

# The CESAR Codesign Center: Early Results

The CESAR Team

March 5, 2012

## Abstract

We describe early results from the Center for Exascale Simulation of Advanced Reactors (CESAR). These flow naturally from our project plan to evolve existing petascale codes to exascale, couple them in order to provide accurate multiphysics simulations of advanced nuclear reactor designs, and carry out code development in a codesign environment that incorporates computer vendors and system software developers in the application development process.

## 1 Introduction

The Center for Exascale Simulation for Advanced Reactors (CESAR) aims to develop a coupled, next-generation nuclear reactor core simulation tool capable of efficient execution on exascale computing platforms. The new code, TRIDENT, will enable tight coupling of extreme-fidelity models in the three physics domains relevant to reactor analysis – neutron transport, structural mechanics, and thermo-hydraulics. These new modeling capabilities have broad relevance across a wide range of disciplines. At the same time they enable new critical reactor design analyses for both existing and next-generation reactors. Three classes of such new problems will serve as the application drivers for CESAR: (1) modeling of full-vessel, coupled neutronics/thermal-hydraulics for systems in natural convection conditions, (2) accurate fuel depletion modeling using coupled highly detailed neutronics and thermal-hydraulics modeling needed in breed/burn concepts, and (3) detailed structural mechanics coupled to both neutronics and thermal-hydraulics to assess core reactivity feedback and fuel assembly structural integrity. Existing industry tools cannot resolve these detailed physics couplings, and the reduced-order models currently in use have significant economic impact on reactor design, licensing, and operation.

Our original plan was to start with three currently well-performing codes (in thermo-hydraulics (Nek5000), neutronics (UNIC, OpenMC), and structural mechanics), use them to understand scaling and performance issues on both near future and future exascale platforms, and coupled them at scale to enable attacks on nuclear engineering problems impossible with conventional methods. Budget downsizing has caused us to put the structural mechanics component on hold for the time being, but we are still architecting the coupled code so as to make possible its later addition. Our initial progress has been in support of our original goals for the thermo-hydraulics and neutronics components. We describe this progress in the following areas.

1. Carry out investigation of our initial codes in depth, focusing on scaling issues and performance on future platforms. This is done with a combination of performance measurement tools, performance modeling tools, and custom kernels and mini-apps. See Section 2.1.
2. Prepare these codes for coupling by developing memory access interfaces. See Section 2.2.
3. Study the application of uncertainty quantification techniques appropriate for nuclear engineering applications. See Section 2.3.

4. Study new hardware developments in the context of our applications, including the many-core many-thread architecture of IBM’s BG/Q, NVRAM, and GPGPU’s. See Section 2.4.
5. Design custom visualization for the project, coupled to the coupling mechanism for the coming merge. See Section 2.5.
6. Plan the demonstration calculation for the first year. See Section 2.6.
7. Finally, we have put in place the infrastructure in which to carry out our code development and our codesign activities in particular. See Section 2.7.

## 2 Initial Progress

In this section we describe the progress we have made in multiple areas across the project.

### 2.1 Analyzing Existing Codes

The CESAR project begins with mature, peta-scale enabled application codes as the starting point for exascale co-design activities. Our foundational codes are *Nek5000* for CFD with heat transfer (thermo-hydraulics), *UNIC* for deterministic neutron transport, and *OpenMC* for Monte Carlo neutron transport. A critical early milestone of CESAR is the abstraction of mini-applications and communication kernels that allow the focused, in-depth study of key performance issues for different aspects of the applications. Below we discuss our kernel application development efforts and early results.

#### 2.1.1 Early Experiments with the Nek-bone Skeleton Application

**Setup:** The accompanying figures show performance results for initial trials with the Nek skeleton code as a function of the number of elements per core ( $E$ ) and number of cores ( $P$ ) on a 2.66 GHz Intel Xeon (the ANL/MCS *churn* compute server).

The test runs through 100 iterations of unpreconditioned conjugate gradient iteration with  $E_t = E \times P$  spectral elements of order  $N_p$ , for a total gridpoint count of  $n \approx E_t N_p^3$ . Each brick-shaped element comprises  $N^3 = (N_p + 1)^3$  points, ordered lexicographically (i.e.,  $i-j-k$  ordering), and the elements themselves are arrayed in a one-dimensional  $1 \times 1 \times E_t$  array, so that any element shares at most two faces and at most two faces per core are exchanged in the direct-stiffness summation process (the finite-element equivalent of a halo-exchange). Tests were made for polynomial order  $N_p=7$  (left) and  $N_p=9$  (right). For each case, two trials ( $a$  and  $b$ ) were run to test repeatability. Tests were performed without interference from other active processes.

