BRINGING PARALLEL ADAPTIVE SIMULATION TO INDUSTRY

Mark S. Shephard

Johnson Professor, Rensselaer Polytechnic Institute, Troy, NY, USA

Cameron Smith

Computational Scientist, Rensselaer Polytechnic Institute, Troy, NY, USA

Jean-Francois Dord

Research Associate, Rensselaer Polytechnic Institute, Troy, NY, USA

Kenneth E. Jansen

Professor, University of Colorado, Boulder, CO, USA

THEME

High Performance Computing

KEYWORDS

High Performance Computing, parallel adaptive analysis, computational fluid dynamics

SUMMARY

Consideration is given to the issues associated with increasing the ability of industry to perform complex simulations using adaptive analysis tools operating on massively parallel computers. After a brief review of some activities underway supporting the ability of industry to apply massively parallel computing, the key components of parallel mesh generation and mesh adaptation tools that operate directly with CAD representations, parallel dynamic load balancing procedures and parallel finite element solvers are overviewed. From that point a description of a set of industrially relevant parallel adaptive flow simulation examples are discussed.

1: Introduction

Advances in hardware and algorithms have provided many orders of magnitude improvement in the ability to perform large-scale simulations. As the plans to move to exascale computing proceed forward (Crosscut, 2010; ExaScale Software Study, 2009; Roadmap, 2010), the inability to effectively increase CPU clock rates requires all truly large-scale computations be performed on massively parallel computers. These future massively parallel computers will be far more heterogeneous and therefore much more complex to program. On the positive side, progress on the development of next generation massively parallel computers is leading to systems that are much more cost effective in terms of computer system, electric power and operational costs. This means an increased ability to cost effectively employ the most computationally intense simulations in engineering design processes assuming we can address the needs to provide the needed software tools and methods of applying them.

A review of the current U.S. Council on Competitiveness publication list (Council, 2010) clearly indicates that the application of simulation using high-performance computing is critical to industrial innovation. The set of high performance computing cases studies listed demonstrates the importance of high performance computing across all industrial sectors. What is not obvious in those case studies is that the level of computation being used for the majority of the engineering simulations is far from that needed and that current high performance computing can provide. Closer examination of the engineering problems being addressed indicates that in an increasing number of cases (e.g., flow, electromagnetics, materials processing, etc.) the resolution of the models and discretizations used is not as high as the engineers would want to ensure the reliability of the simulation results and the simulations being applied are at a single scale ignoring the innovation made possible by performing multiscale simulations. For example in an April 2009 case study (Whirlpool, 2009) a 168-processor system is being applied to support a major manufacturer’s high performance computing needs. Although this case study does demonstrate the impressive gains through the use of this level of computing, that system is less that one-thousandth of the 294,912 processing cores machine that has been applied to a single simulation using an open source unstructured mesh CFD code (Jansen, et al., 2010; Zhou, et al. 2010). Although one may argue that machines with 100,000’s of processing cores are well beyond what industry would obtain (in the near future), industry will easily be able to justify next generations of massively parallel machines with 10,000’s of cores due to the continued dramatic decreases in machine costs and power requirements over the current systems. In addition, through opportunities like the DOE INCITE program, industry will have access to machines with nearly one million processing cores for use on their largest problems.

What is currently missing is robust CAE software that runs effectively on large scale parallel computers, where "runs effectively" means the software maintains excellent scaling to at least 10's of thousands of processors. The gap between the available parallel CAE software and what is needed is particularly large when considering the fact that to support the effective use of simulation-based engineering to support innovation the simulation tools must be fully automated and provide reliable results when used by design engineers. It is well known that adaptive methods are capable of providing greatly increased levels of simulation robustness and reliability. Thus, in addition to the procedures that numerically solve a given discretization (e.g., finite element mesh) running on the parallel computer, the procedures that take the CAD model and create the discretization (e.g., automatic mesh generator) and those that automatically improve the mesh (e.g., adaptive mesh control) must also be able to run on the same parallel computers at acceptable levels of scaling. The need for all the tools to operate on the same parallel computers is driven not only by the size of the data, it is driven by Amdahl's (1967) law which ensures that unless all steps in the process are scaled, an unacceptable bottleneck in the process will form. Note that these sources of such potential bottlenecks include both the execution of specific tasks and the transfer of information between tasks.

