

## **Final Report**

# **HPC Simulation Workflows for Engineering Innovation**

with Application to Complex Industrial Flow Problems

Project Supported by:  
Extreme Science and Engineering Discovery Environment (XSEDE)  
Industrial Challenge Program

Project Executed by:  
Rensselaer Polytechnic Institute  
Scientific Computation Research Center and Center for Computational Innovations  
110 8<sup>th</sup> Street  
Troy, New York 12180-3522

Mark S. Shephard, PI/PD, Rensselaer Polytechnic Institute, [shephard@rpi.edu](mailto:shephard@rpi.edu)

Onkar Sahni, Co-PI, Rensselaer Polytechnic Institute, [sahni@rpi.edu](mailto:sahni@rpi.edu)

Assad A. Oberai, Co-PI, Rensselaer Polytechnic Institute, [oberaa@rpi.edu](mailto:oberaa@rpi.edu)

Cameron W. Smith, Co-PI, Rensselaer Polytechnic Institute, [smithc11@rpi.edu](mailto:smithc11@rpi.edu)

Christopher D. Carothers, Co-PI, Rensselaer Polytechnic Institute, [chris@cs.rpi.edu](mailto:chris@cs.rpi.edu)

John E. Kolb, Co-PI, Rensselaer Polytechnic Institute, [kolbj@rpi.edu](mailto:kolbj@rpi.edu)

## Final Report: HPC Simulation Workflows for Engineering Innovation

M.S. Shephard, O. Sahni, A.A. Oberai, C.W. Smith, C.D. Carothers and J.E. Kolb  
Rensselaer Polytechnic Institute

### Summary of findings

High performance computing hardware and software has advanced to the point that it is possible to address many of the large-scale simulations that are needed to advance science, address national needs and design superior products. The execution of massively parallel simulations at national laboratories and in academia is becoming the norm, while industry, which has a real need to employ these methods, is not applying them on a regular basis. Although the current analysis software is capable of performing the simulations required, it is not integrated into simulation workflows that are of industrial quality and robustness. Thus industrial attempts to use such software tools become labor intensive and end up taking a substantial amount of time. This in-turn makes the process too time consuming and costly for the industry. In addition, industry access to simulation workflows that can be executed on massively parallel computers requires the use of remote systems that are difficult to use.

To help make high performance parallel simulations more viable for industrial application, this XSEDE Industrial Challenge Program addressed the following four areas:

- Optimization of our parallel unstructured mesh-based software components for operation on XSEDE HPC systems [2][3][4][5][8][12]. The key effort involved the implementation of our unstructured mesh infrastructure [2][7] and partition improvement procedures [8] onto the TACC STAMPEDE Intel systems making effective use of the Phi accelerators.
- Definition of an infrastructure to support adaptive unstructured mesh simulation workflows on massively parallel computers and, building on a set of developing components, the implementation of that infrastructure [5][7][9] including the support of the in-memory execution of parallel adaptive simulations [5].
- Development of a science gateway to support the adaptive simulation workflows [9]. This gateway supports the execution of parallel adaptive simulation workflows using the PHASTA CFD code [4] and has been used for simulations executed on TACC STAMPEDE Intel systems. The gateway is currently being installed on the IBM BlueGene/Q at the RPI Center for Computational Innovations.
- Development and demonstration of parallel adaptive simulation workflows of interest working directly with industry [1][4][5][6][9][13].

### Case studies

The focus of the case study efforts was the development of simulation workflows working with various companies. The subsections that follow indicate the results of those efforts. In addition to the uses of XSEDE systems, these case studies employed other parallel computing systems including the IBM Blue Gene/Q and Intel Xeon clusters at the RPI Center for Computational Innovations (CCI), DOE MIRA IBM Blue Gene/Q, and company based clusters.