**Results:** The 8-core Xeon shows clear performance degradation as the number of active cores is increased, particularly as one steps from 6 to 7 cores, at which point the per-core flop rate drops from  $> 2$  Gflops to 0.5–1.0 Gflops. This drop is probably a result of having dual quad-core chips and crossing to the point where one of the chip’s core count is saturated. The performance for 6 cores or less is quite repeatable, but is less so for  $P=7$  or 8. There is also a clear drop in performance for larger memory applications as can be seen when  $E > 60$  for  $N = 8$  or  $E > 30$  for  $N = 10$ .

Finally, we also show results for  $P = 2048$  with BG/P in vn-mode for  $N=8$ . Two interesting features are observed. First, the performance is quite low ( $\sim 300$  Mflops), despite having assembly coded matrix-matrix product routines available that, in principle, exploit the double-hammer floating point unit. (For  $N$  even, all data is quad-aligned, which allows use of the the double-hammer.) With the double-hammer, the 850 MHz

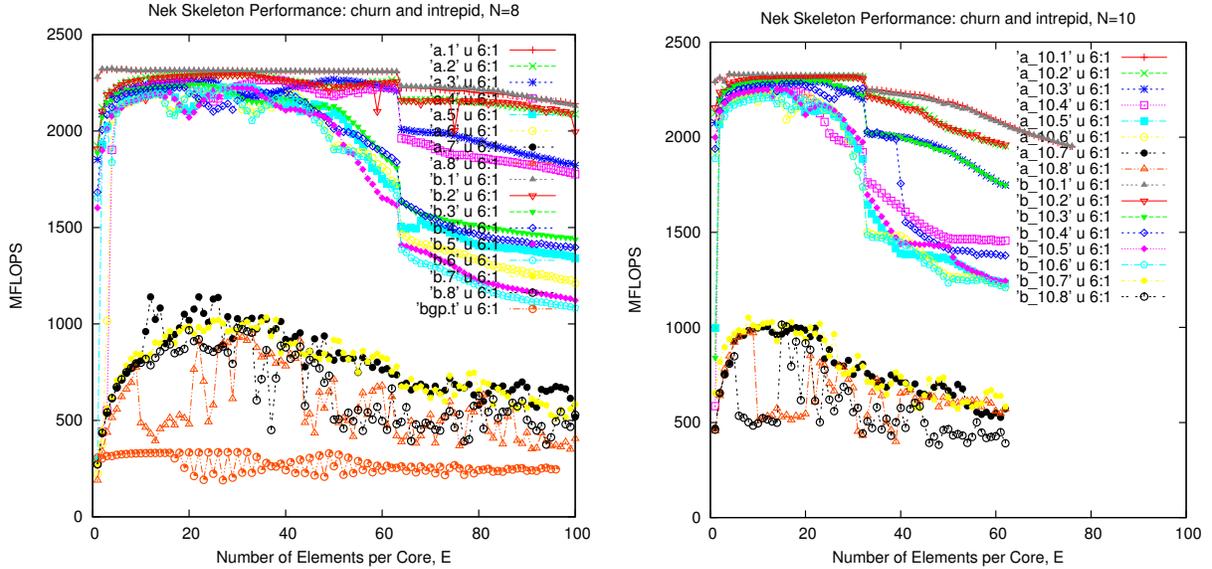


Figure 1: Nek skeleton performance on an 8-core 2.6 GHz Intel Xeon for polynomial order  $N_p=7$  (left) and  $N=9$  (right). (Data array sizes  $N^3 = 8^3$  and  $10^3$ , respectively.)

BG/P core should be capable of a peak 3.4 GFLOPS, in contrast to the 5.3 GHz peak of the Xeon. Secondly, there is an interesting beat phenomenon that shows radical (almost 2x) performance variance as a function of  $E$ . (This latter feature is evident in other tests but not precisely repeatable.) We note that there is no such discrepancy between the BG/P and *Churn* platforms for the full Nek5000 code, which points to a need for full code analysis as well as principal kernels.

### 2.1.2 Early Results with MOCFE Mini-App

For deterministic transport, the method of characteristics approach MOCFE [12] was selected as an initial focus of more detailed study on exascale machines. In MOCFE the neutral particle transport equation (function of space, angle, and energy) is written as a set of  $G \times A \times T$  coupled differential equations, where  $G$ ,  $A$ , and  $T$ , are the number of energy groups, angles, and trajectories used to discretize the domain, respectively. The number of trajectories is strictly dependent upon the angle and the larger the number, the better the accuracy of the solution. Figure 3 shows an example refinement in trajectory distribution (black lines) for a particular angle on a VHTR reactor geometry [12].

