ARL-TR-8618 ● JAN 2019


US Army Research Laboratory

# Quantifying Uncertainty from Computational Factors in Simulations of a Shaped Charge Jet

by Daniel J Hornbaker

## NOTICES

### Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

ARL-TR-8618 ● JAN 2019

US Army Research Laboratory

# Quantifying Uncertainty from Computational Factors in Simulations of a Shaped Charge Jet

**by Daniel J Hornbaker**
*Weapons and Materials Research Directorate, ARL*

# REPORT DOCUMENTATION PAGE

*Form Approved*
*OMB No. 0704-0188*



| 1. REPORT DATE *(DD-MM-YYYY)* | 2. REPORT TYPE | 3. DATES COVERED (From - To) |
|---|---|---|
| January 2019 | Technical Report | August 2017–October 2018 |

**4. TITLE AND SUBTITLE**
Quantifying Uncertainty from Computational Factors in Simulations of a Shaped Charge Jet

**5a. CONTRACT NUMBER**

**5b. GRANT NUMBER**

**5c. PROGRAM ELEMENT NUMBER**

**6. AUTHOR(S)**
Daniel J Hornbaker

**5d. PROJECT NUMBER**
1L162618AH80

**5e. TASK NUMBER**

**5f. WORK UNIT NUMBER**

**7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)**
US Army Research Laboratory
ATTN: RDRL-WMP-E
Aberdeen Proving Ground, MD 21005-5066

**8. PERFORMING ORGANIZATION REPORT NUMBER**
ARL-TR-8618

**9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)**

**10. SPONSOR/MONITOR'S ACRONYM(S)**

**11. SPONSOR/MONITOR'S REPORT NUMBER(S)**

**12. DISTRIBUTION/AVAILABILITY STATEMENT**
Approved for public release; distribution is unlimited.

**13. SUPPLEMENTARY NOTES**

**14. ABSTRACT**
Uncertainty in simulations of a shaped charge jet device is quantified using an approach established in a prior study of a long rod penetrator. The focus of this study is on computational factors, such as domain structure, coordinate system placement, and object orientation. Depth of penetration of the jet into a solid steel block is the primary quantity of interest, with computational time being a secondary metric. Current versions of two different simulation codes, CTH and ALEGRA, are examined.

**15. SUBJECT TERMS**
numerical simulation, uncertainty quantification, CTH, ALEGRA, shaped charge jet

| 16. SECURITY CLASSIFICATION OF: | | | 17. LIMITATION OF ABSTRACT | 18. NUMBER OF PAGES | 19a. NAME OF RESPONSIBLE PERSON |
|---|---|---|---|---|---|
| a. REPORT | b. ABSTRACT | c. THIS PAGE | | | Daniel J Hornbaker |
| Unclassified | Unclassified | Unclassified | UU | 51 | 19b. TELEPHONE NUMBER (Include area code) 410-278-7697 |

## Contents

## List of Figures



## List of Tables

## Acknowledgments

I would like to thank Stephen Schraml, Robert Doney, and George Vunni for their reviews of this work. Special thanks to Stephen Schraml for his assistance in investigating alternate build configurations of CTH, and to Steven Segletes for his illuminating discussions on Jones-Wilkins-Lee equations of state. I would also like to thank the CTH and ALEGRA development teams for their continuing support and collaboration.

This work was carried out under the Terminal Ballistics for Lethality and Protection Sciences Frontier Project, and was supported in part by a grant of computer time from the Department of Defense High Performance Computing Modernization Program at the US Army Research Laboratory Department of Defense Supercomputing Resource Center.

# 1. Introduction

Previously, a simple model system consisting of a tungsten heavy alloy rod impacting a steel block was used to study the sensitivity of simulation outcomes to changes in fundamental computational parameters for two simulation codes commonly employed in the development and evaluation of combat protection systems.[1] The final depth of penetration (DoP) was seen to vary by less than 0.1% of the mean value across the factors examined, apart from domain resolution and the orientation of the system within the computational domain. The small outcome uncertainty in this simple model system is in line with expectations, as these types of simulation codes were initially developed to model ballistic impacts.[2]

Modern battlefield protection systems typically have to account for engagements more complex than simple ballistic impact. One threat type of particular concern is shaped charge jets (SCJs). SCJ devices use high explosives to collapse a cone into a thin stream of high-velocity material capable of penetrating deeply into targets.[3] Such threats can be difficult to simulate as they involve thin shell structures rapidly deforming under explosive pressure with complex flow fields and high velocities.

A complicating factor is that jet formation is sensitive to material variability and imperfections in device fabrication. SCJ devices can exhibit sizeable inherent variability in performance. Variability associated with physical effects is separate from the variability inherent in the numerical techniques of computational simulations, and there is ample research exploring the effects of geometric and material variability in simulations of SCJs.[4–9] Among computational factors, the most well-studied is domain resolution,[6–8] with occasional research including other factors such as material interface tracking and artificial viscosity effects.[9]

This study focuses on the sensitivity of SCJ simulations to computational factors. The objective is to quantify the uncertainty attributable to the fact that simulations are, at heart, numerical approximations. A simple scenario of an SCJ device fired into a steel block target is used, with system geometry and material models fixed to exclude variability arising from physical factors.

The methodology established in the prior study is closely followed. The computational setup is documented in the next section. Next is a description of the SCJ model system and a review of the metrics used to quantify outcomes. The computational study includes five general categories of factors: code variations, simulation execution, domain structure, time step controls, and physical invariance. More than 240 simulations were completed during this research, using over 11 million core-hours of computing time.

## 2. Computational Setup

The simulation codes used for this study are CTH[10] version 12.0 and ALEGRA[11] release 2017.11.06. Both are developed at Sandia National Laboratories and are capable of modeling the solid dynamics and shock physics of multiple deformable materials in up to 3 spatial dimensions.

Unless otherwise noted, simulations were carried out on an SGI ICE XA supercomputer, hostname Centennial, located at the US Army Research Laboratory and managed by its Department of Defense Supercomputing Resource Center. Centennial has 1784 standard compute nodes, each with 40 2.2-GHz cores and 128 GB of memory.

Simulations employed uniform 3-D rectilinear domains with cubic cells. Resolution refers to the length of a cell edge; a simulation with a 0.05-cm resolution has a domain consisting of 0.05- × 0.05- × 0.05-cm cubic cells. With the exception of simulations employing symmetry planes, computational domains were sized to fully contain the entire problem geometry, with boundaries treated as free void having zero pressure and allowing mass to flow out of (but not into) the domain. Apart from the SCJ device and target, the domain is empty void space.

A mixed Lagrangian-Eulerian computational scheme is used in both CTH and ALEGRA. In the Lagrangian portion of each time step, the domain distorts in response to physical forces, and then in the Eulerian portion the resulting material state is mathematically mapped back onto the original undistorted grid.

Unlike the simpler long rod impact problem, material discards are necessary for these simulations to reliably run to completion. The presence of explosive material, as well as rapidly deforming SCJ liner material, can lead to unphysical thermodynamic material states, especially in cells with miniscule volume fractions of material and mixed material cells containing gaseous reactants. Simulations applied the following discard criteria to all cells in the domain:

1. All materials with temperatures exceeding 10,000 K
2. All materials with energies less than –100 kJ
3. Explosive material with pressure less than 1 atm and density less than 0.05 g/cm$^3$ (this removes low pressure gaseous products)
4. Explosive material within 1 cm of the target face (to prevent having a sometimes troublesome layer of mixed material cells across the target face)

A handful of ALEGRA simulations required an additional discard to eliminate instances of small material volume fractions obtaining anomalously high density.

Both codes have an array of parameters available for the user to customize the numeric approach, many of which were explicitly assigned in the input files to provide a thorough accounting of code configuration. A brief description of some of these parameters for each code is provided in the following subsections. This information will be of interest mainly to users of these codes; other readers can proceed directly to Section 3.

## 2.1 CTH Configuration

The default MMP0 option was used to model thermodynamics in cells containing multiple materials. This option partitions volume changes and work energy in proportion to the volume fractions of materials in a cell.

The TBAD parameter was set to 1e30 to allow simulations to continue regardless of the number cells with potentially unrealistic thermodynamics states encountered.

Updated fracture logic was enabled by setting FRAC = 1.

The default energy convection control, in which internal energy is conserved and any resulting discrepancies in kinetic energy are discarded, was selected by setting CONVECTION = 0.

The Sandia Modified Youngs' Reconstruction Algorithm[12] was employed for tracking material interfaces. A special fragment-moving model is implemented in CTH to handle motion of subcell material fragments embedded within a different material. This model was disabled for both void and explosive material.

A zero-velocity threshold value of 0.001 cm/s was used. Material with velocity less than this value had velocity changed to be exactly 0 cm/s.

Time steps were set to 0.55 times the Courant stability limit calculated by the code. The maximum allowed time step ratio for subsequent cycles was capped at 1.068.

## 2.2 ALEGRA Configuration

IGNORE KINEMATIC ERRORS was enabled with the default limit on the stretch tensor eigenvalue (RESET EIGENVALUE = le-5). This is an error-handling scheme for dealing with various numerical problems that can arise in computing the stretch tensor.

PISCES HOURGLASS CONTROL was used with VISCOSITY = 0.05 to stiffen the zero-energy deformation modes of domain cells.

INTERNAL ENERGY ADVECTION was used, in which internal energy is conserved during remap and errors in kinetic energy are discarded.

The Sandia Modified Youngs' Reconstruction Algorithm[12] was used to track material interfaces. The material advection scheme was set to the default MODERATE ADVECTION, which implements a third-order advection method for cells containing only a single material, with various alternatives for cells containing multiple materials. In ALEGRA, momentum is a node-centered quantity that uses a separate advection method from that used for cell-centered quantities. The default Half-Interval Shift method[13] was used for this.

The ISENTROPIC MULTIMATERIAL ALGORITHM was used to determine state variables in multimaterial cells, with the PRESSURE RELAXATION and TEMPERATURE RELAXATION algorithms turned on and THERMAL EQUILIBRIUM set to off.

Like CTH, ALEGRA computes a maximum stable time step based on numerical constraints imposed by the physics of the problem. In addition, a maximum allowed change in cell volume during the domain deformation phase was imposed by setting MAXIMUM VOLUME CHANGE = 0.5. By default, time steps are set to 0.9 of the calculated maximum (TIME STEP SCALE = 0.9). The maximum allowed time step ratio for subsequent cycles was capped at 1.068.

## 3. SCJ System

The model system under study in this report consists of a precision laboratory 65-mm-diameter SCJ device detonated into a solid 15- × 15- × 50-cm target (Fig. 1). This device consists of a conical copper liner with a 1.0-mm wall thickness surrounded by an LX-14-0 explosive charge having a maximum diameter of 65 mm.[4] LX-14-0 consists of 95.5 wt% HMX explosive with a 4.5 wt% polyurethane elastomer binder.[14] The target block is modeled as rolled homogeneous armor (RHA) steel.[15] For reference, at a standoff distance of 2 charge diameters (CDs), real 65-mm SCJ devices penetrate an average of 389.0 mm into stacks of 50.8-mm-thick RHA plates (average of 18 tests, with a standard deviation of 10.5 mm).[16]

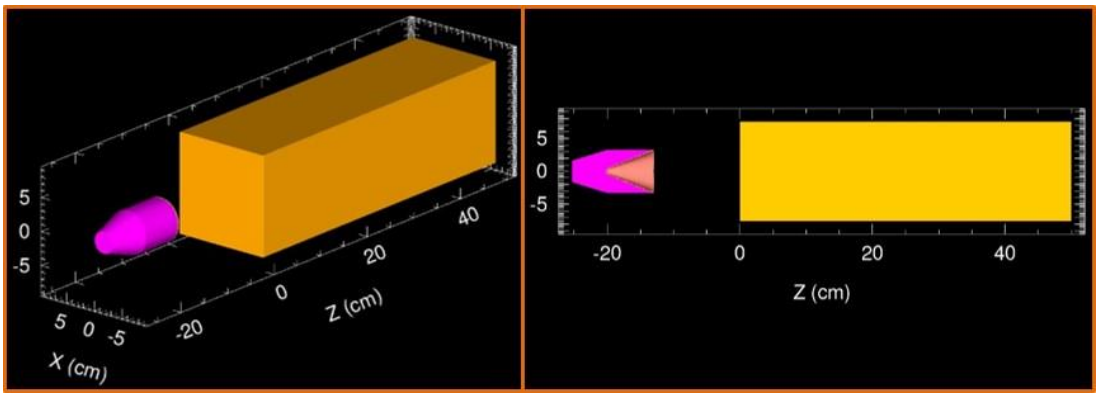


**Fig. 1 Perspective (left) and cut-away (right) rendering of the SCJ system in CTH**

The SCJ device is positioned 2 CDs (130 mm) from the end of the liner to the face of the target block in the model system, with the cylindrical axis of the device oriented normal to the center point of the target face. While standoff distances of SCJ devices are often specified from the base of the liner cone, the 65-mm SCJ has a 3-mm cylindrical lip extending from the cone base, making the distance from the cone base to the target 133 mm. In practice, this lip is often ignored and positioning is measured from the end of the liner to the target.