This paper present a brief status report on a set of efforts to provide industry with large scale parallel simulation tools and technologies. The next section overviews the overall status of these with an emphasis on a set of activities supported by New York State. Two sections after that focus on the specific area of parallel, automated, adaptive unstructured mesh simulation tools and their use on problems of industrial interest.

2: Moving Parallel Simulation from Research to Industry

Computer-aided design (CAD) and computer aided engineering (CAE) technologies have become a cornerstone in the design and manufacture of a broad range of products ranging from consumer products, to aerospace vehicles, to medical devices. The key CAD/CAE tools being used in these processes include geometric design, analysis model generation, engineering analysis, scale linking and visualization tools. However, the further advancement of these technologies for most applications requires the ability to execute coordinated simulations involving multiple physical behaviors acting over multiple spatial and temporal scales. Such simulations must interactive with extensive heterogeneous data sets. A major bottleneck preventing industry from addressing these new simulations is the lack of a full set of simulation components that run on massively parallel computers. This problem has been compounded by the fact that the software from the major CAE ISVs does not yet scale well on massively parallel computers and most of these ISVs continue to employ pricing policies that are simply not realistic for use by industry on problems requiring massively parallel computation.

The DOE, with programs like SciDAC (2010) and PSAAP, has been actively developing new generations of software that can effectively operate on massively parallel computers. These developments include specific simulation tools focused on DOE applications and general tools that support developing new simulations. Four key tools of this include the Portable, Extensible Toolkit for Scientific Computation (PETSc (2010)) equation solvers, the Trilinos (2010) object-oriented software components for constructing solutions to large-scale, complex multi-physics engineering and scientific problems, ZOLTAN (2010) parallel partitioning service, and the ITAPS (2010) interoperable interface components for mesh, geometry, and field access. These primarily open source software packages are beginning to receive increased attention by industry and, to some extent, ISVs.

In addition to the advancement of some of the software technologies, there are a limited number of programs that support industries' use of massively parallel computers. The DOE INCITE (2010) program allows industry to compete for allocations, and associated support, on the major DOE supercomputer systems. In addition a few state governments are working to help industry in applying parallel computing in their product development processes. An example of this type is the High Performance Computing Consortium (HPC^2 (2010)), supported by the NY State Foundation for Science, Technology and Innovation. HPC^2 supports computational scientists to work directly with NY industry to apply massively parallel simulations on supercomputer systems, supported by NY State.

Some CAE oriented ISVs have begun to develop new generations of software that employ data infrastructures and algorithms that allow them to operate on massively parallel computers. In addition, much of this software is designed as components that interact through easy to use interfaces such that users can easily combine components from multiple sources to meet their simulation needs.

3: Parallel Adaptive Unstructured Mesh Simulation

Unstructured mesh finite volume and finite element solvers have the distinct advantage of being able to solve problems over completely general space/time domains of interest to industry using meshes that can be automatically generated and anisotropically adapted to effectively provide the level of solution accuracy desired. Thus, unstructured mesh methods are the emphasis of the parallel adaptive methods discussed in this paper. The steps of an automated adaptive unstructured mesh simulation are: (i) Create a geometry-based problem definition in terms of a non-manifold solid model supplemented with the analysis attributes of loads material properties and boundary conditions. (ii) Create an initial mesh using an automatic mesh generator, meshes are generated directly from CAD representation (Beall, et al. 2004). (iii) Perform the analysis on the given mesh. (iv) Apply error estimation

(Ainsworth and Oden, 2000; Babuska and Strouboulis, 2001) correction indication procedures to determine where and how to improve the mesh. (v) Improve the mesh by altering it to obtain the mesh sizes indicated by the correction indication procedures (Shephard, et al., 2005). (vi) Repeat the steps (iii)-(v) until the desired level of estimated accuracy is obtained.

