Center for Radiative Shock Hydrodynamics

Final Report DOE Cooperative Agreement Number DE-FC52-08NA28616
Date: April 15, 2014
Period Covered: April 15, 2008 through April 14, 2014
Project Director: R Paul Drake,
Atmospheric, Oceanic and Space Sciences, University of Michigan

Co-Principal Investigators:
Marvin L. Adams, Nuclear Engineering, Texas A& M University
James P. Holloway, Nuclear Eng. and Radiological Sciences, University of Michigan
Kenneth G. Powell, Aerospace Engineering, University of Michigan
Quentin F. Stout, Computer Science and Engineering, University of Michigan

Co-Investigators, University of Michigan:
Natasha Andronova, Atmospheric, Oceanic and Space Sciences
Krzysztof J. Fidkowski, Aerospace Engineering
Bruce Fryxell, Atmospheric, Oceanic and Space Sciences
Tamas I. Gombosi, Atmospheric, Oceanic and Space Sciences
Eric Johnsen, Mechanical Engineering and Applied Mechanics
Smadar Karni, Mathematics
Carolyn C. Kuranz, Atmospheric, Oceanic and Space Sciences
Edward W. Larsen, Nuclear Engineering and Radiological Sciences
William R. Martin, Nuclear Engineering and Radiological Sciences
Eric Myra, Atmospheric, Oceanic and Space Sciences
Vijayan Nair, Statistics
Philip I. Roe, Aerospace Engineering
Igor Sokolov, Atmospheric, Oceanic and Space Sciences
Katsuyo Thornton, Materials Science and Engineering
Gabor Toth, Atmospheric, Oceanic and Space Sciences
Bartholomeus van der Holst, Atmospheric, Oceanic and Space Sciences
Bram van Leer, Aerospace Engineering
Co-Investigators, Texas A & M University:
Nancy Amato, Computer Science
Bani Mallick, Statistics
Ryan G. McClarren, Nuclear Engineering
James E. Morel, Nuclear Engineering
Lawrence Rauchwerger, Computer Science
Co-Investigator, Simon Frazier University:
Derek Bingham, Statistics

Executive Summary

The goal of this report is to provide a relatively complete record of the work done by the Center for Radiative Shock Hydrodynamics (CRASH) during its five-plus years of existence. Our intent is as follows. First, to make available a complete set of experimental data, which may prove useful for future predictive science studies. Second, to record enough detail of the CRASH code that the code itself or some of its methods may prove useful to future simulators. Third, to present our predictive studies, which included several methodological developments within the general context of Kennedy-O'Hagan-type models and also included facing quite a few challenges in developing the computer runs and metrics for use in such an analysis. We believe that this too may be useful to future researchers.

The specific focus of the project was radiative shocks, which develop when shock waves become so fast and hot that the radiation from the shocked matter dominates the energy transport. This in turn leads to changes in the shock structure. Radiative shocks are challenging to simulate, as they include phenomena on a range of spatial and temporal scales and involve two types of nonlinear physics - hydrodynamics and radiation transport. Even so, the range of physics involved is narrow enough that one can hope to model all of it with sufficient fidelity to reproduce the data.

CRASH was focused on developing predictions for a sequence of experiments performed in Project Year 5, in which those experiments represented an extrapolation from all previously available data. The previous data involved driving radiative shocks within cylindrical structures, and mainly straight tubes. The Year 5 experiments drove a radiative shock down an elliptical tube. Our long-stated goal for these predictions was that the distribution of predicted values would overlap significantly with the observed distribution. We achieved this goal.

Achieving our goal required the conversion of an established space-weather code to model radiative shocks at high energy density. To obtain reasonable fidelity with respect to the experimental data required implementing a laser absorption package, in addition to a hydrodynamic solver, electron physics and heat conduction, and multigroup diffusive radiation transport. The dedicated experiments provided evidence of experimental variability, validation of the calculation of initial shock wave behavior, and validation data at many observation times using cylindrical shock tubes. Following this were preparatory experiments for and finally the execution of the Year 5 experiments. The predictive science research included a wide range of sensitivity studies to determine which variables were important and a sequence of predictive studies focused on specific issues and sets of data. This led ultimately to predictions of shock location for the Year 5 experiments.

A conclusion from this project is that the serious quantification of uncertainty in simulations is a dauntingly difficult and expensive prospect. Pre-existing codes are unlikely to have been built with attention to what will be needed to quantify their uncertainty. Pre-existing experimental results are even more unlikely to include a sufficiently detailed analysis of the experimental uncertainties. And this will also be true of most experiments that might be used to validate components of the simulation. The analysis of uncertainty in any one of the physical processes (and related physical constants) is a major effort. And addressing model form uncertainty is an even bigger challenge, that may in principle require development of complete, alternative simulation models. We made a start at all of this, and completed almost none of it. But by the end of a project, we finally had all the pieces in place and working that would have enabled a range of important studies and advances in relatively near-term years. But the sponsor terminated the program after only five years. For most of the participants this was a relatively minor development, although for a few of them it proved to be enormously disruptive. We believe that the cost to the nation, in work that was ready be done but now will not be, was much much larger.

Contents

1 Introductory Overview

Physics modeling, realized in complex computer codes, is often used to forecast what would be observed in reality should a physical experiment be conducted. However, such forecasts should include a quantitative understanding of the significant sources of uncertainty affecting the output, including both modeling uncertainty and uncertainty due to parameters in the model. This uncertainty can be understood through quantitative estimates of the sensitivity of the output to variations in the input, and through the construction of a predictive probability distribution of outputs, ideally conditioned upon experimental data.

The goal of predictive modeling can be to provide decision makers with quantitative estimates of uncertainty to factor into decisions, or it can be to provide physicists and modelers with information to help assess and improve physics understanding and the quality of models. It is this second purpose that was the foundation of the Center for Radiative Shock Hydrodynamics, the CRASH center. Our overarching project goal was to develop a simulator—the CRASH code—that can predict in an unexplored region of the experimental input space after being assessed in a different region of input space that has been simultaneously explored by experiments and simulations. That is, we wish to use quantitative measures of sensitivity and uncertainty to guide decisions about experiments and modeling improvements, and ultimately demonstrate that this program has resulted in improvements in predictive power in the unexplored region of the input space. In the short span of this project, our computational sensitivity studies definitely affected our understanding of the system of interest and helped guide the development of improved computational models. But only by the end of the project had we become ready to use further predictive studies to guide our choices in ways that might have produced significant improvements in predictive capability beyond those that would have come naturally to any thoughtful developer of simulation models.

The specific focus of the project was radiative shocks, which develop when shock waves become so fast and hot that the radiation from the shocked matter dominates the energy transport. This in turn leads to changes in the shock structure. Radiative shocks are challenging to simulate, as they include phenomena on a range of spatial and temporal scales and involve two types of nonlinear physics - hydrodynamics and radiation transport. Even so, the range of physics involved is narrow enough that one can seek to model all of it with sufficient fidelity to reproduce the data.

The CRASH project builds upon the basic physical system shown in Figure 1. Ten (0.35 μm wavelength) laser beams from the Omega laser [Boehly et al., 1995] are incident on a 20-μm thick Be disk, at an irradiance of $\sim 7 \times 10^{14}$ W/cm^2 for 1 ns. This shocks the Be and then accelerates the resulting plasma to $>$ 100 km/s. The leading edge of this plasma drives a shock into Xe gas at 1.1 atm pressure with an initial velocity of $\sim$ 200 km/s. This produces the observable structures shown schematically in Figure 1(b) and by a radiograph in Figure 1(c). The radiation from the shocked Xe preheats the unshocked Xe. It also ablates the shock-tube wall, producing a wall shock that drives the Xe gas inward. Where this wall shock meets the primary shock, the shock-shock interaction produces a noticeable deflection of the dense Xe flow (dark in the radiograph). The Xe that flows through both the wall shock and the oblique portion of the primary shock ends up with higher velocity and forms the material described as entrained Xe. On a finer scale than is seen in the radiograph, the shocked Xe ions, which are initially heated to hundreds of eV, cool rapidly as they heat the electrons, and the heated electrons further ionize the Xe and radiate most of their energy away. In response, the shocked Xe layer, which is optically very thick, becomes several times denser. The resulting final temperature in the shocked matter and characteristic radiation temperature is about 40 eV. In contrast, the radiation mean free path in the unshocked Xe is much longer and the radiation transport is not diffusive there. We have published

several papers describing this system and its experimental variability [Doss et al., 2010, 2011a,b, Drake et al., 2011].

The radiograph in Figure 1 shows fundamentally where the dense Xe is. Our goal was to predict the shock location and metrics measuring wall-shock properties. The evidence for these predictions comes from sets of simulations using the CRASH code, described below and in publications [van der Holst et al., 2011, 2012, 2013]. Briefly, CRASH is a 3D radiation-hydrodynamic code employing a Godunov-type hydrodynamic solver, flux-limited multigroup diffusion for radiation transport, flux-limited electron heat transport and the necessary related electron physics, a 3D laser-energy deposition package, and tabular treatment of equations of state and opacities. In the present report, we discuss the path of the project, the experiments, the simulations and several related issues, and our predictive science studies. The publications produced during this project are cited in the appropriate sections of the report and are included as a final appendix.

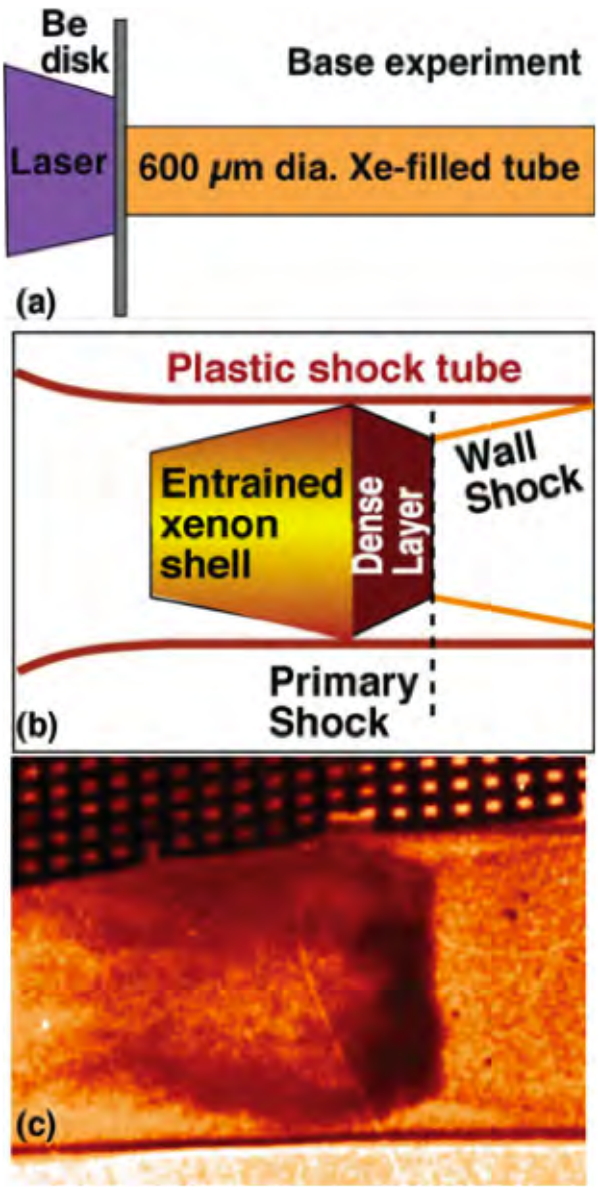

Figure 1: (a) Schematic of a radiative shock experiment. (b) Schematic of features in radiograph. (c) Radiograph. The structure in the dense Xe may be due to a Vishniac-type instability.

1.1 The path of the project

The CRASH project had a mandate to do predictive science throughout, while also doing complex simulations of a target physical system, and providing data to engage the predictive models. We began with a widely used simulation code developed for space-weather applications, BATSRUS [Toth et al., 2005]. This code included a wide variety of MHD models but did not have models for unmagnetized hydrodynamics, radiation, the separate treatement of electrons and their heat transport, or the deposition of laser energy. As our result, our first major priority was to add enough capability to this code to model our physical system of interest with reasonable fidelity. This turned out to take 3.5 years, and dominated the overall manpower expenditures for the project. We call the code run with the features added by our project the CRASH code, although in reality it is BATSRUS run with use of a certain set of subroutines. In order to succeed with BATSRUS, we already had implemented solid

software-verification practices. We continued these practices, and include a discussion of some of the related tests within the report. More detail was provided in our earlier reports on the project.

The CRASH experiments obtained a sequence of data chosen in each year to best advance our ability to conduct predictive studies, and then to provide the data for the novel system that was our year-5 experiment. This turned out to imply that all of them focused on measuring different aspects of radiative shock behavior, in support of assessing and improving code fidelity and later in support of assuring that we could perform the final experiment. Section 2 below discusses their sequence and results. Control of the experiments was essential to obtain the data needed by the project, and also to be able to analyze uncertainties with the necessary depth. Fortunately, the experiments for CRASH were strongly synergistic with the activities of our Center for Laser Experimental Astrophysics Research, and so required only a modest investment of CRASH manpower and funding.