A DIATOM model was constructed for the 65-mm SCJ. DIATOM constructions are a space-filling process for sequentially adding or deleting materials from a computational domain in CTH or ALEGRA. In each step of the construction, the DIATOM algorithm determines which cells intersect a defined object and performs the desired material insertion or deletion. A key aspect is that insertion operations only replace empty void with material, with any pre-existing material in a cell unaffected. For example, if a 10-cm radius sphere of material 1 is inserted at the origin, followed by insertion of an 11-cm radius sphere of material 2 at the origin, the end result is a solid 10-cm-radius sphere of material 1 surrounded by a 1-cm-thick shell of material 2. Reversing the order of operation would simply create an 11-cm-radius sphere of material 2, with no material 1 present since all the target cells were fully occupied.

Forming thin-shell structures such as conical SCJ liners in DIATOM requires some care. Consider the example of the larger material-2 sphere insertion followed by the slightly smaller material-1 sphere, but in a coarsely meshed rectilinear domain. The surface of the first sphere will consist of partially filled cells. If the sphere in the second insertion overlaps some of these partially filled cells, the DIATOM algorithm will replace the available void space with material 1, despite the fact that the specific geometric relationship between the spheres should preclude this. The end result would be an 11-cm-radius sphere of material 2 with patches of material

1 decorating the surface. In DIATOM, geometric surfaces only matter to the extent that they determine the overlap with domain cells; intersections with other geometric objects are not considered. This can lead to "bleed-through" of the surrounding explosive charge through the liner in an SCJ construction. A general approach to avoid this is to construct complex DIATOM objects outward from the center using a temporary placeholder material to fill empty volumes. Inserted components can then be intentionally overlapped with previously placed objects to ensure that surfaces are fully in contact with no stray void spaces remaining. The placeholder material is deleted in the final step of the construction.

The SCJ model system involves three different materials. Both copper and RHA are modeled using the Mie-Grüneisen equation-of-state (EOS)[17] and Johnson-Cook constitutive equation,[18] which are described in the previous report.[1] The Johnson-Cook fracture model is used to model material failure.[4,19] Model parameters for these materials are summarized in Table 1.

**Table 1 Material model parameters for copper and RHA steel**

| Parameter | Symbol | Copper | RHA | Units |
|---|---|---|---|---|
| Initial density | $\rho_0$ | 8.930 | 7.850 | g/cm$^3$ |
| Initial temperature | $T_0$ | 298 | 298 | K |
| Specific heat capacity | $C_v$ | 0.393 | 0.446 | J/gK |
| Grüneisen parameter | $\Gamma_0$ | 1.99 | 1.67 | . . . |
| Sound speed | Cs | 3940 | 4529 | m/s |
| Mie-Grüneisen model parameters | $S_1$ | 1.489 | 1.490 | . . . |
| | $S_2$ | 0 | 0 | . . . |
| Johnson-Cook viscoplastic model parameters | A | 89.7 | 780.0 | MPa |
| | B | 291.9 | 780.0 | MPa |
| | C | 0.025 | 0.004 | . . . |
| | m | 1.09 | 1.00 | . . . |
| | n | 0.310 | 0.106 | . . . |
| Melt temperature | $T_m$ | 1381 | 1783 | K |
| Poisson's ratio | $\nu$ | 0.330 | 0.294 | . . . |
| Johnson-Cook fracture model parameters | D1 | 0.54 | –0.80 | . . . |
| | D2 | 4.89 | 2.10 | . . . |
| | D3 | –3.03 | –0.50 | . . . |
| | D4 | 0.014 | 0.002 | . . . |
| | D5 | 1.12 | 0.61 | . . . |
| Fracture strength | $P_f$ | –1.50 | –2.50 | GPa |

EOS and constitutive model parameters for RHA were previously documented,[1] while fracture model parameters were set to those for 4340 steel from Johnson and Holmquist.[20]

EOS parameters for copper originated from fits of Hugoniot data published by Los Alamos National Laboratory.[21] The specific heat capacity of 0.393 J/gK used here is similar to the value of 0.385 J/gK reported for copper in reference literature.[22] The constitutive and fracture model parameters for oxygen-free high thermal conductivity copper reported in Johnson and Holmquist were used.[20] The melt temperature of 1,381 K was taken from the materials library for CTH and ALEGRA, though it is about 25 K higher than the reference literature value of 1,356–1,358 K.[22,23] Poisson's ratio was set to that of typical C11000 copper.[23] The fracture strength of the material was estimated at –1.50 GPa.[9]

CTH employs an additional temperature-related constitutive model parameter TMELT. If the material temperature in a cell exceeds TMELT, the constitutive model is bypassed and the flow stress is set to zero. For both RHA and copper, TMELT was set to an extremely high temperature (»1e20 K) to ensure the constitutive model was always evaluated.

The remaining material is the LX-14-0 explosive, which provides the pressure wave needed to collapse the liner into a jet. The solid explosive is static and experiences no mechanical deformation prior to detonation, so no EOS or constitutive model was implemented for the material in this form.

Detonation is modeled with a programmed burn routine in which a predetermined initiation geometry is used to compute a detonation time for each cell containing explosive material within a specified burn radius. Detonation was initiated at time zero (the first cycle of the simulation) on a planar disc covering the entire rear surface of explosive charge to force a perfect plane wave. The burn radius was set to 20 cm to ensure detonation times for all cells were directly computed from the initiation geometry, as opposed to being estimated using an alternative detonation front propagation model. The detonation velocity $D_{CJ}$ for LX-14-0 is 8,800 m/s.[24] The rate at which the chemical energy released by detonation is added to the gaseous reactant material is determined by the burn front thickness, which in these simulations is fixed at two cells. Consequently, the physical thickness of the detonation front varies with mesh resolution.

The thermodynamic state of the gaseous explosive products is determined using a version of a Jones-Wilkins-Lee (JWL) EOS,[25] in which the specific energy $E$ and pressure $P$ are given by

$$E(\rho,T) = \frac{1}{\rho_0}\left(\frac{A}{R_1}e^{-R_1\rho_0/\rho} + \frac{B}{R_2}e^{-R_2\rho_0/\rho} - \varepsilon_0\right) + C_V T\,, \qquad (1)$$

$$P(\rho,T) = Ae^{-R_1\rho_0/\rho} + Be^{-R_1\rho_0/\rho} + \omega\rho C_V T\,, \qquad (2)$$

where $\rho$ is the density of the reactant gas, $\rho_0$ is the initial density of the unreacted explosive, $T$ is temperature, $\varepsilon_0$ is the volumetric detonation energy density of the explosive (10.2 GJ/m$^3$ for LX-14-0[24]), and $C_v$ is the specific heat capacity of the products, assumed to be constant. A JWL EOS refers generally to an EOS designed to reproduce the empirically measured behavior of an explosive characterized using a cylindrical expansion test.[26] The specific form shown here is but one example.[27]

Two underlying assumptions of the cylinder expansion experiments used to characterize explosives are that the detonation proceeds through the Chapman-Jouguet (CJ) state, and that the detonation products expand adiabatically. The JWL equations

$$P = A\left[1 - \frac{\omega V_0}{R_1 V}\right]e^{-R_1 V/V_0} + B\left[1 - \frac{\omega V_0}{R_2 V}\right]e^{-R_2 V/V_0} + \frac{\omega U}{V}\,, \qquad (3)$$

$$P(V, S_{CJ}) = Ae^{-R_1 V/V_0} + Be^{-R_2 V/V_0} + C\frac{V_0^{\omega+1}}{V^{\omega+1}}\,, \qquad (4)$$

characterize the pressure $P$ – volume $V$ – energy $U$ relationship for the detonation products, with pressure advancing along the CJ adiabat $S_{CJ}$. The set of parameters [$A$, $B$, $C$, $R_1$, $R_2$, $\omega$] are fit to match the observed outcomes of the experiment. These parameters (excepting $C$, not to be confused with $C_v$) are used in the equations for the JWL EOS. Parameter values for LX-14-0 are listed in Table 2.[24]

A final consideration is determining an appropriate value for the specific heat capacity $C_v$. While $C_v$ can be specified by user, it is typically calculated in the codes by applying the JWL EOS to the CJ state of the explosive:

$$C_V = \frac{P_{CJ} - Ae^{-R_1\rho_0/\rho_{CJ}} - Be^{-R_1\rho_0/\rho_{CJ}}}{\omega\rho_{CJ}T_{CJ}}\,. \qquad (5)$$

The CJ pressure $P_{CJ}$ for LX-14-0 is 37 GPa.[24] The default code value of 4,062 K for the CJ temperature $T_{CJ}$ was used. The CJ density $\rho_{CJ}$ is calculated using the Rayleigh line relation for shock velocity, which is the same as the detonation velocity $D_{CJ}$ in the CJ condition. The initial state for this calculation is the solid explosive, with density $\rho_0$ =1.835 g/cm$^3$ and ambient pressure $P_0 \ll P_{CJ}$:

$$P_{CJ} - P_0 \cong P_{CJ} = D_{CJ}^2\,\rho_0\left(1 - \frac{\rho_0}{\rho_{CJ}}\right). \qquad (6)$$

Solving Eq. 6 with the parameters from Table 2 gives $\rho_{CJ}$ = 2.481 g/cm$^3$, which then yields a value of $C_v$ = 0.513 J/gK for LX-14-0.

**Table 2 Material model parameters for LX-14-0 explosive**

| Parameter | Symbol | LX-14-0 | Units |
|---|---|---|---|
| Initial density | $\rho_0$ | 1.835 | g/cm$^3$ |
| Initial temperature | $T_0$ | 298 | K |
| CJ pressure | $P_{CJ}$ | 37.00 | GPa |
| CJ temperature | $T_{CJ}$ | 4,062 | K |
| Detonation velocity | $D_{CJ}$ | 8,800 | m/s |
| Detonation energy density | $\varepsilon_0$ | 1.02e10 | J/m$^3$ |
| Specific heat capacity (calculated) | $C_v$ | 0.513 | J/gK |
| JWL model parameters | A | 826.1 | GPa |
| | B | 17.24 | GPa |
| | $R_1$ | 4.55 | . . . |
| | $R_2$ | 1.32 | . . . |
| | $\omega$ | 0.38 | . . . |

In the baseline configuration of the SCJ system, the faces of the target are aligned normal to the coordinate axes, and the strike face is on the $z$ = 0 plane and centered on the coordinate origin. The SCJ is oriented to advance in the positive $z$ direction. Domain size is set to provide a minimum 2-cm void buffer around the SCJ device and target. All boundaries are free void, with no symmetry planes used. Simulations of the baseline configuration were run to 500 μs to provide ample time for the penetration process to complete. Table 3 summarizes the baseline configuration parameters.

**Table 3 SCJ system baseline configuration**

| Resolution (cm) | Void buffer (cm) | Jet axis | Domain cells | | | | Stop time (μs) |
|---|---|---|---|---|---|---|---|
| | | | X | Y | Z | Total | |
| 0.0500 | 2.00 | +z | 380 | 380 | 1587 | 229,162,800 | 500 |

Figure 2 shows simulation output for the baseline configuration of the SCJ system in CTH. Images were generated during simulation run time in both CTH and ALEGRA using the integrated Spymaster utility.[28] Images on the left show a perspective view of the system in 100-μs increments of simulated time. On the right

is a cutaway view sectioned along the $x = 0$ plane to reveal the penetration process within the interior of the block. The positive $z$-axis is oriented to the right and the positive $y$-axis is up in these views. Materials are visualized by coloring domain cells that are at least half-filled with a material, with different colors representing different materials. Explosive products were colored to indicate pressure in the cutaway view, which is why the gas is colored various shades of blue in the 100 μs image. All other cells in the domain are transparent. These images reveal that the explosive products reach atmospheric pressure and are fully discarded from the simulation by 200 μs, and the penetration process is largely complete by 400 μs.