#### *Blasch Precision Ceramics*

In an ongoing project with Blasch Precision Ceramics and the NYS Pollution Prevention Institute, the conditions inside a Claus process sulfur recovery unit are being simulated. These high temperature reactors are used within the petroleum refinery industry for removing hydrogen sulfide from waste gas streams. Blasch has introduced a product into the reactor market which replaces the traditional stacked-firebrick baffle walls within these enormous reactors (5 to 15 meters long, 1 to 4 meters in diameter). This product (VectorWall™) is a modular ceramic baffle/static mixer system that allows a degree of influence

over flow direction as gas flows through it. Because the VectorWall product has produced surprisingly large increases in reactor throughput in actual commercial installations (up to 40%), an understanding and optimization of how the diverters on the VectorWall ports affect actual flow conditions in this reacting, highly exothermic reactor is of interest.

We have developed parameterized geometric models to represent a range of reactor configurations and variants on the basic VectorWall module diverter. The parameters defining these models can then be updated automatically over a desired range within COMSOL Multiphysics, in order to explore various configurations of a reactor for a given transport and reaction model.

Initial investigations into these reactors have revealed significant effects of the VectorWalls, and their relative location with respect to side streams, on residence time, and thus the resulting level of chemical conversion. This ongoing project has generated two grants from the NYS Pollution Prevention Institute.

### *Corning*

HPC workflow development for Corning Inc. involved two flow solvers. In the first case we tested parallel scaling of the Fluent software on a multi-core system at the CCI. Fluent has been used by Corning for a preliminary examination of a twin-screw extrusion case. In this effort we identified that CPU binding is necessary to achieve good scalability. In the second case, we added two specific features into the PHASTA flow solver needed for the problems of interest to Corning. This includes non-linear partial-slip boundary condition and non-linear material model based on the non-Newtonian Herschel–Bulkley fluid. The PHASTA solver was selected because it has two attractive features: it can handle highly anisotropic grids and it has been shown to scale on massively parallel machines. We also coordinated with Simmetrix in their development of a thin-section meshing capability for the extrusion case. Furthermore, we delivered the PHASTA flow solver to Corning and created and provided tutorials on it.

The XSEDE Science Gateway developed as part of this project [9] supports the execution of the PHASTA based adaptive simulations. Users have found that the application of the gateway greatly simplified the use of remote parallel computing systems.

Based on the RPI efforts on the development of a 3D non-linear partial slip along with a non-Newtonian material model and Simmetrix meshing efforts, we provided Corning results that they indicate are much more accurate and useful than any they have gotten in the past. Because of the level of new technical development involved our joint work with Corning, it has been presented at several conferences and in proceedings.

Rensselaer computational scientists also provided Corning researchers and engineers support installing, running, and tuning the VASP Ab initio, and GROMACS, and LAMMPS molecular dynamics packages on the Rensselaer CCI Blue Gene/Q for materials research and process analysis. Performance tuning identified installation options, problem sizes, and run time environment options that increased simulation efficiency.

### *IBM*

The long-term goal of this project is massively parallel simulation workflows to address integrated circuit simulation needs for problems with 10's of billions of unknowns. Based on an evaluation of the analysis components available, the ALBANY finite element code from Sandia National Labs was selected to be combined with the RPI SCOREC parallel mesh infrastructure tools and Simmetrix mesh generators.

A tool to go from circuit design layouts to the 3D non-manifold model has been developed to produce the geometries that are simulated.

Typical models can be meshed with quadratic elements. Mechanical analysis has been done with ALBANY, a finite element analysis package out of Sandia National Labs, paired with MeshAdapt, SCOREC's mesh adaptation package, and APF, SCOREC's parallel fields package.

An adaptive loop that includes SCOREC parallel mesh infrastructure, Simmetrix geometry and meshing technologies, SCOREC error estimation procedures, and the Sandia Albany FE analysis code has been implemented and tested in non-linear, time dependent material deformations simulations. To address the nonlinear material problems of interest, temperature-activated version of the Norton model of creep, a viscoplastic model in which strain rate is proportional to a power of Cauchy stress was developed. This model is in line with formalism for viscoplasticity as developed in Simo and Hughes and used for  $J_2$  plasticity in ALBANY. We were able to develop a mathematical model that incorporates both Norton creep and  $J_2$  plasticity for finite deformation in 3D. This model was incorporated into ALBANY and has been added back to the open source ALBANY code base at Sandia. Tests of special cases (pure shear, uniaxial tension, etc) have produced expected results and we have run tests for more complex, industrially relevant systems using experimental data from IBM and the scientific literature.

