
Performance Analysis and Optimisation of the OP2 Framework on Many-core Architectures

M.B. GILES¹, G.R. MUDALIGE¹, Z. SHARIF², G. MARKALL² AND
P.H.J KELLY²

¹*Oxford e-Research Centre, University of Oxford*

²*Dept. of Computing, Imperial College London*

Email: mike.giles@maths.ox.ac.uk, gihan.mudalige@oerc.ox.ac.uk

{zohirul.sharif09,graham.markall08,p.kelly} @imperial.ac.uk

This paper presents a benchmarking, performance analysis and optimisation study of the OP2 “active” library, which provides an abstraction framework for the parallel execution of unstructured mesh applications. OP2 aims to decouple the scientific specification of the application from its parallel implementation, and thereby achieve code longevity and near-optimal performance through re-targeting the back-end to different hardware. Runtime performance results are presented for a representative unstructured mesh application on a variety of many-core processor systems, including traditional X86 architectures from Intel (Xeon based on the older Penryn and current Nehalem micro-architectures) and GPU offerings from NVIDIA (GTX260, Tesla C2050). Our analysis demonstrates the contrasting performance between the use of CPU (OpenMP) and GPU (CUDA) parallel implementations for the solution on an industrial sized unstructured mesh consisting of about 1.5 million edges. Results show the significance of choosing the correct partition and thread-block configuration, the factors limiting the GPU performance and insights into optimizations for improved performance.

Keywords: CFD; Unstructured mesh applications; Performance; GPU; OpenMP; CUDA

Received 00 Month 2010; revised 00 Month 2010

1. INTRODUCTION

Most scientific parallel programs have been and continue to be written by exclusively targeting a parallel programming model or a parallel architecture using extensions to traditional sequential languages such as Fortran, C or C++. This approach increases the cognitive load on programmers and has persisted primarily due to its good performance and the existence of legacy application software that remains critical to the production workloads of scientific organizations. However, the future of such a programming model’s use in an environment of increasingly complex hardware architecture is unsustainable. This problem is becoming compounded as the High Performance Computing (HPC) industry begins to focus on delivering exa-scale systems in the next decade. Thus the current situation is both distracting scientists from investing their full intellectual capacity in understanding the physical systems they model, while also hindering them from exploring the capacity of available hardware. It is therefore clear that a level of abstraction must be achieved so that computational scientists can reach an

increased level of productivity without having to learn the intricate details of new architectures.

Such an abstraction enables users to focus on solving problems at a higher level and not worry about architecture specific optimizations. This splits the problem space into (1) a higher application level where scientists and engineers concentrate on solving domain specific problems and write efficient code that remains unchanged for different underlying hardware architectures and (2) a lower implementation level, that focusses on how a computation can effectively be made faster on a given architecture by carefully analysing the data access patterns. This paves the way for easily integrating support for any future heterogeneous hardware.

OPlus (Oxford Parallel Library for Unstructured Solvers) [1], a research project that had its origins in 1993 at the University of Oxford, provided such an abstraction framework for performing unstructured mesh based computations across a distributed-memory cluster of processors [2]. To this day it is used as the underlying parallelization library for Hydra [3,

4, 5] a production-grade CFD application used in turbomachinery design at Rolls-Royce plc. OP2 is the second iteration of OPlus and builds upon the features provided by its predecessor but develops an “active” library approach with code generation to exploit parallelism on heterogeneous many-core architectures.

The “active” library approach uses program transformation tools, so that the user code is transformed into the appropriate form to be linked against the required parallel implementation (e.g. MPI, OpenMP, CUDA, OpenCL, AVX etc.) enabling execution on different target back-end hardware platforms [6]. Currently OP2 includes support for developing unstructured mesh applications for execution on multi-core and/or multi-threaded (OpenMP) CPUs and CUDA capable GPUs.

This paper presents a performance evaluation of the current OP2 library. Our objective is to provide a contrasting benchmarking and performance analysis study of a representative unstructured mesh application (Airfoil [7]) written using OP2 on a range of systems. These consist of representative current many-core hardware technologies such as the traditional X86 architecture systems from Intel and GPU offerings from NVIDIA. More specifically this paper make the following contributions:

1. We present a performance analysis of the Airfoil unstructured mesh application written using OP2 on a number of multi-core CPU systems. OP2’s code transformation framework is used to generate back-end code targeting multi-threaded executables based on OpenMP for two current multi-core processor systems: an Intel Xeon E5462 based on the older “Penryn” micro-architecture and an Intel Xeon E5540 based on the current Intel “Nehalem” micro-architecture. The end-to-end run times reported in this study are for the execution on an industrial-size problem using an unstructured mesh consisting of about 1.5 million edges.
2. The multi-core, multi-threaded CPU performance of a single OP2 version of the Airfoil source code is compared against equivalent GPU solutions executing back-end code generated by OP2 based on NVIDIA CUDA. Performance on a number of CUDA capable GPUs are presented, including a GTX260 consumer card and a Tesla C2050 based on the new Fermi architecture - NVIDIA’s current flagship GPU offering.
3. Our analysis demonstrates the performance issues that distinguish the use of CPU and GPU architectures executing the Airfoil application. These include the significance of choosing the correct partition and thread-block configuration, the factors limiting CPU and GPU performance and insights into optimizations for improved performance. We implement a number of hardware and software configurations and optimisations as

part of this exercise for the GPU architecture.

The rest of this paper is organized as follows: Section 2 details related work in developing abstraction frameworks for multi-architecture platforms; Section 3 provides a description of the class of applications supported by OP2 and its API; Section 4 details the OP2 framework and key issues related to parallelising unstructured mesh applications; Section 5 and Section 6 present performance figures for the execution of Airfoil on CPU and GPU systems respectively, including comparisons between the two architectures; Section 7 investigates the factors limiting CPU and GPU performance including comparative performance on the GPUs for a number of hardware/software configurations and optimisations. Finally Section 8 concludes the paper.