Figure 2 shows the scaling results on the development version of MOCFE on BlueGene/P for a problem with maximum  $10^{12}$  degrees of freedom (DOF). In Figure 2, the “serial” point is actually a decomposition involving 128, 1, 16, and 512 pieces for strong spatial, strong angle, weak angle, and weak spatial scaling.

The mini-app development greatly simplifies MOCFE to its basic components, yet carries out commu-

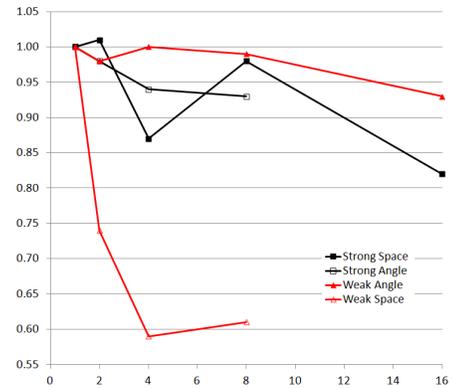


Figure 2: Scaling Results for the MOCFE code

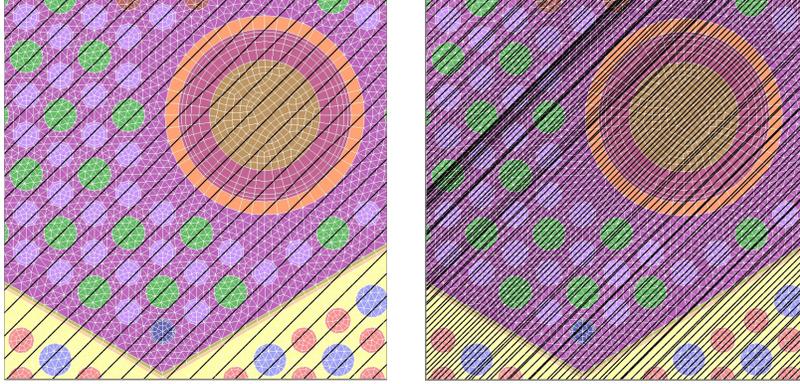


Figure 3: Variation in the Trajectory Density for a Reactor Geometry.

nication in a similar pattern. More complicated communication patterns and the complex load imbalance issues observed in MOCFE can easily be contrived by changing the source. Preliminary results on a single node system indicate behavior similar to that of the MOCFE solver.

As the next steps in this work, we will verify the behavior of the proxy application at large scale on Argonne’s *Intrepid* system using a series of standard performance tools and application specific performance tracking techniques as well as conduct a series of parameter studies highlighting potential bottlenecks in the code as we increase scale.

### 2.1.3 Early results with MCKK

While not historically competitive performance-wise with deterministic methods in reactor physics, Monte Carlo (particle-based) neutronics algorithms potentially become significantly more attractive as we move to candidate exascale-type architectures. Thus, we designed and carried out a series of performance analyses with the mini-application MCKK (Monte Carlo Communication Kernel) and the the OpenMC code. The aim of these analyses was to understand the key performance issues underlying Monte Carlo neutron transport computations on exascale architectures. Our initial studies focused on the feasibility of carrying out domain decomposition (i.e., to distribute the tally memory across processors) to enable realistic computations on architectures with much lower memory per core.

Three key issues have been studied so far: the first involved measuring the cost of communication for domain-decomposed codes. Details are discussed in [9]. The second involved studying the performance *penalties* brought about by load imbalances [10] when using many cores. Finally, we an-

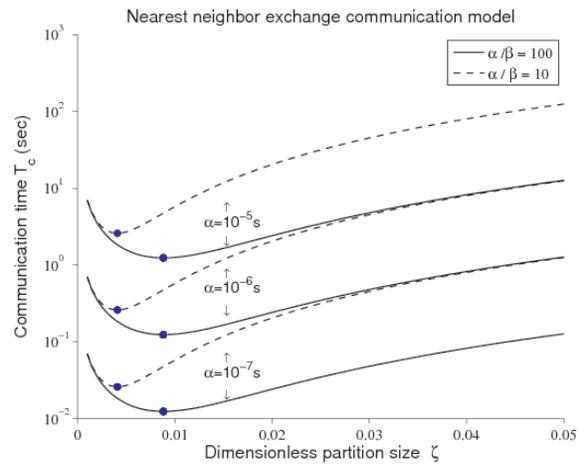


Figure 4: Evaluation of the performance model (Total communication time for values of  $\alpha = [10^{-5}, 10^{-4}, 10^{-3}]$  from bottom to top and  $\beta$ ). Dashed and solid lines indicate identical values of  $\alpha/\beta$ , and the dots indicate the local minimum value predicted by the model.

alyzed the possibility of using in-place many-to-many communication algorithms to save memory for high leakage rate problems [2].