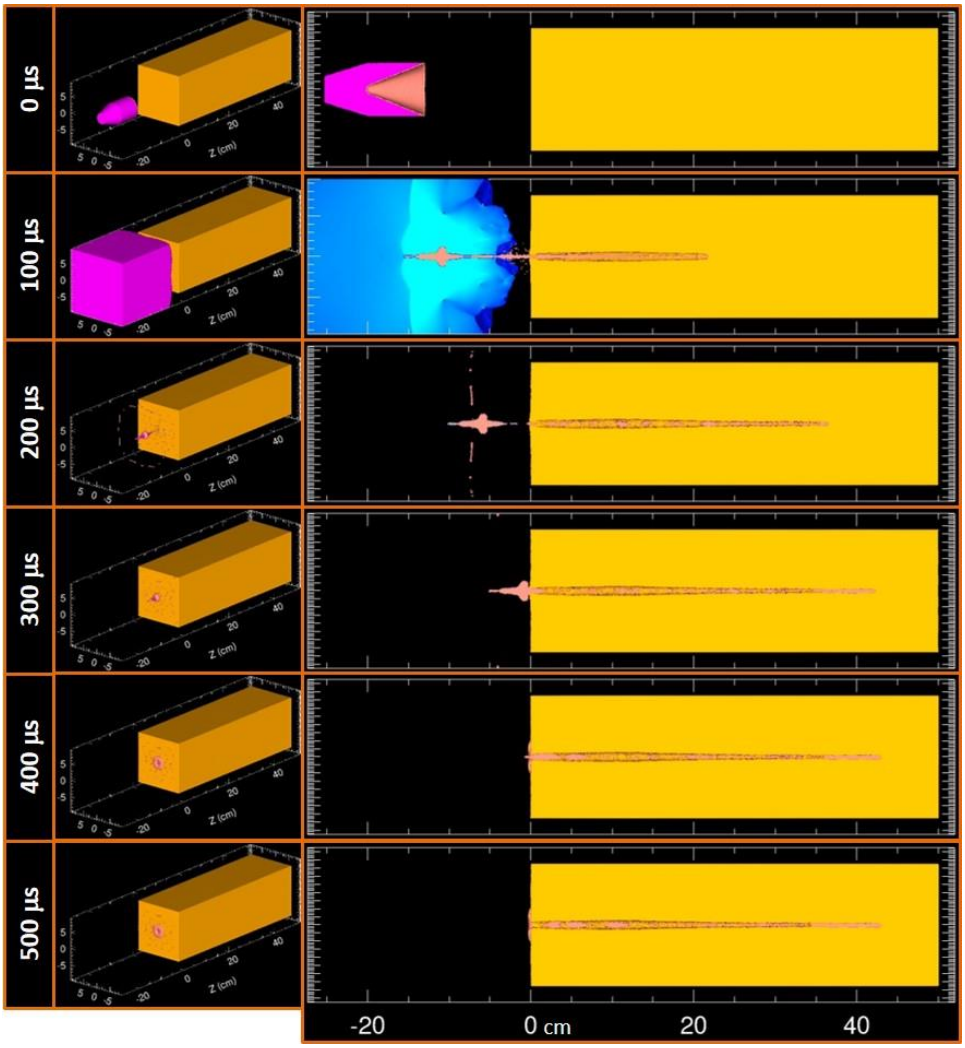


**Fig. 2 Perspective (left) and section (right) views in 100-μs increments for a CTH simulation of the SCJ system in the baseline configuration**

## 4. Quantities of Interest

The primary output in these simulations is the DoP of the jet into the RHA target block. DoP can be calculated by measuring the distance from the farthest point of jet material advancement to the back-face plane of the target and then subtracting from the initial target thickness. Three different methods employed in the prior work for performing this measurement were used here.

The most precise method uses domain cell data. In CTH, positions for all cells having at most 99.95% void by volume were extracted using Spymaster's MPISPYPLT postprocessing program. In ALEGRA, version 5.3.0 of the open-source ParaView application[29] was used to export data for cells containing at least 0.05% by volume of either threat or target material. Jet position is determined by the location of the cell farthest along the threat axis that contains jet material, while the position of the back-face plane of the target is determined with a more complex process described in the previous report.

A second method uses tracer particles to track jet advancement into the target. Tracers are massless virtual objects that can move along with surrounding material and serve as localized point probes. Unlike the previous work with a long rod projectile, it was not possible at the time this study was undertaken to place tracers in a way that produced an array located at predetermined positions along the length of the jet. Recent updates to both the CTH and ALEGRA codes now provide means for embedding tracers at desired positions within a dynamically forming jet, but in this work only a single tracer was placed at the base of the liner on the jet axis. This tracer was constrained to move only along the jet axis (the *z* axis in the baseline configuration) when possible. As the jet advances, this tracer gets "pushed" forward. In principle, the symmetry of the problem should result in a symmetric penetration channel with the tracer located at the deepest point of penetration. The position of the back face of the target block was tracked by placing tracers at each of the rear corners, offset one cell length from each face. The position of the back face at the end of a simulation was computed as the average position along the threat axis of these four tracers, adjusted to account for the offset.

The final method involves examining section images of the target in 1-mm increments along the threat axis to locate both the deepest point of jet advancement and the back face of the target, as described in the previous report. Since positions in this method are only determined to the nearest millimeter, this method carries a measurement uncertainty of ±1 mm.

A secondary quantity of interest in this report is the time needed for a simulation to reach completion, called run time. A related quantity is the computational time for

a simulation, which is the run time multiplied by the total number of cores used. Computational time should be constant for a simulation on an ideal parallel computer. For example, doubling the number of cores used would reduce the run time by half. In reality, communication of information between cores leads to non-ideal scaling.

Simulations of the baseline configuration for the model SCJ system were performed using 512 cores in CTH and 1,024 cores in ALEGRA. The results listed in Table 4 form the point of comparison for variations of the baseline configuration, and will be repeated in subsequent tables of results.

**Table 4 Simulation results for baseline configuration**

| Code | Cores | | | | Cycles | Run time (h) | DoP (cm) | | |
|---|---|---|---|---|---|---|---|---|---|
| | X | Y | Z | Total | | | Cell data | Tracers | Images |
| CTH | 4 | 4 | 32 | 512 | 25,906 | 17.55 | 42.400 | 38.42 | 42.4 |
| ALEGRA | 4 | 8 | 32 | 1,024 | 16,402 | 55.42 | 42.800 | 42.79 | 42.7 |

There is good agreement between the codes regarding final DoP for the baseline configuration. The tracer result for CTH stands out as coming up short of the other measurements, highlighting a common problem with using tracers with SCJ threats. The listed DoP is based on tracer positions at the simulation stop time of 500 μs. The maximum tracer DoP over the course of the simulation was 38.75 cm at 371 μs, still significantly short of the DoP indicated by the other methods.

There are two factors at work here. As seen in Fig. 2, the penetration process is largely completed by 371 μs, but jet material continues to accumulate at the end of the penetration channel. While this material does not advance further into the target, it does create advection in the accumulated jet material, causing the tracer to drift. This is a minor effect compared to the issue illustrated in Fig. 3, which shows the end of the penetration channel deviating from the jet axis. This causes the tracer, whose motion is constrained to the *z* axis, to be stranded short of the channel end. One might question whether removing the constraint on tracer motion would allow it to follow the penetration channel off axis, but in such cases the tracer quickly becomes embedded in the wall of the penetration channel by eroding jet material.

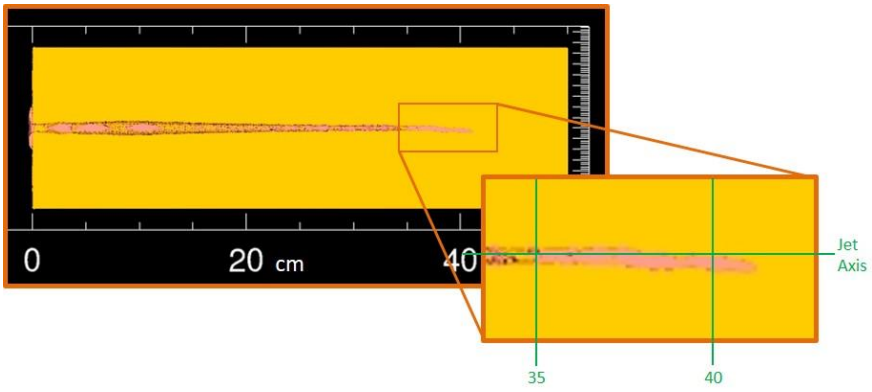


**Fig. 3 CTH simulation of baseline configuration at 500 μs**

Another common feature of SCJ simulations exhibited by these results is that the average simulated DoP of 426 mm is greater than the average experimental DoP of 389 mm for real devices. The greater penetration of the simulated device is a direct consequence of how the model is constructed. The simulated device has a perfectly symmetric geometry and is composed of perfectly homogeneous materials, and the jet is formed from an explosive charge with a perfectly planar detonation front. In other words, the device being simulated is an ideal SCJ device. Real devices have small local variances in geometry and properties, and jet formation is extremely sensitive to these factors. Any departure from symmetry leads to imperfections in the resulting jet alignment, degrading total penetration. That the simulated results are only 10% greater than the average experimental result is a testament to the high quality of these precision lab devices, which is further bolstered by the small observed standard deviation of 10.5 mm.

## 5. Uncertainty from Variations in Computational Factors

The focus of this work is on the effect that changes in fundamental computational parameters, such as those affecting the basic structure of the computational domain, have on outcomes in simulations of an SCJ device. Factors that typically receive close attention, such as material model parameters, are held fixed. Experimental sources of variance, such as uncertainty in device standoff or alignment with the target, are not considered. A prime motivation for this study is to isolate the influence of basic computational factors from these other sources of uncertainty.

This section is subdivided into five general themes. It opens with an excursion examining how successive versions of the CTH and ALEGRA codes perform on the prior long rod problem to explicitly measure any variance due to code changes.

The section on Computational Execution covers factors involved in running a simulation that are external to the simulation input script, such as the choice of

platform and number of cores used. The Computational Domain section examines parameters affecting the structure of the domain and the spatial placement of simulated objects. Time Step Control groups factors affecting the temporal progress of a simulation. Physical Invariance covers tests of code fidelity to invariances in the real world, such as those related to changing frames of reference.

## 5.1 Code Version

The prior study was conducted using CTH version 11.1 and ALEGRA release 21 May 2015, primarily on the Hercules platform, with a few comparison simulations conducted on Excalibur. To test the consistency of results between code versions, the baseline configuration of the model long rod system from the prior report was simulated using more recent code versions on both Excalibur and Centennial.

Table 5 shows the results for various combinations of platform and code version for CTH on 64 cores. The first three runs listed are taken from Table 6 of the previous report. As discussed there, these simulations were completed with the TMELT parameter inadvertently left at the default value of 1,490 K. Run 4 shows the reported results of the baseline long rod configuration on Hercules with TMELT »1e20 K. The remaining simulations are new to this work.

**Table 5 CTH code version study results**

| Run | Code version | Platform | Cores/ node | Cycles | Run time (h) | DoP (cm) | | |
|---|---|---|---|---|---|---|---|---|
| | | | | | | Cell data | Tracers | Images |
| 1* | 11.1 | Hercules | 16 | 7,771 | 17.84 | 8.700 | 8.69 | 8.7 |
| 2* | 11.1 | Excalibur | 16 | 7,778 | 17.92 | 8.700 | 8.70 | 8.7 |
| 3* | 11.1 | Excalibur | 32 | 7,778 | 18.11 | 8.700 | 8.70 | 8.7 |
| 4 | 11.1 | Hercules | 16 | 7,861 | 18.13 | 8.650 | 8.67 | 8.7 |
| 5 | 11.1 | Excalibur | 32 | 7,853 | 18.23 | 8.650 | 8.67 | 8.7 |
| 6 | 11.2 | Excalibur | 32 | 7,846 | 17.42 | 8.650 | 8.67 | 8.7 |
| 7 | 11.2 | Centennial | 32 | 7,851 | 13.23 | 8.650 | 8.67 | 8.7 |
| 8 | 12.0 | Excalibur | 32 | 7,847 | 17.46 | 8.650 | 8.67 | 8.7 |
| 9 | 12.0 | Centennial | 32 | 7,840 | 13.07 | 8.650 | 8.67 | 8.7 |

*TMELT parameter was not changed from the default value in these simulations.