Within a year of the project start date, we completed and reported (in our annual report) simulations enabled by the inclusion of hydrodynamic features and radiation transport by gray diffusion in three dimensions. By 1.5 years into the project we reported results of simulation run sets using CRASH 1.1, which had these features. We released CRASH 2.0 just after two years into the project. It included multigroup radiation transport and electron physics including electron heat transport by flux-limited diffusion. By 2.5 years into the project we reported results of simulation run sets using CRASH 2.1. Our modeling then included all the physics needed to expect to obtain simulations with reasonable fidelity. But we still did not have results with good fidelity to the experimental data. We put a great deal of effort into the pursuit of various hypotheses about why, but in the end the key aspect turned out to be rooted in our approach to simulating the laser-energy deposition.

With the intent of being able to do predictive science studies as early as possible, we had chosen to use an existing Lagrangian code, H2D, to calculate the laser en-

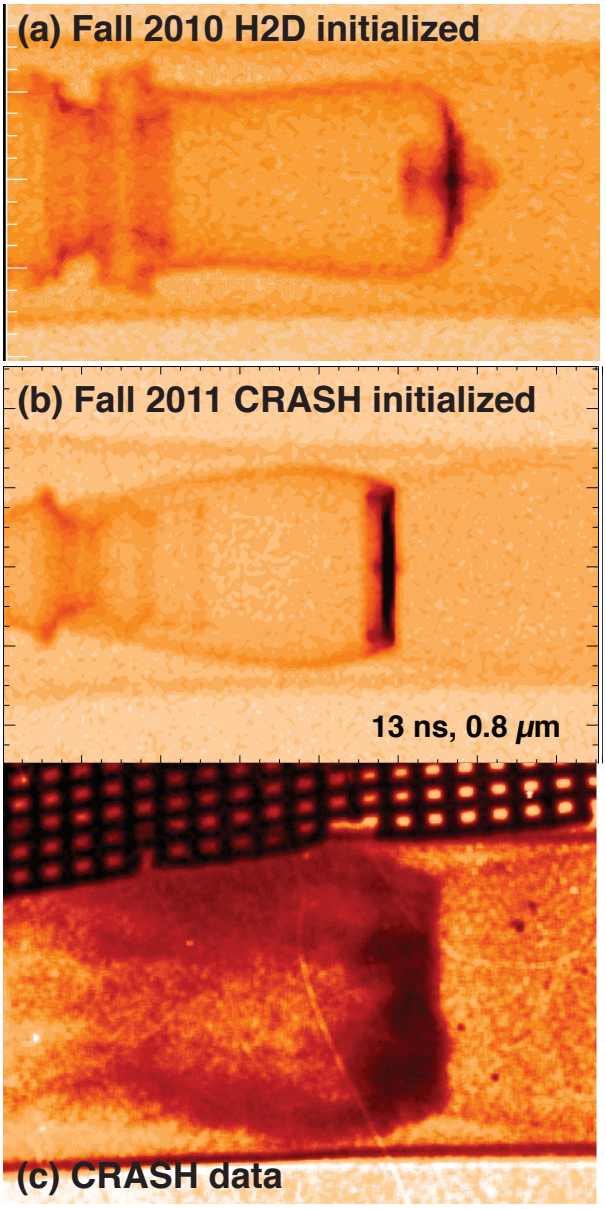

Figure 2: Simulations and data at 13 ns for the base CRASH experiment of Figure 1. Both simulations have 0.8 μm effective resolution with 2 levels of AMR. (a) H2D initialized (b) CRASH laser-package initialized (c) Experimental radiograph, in which the layer is probably thickened somewhat by tilt.

ergy deposition. H2D is a two-dimensional radiation hydrodynamics code that includes the ability to trace laser rays in three dimensions (important for accuracy) along with other typical capabilities needed to model laser-plasma experiments. Using H2D did in fact enable us to accomplish a key predictive study during the middle years of the project, and also enabled us to develop some novel methodology for predictive science, relating to analyzing systems involving coupled models. (See Sec. 6.8.)

By two years into the project it was clear that our combined simulations were not producing results with good fidelity, and we made the decision to develop a "laser package" for CRASH. At that point our completion of the previous software development had freed up the required manpower. By 3.5 years into the project we had tagged CRASH 3.0 and reported results showing that we now had simulations with good fidelity. Figure 2 illustrates this. It is worthwhile to mention that we think H2D is a very good code for applications to which its Lagrangian approach is well suited. This did not prove to include our experimental system, which generates very large vorticity. We discuss this issue further in Sec. 3.5.3.1. It is also worth mentioning that attempts during this project to simulate our experimental system with RAGE at LANL and FLASH at Chicago produced initial results very similar to what we got, and that previous attempts by a graduate student to simulate it using HYDRA at LLNL did not succeed. No doubt an expert user with any of these codes could tune them to produce better outcomes, but one point of pursuing prective science is to stimulate code development that gets good results without expert tuning. Our experimental system is more well-suited to challenging codes in this way than we had imagined, and could be used for further such work in the future.

Our simulation software work in the final phase of the proejct evolved toward making sequence of important though minor improvements, the most significant of which are discussed later in this report. In addition, we undertook several other studies, all discussed below in the report. We did a validation study of the hydrodynamic component. We collaborated in a code-to-code comparison study. We did a study of solution verification in the context of our model. We considered two aspects of model form uncertainty. One of these was an assessment of the importance of non-LTE effects. The other, more extensive effort involved the use of TAMU's PDT code, converted to solve for the transport of thermal radiation, to examine errors resulting from the diffusive model used in CRASH. This was an example of the kind of effort mentioned in the Executive Summary, in which one may need to develop a separate and sophisticated separate code in order to assess model form uncertainty for some aspect of a some workhorse multiphysics simulation code.

One aspect that deserves further comment is our continued work relating to opacities. We took an approach for the project that provided an avenue to analyze uncertainties in opacities and equations of state. The CRASH calculations evaluate these properties using data regarding the atomic levels present in the material to drive calculations based on fundamental statistical physics. This enables one to consider uncertainties in knowledge of the atomic levels as a basis to obtain uncertainties in these properties. One could not properly account for such uncertainties by analyzing only the data in an opacity table. One of our students (at TAMU) completed a demonstration study showing how such an approach could work. Direct comparison of our opacities and equations of state for the low-Z CRASH materials with established results shows that the CRASH model produces results consistent with them. In the case of the Xe opacity, however, the available atomic data provides a description of the actual structure that is so incomplete that the method used by CRASH produces substantially inaccurate opacities. We thus ended up with a conflict between the need for computational fidelity and the goal of evaluating uncertainty. There is some potential to develop statistical methods of analyzing opacity tables to formulate uncertainties, but pursuing this is beyond the scope of what we could undertake. For the purpose of most accurately predicting the year-5 experiment, computational fidelity won out, and we used the

STA Xe opacities produced by some of our team members.

Throughout the project we conducted predictive science studies. These began with sensitivity studies using the one-dimensional version of H2D, Hyades, and followed with one-dimensional studies that combined Hyades and CRASH. We found throughout that the major challenge to doing predictive studies using a complex-multiphysics code (or codes) is to accomplish the sets of simulation runs (run sets) that provide the necessary data. One is forced by computational (and human) limitations to choose a reduced set of variables. Here variables includes not only the physical aspects that define the laboratory experiment of interest, but also the physical parameters whose values are imperfectly known (such as opacities) and the various numerical parameters that control the simulation (grid resolution being a trivial example). This need for a reduced set of variables drives sensitivity studies. One then must define a probability distribution and range to sample for each of the variables selected for the run set, and also specify all other aspects of the simulation. Poor choices here, which may reflect ignorance rather than errors, can lead to a need to redo many simulations. In our case this happened when we had to stop Run Set 11 to redefine the simulation domain. Accomplishing the run set itself is not trivial either, as the large-scale compuational platforms are not stable in time. The second-largest manpower investment for this project was devoted to the run sets and related activities.

These related activities included the extraction of metrics for the statistical portion of the predictive studies. Some of our early activity included seeking to find ways to compare entire radiographic images pixel by pixel, a notion that has some popularity at the labs. That experience showed that the initial simulation result must be extraordinarily close to that of the laboratory experiment if this approach is to prove effective. For realistic predictive studies of complex multiphysics systems, which involve variations of the simulation inputs corresponding to real uncertainties (as opposed to tuning the result using knobs in the code), this now strikes at least some of us as unlikely to succeed. We then turned, during the middle years of our project, to using metrics that were some approximation to taking moments of the data. This might prove to be very effective for simulations with enough fidelity to make several moments useful, and with a team that has time to really learn what they mean. In our case, working with a small number of such metrics was the best we could do while the fidelity of the simulation output is poor. We describe these metrics and their extraction in Sec. 6.2.1.

Throughout the project, shock location proved to be a metric we could extract with well-defined methods, and this formed the basis for most of the statistical predictive studies. Two key aspects of these studies, from a methodological point of view, is that they involved the specification of methods for assessing uncertainty in coupled models and in using models of varying fidelity. These developments and the associated studies are discussed in Section 6. In two cases we withheld all information about an experimental result until the first version of a certain predictive study was complete. The first of these was the extrapolation of shock location to late time (Sec. 6.8), with the result revealed during our third-year review. They agreed, but the uncertainty assigned to the shock location *was* large. The second was the prediction of shock location for the year-5 experiment with the elliptical tube (Sec. 6.10). Here the predicted range of shock locations overlapped with those observed within the experimental variability, which had been our stated goal all along. A cultural note that may be of interest is that the physicists on our team were satisfied and even impressed by this result, while the statisticians were not.

Later on in the project, as our computational fidelity improved, we finally developed robust ways to extract additional metrics that were specific to the structure of the system of interest, as is described in Sec. 6.2.2. It is notable that these results finally provided some evidence of effects one could attribute to the limitations of the physics in the CRASH model. In a system like ours with significant directionality, radiation transport based on diffusion should

transport too much energy laterally. This ought to lead to excessive energy deposited in the walls of our shock tube, so one might expect the wall shock to be further from the wall in the simulations than in the laboratory experiment. We saw some evidence of this using these wall-shock metrics. Also supporting the notion that there remained too much blowoff of wall material is the residual development of axial structure beyond that seen in the data, which (based on studying the simulated results) seems likely to have affected our fidelity if we had extended the studies even later in time. On the other hand, our elliptical tubes were selected because they create some variation in radiation energy deposition with azimuthal angle, and we did not find a resolvable difference in wall shock properties between views of the major or minor axis. When the project ended, we were within a few months of having a predicitve analysis based on the wall shock metrics. For the reasons just discussed, this would have been telling, as would have been further experiments responding to the analysis. Such activities still would be telling.

2 The CRASH Experiments and Data

The selection of experiments for CRASH was driven by the needs of predictive science. These experiments came in the form of one to two "shot days" per year at the Omega laser facililty. It was easy to decide what to do first – a full day of repetition of our basic shock system, in order to assess its experimental vaiability. These data became a primary focus for much of our predictive work, which evolved as did our simulation tools. In the next three years the crucial need was to improve the simulation code until it could model the experiments with sufficient fidelity. This led us to choose experiments that would help evaluate the dynamics in the code, primarily by making measurements at different times. Using various experimental techniques we have obtained data from about 0.5 ns - 30 ns. By the middle of project year 4 our code did produce simulations of reasonable fidelity, and our attention turned to the long-planned variations in the target geometry. We proceeded to do experiments with cylindrical nozzles and eventually with nozzles that directed the shock into an elliptical tube. This inherently three-dimensionial experiment, which forced our predictions to extrapolate beyond systems we had previously studied, provided the key data for the final predictive study. Figure 3 illustrates the experimental sequence of the project.

In the following sections we discuss the experiments, providing details on the experimental parameters and the measured shock locations. We also were able to obtain, in many cases, the location of the "triple point", where the wall shock driven by expansion of the tube wall intersects the primary shock. These results are discussed and tabulated below in Sec. 6.2.2.

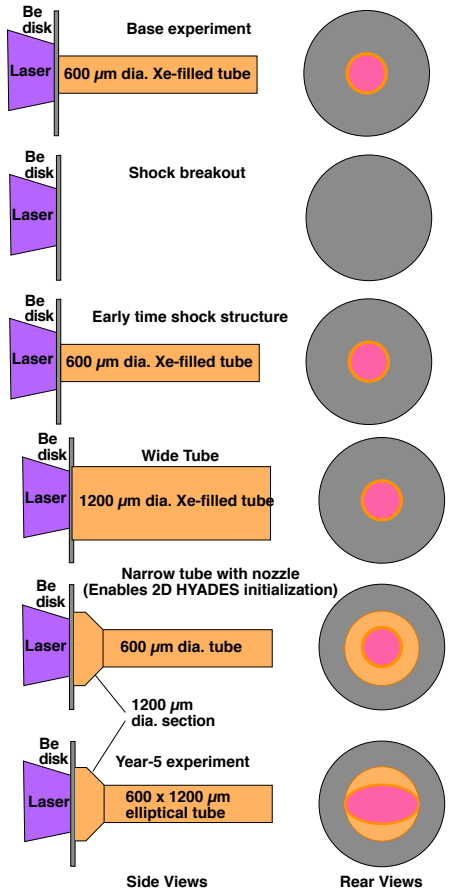

Figure 3: Diagrams of the sequence of targets used in the CRASH experiments. A specific discussion of the reasons for each experiment is provided later in this section.

2.1 Initial CRASH experiment

A 20 μm Be disk was irradiated with ten Omega laser beams that were smoothed with a Distributed Phase Plate (DPP), which creates a spatial profile of a super-Gaussian with an exponent of 4.5 and a laser spot size of about 820 μm FWHM. The laser beams are further smoothed by the temporal beam-smoothing technique Smoothing by Spectral Dispersion (SSD) to produce moving speckles in the beam spot about 5 μm in size. The requested total laser energy was 3.8 kJ. The ablation pressure produced by the laser beams launches a shock into the Be disk. Upon exiting the Be disk the shock moves into Xe gas at nominally 1.1 atm and at a density of 6.5 mg cm^{-3}. The gas is enclosed in a polyimide tube that has an outer diameter of 625 μm with 25 μm thick walls. The

Figure 4: Left: As built basic CRASH target. Right: dimensions of the physically significant components.

