a two-step geometric multigrid that does not require communication for mapping and maintains good load balance in both grids. Here, we use the targeted second-order tetrahedral mesh FEmodel for the inner fine loop (Algorithm 1, line 17), and the same mesh without intermediate nodes of each element for the inner coarse loop (i.e., a first-order tetrahedral mesh FEmodel; Algorithm 1, line 15). We use a robust and efficient CG solver with 3×3 block Jacobi preconditioning to solve the inner fine and coarse loops (Algorithm 2).
- Mixed precision arithmetic. Even for a target problem in double precision, we only need to solve the preconditioning matrix equation roughly. Thus, we compute the whole preconditioner in single precision. This leads to halving the memory access and communication size, and enables

Algorithm 1 Details of Algorithm for solving $\mathbf{K}\mathbf{u} = \mathbf{f}$. The matrix-vector multiplication $\mathbf{K}\mathbf{y}$ is computed using an Element-by-Element method. $diag[\]$, $(\)^c$, and ϵ indicate the 3×3 block Jacobi of $[\]$, single-precision variable, and tolerance for relative error, respectively. $(\)_c$ indicates the calculation related to FEMmodel_c, and the other is the related calculation of the FEMmodel. $\bar{\mathbf{P}}$ is a mapping matrix, from FEMmodel_c¹ to FEMmodel¹, which is defined by interpolating the displacement in each element of FEMmodel_c¹. Solved on CPU except for lines 15 and 17.

```

1: set  $\mathbf{f}$  according to boundary condition
2:  $\mathbf{x} \leftarrow 0$ 
3:  $\bar{\mathbf{B}} \leftarrow diag[\mathbf{K}]$ 
4:  $\bar{\mathbf{B}}_c \leftarrow diag[\mathbf{K}_c]$ 
5:  $\mathbf{r} \leftarrow \mathbf{f}$ 
6:  $\beta \leftarrow 0$ 
7:  $i \leftarrow 1$ 
8: (*outer loop start*)
9: while  $\|\mathbf{r}\|_2 / \|\mathbf{f}\|_2 \geq \epsilon$  do
10:   (*inner loop start*)
11:    $\bar{\mathbf{r}} \leftarrow \mathbf{r}$ 
12:    $\bar{\mathbf{z}} \leftarrow \bar{\mathbf{B}}^{-1} \bar{\mathbf{r}}$ 
13:    $\bar{\mathbf{r}}_c \leftarrow \bar{\mathbf{P}}^T \bar{\mathbf{r}}$ 
14:    $\bar{\mathbf{z}}_c \leftarrow \bar{\mathbf{P}}^T \bar{\mathbf{z}}$ 
15:    $\bar{\mathbf{z}}_c \leftarrow \bar{\mathbf{K}}_c^{-1} \bar{\mathbf{r}}_c$  (* Inner coarse loop: solved on FEMmodelc by
Algorithm 2 with  $\epsilon_c^{in}$  and initial solution  $\bar{\mathbf{z}}_c$  [using GPU] *)
16:    $\bar{\mathbf{z}} \leftarrow \bar{\mathbf{P}} \bar{\mathbf{z}}_c$ 
17:    $\bar{\mathbf{z}} \leftarrow \bar{\mathbf{K}}^{-1} \bar{\mathbf{r}}$  (* Inner fine loop: solved on FEMmodel by
Algorithm 2 with  $\epsilon^{in}$  and initial solution  $\bar{\mathbf{z}}$  [using GPU] *)
18:    $\mathbf{z} \leftarrow \bar{\mathbf{z}}$ 
19:   (*inner loop end*)
20:   if  $i > 1$  then
21:      $\beta \leftarrow (\mathbf{z}, \mathbf{q}) / \rho$ 
22:   end if
23:    $\mathbf{p} \leftarrow \mathbf{z} + \beta \mathbf{p}$ 
24:    $\mathbf{q} \leftarrow \mathbf{K} \mathbf{p}$ 
25:    $\rho \leftarrow (\mathbf{z}, \mathbf{r})$ 
26:    $\alpha \leftarrow \rho / (\mathbf{p}, \mathbf{q})$ 
27:    $\mathbf{q} \leftarrow -\alpha \mathbf{q}$ 
28:    $\mathbf{r} \leftarrow \mathbf{r} + \mathbf{q}$ 
29:    $\mathbf{x} \leftarrow \mathbf{x} + \alpha \mathbf{p}$ 
30:    $i \leftarrow i + 1$ 
31: end while
32: (*outer loop end*)
33:  $\mathbf{u} \leftarrow \mathbf{x}$ 
34: output results

```

the use of high-performance single-precision compute hardware used in most of the current HPC systems.

With Algorithm 1, we can move most of the compute cost of the outer loop to the inner loops by using appropriate threshold values ϵ_c^{in} and ϵ^{in} in the inner CG solvers. Thus, even though this algorithm may seem complicated, the performance of the whole solver is dependent on the matrix-vector multiplication kernel, inner product kernel, or saxpy kernels, which are commonly used with standard CG solvers. Below, we explain the matrix-vector multiplication kernel used in the inner fine loop, which is the most time-consuming kernel that requires algorithmic considerations for efficient computation on GPUs.

C. Design of Element-by-Element Kernel for K computer

1) *Element-by-Element Computation* [7]: As the relative memory transfer capability to floating point computation ca-