In terms of the primary metric, once TMELT is set, the DoP of the rod remains the same across all platforms and code versions tested. But as indicated by the changing number of computational cycles, different code versions do result in mathematically different solution states. What stands out most among these

simulations is that Centennial achieves solutions faster than either Excalibur or Hercules.

Table 6 shows the results for various combinations of platform and code version for ALEGRA on 512 cores. The first three runs listed are taken from Table 7 of the previous report, while the remaining entries are new to this work.

**Table 6 ALEGRA code version study results**

| Run | Code version | Platform | Cores/ node | Cycles | Run time (h) | DoP (cm) | | |
|---|---|---|---|---|---|---|---|---|
| | | | | | | Cell data | Tracers | Images |
| 1 | 21 May 2015 | Hercules | 16 | 8,820 | 23.40 | 8.450 | 8.48 | 8.5 |
| 2 | 21 May 2015 | Excalibur | 16 | 8,813 | 23.33 | 8.450 | 8.47 | 8.5 |
| 3 | 21 May 2015 | Excalibur | 32 | 8,813 | 23.50 | 8.450 | 8.47 | 8.5 |
| 4 | 2016.10.14 | Excalibur | 32 | 8,820 | 23.76 | 8.450 | 8.47 | 8.5 |
| 5 | 2016.10.14 | Centennial | 32 | 8,801 | 23.02 | 8.450 | 8.47 | 8.5 |
| 6 | 2017.11.06 | Excalibur | 32 | 8,808 | 24.24 | 8.450 | 8.47 | 8.5 |
| 7 | 2017.11.06 | Centennial | 32 | 8,796 | 22.44 | 8.450 | 8.47 | 8.5 |

As in the CTH results, DoP remains unchanged across all platforms and code versions tested, but the solution states are mathematically different. Improvements in time-to-solution on Centennial are modest for the ALEGRA code.

Generally speaking, both codes exhibit good stability of solution across different versions. While the end states themselves are not mathematically identical, the primary output of interest does not change by a measurable amount for the long-rod system.

## 5.2 Computational Execution

Only parameters relating to execution of simulations, such as total number of cores and active cores per node, are varied in this subsection. Fixed input files for the baseline configuration were used throughout.

A replication study was performed using multiple identical simulations (Tables 7 and 8). The simulations for each code all ran concurrently, meaning they used different sets of nodes on Centennial.

**Table 7 CTH replication study results**

| Run | Cycles | Run time (h) | DoP (cm) | | |
|---|---|---|---|---|---|
| | | | Cell data | Tracers | Images |
| **BL** | **25906** | **17.55** | **42.400** | **38.42** | **42.4** |
| 2 | 25906 | 17.56 | 42.400 | 38.42 | 42.4 |
| 3 | 25906 | 17.57 | 42.400 | 38.42 | 42.4 |
| 4 | 25906 | 17.59 | 42.400 | 38.42 | 42.4 |
| 5 | 25906 | 17.58 | 42.400 | 38.42 | 42.4 |
| 6 | 25906 | 17.61 | 42.400 | 38.42 | 42.4 |
| 7 | 25906 | 17.59 | 42.400 | 38.42 | 42.4 |

**Table 8 ALEGRA replication study results**

| Run | Cycles | Run time (h) | DoP (cm) | | |
|---|---|---|---|---|---|
| | | | Cell data | Tracers | Images |
| **BL** | **16402** | **55.42** | **42.800** | **42.79** | **42.7** |
| 2 | 16402 | 54.83 | 42.800 | 42.79 | 42.7 |
| 3 | 16402 | 54.72 | 42.800 | 42.79 | 42.7 |
| 4 | 16402 | 54.87 | 42.800 | 42.79 | 42.7 |
| 5 | 16402 | 54.85 | 42.800 | 42.79 | 42.7 |
| 6 | 16402 | 54.98 | 42.800 | 42.79 | 42.7 |
| 7 | 16402 | 54.88 | 42.800 | 42.79 | 42.7 |

Run BL, highlighted in bold, is the baseline configuration simulation from Table 4. The basic expectation that all nodes on the system are identical is met, as all simulations for each code achieved mathematically identical solution states. The average run time for the CTH simulations was 17.58 h, with a standard deviation of just 73 s, which is 0.12% of the mean. For ALEGRA, the mean and standard deviation were 54.94 h and 817 s for all simulations in the table. But the BL simulation is something of an outlier in this group, and the average run time for just the six replicate simulations was 54.86 h, with a much smaller standard deviation of 303 s, 0.15% of the mean.

In the next study, the number of cores used on each compute node was varied. Reducing the number of cores per node affects the computation in two ways. The first is to apportion the available memory on each node among fewer cores. The second effect is to increase the number of nodes the simulation is spread across.

Results for CTH (Table 9) and ALEGRA (Table 10) demonstrate that solution states are not affected by the number of cores used per node. Reductions in run times can be seen as the available memory per core increases.

**Table 9 CTH cores per node study results**

| Run | Cores/ node | Total nodes | Cycles | Run time (h) | DoP (cm) | | |
|---|---|---|---|---|---|---|---|
| | | | | | Cell data | Tracers | Images |
| **BL** | **32** | **16** | **25,906** | **17.55** | **42.400** | **38.42** | **42.4** |
| 2 | 16 | 32 | 25,906 | 15.75 | 42.400 | 38.42 | 42.4 |
| 3 | 8 | 64 | 25,906 | 12.58 | 42.400 | 38.42 | 42.4 |
| 4 | 4 | 128 | 25,906 | 11.52 | 42.400 | 38.42 | 42.4 |

**Table 10 ALEGRA cores per node study results**

| Run | Cores/ node | Total nodes | Cycles | Run time (h) | DoP (cm) | | |
|---|---|---|---|---|---|---|---|
| | | | | | Cell data | Tracers | Images |
| **BL** | **32** | **32** | **16,402** | **55.42** | **42.800** | **42.79** | **42.7** |
| 2 | 16 | 64 | 16,402 | 50.05 | 42.800 | 42.79 | 42.7 |
| 3 | 8 | 128 | 16,402 | 42.54 | 42.800 | 42.79 | 42.7 |

The final study in this section involved varying the total number of cores used for a simulation. The specific domain decomposition used for each total core count was determined by the codes; user-specified decompositions are investigated in the next section. Results for CTH simulations using 64 to 4,096 cores are shown in Table 11.

**Table 11 CTH total cores study results**

| Run | Cores | Cores/ node | Decomposition | | | Cycles | Run time (h) | DoP (cm) | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | | | X | Y | Z | | | Cell data | Tracers | Images |
| 2 | 64 | 32 | 2 | 2 | 16 | 28,026 | 116.62 | 41.650 | 38.54 | 41.7 |
| 3 | 128 | 32 | 2 | 4 | 16 | 25,529 | 56.19 | 44.050 | 44.01 | 44.0 |
| 4 | 256 | 32 | 4 | 4 | 16 | 25,126 | 31.00 | 41.050 | 36.88 | 41.0 |
| **BL** | **512** | **32** | **4** | **4** | **32** | **25,906** | **17.55** | **42.400** | **38.42** | **42.4** |
| 5 | 1,024 | 32 | 4 | 8 | 32 | 26,435 | 9.86 | 44.350 | 41.78 | 44.3 |
| 6 | 2,048 | 32 | 8 | 8 | 32 | 27,515 | 5.52 | 42.100 | 38.49 | 42.1 |
| 7 | 4,096 | 32 | 8 | 8 | 64 | 28,465 | 3.43 | 42.700 | 39.38 | 42.7 |

These outcomes clearly show different solution states resulting from the change in the total number of cores. The underlying cause of this variance, demonstrated more directly in the next section, is the changing decomposition of the problem, which alters the boundaries along which the simulation is parallelized. Put succinctly, CTH 12.0 on Centennial is not decomposition invariant.

This behavior was also observed for CTH 11.1 in the prior report. For version 12.0, the CTH development team implemented changes to the code aimed at making it decomposition invariant. When tested on some systems, such as Hercules, CTH 12.0 does indeed display decomposition invariance. But on other platforms the code is not invariant to decomposition.

The culprit here is a computational factor that is not explicitly studied in this work, but is implicitly related to observed variances in simulation results across different platforms, which is the process for compiling source code into executable form. CTH 12.0 became available in June 2017, and was initially built on various platforms using procedures recommended by the CTH development team. This generated executables that functioned as intended on some systems, but not others.

In consultation with the CTH development team, several alternate build strategies were tried. The only approach found that ensured decomposition invariance was to compile the code with all optimizations disabled. However, this resulted in simulation run times that were 1.5–2.7 times longer than for the standard build. As the development of a satisfactory build strategy remains an open issue, I decided to proceed with the standard build of CTH 12.0 for this study and use this opportunity to illustrate a factor that is not typically considered in routine computational work, but is nevertheless a potential source of variance. Later in this report we will see further consequences of unintended code behavior, where unexpected results will provide a strong clue regarding differences between alternate builds of the code.

In addition to changing solution states, the results in Table 11 also show the first sizeable variances in DoP encountered thus far. The average DoP as measured using cell data is 42.61 cm, with a standard deviation of 1.21 cm, which is 2.8% of the mean. Image measurements of DoP track well with cell data values, but tracer measurements continue to be unreliable, and will be omitted after this section.

Run times for these simulations exhibit the inefficient scaling of nonideal parallel systems. The total computational size of the problem actually decreases slightly from 64 cores (7,464 core-h) to 128 cores (7,192 core-h), but then increases monotonically with core count all the way to 14,049 core-h at 4,096 total cores.

Results for ALEGRA simulations are shown in Table 12. The cores per node were reduced to 16 at 512 total cores to allow enough memory per core for the simulation.

ALEGRA is decomposition invariant, and produced the same solution state regardless of core count. Near-ideal scaling held for run times across the range of total cores examined, which was necessarily narrower than that examined for CTH due to the fact that ALEGRA simulations are roughly 5 times the computational size of comparable CTH simulations.

**Table 12 ALEGRA total cores study results**

| Run | Cores | Cores/ node | Decomposition | | | Cycles | Run time (h) | DoP (cm) | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | | | X | Y | Z | | | Cell data | Tracers | Images |
| 2 | 512 | 16 | 4 | 4 | 32 | 16,402 | 97.65 | 42.800 | 42.79 | 42.7 |
| **BL** | **1,024** | **32** | **4** | **8** | **32** | **16,402** | **55.42** | **42.800** | **42.79** | **42.7** |
| 3 | 2,048 | 32 | 8 | 8 | 32 | 16,402 | 27.51 | 42.800 | 42.79 | 42.7 |
| 4 | 4,096 | 32 | 8 | 8 | 64 | 16,402 | 14.58 | 42.800 | 42.79 | 42.7 |

## 5.3 Computational Domain

To isolate the effects of domain decomposition, simulations were performed using different user-specified decompositions over a fixed number of total cores. Decompositions are specified within the input file, and this is the first study in this report requiring changes to the SCJ baseline configuration input file, as opposed to simply altering the execution parameters of the code.

Results for different decompositions in CTH simulations using 512 cores are shown in Table 13. Run 2 shows that explicitly specifying the default decomposition for the baseline configuration in the input file does not affect the outcome.