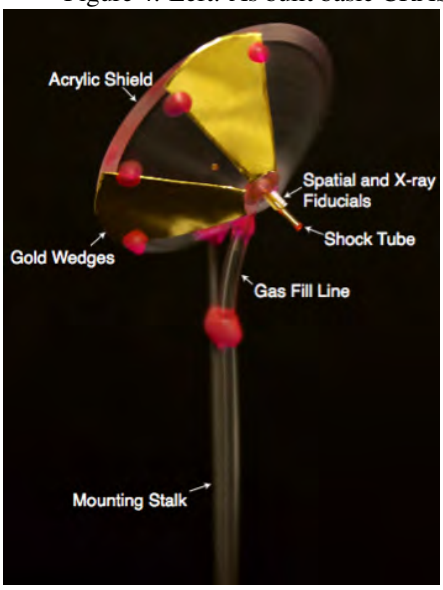

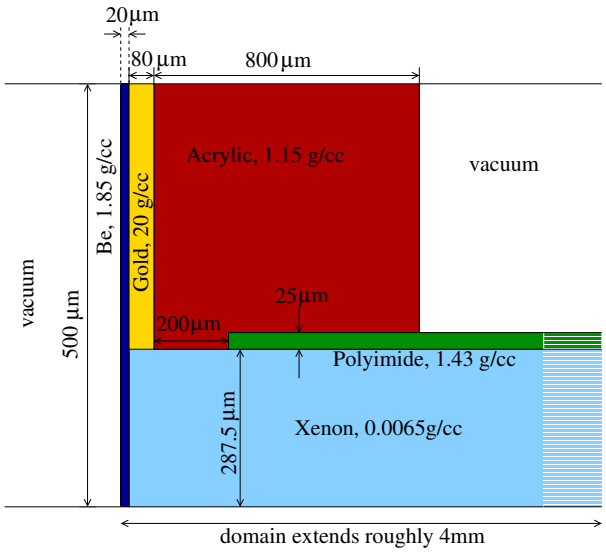

shock accelerates into the lower density gas at velocities of over 100 μm ns^{-1}. It is moving fast enough that radiative fluxes become important in this system.

The first CRASH experimental campaign comprised 11 shots (Omega shots 52660 52671). The targets were all nominally identical replicas of a single design. Figure 4 shows a photograph and the dimensions. The experiments were intended to quantify aspects of experimental repeatability related to target manufacturing and experimental execution. A target consists of a 20 +/- 1 m Be disk mounted on a 625 m ID polyimide tube having 25 μm-thick walls. The tube and disk were mounted into the center of an acrylic shield, approximately 20 mm in diameter and fitted with a 50 m bore hole to accept gas fill into the tube. Two gold wedges were attached to the shield to protect the primary diagnostics of the experiment from x-rays produced during the driving event; this prevents the x-ray film on which data are collected from being prematurely exposed. Preparing for shots, the tube is filled with Xe gas to a pressure of 1.1 atm, inserted into the Omega chamber, and then driven by 10 laser beams delivering 3.8 kJ in a 1 ns pulse to the Be disk surface. Also attached to the shock tubes were a gold grid as a spatial fiducial and an Al strip with stepped density increments as a fiducial for calibrating the diagnostic x-ray emission and assessing whether the x-ray spectrum is as expected.

Primary diagnosis was accomplished by the construction of backlit pinhole targets (shown in Figure 5). Backlighter targets consist of an acrylic frame designed to irradiate the target with x-rays from two emitters, collimated by pinholes. The two x-ray beams are intended to be orthogonal to the target axis and to each other, and to be centered at a point in the shock tube 2 mm from the drive disk. The x-rays were emitted from Va foils stood off from the rear of each pinhole substrate. When each Va foil is illuminated by an additional 5 beams, delivering 70 J each in a 200 ps pulse, it emits 5.2 keV x-rays, suitable for imaging the shocks. The x-rays transmitted through the target were recorded on ungated x-ray film. Image plates were also used as supporting diagnostics, placed behind the film. These captured additional photons to create a second image of each view of each shot. While not of as high a quality data as the data obtained by film, the image plate data had a faster development time and was enabled rapid, well-informed decisions

Figure 5: As built CRASH two-axis backlighter.

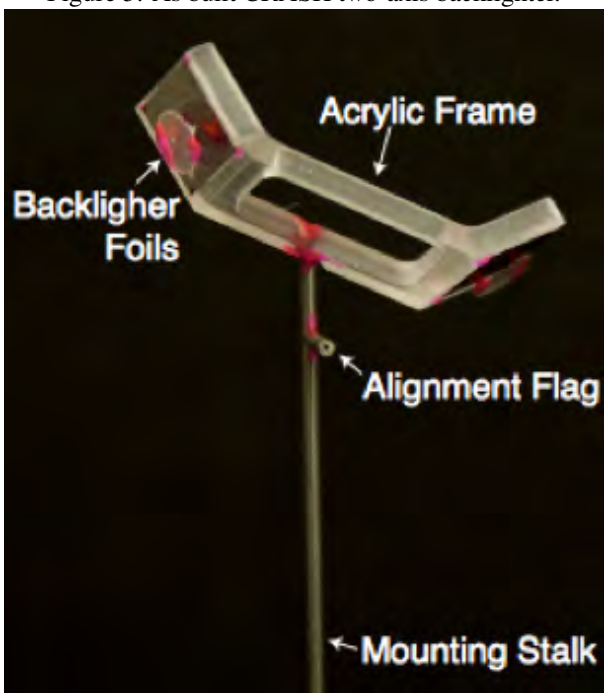

through the day.

Fifteen targets were constructed, of which the 11 that most closely matched the construction specifications were used in the campaign. Backlighter illumination times were varied through the day. First, they were varied in order to center the shock in the data range. After the third shot, they were varied only in some instances to obtain velocity information. Four shots were timed with 1 ns spacing between the two views, and four shots were taken with both views simultaneous. In all of these last eight shots, at least one view was timed for 13 ns after the drive event, in order to obtain the largest set of equivalently timed, well-centered images possible. Figure 6 shows the 16 total radiographic images collected during the campaign. This includes a damaged image from the first shot, which was compromised by a film-loading error, and two images from the sixth shot, which was misaligned in the chamber. Four shots exhibited a failure mode of unknown mechanism in which one view produced data and the other produced only blank film. This could have been a difficulty of aligning the single backlighter frame adequately to properly position both of its two backligher sources.

Figure 7 shows two photographs of the film recording both views of Shot 52667. These are simultaneous, orthogonal views, oriented so that the shock is traveling left to right in each case. The radiographic images show the shock tube and radiating shock. Above the tube is the x-ray fiducial. Below the tube is the gold grid used as the spatial fiducial. The shot numbers and input data from this first set of experiments are shown in Table 1, and the measured shock locations and shocked Xe layer thickness are shown in Table 2. Additional experiment parameters and data are shown in Table 3

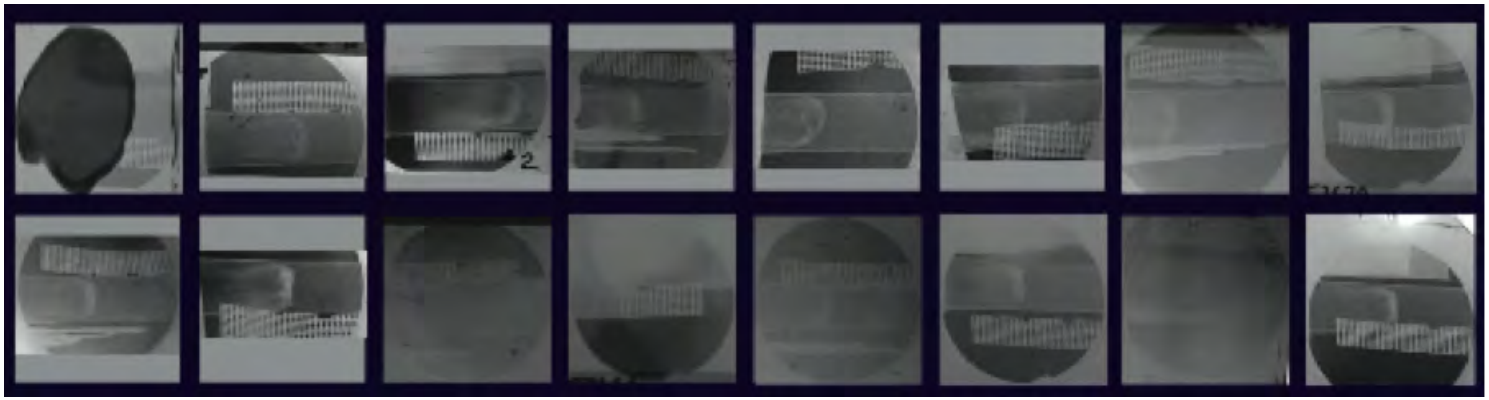

Figure 6: Thumbnails of data collected, obtained by using a digital camera to image the x-ray film on a light table. We are at present awaiting digitization of the film.

Figure 7: Two views of the shock in shot 52667

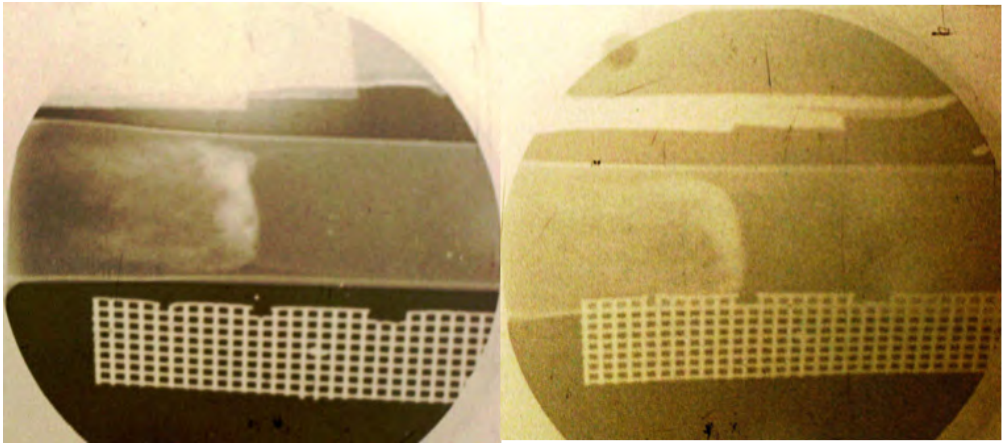

Shot number	Be Thickness (μm)	Laser Energy (J)	Xe Pressure (atm)	t_o (ns)
52661	21	3889.6	1.133	14
52661*	21	3889.6	1.133	16
52663	21	3882	1.167	13
52665	21	3892.4	1.133	13
52665*	21	3892.4	1.133	14
52667	21	3880.2	1.204	13
52667*	21	3880.2	1.204	13
52668	21	3859.8	1.105	13
52668*	21	3859.8	1.105	14
52669	21	3846	1.169	13
52670	21	3841.5	1.17	14
52671	21	3867.4	1.169	13

Table 1: Omega shot numbers and experimentally controlled variables. Note that two views are available for some shots, yielding two observation times for some shots, and one repeated observation time (the two views were taken simultaneously). From a predictive perspective, these 8 shots yield 12 measurements. View 2 data are marked with asterisks (*). The drive disk thickness is uncertain to $\pm 1\mu$m.

Shot number	Shock Position (μm)	Xe layer width (μm)
52661	2308 ± 100	125
52661*	2485 ± 70	117
52663	2030 ± 50	122
52665	2042 ± 25	135
52665*	2178 ± 75	115
52667	2085 ± 25	167
52667*	2077 ± 25	136
52668	2098 ± 65	103
52668*	2310 ± 75	137
52669	1940 ± 75	150
52670	2038 ± 275	137
52671	1943 ± 300	121

Table 2: Omega shot numbers and experimentally observed data. Note that two views are available for some shots, yielding two observation times for some shots, and one repeated observation time (the two views were taken simultaneously). From a predictive perspective, these 6 shots yield 10 measurements. View 2 data are marked with asterisks (*).

2.2 Uncertainties in the CRASH experimental sequence

Some of the uncertainty in quantities of interest for prediction is due to uncertainty in the experimentally "controlled" parameters, referred to in our predictive studies as x. We have characterized many of these during the course of the CRASH project, and we have estimated pdfs of them based on the database of experiments. Figure 8 is an example of a set of laser energy pdfs from several of the CRASH experimental campaigns. These vary systematically from date-to-date due to changes in the Omega facility; during one phase we found and the facility later corrected a tendency for the first shot of the day to come in low. The variability in facility performance in itself raised a challenge: we had to predict before the year-5 experiment is conducted, and so had to estimate in advance the pdf of laser energies on the day of the year-5 experiment.

Figure 9 shows a similar set of distributions for the Xe fill gas pressure; this shows similar changes from experiment to experiment, as the procedures for constructing and filling the shock tubes evolve and as time-delays in firing the laser change uncontrollably from campaign to campaign and shot to shot. The uncertainty of the backlighter timing was significantly reduced for the year-5 experiment, in part in response to our input to the facility.

