

Overset Grid Based Flow Computations for Rotary Wing Flows on GPU Architectures

Dominic Chandar

dchandar@uwyo.edu

Postdoctoral Research Associate

University of Wyoming

Laramie, WY, U.S.A

Jay Sitaraman

jsitaram@uwyo.edu

Assistant Professor

University of Wyoming

Laramie, WY, U.S.A

Dimitri Mavriplis

mavripl@uwyo.edu

Professor

University of Wyoming

Laramie, WY, U.S.A

ABSTRACT

The advent of General Purpose Graphics Processing Units (GPGPUs) has spawned a lot of interest in computing high resolution flows in a shorter span of time. In combination with existing parallel programming techniques such as MPI or openMP, one is able to obtain at least one order increase in speed-up for CFD applications on stationary grids. Due to a high throughput/cost ratio, GPUs are increasingly becoming popular among the CFD community. Flow fields on helicopter systems rank among one of the most challenging and hence are computationally demanding to simulate. Further, presence of multiple bodies moving relative to each other require use of overset grid systems, which in turn require efficient overset grid assembly methods to support the flow solution. In this context, a computational framework for flow simulation across multiple CPU cores and multiple GPUs has been developed and tested. It is shown that the overall wall-clock time for the simulation can be considerably reduced by using multiple GPU cards. The parallel GPU framework is initially tested using both explicit and implicit schemes for the flow past a sphere. Further, a two-blade hovering rotor test case adopted from literature is used to demonstrate the capability of the code towards simulating rotorcraft flows in a short span of time.

INTRODUCTION

Accurate representation of the flow-field around a helicopter is essential in assessing the overall performance especially for new designs and especially rotor-body interactions. Computational methods require higher degrees of freedom to achieve a reasonable level of accuracy for engineering decisions to be made in the design phase. This in turn requires huge computational resources (parallel computers) which then translate into optimizing the cost to accuracy ratio. Currently the capability to simulate rotorcraft flows across parallel processors (CPU cores) do exist (c.f. (Ref. 1), (Ref. 2), (Ref. 3)) and are being widely used for predicting rotor loads. Computer hardware has been evolving rapidly during the last couple of years. With the advent of General Purpose Graphics Processing Units (GPGPUs), the Computational Fluid Dynamics (CFD) community has absorbed this technology and several codes have been developed either from scratch ((Ref. 4), (Ref. 5), (Ref. 6)) or by using an automatic CPU to GPU porting strategy ((Ref. 7), (Ref. 8)). The rotorcraft community also began to embrace the GPU ideology as described by the computations in (Ref. 9), (Ref. 10), (Ref. 11). However, most the existing computational codes that are used to simulate rotorcraft flows (especially using overset/overlapping grids) do not have the capability to run across multiple CPU cores and multiple GPUs simultaneously. Studies on using this hybrid

CPU-GPU framework for CFD applications are limited and hence this paper is geared towards developing such a framework and understanding how effectively one can use all possible computing resources to solve a given problem.

The present framework has been evolved from the authors' previous work CU++ (Ref. 12) which is a framework for solving partial differential equations on structured grids on GPUs. CU++ provides the user a higher level programming framework especially for finite-difference calculations where standard discretized expressions can directly be used in place of standard GPU *kernels* (discussed in the next section) similar to C++. However, for unstructured grid based codes, only the data structures of CU++ are used in conjunction with GPU *kernels* as stencil operations are not straightforward. Also developed using CU++ was GPUINS (Ref. 13), a GPU based three-dimensional incompressible Navier-Stokes solver on moving overset grids. The code runs on a single GPU and was verified using standard incompressible flow test cases such as flow past a sphere and oscillating wing under low Reynolds number flow conditions. The GPU code was able to achieve a speed-up of 5 to 15 compared to a serial CPU code depending on the level of serial CPU code optimization. In this work, a hybrid CPU-GPU framework has been developed for solving the Compressible Euler equations on overset grids. The hybrid feature enables the code to run simultaneously on a GPU and several CPU cores. Although the viscous terms in the current implementation can be turned on, only inviscid test cases using the compressible solver are demonstrated at the moment to verify the capabilities of the GPU based im-

plementation. This paper is organized as follows: In the following section, a short description of the flow algorithm on a per CPU core/GPU basis is described followed by GPU/CPU hardware implementation specifics. Parallel partitioning using METIS and GPU-CPU load balancing is described next followed by results and discussions on parallel performance.