2. RELATED WORK

Although OPlus pre-dates it, OPlus and OP2 can be viewed as an instantiation of the Aecute (access-execute descriptor) [8] programming model that separates the specification of a computational kernel with its parallel iteration space, from a declarative specification of how each iteration accesses its data. The decoupled Access/Execute specification in turn creates the opportunity to apply powerful optimizations targeting the underlying hardware. A number of related research projects have implemented similar programming frameworks including LISZT [9], and the Hybrid Multi-core Parallel Programming (HMPP) [10] workbench.

HMPP allows the user to annotate codelets with HMPP directives that characterise data access in an aggregate (rather than iteration-specific) way. The program is then processed through the tool-chain which uses the hardware vendor specific SDKs to translate it into platform specific code. The resulting executable is run under the “HMPP Runtime” which manages the resources and makes it possible to run a single binary on various heterogeneous hardware platforms. HMPP has no specific support computation on graphs or unstructured meshes.

LISZT is a domain specific language (embedded in Scala) specifically targeting unstructured mesh codes, and is thus more directly comparable. The aim, as with OP2, is to exploit information about the structure of data and the nature of the algorithms in the code, to apply aggressive and platform specific optimizations.

To our knowledge, performance figures for the execution of full scale applications, particularly industrial strength codes developed using the HMPP workbench, have not been published. Preliminary performance figures from the LISZT framework have been presented in [11]. The authors report the performance of Joe, a fluid flow unstructured mesh application using a mesh of 750K cells. Joe is first ported to the LISZT framework and the resulting code

compared to the original code running on a cluster of 4-socket 6-core 2.66GHz Xeon CPUs each with 36GB RAM per node using MPI. Both codes demonstrates equivalent performance (LISZT runtime in fact faster with better scalability) illustrating that no performance loss has resulted due to the use of the LISZT framework. Speed-up figures for the above code running on a Tesla C2050 (implemented using CUDA) against an Intel Core 2 Quad, 2.66GHz processor is provided next, with results showing a speed-up of about $30\times$ in single precision arithmetic and $28\times$ in double precision relative to a single CPU thread.

Related work in the solution of unstructured mesh applications on GPUs, particularly in the CFD domain have also appeared elsewhere. In [12], techniques to implement an unstructured mesh solver on GPUs are described. Implementing three-dimensional Euler equations for inviscid, compressible flow are considered. Average speed-ups of about $9.5\times$ are observed during the execution of the GPU implementation on an NVIDIA Tesla 10 series card against an equivalent optimized OpenMP implementation on an Intel Core 2 Q9450 running 4 threads.

Similarly [13] reports the GPU performance of a Navier-Stokes solver for steady and unsteady turbulent flows on unstructured/hybrid grids. The computations were carried out on NVIDIA's GeForce GTX 285 graphics cards (in double precision arithmetic) and speed-ups up to $46\times$ (vs a single core of two Quad Core Intel Xeon CPUs at 2.00 GHz) are reported.

Research in GPU acceleration often cites speed-ups, relative to a hand-coded CPU implementation - sometimes even comparing to a single-core. In this paper we compare performance on contemporary flagship platforms (NVIDIA C2050, Intel 8-core Penryn and Nehalem). Our goal with OP2 is to generate highly-optimized code for X86 multicore platforms (via OpenMP and the Intel compiler), as well as GPUs, from the same code.

3. BACKGROUND

3.1. Unstructured Mesh Applications

The geometric flexibility of unstructured grids has proved invaluable over a wide area of computational science for solving PDE's (partial differential equations) including: CFD (computational fluid dynamics); CEM (computational electromagnetics); structural mechanics; and general finite element methods. In three dimensions millions of elements are often needed for the required solution accuracy, leading to a large computational cost.

The OPlus approach to the solution of such unstructured mesh problems involves breaking down the unstructured grid algorithms into four distinct parts: (1) sets, (2) data on sets, (3) mappings between sets and (4) operations over sets. These lead to an API through which one can completely and abstractly define




FIGURE 1. An example mesh

any mesh or graph.

Unstructured meshes, unlike structured meshes, use connectivity information to specify the mesh topology. Depending on the application, a set can consist of nodes, edges, triangular faces, or other elements. Associated with these sets are data (e.g. node coordinates, edge weights) and mappings between sets which define how elements of one set connect with the elements of another set. Figure 1 illustrates a simple triangular mesh that we will use as an example to describe the OP2 API.

The mesh illustrated in Figure 1 is defined as two sets, nodes (vertices) and edges, each with their sizes, using the API as follows:

```
op_set set_nodes;
op_decl_set(6, set_nodes, "nodes");
op_set set_edges;
op_decl_set(10, set_edges, "edges");
```

The connectivity is declared through mappings (or pointers) between the sets. The two integer type arrays `pEdge1` and `pEdge2` can be used to represent how an edge is connected to two different vertices.

```
int pEdge1[10] = {1,2,3,3,2,1,4,6,4,1};
int pEdge2[10] = {2,3,4,6,4,4,5,5,6,5};
```

```
op_ptr pointer_edge1;
op_decl_ptr(set_edges, set_nodes, 1, pEdge1,
            pointer_edge1, "pointer_edge1");
op_ptr pointer_edge2;
op_decl_ptr(set_edges, set_nodes, 1, pEdge2,
            pointer_edge2, "pointer_edge2");
```

Each element of `set_edges` is mapped to two different elements in `set_nodes`. In our example, index 0 of `pEdge1` and `pEdge2` maps an edge to its vertices 1 and 2 respectively. Thus the `pEdge1` and `pEdge2` arrays of indices define the connectivity between the two sets. When declaring a mapping we first pass the source (e.g. `set_edges`) then the destination (e.g. `set_nodes`). We then pass the dimension of each element (e.g. 1; as each edge is mapped to 1 node). Once the sets are defined, various data can be associated with them; the following

are some arrays that contain data associated with edges and vertices respectively.

```
float dEdge[10] = {1.2, 3.2, 2.6, 5.2, 5.7,
                  1.3, 1.3, 4.3, 4.7, 3.3};
float dNode[6]  = {0.7, 0.5, 0.9, 2.5, 1.9,
                  6.7};
float *dNode_u =
    (float *)malloc(sizeof(float)*6);

op_dat data_edges;
op_decl_dat(set_edges, 1, "float",
            dEdge, data_edges, "data_edges");

op_dat data_nodes;
op_decl_set(set_nodes, 1, "float",
            dNode, data_nodes, "data_nodes");

op_dat data_nodes_u;
op_decl_dat(set_nodes, 1, "float",
            dNode_u, data_nodes_u, "data_nodes_u");
```

Note that here a single float per set element is declared in this example. A vector of a number of values per set element could also be declared (e.g. a vector with three floats per vertex to store the vertex coordinates).

