



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Development Strategies for Modern Predictive Simulation Codes

A. Koniges, B. Gunney, R. Anderson, A. Fisher,
N. Masters

October 9, 2007

ParCo2007
Aachen/Juelich/Dusseldorf, Germany
September 2, 2007 through September 8, 2007

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.

Development Strategies for Modern Predictive Simulation Codes

Alice Koniges, Brian Gunney, Robert Anderson, Aaron Fisher, Nathan Masters

Lawrence Livermore National Laboratory,
Livermore, CA 94550,
U.S.A.

E-mail: koniges@llnl.gov

Modern simulation codes often use a combination of languages and libraries for a variety of reasons including reducing time to solution, automatic parallelism when possible, portability, and modularity. We describe the process of designing a new multiscale simulation code, which takes advantage of these principles. One application of our code is high-powered laser systems, where hundreds of laser beams are concentrated on centimeter-sized targets to enable the demonstration of controlled fusion and exciting new experiments in astrophysics and high-energy-density science. Each target must be carefully analyzed so that debris and shrapnel from the target will be acceptable to optics and diagnostics, creating new simulation regimes. These simulations rely on a predictive capability for determining the debris and shrapnel effects. Our new three-dimensional parallel code uses adaptive mesh refinement (AMR) combined with more standard methods based on Arbitrary Lagrangian Eulerian (ALE) hydrodynamics to perform advanced modeling of each different target design. The AMR method provides a true multiscale simulation that allows for different physical models on different scales.

We discuss our code development strategies. The code is built on top of the SAMRAI library (structured adaptive mesh refinement application interface) that provides scalable automatic parallelism. During the development phase of this code we have instituted testing procedures, code writing styles, and team coordination applicable to a rapidly changing source code, several dependent libraries, and a relatively small team including university collaborators with their own source and platforms. We use modern coding techniques and open source tools when possible for code management and testing including CppUnit, Subversion (and previously GNU ARCH), TiddlyWiki and group chat rooms. Additionally, we are conducting experiments aimed at providing a data set for validation of the fragmentation models. We describe our tools and our efforts in the area of code verification and validation.

1 Introduction

The high-performance computing environment continues to provide new opportunities for simulations of physical phenomena. Processes that a few years ago were too complex for accurate computer models are now amenable to computer solutions. Yet the simulation environment contains new challenges because of the diversity of tools and methodology that accompany these new simulations. The code developers must know not only the physics, engineering, and mathematics of their simulation, but also some expertise in a variety of programming languages, parallel libraries or languages, and tools for code development is required. We describe our experiences in the design of a new code based on an ALE-AMR method. The code is currently being used to model target fragmentation properties for high-powered laser systems. We start this paper with a discussion of the application environment and the new challenges for this particular area of simulation including the ALE-AMR approach. The next section Sec. 3 overviews some of the basic principles of our code development approach including team development methods, verification and

other issues. Automatic parallelism is obtained by calling parallel libraries and building the code on a parallel framework. We discuss this framework in Sec. 4. Finally, in Sec. 5 we give examples of some simulations.

2 Application Environment

The environment of a high-powered laser chamber contains a good deal of open space from the vacuum that surrounds the small fusion target, often a hohlraum, located at the chamber center. One goal of our simulation code is to model how the target dismantles after being hit by either laser beams or by x-rays that result from the lasers interacting with other various target components. For the purposes of our simulation the "target" includes not only the hohlraum where the laser is focused, but also ancillary target elements such as cooling rings, shields, or appendages that improve diagnostic capabilities. Pieces of the target that are closest to target center where the laser is focused will be vaporized and thus are relatively benign. However target components that are further from the main laser focus point are subject to lower levels of energy and therefore may be fragmented. It is important to determine the size of these fragmented pieces and their velocity vectors after the laser shot so that optics and diagnostics that line the chamber will be protected from damage. Dedicated experiments as well as experience from recent high-powered laser shots provides more information on this environment and the usefulness of mitigation procedures to direct fragments in benign directions.^{1,2}

This environment creates a natural basis for the combination of an ALE method with AMR. An Arbitrary Lagrangian Eulerian (ALE) Method is a standard technique whereby the evolution of the continuum equations is performed by an initial Lagrangian step which allows the mesh to deform according to the fluid motion. After this step, the mesh quality is evaluated and an optional remesh step is performed to recover either the original mesh (fully Eulerian) or more likely an intermediate or "arbitrary" mesh that is optimized for the simulation. Combining this method with Adaptive Mesh Refinement (AMR) allows one to not only gain advantages from the ALE technique but also to concentrate the calculation on areas of the domain where the fragments are being formed.