<i>experiment</i>	$\frac{\beta \ \lambda^{max}\  + \mu}{\beta \ \lambda\  + \mu}$	$\frac{\sum_{i=0}^M \ \lambda^{max}\ ^i}{\sum_{i=0}^M \ \lambda\ ^i}$	$C$	$\frac{p_0^{max}}{P_0}$	$\Delta$
full assembly	1.00	1.13	1.13	3.24	3.67
quarter assembly	1.00	2.58	2.58	3.28	8.47
ninth assembly	1.00	6.98	6.98	3.45	24.15

Table 1: Values of  $\Delta$  and the various terms which contribute to it for each of the three numerical experiments. The tables used values of inverse latency  $\beta = 10^{-8} \frac{sec}{particle}$  and tracking time  $\mu = 5 \times 10^{-4} \frac{sec}{particle}$ . Note that within the precision presented the bandwidth term (second column) is identical in all cases, a manifestation of the fact that bandwidths are much higher than tracking rates. Notice also that the load imbalance penalty is magnified significantly on the finest partition’s grid.

## 2.2 Scalable Coupling of Highly Parallel Codes

The goal of this task is to perform coupled neutronics, heat transfer and fluid flow, and eventually structural mechanics modeling of a reactor core. The underlying coupling technology must be able to link to existing physics codes (UNIC, Nek5000, and Diablo), and treat the various mesh types used in those codes. We will use the solution transfer functionality in MOAB for this purpose [14]. Recently, MOAB was used to couple the PBSM basin modeling code to the FWI3D seismic code [16]. The coupled code was demonstrated to run on 8192 nodes of an IBM BG/P (in SMP mode), with solution transfer accounting for less than 3% of the overall runtime.

Work has started on fully integrating MOAB into Nek5000 and UNIC. Nek5000’s interface to MOAB was enhanced to enable both fluid and conjugate heat transfer calculations. At the same time, input processing was improved to allow specification of boundary conditions and material/physics regions (i.e., fluid or conjugate heat transfer) without requiring recompilation of the code. The UNIC interface to MOAB was updated in preparation for enabling parallel input and output from/to MOAB. Efforts were also undertaken to link the Diablo code to MOAB.

OpenMC is being used as the test bed to evaluate the feasibility and performance in coupling stochastic neutron transport and fluid flow/heat transfer fields by using recently-developed nonlinear low-order operator techniques as the iterative bridge between Monte Carlo and thermal fluids. Initial results for very large-scale steady-state light water reactor (LWR) applications are extremely promising, and effort to extend this coupling to time-domain neutronics is underway. These time domain extensions represent the first credible attempt to push Monte Carlo neutron transport methods into time-domain coupled-physics nuclear reactor applications a truly huge scale simulation problem.

A paper describing this work [11] is being presented at SIAM’s Twelfth Copper Mountain Conference on Iterative Methods (March 25-30, 2012). This paper also presents a detailed specification for the first-of-a-kind massive-scale LWR simulation benchmark of coupled neutronics/fluids analysis methods. This benchmark includes extensive measured reactor data covering two cycles of plant operations that will be extremely useful in V&V/UQ of the reactor simulation tools developed in the CESAR project.

## 2.3 Uncertainty Quantification in Reactor Modeling

The tasks of the uncertainty quantification in the codesign context are two-fold. First, we design and develop the tools that allow uncertainty quantification for nuclear reactor modeling on exascale platforms. Second, we aim to identify simulation of interest *where numerical error is the limiting factor, as opposed to uncertainty in physical data*. Success in this task will ensure that *exascale computing for CESAR is not only possible, but is also necessary* for the so-identified problems. However, the latter is a demanding task, as uncertainties in nuclear data often come in at 10-20%, which tends to dwarf high fidelity concerns.

We made recent progress in both aims, as described in [3]. We developed a novel adjoint approach for differential-algebraic equations, which allows us to estimate both numerical error and uncertainty effects for time-dependent approaches. The critical novelty is that our approach allows for error in the constraints, which then allows for gauging the effects of iterative solvers in the constraints. We demonstrated this in the case of the depletion simulation for the traveling-wave reactor, one of the target application reactors for CESAR, in which case the constraints are the neutronics calculations. A very important finding, which we expect to guide us to achieving our second aim, is that the adjoint variables are  $10^6$  times larger in the beginning of the simulation than at the end of the simulation, as shown in Figure 5. This indicates that the numerical error has to be  $10^6$  times smaller in the beginning of the simulation than at the end for comparable effects. This leads us to hypothesize that beginning cycle calculations may be more limiting than data uncertainty effects.