Algorithm 2 Preconditioned conjugate gradient method [computed on GPU]. This is used to roughly solve $\mathbf{z} = \mathbf{K}^{-1} \mathbf{r}$ and $\mathbf{z}_c = \mathbf{K}_c^{-1} \mathbf{r}_c$ in Algorithm 1. The matrix-vector multiplication $\mathbf{K}\mathbf{y}$ is computed using an Element-by-Element method with voxel: $\sum_i^{N^V} \mathbf{K}_i^V \mathbf{y}_i^V + \sum_i^{N^T} \mathbf{K}_i^T \mathbf{y}_i^T$, where $(\)^V$ indicates values related to merged voxel elements, $(\)^T$ indicates values related to tetrahedral elements, and the subscript i indicates the i -th element. \mathbf{B} and ϵ^{in} are a 3×3 block Jacobi matrix and the tolerance for the relative error, respectively. All of the calculations are made using single-precision variables.

```

1:  $\mathbf{x} \leftarrow \mathbf{z}$  (* CPU to GPU transfer *)
2:  $\mathbf{e} \leftarrow \mathbf{r} - \mathbf{K} \mathbf{x}$  (* update boundary by CPU-GPU transfer *)
3:  $\beta \leftarrow 0$ 
4:  $i \leftarrow 1$ 
5:  $\mathbf{z\_or\_q} \leftarrow \mathbf{B}^{-1} \mathbf{e}$ 
6:  $\rho_a \leftarrow (\mathbf{z\_or\_q}, \mathbf{e})$  (* CPU-GPU transfer for allreduce *)
7:  $r2 \leftarrow \|\mathbf{r}\|_2$  (* CPU-GPU transfer for allreduce *)
8:  $err \leftarrow \|\mathbf{e}\|_2 / r2$ 
9: while  $err \geq \epsilon^{in}$  do
10:    $\mathbf{p} \leftarrow \mathbf{z\_or\_q} + \beta \mathbf{p}$ 
11:    $\mathbf{z\_or\_q} \leftarrow \mathbf{K} \mathbf{p}$  (* update boundary by CPU-GPU transfer *)
12:    $\alpha \leftarrow \rho_a / (\mathbf{p}, \mathbf{z\_or\_q})$  (* CPU-GPU transfer for allreduce *)
13:    $\rho_b \leftarrow \rho_a$ 
14:    $\mathbf{e} \leftarrow \mathbf{e} - \alpha \mathbf{z\_or\_q}$ 
15:    $\mathbf{x} \leftarrow \mathbf{x} + \alpha \mathbf{p}$ 
16:    $\mathbf{z\_or\_q} \leftarrow \mathbf{B}^{-1} \mathbf{e}$ 
17:    $\rho_a \leftarrow (\mathbf{z\_or\_q}, \mathbf{e})$  (* CPU-GPU transfer for allreduce *)
18:    $err \leftarrow \|\mathbf{e}\|_2 / r2$  (* CPU-GPU transfer for allreduce *)
19:    $\beta \leftarrow \rho_a / \rho_b$ 
20:    $i \leftarrow i + 1$ 
21: end while
22:  $\mathbf{z} \leftarrow \mathbf{x}$  (* GPU to CPU transfer *)

```

pability is decreasing, using an algorithm that can reduce memory access is becoming important for short time-to-solution. At the same time, low memory footprint solvers are desirable since the memory capacity per compute capability is decreasing. Thus we use element-by-element (EBE) computation, which is a matrix-free matrix-vector multiplication method, to compute $f = \mathbf{K}u$. Here, u and f are displacement and nodal force vectors, respectively, and \mathbf{K} is the global stiffness matrix generated by superimposing element stiffness matrices K_i^e . K_i^e can be computed using coordinates of nodes x and element material properties. In standard matrix-vector multiplication methods, $f = \mathbf{K}u$ is computed by reading the precomputed global stiffness matrix \mathbf{K} from the memory and multiplying it by u . In contrast, in the EBE computation, matrix-vector multiplications are computed by computing the local matrix-vector multiplications

$$f_i^e = K_i^e u_i^e = K_i^e Q_i^e u, \quad (2)$$

and summing them as

$$f = \sum_i Q_i^e f_i^e. \quad (3)$$

Here, Q_i^e are matrices for mapping local nodal values to global nodal values. Since the coordinates, displacement, and force vectors x, u, f can be kept in cache by reordering nodes and elements, we can drastically reduce the memory access compared to methods that read the global matrix from the

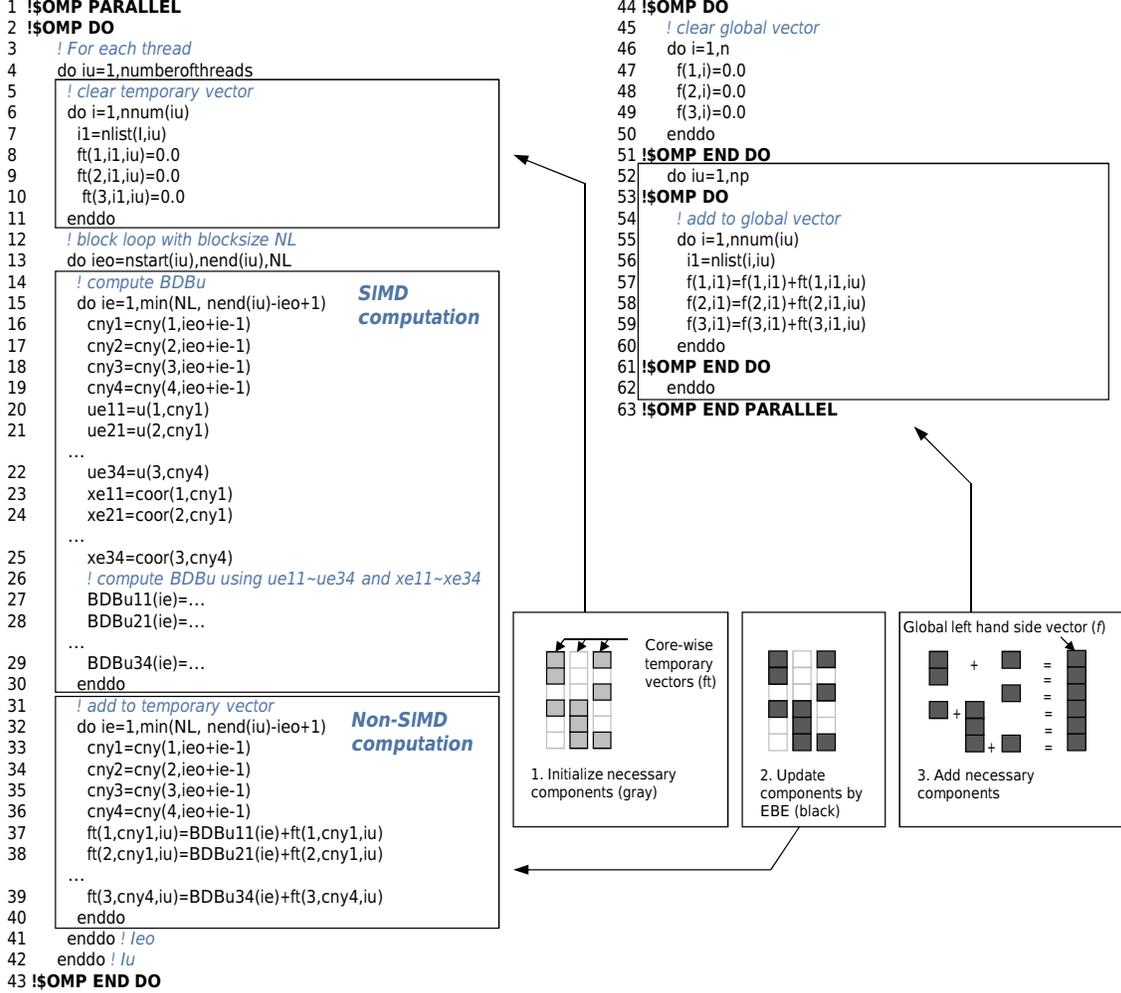


Fig. 1. Parallelization of EBE kernel on a CPU (a linear tetrahedral element case is shown). Data races among multi-cores are avoided by using a temporary vector f_i^{core} (denoted ft). Only parts of the temporary vector updated by each core are initialized and added to the global vector. SIMD parallelization is applied to element loop by loop splitting and loop blocking. Temporary buffer BDBu11-34 stays in cache as a small block size (i.e., NL) is used.

memory. On the other hand, the EBE computation involves more computation, and consequently, transfers memory access cost to computation cost.