FLOW SOLVER ALGORITHM

The compressible Euler equations are discretized in space on an unstructured grid using a cell-centered finite volume method, and a second order Roe’s scheme for interface fluxes. Time marching is performed using (1) an explicit low storage three stage Runge-Kutta scheme (Ref. 14) and (2) an implicit second order backward difference scheme(BDF2) using a Newton-Krylov approach. For the implicit scheme, as a part of the linearization of the fluxes, exact Roe flux Jacobians are used in conjunction with a BiCGSTAB Krylov solver without preconditioning. Typically 4-5 Newton iteration steps per time step are required for convergence of the non-linear system of equations.

GPU Implementation

Implementing the entire algorithm on the GPU has been very straightforward as the framework and relevant data structures were already developed for the Incompressible flow code (Ref. 13) and two-dimensional Euler code (Ref. 6). Programming the GPU using NVIDIA’s Compute Unified Device Architecture (CUDA) (Ref. 15) requires the use of *kernels* which are similar to *functions* in standard programming languages such as Fortran or C++. Each *kernel* call will spawn a number of threads depending on the type of loop involved (loop over faces or cells), and each of this thread will perform its own computation independently thereby maximizing parallel performance. The relevant *kernels* used in the present computation are as follows:

- Gradients using Green-Gauss and Fluxes: One kernel computes the face centered values (loop over faces) and another kernel sums up the face centered values and assigns it to a cell center (loop over cells).
- Runge-Kutta update: One kernel to update the field variables based on different Runge-Kutta stage constants(loop over cells).
- Newton-Krylov update: Usage of NVIDIA Thrust library for computing reductions in the BiCGSTAB procedure and simple DAXPY kernels to update the solution vector (loop over cells)
- Overset connectivity: kernels for donor search and classification(loop over cells).

HYBRID GPU-CPU PARALLEL IMPLEMENTATION

As mentioned previously, this hybrid computational framework utilizes available GPUs and CPU cores for solving a

given problem. Another advantage of this hybrid framework is that the exact same code can run in parallel without GPUs without any modification to the code thereby making it more portable for non GPU based platforms. For proper load balancing, the computational grid must be partitioned in such a way that processes that run on the GPU have a higher load compared to the processes that run on CPU cores. The entire partitioning procedure is shown in Figure 1 and consists of the following steps:

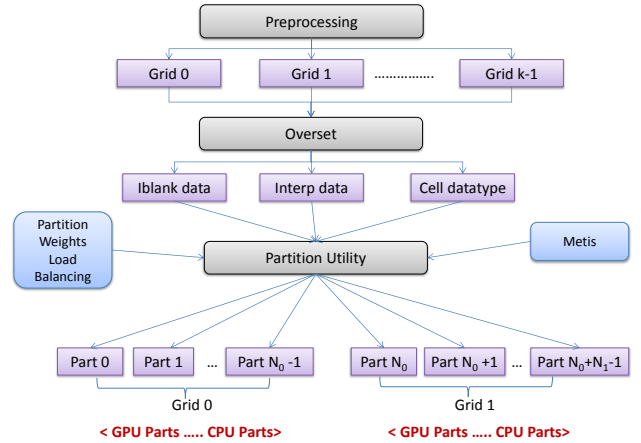


Fig. 1. Grid partitioning procedure

1. For the implementation in the current paper, the overset connectivity is performed offline and the respective data written into different files. For each component grid (*not partitioned grid*), we write three pieces of information. (a) field/fringe/hole cell classification flag into the iblack file, (b) fringe and the corresponding donor cell into the interp file and (c) overset donors and overset fringe cells into the cell datatype file.
2. As there are multiple grids, the domain partitioning must be performed for each of these grids separately. This indicates that no partition will contain more than one grid. The graph partitioning code *METIS* (Ref. 16) is used to classify the cells into various partitions. As GPU partitions require more load, a *weight* file is passed as an input to *METIS* to create weighted partitions. The partition utility calls *METIS* for each of the grids and adds ghost cells across each partition. Figure 2 shows a portion of the partitioned grid for the hovering rotor test case. At the end of this step, the partition utility arranges the cells in a specific order for each partition as follows:
 - Cells that do not need any communication
 - Cells on a regular partition boundary that act as donors to ghost cells on neighboring partitions.
 - Overset donor cells
 - Overset fringe cells (recipient cells)

- Ghost cells that receive data from neighboring partitions.

Grouping the cells in this fashion enables efficient data transfer between the CPU and GPU for MPI based communication.

3. The partition utility then writes out various partition grid files which also have all the necessary information relating to communication between various processes. It also generates a process mapping file which enables mapping the partitioned grids to different processes. For each grid, the partitions are arranged such that the first few partitions are GPU partitions and the remaining are CPU partitions.





Fig. 2. A section (along span and chord) of the partitioned grid for the rotor test case.

When the main flow solver runs on np processes, each process will read its respective grid based on the mapping between process ID and partition ID. Each process would know whether it is a GPU based process or a CPU based process based on the input partition file. For all GPU based processes, the relevant GPU *kernels* are called to perform various computations and for CPU based processes, the corresponding C++ functions are called. When communication is required between various processes, only the GPU based processes would transfer data between the GPU and CPU memory.

Load Balancing GPU and CPU partitions

Since a GPU runs considerably faster than a CPU core, it is important that GPU based partitions have more data to process than CPU based partitions. The ratio of the data between GPU and CPU based partitions depends on how fast a GPU performs with respect to a single CPU core. Let us assume that, T is the total problem size (e.g. number of cells), s the speed-up of a GPU relative to one CPU core, n_g the number of GPUs, n_c the number of additional CPU cores (apart from those controlling a GPU), N_1 the partitioned problem size on the GPU and N_2 the partitioned problem size on the CPU. Based on these variables, the total problem size can be expressed as:

$$T = n_g N_1 + n_c N_2 \quad (1)$$

For the load to be balanced between a CPU core and GPUs, one must have:

$$N_1 = s N_2 \quad (2)$$

The above is only an assumption, as GPU speed-ups may not always vary linearly with grid size as the GPU performance depends on a wide range of parameters such as number of blocks, number of threads per block, number of registers per *kernel* and the GPU hardware itself. Using Eq.(2) in Eq.(1), we obtain

$$N_1 = \frac{sT}{n_g s + n_c} \quad (3)$$

$$N_2 = \frac{T}{n_g s + n_c} \quad (4)$$

Eqs.(3) and (4) represent the load-balanced problem size on respective(GPU/CPU) partitions. To compute that however, one needs to know the quantity s , which represents the speed-up of one GPU relative to one CPU core. Hence, the problem is solved initially without any partitioning to compute the speed-up of the GPU code. Typically for the rotor test case when N (Number of cells per partition) $\geq 2 \times 10^5$, $s \approx 14$. For smaller grid sizes ($N \approx 5 \times 10^4$), $s \approx 10$. One can also estimate the theoretical speed-up gained by using additional CPUs as follows:

If only GPUs are used, then the time spent by each GPU is roughly $t_{GPU1} \approx \frac{T}{n_g}$. By using additional CPU cores with load balanced, the time spent by a GPU is $t_{GPU2} \approx N_1$. Thus

the speed-up gained by using additional CPU cores is given by:

$$t_{GPU1}/t_{GPU2} = 1 + \left(\frac{n_c}{n_g}\right) \frac{1}{s} \quad (5)$$

Note that the above expression does not involve any communication time. Most of the computations reported in this paper were carried out on the NCAR-WY super-computing cluster. A subset of this cluster has 32 NVIDIA M2090 GPUs across 16 compute nodes.