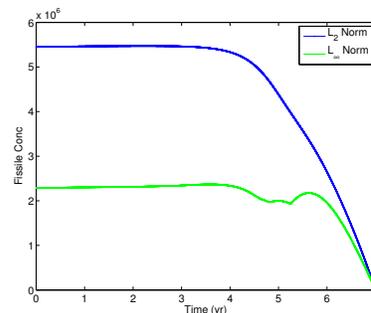


Figure 5: The adjoint variable of fertile material and the end of reactor life

## 2.4 Adapting to New Hardware

In this section we describe some early efforts to evaluate next-generation hardware in the context of our application codes.

### 2.4.1 Next-Generation Blue Gene

A co-design methodology has been developed which incorporates both simulation technology at the small scale and analytic modeling capability at the large scale to explore both the small and large scale system architecture space, as well as potential application designs. This framework is shown in Figure 6.

Parameters derived from the system architecture (using measurement and simulation) are combined with an analytic representation of application behavior to derive a performance model. This model takes as input empirical or simulated single-core or single-socket performance measured using IBM's Mambo simulator and existing Blue Gene/Q hardware and predicts performance at large scale. These predictions are used in turn to guide future architecture and application design decisions.

The Nek5000 application and Nek-bone mini-application

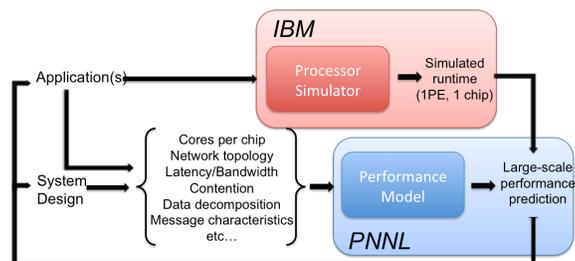


Figure 6: Co-design workflow incorporating both small-scale simulation and large-scale modeling

form the initial focus for this work. At the small scale, these codes are executing on Blue Gene/Q hardware where performance measurements are obtained. Simultaneously, models are under development for these codes and scaling behavior is under study. Completed and validated models will form the basis for future architectural and application studies.

## 2.4.2 NVRAM

In order to meet the power constraints of an exascale system, the CESAR team is exploring the potential use of non-volatile random access memory (NVRAM). In contrast to traditional memory, NVRAM does not require continuous power to retain data. As a consequence, NVRAM can be much more energy efficient for certain memory access patterns. The most common example of NVRAM is the increasingly popular NAND flash memory used in consumer electronics. Our analysis is focused on three types of NVRAM which have promising properties for HPC: spin-transfer torque (STTRAM) memory, phase-change (PCRAM) memory and memristors (MRAM).

Our analysis investigates the CESAR applications using NV-SCAVENGER, a binary instrumentation tool that characterizes an applications memory access pattern and identifies which data structures are a good fit for NVRAM in an exascale hybrid DRAM-NVRAM memory architecture. Our preliminary results from Nek5000 on the 2D eddy problem identify numerous opportunities for the use of byte-addressable NVRAM. In total, 31% of the overall memory footprint is accessed in a pattern that is suitable for NVRAM. Results from simulation suggest a potential 28% reduction in overall power consumption. This analysis also resulted in a set of general rules for what makes an NVRAM-friendly scientific application.

Our preliminary study is focused on understanding the characteristics of Nek and UNIC to determine if they could easily make use of a hybrid DRAM-NVRAM memory architecture by partitioning their data structures into the two different memory spaces. Initial results are positive. However, given NVRAMs device level costs and asymmetric latencies, further analysis is required before making any conclusions about performance and energy efficiency for general exascale applications. Details will be published in [4].

## 2.4.3 GPGPU's

**Setup:** We selected the hexahedral spectral element Laplace operator as an involved and representative kernel that constitutes a significant portion of the operational time spent in the solution of the pressure equation step in the Nek thermofluids solver when using a temporal splitting scheme. In tensor array notation the operation consists of

$$\begin{aligned}
 (\Delta u)_{eijk} &= D_{mi} \left( (G^0 u_r)_{emjk} + (G^1 u_s)_{emjk} + (G^2 u_t)_{emjk} \right) \\
 &+ D_{mj} \left( (G^1 u_r)_{eimk} + (G^3 u_s)_{eimk} + (G^4 u_t)_{eimk} \right) \\
 &+ D_{mk} \left( (G^2 u_r)_{eijm} + (G^4 u_s)_{eijm} + (G^5 u_t)_{eijm} \right) .
 \end{aligned}$$

with  $0 \leq i, j, k \leq N_p$  and  $0 \leq e < K$  where  $K$  is the total number of hexahedral elements. Here  $\mathbf{D}$  is an  $N \times N$  dense matrix, and the line derivatives of the solution are also computed in the kernel with

$$(u_r)_{eijk} = D_{im} u_{emjk}, (u_s)_{eijk} = D_{jm} u_{eimk}, (u_t)_{eijk} = D_{jm} u_{eijm}.$$

The geometric factors for the transformation between the reference element and the physical elements are encoded in the  $\mathbf{G}$  arrays.