The most recent simulation was a model of a flip-chip cool down with a 1.4B element mesh simulation.

### *ITT Gould Pumps*

Hydraulic engineers at ITT: Goulds Pumps are actively using the Simmetrix simulation based engineering technologies as the basis for an automated pump design workflow put together by SCOREC computational scientists and Simmetrix software engineers. This workflow allows ITT hydraulic engineers to define a pump design analysis on a pump's abstract representation, couple the abstraction to a pump design realized in CAD, then generate an ANSYS CFX analysis input deck. With a minimal number of commands the application of the abstraction and subsequent input deck generation can be repeated to support a design study involving multiple CAD instances and multiple design points. This automation allows ITT hydraulic engineers to focus their efforts on pump design rather than tool usage. Recently, at the request of ITT hydraulic engineers, SCOREC computational scientists and Simmetrix software engineers have extended the workflow to support problem definition variables defined in the ANSYS CFX Workbench and assembly models. These features further reduce the engineer's time needed to setup and run a multi-point design study.

ITT hydraulic engineers can run an automated multi-point ANSYS CFX design study in days instead of weeks by using the Simmetrix and SCOREC developed workflows and compute resources at the RPI's CCI. Automated execution is possible through CCI computational scientist's integration of the SLURM job scheduler with ANSYS CFX Workbench. The SLURM-Workbench coupling implementation has been shared with ANSYS engineers for possible integration into their product. Decreased design study run time is through increased computing resources coupled with the job-schedulers ability to concurrently schedule design-point jobs as system resource availability permits. Using the CCI systems, an evaluation design study was completed in 22 hours; greater than a factor of four savings versus ITT systems. In limited product design windows, this quick turn-around of a design change can be critical to success.

### *Pliant Energy Systems*

The objective of the modeling effort is to create a capability for predicting the complex fluid-structure interaction phenomena observed in the energy harvesting devices designed by Pliant Energy Systems. This problem is especially hard to model due to the (1) high-Reynolds number which implies a turbulent flow field, (2) large motion of the highly flexible structure, and (3) the relatively low mass of the structure which implies that most of the inertia is supplied by the moving fluid.

In order to address these challenges, we have:

- Created a fluid-structure interface that couples a well-benchmarked and validated fluid solver (Acusolve) with an in-house complex, reduced-order structural solver.
- Utilized the moving mesh and augmented Lagrangian-Eulerian (ALE) capability within Acusolve to model the large-scale motion of the structure.

- Exercised a reduced order model which models a thin crimped, highly-flexible sheet as a pre-buckled geometrically nonlinear beam.
- Developed and implemented a staggered algorithm that couples the fluid solver with the structural solver.
- Tested the resulting computational scheme on simple benchmark problems.
- Applied the resulting computational methods to a realistic water tunnel experiment in order to validate the computational model. A good agreement was found.
- Used the computational methods to better understand the mechanism of biological propulsion across organisms of different sizes.

Although much more complex than originally expected, we have developed a new fluid structure interaction analysis method that can model the Pliant devices. The results of these developments have advanced the area of fluid/structure interaction analysis and have provided useful insights into the behavior of the Pliant devices.

### *Sikorsky and Boeing*

Active flow control (AFC) simulations were carried out for Sikorsky and Boeing. For Sikorsky, synthetic-jet based actuation was used on a SC1095 airfoil (a common rotorcraft airfoil), where pulse modulation was found to be effective to mitigate/reduce the hysteresis observed during the dynamic stall [13]. For Boeing, dynamic vortex generator (DVG) was applied to control flow separation on a static airfoil with a deflected flap and it was found to be more effective than static vortex generator (SVG) [4].