All the numerically intensive parts of an unstructured mesh application can be described as operations over sets. Within a code this corresponds to a loop over a given set, accessing data either through the mapping arrays (i.e. one level of indirection) or using the data directly from the data arrays, performing some arithmetic, then writing (possibly through the mappings) back to data arrays. The former type of loop is called an indirect loop, while the latter is called a direct loop. The API allows only up to one level of indirection when accessing data arrays.

The OP2 library provides a parallel loop declaration syntax which allows the user to declare the computation over the sets in these loops [14]. For the mesh illustrated in Figure 1 an indirectly accessed loop over all the edges which updates the nodes can be written in the following sequential execution loop:

```
void kernel(int nedge, int *pEdge1,
            int *pEdge2, float *dEdge,
            float *dNode, float *dNode_u) {
    for (int e=0; e<nedge; e++)
        dNode_u[pEdge1[e]] +=
            dEdge[e] * dNode[pEdge2[e]];
}
```

The user declares this loop using the API as follows, leaving the library to handle the parallelization of the loop on a particular architecture.

```
op_par_loop_3(kernel, "kernel", set_edges,
              dEdge, -1, OP_ID, 1, "double", OP_READ,
              dNode, 0, pEdge2, 1, "float", OP_READ,
              dNode_u, 0, pEdge1, 1, "float", OP_INC);
```

This general decomposition of unstructured algorithms imposes no restriction on the actual algorithms, just separates the components of a code. However OP2 makes an important restriction that the order in which the elements are processed must be independent of the final result, that is the order in which we process the elements should not affect the results. This constraint allows the program to choose its own order to obtain maximum parallelism. Moreover the sets and mappings between sets must be static and the operands in the set operations cannot be referenced through a double level of mapping indirection (i.e. a mapping to another set which in turn uses another mapping to data associated with a third set).

Although it might appear that these restrictions are quite severe, the straightforward programming interface, I/O treatment combined with efficient parallel execution makes it an attractive prospect, if the algorithm to be developed falls within the scope of OP2. For example the API could be used for explicit relaxation methods such as Jacobi iteration; pseudo-time-stepping methods; multi-grid methods which use explicit smoothers; Krylov subspace methods with explicit preconditioning; semi-implicit methods where the implicit solve is performed within a set member, for example performing block Jacobi where the block is across a number of PDE's at each vertex of a mesh. However, algorithms based on order dependent relaxation methods such as Gauss-Seidel or ILU (incomplete LU decomposition), lie beyond the capabilities of the API. These are not fundamental restrictions, but we believe they help limit the complexity of the API, and encourage programmers to code in a way that can be made efficient.

The example application used in our analysis, Airfoil, is a non-linear 2D inviscid airfoil code that uses an unstructured grid [7]. It is a much simpler application than the production Hydra [15] CFD application used at Rolls-Royce for the simulation of turbomachinery, but acts as a forerunner for testing the OP2 library for many core architectures. A rendering of a smaller (120×60) unstructured mesh similar to the one used in Airfoil is illustrated in Figure 2. The actual mesh used in our experiments is of size 1200×600, which is too dense to be reproduced here. This consists of over 720K nodes, 720K cells and about 1.5 million edges. The code consists of five parallel loops: `save_soln`, `adt_calc`, `res_calc`, `bres_calc`, `update`. The most compute intensive loop `res_calc` has about 100 floating-point operations performed per mesh edge and is called 2000 times during total execution of the application.

4. OP2

The original OPlus library [1], was developed over 10 years ago for MPI/PVM based distributed memory execution of unstructured mesh algorithms written in Fortran. Its second iteration OP2 is designed to




FIGURE 2. Rendering of a simple 120x60 mesh used in Airfoil




FIGURE 3. Sequential build process

leverage emerging many-core hardware (GPUs, AVX etc.) on top of distributed memory parallelism allowing the user to execute on either a single multi-core/many-core node (including large shared memory systems), or a cluster of multi-core/many core nodes. Currently the OP2 library only supports code development in C/C++. A Fortran API will be developed later with similar functionality.

Since the OP2 specification provides no description of the low-level implementation, re-targeting to future architectures only requires the development of a new code-generation back-end. The current implementation focuses on CUDA (using the NVIDIA CUDA programming model [16]) and OpenMP and will later include code generation for OpenCL and AVX, thus supporting a wide range of CPU and GPU hardware. OP2 will also include support for the above to be executed on distributed memory CPU and GPU clusters in conjunction with MPI.

The OP2 strategy for building executables for different back-end hardware consists of firstly generating the architecture specific code by parsing the user code (which is written using the OP2 API) through a pre-processor and then secondly linking the generated code with the appropriate parallel implementation library. Figures 3 and 4 illustrate and contrast the build process for generating a single threaded CPU executable and a CUDA based executable respectively. For the single thread CPU executable the user’s main program (in this case `jac.cpp`) uses the OP header file `op_seq.h`

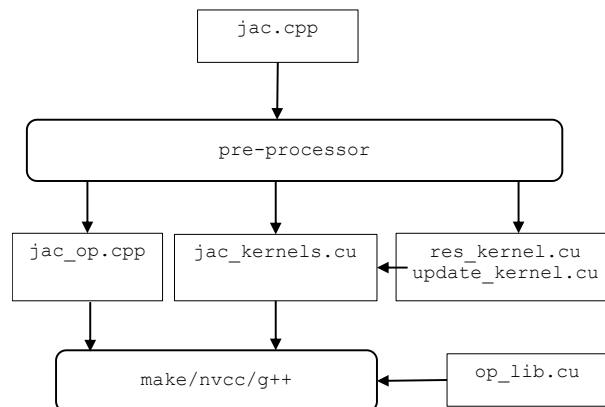


FIGURE 4. CUDA build process

and is linked to the OP routines in `op_seq.c` using `g++`, controlled by a Makefile. For the corresponding CUDA executable, the preprocessor parses the user’s main program and produces a modified main program and a CUDA file which includes a separate file for each of the kernel functions. These are then compiled and linked to the OP routines in `op_lib.cu` using `g++` and the NVIDIA CUDA compiler `nvcc`, again controlled by a Makefile. The pre-processor is currently implemented as a source to source translator using the ROSE compiler framework [17].

