TN

MI

MP

EXASCALE ASTROPHYSICS WITH ENZO–E: DEVELOPMENT OF PHYSICS MODULES FOR GALAXY-SCALE AND COSMOLOGICAL SIMULATIONS

Andrew Emerick, Columbia University
2018–2019 Graduate Fellow

EXECUTIVE SUMMARY

The Enzo–E/Cello project involves a new astrophysical hydrodynamics code written to take advantage of the next generation of exascale computing systems currently under construction. This research has been under development for several years, during which time the code has demonstrated impressive scaling capabilities on test problems utilizing the entire Blue Waters system. In addition, the PI worked to implement key physics modules into Enzo–E/Cello that are required for its use on production-level science runs studying the evolution of individual galaxies and large, cosmological volumes of the universe. This code development is necessary to enable cutting-edge astrophysics research on upcoming next-generation systems.

RESEARCH CHALLENGE

Advancements in our understanding of the universe over the past several decades have largely been due to an investment in ever-improving high-performance computational resources and the development of computational codes to model astrophysical phenomena on these machines. However, much of the current generation of astrophysical hydrodynamics codes—used to study everything from planet formation, star formation, stellar evolution, galaxies, and cosmology—were first written in the 1990s, when scaling to hundreds or thousands of concurrent tasks was the newest technology. Many of these codes are not well suited to take full advantage of the next generation of exascale computing systems that will allow for millions (or more) of concurrent tasks. In particular, Enzo, a grid-based adaptive mesh refinement (AMR) code and one of the most widely used astrophysical hydrodynamics codes, has multiple design shortcomings that prevent it from making efficient use of next-generation computing systems. While continued optimization has improved Enzo's performance over the years, a complete refactoring is required. For this reason, the Enzo–E/Cello project began with the aim of completely redesigning the AMR hierarchy control system (Cello) to scale well to large systems while layering on modern physics modules and algorithms (Enzo–E). Like its predecessor, Enzo, this is an open-source code project that will be available to the astrophysical community for use in a broad range of contexts to tackle new challenging open questions in astronomy. In this project, the PI worked to develop and port the physics modules necessary to conduct the first set of science/production simulations with Enzo–E/Cello, including models for star formation, chemistry, radiative cooling and heating, stellar feedback, analytic gravitational potentials, and isolated galaxy initial conditions.

METHODS & CODES

This project utilized both the well-established AMR, cosmological, hydrodynamics code Enzo and the newly developed code project Enzo–E/Cello. Both codes are written predominantly in C++ with some underlying routines written in Fortran. While Enzo uses MPI for communication across processors/nodes, Enzo–E/Cello uses Charm++, a parallel programming library developed at the University of Illinois at Urbana–Champaign. The change means that Enzo–E/Cello utilizes task-based parallelism rather than domain decomposition or data parallelism.

RESULTS & IMPACT

The Enzo–E/Cello project has made significant strides over the past few years in developing a powerful, scalable AMR hydrodynamics code by demonstrating nearly ideal scaling using the entire Blue Waters computing system on a handful of simple test problems. In this project, the PI successfully implemented physics modules for star formation, stellar feedback, chemistry, radiative cooling and heating, analytic gravitational potentials, and isolated galaxy initial conditions. In addition, he helped improve the code with additional documentation, bug fixes, and error testing.

This work has enabled researchers to begin planning in detail the types of simulations possible with the new code that would not be possible with the current version of Enzo and other similar hydrodynamics codes. These projects will immediately target open questions in galactic evolution and the chemical enrichment of the gas and stars in and around galaxies, the diffuse gas—the circumgalactic and intergalactic media—that fills much of the universe, and the formation of the first- and second-generation of stars in the early universe. Finally, the code is entirely open source and well documented, allowing for its use by the broader astrophysical community.

WHY BLUE WATERS

Blue Waters has provided the computing environment and technical support needed to develop the new code and to conduct a variety of scaling and performance tests. This petascale machine and its support of large simulation runs with large node counts has enabled the PI to explore the limits of Enzo–E/Cello to better target future development and optimization.

As a sixth-year Ph.D. student in astronomy at Columbia University, Andrew Emerick successfully defended his dissertation in May 2019, having worked under the direction of Greg Bryan and Mordecai–Mark Mac Low.