

# A Step towards Energy Efficient Computing: Redesigning A Hydrodynamic Application on CPU-GPU

Tingxing Dong\*, Veselin Dobrev†, Tzanio Kolev†, Robert Rieben†, Stanimire Tomov\*, Jack Dongarra\*

\*Innovative Computing Laboratory, University of Tennessee, Knoxville

†Lawrence Livermore National Laboratory

\*tdong, tomov, dongarra@eecs.utk.edu

†dobrev1,kolev1,riebe1@llnl.gov

## ABSTRACT

Power and Energy consumption are becoming an increasing concern in high performance computing. Compared to multi-core CPUs, GPUs have a much better performance per watt. In this paper we redesigned the most computation intensive part of BLAST, an application that solves the equations for compressible hydrodynamics with high order FEMs [1] on GPUs. In order to exploit the hardware parallelism of GPUs and achieve high performance, we implemented custom linear algebra kernels. We intensively optimized our CUDA kernels by exploiting the memory hierarchy, which exceed the vendor's library routines substantially in performance. We proposed an autotuning technique to adapt our CUDA kernels to the orders of the FEMs. Compared to a previous base implementations, our redesign and optimization lowered the energy consumption of the GPU in two aspects: 60% less time and 10% less watts. Compared to the CPU-only solution, our GPU accelerated BLAST obtained a  $2.5\times$  overall speedup and  $1.42\times$  energy efficiency (greenup) for the high order method, and a  $1.9\times$  speedup and  $1.27\times$  greenup for the low order method.

## Keywords

GPU, hydrodynamics, Power, Energy, FEM

## 1. INTRODUCTION

High performance computing is increasingly becoming power and energy constrained. The average power of TOP 10 supercomputers climbed from 3.2MW in 2010 to 6.6MW in 2013, which is enough to power small towns [2]. DoE has recently set a goal of 20MW for exascale systems, which means 50 GFLOPS per watt; though the fact is No.1 supercomputer Tianhe-2 has already reached 17MW at 0.03 EFLOPS. Limited by the power budget, more and more computing clusters seek to install accelerators, such as GPUs, due to their high floating-point operation capability and energy efficiency advantage over CPUs, as in Figure 1. The trend of accelerated supercomputers is indicated in the latest June 2013 ranking of the TOP 500 and the Green 500. In the TOP 500, 51 are powered by GPUs[2]. Although accelerated systems make up only 10% of the systems, they accomplish 33% of the total computations. In the Green 500, the most efficient systems powered by K20 surpassed 3 GFLOPS per watt, up from 2 GFLOPS per watt in the June 2012 ranking[3].

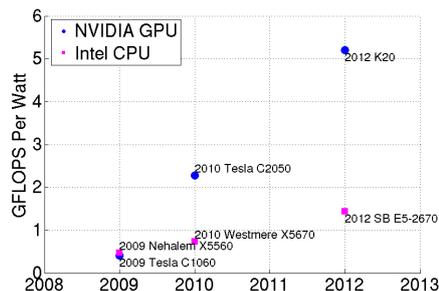


Figure 1: Performance per watt of NVIDIA GPUs versus Intel CPUs in double precision.

Energy consumption can be expressed as

$$Energy = Power \cdot Time$$

Reducing either factor on the right side of the equation will help to reduce energy consumption. Reducing processor power generally involves architecture redesign, such as changing hardware threading, power gating or decreasing frequency [4] [5]. Yet, according to the equation, merely reducing power, like decreasing the frequency of cores, does not necessarily save the energy, since it may lead to longer running time. A more meaningful metric is the performance per watt. For example, the Kepler GPU has better performance per watt and hence is more energy efficient - it doubles the performance of its previous generation, i.e. Fermi, while retaining the same watts [6]. Figure 1 shows the performance per watt of NVIDIA GPUs and that of Intel CPUs in double precision floating-point operations. Most users might be tempted to use TDP (Thermal Design Power) as the power of their application. However, TDP is only an engineering term and can be vastly different from the actual power used by applications[7]. A challenge for most users is that it is impractical to attach a power meter to monitor the power usage of their application. Fortunately, for Intel Sandy Bridge and NVIDIA Kepler users, monitoring and management of power is assisted by using RAPL and NVML, respectively[8][9]. With these tools, computing can be more power aware for users.

In this paper, we redesign a hydrodynamic code on CPU-GPU clusters. BLAST is a software package that solves the equations of compressible hydrodynamics using high order finite element methods (FEM). High order numerical

methods ( $p$ -refinement) and/or high resolution meshes ( $h$ -refinement) are introduced, to reveal more refined physical features. However, for a given number of degrees of freedom, high order methods are more computationally intensive than low order methods, as they couple with more of its neighbors on the mesh. Therefore, the higher the order of the method, the more FLOPs per memory access. The most floating point intensive part of BLAST is the "corner force" part which can take up to 55%-80% of the total run time, increasing with the order of the methods, but only 10% of the code. In our CPU-GPU solution, the FLOP intensive parts, including corner force computation are accelerated on the GPU, while other parts are still on the CPU. In order to exploit the hardware parallelism of GPUs and achieve high performance, we redesigned the CPU code with CUDA. Most of the target CPU code are processed into GPU-efficient parallel batched matrix operations, which are expressed by linear algebra routines(kernels).

Our contributions can be summarized as follows.

1) We reestablish the appeal of high order methods on GPUs. Compared to low order methods, the speedup and energy efficiency (greenup) of the high order FEM are greater, and simulation time grows more slowly.

2) We design custom linear algebra kernels on the GPU and intensively optimize these kernels. Our optimization results in less time and less power. The optimized kernels achieve a substantial improvement in performance compared to the widely used vendor's library routines.

3) We analyze the power and energy consumption of a real application on CPUs and GPUs. Compared to the CPU-only solution, we then present that our hybrid solution obtains both speedup and greenup.

4) We use CUDA and OpenMP inside MPI in order to exploit all the computing resources of multi-core CPU with Fermi GPU. To our best knowledge, this is the first time to apply all the three programming methods in hydrodynamics.

5) We present a good weak scaling on NO.2 supercomputer ORNL Titan to 4096 computing nodes.

## 2. THE BLAST ALGORITHM

The BLAST C++ code uses high order FEM in a moving Lagrangian frame to solve the Euler equations of compressible hydrodynamics. It supports 2D (triangles, quads) and 3D (tets, hexes) unstructured curvilinear meshes.

On a semi-discrete level, the conservation laws of Lagrangian hydrodynamics can be written as:

$$\text{Momentum Conservation: } \mathbf{M}_V \frac{d\mathbf{v}}{dt} = -\mathbf{F} \cdot \mathbf{1}, \quad (1)$$

$$\text{Energy Conservation: } \frac{d\mathbf{e}}{dt} = \mathbf{M}_E^{-1} \mathbf{F}^T \cdot \mathbf{v}, \quad (2)$$

$$\text{Equation of Motion: } \frac{d\mathbf{x}}{dt} = \mathbf{v}, \quad (3)$$

where  $\mathbf{v}$ ,  $\mathbf{e}$ , and  $\mathbf{x}$  are the unknown velocity, specific internal energy, and grid position, respectively. The kinematic mass matrix  $\mathbf{M}_V$  is the density weighted inner product of *continuous* kinematic basis functions and is therefore global, symmetric, and sparse. We solve the linear system of (1) using a preconditioned conjugate gradient (PCG) iterative method at each time step. The thermodynamic mass matrix  $\mathbf{M}_E$  is the density weighted inner product of *discontinuous* thermodynamic basis functions and is therefore symmetric and block diagonal, with each block consisting of a local dense

matrix. We solve the linear system of (2) by pre-computing the inverse of each local dense matrix at the beginning of a simulation and applying it at each time step using sparse linear algebra routines. The rectangular matrix  $\mathbf{F}$ , called the generalized *force matrix*, depends on the hydrodynamic state ( $\mathbf{v}$ ,  $\mathbf{e}$ ,  $\mathbf{x}$ ), and needs to be evaluated at every time step.