RESULTS AND DISCUSSIONS

Inviscid subsonic flow past a sphere at Mach 0.3

To verify the proper implementation of the code, the flow past a sphere at nearly compressible speeds is computed, and the results are compared with inviscid incompressible flow theory with a compressibility correction factor. Figure(3) shows the contours of the pressure coefficient C_p along with a sectional plot at $z = 0$ for a single GPU computation using the explicit RK3 method in comparison with the corresponding theoretical estimates $C_p = (1 - \frac{9}{4}\sin^2\theta)(1 - M_\infty^2)^{-1/2}$. Also shown in Figure (4) is a multi-GPU (24 GPUs mapped by 24 CPUs) computational result using the implicit BDF2 method. Two distinct facts can be stated by looking at these plots. (1) The contours across the overset/partitioned boundary(indicated by the stepped edges) are continuous and (2) a good correlation between computed and theoretical estimates are obtained thereby establishing confidence in the current implementation.

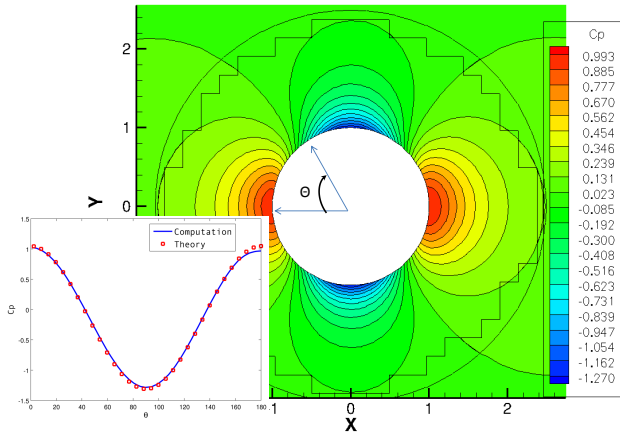


Fig. 3. Contours of the pressure coefficient and sectional pressure coefficient variation (using explicit RK3 time stepping and a single GPU) in comparison with theoretical estimates for the inviscid flow past a sphere at Mach 0.3




Fig. 4. Contours of the pressure coefficient for the inviscid flow past a sphere at Mach 0.3 using 24 GPUs and implicit BDF2 time stepping

Inviscid transonic flow past a rotor in hover

This test case is based on the experimental work of Caradonna and Tung (Ref. 17) and aids in the validation of the code for rotorcraft applications involving moving grids. The rotor consists of two untwisted, untapered wings of NACA0012 profile with an aspect ratio of 6 as shown in Figure (5). For initial computation purposes, there is no hub hence its interaction with the rotors is not accounted for. Several test cases corresponding to various rotational speeds(or tip Mach numbers) and collective pitch settings have been computed. The only difference in the implementation between the sphere test case and the rotor test case lies in the flux expressions where grid speed terms and source terms are added to account for the rotational reference frame. The governing equations are solved in the rotating reference frame but expressed as quantities in the inertial frame; cf. (Ref. 18). All computations in this section were computed using an explicit RK3 time stepping scheme as the grid speed terms have not yet been accounted for in the Roe flux Jacobian. For a collective pitch of 5° we prescribe an angular velocity of $\omega_y = -0.1359$ about the y -axis corresponding to a tip Mach number of 0.815. The entire grid consists of 1.1M cells with a combination of hexahedra and prisms. Using a single GPU, the computed and measured pressures coefficients ($C_p = \frac{p - p_\infty}{0.5\rho_\infty u_t^2}$ (u_t =sectional linear velocity)) at two different sections, $Z = 3$ and $Z = 5.34$ are compared in fig.(6). A reasonable comparison can be obtained with some visible differences due to the absence of viscous effects in the computation. Also plotted in the same figure are (1) surface pressure contours indicating the formation of a shock towards the tip of the rotor and (2) Y -velocity contours showing the downwash from the wing tips. All the preceding comparisons with experimental data are encouraging and leaves us with a lot of room for further development.