We prototyped two families of kernels to gain an understanding of the GPU performance bottlenecks in handling the above sum reduction and scaling operations. The first family uses a natural three-dimensional

OpenCL work group of threads with an extensive local memory footprint, but the maximum order of element is limited by the cap of 256 threads in an OpenCL work group size to 5th order or lower. The second family of kernels uses a two-dimensional work group with the solution  $u$  rolled through registers in an approach adapted from Micikevicius’s three-dimensional finite difference kernels. The second kernel family uses more registers but can be used beyond 5th order elements up to an order determined by the available register file and local memory.

**Results:** In Figure 7 we show the range of performance obtained by the two families of spectral element Laplacian kernels on an NVIDIA Fermi from a GTX 590 GPU and an AMD Tahiti on an AMD Radeon HD 7970. A reference line for each benchmark shows the theoretical maximum performance if we fully saturate the memory bandwidth. The best performing kernels on the AMD GPU scale reasonably close to the bandwidth saturated limit. The reasons for the lower percentage of peak attained on the NVIDIA GPU are still under investigation.

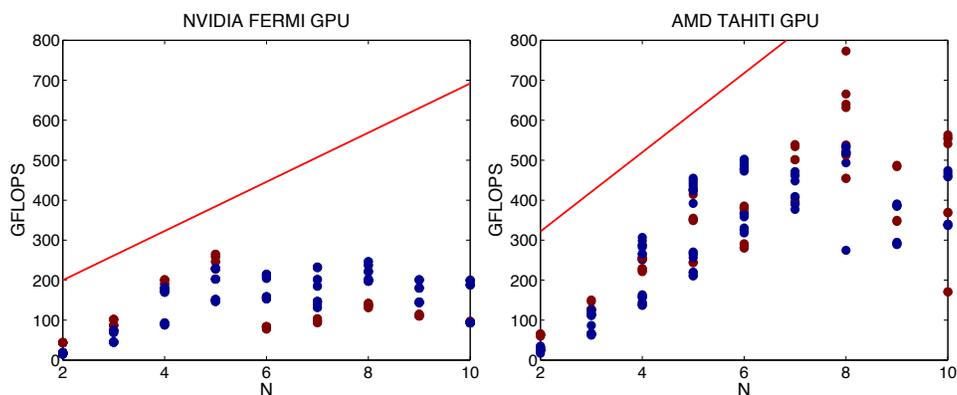


Figure 7: Performance of a sequence of hand-coded, parameterized, OpenCL kernels of slab (red dots) and block (blue dots) type for the hexahedral element spectral element Laplace operation. The red line shows theoretical peak assuming the kernel is memory bandwidth limited. Left: performance on one Fermi GPU chip on the NVIDIA GTX 590. Right: performance on the Tahiti chip of an AMD HD 7970.

## 2.5 Visualization in CESAR

Figure 9 outlines the structure of our activities on three levels: data models, tasks, and tools. Two data model interfaces are being supported: access to the codes’ native data structures and access through the MOAB mesh interface. Data management, analysis, and visualization tasks may be executed in situ, in coprocessing mode, in postprocessing mode, or in combination, and include custom data analyses, production and custom visualization, and I/O. Our research targets all of these tasks, and builds on a number of open-source tools, including VTK [8], ParaView [6], VisIt [1], DIY [7], GLEAN [15], and MOAB [13].

Our accomplishments to date include:

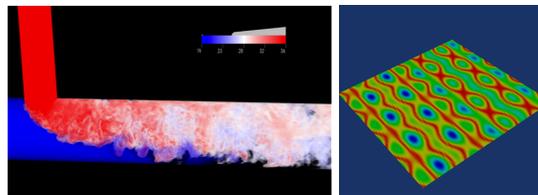


Figure 8: Nek5000 custom visualization. Left: T-junction benchmark visualized in postprocessing using VisIt. Right: Eddy benchmark visualized in situ using VTK.