Evaluation of the matrix  $\mathbf{F}$ , which can be assembled from the generalized *corner force matrices*  $\{\mathbf{F}_z\}$  computed in every zone (or element) of the computational mesh. Evaluating  $\mathbf{F}_z$  is a locally FLOP-intensive process based on transforming each zone back to the reference element where we apply a quadrature rule with points  $\{\hat{q}_k\}$  and weights  $\{\alpha_k\}$ :

$$\begin{aligned} (\mathbf{F}_z)_{ij} &= \int_{\Omega_z(t)} (\sigma : \nabla \bar{w}_i) \phi_j \\ &\approx \sum_k \alpha_k \hat{\sigma}(\hat{q}_k) : \mathbf{J}_z^{-1}(\hat{q}_k) \hat{\nabla} \hat{w}_i(\hat{q}_k) \hat{\phi}_j(\hat{q}_k) |\mathbf{J}_z(\hat{q}_k)|. \end{aligned} \quad (4)$$

In the CPU code,  $\mathbf{F}$  is constructed by two loops: an outer loop over zones (for each  $z$ ) in the domain and an inner loop over the quadrature points (for each  $k$ ) in each zone. Each zone and quadrature point computes a component of the corner forces associated with it independently.

A local corner force matrix  $\mathbf{F}_z$  can be written as

$$\mathbf{F}_z = \mathbf{A}_z \mathbf{B}^T,$$

with

$$(\mathbf{A}_z)_{ik} = \alpha_k \hat{\sigma}(\hat{q}_k) : \mathbf{J}_z^{-1}(\hat{q}_k) \hat{\nabla} \hat{w}_i(\hat{q}_k) |\mathbf{J}_z(\hat{q}_k)|, \text{ where } \mathbf{J}_z = \hat{\nabla} \Phi_z \quad (5)$$

and

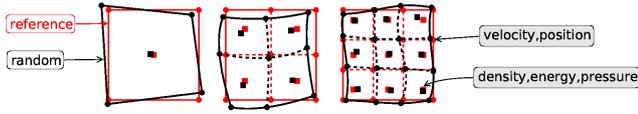
$$(\mathbf{B})_{jk} = \hat{\phi}_j(\hat{q}_k). \quad (6)$$

The matrix  $\mathbf{B}$  contains the values of the thermodynamic basis functions sampled at quadrature points on the reference element  $\hat{\phi}_j(\hat{q}_k)$  and is of dimension number of thermodynamic basis functions by number of quadrature points. The values stored in the matrix  $\mathbf{B}$  are constant in time. The matrix  $\mathbf{A}_z$  contains the values of the gradient of the kinematic basis functions sampled at quadrature points on the reference element  $\hat{\nabla} \hat{w}_i(\hat{q}_k)$  and is of dimension number of kinematic basis functions by number of quadrature points.

This matrix also contains terms which depend on the geometry of the current zone,  $z$ . Finite element zones are defined by a parametric mapping  $\Phi_z$  from a reference zone. The Jacobian matrix  $\mathbf{J}_z$  is non-singular and varies inside each zone. The determinant  $|\mathbf{J}_z|$  can be viewed as local volume. The stress value  $\hat{\sigma}(\hat{q}_k)$  require evaluation at each time step and involve significant amounts of computation including singular value decomposition (SVD), eigenvalue, eigenvector, equation of state (EOS) evaluations, etc., at each quadrature point (see [1] for more details).

A finite element solution is specified by the order of the kinematic and thermodynamic bases. In practice, we choose the order of the thermodynamic basis to be one less than the kinematic basis, where a particular method is designated as  $Q_k$ - $Q_{k-1}$ ,  $k \geq 1$ , corresponding to a continuous kinematic basis in the space  $Q_k$  and a discontinuous thermodynamic basis in the space  $Q_{k-1}$ . High order methods (as illustrated in Figure 2) can lead to better numerical approximations at the cost of more basis functions and quadrature points in

the evaluation of (2). By increasing the order of the finite element method,  $k$ , we can arbitrarily increase the floating point intensity of the corner force kernel of (2) as well as the overall algorithm of (1) - (3).



**Figure 2: Schematic depiction of bilinear ( $Q_1-Q_0$ ), biquadratic ( $Q_2-Q_1$ ), and bicubic ( $Q_3-Q_2$ ) zones.**

Here we summarize the basic steps of the overall BLAST MPI-based parallel algorithm:

- 1) Read mesh, material properties and input parameters;
- 2) Partition domain across MPI tasks and refine mesh;
- 3) Compute initial time step;
- 4) Loop over zones in the sub-domain of each MPI task:
  - (4.1) Loop over quadrature points in each zone;
  - (4.2) Compute corner force associated with each quadrature point and update time step;
- 5) find minimum time step and assemble zone contribution to global linear system;
- 6) Solve global linear system for new accelerations;
- 7) Update velocities, positions and internal energy;
- 8) Go to 4 if final time is not yet reached, otherwise exit.

Step 4 is associated with the corner force calculation of (2) which is a computational hot spot where we focus our effort. Step 6 solves the linear equation (using a simple PCG solver) of (1). Table 1 shows timing data for various high order methods in 2D and 3D. Both of the computational cost of corner force and CG solver increase as the order of the method  $k$  and dimension are increased, though the cost of corner force grows faster than that of CG solver.

Method	Corner Force	CG Solver	Total time
2D: $Q_4-Q_3$	198.6	53.6	262.7
2D: $Q_3-Q_2$	72.6	26.2	103.7
3D: $Q_2-Q_1$	90.0	56.7	164.0

**Table 1: Profile of BLAST on Xeon CPU: The corner force kernel consumes 55% – 75% of total time. The CG solver takes 20% – 34%.**

### 3. HYBRID PROGRAMMING MODEL

Multi-GPU communication relies on CPU-GPU communication on a single node and CPU-CPU communication across nodes. Therefore, a multi-GPU implementation requires CUDA to interact with other CPU programming models like MPI, OpenMP or Pthreads. Our implementation is composed of the following two layers of parallelism: (1) MPI-based parallel domain-partitioning and communication between CPU; (2) CUDA based parallel corner force calculation inside each MPI task.

#### 3.1 CUDA Implementation

We implemented the momentum (1) and energy (2) equations on the GPU. In the CUDA programming guide [10], the term host is used to refer to CPU and device to GPU. Hereafter in this paper, we follow this practice. The CUDA implementation is composed of the following set of kernels:

##### 3.1.1 CUDA Code Redesign

The CPU code loops over the points in each zone and performs operations on the variables, most of which are represented as matrices. The left of Figure 5 shows our base CUDA implementation. `kernel_loop_quadrature_point` is a kernel to unroll  $\mathbf{A}_z$  which loops on quadrature points. The kernel on Fermi is faster than a six core Westmere X5660 CPU. Yet, it is still inefficient and dominated most of the GPU time. We replaced it with six new designed kernels 1-6. A major difference between the CPU code and our newly designed CUDA code is that loops become batch-processed. Thus the challenge is to write GPU-efficient massively parallel batched matrix operations. We modularized our CUDA code as linear algebra routines whose interfaces are very similar to LAPACK's. Thus, it is easy for developers to maintain the code, as well as the users to reuse them.