4.1. Parallelising Unstructured Mesh Applications

A key technical difficulty in the parallel execution of unstructured mesh applications is the data dependency issue encountered when incrementing indirectly-referenced arrays. Thus, for example, a potential problem arises when two edges update the same node. A solution at a coarse grained level would be to partition the nodes such that the *owner* of the nodal data would carry out the computation. The drawback in this case is redundant computation when the two nodes for a particular edge have different *owners*. At the finer grained level, we could assign a “color” for the edges so that no two edges of the same color update the same node. This allows for parallel execution for each color followed by a synchronization. The disadvantage in this case is a possible loss of data reuse and loss of some parallelism. A third method would be to use *atomic* instructions, which combine read/add/write into a single operation. However, atomic operations (especially on doubles) are not present on all the hardware platforms we are interested in (atomics are an optional extension in OpenCL), and performance varies between platforms.

The OP2 design addresses the data dependency problems using the above three methods, at three levels of parallelism. Method 1 will be used in the future at the MPI level. Given a global mesh with sets and data, OP2 will partition the data so that the partition within each MPI process owns some of the




FIGURE 5. An example coloring of edges

set elements i.e. some of the nodes and edges. These partitions only perform the calculations required to update their own elements. However, it is possible that one partition may need to access data which belongs to another partition; in that case a copy of the required data is provided by the other partition. This follows the standard “halo” exchange mechanism used in distributed memory message passing parallel implementations. As the partition size becomes larger, the proportion of “halo” data becomes very small.

For distributed memory architectures the partition size is large. However, within a CPU or a GPU, operations are to be performed on a finer granularity on each processing unit. For a multi-core CPU the processing units are processor cores (each based on a traditional heavy-iron architecture such as x86 or IBM POWER). For NVIDIA GPUs the processing units are a number of relatively low powered stream multiprocessors (SMs) each consisting of a number of stream processors (SPs) that share control logic, an instruction cache and a block of shared memory [16]. The new NVIDIA Fermi architecture also provides an L1 cache per SM.

To solve the data dependency issue on GPUs, indirect data from GPU main memory is loaded into each SM’s shared memory space forming a local mini-block. Each mini-block is assigned to an SM and are executed in parallel. The SPs within an SM, execute the mini-block utilising a number of threads (called a thread block). The threads are executed 32 (called a *warp* in CUDA) at a time in parallel.

During the execution of an indirect loop, the loop receives data from shared memory instead of main memory, maximising data reuse and minimising the traffic between main memory and shared memory. Thus on the GPU, updating the same node could occur either (1) by multiple threads within a single processing unit (an SM) updating data held in its shared memory (i.e. mini-block) or (2) when the result of the shared memory are written back to the main graphics memory which is used by other processing units. In OP2 thread coloring is used for the former and a block coloring is used for the latter.

Edges are colored so that two edges with the same

TABLE 1. CPU node system specifications

Processor	Cores /node	Clock rate	Memory /node
Intel Xeon E5462 (Penryn)	8	2.8GHz	16GB
Intel Xeon E5540 (Nehalem)	8 (16 SMT)	2.5GHz	24GB

color never update the same node (see Figure 5). As a result, the edges with the same color can be processed in parallel by different threads. The coloring is performed very efficiently in a runtime initialization using a bitwise operation on a 32 bit integer for each of the edges [18]. Similarly, a block coloring scheme is used so that results from shared memory, after processing a mini-block, are not used by any other mini-block being processed simultaneously. On a production grade CFD application such as Hydra a single run would consist of over 100K blocks each needing to fit into the shared-memory of a GPU. 10 colors might be needed to avoid data conflicts, suggesting up to 10K blocks per color.

A similar technique is used for multi-core processors. The difference is that now, each mini-block is executed by a single OpenMP thread. The mini-blocks are colored to stop multiple blocks trying to update the same data in the main memory simultaneously. This technique is simpler than the GPU version as there is no need for global-local renumbering (for GPU main memory to shared memory transfer) and no need for low level thread coloring.

5. MULTI-CORE CPU PERFORMANCE

Our first set of experiments is directed at comparing the performance of Airfoil using OpenMP on a single node comprising of multi core, multi-threaded CPUs. This section presents the results for single precision performance on the CPUs. Double precision performance is reported in Section 6. Table 1 details briefly the specifications of each CPU system node. The Intel Xeon E5462 (based on the older Intel Penryn micro-architecture) node consists of two Intel Xeon E5462 quad-core (total of 8 cores) processors operating at 2.8GHz clock rate per core and has access to 16GB of main memory. The Intel Xeon E5540 processor based node, consist of two Intel Xeon E5540 quad-core (total of 8 cores) processors consisting of 2.5GHz per core clock rate and access to 25GB of main memory. These processors are based on Intel’s current flagship Nehalem micro-architecture and have simultaneous multi threading (SMT) enabled for the execution of 16 SMT threads. For brevity and to avoid confusion for the rest of this paper, the Xeon E5462 will be referred to as the Penryn and the Xeon E5540 as the Nehalem. The Airfoil code and OP2 was compiled on both systems using the Intel ICC 11.1 compiler. We used `-parallel -O3 -ipo -vec-report -xSSE2,SSE3,SSSE3,SSE4.1,SSE4.2`