The ALE-AMR method was pioneered by Anderson and Pember in a series of works for hydrodynamics resulting in the basic code that we are using and updating for solution of problems in solid mechanics.³ Features of our ALE-AMR code that are relevant to the solid mechanics implementation include multimaterial interface reconstruction⁴ and material model specific features.⁵

3 Team Development and Tools

There are a variety of free or open source tools that make code development much easier than previously. We briefly discuss a few of these that our team has found useful. We also discuss our verification ideas.

Subversion⁶ is a version control system designed for the open source community, released under an Apache/BSD-style opens source license. Our group found the version control system particularly easy to use, and given that some of our developers had previous experience with systems like CVS, using this system was quite straight forward.

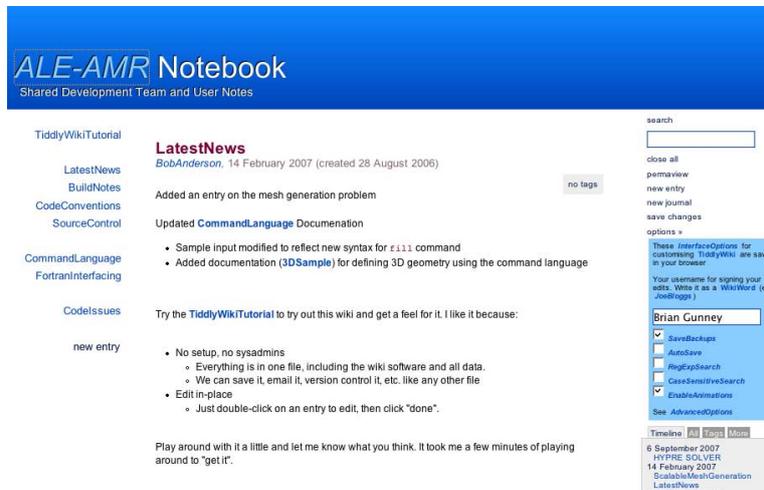


Figure 1. The ALE-AMR notebook demonstrates the use of a wiki for fast code documentation.

The term Wiki, as used in computer science is defined as a collaborative website, in which participants can directly edit the text of the "living" document. Perhaps the most famous implementation of a computer science Wiki is Wikipedia, the online free encyclopedia that is used written collaboratively by people throughout the world. Wiki in computer science is derived from the Hawaiian word Wiki which means to hurry. The notion of "hurry" or fast, is particularly important for us, since it encourages people to enter their documentation directly and quickly without being bound by a particular format. While the wiki itself is not a replacement for a traditional "user's manual," it serves as a convenient place to store daily or weekly information that is easily and quickly accessible. The particular wiki format we have chosen is TiddlyWiki⁷. A picture of our TiddlyWiki is given in Fig. 1. For reference, this figure shows the help message from the basic Wiki available online.

Another resource that our code team has used to great advantage is a "chat room." Use of the chat room allows developers who reside in physically different offices to discuss code changes, machine changes, strategies, etc. For the purposes of chat room security, the following policies are implemented: 1) Off-site chat does not work unless you set up an SSH tunnel to the server, 2) Not all rooms are "public"; a user must be invited before they can successfully join a room, 3) There may be more rooms in existence that what you see: rooms may be defined as invisible to outsiders. Our team has found that this strategy eliminates slow email question and answer periods and also encourages answers from developers who might not be queried on a particular issue. We have noticed a marked increase in the productivity of new team members when their chat room configuration was implemented.