The purpose of kernel 1-6 is to compute  $\mathbf{A}_z$  in (5). These kernels take the new state data  $(\mathbf{v}, \mathbf{e}, \mathbf{x})$  as input.

**Kernel 1,2** solve SVD, eigenvalue, EOS function. They are used in evaluations of  $\hat{\sigma}(\hat{q}_k)$ , and adjugate of  $\mathbf{J}_z$ . Independent operations are performed on each quadrature point (thread). Each thread implements formula to solve SVD, eigenvalue for  $DIM \times DIM$  ( $2 \times 2$  or  $3 \times 3$ ) matrices. They are data streamed through kernels.

**Kernel 3,4** Evaluation of  $\hat{\nabla} \hat{w}_i(\hat{q}_k)$ ,  $\mathbf{J}_z(\hat{q}_k)$ . This kernel can be expressed by a special batched DGEMM  $\mathbf{G}_k = \mathbf{D}_z^T \mathbf{W}$  where  $\mathbf{D}_z$  is related to the gradient operator,  $\mathbf{W}$  is the basis function.

**Kernel 5,6** Auxiliary kernels batched DGEMM, where all matrices are  $DIM \times DIM$ . These kernels multiply Jacobin  $\mathbf{J}_z$ , gradient of basis functions  $\hat{\nabla} \hat{w}_i$ , stress values  $\hat{\sigma}$  together.

**Kernel 7** One thread block works on one zone. Each thread block does a matrix-matrix transpose multiplication  $\mathbf{F}_z = \mathbf{A}_z \mathbf{B}^T$ , where  $\mathbf{A}_z$  is the output of last kernel. Therefore, this kernel can be also expressed as a batched DGEMM, with the number of batches being the number of zones.

**Kernel 8** and **Kernel 10** compute  $-\mathbf{F} \cdot \mathbf{1}$  from (1) and  $\mathbf{F}^T \cdot \mathbf{v}$  from (2), respectively. Each thread block does a matrix-vector multiplication (DGEMV) and computes part of a big vector. All thread blocks assemble the result vector. The two kernels can be expressed as batched DGEMV.

**Kernel 9** is not a single kernel function in CUDA. It is a custom conjugate gradient solver for (1) with a diagonal preconditioner (PCG) [15]. It is constructed with CUBLAS/CUSPARSE routines [13]. When enabled, this custom CUDA solver performs step 6 from Section 2 and is outside of the corner force calculation. Currently, it only runs on a single GPU (see discussion in Section 3.4). The result vector  $\frac{d\mathbf{v}}{dt}$  is transferred back to the CPU and used in next time step if time is not out.

**Kernel 11** is a sparse (CSR) matrix multiplication to compute  $\mathbf{M}_\varepsilon^{-1}(\mathbf{F}^T \cdot \mathbf{v})$  by calling a CUSPARSE SpMV routine [13]. The reason for calling SpMV routine instead of using a CUDA-PCG solver as in kernel 9 is that the matrix  $M_\varepsilon$  is block diagonal as described in Section 2. The inverse of  $M_\varepsilon$  is only computed once at the initialization stage. The result vector  $\frac{d\mathbf{e}}{dt}$  is transferred back to the CPU and used in next time step.

##### 3.1.2 Memory Transfer and CPU Work

Input vectors  $(\mathbf{v}, \mathbf{e}, \mathbf{x})$  are transferred from the host to the device before kernel 1, and output vectors  $\frac{d\mathbf{e}}{dt}$  are transferred

No.	Kernel Name	Purpose
1	kernel_CalcAjugate_det	SVD,Eigen,Ajugate
2	kernel_loop_grad_v	EoS, $\hat{\sigma}(\hat{q}_k)$
3	kernel_PzVz_Phi_F_Batched	$\hat{\nabla} \hat{w}_i(\hat{q}_k), \mathbf{J}_z(\hat{q}_k)$
4	kernel_Phi_sigma_hat_z	$\hat{\sigma}(\hat{q}_k)$
5	kernel_NN_dgemmBatched	Auxiliary
6	kernel_NT_dgemmBatched	Auxiliary
7	kernel_loop_zones	$\mathbf{A}_z \mathbf{B}^T$
8	kernel_loop_zones_dv_dt	$-\mathbf{F} \cdot \mathbf{1}$
10	kernel_dgemvt	$\mathbf{F}^T \cdot \mathbf{v}$
9	CUDA_PCG	Solve linear system(1)
11	SpMV	Solve linear system(2)

**Table 2: Implementation on Kepler. Kernel 9 is a set of kernels instead of one single kernel.**

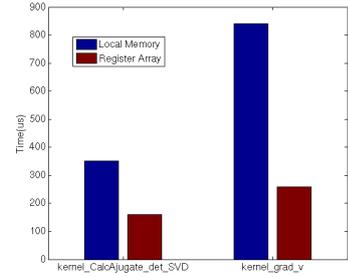
back from the device to the host after kernel 11. Whether vector  $\frac{d\mathbf{v}}{dt}$  after kernel 9 or vector  $-\mathbf{F} \cdot \mathbf{1}$  after kernel 8 is transferred to host depends on turning on/off the CUDA-PCG solver. The time integration of the output right-hand-side vectors in the momentum (1) and energy (2) equations, together with the motion (3) equation are still done on CPU to get new  $(\mathbf{v}, \mathbf{e}, \mathbf{x})$  states.

Because of kernel 8,10, the two DGEMV kernels, we avoid transferring the full matrix  $\mathbf{F}$  which has large number of non-zeros due to its high-order origin. This leads to significant reduction in the amount of data transferred between the CPU and GPU via the relatively slow PCI-E bus.

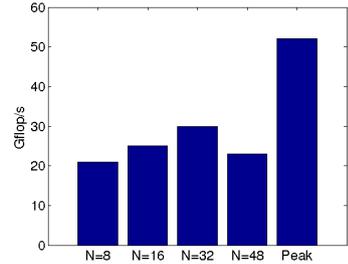
### 3.2 Optimization and Autotuning

Data streamed **kernel 1,2**. Each thread maintains a  $DIM \times DIM$  workspace for each matrix and a number of scalar variables. In the base implementation, the workspace related to the two kernels will be put in local memory by compiler, even declared as registers. The register spill issue is serious by inspecting the PTX code, especially on Fermi(computing ability 2.0) whose registers are rather limited [10]. After separating the two kernels out, there is no register spill issue related to their workspaces, especially on Kepler(computing capability 3.5) which doubles the number of physical registers per SMX. Registers is critical to the two kernels performance, because there are a large amount of scalar operations related with the workspace. Figure 3 compared the performance of using register array and that of being forced using local memory for workspace on Kepler. By taking advantage of more registers of Kepler, kernel 2 achieved 4x faster.