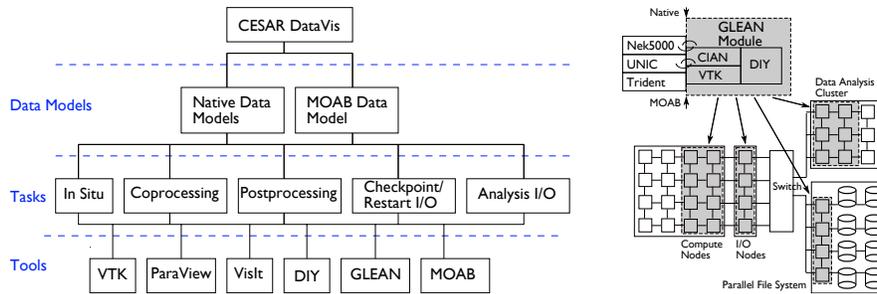


Figure 9: Left: The structure of the CESAR data management, analysis, and visualization research. Right: Custom software organized into modules that can be executed at various locations.

- Demonstrated accessing in situ data in Nek5000 via both native and MOAB data models
- Demonstrated calling back into native Nek5000 during in situ analysis and using its functions as needed for analysis
- Demonstrated computing lambda-2 vorticity in situ in Nek5000 and filtering it to a small subset of values
- Captured and replicated the current VisIt visualization workflow for Nek5000
- Repaired and upgraded Nek5000 reader in ParaView
- Ran VTK in situ in Nek5000 via native data model

Two examples of these success are shown in Figure 8. We have successfully replicated the existing VisIt workflow used by previous projects to generate the T-junction visualization shown at left, run in parallel on the *Eureka* machine at the Argonne Leadership Computing Facility. At right, for the first time, a prototype visualization using VTK was generated in situ in Nek5000 using its native data model.

## 2.6 A Reactor Engineering Calculation

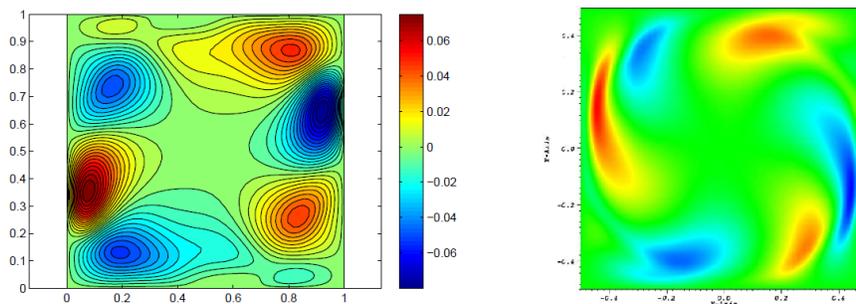


Figure 10: Most unstable eigenmode for a differentially heated cavity, left - modal analysis, right - adjoint method

The first task undertaken has been determining the focus problem and procuring the data to produce a meaningful simulation. The proposed reactor simulation is based on the MASLWR experimental facility at

Oregon State University. The experiment is designed to be in-scale with an SMR design. Drawings, steady-state data, and some specific transient data are available. At the moment we are focused on reproducing the steady state by the means of a single phase simulation with Nek5000. A CAD file from the drawings is under preparation. The next step will be simulating a start-up transient. The size of the problem is ideal for beyond-petascale simulations.

Concurrently with the preparation of the simulations, we are working on implementing in Nek5000 tools that are able to predict the stability of a given flow solution in natural convection. The methodology would be based on the solution of the adjoint problem as well as the direct problem. The method is currently implemented and tested on differentially heated cavities and thermosiphons [5]. The results obtained are largely consistent with those available in the literature for traditional modal linear stability analysis. The differences between the results of the two methodologies are currently under investigation (See Figure 10).

## **2.7 Crosscutting Activities**

A project of this size requires a certain amount of publicly accessible infrastructure in order to facilitate both internal and external communication.

### **2.7.1 Web Site**

In order to provide public visibility as well as access to assorted documents for the ASCR Program Office, we have created a web site, <https://cesar.mcs.anl.gov>, at Argonne. This site is still under construction, but is beginning to take shape. It will be a single point of reference for highlights, reports, publications, personal contacts, etc., for the CESAR project. It contains a “internal” section as well, available only to project members, to facilitate communication within the project. It currently contains presentations given at our project meetings.

### **2.7.2 Code Repository**

All members of the Project (See CESAR web page under “People”) have access to a code repository containing both our current full applications and the “mini-apps” or “skeleton apps” described in Section 2.1. These are our internal working copies. As the Codesign Center Consortium becomes established, we will package certain of these codes and make them available to the Consortium and to vendors other than our vendor partner.