2) *Parallel Computation of Element-by-Element Kernel:* As seen above, the EBE method has low algorithmic byte/FLOP and is suitable for current computers. However, it is not straightforward for parallel computation due to the data recurrence required for adding local force vectors of elements with shared nodes in Eq. (3). Thus, we have developed an algorithm for parallel computation on multi-core and SIMD-based CPUs (Fig. 1).

- Multi-core parallelization. We allocate and initialize temporary left-hand side vectors f_j^{core} (Fig. 1, lines 5-11), sum element-wise results for each core in parallel (Fig. 1, lines 31-40), and sum core-wise results into the global vector f (Fig. 1, lines 52-62).

- SIMD parallelization. Data recurrence may occur for elements allocated to each SIMD lane of compute cores. Thus, we split the innermost loop into two loops, i.e., the computation part of $f_i^e = K_i^e u_i^e$ (Fig. 1, lines 14-30) and the addition part of f_i^e to f_j^{core} (Fig. 1, lines 31-40). This enables SIMD computation of the computationally rich first loop. We reduce the loop splitting overhead by blocking the loop with a small block size (NL), which maintains the temporary buffers BDBu11-34 in cache, which reduces the cost involved in loop splitting.

3) *Mixed Structured/Unstructured Computation of Element-by-Element Kernel:* On the K computer, the bottleneck of multi-core and SIMD-enabled EBE kernels becomes random register-L1 cache access (Fig. 1, lines 20-25 and 37-39). Thus, we have developed an integrated modeling-analysis environment for problem solving to reduce random access. Here, structured voxel elements are used to model homogeneous

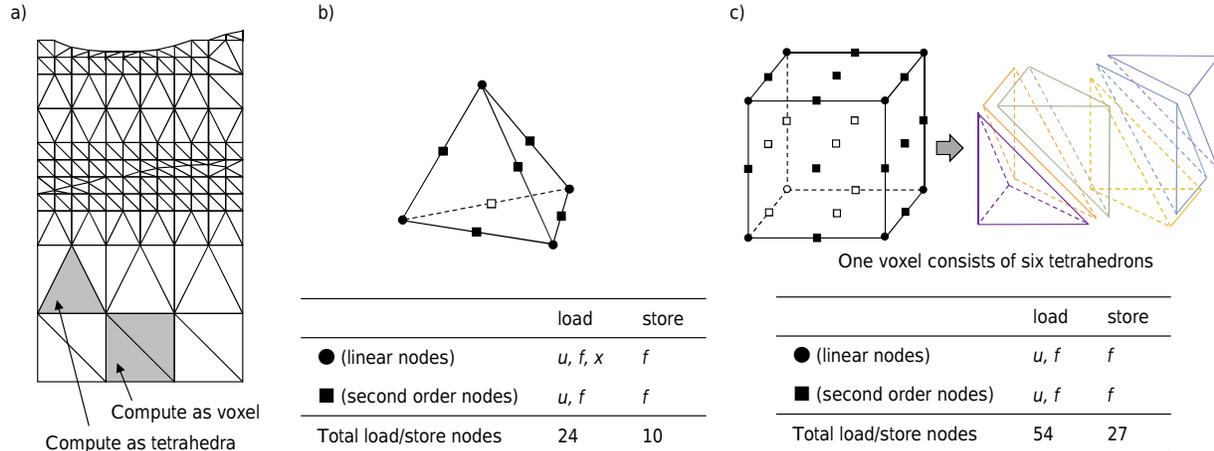


Fig. 2. Use of mixed structured/unstructured mesh for EBE computation in finite-element analysis. a) Generated mesh for two-layered crust structure. Voxels are used in homogeneous parts of the domain, and decomposed into six tetrahedrons with fixed geometry. Unstructured tetrahedral elements are used for the remaining parts of the domain with complex geometry. b) Ten nodes are accessed during EBE computation of a tetrahedron element. c) Twenty-seven nodes are accessed during EBE computation of a voxel. Since each voxel is equivalent to six tetrahedrons, random access per tetrahedron can be reduced by $27/(6 \times 10) = 0.45$ times.

parts of the target domain, while unstructured tetrahedral elements are used to model the remaining complex parts (Fig. 2a). Since the voxel elements are equivalent to six tetrahedrons in a pure unstructured mesh, we can halve the amount of random access compared to a pure unstructured mesh by computing them together. In addition, because the geometry is fixed, the FLOP per element required to compute the local matrix-vector multiplication is reduced. Thus, the same matrix-vector multiplication results can be computed with less random access and less computation cost using mixed structured/unstructured computation. As the gap between random access performance and computation performance widens on GPUs, we can expect effective speedup by using mixed structured/unstructured algorithms on GPUs.

III. ACCELERATING FINITE-ELEMENT EARTHQUAKE SIMULATION ON GPU

A. Solver Overview

We port the solver, both with and without mixed structured/unstructured computation, to a CPU-GPU heterogeneous compute environment. Most of the compute time is algorithmically designed to be spent on the inner fine and inner coarse loops; therefore, we can expect drastic reduction of overall compute time by applying a GPU to these preconditioners. Performance for double-precision computation is generally lower than that for single-precision computation on GPUs, and thus, the speedup of double-precision computation is small considering the cost of using large amounts of device memory. On the K computer, only 11% of the overall elapsed time is spent for this double-precision computation part; therefore, we compute the outer CG loop on the CPU. Thus, part of Algorithm 1 (indicated as [using GPU]) and all parts of Algorithm 2 are offloaded to the GPU.

Since both MPI communication and CPU-GPU transfer are slow, these data transfers become bottlenecks when porting applications to GPUs. Thus, reducing the amount and frequency of data transfer becomes important to obtain high performance with GPU-CPU heterogeneous compute environments. In the default OpenACC settings, CPU-GPU data transfers on slow PCI Express are called before and after GPU kernel calls. We first use OpenACC directives, which explicitly designate data transfer such that only a minimum amount of data is transferred between the CPU and GPU (Algorithm 2 shows the data transfer parts). On the other hand, communication for synchronizing domain boundaries during matrix-vector multiplication and computing allreduce operations during vector inner product operations must be passed between GPUs in the CG solvers. Since some GPU-CPU heterogeneous environments do not have direct communication capability between GPUs, we use CPU-GPU transfer via PCI Express and MPI communication to increase the range of available compute environments. Reduction of communication size during domain boundary synchronization has been implemented in the CPU code; nodal values required for synchronization are packed into a buffer before it is sent/received through an MPI function call. Thus, we can obtain coalesced data transfer with minimal size by implementing OpenACC directives for such buffer packing/unpacking operations and setting CPU-GPU data transfer directives before and after MPI calls.