**Kernel 5,6:** Batched DGEMM of  $DIM \times DIM$  matrices. Each thread block performed multiple matrices operations. This avoided unaligned memory access problem in the case of one thread block reading one matrix size of 4 or 9. Threads inside block are configured flexibly. When reading and writing, threads can be organized as one dimension to access linearly stored data in global memory. While performing multiplication, they are configured as two-dimension to naturally fit the matrix indexing. The number of matrix performed per thread block can be tuned to find an optimal occupancy. We use autotuning(Section 3.2.1) to tune this number. We find 32 delivered the best performance with an occupancy 98.3%. Figure 4 show the effect of tuning and we are able to achieve 60% of the theoretic



**Figure 3: Performance of kernel 1,2 with local memory and register arrays, respectively. Test is performed on K20 for a 3D  $Q_2-Q_1$  case.**



**Figure 4: Tuning of kernel 3 which achieved 60% of theoretical peak performance on K20. N is the number of matrices performed in each thread block. The optimal N is found by auto tuning.**

cal peak performance of batched DGEMM on K20. Notice the peak performance of batched is lower than that of the regular DGEMM, since batches of small matrices can not achieve the same GFLOPS as one large matrix. One  $n^2$  matrix performs  $n^3$  operations, but  $k^2$  small  $(\frac{n}{k})^2$  matrices only perform  $k^2(\frac{n}{k})^3 = \frac{n^3}{k}$  operations with the same input size [12]. The Batched DGEMM flop per element is  $\frac{2k * DIM^3}{3k * DIM^2} = \frac{2DIM}{3}$ . The bandwidth of K20 is 208GB/s, which means it is able to get 26G data in double precision per second. Since each element will perform 4/3, 2 operations, the theoretical peak performance on K20 is 35, 52 Gflop/s for  $DIM = 2,3$ , respectively. `cublasDgemmbatched` has exactly the same purpose but only achieves 1.3Gflop/s.

**Kernel 3,4** implement custom batched DGEMM  $\mathbf{C} = \mathbf{A}\mathbf{B}$ . Notice here  $\mathbf{A}$  and  $\mathbf{B}$  are different from that in (5) and (6). They are described in Table 3. In kernel 3,4, since number of quadrature points  $\ll$  zones, the number of matrices  $\mathbf{B}$  is much smaller compared to that of  $\mathbf{A}$ . Therefore we choose to use texture memory to access matrix  $\mathbf{B}$  in version 1(v1), because we hope they fit the cache. We always use shared memory to read  $\mathbf{A}$ . It seems that reading  $\mathbf{B}$  via cached texture memory is still not as fast as shared memory as in v2, though shared memory will introduce synchronization overhead. To reuse  $\mathbf{A}$ , each thread block will loop over all of smaller matrices  $\mathbf{B}$ , because if this thread block only loop partial of  $\mathbf{B}$ ,  $\mathbf{A}$  will be picked up in other thread block who loops the remaining part of  $\mathbf{B}$ . Similar to the optimiza-

tions in **kernel 5,6**, to increase occupancy, we fit multiple **A** in one thread block. This, in turn, also helps the reuse of **B**, for the same reason as above. Yet, too many matrices will overflow shared memory and reduce the occupancy which offsets the benefits of data reuse. But we can always use auto tuning to find the balance point. v3 is the tuning result, as shown in Figure 6.

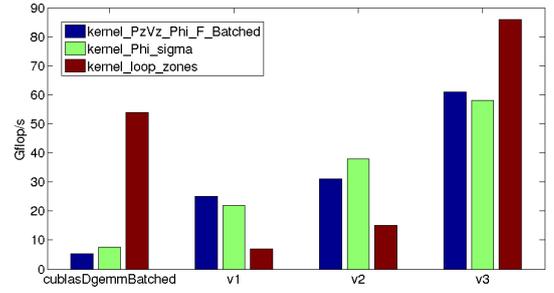
Name	Num of A	Num of B	Num of C
kernel 3	zones	points	zones*points
kernel 4	zones*points	points	zones*points
kernel 7	zones	1	zones

**Table 3: Points refers to the quadrature points. For example, in kernel 3 each quadrature point corresponds a matrix B and each zone corresponds a matrix A.**

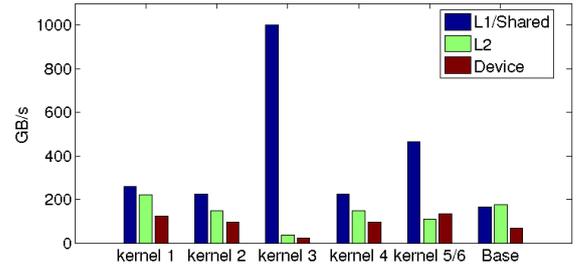
Optimization of **kernel 7**. v1 is a naive implementation.  $A_z$  and **B** were loaded directly from global memory, as shown in Figure 6. In v2, we use shared memory to read  $A_z$ , while **B** is read in constant memory, since **B** is globally shared by all thread blocks. v2 is a substantial improvement, but still not satisfactory. As a further optimization, v3 uses blocking technique. Blocking is the process of dividing a large matrix into smaller matrices to solve. Blocking is widely adopted in LAPACK. The main purpose of blocking is to increase data locality and thus improve cache performance. On GPU, blocking can deliver a second benefit: reducing the amount of shared memory for each thread block and allowing more thread blocks to reside on streaming multiprocessors and thus enhance the parallelism. Blocking can be done in different patterns. We found that accessing columns in blocks by 1D dimension proved to be most effective, while blocking in rows has little benefits, probably because the data layout is in column major. Again, the optimal blocking size can be found by our autotuning technique. An alternative implementation of these batched DGEMM kernels is to call `cublasDgemmbatched` [12] as shown in Figure 6.

**Bandwidth** Because of the "memory wall", most applications are bandwidth bounded nowadays. Here we chose to profile the bandwidth of the base and optimized code on K20. Figure 7 shows the bandwidth of all the 3 level memory, from on-chip memory L1/Shared to off-chip L2 and device memory. All the optimized kernels exceeded the base implementation in bandwidth of L1/Shared and device memory except kernel 3 in device memory, which instead has very high bandwidth in L1/shared memory. Optimized kernels achieved much higher bandwidth in L1/shared memory, because they exploited shared memory and/or registers a lot. Kernel 1,2 have streaming data which is more likely cached by L2. Their bandwidth in L2 tend to be higher compared to other kernels. Because on-chip memory is much faster than off-chip memory, the bandwidth of on-chip memory has a greater impact on performance.

**Kernel 8 and 10** can be viewed as batched DGEMV. In our implementation, each thread block (zone) does a DGEMV operation. An implementation involving CUBLAS is to put `cublasDgemv` in streams of number zones, as recommended in the User Guide [12], since there is no batched DGEMV routine in CUBLAS. However, the performance is very poor, as shown in Table 4. Our custom kernel is 90x faster than that of `cublasDgemv`, achieving 50 % of theoretical peak performance of batched DGEMV on C2050.



**Figure 6: The performance of kernel 3,4,7 on K20. v1 is the straightforward implementation. v3 is the optimized and tuned result.**



**Figure 7: Memory bandwidth of previous(base) and optimized kernels. The theoretical peak bandwidth of device memory of K20 is 208GB/s.**

We implemented a custom CUDA-PCG solver(**kernel 9**) from scratch. CUDA-PCG contains a SpMV and a dot product routine only where we call CUSPARSE SpMV and `cublasDdot` as shown in Figure 5. Kernel 11 is a sparse matrix multiplication routine in CUSPARSE. Notice, this SpMV routine is also needed in **kernel 9**. But kernel 11 is only called once per time step. From Figure 5 We can see the performance of SpMV is critical to the CUDA-PCG, since it is the biggest component of CUDA-PCG. The CUDA-PCG solver is outside of corner force. By comparing Figure 5 and timing data in Table 1, we can see the optimized corner force time becomes very small in overall time.