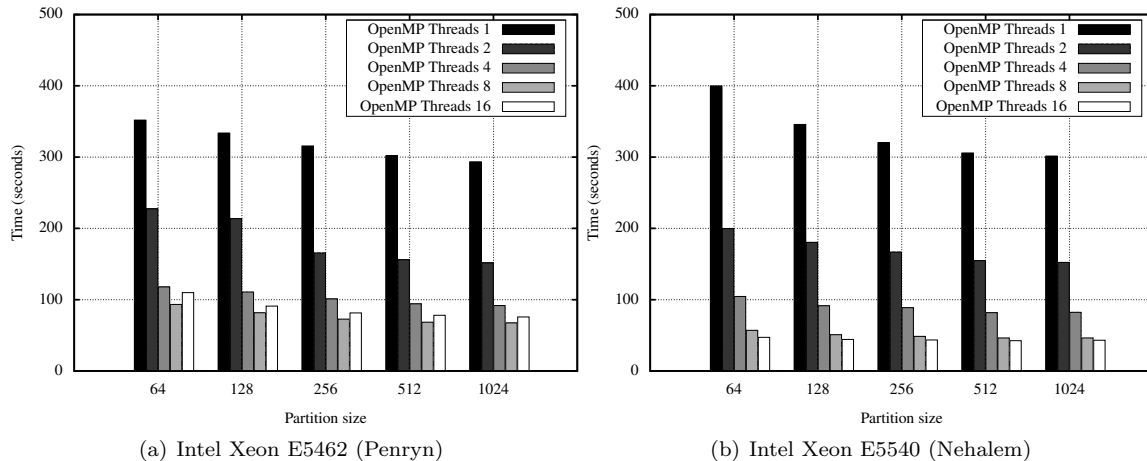


FIGURE 6. Runtime of Airfoil on the Intel processors Compiled with Intel CC 11.1 with up to 16 OpenMP threads - Single Precision (1000 Iterations)

TABLE 2. GPU node system specifications

GPU	Cores	Clock GHz	Global Mem	Shared Mem /SM	Driver /Comp. Cap.
GeForce GTX260	216	1.4	0.8GB	16kB	3.2/1.3
Tesla C2050	448	1.15	3.0GB	48kB	3.2/2.0

-funroll-loops compiler flags for generating optimised code on both systems. Thus for the Nehalem micro-architecture the compiler makes use of the latest SSE4.2 instructions.

As mentioned previously OpenMP parallelism is achieved by OP2 on multi-core processors by partitioning the unstructured mesh assigned to the multi-core node and using one OpenMP thread for each mini-block. Coloring is used to stop multiple mini-blocks interfering with the same data. Thus, a key parameter in our study will be to investigate the mini-block size that provides the best runtime for the Airfoil application for a given unstructured mesh.

Figure 6 presents the total runtime of Airfoil on the Penryn and Nehalem based nodes, compiled using the ICC compiler, for a range of partition sizes using up to 16 OpenMP threads. There is only a marginal difference between the performance of the two systems. Due to the higher clock rate on the Penryn, it exhibits better single core (single thread) performance. However when using all 8 cores we see about 30% better performance from the Nehalem with a best runtime of about 42 seconds on a partition size of 512.

Regardless of the partition size increasing the number of threads from 1 to 8 provides diminishing returns. 16 threads provides an even smaller performance benefit (if any) as two threads share a core using SMT. It appears that other factors of multi-core chips may be limiting their scalability. We have observed a bandwidth utilization of over 20 GB/s on the Nehalem system during the execution of Airfoil. Given that the maximum available bandwidth of these processors

is 25.6GB/s [19] the code appears to be saturating the processor’s bandwidth capacity. We investigate memory bandwidth utilisation in more detail later in Section 7.

6. GPU PERFORMANCE

Next, we explore the performance of the Airfoil code on two NVIDIA GPUs - the consumer grade GTX260 and the HPC-capable Tesla C2050 based on NVIDIA’s current Fermi GPU architecture. The OP2 code transformation framework in this case generates CUDA to be executed on the GPUs. Table 2 details the specifications of each system. The host CPU for the GTX260 is an AMD Athlon X2 dual core processor at 2GHz, while for the Tesla the host is a quad-core Intel Xeon E5530 processor operating at 2.4GHz. In both GPU systems NVIDIA’s CUDA/C compiler nvcc was built using the GNU C compiler 4.4.5.

Figure 7 presents the total runtime of Airfoil (for single precision arithmetic) on the two NVIDIA cards. For these runs the number of CUDA threads allocated per mini-block provides an additional configuration parameter. The GTX260 could only execute partition sizes up to 256 due to its limited memory. The GTX260 performs only about two times slower than the Tesla C2050 due to their comparable single precision floating-point performance. The best performance- just over 12 seconds - on the C2050 is achieved at a partition size of 256 running a thread-block size of 256. This is a speed-up of just under 4× compared to the Intel Nehalem processor systems performance on 16 OpenMP threads.

Given the mesh size, we can approximately compute the single precision floating point performance achieved on both the Nehalem and the C2050 during the most compute intensive loop, `res_calc`. The mesh consists of approximately 1.5 million edges each responsible for 100 floating-point operations in `res_calc`. This routine is in turn called 2000 times giving 30×10^{10} floating-point operations in total. This translates to 15 GFlops




FIGURE 7. Runtime of Airfoil on NVIDIA GPUs on a range of partition sizes and thread-block size configurations - Single Precision (1000 Iterations)