**Table 13 CTH decomposition study results**

| Run | Decomposition | | | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|---|---|
| | X | Y | Z | | | Cell data | Images |
| **BL** | **4** | **4** | **32** | **25,906** | **17.55** | **42.400** | **42.4** |
| 2 | 4 | 4 | 32 | 25,906 | 17.55 | 42.400 | 42.4 |
| 3 | 8 | 8 | 8 | 27,321 | 18.78 | 41.650 | 41.6 |
| 4 | 1 | 2 | 256 | 27,330 | 31.37 | 41.900 | 41.9 |
| 5 | 2 | 1 | 256 | 25,865 | 30.63 | 43.500 | 43.5 |
| 6 | 1 | 4 | 128 | 26,192 | 21.97 | 45.050 | 45.0 |
| 7 | 2 | 2 | 128 | 26,440 | 22.27 | 42.900 | 42.9 |
| 8 | 4 | 1 | 128 | 26,793 | 24.18 | 42.050 | 42.0 |
| 9 | 1 | 8 | 64 | 28,028 | 19.81 | 44.250 | 44.2 |
| 10 | 2 | 4 | 64 | 32,947 | 23.02 | 41.550 | 41.5 |
| 11 | 4 | 2 | 64 | 27,854 | 20.28 | 43.300 | 43.3 |
| 12 | 8 | 1 | 64 | 26,647 | 22.06 | 44.700 | 44.7 |
| 13 | 1 | 16 | 32 | 27,242 | 19.35 | 44.800 | 44.8 |
| 14 | 2 | 8 | 32 | 26,308 | 17.03 | 40.400 | 40.4 |
| 15 | 4 | 32 | 4 | 28,172 | 24.94 | 44.400 | 44.4 |
| 16 | 32 | 4 | 4 | 26,992 | 29.86 | 44.500 | 44.5 |
| 17 | 2 | 16 | 16 | 28,082 | 18.75 | 43.250 | 43.2 |
| 18 | 16 | 2 | 16 | 27,105 | 24.23 | 41.050 | 41.0 |
| 19 | 16 | 16 | 2 | 32,394 | 33.80 | 42.900 | 42.9 |
| 20 | 4 | 8 | 16 | 26,071 | 17.07 | 41.200 | 41.2 |
| 21 | 4 | 16 | 8 | 25,764 | 18.08 | 45.650 | 45.6 |

The average DoP in these simulations (excluding the original baseline simulation to avoid double counting) is 43.07 cm, with a standard deviation of 1.52 cm, 3.5% of the mean. The absolute range of outcomes is fairly sizable, from 40.40 cm to 45.65 cm. This stands in stark contrast to the results seen for the long rod model system, which exhibited no measurable variation in DoP with decomposition. However, this variance would not occur for decomposition invariant builds of the code, and the DoP would have a constant value.

There is a considerable spread in simulation run times, from 17 h to nearly 34 h. The default decomposition is within 0.5 h of the fastest run time, with most alternative decompositions taking significantly longer to complete.

The ALEGRA results for various decompositions using 1,024 cores are shown in Table 14. The decomposition invariance produces a constant DoP. Recalling that the baseline simulation run time was something of an outlier, run 2 is more indicative of typical default run time. Only two other decompositions clock in with faster run times, with run 13 besting it by 1 h.

**Table 14 ALEGRA decomposition study results**

| Run | Decomposition | | | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|---|---|
| | X | Y | Z | | | Cell data | Images |
| **BL** | **4** | **8** | **32** | **16,402** | **55.42** | **42.800** | **42.7** |
| 2 | 4 | 8 | 32 | 16,402 | 54.86 | 42.800 | 42.7 |
| 3 | 2 | 2 | 256 | 16,402 | 65.64 | 42.800 | 42.7 |
| 4 | 2 | 4 | 128 | 16,402 | 59.28 | 42.800 | 42.7 |
| 5 | 4 | 2 | 128 | 16,402 | 58.41 | 42.800 | 42.7 |
| 6 | 4 | 4 | 64 | 16,402 | 56.31 | 42.800 | 42.7 |
| 7 | 1 | 32 | 32 | 16,402 | 60.26 | 42.800 | 42.7 |
| 8 | 2 | 16 | 32 | 16,402 | 54.74 | 42.800 | 42.7 |
| 9 | 8 | 8 | 16 | 16,402 | 53.90 | 42.800 | 42.7 |
| 10 | 8 | 16 | 8 | 16,402 | 55.67 | 42.800 | 42.7 |
| 11 | 16 | 8 | 8 | 16,402 | 55.99 | 42.800 | 42.7 |
| 12 | 8 | 4 | 32 | 16,402 | 54.49 | 42.800 | 42.7 |
| 13 | 16 | 2 | 32 | 16,402 | 53.84 | 42.800 | 42.7 |
| 14 | 32 | 1 | 32 | 16,402 | 61.16 | 42.800 | 42.7 |

The computational domain of the model SCJ system was sized such that objects were fully embedded within the domain and did not interact with the boundaries. Cells at the boundary are subject to special user-selected conditions to compensate for having less than a full complement of neighboring cells. Maintaining a buffer of void cells limits the influence of these conditions on the simulation solution, at the cost of increasing the size of the domain. Simulations with void buffers of different size were performed.

In CTH (Table 15), the total computational time was fairly proportional to the total number of cells in the domain, roughly 42 core-h per million cells. Different domain sizes generated different solution states, with the DoP varying from 39.4 cm to 44.6 cm, about a 5-cm range of outcomes.

**Table 15 CTH void buffer study results**

| Run | Void buffer (cm) | Cores | Domain cells | | | | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|---|---|---|---|---|
| | | | X | Y | Z | Total (M) | | | Cell data | Images |
| 2 | 0.50 | 512 | 320 | 320 | 1,527 | 156.4 | 25,840 | 13.64 | 39.350 | 39.4 |
| 3 | 1.00 | 512 | 340 | 340 | 1,547 | 178.8 | 26,104 | 15.33 | 42.500 | 42.5 |
| BL | 2.00 | 512 | 380 | 380 | 1,587 | 229.2 | 25,906 | 17.55 | 42.400 | 42.4 |
| 4 | 4.00 | 512 | 460 | 460 | 1,667 | 352.7 | 27,073 | 28.93 | 42.250 | 42.2 |
| 5 | 8.00 | 512 | 620 | 620 | 1,827 | 702.3 | 28,988 | 56.97 | 44.600 | 44.6 |

In ALEGRA (Table 16), there was not enough available memory to run simulations 4 and 5 on only 1,024 cores, so the total number of cores was increased for these. As with CTH, the total computational time was approximately proportional to domain size at about 252 core-h per million cells. Solution states also varied, with DoP values all falling within a 3-mm range.

**Table 16 ALEGRA void buffer study results**

| Run | Void buffer (cm) | Cores | Domain cells | | | | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|---|---|---|---|---|
| | | | X | Y | Z | Total (M) | | | Cell data | Images |
| 2 | 0.50 | 1,024 | 320 | 320 | 1,527 | 156.4 | 16,489 | 39.94 | 42.943 | 42.9 |
| 3 | 1.00 | 1,024 | 340 | 340 | 1,547 | 178.8 | 16,511 | 44.90 | 43.149 | 43.0 |
| BL | 2.00 | 1,024 | 380 | 380 | 1,587 | 229.2 | 16,402 | 55.42 | 42.800 | 42.7 |
| 4 | 4.00 | 2,048 | 460 | 460 | 1,667 | 352.7 | 16,467 | 42.28 | 42.949 | 42.9 |
| 5 | 8.00 | 4,096 | 620 | 620 | 1,827 | 702.3 | 16,462 | 42.93 | 42.900 | 42.9 |

Simulation run times can be significantly reduced for geometries having mirror plane symmetry by truncating the computational domain at the symmetry plane and imposing a mirror boundary condition. The SCJ baseline configuration has mirror plane symmetry on both the $x = 0$ and $y = 0$ planes.

Simulations using *x*- and/or *y*-symmetry were performed (Tables 17 and 18). The total number of cores were scaled along with symmetry to keep the partitioning of the problem fixed.

**Table 17 CTH symmetry study results**

| Run | Symmetry | Cores | Domain cells | | | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|---|---|---|---|
| | | | X | Y | Z | | | Cell data | Images |
| **BL** | **None** | **512** | **380** | **380** | **1,587** | **25,906** | **17.55** | **42.400** | **42.4** |
| 2 | X | 256 | 190 | 380 | 1,587 | 25,900 | 16.95 | 43.000 | 43.0 |
| 3 | Y | 256 | 380 | 190 | 1,587 | 28,835 | 19.01 | 40.650 | 40.7 |
| 4 | X&Y | 128 | 190 | 190 | 1,587 | 26,214 | 16.72 | 43.900 | 43.9 |

**Table 18 ALEGRA symmetry study results**

| Run | Symmetry | Cores | Domain cells | | | Cycles | Run Time (h) | DoP (cm) | |
|---|---|---|---|---|---|---|---|---|---|
| | | | X | Y | Z | | | Cell data | Images |
| **BL** | **None** | **1,024** | **380** | **380** | **1,587** | **16,402** | **55.42** | **42.800** | **42.7** |
| 2 | X | 512 | 190 | 380 | 1,587 | 16,198 | 52.78 | 42.745 | 42.7 |
| 3 | Y | 512 | 380 | 190 | 1,587 | 16,210 | 52.99 | 42.600 | 42.5 |
| 4 | X&Y | 256 | 190 | 190 | 1,587 | 15,987 | 51.33 | 43.484 | 43.4 |

Exploiting symmetry planes changes the solution states, with DoP values falling within a 32-mm range in CTH and a 9-mm range in ALEGRA. There is a small reduction in total computational time in most cases, consistent with the previously observed scaling with total numbers of cores in Tables 11 and 12. The increased computational time for the *y*-symmetry problem in CTH is tied to the marked increase in total time step cycles needed for completion.

The Cartesian coordinate system in the simulation codes is oriented along the edges of the cells, with the location of the origin set by the user. Since object geometries are defined relative to this origin, a shift in the origin point will translate objects within the domain. The origin position is fixed by specifying the minimum coordinate values in the domain. The minimum coordinate values in the baseline configuration were $x_{min} = y_{min} = -9.50$ cm and $z_{min} = -27.35$ cm. This placed the origin between cells 190 and 191 in *x* and *y*, and between cells 547 and 548 along *z*. This location corresponds to a corner of a cell, which is a node point of the mesh.

A study was conducted to examine the effect of shifting the origin small distances in different directions away from this node (Tables 19 and 20). "Shift Type" denotes the displacement of the origin from the baseline nodal position, with "dx" denoting one cell length in the *x*-direction, "dy" one cell length in the *y*-direction,

and so on. The origin was placed at different cell symmetry points, such as corners and face centers, as well as at an arbitrary point.

**Table 19 CTH origin shift study results**

| Run | Shift type | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|
| | | | | Cell data | Images |
| **BL** | **None** | **25,906** | **17.55** | **42.400** | **42.4** |
| 2 | –dx | 28,471 | 19.31 | 42.000 | 42.0 |
| 3 | –dy | 28,219 | 19.13 | 43.400 | 43.4 |
| 4 | –dz | 26,735 | 18.12 | 43.950 | 44.0 |
| 5 | –dx/2 | 26,003 | 17.68 | 45.750 | 45.8 |
| 6 | –dy/2 | 27,180 | 18.46 | 44.700 | 44.7 |
| 7 | –dz/2 | 26,103 | 17.74 | 45.450 | 45.5 |
| 8 | –(dx+dy)/2 | 25,753 | 17.52 | 42.750 | 42.7 |
| 9 | –(dx+dz)/2 | 28,378 | 19.24 | 41.100 | 41.1 |
| 10 | –(dy+dz)/2 | 26,918 | 18.33 | 45.600 | 45.6 |
| 11 | –(dx+dy+dz)/2 | 27,307 | 18.56 | 40.450 | 40.5 |
| 12 | –(dx+dy+dz)/4 | 28,750 | 19.54 | 42.900 | 42.9 |
| 13 | –0.177dx-0.571dy-0.601dz | 26,467 | 17.99 | 42.700 | 42.7 |

**Table 20 ALEGRA origin shift study results**

| Run | Shift type | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|
| | | | | Cell data | Images |
| **BL** | **None** | **16,402** | **55.42** | **42.800** | **42.7** |
| 2 | –dx | 16,622 | 55.09 | 42.900 | 42.9 |
| 3 | –dy | 16,455 | 54.50 | 42.696 | 42.7 |
| 4 | –dz | 16,515 | 54.74 | 42.900 | 42.9 |
| 5 | –dx/2 | 16,474 | 54.62 | 41.847 | 41.8 |
| 6 | –dy/2 | 16,459 | 54.56 | 42.650 | 42.6 |
| 7 | –dz/2 | 16,962 | 56.14 | 43.100 | 43.1 |
| 8 | –(dx+dy)/2 | 16,487 | 54.63 | 41.200 | 41.2 |
| 9 | –(dx+dz)/2 | 16,430 | 54.43 | 42.350 | 42.3 |
| 10 | –(dy+dz)/2 | 16,541 | 55.50 | 41.750 | 41.7 |
| 11 | –(dx+dy+dz)/2 | 16,434 | 54.43 | 41.800 | 41.8 |
| 12 | –(dx+dy+dz)/4 | 16,531 | 54.83 | 41.500 | 41.5 |
| 13 | –0.177dx-0.571dy-0.601dz | 16,459 | 54.59 | 41.400 | 41.4 |

There is about a 2-h spread in run times for both codes, with variations attributable to the number of cycles for each simulation to reach completion. The average computational rates were steady at 1,468–1,476 cycles/h for CTH and 296–302 cycles/h for ALEGRA.

The average DoP for the CTH simulations was 43.3 cm, with outcomes spanning a 5.3-cm range. The ALEGRA results were more tightly grouped, with an average DoP of 42.2 cm and a range of just 1.9 cm.