The primary GPU computations becomes vector inner products, vector addition, vector multiplication, and matrix-vector multiplication. The vector inner product, vector addition, and vector multiplication kernels, which are memory bandwidth-limited kernels with simple loop structures, are ported directly by adding OpenACC directives to the Fortran77 (and partly Fortran90) CPU code. To instruct parallel computation,

```

a) Coloring add
1  !$ACC PARALLEL LOOP
2  ! clear global vector
3  do i=1,n
4  f(1,i)=0.0
5  f(2,i)=0.0
6  f(3,i)=0.0
7  enddo
8  do icolor=1,ncolor
9  !$ACC PARALLEL LOOP
10 do ie=ns(icolor),ne(icolor)
11 ! compute BDBu
12 cny1=cny(1,ie)
13 cny2=cny(2,ie)
14 cny3=cny(3,ie)
15 cny4=cny(4,ie)
16 ue11=u(1,cny1)
17 ue21=u(2,cny1)
18 ...
19 ue34=u(3,cny4)
20 xe11=coor(1,cny1)
21 xe21=coor(2,cny1)
22 ...
23 xe34=coor(3,cny4)
24 ! compute BDBu using ue11~ue34 and xe11~xe34
25 BDBu11=...
26 BDBu21=...
27 ...
28 BDBu34=...
29 ! add to global vector
30 f(1,cny1)=BDBu11+f(1,cny1)
31 f(2,cny1)=BDBu21+f(2,cny1)
32 ...
33 f(3,cny4)=BDBu34+f(3,cny4)
34 enddo
35 enddo

b) Atomic add
1  !$ACC PARALLEL LOOP
2  ! clear global vector
3  do i=1,n
4  f(1,i)=0.0
5  f(2,i)=0.0
6  f(3,i)=0.0
7  enddo
8  !$ACC PARALLEL LOOP
9  do ie=1,ne
10 ! compute BDBu
11 cny1=cny(1,ie)
12 cny2=cny(2,ie)
13 cny3=cny(3,ie)
14 cny4=cny(4,ie)
15 ue11=u(1,cny1)
16 ue21=u(2,cny1)
17 ...
18 ue34=u(3,cny4)
19 xe11=coor(1,cny1)
20 xe21=coor(2,cny1)
21 ...
22 xe34=coor(3,cny4)
23 ! compute BDBu using ue11~ue34 and xe11~xe34
24 BDBu11=...
25 BDBu21=...
26 ...
27 BDBu34=...
28 ! add to global vector
29 !$ACC ATOMIC
30 f(1,cny1)=BDBu11+f(1,cny1)
31 !$ACC ATOMIC
32 f(2,cny1)=BDBu21+f(2,cny1)
33 ...
34 !$ACC ATOMIC
35 f(3,cny4)=BDBu34+f(3,cny4)
36 enddo

```

Fig. 3. Parallelization of EBE kernel on GPU (a linear tetrahedral element case is shown). a) **Coloring add**: Elements are divided into groups with non-overlapping nodes. Threads update elements in each color simultaneously, such that data race does not occur when adding local results $Q_i^e f_i^e$ to global left-hand side vector f . b) **Atomic add**: Atomic operations are used to avoid data race. Data-locality can be exploited by using on-chip cache in this case.

OpenACC requires three classes that are called gang, worker, and vector, respectively. Regarding NVIDIA GPUs, one gang corresponds to one block and one vector corresponds to one thread, so we need to specify only the number of gangs and the length of vectors to conduct parallel computation. The porting cost for this part is a few hours; thus, actual applications can be ported with small cost by using OpenACC. For the complex matrix-vector multiplication kernel, we modify the detailed algorithm as shown in the next subsection.

B. Design of Element-by-Element Kernel for GPU

Since the GPU is designed to hide memory latency by launching many threads on each of its 10^3 physical cores, magnitudes of more threads are used for computation compared to the CPU. In the case of the EBE kernel, we launch threads for each element in the mesh, and conduct local matrix-vector multiplication in parallel. Although such computation is suitable for parallel processing, allocating temporary arrays for each thread for addition of local matrix-vector multiplication results becomes infeasible when considering the number of threads and the device memory capacity. Thus, we must change the algorithm to add the local matrix-vector multiplication results to a global result vector in EBE kernels on GPUs. Since the performance of cache and atomics to global device

memory has been relatively low until the development of recent GPUs, GPU computation of EBE-based finite-element analyses has often been conducted using coloring, which is an algorithm that does not depend on cache and atomics [8]. During coloring, elements that do not share nodes are grouped, and each group of elements is computed in parallel (Fig. 3a). On the other hand, cache and atomic hardware has improved in recent GPUs, e.g., NVIDIA Pascal architecture. Thus, in this study, we use an algorithm capable of taking advantage of these improved hardware capabilities. Here, we launch threads for the element loop i and use atomics when adding $Q_i^e f_i^e$ to the global result vector f (Fig. 3b). Although random data access is inevitable in unstructured mesh computation, its cost is reduced by taking advantage of the data locality in the element node connectivity by effective use of cache. At the same time, data access for element material properties and node connectivity information is coalesced; we expect hiding memory access latency by using a sufficient number of threads. Since the generated threads are computed automatically in parallel by SIMT on GPUs, the loop splitting and loop blocking used in the CPU algorithm are not required.

TABLE I
COMPARISON OF HARDWARE CAPABILITIES OF K COMPUTER, GPU CLUSTER, AND NVIDIA DGX-1.

	K computer	K40 GPU cluster	NVIDIA DGX-1 (P100)
# of nodes	8	8	1
CPU/node	1 x eight-core SPARC64 VIIIfx	2 x twelve-core Intel Xeon E5-2695 v2	2 x twenty-core Intel Xeon E5-2698 v4
Accelerator/node	-	1 x NVIDIA K40	8 x NVIDIA P100
# of MPI processes/node	1	1	8
hardware peak FLOPS/process	128 GFLOPS	4.29 TFLOPS (GPU only)	10.6 TFLOPS (GPU only)
Bandwidth/process	64 GB/s	288 GB/s (GPU only)	720 GB/s (GPU only)
Interconnect	Tofu (4 lanes x 5GB/s in both directions)	InfiniBand FDR + PCI Express 3.0 x 16 (16GB/s in both directions)	4 x InfiniBand EDR + NVLink (20 GB/s)
Compiler	Fujitsu Fortran Driver Version 1.2.0	PGI compiler 16.7 -ta=tesla:cc3+,loadcache:L1 -acc -memodel=medium	PGI compiler 16.7 -ta=tesla:cc60,loadcache:L1 -acc -memodel=medium
Compiler option		-Minline=levels:10 -Mvect=simd:256 -Mipa=fast -fastsse -O3 -Mcuda=cuda8.0	-Minline=levels:10 -Mvect=simd:256 -Mipa=fast -fastsse -O3
MPI	custom MPI	OpenMPI 1.8.3	OpenMPI 1.10.2