A large portion of our code is based on libraries. The major library/framework that encompasses all of the parallelism is described in the next section. It is a patch-based Adaptive Mesh Refinement (AMR) library/framework. We also use the visualization package called VisIt⁸. VisIt, which can be downloaded from the web, is particularly well suited to

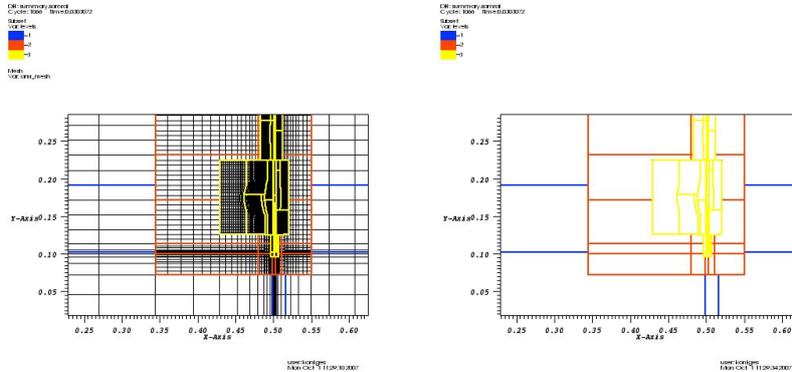


Figure 2. The left figure shows the mesh as well as the patch boundaries. Yellow contains the finest grid lines, which are fully black in some cases because of the refinement. The right most figure shows just the three levels of refined patches with grid turned off.

moving hydrodynamic simulations. Enhanced capabilities for parallelism and AMR grids are also included. Figure 2 shows example output from visIt. On the left, the full mesh and the patch boundaries are shown. On the right, just the patch boundaries are shown. Controls in visIt allow one to selectively look at patches and mesh pieces for a good understanding of the simulation. Automatic movie making commands are also included. VisIt also automatically handles groups of parallel files that are generated by the parallel code. Finally the visualization engine may be run remotely and/or in parallel. This visualization toolkit also helps in finding code errors and depicting how the automatic refinement works in the AMR scheme. This capability is useful in both finding bugs and also in determining the appropriate trigger for mesh refinement and de-refinement (coarsening).

4 The SAMRAI Framework

SAMRAI is an object-oriented structured AMR (SAMR) library supporting multiscale multiphysics applications and serving as a platform for SAMR algorithms research at LLNL.

SAMRAI factors out most of the common components and tasks in SAMR applications, allowing the application writers to focus on the specifics of their application. Among the most prominent tasks are mesh generation, mesh adaptation, data management, data transfer, high-level math operations, plotting interfaces, restart capabilities and parallelism. Less prominent time-saving services include computing mesh dependencies and mesh boundaries.

SAMRAI's interface design⁹ is flexible and expandable while supporting powerful features. Typical mesh data (cell-centered, node-centered and face-centered) are provided. One can seamlessly integrate arbitrary user-defined data without having to modify SAMRAI. data, such as arbitrary data sparsely residing in cells. For example, arbitrary data residing in sparse cells has been used to represent surfaces cutting through the

mesh. Application writers implement a small set of operations whose interfaces are defined by SAMRAI, and SAMRAI integrates the new data as if it were natively supported data. Data interpolation and copying is similarly integrated. SAMRAI natively supports Cartesian meshes, though its interface allows user-defined meshes such as deforming meshes such as that used by our ALE-AMR code.

Parallelism is largely hidden, allowing application developers to focus on serial code development and still build a parallel applications. All mesh data communications occur outside the interface seen by the application developers. User codes required for supporting user-defined data are all serial. Techniques such as combining messages are employed to enhance communication efficiency.

SAMRAI has been serving as a research platform in new SAMR algorithms, resulting in significant recent advances such as box search algorithms¹⁰, clustering algorithms¹¹, load balancing algorithms¹² and a new approach for managing SAMR meshes on distributed memory computers¹². Figure 4 shows the differences between traditional and new approaches and their associated weak scaling characteristics. The baseline results, from current widely used SAMR algorithms shows that although the serial physics engine scales well, other costs increase rapidly and quickly overwhelms the simulation. The new results gives much more ideal scaling behavior.