Cell resolution is typically the most important computational factor for a simulation. When certain mathematical criteria are met, the numerical approximation errors in a simulation become vanishingly small as the cell size goes to zero, and the solution state converges to the unique solution for the original continuum equations describing the problem. If these criteria are not satisfied, a simulation can fail to produce a reasonable solution, or any solution at all. In principle, error bounds for a simulation can be calculated, but this is generally not feasible in practice. Furthermore, no number of simulations can actually prove that solutions are “close” to a continuum solution. The best that can typically be accomplished is to demonstrate that at some resolution differences in simulation outcomes at nearby resolutions are within some acceptable tolerance.

A study was performed in which cell size was decreased from a starting value of 0.1 cm by factors of the cube root of 2 until simulations became impractical to run (Tables 21 and 22). The scaling factor corresponds to a doubling of the total number of cells in the domain at each step. The total number of cores used was scaled with domain size.

**Table 21 CTH resolution study results**

| Run | Resolution (cm) | Cores | Domain cells | | | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|---|---|---|---|
| | | | X/Y | Z | Total(M) | | | Cell data | Images |
| 2 | 0.1000 | 64 | 190 | 794 | 28.7 | 10,772 | 6.40 | 32.000 | 32.1 |
| 3 | 0.0794 | 128 | 240 | 999 | 57.5 | 14,629 | 8.96 | 36.422 | 36.6 |
| 4 | 0.0630 | 256 | 302 | 1,259 | 114.8 | 19,537 | 12.67 | 42.945 | 43.0 |
| **BL** | **0.0500** | **512** | **380** | **1,587** | **229.2** | **25,906** | **17.55** | **42.400** | **42.4** |
| 5 | 0.0397 | 1,024 | 479 | 1,998 | 458.4 | 35,268 | 23.81 | 40.353 | 40.3 |
| 6 | 0.0315 | 2,048 | 604 | 2,518 | 918.6 | 44,378 | 31.30 | 38.219 | 38.2 |
| 7 | 0.0250 | 4,096 | 760 | 3,173 | 1832.7 | 62,119 | 48.14 | 39.000 | 39.0 |
| 8 | 0.0198 | 8,192 | 960 | 4,006 | 3691.9 | 83,642 | 64.82 | 36.893 | 36.9 |

**Table 22 ALEGRA resolution study results**

| Run | Resolution (cm) | Cores | Domain cells | | | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|---|---|---|---|
| | | | X/Y | Z | Total(M) | | | Cell data | Images |
| 2 | 0.1000 | 128 | 190 | 794 | 28.7 | 7,802 | 23.54 | 22.500 | 22.4 |
| 3 | 0.0794 | 256 | 240 | 999 | 57.5 | 9,903 | 31.60 | 28.483 | 28.4 |
| 4 | 0.0630 | 512 | 302 | 1,259 | 114.8 | 12,750 | 41.44 | 34.943 | 34.9 |
| **BL** | **0.0500** | **1,024** | **380** | **1,587** | **229.2** | **16,402** | **55.42** | **42.800** | **42.7** |
| 5 | 0.0397 | 2,048 | 479 | 1,998 | 458.4 | 21,615 | 72.31 | 45.752 | 45.7 |

Total computational time (Fig. 4) increases proportionally to the inverse of the resolution to a power a little more than 4 (4.42 for CTH and 4.22 for ALEGRA), reflecting the discretization of the problem in 4 dimensions (3 spatial and 1 time). In addition to the doubling of the number of cells at each increment, smaller cells necessitate smaller time steps, and thus more cycles are needed to obtain a solution. The increasing number of cores required to run simulations at higher resolution also imposes additional costs due to inefficient scaling.

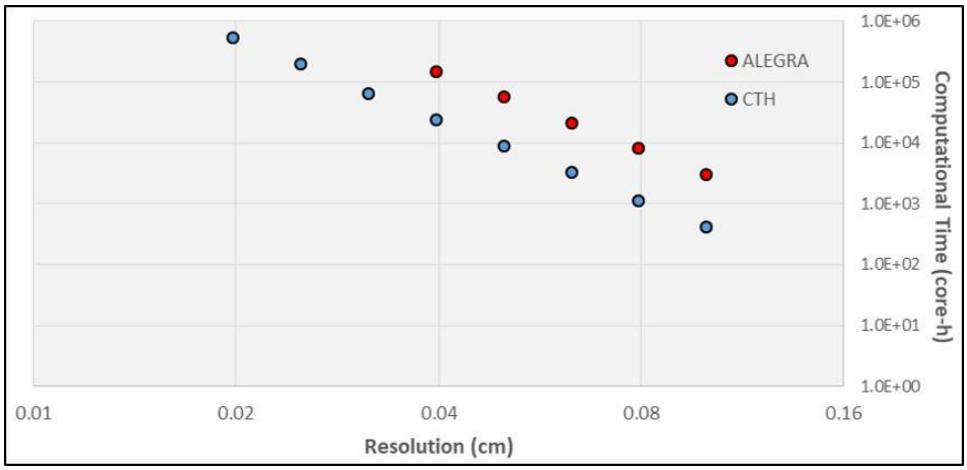


**Fig. 4 Log-log plot of computational time vs. resolution**

DoP results change considerably with resolution (Fig. 5). CTH outcomes initially increase as resolution decreases from 0.1000 cm to 0.0630 cm, then generally decrease over the remainder of the range, excepting a slight uptick at 0.0250 cm. ALEGRA DoP results monotonically increase as resolution decreases over the full range investigated. In neither case is there unambiguous evidence of a limiting value being approached. By coincidence, the DoP values for the two codes happen to be exceptionally close at the baseline resolution of 0.0500 cm chosen for this work.

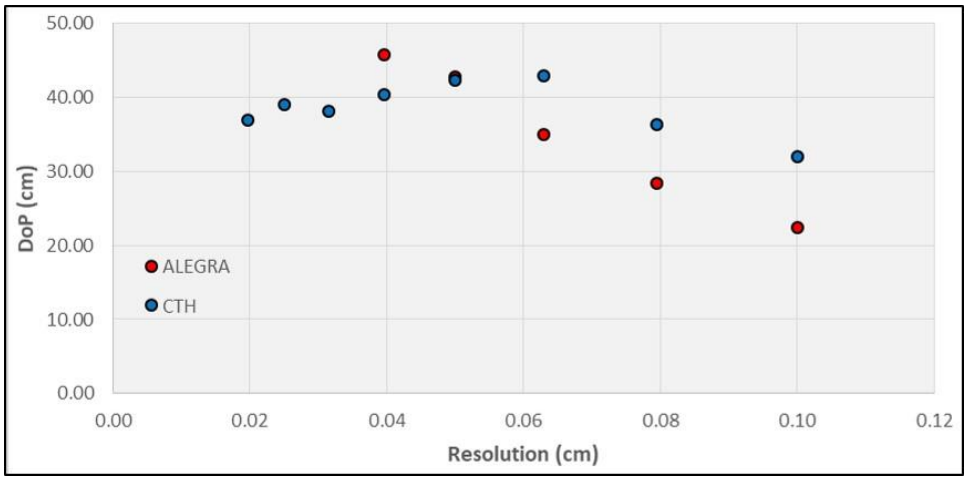


**Fig. 5 Plot of DoP vs. resolution**

## 5.4 Time Step Control

The baseline configuration simulation was run to 500 μs to ensure ample time for jet penetration to achieve a maximal DoP. To investigate the effect of selecting different stop times, simulations were run using stop times ranging from 300 μs to 600 μs in 50-μs increments (Tables 23 and 24).

**Table 23 CTH stop time study results**

| Run | Stop time (μs) | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|
| | | | | Cell data | Images |
| 2 | 300 | 16,087 | 11.02 | 42.050 | 42.1 |
| 3 | 350 | 18,339 | 12.56 | 42.400 | 42.4 |
| 4 | 400 | 20,756 | 14.18 | 42.400 | 42.4 |
| 5 | 450 | 23,309 | 15.87 | 42.400 | 42.4 |
| **BL** | **500** | **25,906** | **17.55** | **42.400** | **42.4** |
| 6 | 550 | 28,503 | 19.29 | 42.350 | 42.4 |
| 7 | 600 | 31,099 | 20.99 | 42.350 | 42.4 |

**Table 24 ALEGRA stop time study results**

| Run | Stop time (μs) | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|
| | | | | Cell data | Images |
| 2 | 300 | 10,236 | 34.57 | 42.800 | 42.7 |
| 3 | 350 | 11,803 | 39.58 | 42.800 | 42.7 |
| 4 | 400 | 13,348 | 44.52 | 42.800 | 42.7 |
| 5 | 450 | 14,872 | 49.45 | 42.800 | 42.7 |
| **BL** | **500** | **16,402** | **55.42** | **42.800** | **42.7** |
| 6 | 550 | 17,945 | 59.39 | 42.750 | 42.7 |
| 7 | 600 | 19,465 | 64.29 | 42.750 | 42.7 |

The penetration process terminates sometime between 300–350 μs in CTH, while the ALEGRA results show maximum penetration by 300 μs. In both codes, a small 1-cell-length decrease in DoP is seen at long times, likely due to the end of the penetration channel crossing a cell boundary as the target drifts through the mesh.

Time steps in both codes are calculated using stability criteria relating to the requirements for solution convergence, as discussed for the cell resolution study. Users can also set a maximum limit on the size of a time step. Simulations were run across range of maximum time steps to investigate effects on outcomes.

The largest time step in CTH simulations of the baseline configuration is less than 25 ns, so no changes in outcome are observed until the maximum time step is set to 20 ns (Table 25). DoP values increase when time steps are limited, with the maximum value of 45.6 cm observed in this study being 22 mm deeper than the baseline configuration.

**Table 25 CTH maximum time step study results**

| Run | Max time step (ns) | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|
| | | | | Cell data | Images |
| **BL** | **None** | **25,906** | **17.55** | **42.400** | **42.4** |
| 2 | 50 | 25,906 | 17.58 | 42.400 | 42.4 |
| 3 | 30 | 25,906 | 17.58 | 42.400 | 42.4 |
| 4 | 25 | 25,906 | 17.57 | 42.400 | 42.4 |
| 5 | 20 | 27,590 | 18.70 | 44.350 | 44.3 |
| 6 | 15 | 34,348 | 23.31 | 43.500 | 43.5 |
| 7 | 10 | 50,498 | 34.28 | 43.800 | 43.8 |
| 8 | 5 | 100,004 | 68.59 | 45.600 | 45.6 |
| 9 | 2.5 | 200,001 | 136.35 | 44.350 | 44.4 |

The maximum time step in ALEGRA simulations of the baseline configuration is between 30–50 ns (Table 26). Variance in DoP values is smaller than in CTH, with outcomes all falling within a 7-mm range.

**Table 26 ALEGRA maximum time step study results**

| Run | Max time step (ns) | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|
| | | | | Cell data | Images |
| **BL** | **None** | **16,402** | **55.42** | **42.800** | **42.7** |
| 2 | 50 | 16,402 | 54.33 | 42.800 | 42.7 |
| 3 | 30 | 17,356 | 57.47 | 42.900 | 42.9 |
| 4 | 25 | 20,302 | 67.16 | 42.650 | 42.6 |
| 5 | 20 | 25,071 | 82.80 | 42.800 | 42.7 |
| 6 | 15 | 33,347 | 110.71 | 42.500 | 42.4 |
| 7 | 10 | 50,001 | 166.22 | 42.300 | 42.2 |

Sometimes random events, like hardware failure, can cause a simulation to fail. Both codes allow for periodic saving of a calculation for resumption at a later time to avoid the need to repeat a failed simulation from the beginning. To examine the effects of restarting an interrupted calculation, simulations were run as a series of restarts in 100 μs increments.

In CTH (Table 27), the solution state at the conclusion of the fourth restart simulation, spanning 400–500 μs of simulated time, does not match the solution of the baseline configuration simulation, with the DoP coming up 3 mm short of the baseline result. This is unexpected, as the restart process is intended to preserve a calculation in a state that can be resumed at a later time as though it had never stopped. Indeed, solution states were unaffected by the restart process in the earlier version of CTH studied in the prior work.

Close examination of the calculation diagnostic output data for the first restart (run 3) shows the first divergence from the baseline output, which is reported in scientific notation with 12-digit precision, is at cycle 7,223, which is 106.096 μs in simulation time. This provides an important clue regarding the issues with building the CTH code discussed in Section 5.2.