Effective active flow control methods were investigated. Pulse modulation of synthetic jets were found to be effective for dynamic stall control and dynamic or periodic actuation of vortex generators was found to be effective for separation control over the flap.

## **Resources**

SCOREC developed open source component tools used in the development of the simulation workflows which include:

- PCU – Parallel Control Utility
- APF – Attached Parallel Fields
- PUMI – Parallel Unstructured Mesh Infrastructure
- GEOM – Geometry Interface
- MeshAdapt – Mesh Adaptation
- ParMA – Partitioning using Mesh Adjacencies

For more information, see the SCOREC software web page (<http://www.scorec.rpi.edu/software.php>) with documentation on the SCOREC GitHub page (<https://github.com/SCOREC/core/wiki>).

Simulation automation components from Simmetrix Inc. used in the development of the workflows include:

- SimModeler – GUI for mesh generation and problem definition
- SimModSuite – C++ library for parallel mesh generation and adaptation
- Abstract – creation of mesh generation and problem definition templates

Analysis codes used in the development of the simulation workflows include:

- PHASTA - Parallel Hierarchic Adaptive Stabilized Transient Analysis

- ALBANY - implicit, unstructured grid, finite element code for the solution and analysis of PDEs

## Support utilized

Parallel computers from NSF XSEDE, the DOE and RPI's Center for Computational Innovation (CCI) were used in the execution of this project. Approximately 250,000 SUs were used on the TACC Stampede system for principal workflow development and execution (through an XSEDE XRAC allocation), over 1,000,000 core-hours on the Argonne Mira and Cetus Blue Gene /Q systems (through an INCITE grant) for workflow component development and testing, and over 5,000 core-hours on the Rensselaer Blue Gene /Q for portal developments.

The XSEDE support utilized on this project consisted of detailed interactions with TACC HPC scientists and support staff for porting and tuning workflow execution on the Intel Xeon Phi and the XSEDE Science Gateway team for development and testing of the PHASTA portal (GATEWAY WIKI PAGE). Critical to both interactions was support from David O'Neal at the Pittsburgh Supercomputing Center.

At the RPI Scientific Computation Research Center (SCOREC) the software that was developed is supported by a combination of DOE, joint SBIR/STTRs with Simmetrix, and Industrial grants. The key DOE grant is SCOREC's portion of the FASTMath Institute. Simmetrix, a spinoff from SCOREC 18 years ago, has DOE and DoD SBIR/STTR grants on the development of parallel simulation based engineering technologies that SCOREC is a sub-contractor on. The key industrial grant supporting software development is provided by IBM. Corning and IBM provided support for the development and application of specific simulation workflows. New York State HPC-NY and an NSF PIF project supported SCOREC a computational scientist working on the other industrial simulation workflows.

## Science findings

The execution of the industrially relevant applications for Corning, IBM, Sikorsky and Boeing outlined above indicate that the parallel unstructured mesh infrastructure discussed in references [5][7][9] can effectively address complex simulation workflows over fully general geometries in a manner that meets the cost and time constraints of industry.

The development of in-memory coupling to eliminate file I/O between HPC simulation components is now done through data structure coupling of the primary components, thus supporting the integration of new analysis codes with a minimum of modification. Coupling the components is done using functional interfaces for instantiation and execution, information passing, and data transformation. The coordination of these interfaces is provided by a high-level driver function. For a given component, specific questions are answered by another component provided that the information exchange can be based on abstractly defined descriptor handles and transformation procedures. For example, the analysis component requires mesh coordinates and connectivity information, and similarly, solution and mesh-size fields are required by the mesh adaptation component.