5 Code Applications

The code is being used to model the process of fragmentation and dismantling of high-powered laser facility targets. Figure 4 shows spall off the back of a thin target element that has energy deposition at the front surface of the thin plate in a localized region. This simulation is similar to Fig. /refpatchfig that shows the ALE-AMR grid and VisIt options. Here, the energy deposition first causes a plasma blow-off at the front surface of the plate. The top half of the figure is color-coded to show density contours and the bottom half of the figure is color-coded to show temperature contours. The plasma blow-off boundary is reconstructed via the multimaterial scheme described in Masters, et al.⁴ The interface between the metal and the gas is shown black. AMR patch boundaries are given in light blue. If this shield were part of a target configuration, we would use the code to determine the direction and quantity of plasma blow-off to the left, and the size and velocity of particle spall off the back. Spall off the back comes from the pressure wave that propagates horizontally through the material following the energy deposition at the front face.

Another simulation of interest to high-powered laser facilities is what happens to material that surrounds a can-shaped hohlraum target. For instance, for the target to be cryogenically cooled, it is often surrounded by metal washer-shaped cooling rings. Figure 5 shows a simulation of the failure and fragmentation of a typical Al cooling ring subject to an impulsive loading. Annular shaped spall planes break off from the main cooling ring. Finally, Fig. 6 shows a modification of a cooling ring simulation as described also in Koniges, et al.² Here the ring was not a complete annulus, as in Fig. ??, but instead was notched for experimental purposes. Simulations of this configuration showed how potentially damaging pieces could form when the stress in the ring were not distributed uniformly in the annulus.