For parallel computations, the total number of cells was in-

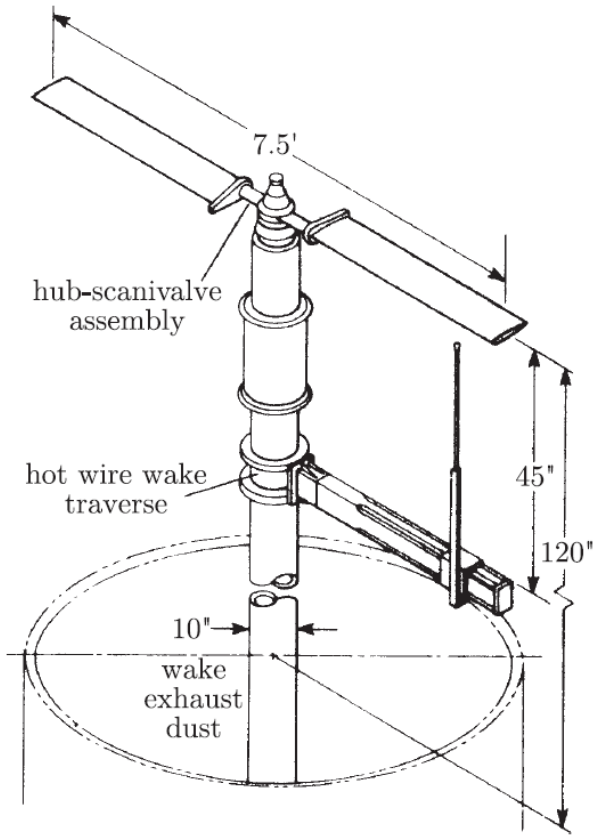


Fig. 5. Experimental setup for the hovering rotor test case, adapted from Caradonna and Tung (Ref. 17)

creased to $2M$ so that the partitioned grids that run on the GPU have sufficiently larger load. Figures (7) and (8) show the pressure contours on a section along the wing tip and the inviscid rotor wake respectively. The partition boundaries are also shown to demonstrate that the contours are continuous across these boundaries and that the communication between various partitions has taken place without any issues. As the background grid is relatively coarse in comparison with the rotor grid (Figure 2), the contours across the overset boundary have some discontinuities.

PARALLEL GPU PERFORMANCE

To assess the performance of the parallel code, computations are performed across a range of CPU cores (8, 12, 16, and 20) and GPUs (1, 4, 8, 12, 16, and 20). For example, if 12 GPUs are used in conjunction with 16 CPU cores, there are 12 CPU cores controlling 12 GPUs and 4 CPU cores running independent tasks. Hence n_g and n_c in Eq. 1 are 12 and 4 respectively. The compressible flow past a sphere using the explicit RK3 scheme described earlier is considered for computing the wall-clock times under different conditions. Overall performance is assessed based on the parallel performance of the code relative to 8 CPU cores. Figure (9) shows contours of