Most of the effort to date has focused on the development of scalable analysis procedures, step (iii), considering both explicit and implicit methods for the solution to the global equation systems. Although more difficult to scale, implicit solvers are of particular interest because of their relative advantage on many problems of interest to industry. Recently it has been shown that, when given well balanced mesh partitions, such solvers can scale well on massively parallel computers. In particular the PHASTA code has shown excellent strong scaling to 294,912 cores on Blue Gene/P and 98,304 cores on Cray XT5 on general anisotropic meshes (Jansen, et al., 2010; Sahni, et al., 2009). Parallelization of the error estimation and correction procedures is reasonably straight forward since they are dominated by element level information and at most require communication of information for neighboring elements.

Parallel initial mesh generation, particularly directly from CAD geometry, is particularly challenging due to the inherent complexity of parallel algorithms for those operations (deCougny and Shephard, 1999) and the complexity of interacting with the CAD representation. However, since there are likely times where even the initial mesh will need billions of elements, having parallel mesh generators that operate in parallel is critical. Although the number of efforts in this area are limited, there are commercial parallel automatic mesh generators (Simmetrix 2010) that have generated meshes of 180 million in under 10 minutes on 64 cores (file writing time not included).

Methods for the parallel mesh adaptive modification of unstructured meshes (deCougny and Shephard, 1999a) based on interaction with mesh partitions and dynamic load balancing (ZOLTAN 2010) have been defined. These procedures have been executed in parallel on up to 65,536 cores (Zhou, et al., 2010).

A critical aspect of the development of effective parallel simulation tools is a software infrastructure that will support the effective combinations of simulation components. Neither library or framework approaches are well-suited to supporting the use of existing simulation components, such as an existing fixed mesh PDE solver, in the development of complete simulation procedures such as an automated adaptive simulation procedure. An alternative approach that has been gaining momentum is the use of interoperable components that interact through functional interfaces. An example of the use of functional interfaces in support of unstructured mesh simulations is the development of automatic mesh generation and mesh adaptation procedures that interact directly with the geometric modeling kernels of CAD systems (Beall, et al., 2004). Another example is the interoperable interfaces for unstructured meshes being developed by the ITAPS center (Devine, et al. 2009; ITAPS, 2010). The interfaces being developed by ITAPS include access to CAD geometry, unstructured meshes and simulation fields. In addition to the basic interfaces the ITAPS center provides a number of services useful to unstructured mesh operations including element shape optimization, mesh adaptation, front tracking and mesh-to-mesh solution transfer.

4: Application of Parallel Adaptive Simulation on Industrially Relevant Problems

Each of the examples in this section are being jointly developed by computational scientists and researchers from RPI for industry as part of the NY State High Performance Computing Consortium.

Flow over an unmanned air vehicle (UAV) can produce very complex flows. The ability to simulate flows is further complicated by the introduction of synthetic jets that are used to improve the flow by either reducing or eliminating the separation. Because the flow is laminar over much of the airfoil and because the jets are unsteady, large eddy simulation is necessary to capture the effect of the jets. In addition, the simulation must fully account for the detailed geometry of each of the individual synthetic jets of which there are hundreds in specific applications of interest.