### **2.7.3 Codesign Activities**

A distinguishing feature of our Center is the (funded) partnership with IBM to facilitate the sort of codesign activities that are best accomplished with an IBM employee as part of our project. Initial activities in this area have been collaboration with our program modeling activities (see Section 2.4.1) and in the area of data analysis hardware characteristics (see our position paper on the topic).

## **3 Acknowledgments**

This work was supported in part by the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract #DE-AC02-06CH11357 and award #DE-FG02-08ER25835.

## References

- [1] Hank Childs, Mark Duchaineau, and Kwan-Liu Ma. A scalable, hybrid scheme for volume rendering massive data sets. In *Proceedings of Eurographics Symposium on Parallel Graphics and Visualization 2006*, pages 153–162, Braga, Portugal, 2006.
- [2] K. Felker, A. R. Siegel, and S.F. Siegel. Optimizing memory constrained environments in Monte Carlo nuclear reactor simulations. *International Journal of High Performance Computing Applications*, 2012. (under revision).
- [3] M. L. Adams H. F. Stripling, M. Anitescu. A generalized adjoint framework for sensitivity and global error estimation in time-dependent nuclear reactor simulations. *Annals of Nuclear Engineering*, 2012. Submitted. Also, Preprint ANL/MCS-P1963-1011.
- [4] Dong Li, Jeffrey Vetter, Gabriel Marin, Collin McCurdy, Cristi Cira, Zhuo Liu, and Weikuan Yu. Identifying opportunities for byte-addressable non-volatile memory in extreme-scale scientific applications. In *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS)*, 2012.
- [5] E. Merzari, P. Fischer, and W.D. Pointer. Optimal disturbances in three-dimensional natural convection flows. In *Proceedings of the ASME 2012 Fluids Engineering Summer Meeting*. (submitted).
- [6] Kenneth Moreland, Lisa Avila, and Lee Ann Fisk. Parallel unstructured volume rendering in ParaView. In *Proceedings of IS&T SPIE Visualization and Data Analysis 2007*, San Jose, CA, 2007.
- [7] Tom Peterka, Robert Ross, Wesley Kendall, Attila Gyulassy, Valerio Pascucci, Han-Wei Shen, Teng-Yok Lee, and Abon Chaudhuri. Scalable parallel building blocks for custom data analysis. In *Proceedings of LDAV'11*, Providence, RI, 2011.
- [8] William J. Schroeder, Kenneth M. Martin, and William E. Lorensen. The design and implementation of an object-oriented toolkit for 3D graphics and visualization. In *Proceedings of the 7th conference on Visualization '96, VIS '96*, pages 93–ff., Los Alamitos, CA, USA, 1996. IEEE Computer Society Press.
- [9] A. Siegel, K. Smith, P. Fischer, and V. Mahadevan. Analysis of communication costs for domain decomposed Monte Carlo methods in nuclear reactor analysis. *Journal of Computational Physics*, 231(8):3119–3125, April 2012.
- [10] A. R. Siegel, K. Smith, P. K. Romano, B. Forget, and K. Felker. The effect of load imbalances on the performance of Monte Carlo codes in LWR analysis. *Journal of Computational Physics*, 2012. (submitted).
- [11] Kord Smith, Benoit Forget, Paul Romano, Bryan Herman, and Andrew Siegel. Challenges and approaches to coupled Monte Carlo neutronic/fluid flow analysis of commercial nuclear reactors.
- [12] M. A. Smith, A. Marin-Lafleche, W. S. Yang, D. Kaushik, and A. Siegel. Method of characteristics development targeting the high performance Blue Gene/P computer at Argonne National Laboratory. In *International Conference on Advances in Mathematics, Computational Methods, and Reactor Physics (M&C)*, Rio de Janeiro, Brazil, May 2011.

- [13] Timothy Tautges, Corey Ernst, Clint Stimpson, Ray Meyers, and Karl Merkley. MOAB: A mesh-oriented database. Technical Report SAND2004-1592, April 2004.
- [14] Timothy J. Tautges and Alvaro Caceres. Scalable parallel solution coupling for multiphysics reactor simulation. *Journal of Physics: Conference Series*, 180, 2009.
- [15] Venkatram Vishwanath, Mark Hereld, Vitali Morozov, and Michael E. Papka. Topology-aware I/O acceleration for IBM Blue Gene/P supercomputing applications. In *Proceedings of SC11*, Seattle, WA, 2011.
- [16] Mi Yan, Kirk Jordan, Dinesh Kaushik, Michael Perrone, Vipin Sachdeva, Timothy J. Tautges, and John Magerlein. Coupling a basin modeling and a seismic code using MOAB. In *International Conference on Computational Science*, Omaha, NE, June 2012. (to appear).