The impact of redesign and optimization is shown in the right side of Figure 5. `kernel_loop_quadrature_point` is replaced by kernel 1-6. Its percentage decreases to 25% from 65% after optimization. The actual time of `kernel_csrMv_ci_kernel` (SpMV) does not change, but its percentage increases from 30% to 65%, because the total time is reduced. Its percentage is big, because it is called many times in one time step. Other kernels will only be called once, except kernel 5 twice.

### 3.2.1 Autotuning CUDA kernels

	streamed <code>cublasDgemv</code>	kernel 8	theoretical
Gflop/s	0.2	18	35.5

**Table 4: Custom kernel 8 and streamed `cublasDgemv` implement batched DGEMV on one C2050. In this test case, each small matrix is 81 by 8 and each vector is 8. Number of batches (streams) is 4096.**

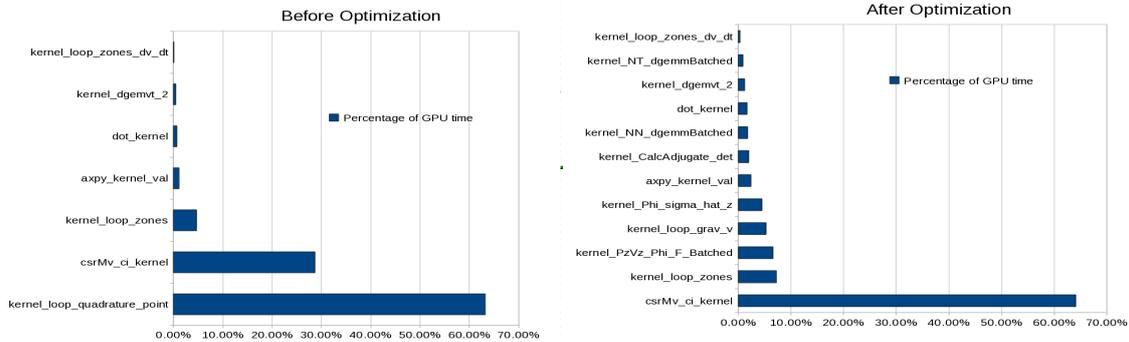


Figure 5: Kernel time distribution of the base versus redesigned and optimized implementation

Except the Jacobin matrix  $\mathbf{J}$  and sigma  $\hat{\sigma}$  only depending on the dimension, the others change a lot with the orders of FEMs. For example,  $\hat{w}_i(\hat{q}_k)$  is  $81 \times 64$  for  $Q_2$ - $Q_1$ ,  $375 \times 512$  for  $Q_4$ - $Q_3$  in 3D. We propose an autotuning technique to adapt our kernels to the order of the method. Tuning can cause huge difference in performance, as shown in the figures above.

Our autotuning is based on the iterative feature of CFD application. First, we parameterize every kernel as far as possible. Parameters include the blocking size in kernel 7, number of matrices in kernel 5,6, etc. Second, we set up a range of possible values for the parameters we want to tune. Artificial values, like those exceeding the shared memory, will be eliminated. The second step usually takes most efforts. In each sampling period, the scheduler picks up a candidate value and times it. After comparing all the candidates, the scheduler will give an optimal one. In our test, one sampling period consists of forty time steps which will be averaged to eliminate the noise.

### 3.3 CUDA + OpenMP Implementation of Corner Force

CUDA and OpenMP is used on Kepler K10 and Fermi clusters, where only one MPI process is allowed to connect to the GPU at a time. Multiple MPI processes will be forced into a serialization if the GPU is in shared compute mode. They will be prohibited if it is in exclusive compute mode. This limitation usually requires the number of GPU to match CPU cores, otherwise GPU and CPU utilization will be limited. However, the number of CPU cores is usually far more than that of GPU. Therefore, we adopt OpenMP to deal with multi-core and CUDA to deal with one K10 or Fermi in one MPI process. On Kepler K20, we may turn OpenMP off, because MPI processes from multi-core can use a subset of a shared K20 simultaneously due to Hyper-Q (see Section 4.2) without causing the CPU core idle.

In the corner force computation, after the launch of CUDA kernels, control can return to a host thread prior to the GPU completing work. The host thread will spawn OpenMP threads and distribute the a portion of zones among the threads. Each thread allocates private working space and executes like normal serial code. There is no synchronization between threads unless they exit the parallel region. A synchronization between the CPU and the GPU is required to complete the corner force calculation, because there is data dependency in the following code. We use auto-balance to

find the ratio between CPU and GPU to ensure load balance. The idea of auto-balance of CUDA and OpenMP is the same with autotuning. Scheduler will compare their time to decide to move more or less work to each processor. After a few sampling periods, the scheduler will converge to an optimal ratio. Our tests shows that the convergence only takes a few time periods as shown in Table 5.

Problem	Optimal ratio	Convergence periods
2D: Sedov	75%	14
2D: Triple-pt	77%	12

Table 5: Ratio refers to the percentage of zones on GPU. Zones are allocated on a six core X5560 CPU and a C2050 GPU.

Auto tuning is a convenient and robust tool. When the code is ported on another architecture, the changes will be detected and load will be rebalanced automatically.

### 3.4 MPI Level Parallelism

The MPI level parallelism in BLAST is based on MFEM which is a modular C++ finite element library [14]. At the initialization stage (Step 2 in Section 2), MFEM takes care of the domain splitting and parallel mesh refinement as shown in Figure 8. Each MPI task is assigned a sub-domain consisting of a number of elements (zones). Finite element degrees of freedom (DOFs) shared by multiple MPI tasks are grouped by the set (group) of tasks sharing them and each group is assigned to one of the tasks in the group (the master), see Figure 9. This results in a non-overlapping decomposition of the global vectors and matrices and typical FEM and linear algebra operations, such as matrix assembly and matrix-vector product, require communications only within the task groups.

After computing the corner forces, a few other MPI calls are needed to handle the translation between local finite element forms and global matrix / vector forms in MFEM (Step 5 in Section 2). An MPI reduction is used to find the global minimum time step.

Because computing the corner forces can be done locally, the MPI level and the CUDA/OpenMP parallel corner force level are independent. Each module can be enabled or disabled independently. However, the kinematic mass matrix  $\mathbf{M}_V$  in (1) is global and needs communication across processors, because the kinematic basis is continuous and components from different zones overlap. The modification of MFEM's PCG implementation needed to enable the CUDA-PCG solver to work on multi-GPU, is beyond the scope of

the present work. With the higher order of the methods, CG time will be less significant compared to corner force time. Therefore, we only consider the CUDA-PCG solver for (1) on a single GPU.

