TECHNOLOGIES FOR THE CONSTRUCTION OF PARALLEL SIMULATION WORKFLOWS

Mark S. Shephard

Professor, Rensselaer Polytechnic Institute, USA

Cameron W. Smith

Computational Scientist, Rensselaer Polytechnic Institute, USA

Mark W. Beall

President, Simmetrix Inc., USA

SUMMARY

Increasingly industry is relying on massively parallel computers to perform high-fidelity simulations. Although substantial progress has been made on the development of highly scalable methods to model behaviors of interest, the currently available tools fall short of providing engineers with the capabilities needed to execute the massively parallel simulation workflows they want to perform on a daily basis. A key aspect of properly addressing the needed capabilities, particularly on the new generations of heterogeneous parallel computers, is moving from a focus on flop rates to consideration of time of execution and energy consumption for a given level of simulation accuracy.

Although the majority of the flops in mesh-based analysis is consumed by the forming of matrix contributions and solving the global system, the process of setting-up and controlling the analysis model is substantial. For example the common estimate in the case of mesh-based methods on workstations is that 85% of the cost is involved with the generation and control of the spatial decomposition of the domain that constitutes the input mesh. In the case of massively parallel simulations with meshes of millions, to billions, of elements this percentage can be higher due to the bottlenecks associated with the data transfer. This clearly indicates that the meshing functions must be fully automated and execute on the same massively parallel computers as the mesh-based analysis procedure. In addition, as the complexity of the simulations increase it is also becoming clear that the approximations associated in going from the given problem to the one numerically solved should be automatically controlled to ensure simulation reliability. In most all cases, *a posteriori* methods that iteratively execute analysis steps followed by model adaptation are required to provide the simulation reliability desired. In addition to providing increased simulation result reliability, adaptive methods can yield the desired levels of accuracy with much less total computation. For example, adaptively defined finite element meshes will regularly have two orders of

magnitude fewer unknowns than non-adaptive meshes for the same level of solution resolution. However, adaptive methods require more complex, and less structured, data sets and operations and the process of adapting the numerical models is constantly changing the distribution of the data and computations leading to load imbalance that will completely destroy parallel scalability without the regular application of dynamic load balancing.

This paper discusses a set on ongoing programs developing parallel adaptive industrial workflows using software components interacting through well-defined open interfaces.

1: Industrial Programs

The development of interoperable components for massively parallel adaptive simulations is supported by the NSF PICPES and DOE SCIDAC FASTMath [FASTMATH_WEB] institute projects, while the development of industrial parallel simulation workflows is supported by the NYSTAR funded NY State High Performance Computing Consortium, HPC2 (Shephard et al., 2013), the XSEDE Industrial Challenge (XSEDE, 2014) and industry. Simulation components interact through well-defined interfaces (Seol et al., 2014 Simmetrix Inc., 2014) allowing the effective construction of component-based simulation workflows. This approach effectively enables the integration of current best-in-class components for CAD modeling, automated parallel mesh generation, PDE analysis, adaptive mesh and model control, dynamic load balancing, and results visualization.

HPC2 is working with companies to apply parallel simulations to problems of importance to them. Parallel workflows, workflow components, and advanced models, have been developed for companies that include Corning, GE, Blasch Ceramics, Boeing, ITT Fluid Technologies: Goulds Pumps, Pliant Energy Systems, IBM and Xerox (Shephard et al., 2013). Compute resources used in the workflows include the Rensselaer CCI that operates a one petaflop 81,920 core IBM BlueGene/Q system, AMOS, and an Intel Xeon cluster. The goal of the XSEDE Industrial Challenge project is to apply component technologies to make massively parallel PDE-based and multiscale simulations an integral part of the design of innovative products and processes for key Sikorsky and Corning applications (XSEDE, 2014). XSEDE FTEs support this effort with expertise in scientific workflows and component tuning.

2: Application Workflows

Three approaches have historically been taken in the development of parallel adaptive workflows. The first two approaches tightly integrate the PDE solver and adaptivity components through in-memory operations on either a single set of data structures or through embedded low-level component interface calls in coupled components. These methods have the distinct advantage of being highly efficient but require substantial development time and lack flexibility. A third approach keeps the interactions between existing software components at a higher-level so minimal modifications are required in any component. File-based implementations of this approach (Shephard et al., 2011; 2013) are the fastest to develop, but are the least efficient due to the file I/O bottleneck. In-memory parallel adaptive simulation workflows linking specific combinations of Simmetrix and/or RPI SCOREC CAD interrogation, parallel mesh generation, and adaptation tools have been developed for CFD using PHASTA or FUN3D, electromagnetics using ASC3P, MHD using M3D-C1, and non-linear solids using the Albany multiphysics code.