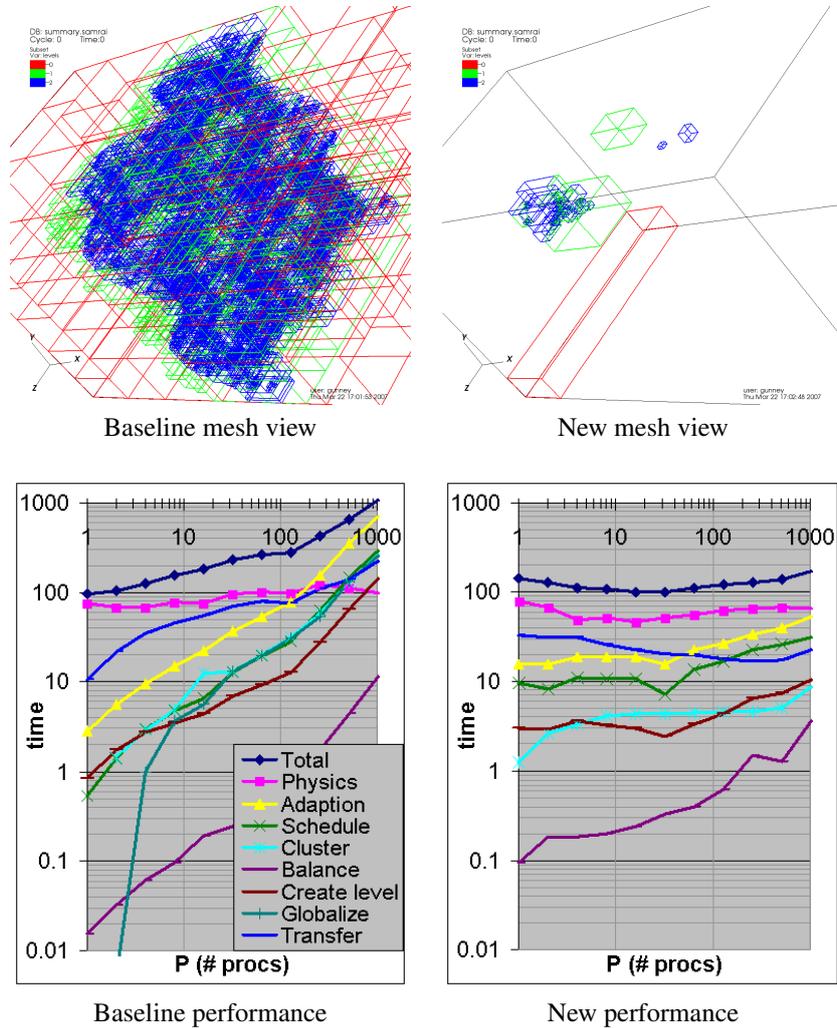


Figure 3. The top figures show the mesh as seen by the baseline approach and the new approach. The baseline approach stores and manages the global mesh. The new approach distributes the mesh information to reduce the amount of work per processor. The bottom figures shows the weak scaling characteristics for common tasks in an SAMR applications. For ideal weak scaling, the time for a given task does not vary with the number of processors. The new approach achieves more ideal scaling.

6 Summary

In this paper we have described the development of a new simulation code by a team of roughly 4 - 6 people. The code uses a variety of languages and some of the modern interaction techniques such as a wiki and a chat room are described. The code is scalable and parallel, based on the use of an advanced library framework known as SAMRAI. The use of SAMRAI tends to isolate the code developers working on the physics from code

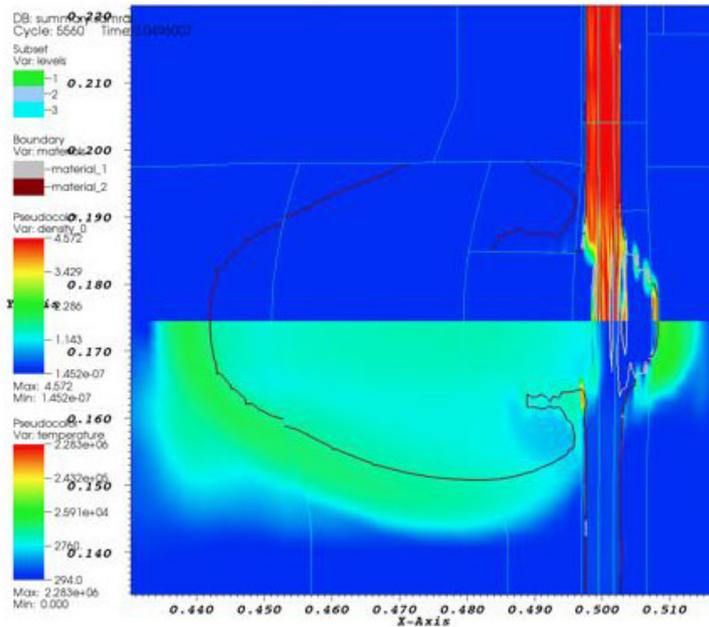


Figure 4. Formation of spall planes off the back of a thin target is shown. Additionally, contours of density (top) and temperature (bottom) are given. Interface reconstruction between the thin target metal and the gas surrounding the target is shown in black. AMR patch boundaries are given in light blue.

developers working on the parallel scalability and parallel programming. The code is providing new results for an upcoming area of physics/engineering, namely the protection of high-powered laser facility target chambers.

Acknowledgments

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

References

1. A. E. Koniges, R. W. Anderson, P. Wang, B. T. N. Gunney, R. Becker, *Modeling NIF Experimental Designs with Adaptive Mesh Refinement and Lagrangian Hydrodynamics* J. Physics IV **133**, 587–593 (2006)
2. A. E. Koniges, et al., J. Physics, submitted 2007.
3. R. W. Anderson, N. S. Elliott, and R. B. Pember *An arbitrary Lagrangian–Eulerian method with adaptive mesh refinement for the solution of the Euler equations*, J. Comp. Phys. **199**, 598–617 (2004).
4. N. Masters, et al., J. Physics, submitted 2007.