To date, in-memory interfaces between NASA's FUN3D, ERDC's Proteus, Sandia's Albany, PHASTA, SLAC's ACE3P, and PPPL's M3DC1 unstructured mesh adaptation have been constructed. Initial PHASTA results on an 85 million element mesh partitioned to 512 and 2048 parts indicate greater than an order of magnitude performance improvement of the in-memory information transfer and transformation over the file-based approach. Performance comes from the elimination of file I/O for the passing and transformation of information, and the elimination of the serial bottleneck when passing control from one component executable to the next.

The gateway technologies (<https://github.com/apache/airavata-php-gateway>) support development of an easy to use interface for running automated parallel simulations that abstracts away complexities such as data management, job schedulers, and run-time environment setup. An industrially relevant viscous fluid flow computational fluid dynamics workflow is supported with the PHASTA science gateway developed

in this project. Simulations are currently run on the TACC Stampede cluster, and in the coming month, through continued support from the XSEDE Industrial Challenge Program, the Rensselaer CCI Blue Gene/Q will host jobs submitted by Corning.

As part of the project started with IBM to do detailed modeling of fully featured integrated circuits, it was determined that new methods to construct geometric models that have millions of model features will be needed. The electronic industry circuit design data is closely linked to the circuit manufacturing processes and are limited to sets of 2-D layers and process information. Such information is not sufficient to support 3-D analysis that, for circuits, requires a non-manifold solid model. Thus we have been developing a set of procedures that use the 2-D layer and process information and execute solid modeling operations to create the non-manifold models using CAD systems. Although these procedures do create geometric models appropriate for automated adaptive analysis, the computational effort and model sizes for circuits with the desired number of features is becoming a serious bottleneck. Working with Simmetrix we defined, and Simmetrix implemented, methods to distribute the model in a manner consistent with the distribution of the mesh of that domain at any point in the process. Consideration of parallel model generation is a much more complex problem and we are looking for funding opportunities to support the development of an approach that we have outlined.

The performance and scalability of PHASTA on the Stampede Intel Xeon Phi based system was examined and compared against Mira BGQ system results on up to 3840 cores as shown in Table 1 [10]. One MPI rank was assigned to each of 60 cores on the Phi; the 61st core running the OS was avoided. Likewise, one MPI rank was run per core on BGQ. On Phi core pinning (`I_MPI_PIN_DOMAIN=core`) was used to ensure processes did not move between cores and introduce jitter. Processes are pinned by default on the BGQ. To avoid excessive MPI memory usage on the Phi the rendezvous and connectionless DAPL protocols are used. Unlike the eager protocol which copies data with the assumption that the receiver is ready, rendezvous uses a more memory efficient handshaking model where data is not copied until the receiver declares itself as ready. The connectionless DAPL protocol, UD, reduces memory by using a fixed number of connection pairs as opposed to the default protocol which requires each pair of processes to setup a one-to-one connection. ECSS and TACC support was critical to understanding and tuning these environment controls.

Number of cores	Elements per core	Phi: Avg. Time (s)	Phi: Scaling	BGQ: Avg. Time (s)	BGQ: Scaling	Phi Scaling/ BGQ Scaling
480	45,833	84.60	1	73.7	1.00	1.15
960	22,917	37.87	1.12	37.2	0.99	1.02
1,920	11,458	26.12	0.81	18.8	0.98	1.39
3,840	5,729	11.60	0.91	9.60	0.96	1.21

Table 1: PHASTA strong scaling results on Stampede Intel Xeon Phi (native mode) and Mira BlueGene/Q systems for a mesh with 22M elements.

PHASTA maintains strong scaling over three Phi processor doublings from 480 cores out to 3840 on Stampede; Column 3 of Table 1. Attempts to run with 240 cores for this case failed due to lack of memory. This is not a great concern given that typical PHASTA production runs are in this lightly loaded range (on the order of 5k elements per part) to exploit the strong scaling and compress the time-to-solution. While the dynamic scaling range has been shown to be higher on the BGQ (90% scaling through nine processor doublings) this is primarily due to the larger per-core memory. The proposed future path to

exascale machines (Stampede 1.5, Theta, and Aurora) for Phi architecture double the fast memory (8GB to 16GB) together with lighter weight MPI implementations which will significantly extend the strong scaling in the lower core count direction while faster interconnects will extend it in the higher core count direction. That said, the 91% scaling at the typical production run level already shows great promise for PHASTA on this architecture.