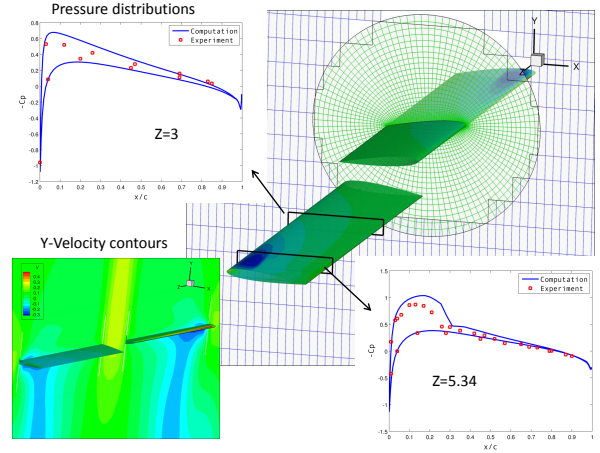


Fig. 6. Computed pressure distributions and Y-velocity contours for the hovering rotor test case on a single GPU

speed-up relative to the performance on 8 CPU cores. The plot is very interesting in the sense that it gives us an idea on how many CPU cores can replace a GPU. For example, take a base case with 8 CPU cores and 0 GPUs. If we need a 2x improvement over 8 CPU cores, we would move along the x-axis and locate a point with the same contour value. This indicates that we need roughly 17 CPU cores. However one may also traverse vertically along the Y-axis and attain the same contour value with 2 GPUs and 8 CPU cores. The same can be said for other combinations of CPU cores and GPUs.

Best CPU Performance Vs. Best GPU Performance

For a given grid size, it would be beneficial to know how much potential one could extract from an array of GPUs in comparison with a standard parallel implementation that uses only CPU cores for a realistic test case such as the hovering rotor described earlier. To estimate this, we first run the code without any GPUs and estimate the wall-clock time for different number of processes. Following this run, we use as many GPUs as the number of processes so that each GPU is mapped by one CPU core. Although the maximum number of GPUs present in the NCAR-WY Supercomputing Caldera cluster is 32, only 24 could be used at a time due to job scheduling restrictions. Figure (10) shows a summary of the wall-clock times for ten iterations in *seconds* for various runs. We can see that the best CPU performance is attained using 126 CPU cores (14 cores \times 9 nodes) and the best GPU performance using 24 GPUs (2 cores \times 12 nodes). If cost for each run is computed as $Number\ of\ Cores \times wall\ -\ clock\ time \times (cost(\$)\ per\ core = k)$, we can see that the best GPU performance is 13 times economical (in terms of cost) than the best CPU performance. The hybrid GPU-CPU implementation is not included due to the fact that when the number of GPUs used are high, a higher number of CPU cores must be used for proper load balancing (according to Eq.(5)). For 24 GPUs, to get at least 50% increase in the speed-up, we need an additional 120




Fig. 7. Computed pressure contours along $Z = -5.8$ for the hovering rotor test case on 18 GPUs

CPU cores. Also for a $2M$ grid the partitioned size of the grid would be small for the GPU to produce any significant speed-up and that communication times may be in the same order of magnitude as that of computation times. Hence for the hybrid GPU-CPU framework to produce meaningful results, the grid has to be very large.

CONCLUSIONS

We have presented an overset grid based flow computation strategy on multi-GPU, multi-core architectures. Successful application and validation was demonstrated for the (1) sphere test case and (b) Caradonna and Tung rotor. It is also demonstrated how the hybrid nature of the code, where both GPUs and CPU cores are active provide additional speed-up than cases that involve only GPUs. In future, we plan to:

- Extend this framework so that the overset connectivity would be performed in parallel. This would pave way for solving problems that involve relative mesh motion such as rotor-body/rotor-rotor interaction.
- Improve the discretization procedure for increased accuracy of the simulations
- Implement viscous effects and
- Run the codes on the Kepler GPU architecture (Ref. 19) to take advantage of NVIDIA *GPUDirectTM* - a framework to transfer data between various GPUs that lie across different compute nodes in a cluster without the CPU interference.