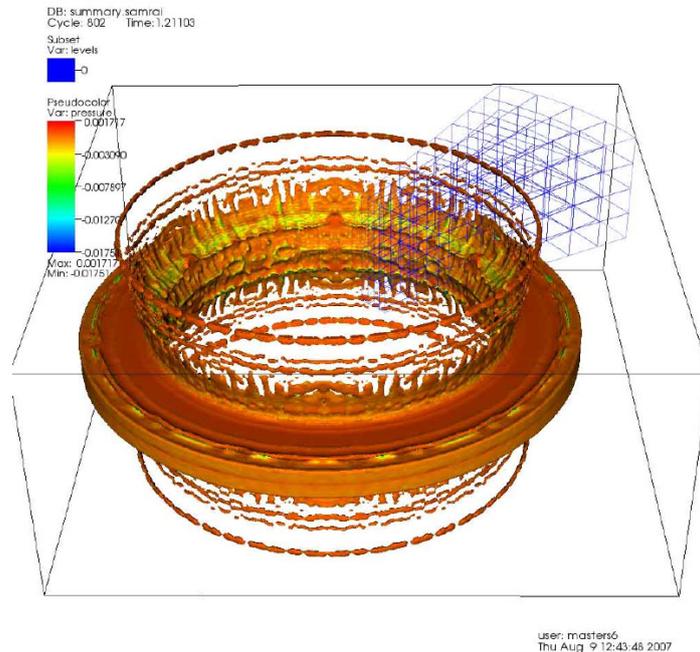


Figure 5. A ring of metal is impulsively loaded from the center to demonstrate the fragmentation pattern are target cooling rings.

5. A. Fisher, et al., J. Physics, submitted 2007.
6. <http://subversion.tigris.org/>
7. <http://www.tiddlywiki.org/>
8. <http://www.llnl.gov/visit/doc.html>
9. Richard D. Hornung and Scott R. Kohn. Managing application complexity in the SAMRAI object-oriented framework. *Concurrency and Computation: Practice and Experience*, 14:347–368, 2002.
10. Andrew M. Wissink, David Hysom, and Richard D. Hornung. Enhancing scalability of parallel structured AMR calculations. In *Proceedings of the 17th ACM International Conference on Supercomputing (ICS03)*, pages 336–347, San Francisco, June 2003.
11. Brian T. N. Gunney, Andrew M. Wissink, and David A. Hysom. Parallel clustering algorithms for structured AMR. *Journal of Parallel and Distributed Computing*, 66(11):1419–1430, November 2006.
12. Brian T. N. Gunney. Scalable metadata and algorithms for structured amr. SIAM Computational Science and Engineering Conference, February 2007. UCRL-PRES-228081.

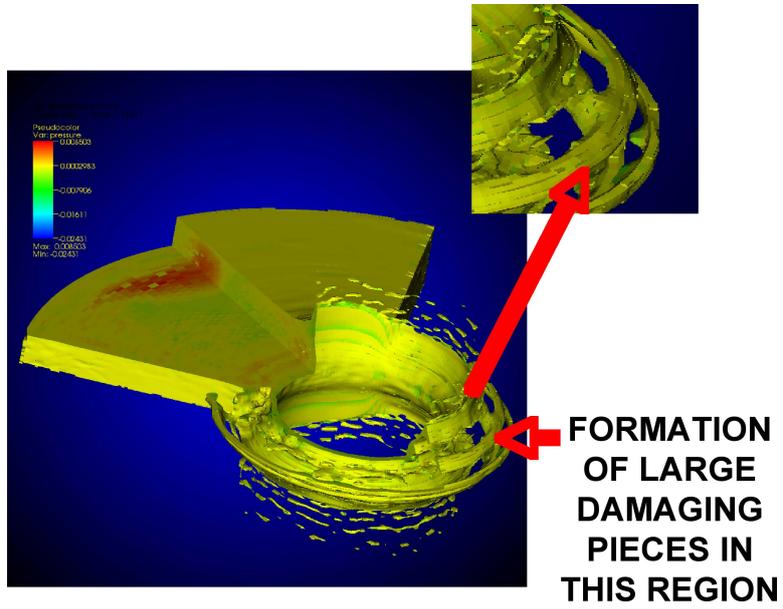


Figure 6. A notched ring simulation shows the formation of potentially damaging pieces during its fragmentation process.