2.3 Shock Breakout Experiments

A key concern in modeling a laser-driven radiative shock is how accurate the deposition of energy and momentum is. The fact that this is problematic is evident because simulations generally need to reduce the laser energy to obtain correct results, especially in 1D but even in 2D. This makes it very worthwhile to obtain experimental data that can

Date	Shot number	Laser energy (J)	disk thickness (μm)	gas pressure (atm)	time (ns)	shock location (μm)
2008-Oct	52660	3853	21	1.33		
2008-Oct	52661	3890	21	1.13	14	2286
2008-Oct	52661	3890	21	1.13	16	2500
2008-Oct	52663	3890	21	1.17	13	2030
2008-Oct	53664	3820	21	1.09	13	1798
2008-Oct	53664	3820	21	1.09	13	1748
2008-Oct	53665	3892	21	1.13	14	2200
2008-Oct	53665	3892	21	1.13	13	2075
2008-Oct	53666	3859	21	1.2		
2008-Oct	56667	3880	21	1.2	13	2096
2008-Oct	53667	3880	21	1.2	13	2073
2008-Oct	53668	3860	21	1.11	14	2315
2008-Oct	53668	3860	21	1.11	13	2105
2008-Oct	53669	3846	21	1.17	13	1970
2008-Oct	53670	3842	21	1.17	14	2043
2008-Oct	53671	3867	21	1.17	13	1955
2010-Aug	59023	3744	21	1.16		
2010-Aug	59024	3730	21	1.09		
2010-Aug	59025	3801	21	1.24		
2010-Aug	59026	3830	21	1.28		
2010-Aug	59027	3787	21	1.31	20	2737
2010-Aug	59029	3750	21	1.36	26	3489
2010-Aug	59029	3750	21	1.36	26	3434

Table 3: Experimental parameters and shock location for CRASH experiments performed with the nominal tube over multiple years. In some instances an experiment produces data from 2 views and provides 2 data points. Experiments are listed that did not produce data. The uncertainty for the observation time is $\pm$ 0.5 ns and $\pm$ 100 μm for the shock location.

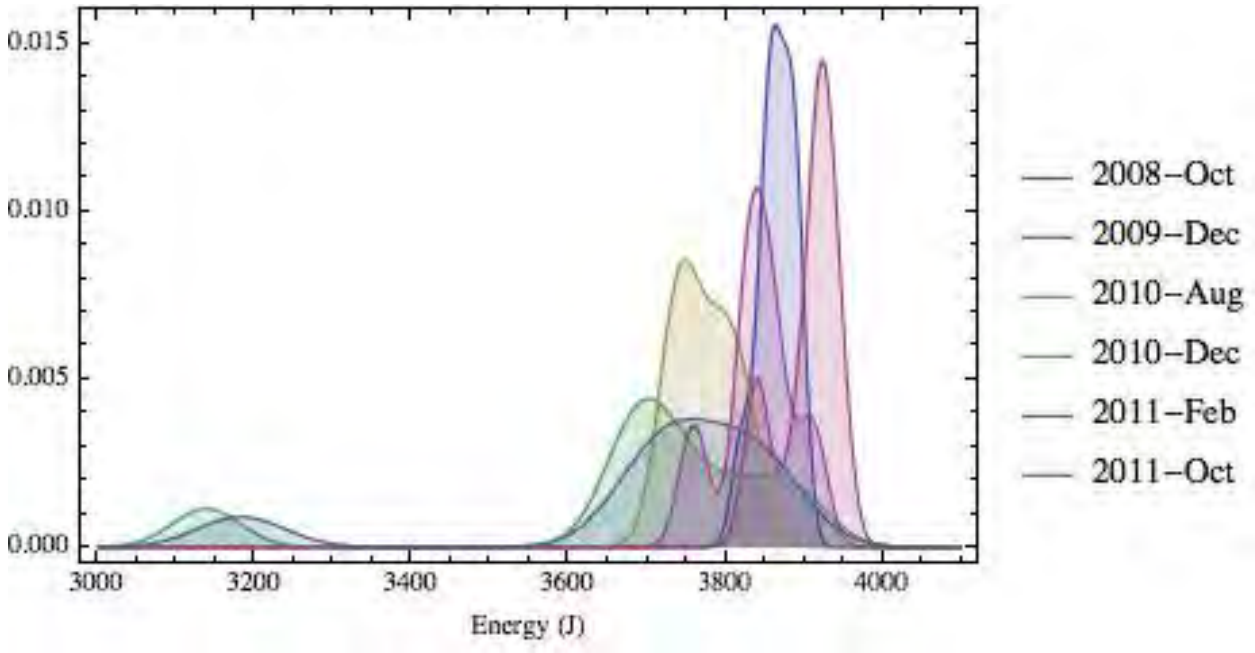

Figure 8: PDFs of the laser energy distribution at Omega.

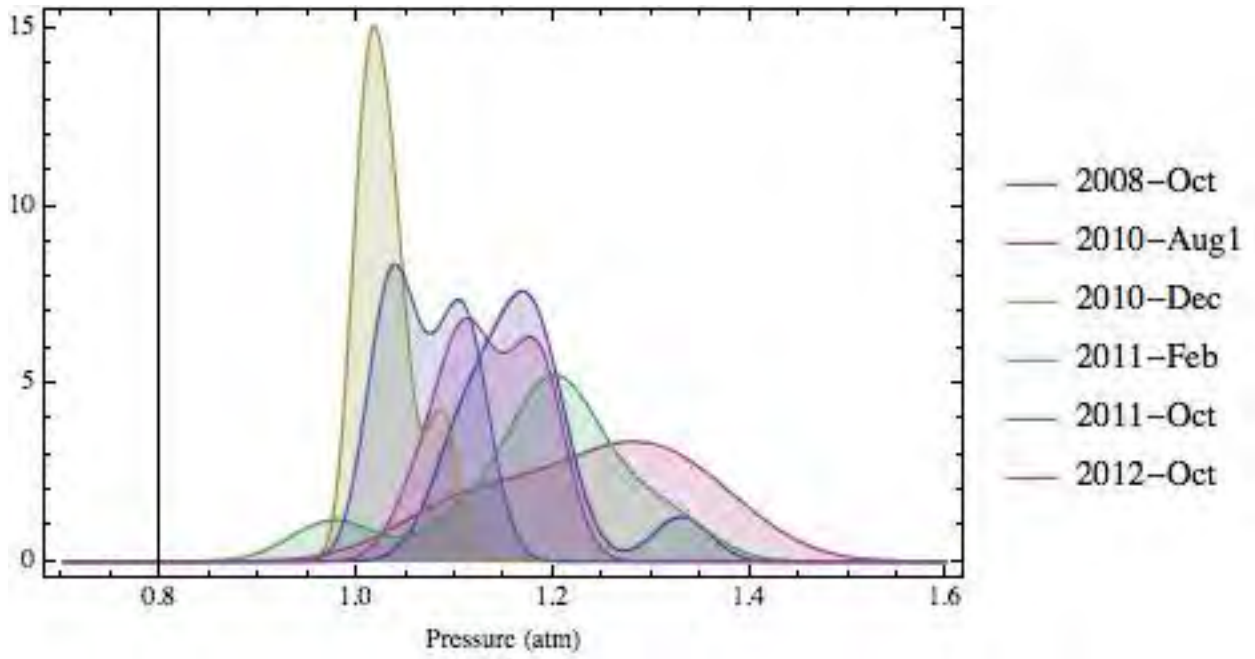

Figure 9: PDFs of the Xe fill pressure distribution in the CRASH targets.

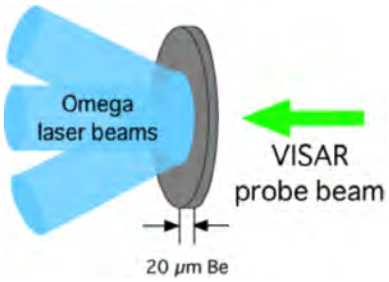

Figure 10: A schematic of the target. A nominally 20 μm Be disk is irradiated with several laser beams. The diagnostics (2 VISARs and SOP) view the rear surface of the target and measure the shock breakout time.

validate the early behavior in the simulation. For this reason we pursued shock breakout measurements.

To characterize the initial state of the radiative shock experiment we performed 12 experiments to measure the shock breakout from a Be disk over three experimental days. The majority of the Be disk thicknesses ranged from 19 - 21 μm $\pm$ 0.5 μm where the average disk thickness was 19.9 μm with a standard deviation of 0.6 μm. Two experiments were performed with a 10 μm $\pm$ 0.5 μm disk. For the 12 experiments performed, the on-target total laser energy averaged 3.841 kJ $\pm$ 0.001 kJ. The range of laser energy for these experiments was 3.403 kJ - 3.946 kJ with a standard deviation of 0.152 kJ. The laser pulse was a 1 ns square pulse with about 100 ps of rise and fall time. The nominal on-target laser irradiance was about 7×10^{14} W/cm^2. Figure 10 shows a sketch of this experiment.

When the shock emerges from the rear surface it is recorded on multiple experimental diagnostics. These diagnostics are calibrated and the amount of time it takes for the shock to move through the Be disk is inferred. Three instruments were used on each experiment to make these measurements and in the majority of experiments all three collected data. Two of the instruments were a Velocity Interferometer System for Any Reflector (VISAR) [Barker and Hollenback, 1972] set to different sensitivities. A VISAR uses a laser with a wavelength of 532 nm to probe a surface and detect the rate of change in the optical path to a surface. This can lead to a measurement of a velocity profile of a surface from which one can infer average pressures. For the experiment reported here, the probe laser is reflected off of the rear (non laser-irradiated side) of the Be disk. Since the Be disk is opaque to the probe laser light, only movement on the rear surface, such as the shock exiting the rear surface of the disk, is recorded. This is referred as shock breakout and the time at which it occurs is referred to as the shock breakout time.

The third diagnostic used to measure the shock breakout time was a Streaked Optical Pyrometer (SOP)[Miller et al., 2007]. A SOP is a passive detector that records thermal emission on a streak camera which results in a 2D image showing the surface emission in space and time. SOP views the rear surface of the target and as the hot shock emerges from the rear of the Be disk its emission will be recorded with the SOP, from which the breakout time can be inferred.

Examples of the typical VISAR and SOP data from this experiment are shown in Figure 11. The SOP data is shown in the left panel with time increasing to the right. There is no detectable emission from the rear surface of the target until the shock emerges from the rear surface of the disk. The thin bright line on the left of this image occurs as the laser pulse is initiated although for unknown reasons. It should be noted that the SOP (and VISAR) data show curvature in the breakout feature. This is due to the profile of the laser spot, which results in a curved shock as the Be disk is larger than the laser spot. For the radiative shock experiment the shock tube is 575 m in diameter; the shock breakout feature is relatively planar over that region.

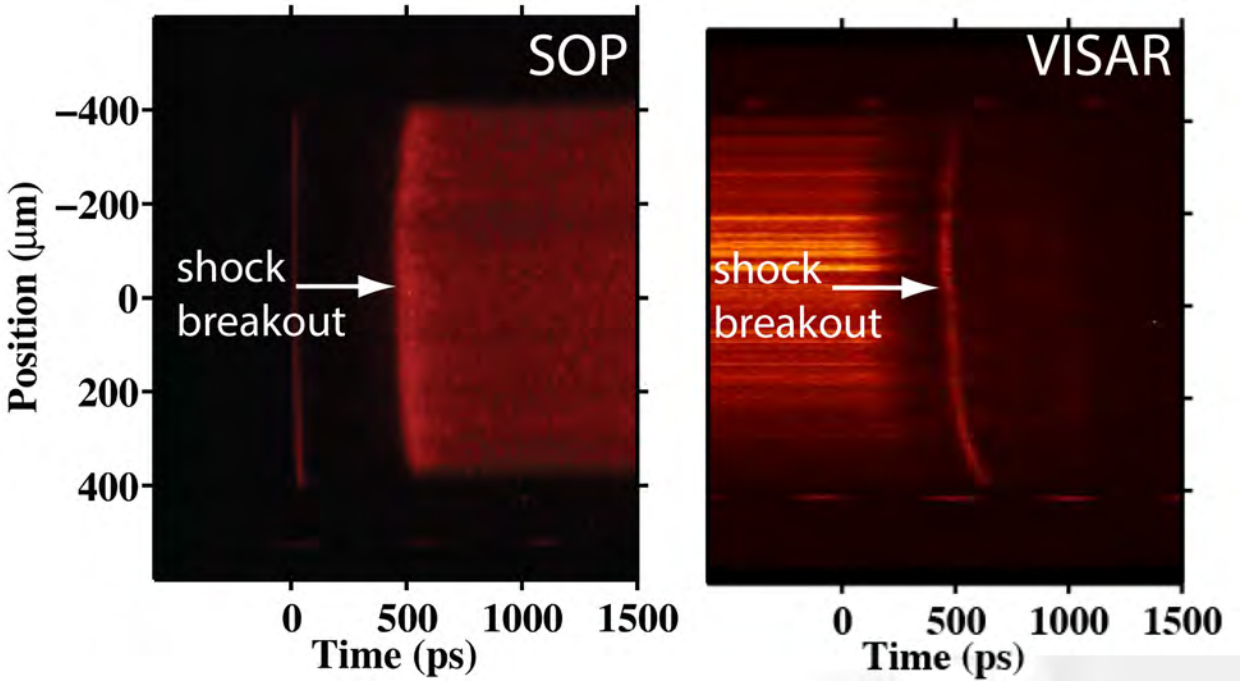

Figure 11: Typical data from experiments performed with a nominally 20 μm thick Be disk from the SOP (left panel) and VISAR (right panel). Both diagnostics view an 800 μm high region on the rear surface of the Be disk and record when the shock breaks out of the disk. The fiducials on the bottom of each image are spaced 548 ps apart.