ACKNOWLEDGMENTS

We gratefully acknowledge the support from the Office of Naval Research under ONR grant N00014-09-1060 and the

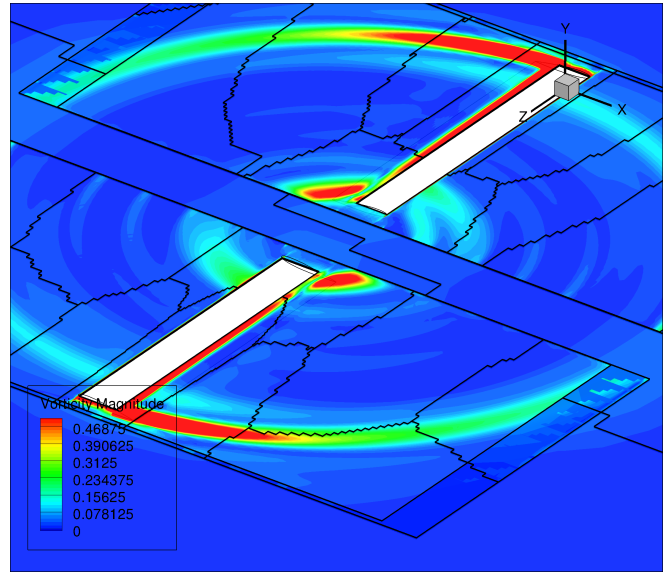


Fig. 8. Computed vorticity(magnitude) contours along $Y = 0$ for the hovering rotor test case on 18 GPUs

NCAR-Wyoming Supercomputing facility for the computational resources.

REFERENCES

- ¹Sankaran, V., Sitaraman, J., Wissink, A., Datta, A., Jayaraman, B., Potsdam, M., Mavriplis, D., Yang, Z., O'Brien, D., Saberi, H., Cheng, R., Hariharan, N., and Strawn, R., "Application of the Helios Computational Platform to Rotorcraft Flowfields", Paper AIAA 2010-1230, 48th AIAA Aerospace Sciences Meeting and Exhibit, Orlando, FL, Jan 2010.
- ²Schwarz, T., "Simulation of the Unsteady Flow Field Around a Complete Helicopter with a Structured RANS Solver", *High Performance Computing on Vector Systems*, Springer Verlag, Heidelberg, 2006, pp. 125–137.
- ³Allen, C. B., "Parallel Flow-Solver and Mesh Motion Scheme for Forward Flight Rotor Simulation", Paper AIAA 2006-3476, 24th AIAA Applied Aerodynamics Conference, San Francisco, CA, 2006.
- ⁴Cohen, J. M., and Molemaker, M. J., "A Fast Double Precision Code using CUDA ", *Proceedings of Parallel CFD*, Moffett Field, CA, 2009.
- ⁵Hagen, T. R., Lie, K-. A., and Natvig, J. R., "Solving the Euler Equations on Graphics Processing Units", *Lecture Notes in Computer Science*, Vol. 3994, 2006, pp. 220–227.
- ⁶Soni, K., Chandar, D., and Sitaraman, J., "Development of an Overset Grid Computational Fluid Dynamics Solver on Graphical Processing Units", *Computers and Fluids*, Vol. 58, 2012, pp. 1–14.