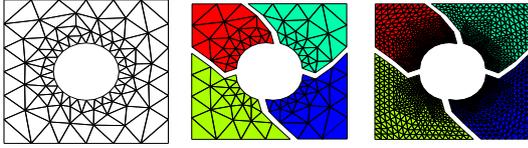


Figure 8: Parallel mesh splitting and parallel mesh refinement

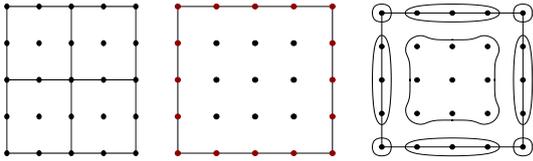


Figure 9: Zones assigned to one MPI task and associated  $Q_2$  DOFs (left); the DOFs at the boundary of this subdomain are shared with neighboring tasks (middle); groups of DOFs, including the local group of internal DOFs (right).

## 4. TESTING RESULTS AND DISCUSSION

For our test cases we consider the 3D Sedov blast wave problem (see [1] for further details on the nature of these benchmarks). In all cases we use double precision. The gcc compiler and NVCC compiler under CUDA v5.0 are used for the CPU and GPU codes, respectively.

### 4.1 Validation of CUDA code

We get consistent results on the CPU and the GPU. Both the CPU and the GPU code preserved the total energy of each calculation to machine precision, as shown in Table 6. The density distributions of a 3D Sedov problem in GPU code are shown in Figure 10.

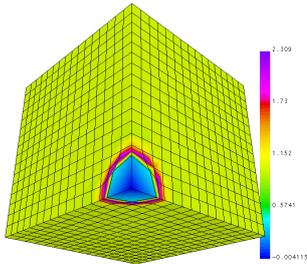


Figure 10: 3D Sedov problem

### 4.2 Performance on A Single Node

Due to the new feature Hyper-Q of Kepler, multiple MPI processes can run on a K20 GPU simultaneously. K20 GPU is able to set up up to 32 work queues between the host and the device. Each MPI process will be assigned to a different hardware work queue, and thus enabling them to run concurrently on the GPU.

In our test, the CPU is a 8 core Sandy Bridge E5-2670 and the GPU is a K20. Unless explicitly noted, we always use

them to perform our tests in the following sections. In our configuration, 8 MPI tasks share one K20. Only corner force is accelerated on the GPU. Figure 11 shows the speedup achieved by CPU-GPU over the CPU. We compared two order of methods  $Q_2-Q_1$  and  $Q_4-Q_3$ . When the order is higher, the percentage of corner force increases, and GPU acceleration benefit BLAST more. The overall simulation time of  $Q_4-Q_3$  than  $Q_2-Q_1$  is 3.2x on the CPU, but only 2x on CPU-GPU.

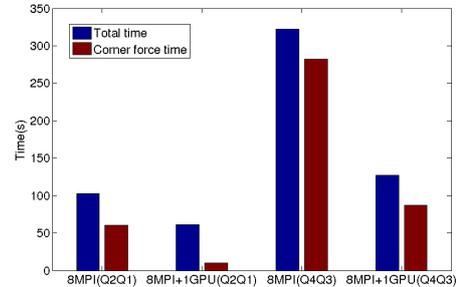


Figure 11: Speedup of CPU-GPU code over the CPU. A 1.9x overall speedup is obtained on  $Q_2-Q_1$ ; 2.5x on  $Q_4-Q_3$ .

### 4.3 Performance on Distributed Systems: Strong and Weak Scalability

We tested our code on ORNL Titan which has 16 AMD cores and 1 K20m per node. We scaled it up to 4096 computing nodes. 8 nodes is the base line. For a 3D problem, one refinement level will cause the domain size 8x bigger. We fixed a domain size 512 for each computing node, and used 8x more nodes for every refinement. From 8 nodes to 4096 nodes, 5 cycle steps increase from 0.85 to 1.83 seconds. The limiting factor is the MPI global reduction to find the minimum time step after corner force computation and MPI communication in MFEM (Step 5 in Section 2).

We also tested the strong scalability on a small cluster, SNL Shannon machine. It has 30 computing nodes, with two K20m and two sockets of Intel E5-2670 CPU per node. Figure 13 shows the linear strong scaling on this machine. The domain size is  $32^3$ .

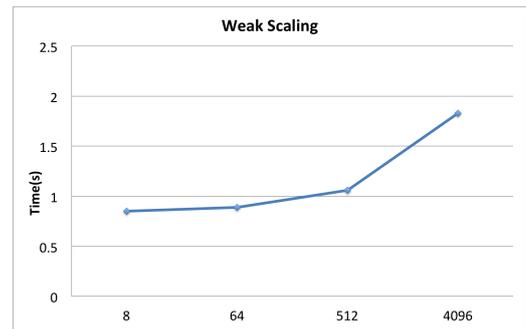


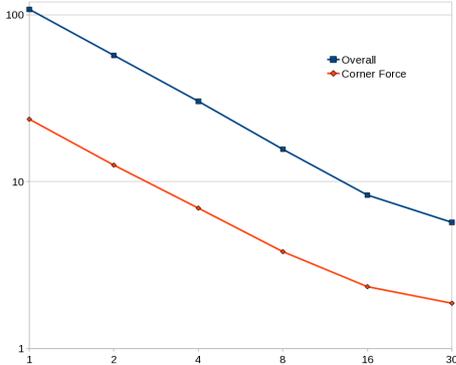
Figure 12: Weak scaling. Time is for 5 cycles of time steps.

## 5. POWER AND ENERGY ANALYSIS

Generally, there are two ways to measure power. First, attaching an external power meter to the machine. This is

Platform	Final Time	Kinetic	Internal	Total	Total Change
CPU	0.6	5.0423596813598e-01	9.5457640318651e+00	1.005000000001e+01	-9.2192919964873e-13
GPU	0.6	5.0418618040297e-01	9.5458138195986e+00	1.005000000002e+01	-4.9382720135327e-13

**Table 6: Results of CPU and GPU code for 2D triple-pt problem using a  $Q_3$ - $Q_2$  method; the total energy includes kinetic energy and internal energy. Both CPU and GPU results preserve the total energy to machine precision.**



**Figure 13: Strong scaling.**

from a hardware aspect. It is accurate, but it can only measure the power of the whole machine. It is not able to profile power usage of individual processor or memory. The other way is estimation from an software aspect. We adopt this way in our measurement.

### 5.1 Profiling CPU Power with Intel RAPL

From SandyBridge, Intel CPUs support on board power measurement via the Running Average Power Limit (RAPL) interface [8]. The internal circuitry can estimate current energy usage based on a model accessing Model (or sometimes called Machine) Specific Registers(MSRs), with an update frequency on the order of milliseconds. The power model has been validated by Intel [16] to actual energy.

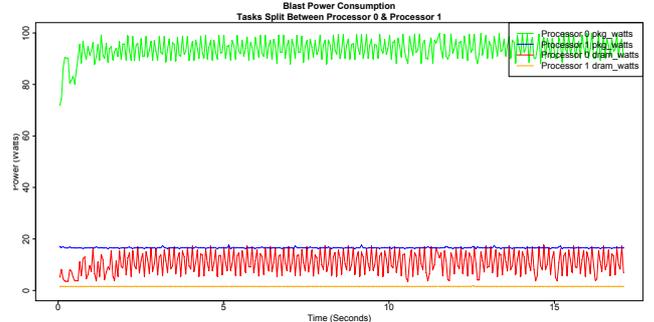
RAPL provide measurement of

- 1) The total package domain.
- 2) PP0 (Power Plane 0) which refers to the processor cores in a package.
- 3) Memory which includes the directly-attached DRAM.