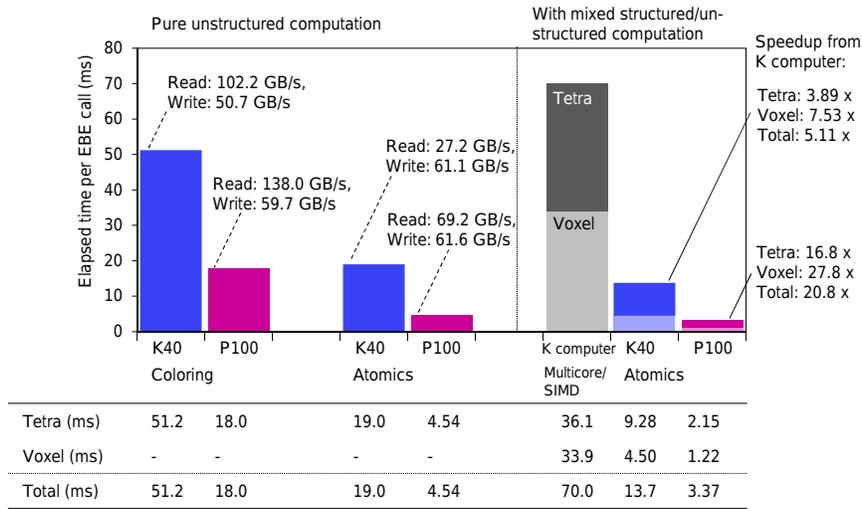


Fig. 4. Performance of EBE kernel. Insets indicate device memory access throughput of kernel. Here, one node of K computer with single SPARC64 VIII fx CPU, one K40 GPU, and one P100 GPU is compared.

C. Performance Measurement

We measure the performance of the developed EBE kernel on a single GPU. Here, we use NVIDIA K40 [9] and NVIDIA P100 [5] GPUs. Table I summarizes the compute environment specifications. The memory bandwidth of a K40 GPU is 288 GB/s, with peak single-precision compute performance of 4.29 TFLOPS. The memory bandwidth of a P100 GPU is 720 GB/s, with peak single-precision compute performance of 10.6 TFLOPS. Since these GPUs have read-only cache, L1 cache, and L2 cache, using this fast memory considering data locality is expected to lead to fast computation.

First, we measure the performance of an EBE kernel for unstructured second-order tetrahedral elements. Here, a problem with 10,427,823 degrees-of-freedom and 2,519,867 tetrahedral elements is used. Figure 4 shows the elapsed time for a single

matrix-vector multiplication using the coloring version and atomic version, both implemented using OpenACC. As can be seen in Fig. 4, the atomics version is faster than the coloring version for both K40 and P100 GPUs. This could be due to the high data locality of nodal values in the atomics version, which leads to high cache utilization and less memory access. The total amount of memory read from the device memory is reduced from 2.48 GB to 0.31 GB by switching from the coloring version to the atomics version for the P100 GPU case. Thus, it is evident that more data are reused in the atomics version. When comparing the P100/K40 performance ratio for the atomics version, we can see a 4.2 times speedup, which is higher than the increase in theoretical peak FLOPS ratio of 2.5. This could be due to the improvement in atomics by the new native 32-bit atomic add instructions or the increase in

cache capacity [5]. Another reason may be the increase in the number of registers per CUDA core. This means that more threads can be executed simultaneously on each CUDA core, expecting to lead to better coverage of memory access latency. Although we try changing the number of threads per block or swapping the dimensions of arrays to achieve additional performance improvement, we find that the algorithmic change between coloring and atomics is the most effective. Details of the performance change due to the change in the number of threads per block and array dimension swapping are given in the Appendix.

One concern for application developers is choosing between OpenACC and CUDA for implementing performance critical kernels on GPUs. On the K40 GPU, the performance of the atomics version implemented using CUDA Fortran is nearly the same as that of OpenACC (CUDA Fortran: 18.06 ms; OpenACC: 18.96 ms). Here, we use 32 threads per block, which perform the best for both implementations. As can be seen, the OpenACC directive-based implementation can perform as well as the more general CUDA Fortran, which has high development cost, for this complex EBE kernel with a large loop body (504 lines) and many temporary variables (216 variables). In the subsequent sections, we will measure performance based on OpenACC implementation with atomic addition.

Next, we measure the effectiveness of mixed structured/unstructured computation. The models used in the previous measurements correspond to a mixed problem with 204,185 structured voxels and 1,294,757 unstructured tetrahedral elements. Figure 4 shows the elapsed time for each method using atomics. For the P100 GPU case, the elapsed time for the tetrahedral part is 2.15 ms and the elapsed time for the voxel part is 1.22 ms (total of 3.37 ms). Since one voxel is equivalent to six tetrahedral elements, structured computations results in 1.81 times speedup per tetrahedral element. The speedup ratio for the K40 GPU case is also approximately the same. Thus, it is evident that mixed structured/unstructured computation is effective for a variety of architectures.

Last, we compare the performance of the mixed structured/unstructured EBE kernel with one node of the K computer (Table I). The single-precision compute performance and memory bandwidth capability of the K40 GPU is 33.5 and 4.5 times that of the K computer, while the single-precision compute performance and memory bandwidth capability of P100 GPU is 82.8 and 11.3 times that of the K computer, respectively. On the K computer, the elapsed time for the tetrahedral part is 36.1 ms and the elapsed time for the voxel part is 33.9 ms, which totals to 70.0 ms (Fig. 4). Thus, the EBE kernels are faster by 16.8 times and 27.8 times for the tetrahedral and voxel kernels on the P100 GPU, respectively. Although this is less than the hardware peak performance ratio of 82.8 times, we consider this a decent performance because a highly tuned K computer implementation is used as a baseline for this complex kernel, which naturally suits CPUs. Indeed, the CPU implementation on the K computer obtains performance of 20.6% of hardware FLOPS peak for the total

EBE kernel, which is considered high for an unstructured, low-order finite-element method. When comparing the speedup for tetrahedral and voxel kernels, it can be seen that the speedup for the voxel kernel is greater than that for the tetrahedral kernel. This is consistent with our expectation that the reduction in random access due to structure computing will be more effective on the GPU. We can see that the proper choice of algorithm is leading to large speedup.

IV. APPLICATION EXAMPLE ON GPU CLUSTER

Estimating the distribution of earthquake fault slip is important for analyzing past earthquakes and predicting future earthquakes. Although conducting 10^2 to 10^3 simulations of 10^8 degrees-of-freedom high-resolution finite-element crust-deformation simulations are required for precise estimation of underground fault-slip distribution from the data observed at the surface, the huge computational cost involved in such simulations has hindered realization of such estimations. In this section, we show the effectiveness of the accelerated earthquake simulation method by conducting many-case high-resolution crustal-deformation analyses reflecting the available crust data, and use these results to estimate the slip distribution of the 2011 Tohoku Earthquake. Here, we model a 784 km (EW, x) \times 976 km (NS, y) \times 400 km (UD, z) region that includes the fault plane and observation points with a finite-element model at a 2000 m resolution (Fig. 5 and Fig. 6). Figures 6a and 6b show the generated 82,196,106 degrees-of-freedom and 19,921,530 second-order tetrahedral finite-element model. The material properties of this four-layered model are set according to a previous study [11] (Fig. 6d). This model is subdivided into eight partitions using METIS [12].