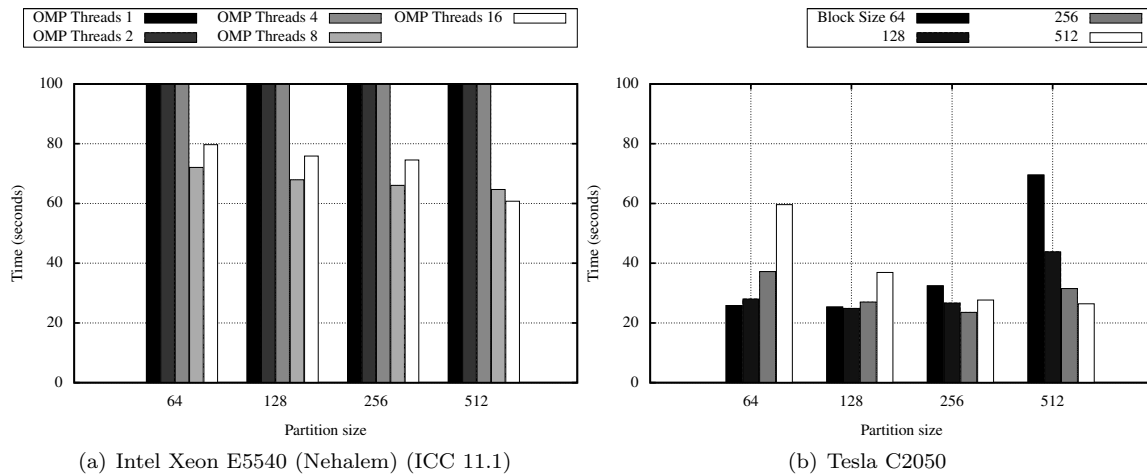


FIGURE 8. Runtime of Airfoil on Intel Xeon E5540 (Nehalem) and NVIDIA Tesla C2050 - Double Precision (1000 Iterations)

per second on the Intel Nehalem processor based system and about 30 GFlops per second on the Tesla C2050. Thus we see that only a fraction of the peak single-precision floating-point performance on the GPU is achieved [20].

A key concern in determining whether GPUs are suitable for main-stream HPC and production scientific work is how their performance compares against the traditional processors when executing double precision floating-point codes. Because this has been an increasing concern for the adoption of GPUs, NVIDIA has invested heavily in improved double-precision floating-point performance on their current Fermi based GPUs. The key benchmarking study for us therefore is to investigate the Airfoil execution in double-precision. Figure 8 details the double precision performance of Airfoil on the Intel Nehalem and Tesla C2050. The results demonstrate a speed-up of approximately $2.5\times$ on the Tesla C2050 compared to the best runtime on the Intel Nehalem.

The best overall runtimes according to Figure 8 are achieved when the problem is configured to be executed with a partition size of 512 on the Nehalem and a

partition size of 256 using a thread block size of 256 on the C2050. However breaking down the runtime into the time taken by the five parallel loops reveal that the optimum partition size and the thread block size differs for each loop. Thus for example `res_calc` routine runs best when configured with a partition size of 256 and a thread block of 256 on the C2050, while `update` is optimised at a partition size of 64 and a thread block of 64. Similar behaviour can be observed for the Nehalem runs. Thus it is apparent that further runtime improvements could be gained by simply configuring each parallel loop to be executed on its optimum partition size and thread block size settings. The ability to infer the optimum configuration could be gained through historical runtime observation, through a performance model or utilising an auto-tuning mechanism. We are currently investigating the implementation of an auto-tuning mechanism within OP2.

7. PERFORMANCE ANALYSIS AND OPTIMISATION

Considering the best execution time for each parallel loop the total runtime for the Airfoil code on the

TABLE 3. Best runtimes for each parallel loop in Airfoil

Nehalem			
Loop	Part.	Time (S)	BW(GB/s)
save_soln	128	2.62	24.57
adt_calc	64	14.80	8.75
res_calc	512	28.11	14.52
bres_calc	64	0.17	4.35
update	64	10.70	21.07
Total		56.40	-

C2050				
Loop	Part.	Block.	Time (S)	BW(GB/s)
save_soln	512	64	0.80	80.68
adt_calc	256	256	3.60	45.20
res_calc	256	256	13.14	19.20
bres_calc	64	64	0.12	6.18
update	64	64	4.41	50.10
Total			22.07	-

Nehalem is about 56.4 seconds, while for the C2050 this is about 22.07 seconds. Table 3 details the time for each loop in conjunction with their observed memory bandwidth utilization for this case. The bandwidth figures were obtained by counting the total amount of data transferred with main memory during the execution of a parallel loop and dividing it by the runtime of the loop. Currently, only the total amount of useful data transferred is counted. In the future this will be augmented to include bandwidth utilisation due to cache line loading for specific CPU/GPU architectures.

As mentioned before, from observing the memory bandwidth figures on the Nehalem it is apparent that during the most time consuming parallel loop (`res_calc`) over 50% of the maximum available bandwidth (25.6GB/s [19]) on the processor is utilised. Other loops such as `update` gets much closer to saturating the available memory bandwidth on the CPU. Thus we suspect that the memory bandwidth of the single node system may become the bottleneck limiting future thread scalability of multi-core CPUs. The memory bandwidth utilisation on the Tesla C2050, also comes close to the available upper limit of 144GB/s [20], on `save_sol` but remains relatively low on `res_calc` suggesting that the higher compute intensity of this loop eases bandwidth saturation.

Given that each element on the unstructured mesh could be computed independently, it is not surprising that the GPU architecture out-performs the traditional multi-core processors. But it is interesting that only about a factor of $4\times$ speed-up is achieved in single precision and only about $2.5\times$ speed-up gained in double precision. One potential cause may be due to the cost of integer pointer arithmetic in computing indirect references in unstructured mesh computations. As there is no separate integer pipeline on the simple GPU cores, an integer operation costs as much as a floating-point operation (at least in single precision). The remainder of this paper investigates

TABLE 4. C2050: Thread coloring & L1 cache performance

Loop	Original Time (S)	Coloring OFF (% diff)	L1 OFF (% diff)
save_soln	0.80	0.49	-6.47
adt_calc	3.60	1.37	-4.87
res_calc	13.14	23.52	-1.91
bres_calc	0.12	2.41	-3.16
update	4.41	0.88	-2.97
Total	22.07	14.43	-2.78

further the performance of the Airfoil code in an attempt to uncover insights into optimising the OP2 framework. We target the GPU architecture for our analysis/efforts in order to demonstrate the expert techniques required to maintain performance on this emerging architecture. Our goal is to demonstrate that using a framework such as OP2 does not sacrifice performance improvements that can be gained by low-level optimisations/configurations available for a target back-end architecture. At the same time we demonstrates the significant benefits of such a framework which generates these improvements without the intervention of the applications programmer.