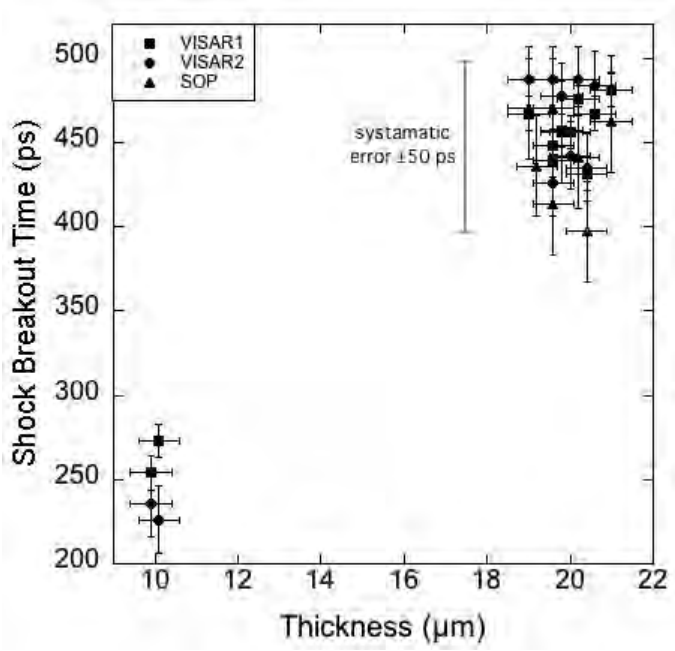

Figure 12: Shock breakout time for Be disks of 19, 20 and 21μm. Data points are offset in thickness to discern between individual experiments. The error bars on individual points indicate the uncertainty due to that diagnostic and the larger error bar on the right represents the systematic error in the experiment due to the uncertainty in the timing of the laser pulse irradiating the Be disk and the laser pulse that creates the fiducial for the experimental data.

Date	Shot number	Laser energy (J)	disk thickness (μm)	gas pressure	time (ns)	shock location
2009-Dec	56301	3838	21	481	481	462
2009-Dec	56302	3925	20	476	487	441
2009-Dec	56303	3938	20	431	435	397
2009-Dec	56304	3888	19	457	477	456
2009-Dec	56305	3915	20	448	487	470
2009-Dec	56306	3913	20			436
2009-Dec	56307	3923	19	467	487	470
2009-Dec	56313	3946	19	439	426	413
2010-Aug	59010	3403	10	273	226	
2010-Aug	59011	3720	20	467	484	
2010-Aug	59014	3844	10	254	236	

Table 4: Experimental parameters and shock breakout time for CRASH experiments. Times are reported for each of the threes diagnostics used. Each diagnostic time has an uncertainty of $\pm$ 50 ps.

The VISAR data is shown in the right panel in Figure 11 with time increasing to the right. In this case, the probe beam is reflected off of the rear surface of the Be disk, which produces the bright signal on the left. The thin, bright line in the image is the shock breaking out of the Be disk. Prior to shock breakout, there is an area on the image with no signal. This occurs after the initial laser pulse has begun and the VISAR probe beam is absorbed, which often occurs at high laser irradiances [Boehly] until the shock breaks out of the disk and the shock itself is reflective.

Both of these VISAR and SOP measurements can be calibrated in time using the timing fidicials, which are created with an optical laser, seen at the bottom of both of the images shown in Figure 11. Figure 12 shows the results. The shock breakout time was measured to the 50% rise of the feature indicating breakout. To discern individual experiments, the thickness of each disk has been offset 0.2 μm. The 3 data points for each disk are from the 2 VISARs and the SOP instruments. In some cases, 1 or 2 of the diagnostics did not produce usable data. The vertical error bars on each point are due to the error in each diagnostic measurement, which is due to the sensitivity of the measurement. The VISARs (set to different sensitivities) were the most sensitive and had errors of $\pm$ 10 ps and $\pm$ 20 ps while the SOP had a larger error of $\pm$ 30 ps. The error in the disk thickness is $\pm$ 0.5 μm. Experimental parameters and shock breakout times are shown in Table 4. There exists a larger systematic error in the time calibration due to the timing of the fiducial laser relative to the laser pulse used to irradiate the disk. The timing of these 2 lasers is known to $\pm$ 50 ps. This timing should be consistent on a single experimental day, but the data shown in Figure 12 was taken over 3 experimental days so the maximum credible offset for the data set is $\pm$ 50 ps. The systematic timing error is the largest error for the data and encompasses the entire timing range observed, which was 397 - 487 ps. This is indicated by a $\pm$ 50 ps error bar on the plot. Under the experimental conditions described above the average shock breakout time for the nominally 20 μm disk was 457 ps.

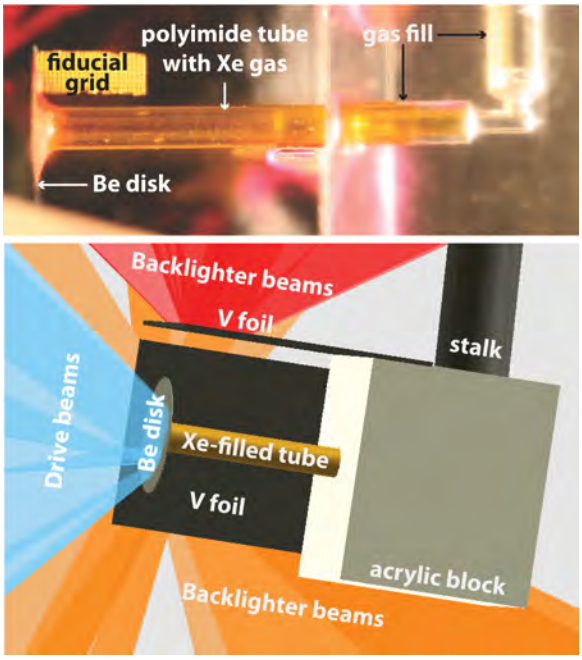

Figure 13: A photograph (top) and model showing laser the interaction of laser beams (bottom) with a target used to create a radiative shock in laboratory experiments. A Be disk behaves as a piston to shock and compress Xe gas in a polyimide tube. The tube and target stalk are inserted into an acrylic block. Two V foils are attached to the target to create x-rays to enable observation of the shocked Xe.

2.4 Early-time radiative shock experiments

As the project moved into year 3, our simulations continued to produce relatively poor fidelity, later overcome by the addition of the laser package. We came to understand that material released early in time from the walls of the cylindrical shock tube was producing a pressure spike on axis that pushed material forward and altered the shape of the shocked Xe layer. This led us to choose to pursue early time measurements to assess whether the simulations were badly wrong about the early behavior.

The target used in these experiments is shown in the top panel of Figure 13. The tube and stalk are inserted into a machined acrylic block. Gas flows through the target stalk and through channels machined into the block to the polyimide tube. The gold grid is used to spatially calibrate the image. There are 2 diagnostic techniques used for this experiment: area radiography and streaked radiography, which will be discussed later in this paper. The top panel of Figure 13 is the diagnostic view for the area radiography technique. The view for the streaked radiography technique would be similar except there would be a wire placed across the tube axis to be used an as a spatial fiducial. The two diagnostic views are orthogonal.

In this experiment, two radiographic techniques (area and streaked radiography) were utilized on a single experiment in an orthogonal geometry. This results in a measurement of velocity over several nanoseconds and an image of the shock from an orthogonal view. Both of these techniques image the x-ray transmission through the experimental target. In both cases, a 5 μm V foil is placed about 2 mm from the target. The placement of these foils is shown in

Date	Shot number	Laser energy (J)	disk thickness (μm)	gas pressure (atm)	time (ns)	shock location (μm)
2010-Dec	60395	3139	21			
2010-Dec	60397	3684	21			
2010-Dec	60398	3847	21	1.01		
2010-Dec	60399	3683	21	1.09		
2010-Dec	60400	3728	21	1.02		
2010-Dec	60402	3781	21	1.04		
2010-Dec	60404	3895	21	1.03	3.9	540
2010-Dec	60404	3895	21	1.03	4.1	578
2010-Dec	60404	3895	21	1.03	4.3	600
2010-Dec	60404	3895	21	1.03	4.5	642
2010-Dec	60405	3699	21	1.01	2.9	302
2010-Dec	60405	3699	21	1.01	3.1	339
2010-Dec	60405	3699	21	1.01	3.3	364
2010-Dec	60405	3699	21	1.01	3.5	397
2011-Feb	61212	3183	22	1.25	1.9	396
2011-Feb	61212	3183	22	1.25	2.4	448
2011-Feb	61212	3183	22	1.25	2.6	471
2011-Feb	61212	3183	22	1.25	2.8	488
2011-Feb	61212	3183	22	1.25	3.0	513
2011-Feb	61212	3183	22	1.25	3.3	544
2011-Feb	61212	3183	22	1.25	3.6	573
2011-Feb	61212	3183	22	1.25	4.1	624
2011-Feb	61212	3183	22	1.25	4.2	642
2011-Feb	61212	3183	22	1.25	2.0	490
2011-Feb	61212	3183	22	1.25	3.9	490
2011-Feb	61214	3759	22	1.21	6.5	917
2011-Feb	61214	3759	22	1.21	6.0	853
2011-Feb	61214	3759	22	1.21	5.5	779
2011-Feb	61215	3695	22	1.20	2.5	401
2011-Feb	61215	3695	22	1.20	2.9	440
2011-Feb	61215	3695	22	1.20	3.3	492
2011-Feb	61215	3695	22	1.20	3.6	533
2011-Feb	61215	3695	22	1.20	4.1	581

Table 5: Experimental parameters for early time CRASH experiments. The error in the shock location is $\pm$ 30 μm and the error in the observation time is $\pm$ 0.2 ns.

Date	Shot number	Laser energy (J)	disk thickness (μm)	gas pressure (atm)	time (ns)	shock location (μm)
2011-Feb	61220	3695	22	1.20	4.5	616
2011-Feb	61216	3844	22	1.32	6.0	812
2011-Feb	61216	3844	22	1.32	5.5	769
2011-Feb	61216	3740	22	1.11	1.6	389
2011-Feb	61216	3740	22	1.11	1.8	410
2011-Feb	61216	3740	22	1.11	2.0	438
2011-Feb	61216	3740	22	1.11	2.2	471
2011-Feb	61216	3740	22	1.11	2.6	510
2011-Feb	61219	3856	22	0.98	3.1	560
2011-Feb	61219	3856	22	0.98	3.6	602
2011-Feb	61219	3856	22	0.98	4.0	651
2011-Feb	61219	3856	22	0.98	4.6	700
2011-Feb	61219	3856	22	0.98	5.2	767
2011-Feb	61219	3856	22	0.98	5.5	793
2011-Feb	61219	3856	22	0.98	4.5	703
2011-Feb	61219	3856	22	0.98	4.0	636
2011-Feb	61219	3856	22	0.98	3.5	560
2011-Feb	61219	3856	22	0.98	4.5	703
2011-Feb	61219	3856	22	0.98	4.0	636
2011-Feb	61219	3856	22	0.98	3.5	560
2011-Feb	61221	3821	22	1.19	3.0	499
2011-Feb	61221	3821	22	1.19	3.8	559
2011-Feb	61221	3821	22	1.19	4.1	589
2011-Feb	61221	3821	22	1.19	5.7	748
2011-Feb	61221	3821	22	1.19	6.2	787
2011-Feb	61221	3821	22	1.19	6.6	830
2011-Feb	61221	3821	22	1.19	6.9	869
2011-Feb	61221	3821	22	1.19	4.0	537
2011-Feb	61221	3821	22	1.19	3.5	485
2011-Feb	61221	3714	22	1.19	3.1	542
2011-Feb	61221	3714	22	1.19	3.5	575
2011-Feb	61221	3714	22	1.19	3.9	616
2011-Feb	61221	3714	22	1.19	4.4	666
2011-Feb	61221	3714	22	1.19	4.8	712
2011-Feb	61221	3714	22	1.19	5.1	748
2011-Feb	61221	3714	22	1.19	6.0	828

Table 6: Experimental parameters for early time CRASH experiments. The error in the shock location is $\pm$ 30 μm and the error in the observation time is $\pm$ 0.2 ns.

Date	Shot number	Laser energy (J)	disk thickness (μm)	gas pressure (atm)	time (ns)	shock location (μm)
2011-Feb	61221	3714	22	1.19	6.4	860
2011-Feb	61221	3714	22	1.19	6.7	891
2011-Feb	61221	3714	22	1.19	4.5	677
2011-Feb	61221	3714	22	1.19	4.0	606
2011-Feb	61221	3714	22	1.19	3.5	536
2011-Feb	61221	3714	22	1.19	4.5	677
2011-Feb	61221	3714	22	1.19	4.0	606
2011-Feb	61221	3714	22	1.19	3.5	536

Table 7: Experimental parameters for early time CRASH experiments. The error in the shock location is $\pm$ 30 μm and the error in the observation time is $\pm$ 0.2 ns.

the bottom panel of Figure 13. The foil is irradiated by additional laser beams at some time after the experiment is initiated by the main laser pulse. This generates He-α like x-rays that pass through the target. This x-ray energy was chosen because it is strongly absorbed in the shocked Xe layer, but transmissive in the unshocked layer. This provides contrast in the resulting image.