Figure 14 shows the power of two CPU packages and their DRAM. For comparison purpose, we let one package(processor) busy, while the other keep idle. The full loaded package power is 95W with DRAM at 15W. The idle power is slightly lower than 20W with DRAM almost at 0. The TDP of the E5-2670 is 115W. Our observation 95W (82%) confirms the AMD reports of the normal range of Average CPU Power (ACP) in [7]. The test is a 3D  $Q_2$ - $Q_1$  case with 8 MPI tasks without GPU.

### 5.2 Profiling GPU Power with NVML

Recent NVIDIA GPU support power management via the NVIDIA Management Library (NVML). NVML provides API to developers. NVIDIA also provides a high level utility nvidia-smi which calls the same interface. It only reports the entire board power, including GPU and its memory. It has milliwatt resolution within  $\pm 5$  W and is updated per millisecond. Our CUDA kernels time is around several to tens



**Figure 14: Power of two packages of Sandy Bridge CPU. Package 0 is full loaded. Package 1 is idle.**

milliseconds for our target problem, so the computation will not be missed by NVML.

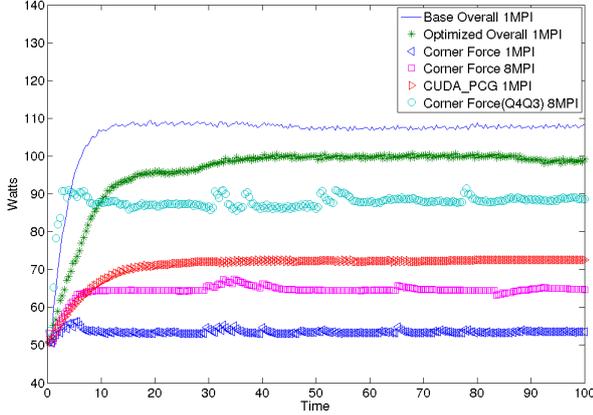
We test the power usage of one K20 GPU in six scenarios. 1,2) Previous base versus optimized implementation with both corner force and CUDA\_PCG solver enabled with 1 MPI task. 3) Optimized corner force ( $Q_2$ - $Q_1$ ) only with 1MPI task. 4,5) Optimized corner force ( $Q_2$ - $Q_1$  and  $Q_4$ - $Q_3$ ) with 8 MPI tasks running on the same GPU. 6) CUDA\_PCG ( $Q_2$ - $Q_1$ ) only with 1 MPI task. The test case is a 3D Sedov problem with the domain size  $16^3$ , which is the maximum size we were able to allocate with  $Q_4$ - $Q_3$  because of the limitation of memory of K20. The GPU is warmed up by a few runs to reduce noise. Our test shows that the startup power is around 50W by launching any kernel. The idle power is 20W if the GPU is doing nothing for a long time. The TDP of K20 is 225W.

From Figure 15 we can see, the optimized code not only runs faster, but also lowers the power cost than base implementation. Since both perform the same FLOPs, their main difference is how to exploit the memory. The memory power consumes around 25% of total GPU power in various reports [5]. Studies have examined the power consumption of different components of GPU [17] [18]. The power consumption of device memory is much higher than on-chip memory. In the micro-benchmarks of [18], the device memory power is 52, while shared memory is 1 with FP and ALU only 0.2(normalized unit). Their difference can be explained by their accessing cost. Accessing on-chip shared memory can only take 1 cycle, while accessing device memory may take 300 cycles [10]. If the memory utilization (the cycles of data-bus transferring data for read or write) is lowered, the overall power will also be lowered. Because the optimized kernels effectively exploited on-chip memory and significantly improved the bandwidth, as shown in Figure 7, the memory utilization efficiency is improved and power is reduced.

When the GPU is shared by 8 MPI tasks, its power will be higher than 1 MPI (with the same domain size and problem). We did not find any previous report about this situation,

but obviously this additional power cost should come from the overhead of Hyper-Q. 1MPI corner force  $Q_2-Q_1$  has not saturated GPU, therefore its power is low. Compared to  $Q_2-Q_1$ , because the FLOPs and data of  $Q_4-Q_3$  method is much bigger, its utilization and power consumption of GPU is also higher.

The power of solver (CUDA-PCG) is higher than that of corner force. Partly because CG(SpMV) is very memory bound due to its sparse structure, it is very hard to achieve a memory bandwidth efficiency comparable with dense matrices. While kernels in the corner force are dense matrix operations.



**Figure 15: One K20 power. Except denoted, all are with a  $Q_2-Q_1$  method.**

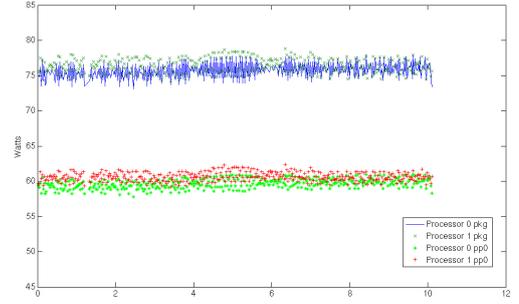
The CPU power with corner force accelerated on GPU is shown in Figure 16. Both of the two processors are busy. The total package power is around 75W and PPO at 60W. Their difference is mainly the DRAM power. Compare to Figure 14, CPU power is reduced by 20W. We tested various orders of methods, but did not see any obvious difference.

### 5.3 Greenup

Similar to the notion of speedup which is usually used to describe the performance boost, we define a notion of greenup to quantify the energy efficiency [31].

$$\begin{aligned}
 \text{Greenup} &= \frac{CPU_{energy}}{(CPU + GPU)_{energy}} = \\
 &= \frac{CPU_{power} \cdot CPU_{time}}{(CPU + GPU)_{power} \cdot (CPU + GPU)_{time}} \\
 &= \frac{CPU_{power}}{(CPU + GPU)_{power}} \cdot \text{Speedup} \\
 &= \text{Powerup} \cdot \text{Speedup}
 \end{aligned}$$

where powerup and speedup are larger than 0. Powerup may be less than 1, since CPU+GPU power may exceed that of CPU only. Yet, the speedup is greater than 1. Therefore the greenup will be larger than 1. Table 7 outlines the greenup, powerup and speedup of BLAST code of Sedov problem. Speedup is from Figure 11. The CPU+GPU power we used in Table 7 is by adding data in Figure 15 and Figure 16 together. The hybrid CPU-GPU solution make BLAST greener. It saved 27% and 42% of energy, respectively for the two methods, compared to CPU only solution. However, The use of GPU is more than energy and speedup.



**Figure 16: CPU power with GPU accelerating.**

Because the CPU power decreases, the power leakage and failure rate of cores are also reduced. Applications are more fault tolerant and runs faster, since the frequency of checking points can be reduced.

Method	Power Efficiency	Speedup	Greenup
$Q_2-Q_1$	0.67	1.9	1.27
$Q_4-Q_3$	0.57	2.5	1.42

**Table 7: The CPU-GPU greenup over CPU for BLAST code in 3D Sedov problems.**

## 6. RELATED WORK

The incompressible flow computations on GPUs with stencil computing were studied in [19],[20]. Strategies of assembly of galerkin methods on GPUs were discussed in [21],[24]. [22] [23] focused on solving sparse linear system resulting from FEM discretization. A High order discontinuous galerkin FEM applied in electromagnetic on GPUs was discussed in [25]. L.Wang et al examined the limiting factors of scaling on big CPU-GPU clusters in [26].