The type of slow, growing divergence seen here points to a mismatch in the precision with which data is written to or read from a file, and the precision with which calculations are performed. This creates a tiny, non-zero rounding error in the first restart calculation, as compared to the uninterrupted baseline simulation, that slowly grows with each successive cycle. The first CTH restart begins at cycle 6,940, so about 80 iterations pass before the perturbation is large enough to affect the 12th reported digit in the output file data.

The specifics of how floating-point operations are performed vary across different processor architectures, which is one reason why building the same code on different platforms can lead to variance. A precision mismatch between different parts of the code can result from optimizations performed during the build process. This could also explain the decomposition variability seen in Section 5.3, by affecting the precision with which border cell data is communicated across cores.

**Table 27 CTH restart study results**

| Run | Restart | Stop time (μs) | Cycles | Run time (h) | Cumulative (h) | DoP (cm) | |
|---|---|---|---|---|---|---|---|
| | | | | | | Cell data | Images |
| 2 | 0 | 100 | 6,940 | 4.74 | 4.74 | 22.050 | 22.1 |
| 3 | 1 | 200 | 11,546 | 3.16 | 7.90 | 37.000 | 37.1 |
| 4 | 2 | 300 | 15,777 | 2.87 | 10.77 | 41.850 | 42.0 |
| 5 | 3 | 400 | 20,462 | 3.16 | 13.93 | 42.100 | 42.1 |
| 6 | 4 | 500 | 25,657 | 3.38 | 17.31 | 42.100 | 42.1 |
| **BL** | **. . .** | **500** | **25,906** | **17.55** | **. . .** | **42.400** | **42.4** |

The same study in ALEGRA (Table 28) shows the expected behavior of the restart process, with the solution state of the final restart identical to the baseline configuration result.

**Table 28 ALEGRA restart study results**

| Run | Restart | Stop time (μs) | Cycles | Run time (h) | Cumulative (h) | DoP (cm) | |
|---|---|---|---|---|---|---|---|
| | | | | | | Cell data | Images |
| 2 | 0 | 100 | 3,896 | 13.44 | 13.44 | 22.546 | 22.4 |
| 3 | 1 | 200 | 7,099 | 10.55 | 23.99 | 38.600 | 38.5 |
| 4 | 2 | 300 | 10,236 | 10.13 | 34.12 | 42.800 | 42.7 |
| 5 | 3 | 400 | 13,348 | 10.42 | 44.54 | 42.800 | 42.7 |
| 6 | 4 | 500 | 16,402 | 10.26 | 54.80 | 42.800 | 42.7 |
| **BL** | **. . .** | **500** | **16,402** | **55.42** | **. . .** | **42.800** | **42.7** |

## 5.5 Physical Invariance

The final collection of studies examines how closely simulation codes adhere to several fundamental physical invariances.

Large translations in space were tested by shifting the model system geometry in various directions within the domain. To provide sufficient space for such operations, a domain spanning 620 × 620 × 1827 cells was employed, identical to the largest domain in the void buffer study (run 5 in Tables 15 and 16). Note that all translations tested here equate to shifting the geometry by an integer number of cells. Translations across a noninteger number of cells should include behavior similar to that documented in the origin shift study (Tables 19 and 20).

CTH simulations (Table 29) were run on 512 cores, with run times spanning 44.90 to 58.42 h. The mean DoP over all runs was 42.26 cm, with a standard deviation of 1.51 cm, 3.6% of the mean.

**Table 29 CTH translation study results**

| Run | Translation (cm) | | | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|---|---|
| | X | Y | Z | | | Cell data | Images |
| 1 | 0 | 0 | 0 | 26,912 | 48.80 | 42.450 | 42.4 |
| 2 | 6 | 0 | 0 | 27,181 | 49.40 | 39.800 | 39.8 |
| 3 | –6 | 0 | 0 | 26,991 | 49.13 | 40.850 | 40.9 |
| 4 | 0 | 6 | 0 | 32,424 | 58.42 | 42.800 | 42.8 |
| 5 | 0 | –6 | 0 | 24,834 | 44.90 | 43.250 | 43.3 |
| 6 | 0 | 0 | 6 | 27,070 | 49.15 | 44.750 | 44.7 |
| 7 | 0 | 0 | –6 | 26,621 | 48.28 | 43.450 | 43.4 |
| 8 | 6 | 6 | 0 | 25,909 | 47.18 | 42.550 | 42.5 |
| 9 | 6 | 0 | 6 | 28,183 | 51.34 | 43.250 | 43.3 |
| 10 | 0 | 6 | 6 | 25,074 | 45.45 | 41.550 | 41.5 |
| 11 | 6 | 6 | 6 | 26,332 | 48.08 | 40.200 | 40.2 |

ALEGRA simulations (Table 30) were run on 4096 cores. Run times were highly uniform, with all simulations finishing within 13 min of the mean. DoP values exhibited little variance as well, with an average value of 42.86 cm and a standard deviation of just 0.9 cm, 0.2% of the mean.

**Table 30 ALEGRA translation study results**

| Run | Translation (cm) | | | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|---|---|
| | X | Y | Z | | | Cell data | Images |
| 1 | 0 | 0 | 0 | 16,481 | 41.18 | 42.950 | 42.9 |
| 2 | 6 | 0 | 0 | 16,396 | 40.88 | 42.700 | 42.7 |
| 3 | –6 | 0 | 0 | 16,464 | 41.11 | 42.843 | 42.8 |
| 4 | 0 | 6 | 0 | 16,532 | 41.02 | 42.850 | 42.8 |
| 5 | 0 | –6 | 0 | 16,456 | 41.02 | 42.996 | 42.9 |
| 6 | 0 | 0 | 6 | 16,428 | 41.03 | 42.800 | 42.7 |
| 7 | 0 | 0 | –6 | 16,485 | 41.21 | 42.850 | 42.8 |
| 8 | 6 | 6 | 0 | 16,470 | 40.84 | 42.800 | 42.7 |
| 9 | 6 | 0 | 6 | 16,511 | 41.24 | 42.890 | 42.8 |
| 10 | 0 | 6 | 6 | 16,471 | 40.94 | 42.797 | 42.7 |
| 11 | 6 | 6 | 6 | 16,506 | 41.05 | 42.949 | 42.9 |

Rotation about the symmetry axis of the SCJ device was investigated by rotating the model system about the *z*-axis in 15° increments from 0° to 90°. The expanded 620- × 620- × 1827-cell domain was used to accommodate the rotated geometry.

In CTH (Table 31), the average DoP across all simulations was 44.20 cm, with a standard deviation of 0.92 cm, 2.1% of the mean. Rotating the system by 90° produces a different solution state from the initial system at 0°, which may be due to round off errors in the geometric transformations performed by the code.

**Table 31 CTH axial rotation study results**

| Run | Angle (°) | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|
| | | | | Cell data | Images |
| 1 | 0 | 26912 | 52.33 | 42.450 | 42.4 |
| 2 | 15 | 25359 | 49.89 | 43.550 | 43.5 |
| 3 | 30 | 27339 | 53.29 | 45.000 | 45.0 |
| 4 | 45 | 27285 | 52.97 | 44.650 | 44.6 |
| 5 | 60 | 29326 | 57.23 | 44.800 | 44.8 |
| 6 | 75 | 25749 | 50.69 | 44.800 | 44.8 |
| 7 | 90 | 25586 | 49.82 | 44.150 | 44.2 |

ALEGRA simulations (Table 32) exhibited very little outcome variance, with a standard deviation of only 0.04 cm on a 43.00-cm average DoP. In contrast to CTH, rotating the system by 90° returns the same solution state as the unrotated geometry.

**Table 32 ALEGRA axial rotation study results**

| Run | Angle (°) | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|
| | | | | Cell data | Images |
| 1 | 0 | 16,481 | 44.59 | 42.950 | 42.9 |
| 2 | 15 | 16,474 | 44.94 | 43.000 | 42.9 |
| 3 | 30 | 16,681 | 45.66 | 43.050 | 43.0 |
| 4 | 45 | 16,480 | 45.67 | 43.050 | 43.0 |
| 5 | 60 | 16,391 | 44.62 | 43.000 | 42.9 |
| 6 | 75 | 16,395 | 44.75 | 43.000 | 42.9 |
| 7 | 90 | 16,481 | 44.64 | 42.950 | 42.9 |

More complex rotations were performed by varying the direction of the jet axis within the domain. Custom domain sizes were used to accommodate the geometry while keeping total simulation size manageable.

In both codes there is a marked difference between simulations with the jet oriented along one of the principle coordinate axes (runs 1–4) and those oriented in other directions (runs 5+). In the first case, the average DoP is 42.90 cm in CTH (Table 33) and 43.01 cm in ALEGRA (Table 34). While the close agreement between the codes is a coincidence of resolution (see Fig. 5), the marked change in outcome when the jet axis is oriented away from a coordinate axis is not, with the average DoP dropping to 33.04 cm in CTH and 34.21 cm in ALEGRA. This decrease can be attributed to the effects of material advection, which is generally more accurate for objects moving parallel to a principle coordinate axis than those moving in non-axial directions. This effect has important implications for the structuring of problems involving SCJ devices. Any angular relationship between a jet and a target, including yaw in the SCJ device, should be exclusively applied to the target geometry, with the jet always left oriented parallel to a coordinate axis.

**Table 33 CTH jet axis direction study results**

| Run | Jet axis | | | Cores | Domain cells | | | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | X | Y | Z | | X | Y | Z | | | Cell data | Images |
| 1 | 0 | 0 | 1 | 512 | 380 | 380 | 1,587 | 25,906 | 18.88 | 42.400 | 42.4 |
| 2 | 0 | 0 | –1 | 512 | 380 | 380 | 1,587 | 28,256 | 20.51 | 43.500 | 43.5 |
| 3 | 1 | 0 | 0 | 512 | 1,587 | 380 | 380 | 27,322 | 20.52 | 43.300 | 43.3 |
| 4 | 0 | 1 | 0 | 512 | 380 | 1,587 | 380 | 25,343 | 18.11 | 42.400 | 42.4 |
| 5 | 1 | 0 | 1 | 1,024 | 1,392 | 380 | 1,392 | 25,667 | 30.74 | 31.615 | 31.6 |
| 6 | 0 | 1 | 1 | 1,024 | 380 | 1,392 | 1,392 | 24,576 | 27.88 | 31.333 | 31.3 |
| 7 | 1 | 1 | 0 | 1,024 | 1,392 | 1,392 | 380 | 26,599 | 31.04 | 30.272 | 30.3 |
| 8 | 1 | 1 | 1 | 8,192 | 1,341 | 1,227 | 1,341 | 22,551 | 11.22 | 36.519 | 36.5 |
| 9 | 1 | √π | π | 4,096 | 842 | 1,087 | 1,620 | 30,598 | 18.25 | 35.479 | 35.4 |

**Table 34 ALEGRA jet axis direction study results**

| Run | Jet axis | | | Cores | Domain cells | | | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | X | Y | Z | | X | Y | Z | | | Cell data | Images |
| 1 | 0 | 0 | 1 | 1,024 | 380 | 380 | 1,587 | 16,402 | 59.32 | 42.800 | 42.7 |
| 2 | 0 | 0 | –1 | 1,024 | 380 | 380 | 1,587 | 17,341 | 62.65 | 42.897 | 42.9 |
| 3 | 1 | 0 | 0 | 1,024 | 1,587 | 380 | 380 | 16,499 | 59.63 | 43.348 | 43.2 |
| 4 | 0 | 1 | 0 | 1,024 | 380 | 1,587 | 380 | 16,466 | 59.78 | 43.000 | 42.9 |
| 5 | 1 | 0 | 1 | 4,096 | 1,352 | 360 | 1,352 | 15,952 | 38.68 | 34.196 | 34.2 |
| 6 | 0 | 1 | 1 | 4,096 | 360 | 1,352 | 1,352 | 15,920 | 38.42 | 34.302 | 34.3 |
| 7 | 1 | 1 | 0 | 4,096 | 1,352 | 1,352 | 360 | 15,953 | 38.70 | 34.125 | 34.1 |

In the final study, the choice of reference frame is examined. A basic assumption in all simulations thus far is that objects are initial stationary. A fundamental principle of physics is that outcomes are invariant to uniform motion of the system (for velocities much less than the speed of light, relativistic frames of reference are not implemented in the codes). To test this, simulations were performed in which both the SCJ device and target are initially moving at some fixed velocity along the *z*-axis. The size of the domain along the *z*-axis was expanded to allow sufficient space for objects to move without exiting the domain. CTH simulations were performed on 512 cores, while ALEGRA simulations used 2048 cores. The CTH results (Table 35) show a 7.25-cm range of DoP values, while the ALEGRA results (Table 36) have a range of 1.89 cm.