An example of an x-ray streaked radiograph is shown in the left panel of Figure 14. Position is increasing away from the Be disk to the right and time is increasing as one moves upwards in the image. Typically, dark regions are associated with regions of low transmission due to high opacity. The shock is the dark vertical feature moving from the bottom left to the top center of the image. This indicates that the shock is approaching the fiducial wire, which is the thick, dark line on the right side of the image. This particular radiograph was imaged from about 2 ns to 4 ns and the shock moved from about 400 μm to 650 μm.

An example of area radiography data (from 1 of the 16 pinholes) is shown in the right panel of Figure 14. In this image the shock tube walls are the horizontal dark lines on the upper and lower regions of the image. The shocked Xe is flowing to the right and the fiducial grid is in the upper portion of the image. The Be is to the left of the shocked Xe, but is transparent to the x-ray energy used to create this image. This radiograph was observed at 4.5 ns after the initial laser pulse began. The shock is about 600 μm from the initial surface of the Be disk and is about 50 μm thick yielding an apparent compression ratio of about 12. This is a lower limit because the layer is almost certainly tiled along the line of sight of the instrument. Analysis of tilting in a collection of data at later times (near 13 ns) found that the actual inferred compression was $\sim$ 22 [Doss et al., 2011b]. A list of experimental parameters and shock locations for these early time experiments are shown in Tables 5, 6 and 7.

2.5 Preparations for tubes with elliptical profiles

In October 2011 we executed experiments that began to explore the effects of geometry on the radiative shock system. This day of experiments served as a transition from the cylindrical tubes of 600 μm diameter we have used previously toward the year 5 targets, which have an initial cylindrical section of 1200 μm diameter that then necks down to an elliptical tube having axes of 1200 μm by 600 μm. We choose to study 3 variations in geometry: wide tube, circular

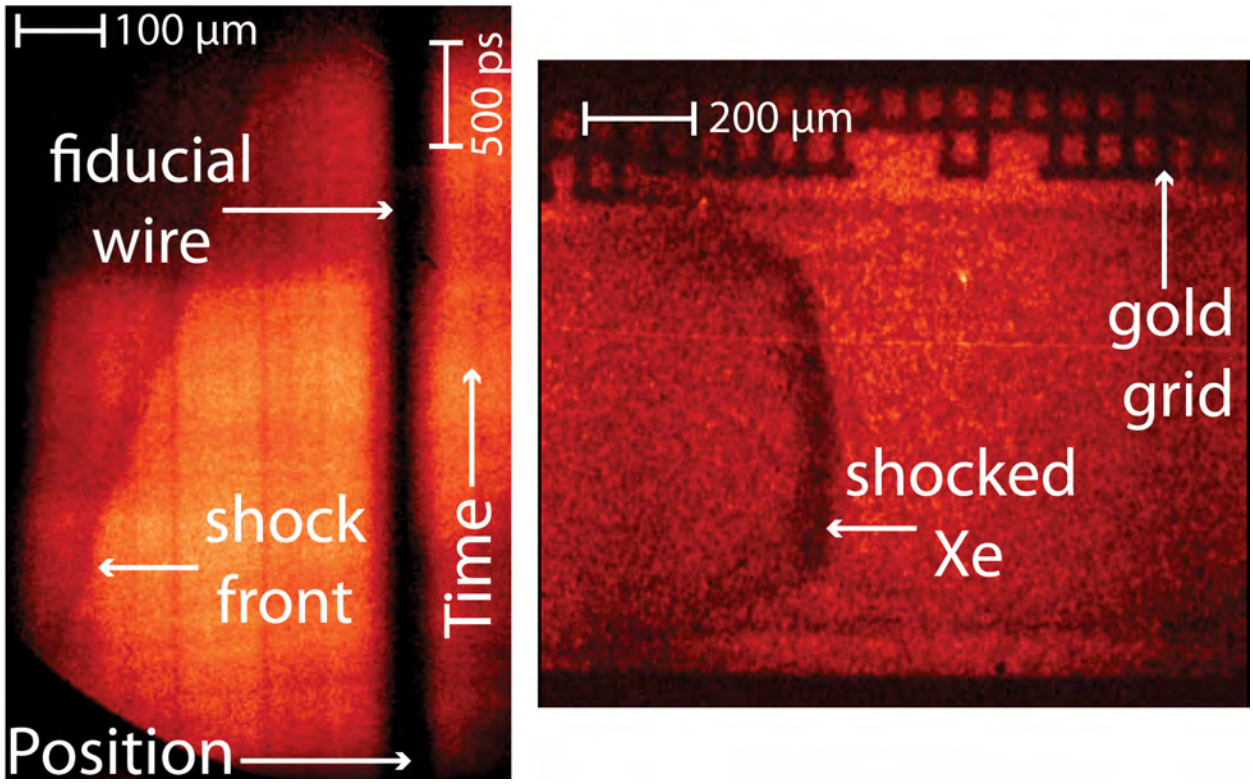

Figure 14: Typical x-ray streaked radiographic (left) and area radiographic (right) data. The shocked Xe is labeled in both images and is moving to the right. In the streaked radiograph the data was observed between about 2 and 4 ns and the shock moved from about 400 μm to 650 μm. In the area radiograph the shock has moved about 600 μm from the laser-irradiated surface of the Be disk at the observation time of 4.5 ns.

nozzle tube and elliptical nozzle tube. Side and rear views of these tubes and their dimensions are shown in Figure 15.

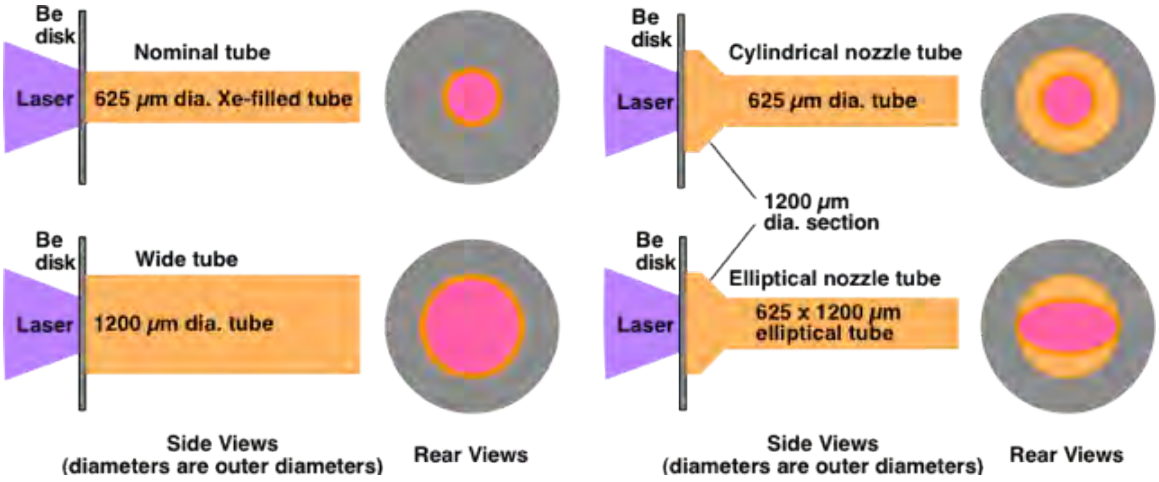

Figure 15: Schematics of tube geometries for the nominal experiments (upper right) and the three geometries used in October 2011, the wide tube(lower right), the cylindrical nozzle (upper left) and the elliptical nozzle (lower left). The elliptical nozzle will be used in the Year 5 experiment.

We shot targets having a straight, cylindrical tube of 1200 μm diameter and targets having an initial cylindrical section of 1200 μm diameter and a circular nozzle that reduced the diameter to 600 μm. Data from these targets were used with sets of simulations as part of the assessment of predictive capability for our year-5 experiments. We also shot targets in the year 5 configuration, as recommended by the Year 3 Review Team report. The purpose of these experiments was to develop and demonstrate the capability to perform such experiments. This was particularly important in the context of target fabrication. We pursued three avenues for production of these targets. One of them succeeded. Figure 16 shows one of these targets. Data from this type of target was not used in the predictions in advance for the year-5 experiments.

Orthogonal ungated radiography was used to diagnose these experiments. X-ray radiographs from an experiment performed with an elliptical nozzle tube are shown in Figure 17. The left panel shows the narrow view of the elliptical tube and was taken at t = 30 ns. The right panel in Figure 17 shows the wide view of the target at t = 28 ns. The most noticeable difference from the nominal experiment in the structure of the shock is the curvature of the shocked Xenon in the wide view. This is expected as the laser spot is smaller than the wide dimension of the tube. A list of experimental parameters and shock locations for these the preparatory experiments with varying tube geometry are shown in Table 8.

2.6 The Year 5 Experiments

The 5^{th} year CRASH experiment was completed on October 11, 2012 at the Omega Laser Facility. The final experiment for the project focused on the effects of tube geometry on the evolution of a radiative shock. A schematic of the experiment is shown in Figure 18. The 20 μm thick Be disk is followed by an nozzle-shaped polyimide tube, tapered to an elliptical cylinder. The tube is initially a circular shape with an outer diameter of 1200 μm. After a length of 500 μm, the cylinder tapers to an elliptical tube with a major axis of 1200 μm and a minor axis of 625 μm. The side and

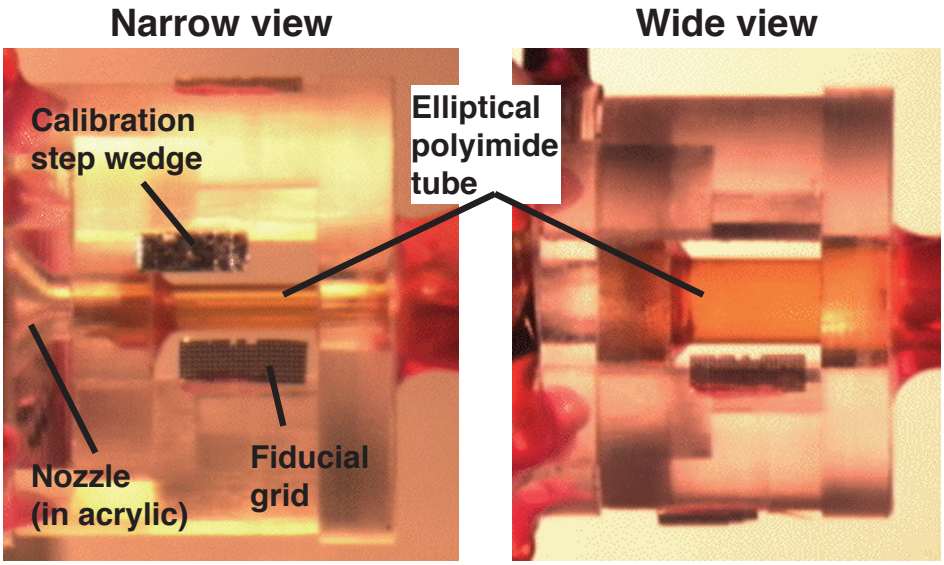

Figure 16: Photographs of a target for the year-5 experiment. An acrylic superstructure supports the nozzle, which mates with a polyimide tube that has been shaped by heating and cooling while distorted.

Date	Shot number	Target type	Laser energy (J)	disk thickness (μm)	gas pressure (atm)	time (ns)	shock location (μm)
2011-Oct	63766	elliptical	3910		1.1		
2011-Oct	63768	circular	3827	21.6	1.1	26	3558
2011-Oct	63768	circular	3827	21.6	1.1	26	3665
2011-Oct	63769	wide	3832	22.2	1.0	26	3347
2011-Oct	63769	wide	3832	22.2	1.0	26	3449
2011-Oct	63771	elliptical	3759		1.1		
2011-Oct	63772	circular	3848		1.0		
2011-Oct	63773	wide	3866	22.4	1.0	28	3413

Table 8: Experimental parameters and shock location for CRASH experiments performed with varying tube geometry. Experiments are listed that did not produce data. The uncertainty in the shock location is $\pm$ 100 μm for the observation time is $\pm$ 0.1 ns.

rear view of the target are shown in Figure 18.

A picture of the target used in the 5^{th} year experiment is shown in Figure 19. The large acrylic super structure is used for multiple reasons. The conical part with gold wedges serves as a diagnostic shield. The base of the acrylic structure contains the elliptical nozzle shape that has been machined into the acrylic. An elliptical polyimide tube is inserted into the acrylic, forming the shape shown in Figure 18. The bottom of the acrylic component is also used to anchor the rear end of the tube and attach the gas fill hose, which is used to fill the tube with Xe gas.

The main diagnostic for this experiment was dual, orthogonal x-ray radiography. This allows us to obtain 2 radiographic images from a single experiment. In this case, we looked at the minor and major axes of the target. Two radiographs are shown in Figure 20 where the left panel shows the minor axis of the target and the major axis is shown on the right. In both images the shock is moving to the right and has been spatially calibrated using the fiducial grid in the lower portion of each image. Each image was taken at 26 ns after the initial laser pulse irradiated the Be disk. Wall shocks are visible in each image as is a large amount of structure in the shocked Xenon. It should be noted that in the wide view the shock front appears to be more curved and has collapsed to a very thin shell.

For this experiment data were obtained for 9 experimental shots with dual radiographs acquired on all but 1 shot. The laser energy for the shot day varied from 3.807 kJ to 3.925 kJ and the gas pressure varied from 1.06 to 1.19 atm. All Be disks were within 0.5 μm of the nominal thickness of 20 μm. The data have been analyzed to extract the shock position and other metrics. A plot of shock location for this data set is shown in Figure 21 with appropriate error bars for shock position and observation time. Figure 22 shows an overall summary of all shock location data. Section 6 below shows a variety of comparisons of these data with simulations.