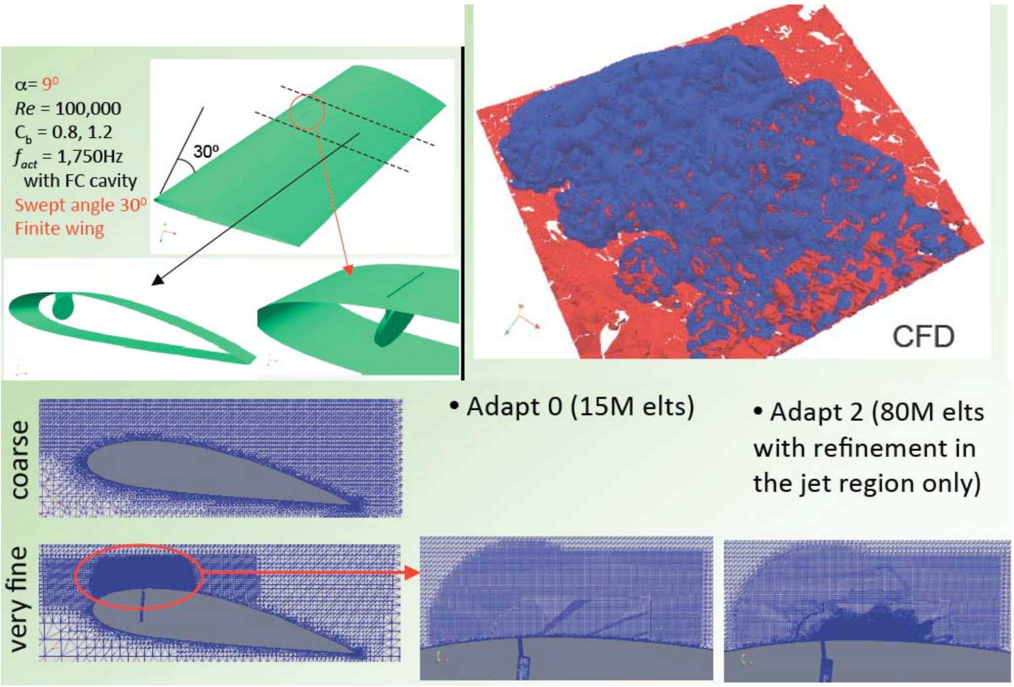

Figure 1: A swept wing at 9 degrees angle of attack produces a separated flow at Reynolds number 100k. Synthetic jets are used to improve the flow. The model, initial mesh, adapted mesh and isosurfaces of vorticity from the unsteady solution are shown.

Figure 1 shows some results and meshes for a simulation performed of flow over a full, swept wing at 9 degrees angle of attack. At this angle of attack the wing produces a separated flow at Reynolds number 100,000. The model (upper left), initial mesh (middle left), adapted meshes with close-up (lower row) and isosurfaces of vorticity (upper right) from the unsteady solution are shown in Figure 1. The initial simulations were performed on 2048 processors but, as the adaptivity made the mesh finer, the number of processors was increased in proportion to the mesh size with the final solution performed on 12,288 processors. Since the solver shows nearly perfect scaling the simulations times in both cases is nearly identical.

A second example being done for industry involves the modeling and simulation of fluid ejected from a nozzle and the subsequent formation and evolution of fluid droplets. A piezo-electric device drives the ejection of the fluid from a fluid filled chamber. A goal of this project is to demonstrate the ability to employ parallel adaptive simulation in a mode where fast design iterations can be carried out. As a demonstration of the workflow that is executed to perform these simulations, consider the geometric model depicted in Figure 2 where a fluid will be injected at the left end of the structure through a nozzle. In the initial state of the simulation the nozzle is filled with fluid up to the air domain (flat interface); both the fluid and the air are still in the initial state.

