



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Preparing the HARVEY Circulatory Modeling Code for Exascale

E. W. Draeger

October 25, 2019

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Proposal PI: Draeger, Erik W.; draeger1

PI Organization(s): Computation

Proposal Title: Preparing the HARVEY Circulatory Modeling Code for Exascale
Applicable Mission Focus Areas and/or Core Competencies (select a maximum of three):

- High Performance Computing
- Predictive, patient-specific computational biology

Abstract: Atherosclerotic cardiovascular disease is the leading global cause of death, accounting for more than 17.3 million deaths per year, a number that is expected to grow to more than 23.6 million by 2030. Personalized computational models of the coronary arteries can be used to analyze the influence that a narrowing of the blood vessel can have on the associated hemodynamics. To achieve the necessary scale and precision in circulatory models with explicit fluid coupled to deformable cell models, our simulation code HARVEY must be designed to efficiently use current and next-generation HPC architectures. This Feasibility Study will involve a careful assessment and redesign of the code's core data structure to prepare for GPU-accelerated systems like Sierra as well as next-generation architectures like the A21 exascale system.

Description: HARVEY is a Lattice Boltzmann fluid dynamics code that has been developed in collaboration with the Randles research group at Duke University for the past several years to push the boundaries of circulatory modeling in arterial geometries. We have demonstrated unprecedented scalability with high resolution simulations [1] and have since continued to optimize the time to solution and now achieve subsecond simulation times on large, high resolution geometries at full scale. HARVEY was a finalist for the 2015 Gordon Bell Prize [2] based on runs performed on the full Sequoia Blue Gene/Q machine at LLNL. The code has since been extended to include deformable cells, coupled to the fluid with the Immersed Boundary Method. This has greatly expanded the predictive capability of the code but requires a careful redesign to avoid bottlenecks in load balance, data motion, memory usage, and I/O for both checkpointing and visualization and analysis.

This Feasibility Study will consist of three activities: (1) Replacing the existing data structure with a performance portable version, capable of working with heterogeneous memory spaces and accelerators such as GPUs; (2) Improving the load balance for simulations with and without explicit cells; and (3) Identifying synergistic efforts that can be leveraged for I/O and *in situ* visualization and analysis. The first activity will explore potential abstractions to the existing data structure that will allow for algorithm- and architecture-specific implementations, as well as a hierarchical subblocking approach to improve load balance and increase concurrency on GPUs. The second will expand a prototype serial load balance implementation developed by Duke University graduate student Jeff Ames using the Metis partitioning library to be more scalable and memory efficient. This work will be done in close collaboration with Duke. The final piece will consist of evaluating and integrating libraries developed by LLNL and other DOE laboratories as part of the Exascale Computing Project designed to create common exascale components for use in a variety of HPC applications.

FY19 LDRD Feasibility Study (FS) Whitepaper

Expected Outcomes: This study will enable HARVEY to remain at the forefront of high-resolution circulatory simulation modeling. It will also allow the PI to expand his skills in preparing HPC applications for the next generation of exascale architectures, thereby allowing him to be a more effective leader of scalable application development projects. The results of this Feasibility Study will allow for new capability calculations to be carried out on Sierra, potentially resulting in a Gordon Bell award submission. They will also result in future Grand Challenge proposals, as well as a potential full LDRD submission to explore the new science that the enhanced code will enable.

Anticipated Challenges and Risk-Mitigation Steps: The primary challenge will be ensuring close collaboration and frequent communication with the Randles research group at Duke University, as the proposed changes to the core data structures and communication libraries of HARVEY will be pervasive and will need to be implemented in a coordinated fashion with the other development activities of the group. To mitigate this, we will hold bi-weekly code development calls. The PI will also travel to Duke University multiple times during the year to hold week-long development sprints. Other risks include data structure changes requiring more substantial refactoring of the code than currently anticipated, which will be mitigated by leveraging the efforts of the students and postdocs in the Randles group and delegating refactoring tasks where possible. A graduate student could be brought to LLNL as a summer intern if necessary.

Relevance to Mission Focus Area(s) and Core Competencies (1 paragraph):

Development of highly-scalable, performance-portable applications for next-generation architectures is a key focus area for LLNL.

Key Team members:

- Professor Amanda Randles, Duke University
- Jeff Ames, Duke University

Estimate of Proposal Budget Request:

- Early and Mid-Career Award, 0.2 FTE for PI time and travel.

[1] A. Randles, E. W. Draeger, and P. E. Bailey, "Massively parallel simulations of hemodynamics in the primary large arteries of the human vasculature," *J. Comp. Sci.*, vol. 9, pp. 70-75, 2015.

[2] A. Randles, E. W. Draeger, T. Ooppelstrup, L. Krauss, and J. A. Gunnels, "Massively parallel models of the human circulatory system," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, p. 1, ACM, 2015. (Gordon Bell finalist)