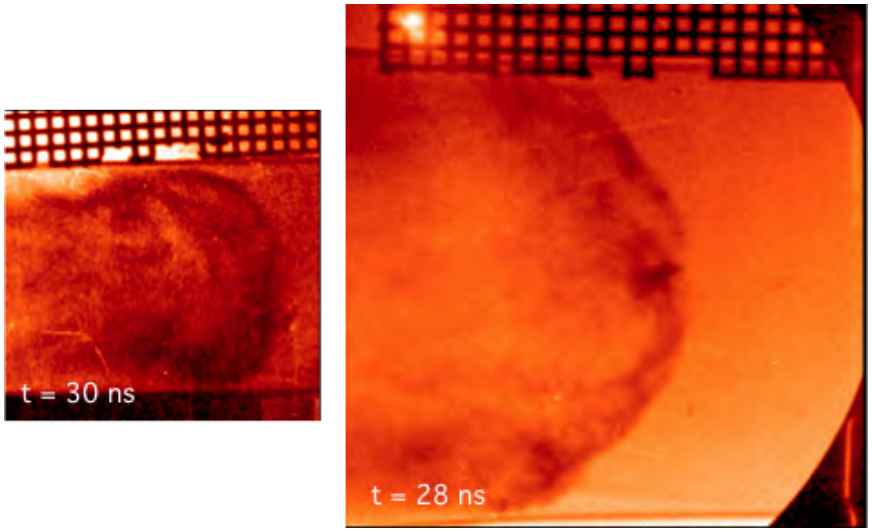

Figure 17: Radiographic images from the October 2011 experiment that explored the effect of variation in tube geometry on the radiative shock system. The images shown here were from an experiment that used an elliptical nozzle tube. The image on the left is a view of the narrow view of the elliptical tube at t = 30 ns and the image on the right is a view of the wide part of the tube at t = 28 ns.

Date	Shot number	Laser energy (J)	disk thickness (μm)	gas pressure (atm)	time (ns)	shock location (μm)
2012-Oct	67701	3925	20.6	1.16	26	3818
2012-Oct	67702	3902	20.8	1.12	26	3778
2012-Oct	67702	3902	20.8	1.12	26	3725
2012-Oct	67703	3850	20.8	1.1	26	3337
2012-Oct	67703	3850	20.8	1.1	26	3166
2012-Oct	67704	3899	20.6	1.1	26	3593
2012-Oct	67704	3899	20.6	1.1	26	3686
2012-Oct	67705	3855	20.5			
2012-Oct	67706	3875	20.5	1.19	26	3612
2012-Oct	67706	3884	20.6	1.19	26	3581
2012-Oct	67706	3884	20.6	1.18	26	3548
2012-Oct	67707	3884	20.6	1.18	26	3500
2012-Oct	67709	3807	20.8			
2012-Oct	67710	3903	20.6	1.06	26	3653
2012-Oct	67710	3903	20.6	1.06	26	3556
2012-Oct	67711	3874	20.6	1.12	26	3645
2012-Oct	67711	3874	20.6	1.12	26	3646

Table 9: Experimental parameters and shock location for CRASH experiments performed with elliptical tubes. Experiments are listed that did not produce data. The uncertainty in the shock location is $\pm$ 100 μm for the observation time is $\pm$ 0.1 ns.

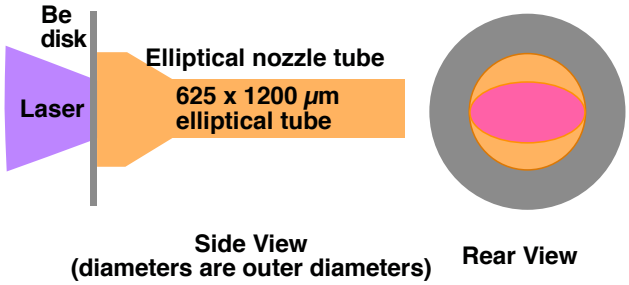

Figure 18: Side and rear schematic of an elliptical nozzle target. The Be disk is nominally 20 μm thick. The outer diameter of the polyimide tube is initially 1200 μm. After 500 μm there is a 500 μm taper to an elliptical shape. The major axis of the ellipse is 1200 μm and the minor axis is 625 μm.

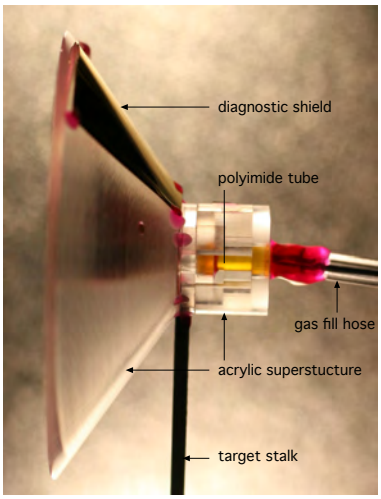

Figure 19: Picture of CRASH target for the Year 5 experiment. The elliptical nozzle is machined into the acrylic super structure and an elliptical-shaped polyimide tube is inserted into the acrylic.

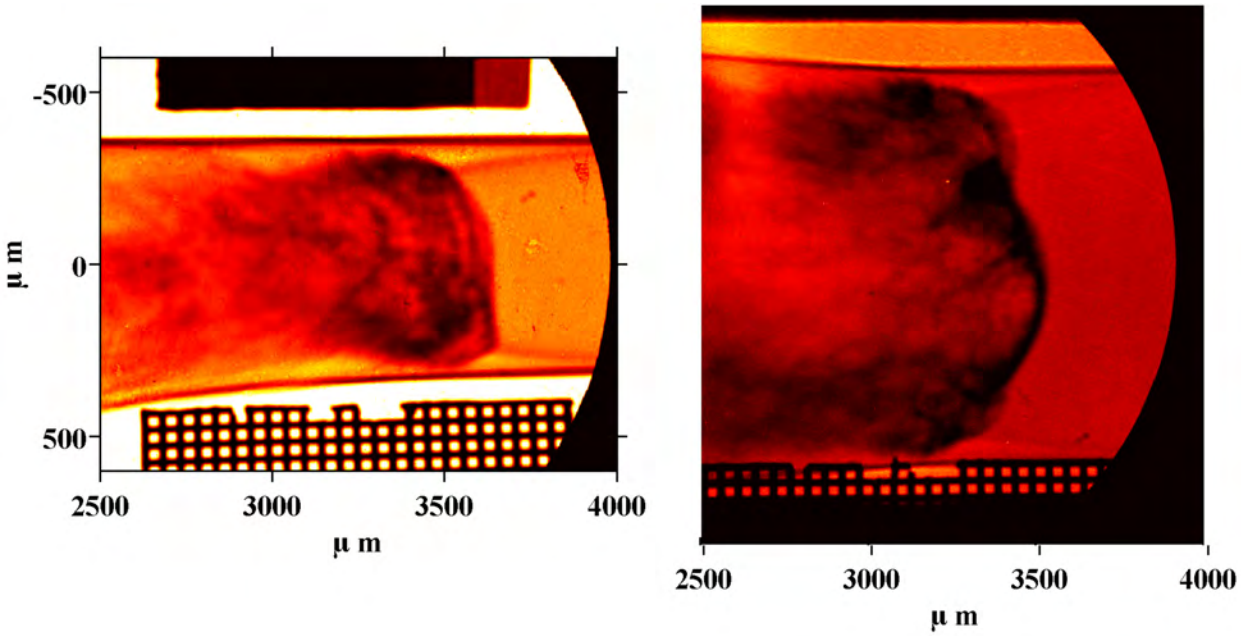

Figure 20: Radiographic images of a single experiment from orthogonal views. The radiative shock is moving to the right and the minor axis of the ellipse (left) and the major axis (right) were imaged. Both images were taken 26 ns after the initial laser pulse began.

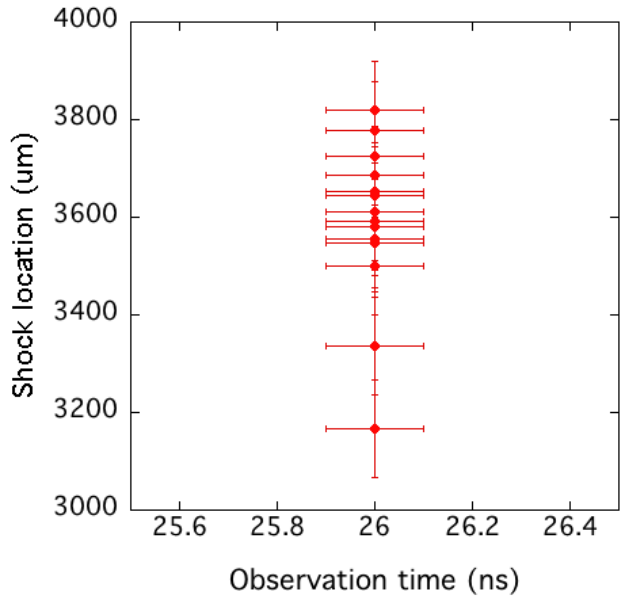

Figure 21: The shock location v. time for multiple experiments. The experimental error is $\pm$ 100 μm for the shock location and $\pm$ 100 ps in the observation time.

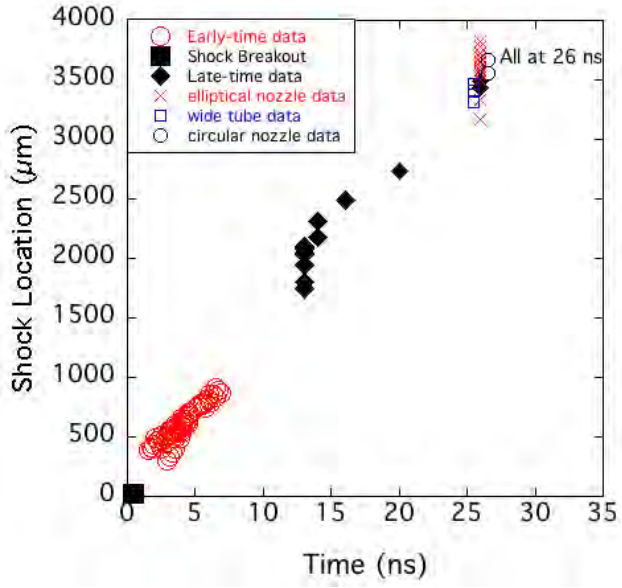

Figure 22: The shock location vs. time for multiple experiments. The experimental error is $\pm 100 \mu$m for shock location and ± 100 ps in the observation time.

3 The CRASH code

We began the CRASH project with an existing code for compressible plasmas: the BATS-R-US code [Powell et al., 1999, Toth et al., 2012]. We extended it with a new radiation-transfer and heat-conduction capability. This new combination, together with the equation-of-state and multigroup opacity solver, is called the CRASH code. This code uses the recently developed parallel Block Adaptive Tree Library (BATL) [Toth et al., 2012] to enable highly resolved radiation hydrodynamic solutions. The implemented radiation hydrodynamic schemes solve for gray or multigroup radiation diffusion models in the flux-limited diffusion approximation. Since radiative shocks are the main application for CRASH, we have also implemented a separate electron pressure equation with electron thermal heat conduction. For the electron heat conduction, we have added the option of a flux limiter to limit the thermal flux with the free-streaming heat flux.

The multi-material radiation hydrodynamic equations are solved with an operator split method that consists of three substeps: (1) solving the hydrodynamic equations with standard finite volume shock-capturing schemes, (2) the linear advection of the radiation in frequency-logarithm space, and (3) the implicit solution of the radiation, heat conduction, and energy exchanges. For the implicit solver, standard Krylov solvers are used together with a Block Incomplete Lower-Upper decomposition (BILU) preconditioner. This preconditioner scales well up to 1000 processors on present-day parallel computers. The pieces of the scheme are described below.

In the following, Section 3.1 introduces the radiation hydrodynamic equations for multi-material plasmas, in a form general enough to apply at high energy density. Section 3.2 describes the numerical algorithms to solve these equations.

3.1 Equations of radiation hydrodynamics in dense plasmas

The equations of radiation hydrodynamics describe the time evolution of both matter and radiation. For the applications that supported the work reported here, the code must be able to model matter as a high energy density plasma that is in LTE so that the population of all atomic and ion states can be obtained from statistical physics (see for instance Landau and Lifshitz [1980]). We allow for multiple materials throughout the spatial domain of interest, but restrict the analysis to plasma flows that are far from relativistic. The materials can be heated to sufficiently high temperatures so that they can ionize and create free electrons, introducing the need for a time evolution equation for the electron energy density. The electrons transfer heat by thermal heat conduction and emit and absorb photon radiation. The radiation model discussed in this section is non-equilibrium diffusion, in which the electron and radiation temperature can be different. We approximate the radiation transfer with a gray or multigroup flux-limited diffusion (FLD). This model is also of interest for application to a number of astrophysical problems.

In the following subsections, we will describe the radiative transfer equations for the evolution of the multigroup radiation energy densities (Section 3.1.1) in the FLD approximation (Section 3.1.5). The coupling of the radiation field to the two species hydrodynamic equations of electrons and ions are discussed in Section 3.1.2. In Section 3.1.3, the method for tracking the different materials is treated, while the lookup tables used for of the EOS and opacities are mentioned in Section 3.1.4.