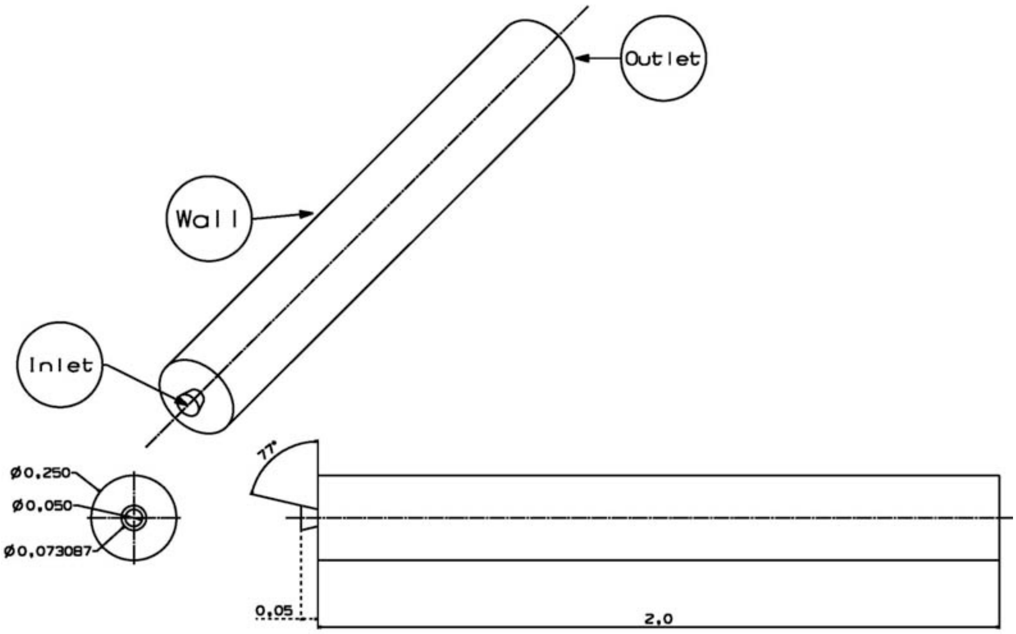

Figure 2: Geometric model for workflow demonstration of two-phase flow problem. Units are in mm.

The fluid surface is modeled by a level set method implemented in the computational fluid dynamics software PHASTA (Nagrath, et al, 2005). At first the flow is pushed uniformly until the volume of fluid outside the nozzle is large enough at which point the inlet surface flow is stopped and the ejected fluid moves due to its inertia.

Figure 3. shows half the 3D domain at four points in the simulation. Close to the surface, the anisotropy factor in the mesh at the fluid/air interface is 8 (the characteristic length of the element in the normal direction is 8 times smaller than that in the tangent directions); this is shown by the elongated elements close to the interface. The simulation calls for adaptive meshing and, in this case, the mesh was adapted every 100 iterations of the solver (that is approximately after the interface moved by 12 layers of anisotropic elements).