First we show the performance of the ported solver using the developed EBE kernels explained in the previous sections, for one case of a crustal-deformation simulation on this finite-element model. We compare eight nodes of the K computer with an eight-node GPU cluster with a single K40 GPU per node and the NVIDIA DGX-1 [13] with eight P100 GPUs (Table I). First we compare the performance of the inner fine loop (Table II). The speedup for the non-memory bandwidth bound EBE tetra kernel is 18.9 times and that for the EBE voxel kernel is 28.9 times for DGX-1 (23.1 times speedup for total EBE kernel). These speedups are similar to the speedups obtained in the previous section; we can see that the developed EBE kernels attains performance on actual application problems with complex mesh. For the memory bandwidth bound inner product kernel, near-linear speedup with increased hardware memory bandwidth is obtained by both K40 GPU cluster and DGX-1. Thus, the OpenACC implementations enable increased performance relative to the hardware improvements for this kernel with both compute environments. These increase in performance enables a 16.7 times speedup for the whole inner fine loop. As the inner fine loop is a standard CG finite-element solver with 3×3 block Jacobi preconditioning, by using the developed EBE kernel, we expect a similar speedup for solvers that are widely

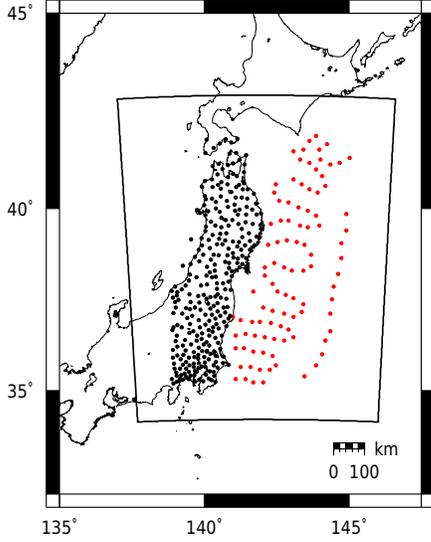


Fig. 5. Target area of application example (black line) and location of observation points. The analysis region is almost identical as that of a previous report [10]. x , y , and z components of GEONET (black points) and z component of S-Net (red points) are used for fault-slip estimation in this study.

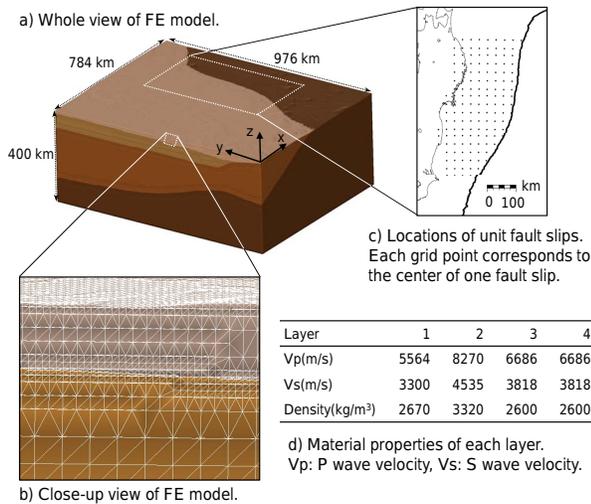


Fig. 6. Finite-element (FE) model and location of unit fault slips used in the application example.

used in a variety of finite-element applications. Although not implemented at present, we can overlap EBE kernel computation and MPI communication by first computing the elements connected to inter-MPI process boundary nodes, and conducting CPU-GPU data transfer and MPI communication while computing the elements in the inner domain. This is expected to hide costs of CPU-GPU data transfer and MPI communication and lead to further performance improvement of the solver.

Next we compare the performance of the whole solver with

multi-grid and outer loops. Although a decent speedup of 16.0-23.7 times is obtained for the EBE kernels of inner coarse loops on DGX-1, the speedup of the whole inner coarse loop was limited to 10.3 times. This may be due to the latency of MPI communication or CPU-GPU data transfer because the degrees-of-freedom computed in each loop are reduced to one-eighth when compared to the inner fine loop. Performance of this part may also benefit from overlapping of data transfer and communication with computation in future works. Since only 1.18 GB out of 16 GB device memory is used on each P100 GPU in this case, we can expect further performance improvement by increasing the problem size per process. Since fewer CPU-cores are used per process on DGX-1, the elapsed time for outer loop is slightly longer on DGX-1 in comparison with K computer and K40 GPU cluster. Thus, the total speedup of the solver is 3.49 and 5.10 times for the K40 GPU cluster and DGX-1, respectively. As the peak performance of double-precision computation on P100 is improved from previous generations of NVIDIA GPUs, we can expect a large speedup of the outer loop by porting to a GPU. A conservative estimate of speedup of the outer loop running on a DGX-1 is eight times, which will lead to at least 11 times speedup of the whole solver.

Using this accelerated finite-element application on the eight-node K40 GPU cluster, we conducted 360 crust-deformation simulations to obtain Green's functions for unit slip in the x and y directions for 180 grid points, as shown in Fig. 6c. The total elapsed time for solving these Green's functions is 4,731 s. This corresponds to a speedup of 186 times compared to a previous study [14], which computed a single case of 157,897,032 degrees-of-freedom (approximately 1.9 greater than the present study) crustal deformation problem in 4,746 s using two Intel Xeon X5680 CPUs. Figure 7 shows the estimated fault-slip distribution using the computed Green's functions and observed data at the surface. [15] and [16] noted that simplification of the three-dimensional heterogeneity of crustal structure possibly has a significant impact on the analysis results in some cases; we can expect higher reliability in estimation results by using Green's functions that better reflects the crust structure. Since we can compute crust-deformation simulations in a very short time, we enable to introduce methods which evaluate uncertainties quantitatively by conducting multiple computation cases (e.g., Monte-Carlo method). We expect a more reliable slip distribution estimation considering uncertainties in crust and observation data through 10^6 computations within feasible time by using the proposed method on leading supercomputer systems.

V. RELATED WORK

Many studies into sparse matrix-vector multiplications on GPUs have been reported. For example, NVIDIA is developing and supplying the cuSPARSE library [17], which stores global matrices in formats suitable for GPUs and computes matrix-vector multiplications efficiently. The method used in this study targets a larger degrees-of-freedom problem per GPU; thus, a matrix-free method, i.e., EBE computation, is used.

