Bellerophon: A Computational Workflow Environment for Real-time Analysis, Artifact Management, and Regression Testing of Core-Collapse Supernova Simulations

Eric Lingerfelt
Oak Ridge National Laboratory
Oak Ridge, TN
lingerfeltej@ornl.gov

Bronson Messer
Oak Ridge National Laboratory
Oak Ridge, TN and
University of Tennessee
Knoxville, TN
bronson@ornl.gov

ABSTRACT

We present an overview of the Bellerophon software system, which has been built to support CHIMERA, a production-level HPC application that simulates the evolution of core-collapse supernovae. Developed over the last 5 years at ORNL, Bellerophon enables CHIMERA's geographically dispersed team of collaborators to perform job monitoring and real-time data analysis from multiple supercomputing resources, including platforms at OLCF, NERSC, and NICS. Its multi-tier architecture provides an encapsulated, end-to-end software solution that enables the CHIMERA team to quickly and easily access highly customizable animated and static views of results from anywhere in the world via a cross-platform desktop application. Bellerophon's robust artifact management system enables complete provenance tracking and provides direct access to all data files, renderings, and associated metadata though the client-side user interface. Bellerophon has quickly evolved into the CHIMERA team's de facto workflow environment for real-time analysis, artifact management, regression testing, and other workflow tasks.

Keywords

Real-time data analysis, core-collapse supernova simulations, HPC workflows, data management, software engineering.

1. INTRODUCTION

Running large-scale simulations is not enough to gain scientific insight from HPC applications. The development and scientific productivity of modern leadership-class HPC applications requires near-constant code development, with all code modifications motivated by well-planned scientific targets. The need to manage enterprise-scale code development and workflow, along with the need to perform custom data analysis, visualization, and robust data management with provenance tracking lead to a daunting set of additional requirements beyond the simulation code itself. The ongoing development of CHIMERA [1], a multi-dimensional, multi-physics code designed to study core-collapse supernovae, has alleviated these issues through the Bellerophon software system. Since going into production in 2010, Bellerophon has enabled the CHIMERA team to accelerate code development and scientific discovery of core-collapse supernova phenomena through real-time analysis and visualization, data and artifact management, automated software engineering tasks, and general workflow management.

2. CORE CAPABILITIES

The suite of software tools comprising Bellerophon's end user workflow environment is a mix of analysis and code development capabilities. Once an ongoing CHIMERA simulation is bound to Bellerophon, users are instantly provided the capability to robustly analyze and visualize data in real-time using highly customizable views [2]. Simulation data and all related artifacts are fully searchable by metadata stored in a MySQL database by user, HPC resource, date, and model. Besides analysis tasks, Bellerophon's automated regression test framework attempts to compile and execute the latest revision of CHIMERA every 24 hours on selected HPC resources [3]. The results of each test are transmitted to the central hub and processed for investigation through a specialized software tool. Bellerophon also provides workflow management through on-demand code repository statistics generation and integrates smoothly with other external workflow tools used by the team.

3. SYSTEM ARCHITECTURE AND DESIGN

Bellerophon's multi-tier architecture is realized by a system of distributed software components and services that communicate over the network. At the core of the new system is a web and data server enabling multiple, concurrent users to securely execute and access the system's core capabilities through Bellerophon's client-side desktop application. In addition, software components installed on HPC resources (e.g. at OLCF, NERSC, and NICS) enable CHIMERA simulations to perform co-scheduled real-time data analysis and transmit the results to the web and data server for processing, archiving, and rendering. Besides data and artifact management, the server allows users to execute a variety of programs and utilities as on-demand backend services. Bellerophon's user interface is deployed as a web-deliverable, cross-platform Java application and is implemented as a dashboard of easy-to-use software tools.

4. IMPACT AND UTILIZATION

Bellerophon has become the CHIMERA team's de facto work environment for analysis, artifact management, regression testing, and other workflow tasks. The system allows geographically distributed team members to share analysis results and tasks, test a new piece of code, or quickly generate a new visualization artifact. It also allows them to monitor long-running simulations (simulation runs can last many months of wallclock time). Use by the team has led to a Bellerophon-managed database containing 40 unique core-collapse supernova models with a total of 850 animations, 850,000 PNG images, and 110,000 binary and tabular data files under data management with provenance.

A significant impact of the Bellerophon system on CHIMERA development has been as a visual debugger. Multiple issues have been identified and resolved using Bellerophon. Those issues include memory errors, compiler issues, implementation errors, third-party module issues, and implementation bugs. Indeed, in many cases members of the CHIMERA team who were not directly overseeing the running simulations discovered the bugs via Bellerophon. This Bellerophon feature has directly impacted multiple scientific publications [e.g., 4,5].

5. ACKNOWLEDGMENTS

This work has been supported by the U.S. Department of Energy Offices of Nuclear Physics and Advanced Scientific Computing Research as well as the National Center for Computational Sciences at Oak Ridge National Laboratory. ORNL is managed by UT-Battelle, LLC for the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. This material is based upon work supported by the National Science Foundation under Grant numbers 0711134, 0933959, 1041709, and 1041710 and the University of Tennessee through the use of the Kraken computing resource at the National Institute for Computational Sciences. An award of computer time was provided by the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725, and resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

REFERENCES

[1] O. E. B. Messer, S. W. Bruenn, J. M. Blondin, W. R. Hix, A. Mezzacappa and C. J. Dirk. "Petascale supernova simulation with CHIMERA", J. Phys.: Conf. Ser., 78 012049 2007.

[2] E. J. Lingerfelt, O. E. B. Messer, S. S. Desai, C. A. Holt, E. J. Lentz. "Near Real-time Data Analysis of Core-Collapse Supernova Simulations With Bellerophon", ICCS 2014, Procedia Computer Science, Vol. 29, Pages 1504-1514, 2014.

[3] E.J. Lingerfelt, O.E.B. Messer, J.A. Osborne, R.D. Budiardja, A. Mezzacappa. "A Multitier System for the Verification, Visualization and Management of CHIMERA", ICCS 2011, Procedia Computer Science, Vol 4, Pages 2076-2085, 2011.

[4] S.W. Bruenn, A. Mezzacappa, W.R. Hix, E.J. Lentz, O.E.B. Messer, E.J. Lingerfelt, J.N. Blondin, E. Endeve, P. Marronetti, K. N. Yakunin. "Axisymmetric Ab Initio Core-Collapse Supernova Simulations of 12-25 MSolar Stars", ApJ 767, no. 1, L6, Apr. 10 2013.

[5] S.W. Bruenn, E.J. Lentz, E.J. Lingerfelt, A. Mezzacappa, W.R. Hix, J.N. Blondin, O.E.B. Messer, P. Marronetti. "Neutrinos and Supernovae", Hamburg Neutrinos from Supernova Explosions (HAvSE 2011) Proceedings, DESY, Hamburg, Germany, Dec 2011.