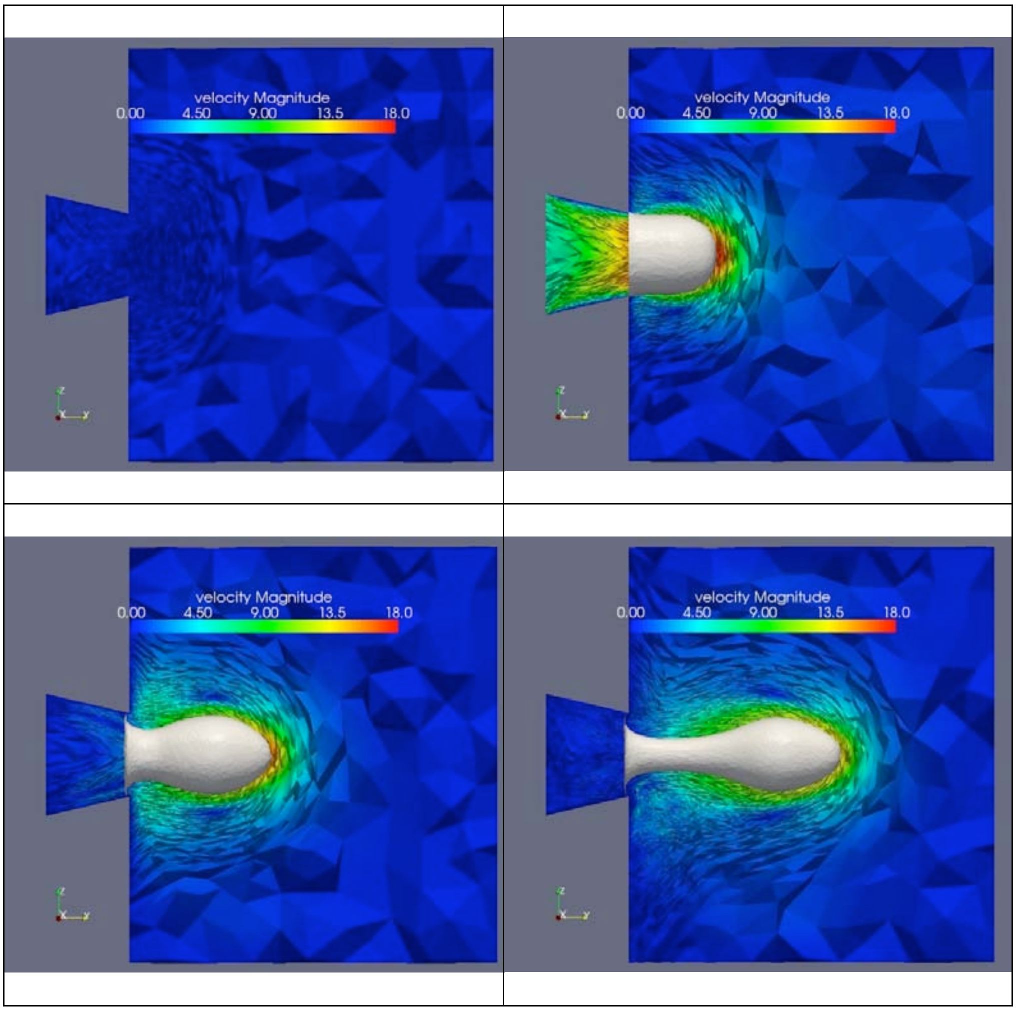

Figure 3: Four steps in the fluid ejection simulation. An axial slice of the mesh is colored by the magnitude of velocity with the interface between the fluid ejecting from the nozzle and air depicted by the gray surface.

The adapted meshes in this example included up to 480,129 elements. A uniform mesh with similar resolution would have over 14,000,000 elements. The current automated adaptive simulation executes in under 8 hours on the order of 200 processors which is easily in the range of what industry can provide its users. Since the analysis program shows nearly perfect scaling and the mesh adaptation procedures scale have reasonable scaling to at least 20,000 processors, the same simulation can easily be performed in half an hour.

Another ongoing industrial interaction is focused on the modeling and simulation of free surface flows that are a part of a manufacturing process. A sheet of viscous fluid is flowing in the direction of gravity while being pulled near the bottom of the sheet. The ACUSIM (1010) software suite's problem definition, execution, and post processing tools, and its 3D free surface modelling capability, have successfully demonstrated the ability to effectively simulation this flow.