**Table 35 CTH reference frame study results**

| Run | Velocity (m/s) | Z Cells | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|---|
| | | | | | Cell data | Images |
| 2 | –1,000 | 2,587 | 25,357 | 27.23 | 41.750 | 41.8 |
| 3 | –100 | 1,687 | 23,686 | 17.08 | 39.300 | 39.3 |
| BL | 0 | 1,587 | 25,906 | 17.55 | 42.400 | 42.4 |
| 4 | 100 | 1,687 | 41,076 | 29.12 | 37.050 | 37.1 |
| 5 | 1,000 | 2,587 | 24,154 | 25.08 | 35.150 | 35.1 |

**Table 36 ALEGRA reference frame study results**

| Run | Velocity (m/s) | Z Cells | Cycles | Run time (h) | DoP (cm) | |
|---|---|---|---|---|---|---|
| | | | | | Cell data | Images |
| 2 | –1,000 | 2,587 | 16,461 | 47.56 | 41.500 | 41.4 |
| 3 | –100 | 1,687 | 16,608 | 31.31 | 43.389 | 43.4 |
| BL | 0 | 1,587 | 16,402 | 55.42 | 42.800 | 42.7 |
| 4 | 100 | 1,687 | 16,320 | 31.70 | 42.681 | 42.6 |
| 5 | 1,000 | 2,587 | 16,104 | 43.31 | 41.814 | 41.8 |

# 6. Conclusion

This report examined the influence of computational factors—those parameters relating strictly to the numeric aspect of the calculations and not the physics of the materials being modeled—on variability in simulation outcomes for a model system based on an SCJ device. The CTH and ALEGRA multiphysics codes were used. DoP of the jet into an RHA target block was the primary quantity of interest, with run time and computational size also studied. This investigation was structured identically to a previous study involving a long rod penetrator.

The overall results are summarized in Table 37, which lists the observed variances in DoP due to various computational factors in terms of standard deviation as a percentage of the mean (commonly termed the coefficient of variation). Note that the values shown for stop time only include times beyond which the penetration process has completed.

**Table 37 Coefficients of variation for DoP due to various computational factors in simulations of two model systems using CTH and ALEGRA**

| | Long Rod | | SCJ Device | |
|---|---|---|---|---|
| Factor | CTH 11.1 | ALEGRA 21 May 2015 | CTH 12.0 | ALEGRA 2017.11.06 |
| Cores/node | 0 | 0 | 0 | 0 |
| Total cores | < 0.1% | 0 | 2.8% | 0 |
| Decomposition | < 0.1% | 0 | 3.5% | 0 |
| Void buffer | < 0.1% | < 0.1% | 4.4% | 0.3% |
| Symmetry | < 0.1% | < 0.1% | 3.2% | 0.9% |
| Origin placement | < 0.1% | 0.4% | 3.9% | 1.6% |
| Stop time | < 0.1% | 0.3% | 0.3% | 0.1% |
| Max time step | < 0.1% | < 0.1% | 2.4% | 0.5% |
| Restarts | 0 | 0 | > 0 | 0 |
| Translation | < 0.1% | < 0.1% | 3.6% | 0.2% |
| Axial rotation | < 0.1% | < 0.1% | 2.1% | 0.1% |
| Threat direction | 2.2% | 0.4% | 14.9% | 12.0% |
| Reference frame | 0.3% | 0.4% | 7.9% | 1.8% |

Absent from this table is the effect of domain resolution, which is overtly tied to outcome variance via the computational methods employed by the codes. In properly structured simulations, the approximation errors of the numeric method should become vanishingly small as the resolution goes to zero. Evidence of such limiting behavior was only observed for the long-rod model system in CTH, while results for the SCJ model system in CTH were ambiguous (Fig. 5). In ALEGRA, both systems exhibiting monotonically increasing DoP with decreasing cell size, with no clear limiting behavior over the accessible range of resolutions.

One conclusion of this work is that the inherent variability in simulations involving SCJ devices is appreciably larger than those involving long rod penetrators. This is before considering variability due to material models and parameters, and physical effects like geometric variations in devices and uncertainty in device position and alignment. The variability observed in CTH simulations was roughly 3 times larger than in ALEGRA simulations, but the computational size of problems in ALEGRA is 5–6 times larger.

If one assumes that the SCJ model system behavior can be generalized to include all SCJ simulations, then DoP results can be expected to vary with a standard deviation about 3%–4% of the mean value in CTH, and about 1%–2% of the mean

in ALEGRA. This is the fundamental variation inherent in the numerical methods being applied in these codes, which is separate from any variances associated with uncertainties or variability in physical properties. Furthermore, these variances are for the simplest types of SCJ system; simulations involving complex targets may have greater variances, especially if the interactions between the jet and target are highly sensitive to changes in jet dynamics.

It is one thing to identify and quantify the uncertainty associated with simulations, and another to know what to do with this information. From the perspective of the design and evaluation of protection systems, our goal is to ensure that systems in the field provide the expected level of protection while minimizing undesirable attributes, such as size, weight, and cost. This requires accounting for a wide variety of possible engagement scenarios. Ideally, we would determine all these outcomes experimentally, with multiple repeat testing to account for the variability in the performance of real SCJ warheads. But this would be prohibitively expensive and time consuming, so we turn to numerical simulations to supplement what limited testing we can accomplish. Quantifying the variability in these simulations improves our ability to use computational modeling to design and evaluate the systems that protect our armed forces.

## 7. References


1. Hornbaker DJ. Quantifying uncertainty from computational factors in simulations of a model ballistic system. Aberdeen Proving Ground (MD): Army Ballistic Research Laboratory (US); 2017 Aug. Report No.: ARL-TR-8074.

2. Zukas JA. Introduction to hydrocodes. Boston (MA); Elsevier; 2004.

3. Walters WP. The shaped charge concept, part I. Introduction. Aberdeen Proving Ground (MD): Army Ballistic Research Laboratory (US); 1990 Aug. Report No.: BRL-TR-3142.

4. Schraml S. Simulation of shaped-charge jet formation and penetration using ALE3D. Aberdeen Proving Ground (MD): Army Ballistic Research Laboratory (US); 2016 Aug. Report No.: ARL-TR-7744.

5. Portone TN, Niederhaus JHJ, Sanchez JJ, Swiler LP. Application of Bayesian model selection for metal yield models using ALEGRA and Dakota. Albuquerque (NM): Sandia National Laboratories; 2018 Feb 24. Report No.: SAND2018-1889.

6. Ferman-Coker M. A 2-D axisymmetric model of a stretching jet in ALEGRA. Aberdeen Proving Ground (MD): Army Ballistic Research Laboratory (US); 2004 Sep. Report No.: ARL-TR-3309.

7. Attaway SW, Hills RG, Giunta AA. CTH shaped charge simulations. Albuquerque (NM): Sandia National Laboratories; 2009 July. Report No.: SAND2009-0726.

8. Niederhaus JHJ. Analysis and validation for shaped charge jet simulations using ALEGRA. Albuquerque (NM): Sandia National Laboratories; 2011 Apr. Report No.: SAND2011-2819.

9. Kmetyk LN, Yarrington P, Vigil MG. CTH analyses of Viper conical shaped charges with comparisons to experimental data. Albuquerque (NM): Sandia National Laboratories; 1991 Jan. Report No.: SAND90-2604.

10. Schmitt RG, Crawford DA, Harstad EN, Hensinger DM, Ruggirello K. CTH user's manual and input instructions: CTH development project. Albuquerque (NM): Sandia National Laboratories; 2017 Apr 20.

11. Niederhaus JHJ, Hansen GA, Labreche DA, Love E, Luchini CB, Roberts NV, Robinson AC, Sanchez JJ, Siefert C, Voth TE, et al. ALEGRA user manual. Albuquerque (NM): Sandia National Laboratories; 2017 Dec 20. Report No.: SAND2017-13760.

12. Bell RL, Hertel ES Jr. An improved material interface reconstruction algorithm for Eulerian codes. Albuquerque (NM): Sandia National Laboratories; 1992 Sep. Report No.: SAND92-1716.

13. Benson DJ. Momentum advection on a staggered mesh. Journal of Computational Physics. 1992;100(1):143–162.

14. MIL-H-48358 (AR). HMX/Resin explosive composition LX-14-0 (for use in ammunition). Picatinny Arsenal (NJ): Army Armament Research, Development and Engineering Center (US); 1977 Sep 30 (with Amendment 7 dated 1999 Apr 14).

15. MIL-DTL-12560K. Armor plate, steel, wrought, homogeneous (for use in combat-vehicles and for ammunition testing). Aberdeen Proving Ground (MD): Army Research Laboratory (US); 2013.

16. Schraml SJ, Wagoner VS. Ballistic performance of HY-80 steel relative to rolled homogeneous armor. Aberdeen Proving Ground (MD): Army Ballistic Research Laboratory (US); 2014 July. Report No.: ARL-TR-6983.

17. Segletes SB. An analysis on the stability of the Mie-Grüneisen equation of state for describing the behavior of shock-loaded materials. Aberdeen Proving Ground (MD): Army Ballistic Research Laboratory (US); 1991 Mar. Report No.: BRL-TR-3214.

18. Johnson GR, Cook WH. A constitutive model and data for metals subjected to large strains, high strain rates and high temperatures. Proceedings of the 7th International Symposium on Ballistics. 1983;541–547.

19. Johnson GR, Cook WH. Fracture characteristics of three metals subjected to various strains, strain rates, temperatures and pressures. Engineering Fracture Mechanics. 1985;21(1):31–48.

20. Johnson GR, Holmquist TJ. Test data and computational strength and fracture model constants for 23 materials subjected to large strains, high strain rates, and high temperatures. Los Alamos (NM): Los Alamos National Laboratories; 1989 Jan. Report No.: LA-11463-MS.

21. Group GMX-6. Selected Hugoniots. Los Alamos (NM): Los Alamos National Laboratories; 1969 May. Report No.: LA-4167-MS.

22. Thermal and physical properties of pure metals. In: Haynes WM, editor. CRC handbook of chemistry and physics, 97th ed. (Internet version 2017). Boca Raton (FL): CRC Press/Taylor & Francis; 2017.

23. Electrolytic tough pitch copper, UNS C11000, H00 temper flat products. MATWEB material property data [accessed 2018 Jul 27]. www.matweb.com/search/datasheet.aspx?matguid=3c78d450e90f48c48d68e2d17a8e51f7.

24. Dobratz BM, Crawford PC. LLNL explosives handbook. Livermore (CA): Lawrence Livermore National Laboratory; 1985 Jan. Report No.: UCRL-52997-Chg.2.

25. Hertel ES Jr, Kerley GI. CTH reference manual: The equation of state package. Albuquerque (NM): Sandia National Laboratories; 1998 Apr. Report No.: SAND98-0947.

26. Lee EL, Hornig HC, Kury JW. Adiabatic expansion of high explosive detonation products. Livermore (CA): Lawrence Livermore National Laboratory; 1968 May. Report No.: UCRL-50422.

27. Segletes, SB. An examination of the JWL equation of state. Aberdeen Proving Ground (MD): Army Ballistic Research Laboratory (US); 2018 July. Report No.: ARL-TR-8403.

28. Crawford DA, Schmitt RG, Butler RJ. Spymaster user's guide, version 6.0. Albuquerque (NM): Sandia National Laboratories; 2016 June.

29. Ayachit U. The ParaView guide, community edition. Clifton Park (NY): Kitware Inc.; 2016 June 17.

## List of Symbols, Abbreviations, and Acronyms

| | |
|---|---|
| 3-D | 3-dimensional |
| CD | charge diameter |
| CJ | Chapman-Jouguet |
| DoP | depth of penetration |
| EOS | equation of state |
| JWL | Jones-Wilkins-Lee |
| RHA | rolled homogenous armor |
| SCJ | shaped charge jet |

| | |
|---|---|
| 1 (PDF) | DEFENSE TECHNICAL INFORMATION CTR DTIC OCA |
| 2 (PDF) | DIR ARL IMAL HRA RECORDS MGMT RDRL DCL TECH LIB |
| 1 (PDF) | GOVT PRINTG OFC A MALHOTRA |
| 1 (PDF) | ARL RDRL WMP E D HORNBAKER |