TABLE II
COMPARISON OF ELAPSED TIME OF SOLVER FOR K COMPUTER, K40 GPU CLUSTER, AND NVIDIA DGX-1 (MEAN VALUES OF ALL RANKS ARE INDICATED).

	Total number of iterations	K computer			K40 GPU cluster		NVIDIA DGX-1	
		Elapsed time	FLOPS/peak	Memory bandwidth utilization	Elapsed time	Speedup from K computer	Elapsed time	Speedup from K computer
Inner fine loop	126	16.4 s	11.7%	25.2%	3.02 s	5.4	0.98 s	16.7
(EBE tetra kernel)	126	3.97 s	28.6%	15.8%	1.11 s	3.6	0.211 s	18.9
(EBE voxel kernel)	126	4.51 s	11.9%	22.4%	0.82 s	5.5	0.156 s	28.9
(Inner product kernel)	126	0.63 s	10.9%	77.7%	0.16 s	3.9	0.066 s	9.5
Inner coarse loop	1387	28.5 s	13.5%	23.0%	7.02 s	4.1	2.78 s	10.3
(EBE tetra kernel)	1387	11.7 s	23.5%	12.8%	1.85 s	6.3	0.731 s	16.0
(EBE voxel kernel)	1387	7.49 s	9.48%	19.4%	0.90 s	8.3	0.315 s	23.7
(Inner product kernel)	1387	1.38 s	7.40%	66.0%	0.27 s	5.1	0.125 s	11.0
Outer loop [computed on CPU in all systems]	6	5.60 s	3.59%	15.7%	4.43 s	1.3	6.23 s	0.91
Total of solver	-	50.5 s	11.8%	22.9%	14.5 s	3.49	10.4 s	5.10

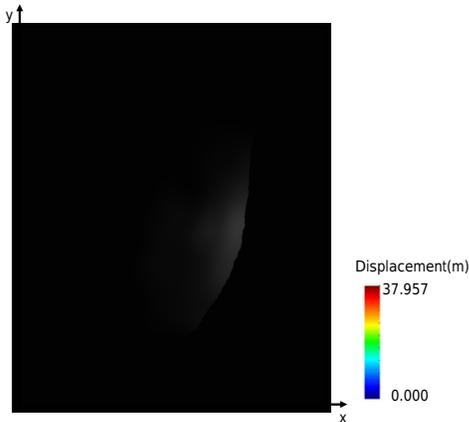


Fig. 7. Estimated co-seismic fault slip distribution (m).

Targeting GPUs with small cache and slow atomics, an EBE kernel that uses coloring to sum the element-wise results has been developed [8]. The present study can be considered an update targeting newer GPUs with improved cache and atomics. Another study [18] measured the performance of atomic/coloring addition in EBE computation on NVIDIA K20 GPUs. In this case, the elapsed time for coloring and atomics implementations were nearly the same. This is substantially different from the results of the present study, which can be due to the difference in the complexity of the mesh used. A structured mesh with eight colors has been used in [18], whereas the present study uses a three-dimensional unstructured mesh with up to 23 elements sharing one node, leading to 72 colors. It is evident that the atomics used in this study are more robust for complex three-dimensional problems.

The study by Gutzwiller et al. [19] is an example of porting structured mesh-based solvers to a GPU-CPU heterogeneous environment. By decomposing the target domain with a set of structured meshes, computation can be performed efficiently on GPUs. However, its applicability to complex shapes is

limited because it cannot be mixed with an unstructured mesh. By using a mesh generator capable of generating large mixed structured/unstructured meshes for complex shaped domains [20], the present method is capable of fast analysis for a wider range of problems. Thus, the present method is characterized as a fast finite-element solver with low memory footprint and wide applicability to general problems with complex shapes.

VI. CLOSING REMARKS

We accelerated the EBE kernel, which is the key kernel for attaining high-performance on unstructured implicit low-order finite-element solvers, by developing suitable algorithms for GPUs and porting to GPU-CPU heterogeneous environment by OpenACC. We also ported other parts of the CPU-based code by directly inserting OpenACC directives. When comparing eight K computer nodes and eight Pascal GPUs, we achieved 23.1 times speedup for the EBE kernel, which lead to 16.7 times speedup on a 3×3 block Jacobi preconditioned CG finite-element solver (i.e., the inner fine loop). This is expected to lead to at least 11 times speedup of whole solver. From these results, it is evident that by using OpenACC and minimal algorithmic development for performance-sensitive kernels, high performance can be achieved in a GPU-CPU heterogeneous compute environment with low development costs. Once suitable algorithms for GPUs were selected, we confirmed that complex kernels, such as the EBE kernel, can be ported to a GPU using OpenACC directives with performance that is nearly equivalent to that of CUDA Fortran implementations. By taking the same approach to design suitable EBE kernels considering the underlying hardware, we can expect good performance with other architectures, such as the Post K supercomputer [21] or Intel's Xeon Phi Knights Landing processors [22].

By using the developed method on a GPU cluster, we enabled many-case crust-deformation simulations within short time. Such acceleration in finite-element simulation using GPUs will improve the applicability of earthquake simulations to wider compute environments and also help improve the reliability of simulation results. It is also expected to accelerate

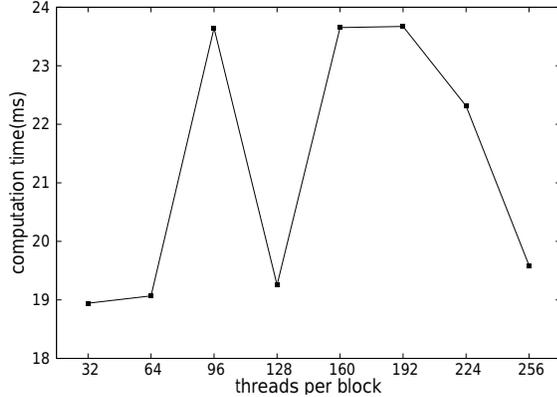


Fig. 8. Comparison of computation time by varying the number of threads per block (case for second-order tetrahedral EBE kernel computed using atomics).

other finite-element simulations used in the industry. In the future, we hope for advances in programming environments that can exploit automatic insertion of host-device data transfer and automatic tuning of performance-critical parameters (e.g., number of threads), which will help accelerate development and porting of high-performance applications.

ACKNOWLEDGMENTS

We thank Mr. Craig Toepfer (NVIDIA) and Mr. Yukihiro Hirano (NVIDIA) for the generous support concerning the use of NVIDIA DGX-1 (Pascal P100 GPU) environment. We also thank Mr. Shigeki Matsuo (NEC) for the generous support concerning the use of GPU environment. Part of the results were obtained using the K computer at the RIKEN Advanced Institute for Computational Science (Proposal numbers: hp160221, hp160160, and 160157). This work was supported by Post K computer project (priority issue 3: Development of Integrated Simulation Systems for Hazard and Disaster Induced by Earthquake and Tsunami), Japan Society for the Promotion of Science (KAKENHI Grant Numbers 15K18110, 26249066, 25220908) and FOCUS Establishing Supercomputing Center of Excellence.