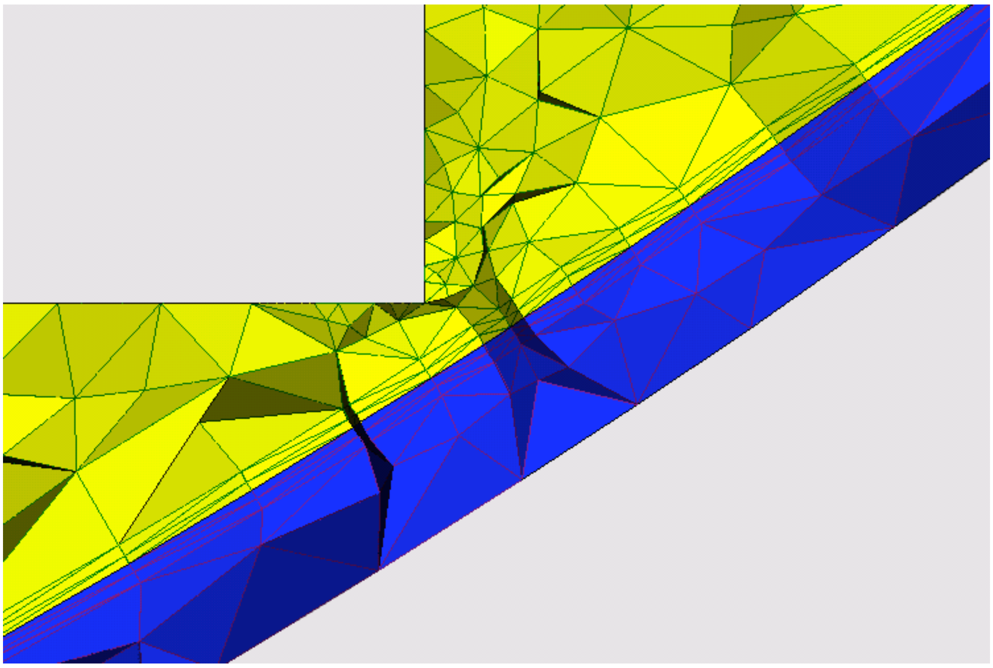

Figure 4: 2D slice of 3D mesh; the intersected mesh entities are shown. Mesh size controls on geometric model edges and boundary layer growth from geometric model faces.

The final example demonstrates the ability to apply these methods to complex geometric models with rotating geometry and extreme characteristic dimensions. For instance and besides the standard requirements for reasonable-sized and smooth-transitioned mesh, the fluid domain had to be meshed knowing that at some point the gap between the impeller and the housing is 1/1000th the largest dimension of the pump. To this end, the Simmetrix MeshSim and SimAppS tools have been used as they support mesh size controls defined on geometric model entities and mesh controls for growing and blending boundary layers from geometric model faces. This combination of controls facilitates the creation of meshes with the requisite resolution for high fidelity simulations. Figure 4 depicts a coarse mesh generated by both growing boundary layers on the interface between rotating components and setting mesh controls on the nearby geometric model edges and surfaces. (The actual mesh required for accurate simulations is substantially finer, but not shown because the local mesh features are not easily discerned.)

5: Closing Remarks

With recent advances in computing hardware the cost to purchase and operate massively parallel computer systems is dropping to the point that industry will be able to perform reliably simulations that provide results of the level of accuracy needed for many applications that can not be properly addressed today. As these systems continue to become available the key problem areas will become the lack of industrial grade software and the people that can effectively apply that software. This paper has provided a brief overview of a set of efforts on the development and application of the types of parallel adaptive simulation tools that will be needed to allow industry to take the next step forward in the application of simulation in the design of next generation products. The software discussed is aimed at allowing design engineers to perform reliable automated simulations, starting from the CAD design, with software that runs effectively on massively parallel computers.

The example applications presented in this paper are industrially relevant flow simulations that companies have specifically indicated that are unable to reliably execute to the required degree of accuracy with most currently available CAE software on the computing systems to which they now have access. Through the combined use of new generations of parallel CAE tools and access to larger parallel computers, each of the examples can reliably solved to the desired levels of resolution in timeframes where designers can use the results to make design decisions. These examples are only the start to what can be done if new generations of parallel adaptive simulation tools for multiscale/multiphysics simulations are developed.

REFERENCES

ACUSIM (2010) Software web page, http://www.acusim.com/

Ainsworth, M. and Oden, J.T. (2000) *A Posteriori Error Estimation in Finite Element Analysis*, Wiley.

Amdahl, G.M. (1967) 'Validity of the single-processor approach to achieving large scale computing capabilities', *AFIPS Conference Proceedings* vol. 30, AFIPS Press, Reston, Va., pp. 483-485.