7.1. Thread Coloring

We first quantify the performance degradation (if any) due to the thread coloring mechanism in OP2. Recall that thread coloring is an important method for avoiding dependency conflicts on both the CPU and GPU. However executing partitions via thread coloring has its disadvantages, such as possible loss of data reuse and loss of some parallelism. By switching off the coloring mechanism within a thread block for execution on the GPU we can observe the idealised scenario of an execution on the GPU without the above potential disadvantages. However it should be noted that in this case, the mathematical solution of the CFD code becomes inaccurate.

Column 3 of Table 4 presents the percentage performance difference when the Tesla C2050 executes the Airfoil code without thread coloring. The maximum performance difference is in the `res_calc` loop which execute 23% faster without thread coloring. The overall performance difference is about 14%, which suggests that there is further room for improvement of the thread coloring algorithm for the GPU architecture.

7.2. The L1 Cache on the C2050

The inclusion of an L1 cache per SM is an important addition to the NVIDIA GPU architecture. Column 4 of Table 4 presents the performance difference of the runtime when L1 is switched off (negative values represents worse performance). The performance effects appears to be minor, due to the fact that the current OP2 library utilise the shared memory block per SM during the execution of indirect loops such as `res_calc`. For the solution of production grade codes such as

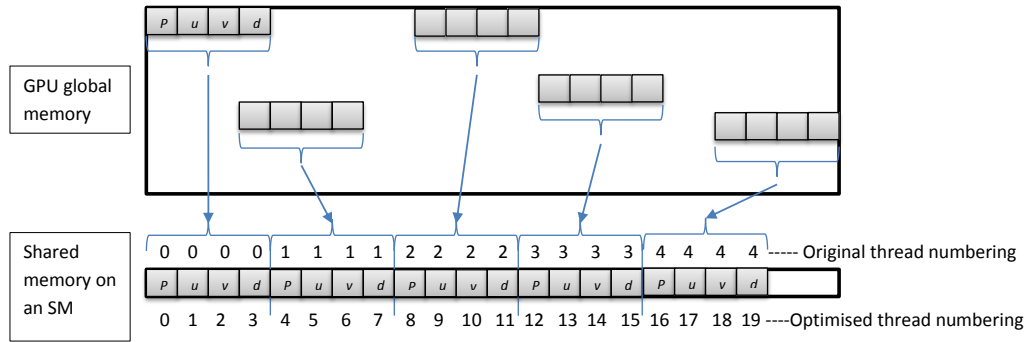


FIGURE 9. Indirect loop optimisation

Hydra on GPUs the number of registers per thread required will significantly increase causing register spillage into main memory. Thus we believe that in the long term the L1 cache is probably best left unused so that the memory can instead be used for register spillage in place of the slower main memory. This is a particular feature of the Fermi GPU and the above results reveal that the current OP2 implementation is already in an optimum set up to handle such a scenario.

7.3. Improving Indirect and Direct Loop Execution

In this final section we implement two optimisations to improve the performance of the execution of direct and indirect loops and explore their performance. For direct loops, OP2 does not need to handle dependency issues and currently are fetched directly from main memory of the GPU. Thus utilising shared memory for direct loops could provide some performance gains, particularly when the L1 cache is switched off. This forms our first optimisation.

On the other hand, recall that the execution of indirect loops on the GPU is done by firstly loading indirect data from GPU main memory into an SM's shared memory. The second optimisation looks into the benefits of using a different thread numbering scheme to implement coalescence for fetching data from GPU main memory in to shared memory.

Figure 9 illustrates the current and optimised thread numbering schemes in relation to the execution of `res_calc`. The *global* data array consists of the indirectly accessed data, that is used in the computation within `res_calc` where each element consists of four floating-point values representing the pressure (P), velocities in the 2 dimensions (u, v) and density (d). The data used within a particular mini-block of the unstructured mesh is scattered across the GPU global memory as illustrated in the figure. The current implementation utilises one thread per element to access the data. Thus thread 0 will access all four values, P, u, v, d of the first element, thread 1 will access all four values, P, u, v, d of the next element and so on. Such an access pattern causes a performance loss due to strided memory fetches by each thread. The

TABLE 5. C2050: Optimised direct and indirect loops

Loop	Original Time (S)	L1 ON (% diff)	L1 OFF (% diff)
<code>save_soln</code>	0.80	26.29	26.22
<code>adt_calc</code>	3.60	3.28	-1.15
<code>res_calc</code>	13.14	12.10	11.71
<code>bres_calc</code>	0.12	13.98	15.14
<code>update</code>	4.41	15.44	15.55
Total	22.07	11.85	10.92

optimised version will use a different thread (up to the maximum number of threads in the thread block with wrap around) per floating-point value to fetch data. Now, each element is copied as a single coalesced memory access in to the shared memory, followed by a call to `_syncthreads` for synchronising the threads.

Table 5 details the combined performance gains of both the above optimisations. The percentage improvement over the original implementation is reported both with and without the L1 cache. We see over 25% and 15% performance gain for the direct loops `save_soln` and `update` respectively, while the indirect loops `res_calc` and `bres_calc` shows about 11% and 15% gains respectively. The effect of the L1 cache remains negligible.

We believe that further improvements could be made to indirect and direct loops. Future work will look into the use of the `float4` data type to load four floating-point values directly to registers within an SM bypassing the shared memory and removing the `_syncthreads` statement to eliminate thread synchronisation overheads. Padding the shared memory to avoid memory bank conflicts when a CUDA warp accesses shared memory will also be investigated.

8. CONCLUSIONS

We have presented an early performance analysis of the OP2 “active” library, which provides an abstraction framework for the solution of unstructured mesh applications. OP2 aims to decouple the scientific specification of the application from its parallel implementation to achieve code longevity and near-optimal performance through re-targeting the back-end to different hardware. OP2's code transformation framework was used to generate back-end code for a

significant CFD application, targeting multi-threaded executables based on OpenMP and NVIDIA CUDA. The performance of this code was benchmarked during its solution of a mesh consisting of 1.5 million edges on Intel multi-core/multi-threaded CPUs (Penryn and Nehalem) and NVIDIA GPUs (GTX260 and Tesla C2050).