There are some reports comparing the power consumption of GEMM routines on CPU/GPU[27], [28]. They were using power meters and measuring the whole system power. [17] studied the power and energy consumption of BLAS2 and MAGMA Hesseberg kernel with NVML. PAPI provide a uniform interface to measure AMD and Intel CPU power, which in turn calls RAPL on Intel CPUs [29]. [30] used RAPL to study the HPC system power allocation of CPU.

## 7. CONCLUSIONS

The BLAST code uses high order finite element methods to solve compressible hydrodynamics problems. It requires the assembling of corner force matrices and the solution of both sparse and dense linear algebra problems. In this paper, we redesigned the most computational intensive part of BLAST for CPU-GPU clusters. Our goal is to maximize the performance and lower the energy consumption of BLAST.

We redesigned a base CUDA implementation and modularized our new implementation into a set of linear algebra routines. Our optimized routines exceed CUBLAS library routines substantially in performance. Besides, compared to the base implementation, our redesign and optimization achieved better performance per watt: it reduced 60% in executing time and 10% in watts. To adapt to high order methods, we further introduced an autotuning technique to tune our CUDA kernels. Due to their general purpose, these linear algebra routines can be utilized by other applications.

Our proposed hybrid solution has been proved to be very beneficial in performance and energy efficiency, especially for high order finite element simulations. We reestablished the appeal of the high order method on GPUs. Compared to low order methods, the speedup and greenup of the high order FEMs are greater, and simulation time grows more slowly.

## Acknowledgments

Implementation were done during an internship at LLNL. Some optimizations were done at University of Tennessee, Knoxville. The authors thank Barry Rountree for providing the CPU power graph. The authors would like to thank the NSF and NVIDIA for supporting. The authors thank the testing support of the Performance End Station PEAC Project sponsored by DoE under Contract No. DE-AC05-00OR22725.

## 8. REFERENCES

- [1] V.A.Dobrev, Tz.V.Kolev, R.N.Rieben. *High order curvilinear finite element methods for Lagrangian hydrodynamics*, SIAM J. Sci. Comp., 34(5), 2012, 606-641.
- [2] <http://www.top500.org>, 2013
- [3] <http://www.green500.org/>
- [4] P.Wang, C.Yang, Y.Chen, Y.Cheng *Power Gating Strategies on GPUs*, ACM Transactions on Architecture and Code Optimization, Volume 8 Issue 3, October 2011.
- [5] J.Zhao, G.Sun, G.H.Loh, Y.Xie *Energy efficient GPU Design with Reconfigurable Inpackage Graphics Memory*, ISLPED,12
- [6] Whitepaper: NVIDIA Next Generation CUDA Compute Architecture: Kepler GK110. <http://www.nvidia.com/content/PDF/kepler/NVIDIA-Kepler-GK110-Architecture-Whitepaper.pdf>
- [7] ACP The Truth About Power Consumption Starts Here: <http://www.amd.com/us/Documents/43761D-ACP-PowerConsumption.pdf>
- [8] Intel 64 and IA-32 Architectures Software Developer's. <http://download.intel.com/products/processor/manual/>
- [9] <https://developer.nvidia.com/nvidia-management-library-nvml>
- [10] NVIDIA CUDA C Programming Guide v4.2, <http://developer.nvidia.com/cuda/nvidia-gpu-computing-documentation>.
- [11] MAGMA: <http://icl.cs.utk.edu/magma/>
- [12] CUBLAS User Guide, <http://developer.nvidia.com/cuda/nvidia-gpu-computing-documentation>.
- [13] CUSPARSE User Guide, <http://developer.nvidia.com/cuda/nvidia-gpu-computing-documentation>.
- [14] MFEM: <http://mfem.googlecode.com/>
- [15] M.Naumov, *Incomplete-LU and Cholesky Preconditioned Iterative Methods Using CUSPARSE and CUBLAS*, June 21, 2011.
- [16] Rotem, E., Naveh, A., Rajwan, D., Ananthakrishnan, A, E, Weissmann. *Power-management architecture of the Intel microarchitecture codenamed Sandy Bridge*, IEEE Micro, vol. 32, no. 2, pp. 20-27, 2012
- [17] K.Kasichayanula, D.Terpstra, P.Luszczek, S.Tomov, S.Moore, G.D.Peterson. *Power Aware Computing on GPUs*, SAAHPC, 2012.
- [18] S.Hong, H.Kim. *An Integrated GPU Power and Performance Model*, ISCA, 2010.
- [19] D.A.Jacobsen, J.C.Thibault, I.Senocak. *An MPI-CUDA Implementation for Massively Parallel Incompressible Flow Computations on Multi-GPU Clusters*, 48th AIAA Aerospace Sciences Meeting and Exhibit, 2010.
- [20] J.C.Thibault, I.Senocak. *Accelerating Incompressible Flow Computations with a Pthreads-CUDA Implementation on Small-Footprint Multi-GPU Platforms*, The Journal of Supercomputing, 59(2), 693-719.
- [21] C.Cecka, A.Lew, E.Darve. *Assembly of finite element methods on graphics processors*, Numerical Methods in Engineering. Volume 85, Issue 5, 4 February 2011.
- [22] J.Bolz, I.Farmer, E.Grinspun, P.Schroder. *Sparse Matrix Solvers on the GPU: Conjugate Gradients and Multigrid*. ACM Transactions on Graphics, 2003;
- [23] M.Geveler, D.Ribbrock, D.Goddeke, P.Zajac, S.Turek. *Towards a complete FEM-based simulation toolkit on GPUs: Unstructured grid finite element geometric multigrid solvers with strong smoothers based on sparse approximate inverses*, Computers & Fluids, Feb,2012.
- [24] A.Klocknera, T.Warburtonb, J.Bridgeb, J.S.Hesthavena, *Nodal discontinuous Galerkin methods on graphics processors*, Journal of Computational Physics Volume 228, Issue 21, 20 November 2009.
- [25] Godel, N. Nunn, N. Warburton, T. Clemens, M. *Scalability of Higher-Order Discontinuous Galerkin FEM Computations for Solving Electromagnetic Wave Propagation Problems on GPU Clusters*. Magnetics, IEEE Transactions, Aug 2010.
- [26] L.Wang, W.Jia, X.Chi, Y.Wu, W.Gao, L.Wang. *Large Scale Plane Wave Pseudopotential Density Functional Theory Calculations on GPU Clusters*, SC11, 2011.
- [27] S.Huang, S.Xiao, W.Feng. *On the Energy Efficiency of Graphics Processing Units for Scientific Computing*, IPDPS, 2009
- [28] Y.Abe, H.Sasaki, M.Peres, K.Inoue, K.Murakami, S.Kato. *Power and Performance Analysis of GPU-Accelerated Systems*, USENIX, 2012 Workshop on Power-Aware Computing and Systems
- [29] V.M. Weaver, M.Johnson, K.Kasichayanula, J.Ralph, P.Luszczek, D.Terpstra, S.Moore. *Measuring Energy and Power with PAPI*, Parallel Processing Workshops, 2012 41st International Conference on Sep, 2012
- [30] O.Sarood, A.Langer, L.Kale, B.Rountree, B.Supinski *Optimizing Power Allocation to CPU and Memory Subsystems in Overprovisioned HPC Systems*, IEEE Cluster 2012
- [31] D.Lukarski, T.Skoglund. *A priori power estimation of linear solvers on multi-core processors*, UPMARC Winter Meeting 2013