Babuska, I. and Strouboulis, T. (2001) *The Reliability of the FE Method*, Oxford Press.

Beall, M.W., Walsh, J. and Shephard, M.S (2004) 'A comparison of techniques for geometry access related to mesh generation', *Engineering with Computers*, 20(3):210-221.

Council (2010) Council on Competitiveness publication web page, http://www.compete.org/publications/

Crosscut (2010) 'Crosscutting Technologies for Computing at the Exascale', DOE, http://extremecomputing.labworks.org/crosscut/report.stm, June 2010.

de Cougny, H.L. and Shephard, M.S. (1999) 'Parallel volume meshing using face removals and hierarchical repartitioning', *Comp. Meth. Appl. Mech. Engng.*, 174(3-4):275-298.

de Cougny, H.L. and Shephard, M.S. (1999a) 'Parallel Refinement and Coarsening of Tetrahedral Meshes', *Int. J. Numer. Meth. Eng.,* 46:1101-1125.

Devine, K., et al. (2009) 'Interoperable Mesh Components For Large-Scale, Distributed-Memory Simulations', *J. Phys.: Conf. Serie*s, 180, 012011, 11pg.

DOE INCITE (2010) http://www.science.doe.gov/ascr/incite/index.html

DOE SciDAC (2010) Scientific Discovery through Advanced Computing, http://www.scidac.gov/.

ExaScale Software Study (2009), "ExaScale Software Study: Software Challenges in Extreme Scale System, DARPA Information Processing Techniques Office (IPTO). http://users.ece.gatech.edu/mrichard/ExascaleComputingStudyReports/ecss%20report%20101909.pdf, Sept. 2009.

HPC^2 (2010) NY State High Performance Computing Consortium, http://hpc2.org/

ITAPS (2010) 'Interoperable Technologies for Advanced Petascale Simulations (ITAPS) Center', http://www.itaps-scidac.org/

Jansen, K.E., Sahni, O., Ovcharenko, A., Shephard, M.S. and Zhou, M. (2010) 'Adaptive Computational Fluid Dynamics: Petascale and Beyond', *Journal of Physics: Conference Serie*s.

Nagrath, S., Jansen, K.E. and Lahey, R.T. (2005) 'Computation of incompressible bubble dynamics with a stabilized finite element level set method', *Computer Methods in Applied Mechanics and Engineering*, 194, n42-44 pp. 4565-4587.

PETSc (2010) http://www.mcs.anl.gov/petsc/petsc-as/features/index.html

Roadmap (2010), 'International Exascale Software Project Roadmap', DOE ASCR, http://www.exascale.org/mediawiki/images/a/a1/Iesp-roadmap-draft-0.93-complete.pdf, Jan. 2010.

Sahni, O., Carothers, C.D, Shephard, M.S. and Jansen, K.E. (2009) 'Strong Scaling Analysis of an Unstructured, Implicit Solver on Massively Parallel Systems', *Scientific Programming*, 17:261-274.

Shephard, M.S., Flaherty, J.E., Jansen, K.E., Li, X., Luo, X.-J., Chevaugeon, N., Remacle, J.-F., Beall, M.W. and O'Bara, R.M. (2005) 'Adaptive mesh generation for curved domains', *J. Applied Num. Math.*, 53(2-3)251-271.

Simmetrix (2010), http://www.simmetrix.com/

Trilinos (2010) http://trilinos.sandia.gov/

Whirlpool (April 2009) 'Whirlpool's Home Appliance Rocket Science: Design to Delivery with High Performance Computing', Case Study Council on Competitiveness, http://www.compete.org/publications/.

Zhou, M., Sahni, O., Shephard, M.S., Devine, K.D. and Jansen, J.E. (2010) 'Controlling unstructured mesh partitions for massively parallel simulations', to appear, *SIAM J. Sci. Comp*.

ZOLTAN (2010) http://www.cs.sandia.gov/Zoltan/Zoltan_phil.html