PHASTA partitions are load balanced using a combination of the ParMETIS (via Zoltan) graph partitioner and ParMA. ParMA can run with MPI only and with MPI+Pthreads using MPI\_THREAD\_MULTIPLE. Partitioning from two parts to 128 parts, one part per thread or process, the memory usage increase for an MPI only run is nearly 30x (51Mb to 1500Mb) while the MPI+Pthread run only has a 2x memory increase (65Mb vs 140Mb). On the Phi the significant memory savings come at the cost of MPI shared memory inefficiencies when the Pthreads call MPI routines [11]. Table 2 lists the ParMA partition improvement timing for runs with MPI only and MPI + Pthreads. At 4,096 parts on 69 Phi cards (at most 60 parts per card) the MPI + Pthreads run is over an order of magnitude slower than the MPI only run. Note, the eager communication protocol was used to avoid an Intel MPI 14 failure seemingly associated with the default dynamic switching between the eager and rendezvous protocol based on message length.

Parts	Phi Cards	Vertex Imbalance Reduction (%)	MPI Balance (s)	MPI + Pthreads Balance (s)
1024	18	3	13.91	23.59
2048	35	6	12.66	48.46
4096	69	15	12.52	133.58

Table 2. ParMA load balancing performance comparison of MPI only versus MPI + Pthreads on Stampede Intel Xeon Phi in native mode.

## Business findings

In each of the industrial workflows considered, new simulations capabilities were developed that addressed a clear industrial need. The level of HPC capabilities needed varied from application to application.

In a couple of cases, the simulation needs can be met on small to medium sized clusters. However, even in those cases the extension to engineering design processes that include the application of uncertainty quantification (UQ) technologies will dramatically increase the computational demands. Although UQ technologies are just beginning to mature, and there is a reasonable amount of work left before they can be easily applied by industry, the increased reliability they can introduce is very attractive to industry.

The IBM, Sikorsky and Boeing simulations are highly demanding and require the application of massively parallel computing technologies from the beginning.

RPI's work on simulating the behavior and operation of Pliant Energy Systems mechanisms was an element of the company's proposal to Office of Naval Research for a novel marine propulsion system. Outlining RPI's work with Pliant Energy to date was important for (1) validating the technology's theory of operation and (2) demonstrating that numerical modelling can be a design tool for the further development and optimization of the technology. Partly as a result of this work with RPI, Pliant Energy Systems has received over \$1.46 million in federal funding and retained three jobs.

Simmetrix technologies are being used in almost every RPI industrial project to support various simulation workflow steps. RPI SCOREC continues to work with Simmetrix on the development and execution of SBIR/STTR projects. Simmetrix and RPI have recently begun a \$990,000 DOE Phase II STTR project entitled “Automated Simulation of Selective Laser Melting Additive Manufacturing for Process Design”. In addition, Simmetrix and five RPI faculty have recently begun one DOE Phase I SBIR and two DOD Phase I STTRs.

Corning has collaborated with faculty and postdocs to incorporate physical properties including partial-slip boundary conditions into an open source CFD code PHASTA which scales well on systems like the CCI Blue Gene Q. This enables modeling the complex geometry and fluid mechanics of the twin-screw extrusion of ceramic pastes to make thin walled catalytic converter substrates – an important Corning business and one central to the successful reduction of automobile air pollution in the U.S. Corning has also made significant use of the computing facilities at CCI to leverage our internal computing, using molecular dynamics codes to model materials.

## Assessment of impact

The ability to address multiple industrial simulation workflows represents the primary measure of impact of the methods developed in this project. In this project simulation workflows were developed for seven companies. Examples of the impact to the companies include:

- The fact that active flow control devices will be used in a next generation aircraft.
- A company had been able to dramatically reduce their design customization cycle from weeks to days greatly improving their ability to respond to the requests from potential customers.
- An increased understanding of the processes associated with a specific flow control device and subsequent design improvement.