APPENDIX

Here, we present some details of EBE kernel implementations for GPUs.

a) Number of Threads per Block: The number of simultaneously running threads on GPUs and register pressure can be adjusted by changing the number of threads per block. Here, we compare the performance of an atomics based unstructured second-order tetrahedral element EBE kernel with various numbers of threads per block. Since 32 threads are grouped into warps in NVIDIA GPUs for SIMT computation, we vary the number of threads by units of 32 threads. Figure 8 shows the performance on a K40 GPU. We can see that (divisors of 8) \times 32 threads obtain high performance as occupancy is obtained. Based on these observations, we use 32 threads per block for all EBE kernels in this study.

TABLE III
COMPARISON OF COMPUTATION TIME ON K40 GPU FOR DIFFERENT INDEX ORDERING OF TWO-DIMENSIONAL ARRAYS (CASE FOR MIXED STRUCTURED/UNSTRUCTURED EBE KERNEL COMPUTED USING ATOMICS).

Computation time(ms)	Tetra part	Voxel part
coor(3,n), etc.	9.28	4.50
coor(n,3), etc.	9.49	4.70

b) Effect of Array Dimension swapping: Table III shows the performance of array swapping for arrays coor, f, u in Fig. 3b. Here we measure the performance for an atomics-based mixed structured/unstructured inner fine EBE kernel on a K40 GPU. It is evident that ordering coor(3,n) is slightly better than ordering coor(n,3).

REFERENCES

- [1] T. Ichimura, K. Fujita, S. Tanaka, M. Hori, L. Maddeggedara, Y. Shizawa, and H. Kobayashi, "Physics-based urban earthquake simulation enhanced by 10.7 BlnDOF \times 30 K time-step unstructured FE non-linear seismic wave simulation," Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, (SC'14), pp. 15–26, 2014.
- [2] T. Ichimura, K. Fujita, P.E.B. Quinay, L. Maddeggedara, M. Hori, S. Tanaka, Y. Shizawa, H. Kobayashi, and K. Minami, "Implicit nonlinear wave simulation with 1.08T DOF and 0.270T unstructured finite elements to enhance comprehensive earthquake simulation," Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, (SC'15), 2015.
- [3] H. Miyazaki, Y. Kusano, N. Shinjou, F. Shoji, M. Yokokawa, and T. Watanabe, "Overview of the K computer system," FUJITSU Sci. Tech. J., vol. 48, no. 3, pp. 302–309, 2012.
- [4] OpenACC, [Online]. <http://www.openacc.org/>
- [5] NVIDIA Pascal GPU [Online]. <http://www.nvidia.com/object/tesla-p100.html>
- [6] G. H. Golub and Q. Ye, "Inexact conjugate gradient method with inner-outer iteration," SIAM, Journal on Scientific Computing, vol. 21, no. 4, pp. 1305–1320, 1997.
- [7] J. M. Winget and T. J. R. Hughes, "Solution algorithms for nonlinear transient heat conduction analysis employing element-by-element iterative strategies," Computer Methods in Applied Mechanics and Engineering, vol. 52, pp. 711–815, 1985.
- [8] I. Kiss, Z. Badics, S. Gyimothy and J. Pavo, "High locality and increased intra-node parallelism for solving finite element models on GPUs by novel element-by-element implementation," High Performance Extreme Computing (HPEC), 2012 IEEE Conference on, Waltham, MA, pp. 1–5, 2012.
- [9] NVIDIA Kepler GPU, [Online]. <http://www.nvidia.com/object/tesla-k80.html>
- [10] K. Koketsu, Y. Yokota, N. Nishimura, Y. Yagi, S. Miyazaki, K. Satake, Y. Fujii, H. Miyake, S. Sakai, T. Yamanaka, and T. Okada, "A unified source model for the 2011 Tohoku earthquake", *Earth and Planetary Science Letters*, vol. 310 (3–4), pp. 480–487, 2011.
- [11] M. Sato, N. Minagawa, M. Hyodo, T. Baba, T. Hori, and Y. Kaneda, "Effect of elastic inhomogeneity on the surface displacements in the north-eastern Japan: based on three-dimensional numerical modeling," *Earth Planets Space*, vol. 59, pp. 1083–1093, 2007.
- [12] METIS 5.1.0, [Online]. <http://glaros.dtc.umn.edu/gkhome/metis/metis/overview>
- [13] NVIDIA DGX-1, [Online]. www.nvidia.com/dgx1
- [14] T. Ichimura, R. Agata, T. Hori, K. Hirahara, and M. Hori, "Fast numerical simulation of crustal deformation using a three-dimensional high-fidelity model," *Geophysical Journal International*, vol. 195, no. 3, pp. 1730–1744, 2013.
- [15] T. Masterlark, "Finite element model predictions of static deformation from dislocation sources in a subduction zone: Sensitivities to homogeneous, isotropic, Poisson-solid, and half-space assumptions," *Journal of Geophysical Research*, vol. 108(B11), 2003.

- [16] K. Hughes, T. Masterlark, and W. Moonry, "Poroelastic stress-triggering of the 2005 M8.7 Nias earthquake by the 2004 M9.2 Sumatra-Andaman earthquake," *Earth planet. Sci. Lett.*, vol. 293, pp. 289–299, 2010.
- [17] The NVIDIA CUDA Sparse Matrix library (cuSPARSE), [Online]. <https://developer.nvidia.com/cuspars>
- [18] K. Ljungkvist, "Matrix-Free Finite-Element Operator Application on Graphics Processing Units," In *Euro-Par 2014: Parallel Processing Workshops*, volume 8806 of *Lecture Notes in Computer Science*, pp. 450–461, 2014.
- [19] D. Gutzwiller, R. Srinivasan, and A. Demeulenaere, "Acceleration of the FINE/Turbo CFD solver in a heterogeneous environment with OpenACC directives," In *Proceedings of the Second Workshop on Accelerator Programming using Directives (WACCPD '15)*. ACM, New York, NY, USA, Article 6, 2015.
- [20] K. Fujita, K. Katsushima, T. Ichimura, M. Hori, and L. Maddegada, "Octree-based Multiple-material Parallel Unstructured Mesh Generation Method for Seismic Response Analysis of Soil-Structure Systems," *Procedia Computer Science*, vol. 80, pp. 1624–1634, 2016.
- [21] Y. Ishikawa, "System software in post K supercomputer," SC15: International Conference for High Performance Computing, Networking, Storage and Analysis (Invited Talk), 2015.
- [22] HotChips 2015, Knights Landing: 2nd Generation Intel Xeon Phi Processor, [Online]. http://www.hotchips.org/wp-content/uploads/hc_archives/hc27/HC27.25-Tuesday-Epub/HC27.25.70-Processors-Epub/HC27.25.710-Knights-Landing-Sodani-Intel.pdf