Performance results show that for this application the Tesla C2050 performs about $4\times$ and $2.5\times$ better in single precision and double precision mathematics respectively compared to two high end Intel multi-core processors executing 16 OpenMP threads. These results suggest competitive performance by the GPUs for this class of applications at a production level, but we have also highlighted key concerns, such as memory bandwidth limitations on multi-core/many core architectures at increasing scale, which can limit the achievable performance.

Finally, we presented performance comparisons for a number of low-level hardware and software re-configurations and optimisations on the GPU architecture. Our efforts demonstrate the skilled techniques needed to maintain performance on an emerging architecture such as GPUs and point to the significant benefits of utilising the OP2 expert framework which does this for novice programmers developing unstructured mesh applications.

The full OP2 source and the Airfoil test case code are available as open source software [6] and the developers would welcome new participants in the OP2 project.

ACKNOWLEDGEMENTS

This research used HPC resources at the Oxford Supercomputing Centre (OSC). We are thankful to Leigh Lapworth, Yoon Ho and David Radford at Rolls-Royce, Carlo Bertolli at Imperial College London, Jamil Appa and Pierre Moinier at BAE Systems, Tom Bradley at NVIDIA and Nick Hills at the University of Surrey for their contributions to the OP2 project.

REFERENCES

- [1] Crumpton, P. I. and Giles, M. B. Multigrid Aircraft Computations Using the OPlus Parallel Library. *Parallel Computational Fluid Dynamics: Implementations and Results Using Parallel Computers*. 339-346, A. Ecer, J. Periaux, N. Satofuka, and S. Taylor, editors, North-Holland, 1996.
- [2] Burgess, D. A., Crumpton, P. I., and Giles, M. B. (1994) A Parallel Framework for Unstructured Grid Solvers. In Wagner, S., Hirschel, E., Periaux, J., and Piva, R. (eds.), *Computational Fluid Dynamics '94: Proceedings of the Second European Computational Fluid Dynamics Conference*, pp. 391–396. John Wiley and Sons.
- [3] Campobasso, M. S. and Giles, M. B. (2003) Effect of flow instabilities on the linear analysis of turbomachinery aeroelasticity. *AIAA Journal of Propulsion and Power*, **19**, 250–259.
- [4] Moinier, P., Muller, J. D., and Giles, M. (2002) Edge-based multigrid and preconditioning for hybrid grids. *AIAA Journal*, **40**, 1954–1960.
- [5] Giles, M. B., Duta, M. C., Muller, J. D., and Pierce, N. (2003) Algorithm developments for discrete adjoint methods. *AIAA Journal*, **42**, 198–205.
- [6] Giles, M. OP2 for Many-Core Platforms. <http://people.maths.ox.ac.uk/gilesm/op2/>.
- [7] Giles, M. B., Ghate, D., and Duta, M. C. (2008) Using Automatic Differentiation for Adjoint CFD Code Development. *Computational Fluid Dynamics Journal*, **16**, 434–443.
- [8] Howes, L. W., Lokhmotov, A., Donaldson, A. F., and Kelly, P. H. J. (2009) Deriving efficient data movement from decoupled access/execute specifications. In Seznec, A., Emer, J., O’Boyle, M., Martonosi, M., and Ungerer, T. (eds.), *High Performance Embedded Architectures and Compilers*, Lecture Notes in Computer Science, **5409**, pp. 168–182. Springer Berlin/Heidelberg.
- [9] Chafi, H., DeVito, Z., Moors, A., Rompf, T., Sujeeth, A. K., Hanrahan, P., Odersky, M., and Olukotun, K. (2010) Language Virtualization for Heterogeneous Parallel Computing. *Proceedings of the ACM international conference on Object oriented programming systems languages and applications*, New York, NY, USA OOPSLA ’10, pp. 835–847. ACM.
- [10] HMPP workbench. <http://www.caps-entreprise.com/>.
- [11] DeVito, Z., Joubert, N., Medina, M., Barrientos, M., Oakley, S., Alonso, J., Darve, E., Ham, F., and Hanrahan, P. Liszt: Programming Mesh Based PDEs on Heterogeneous Parallel Platforms. Presentation given by the Stanford PSAAP Center, Oct 2010 <http://psaap.stanford.edu>.
- [12] Corrigan, A., Camelli, F. F., Löhner, R., and Wallin, J. (Jun 2009) Running Unstructured Grid-based CFD Solvers on Modern Graphics Hardware. *19th AIAA Computational Fluid Dynamics Conference*, Grand Hyatt Hotel, San Antonio, Texas, pp. 1–11.
- [13] Asouti, V. G., Trompoukis, X. S., Kampolis, I. C., and Giannakoglou, K. C. (to appear) Unsteady CFD Computations Using Vertex-Centered Finite Volumes for Unstructured Grids on Graphics Processing Units. *International Journal for Numerical Methods in Fluid*, **n/a**.
- [14] Giles, M. OP2 User Manual. <http://people.maths.ox.ac.uk/gilesm/op2/user.pdf>.
- [15] Giles, M. Hydra. <http://people.maths.ox.ac.uk/gilesm/hydra.html>.
- [16] CUDA C Programming guide. http://developer.download.nvidia.com/compute/cuda/3_1/toolkit/docs/NVIDIA_CUDA_C_ProgrammingGuide_3.1.pdf.
- [17] The ROSE Compiler. <http://www.rosecompiler.org/>.
- [18] Giles, M. OP2 Developer’s Manual. <http://people.maths.ox.ac.uk/gilesm/op2/dev.pdf>.
- [19] Intel Xeon Processor E5540 specifications. <http://ark.intel.com/Product.aspx?id=37104>.
- [20] NVIDIA Tesla C2050 / C2070 GPU Computing Processor. http://www.nvidia.com/object/product-tesla_C2050_C2070_us.html.