Experience with the PHASTA Science Gateway indicates that the application of these technologies greatly reduces the complexity of execution of parallel simulation on remote systems.

The ability to integrate the unstructured mesh infrastructure tools with multiple parallel simulation codes substantially increases the ability to address industries’ simulation needs.

## References

- [1] K.C. Chitale, O. Sahni, M.S. Shephard, S. Tendulkar and K.E. Jansen, “Anisotropic adaptation for transonic flows with turbulent boundary layers”, *AIAA Journal*, 53(2):277-295, 2015.
- [2] D.A. Ibanez, E.S. Seol, C.W. Smith and M.S. Shephard, “PUMI: Parallel Unstructured Mesh Infrastructure”, to appear, *ACM Transactions on Mathematical Software*, 2015.
- [3] O. Klaas, M.W. Beall and M.S. Shephard, “Generation of Geometric Models and Meshes from Segmented Data”, *Proc. NAFEMS Americas Conference 2014*, 2014.
- [4] M. Rasquin, C. Smith, K. Chitale, E.S. Seol, B. Matthews, J. Martin, O. Sahni, R. Loy, M.S. Shephard and K.E. Jansen, “Scalable fully implicit finite element flow solver with application to high-fidelity flow control simulations on a realistic wing design”, *Computing in Science and Engineering*, 16(6):13-21, 2014.
- [5] M.S. Shephard, C.W. Smith, D.A. Ibanez, B. Granzow, M.W. Beall, and S. Tendulkar, “Developing scalable components for massively parallel adaptive simulations”, *Proc. NAFEMS World Congress*, San Diego, CA, 10 pages, 2015.
- [6] M.S. Shephard and C.W. Smith, “HPC Simulation Workflows for Engineering Innovation”, *XSEDE '14, Jul 13-18 2014*, Atlanta, GA, USA ACM 978-1-4503-2893-7/14/07. <http://dx.doi.org/10.1145/2616498.2616556>

- [7] M.S. Shephard, C.W. Smith and M.W. Beall, “Technologies for the Construction of Parallel Simulation Workflows”, Proc. NAFEMS Americas Conference 2014, 2014.
- [8] C.W. Smith, M. Rasquin, D. Ibanez, K.E. Jansen and M.S. Shephard, “Application Specific Mesh Partition Improvement”, *Software and High-Performance Computing*, 2015.
- [9] C.W. Smith, S. Tran, O. Sahni, F. Behafarid, M.S. Shephard and R. Singh, “Enabling HPC Simulation Workflows for Complex Industrial Flow Problems”, XSEDE '15 Jul 2015, St. Louis, MO, USA ACM
- [10] Cameron W. Smith, Ben Matthews, Michel Rasquin, and Kenneth E. Jansen, “Performance and Scalability of Unstructured Mesh CFD Workflow on Emerging Architectures”, SCOREC Reports. [http://www.scorec.rpi.edu/reports/view\\_report.php?id=663](http://www.scorec.rpi.edu/reports/view_report.php?id=663). 2015.
- [11] Cameron W. Smith. commkernel: communication kernel for performance testing on Stampede, <https://github.com/cwsmith/commKernel>. 2015.
- [12] S. Tendulkar, M.W. Beall, R. Nastasia, O. Sahni, S. Tran and M.S. Shephard. “Advanced meshing and mesh adaptation for complex flow problems”, *Proc. NAFEMS World Congress*, San Diego, CA, 16 pages, 2015.
- [13] S.A. Tran, A.E. Fisher, D. Corson and O. Sahni, “Dynamic Stall Alleviation for an SC1095 Airfoil using Synthetic Jet Actuation”, *Proceedings of the AIAA SciTech 2015: Science and Technology Forum and Exposition*, Kissimmee, FL, USA, Jan. 2015.