In order to perform interactions at a higher-level, a clearly defined set of procedures, and associated interactions, are defined to support adaptive analysis following geometry-based simulation concepts that emphasize storing and retrieving information at the right level in the hierarchy of abstractions (Delalondre et al., 2010). The development cost to interface components varies greatly with the implementation of the component. Monolithic legacy codes operating through file structures, such as PDE analysis written in procedural FORTRAN 77/90, are most easily interfaced by replacing the file procedures, but often require careful attention to memory management. Conversely, modular codes with existing functional interfaces require a thinner interface layer and often already have the necessary memory management mechanisms.

A massively parallel in-memory integrated circuit physics application workflow is being developed for a semi-conductor manufacturer by coupling Simmetrix parallel mesh generation and CAD interrogation tools, SCOREC mesh adaptation and load balancing tools, and Albany multiphysics code (Chen et al., 2014). Critical to the fidelity of the adaptive workflow is Simmetrix GeomSim's support for interrogating a parametric representation of the complex geometric model on leadership class machines where the commercial CAD kernels are not supported. Figure 1 depicts a geometric model automatically generated from design layout data and a 1.1 billion element mesh generated in parallel with Simmetrix SimModSuite.

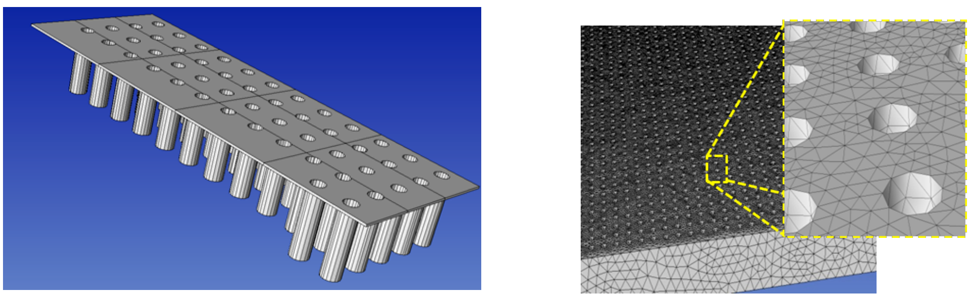

Figure 1. Geometric model of liner film (left) and a 1.1 billion element mesh generated in parallel with Simmetrix SimModSuite on AMOS BG/Q (right).

3: Closing Remark

The tools and methods needed to develop large scale parallel simulation workflows have evolved to the point where it is practical, in terms of time and cost, to construct such workflows to directly meet the needs of industry. With programs like PICPES, HPC2 and XSEDE and expertise such as that at RPI SCOREC and Simmetrix we welcome industry to work with us to address their large scale parallel simulation workflow needs.

REFERENCES

Delalondre, F., Smith, C.W., Shephard, M.S.. 2010. Collaborative software infrastructure for adaptive multiple model simulation. *Computer Methods in Applied Mechanics and Engineering*. **199** (21-22). pp. 1352-1370.

FASTMath. 2014. *FASTMath: Applied mathematics algorithms, tools, and software for HPC applications*. [ONLINE]. [Accessed 24 April 2014]. Available from: http://www.fastmath-scidac.org/.

HPC2. 2014. *NY State High Performance Computing Consortium*. [ONLINE]. [Accessed 24 April 2014]. Available from: http://hpc2.org/.

Chen, Q., Ostien, J., Hansen, G.. 2014. Development of a used fuel cladding damage model incorporating circumferential and radial hydride responses. *Journal Of Nuclear Materials*. **447**(1-3). pp. 292-303.

Seol, S., Ibanez, D.A., Smith, C.W., Shephard, M.S.. 2012. PUMI: Parallel Unstructured Mesh Infrastructure. [ONLINE]. [Accessed 24 April 2014]. Available from: http://www.scorec.rpi.edu/reports/view_report.php?id=625.

Shephard, M.S., Smith, C.W., Dord, J. and Jansen, K.E.. 2011. *Bringing Parallel Adaptive Simulation to Industry*, In: *Proceedings of the NAFEMS World Congress*, 23-26 May 2011, Boston, MA.

Shephard, M.S., Smith, C.W., Kolb, J.E.. 2013. Bringing HPC to Engineering Innovation. *Computing in Science and Engineering: Digital Manufacturing*. **15** (1), pp.16-25.

Simmetrix Inc.. 2014. *Simmetrix Inc. - Mesh Generation, Geometry Access*. [ONLINE]. [Accessed 24 April 2014]. Available from: http://www.simmetrix.com/.

XSEDE. 2014. *XSEDE | Industry Challenge Program*. [Online]. [Accessed 24 April 2014]. Available from: https://www.xsede.org/industry-challenge-program.