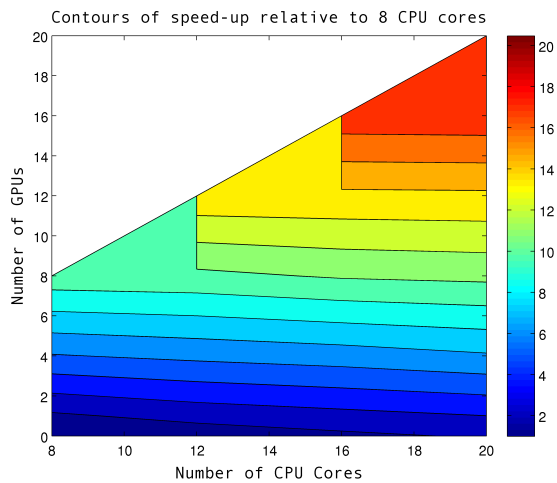


Fig. 9. Performance of the parallel GPU code relative to 8 CPU cores

| Number of CPU Cores/GPUs | 6 | | 12 | | 24 | | 48 | | 96 | | 126 | | 150 | |
|--------------------------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|
| | Time | Cost | Time | Cost | Time | Cost | Time | Cost | Time | Cost | Time | Cost | Time | Cost |
| CPU Only | 14.1 | 85k | 7.4 | 88k | 4 | 96k | 2.4 | 115k | 1.28 | 122k | 1 | 126k | 1.35 | 202k |
| GPU | 1.33 | 8k | 0.73 | 8.7k | 0.4 | 9.7k | - | - | - | - | - | - | - | - |

Fig. 10. Wall-Clock times and total cost under different run conditions

⁷Corringan, A., Camelli, F., Lohner, R., and Mut, F., “Semi-Automatic Porting of a Large-Scale Fortran CFD Code to GPUs”, *International Journal for Numerical Methods in Fluids*, Vol. 69 (6), 2011, pp. 314–331.

⁸Poole, D., “Introduction to OpenACC Directives”, *NVIDIA GPU Technology Conference*, 2012.

⁹Lohry, M. W., Ghosh, S., and Rajagopalan, R. G., “Graphics Hardware Acceleration for Rotorcraft Brownout Simulation”, Paper AIAA 2011-3224, 20th AIAA Computational Fluid Dynamics Conference, Honolulu, HI, 2011.

¹⁰Thomas, S., Amiraux, M., and Baeder, J. D., “A GPU Accelerated Navier-Stokes Solver for Rotorcraft Applications”, American Helicopter Society 69th Annual Forum, Phoenix, AZ, May 2013.

¹¹Stock, M. J., Gharakhani, A., and Stone, C. P., “Modeling Rotor Wakes with a Hybrid OVERFLOW-Vortex Method on a GPU Cluster”, Paper AIAA 2010-4553, 28th AIAA Applied Aerodynamics Conference, Chicago, IL, 2010.

¹²Chandar, D., Sitaraman, J., and Mavriplis, D., “CU++ET: An Object Oriented Tool for Accelerating Computational Fluid Dynamics Codes Using Graphical Processing Units”, Paper AIAA 2011-3222, 20th AIAA Computational Fluid Dynamics Conference, Honolulu, HI, 2011.

¹³Chandar, D., Sitaraman, J., and Mavriplis, D., “Dynamic Overset Grid Computations for CFD Applications on Graphics Processing Units”, Paper ICCFD7-12-2, Proceedings of the International Conference on Computational Fluid Dynamics, Big Island, Hawaii, 2012.

¹⁴Kennedy, C. A., Carpenter, M. H., and Lewis, R. M., “Low-Storage, Explicit Runge-Kutta Schemes for the Compressible Navier-Stokes Equations”, NASA/CR 1999 209349, 1999.

¹⁵NVIDIA CUDA C programming Guide 5.0, <https://developer.nvidia.com/what-cuda>

¹⁶Karypis, G., and Kumar, V., “A Fast and Highly Quality Multilevel Scheme for Partitioning Irregular Graphs”, *SIAM Journal on Scientific Computing*, Vol. 20 (1), 1999, pp. 252–399.

¹⁷Caradonna, X. C., and Tung, C., “Experimental and Analytical Studies of a Model Helicopter Rotor in Hover”, NASA TM 81232, 1981.

¹⁸Biedron, R. T., and Thomas, J. L., “Recent Enhancements to the FUN3D Flow Solver for Moving-Mesh Applications”, Paper 2009-1360, AIAA Aerospace Sciences Meeting and Exhibit, Orlando, FL, 2009.

¹⁹NVIDIA’s Next Generation CUDA Compute Architecture, KEPLER G110, <http://www.nvidia.com/content/PDF/kepler/NVIDIA-Kepler-GK110-Architecture-Whitepaper.pdf>