3.1.1 Radiation Transport

In this section, we will build up the form of the radiation transport in the multigroup diffusion approximation that is used for the implementation in the CRASH code. The spectral pressure tensor is often approximated in the form [Mihalas and Weibel-Mihalas, 1984]

$$\mathbf{P}_\nu(\mathbf{r},t,\nu) = E_\nu \mathbf{T}_\nu, \tag{1}$$

where

$$\mathbf{T}_\nu(\mathbf{r},t,\nu) = \frac{1}{2}(1-\chi_\nu)\mathbf{I} + \frac{1}{2}(3\chi_\nu - 1)\frac{\mathbf{F}_\nu \mathbf{F}_\nu}{|\mathbf{F}_\nu|^2}, \tag{2}$$

is the spectral Eddington tensor, χ_ν is the Eddington factor, and $\mathbf{I}$ is the identity matrix. The second term on the right hand side is a dyad constructed from the direction of the spectral radiation flux. The pressure tensor can be used to arrive at a time evolution equation for the solid angle integrated spectral radiation energy [Buchler, 1983]

$$\frac{\partial E_\nu}{\partial t} + \nabla\cdot(E_\nu \mathbf{u}) - \nu\frac{\partial}{\partial \nu}(\mathbf{P}_\nu : \nabla\mathbf{u}) = \text{diffusion} + \text{emission} - \text{absorption}, \tag{3}$$

which contains the velocity $\mathbf{u}$ of the background plasma. Here the colon denotes the contraction of the two tensors $\mathbf{P}_\nu$ and $\nabla\mathbf{u}$. The processes described by the symbolic terms on the right hand side of equation (3) will be described below.

Setting the Eddington factor $\chi_\nu = 1/3$ corresponds to the radiation diffusion model. In this case the radiation is assumed to be effectively isotropic and the spectral radiation pressure can be described by the scalar

$$p_\nu = \frac{1}{3}E_\nu = (\gamma_r - 1)E_\nu, \tag{4}$$

where we have introduced the adiabatic index of the radiation field, which in this case has the relativistic value $\gamma_r = 4/3$. The time evolution for the spectral energy density can then be simplified to

$$\frac{\partial E_\nu}{\partial t} + \nabla\cdot(E_\nu \mathbf{u}) - (\gamma_r - 1)(\nabla\cdot u)\nu\frac{\partial E_\nu}{\partial \nu} = \text{diffusion} + \text{emission} - \text{absorption}. \tag{5}$$

The second and third terms on the left hand side of equation (5) express the change in the spectral energy density due to the advection and compression of the background plasma, which moves with the velocity $\mathbf{u}$, as well as the frequency shift due to compression. In the free-streaming limit where the radiation hardly interacts with the matter, χ_ν approaches one. In this section we will keep $\chi_\nu = 1/3$ and at the same time use a flux-limited diffusion for the free-streaming regime whenever needed (see Section 3.1.5).

The set of equations for the spectral energy density (5) still consists of an infinite number of equations, one for each frequency. A finite set of governing equations to describe the radiation transport in the multigroup diffusion approximation is obtained when we choose a set of frequency groups. Here we enumerate groups with the index, $g = 1,\ldots,G$. The interval of the photon frequencies, relating to the gth group is denoted as $[\nu_{g-1/2}, \nu_{g+1/2}]$. A discrete set of group energy densities, E_g, is introduced in terms of the integrals of the spectral energy density of the frequency group interval:

$$E_g = \int_{\nu_{g-1/2}}^{\nu_{g+1/2}} E_\nu d\nu. \tag{6}$$

Now we can integrate equation (5) to arrive at the desired set of the multigroup equations:

$$\begin{aligned} \frac{\partial E_g}{\partial t} \quad &+ \quad \nabla\cdot(E_g\mathbf{u}) + (\gamma_r - 1)E_g\nabla\cdot u - (\gamma_r - 1)(\nabla\cdot u)\int_{\nu_{g-1/2}}^{\nu_{g+1/2}} \frac{\partial(\nu E_\nu)}{\partial \nu} d\nu \\ &= \quad \int_{\nu_{g-1/2}}^{\nu_{g+1/2}} (\text{diffusion} + \text{emission} - \text{absorption}) d\nu. \end{aligned} \tag{7}$$

The fourth term on the left hand side is a frequency shift due to the plasma compression. This term is essentially a conservative advection along the frequency axis.

In the context of the multi-group radiation diffusion, a discussion about the stimulated emission is not less important than LTE. Excellent accounts on the stimulated emission exist in the literature, see for instance Zeldovich and Raizer [1966]. Here, we merely summarize how the stimulated emission modifies the absorption opacity κ_ν^a obtained from, e.g., opacity tables. This is important when dealing with externally supplied opacity tables, since the CRASH code assumes that the absorption opacities are corrected. Integrating the total absorption and emission over all directions and summing up the two polarizations of the photons, the following expression can be derived for the emission and absorption

$$\text{emission} - \text{absorption} = c\kappa_\nu^{a\prime}\left(B_\nu - E_\nu\right), \tag{8}$$

where the effective absorption coefficient, $\kappa_\nu^{a\prime}$, is introduced to account for the correction due to stimulated emission:

$$\kappa_\nu^{a\prime} = \kappa_\nu^a\left(1 - \exp\left[-\frac{\varepsilon}{k_B T_e}\right]\right), \tag{9}$$

in which $\varepsilon = h\nu$ is the photon energy, k_B is the Boltzmann constant, and T_e is the electron temperature. We also introduced the spectral energy density distribution of the black body radiation (the Planckian)

$$B_\nu = \frac{8\pi}{h^3c^3}\frac{\varepsilon^3}{\exp[\varepsilon/(k_B T_e)] - 1}. \tag{10}$$

The total energy density in the Planck spectrum equals $B = \int_0^\infty d\nu B_\nu = aT_e^4$, where $a = 8\pi^5 k_B^4/(15h^3c^3)$ is the radiation constant.

We use the standard definition of the group Planck mean opacity κ_{Pg} and group Rosseland mean opacity κ_{Rg} [Mihalas and Weibel-Mihalas, 1984]

$$\kappa_{Pg} = \frac{\int_{\nu_{g-1/2}}^{\nu_{g+1/2}} d\nu \kappa_\nu^{a\prime} B_\nu}{B_g}, \qquad \kappa_{Rg} = \frac{\frac{\partial B_g}{\partial T_e}}{\int_{\nu_{g-1/2}}^{\nu_{g+1/2}} d\nu \frac{1}{\kappa_\nu^t}\frac{\partial B_\nu}{\partial T_e}}, \qquad B_g = \int_{\nu_{g-1/2}}^{\nu_{g+1/2}} d\nu B_\nu \tag{11}$$

in which κ_ν^t is the spectral total opacity. The right hand side of equation (7) can now be written as (see for instance Mihalas and Weibel-Mihalas [1984], Pomraming [2005])

$$\begin{aligned}\frac{\partial E_g}{\partial t} &+ \nabla\cdot(E_g\mathbf{u}) + (\gamma_r - 1)E_g\nabla\cdot u - (\gamma_r - 1)\nabla\cdot u\int_{\nu_{g-1/2}}^{\nu_{g+1/2}}\frac{\partial(\nu E_\nu)}{\partial\nu}d\nu \\ &= \nabla\cdot(D_g\nabla E_g) + \sigma_g(B_g - E_g),\end{aligned} \tag{12}$$

where $D_g = c/(3\kappa_{Rg})$ is the radiation diffusion coefficient for radiation group g in the diffusion limit. The absorption and emission uses the coefficient $\sigma_g = c\kappa_{Pg}$. These group mean opacities are either supplied by lookup tables or by an opacity solver.

In a single group approximation (gray diffusion), the spectral energy density is integrated over all photon frequencies and the total radiation energy density is obtained by

$$E_r(\mathbf{r},t) = \int_0^\infty E_\nu d\nu. \tag{13}$$

This amounts to summing up all groups $E_r = \sum_g E_g$. The gray radiation diffusion equation can be derived as (see for instance Mihalas and Weibel-Mihalas [1984], Pomraming [2005], Drake [2006])

$$\frac{\partial E_r}{\partial t} + \nabla\cdot(E_r\mathbf{u}) + (\gamma_r - 1)E_r\nabla\cdot u = \nabla\cdot(D_r\nabla E_r) + \sigma_r(B - E_r), \tag{14}$$

where the diffusion coefficient D_r in the diffusion limit is now defined by the single group Rosseland mean opacity κ_R as $D_r = c/(3\kappa_R)$, and the absorption coefficient σ_r is defined in terms of the single group Planck mean opacity κ_P as $\sigma_r = c\kappa_P$.

3.1.2 Hydrodynamics

In the CRASH code, a single fluid description is used, so that all of the atomic and ionic species as well as the electrons move with the same bulk velocity $\mathbf{u}$. The conservation of mass

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0, \tag{15}$$

provides the time evolution of the mass density ρ of all the materials in the simulation. The plasma velocity is obtained from the conservation of momentum

$$\frac{\partial \rho \mathbf{u}}{\partial t} + \nabla \cdot [\rho \mathbf{u}\mathbf{u} + I(p + p_r)] = 0. \tag{16}$$

The total plasma pressure is the sum of the ion and electron pressures: $p = p_i + p_e$. The net force of the radiation on the plasma is given by the gradient of the total radiation pressure $-\nabla p_r$, where the total radiation pressure is obtained from the group radiation energies: $p_r = (\gamma_r - 1)\sum E_g$.

In a high density plasma, the electrons are very strongly coupled to the ions by collisions. However, for higher temperatures, the electrons and ions get increasingly decoupled. At a shock front, where ions are preferentially heated by the shock wave, the electrons and ions are no longer in temperature equilibrium. Ion energy is transferred to the electrons by collisions, while the electrons in turn radiate energy. We therefore solve separate equations for the ion/atomic internal energy density E_i and the electron internal energy density E_e:

$$\frac{\partial E_i}{\partial t} + \nabla \cdot (E_i \mathbf{u}) + p_i \nabla \cdot u = \sigma_{ie}(T_e - T_i), \tag{17}$$

$$\frac{\partial E_e}{\partial t} + \nabla \cdot (E_e \mathbf{u}) + p_e \nabla \cdot u = \nabla \cdot (C_e \nabla T_e) + \sigma_{ie}(T_i - T_e) + \sum_{g=1}^{G} \sigma_g (E_g - B_g). \tag{18}$$

The coupling coefficient $\sigma_{ie} = n_a k_B / \tau_{ie}$ in the collisional energy exchange between the electrons and ions depends on the ion-electron relaxation time $\tau_{ie}(T_e, n_a, m)$ and the atomic number density n_a. Energy transfer depends also on the difference between ion temperature T_i and electron temperature T_e. In equation (18), we have included electron thermal heat conduction with conductivity $C_e(T_e, n_a, m)$. Since electrons are the species are responsible for radiation emission and absorption, the energy exchange between the electrons and the radiation groups is added to equation (18).

For the development of the numerical schemes in Section 3.2, we will use an equation for the conservation of the total energy density

$$e = \frac{\rho u^2}{2} + E_i + E_e + \sum_{g=1}^{G} E_g, \tag{19}$$

instead of the equation for ion internal energy (17). This is especially important in regions of the computational domain where hydrodynamic shocks can occur, so that we can recover the correct jump conditions. Conservation of total energy can be derived from equations (12) and (15)–(18) as

$$\frac{\partial e}{\partial t} + \nabla \cdot [(e + p + p_r)\mathbf{u}] = \nabla \cdot (C_e \nabla T_e) + \sum_{g=1}^{G} \nabla \cdot (D_g \nabla E_g). \tag{20}$$

The frequency shift term in equation (12), due to the plasma compression, does not show up in the conservation of total energy if we use energy conserving boundary conditions at the end points of the frequency domain, i.e. at $\nu = 0$ and $\nu = \infty$ in the analytical description or at the end points of the numerically truncated finite domain.

3.1.3 Level Sets and Material Identification

In many CRASH applications, we need a procedure to distinguish between different materials that may be present. We assume that materials do not mix, but differ from each other by their properties such as the equation of state and opacities. If we use M different materials, then we can define for each material $m = 1, \ldots, M$ the level set function $d_m(\mathbf{r}, t)$ (see for instance Kreiss and Olsson [2005], Olsson et al. [2007], Sussman and Puckett [2000]) that is initially set to zero at the material interface, while positive inside the material region and negative outside. Generally, we use a smooth and signed distance function in the intial state. At later times, the location of material m is obtained by means of a simple advection equation

$$\frac{\partial d_m}{\partial t} + \nabla \cdot (d_m \mathbf{u}) = d_m \nabla \cdot u. \tag{21}$$

For any given point in space and time, we can determine what the material is, since analytically only one of the level set functions d_m can be positive at any given point. Numerical errors will create regions where this is not true in solutions to the discretized form of equation (21). Whenever this happens, the material having the largest d_m is assigned as the material in that region. This is a simple approximation, and we may explore more sophisticated approaches in the future. The number of material levels M is configured at compile time.

3.1.4 Equation of State and Opacities

We have implemented EOS solvers and a code to calculate the frequency-averaged group opacities. This implementation will be reported elsewhere, but it is important to note that in the EOS and opacity solver, the temperature is assumed to be well below relativistic values: $T \ll 10^5$ eV. A non-relativistic speed of motion is also assumed, which simplifies the radiation transport equation and allows neglect of relativistic corrections for opacities. In this section, we will assume that all necessary quantities are calculated and stored in lookup tables. Our EOS solver assumes that the corrections associated with ionization, excitation, and Coulomb interactions of partially ionized ion-electron plasma are all added to the energy of the electron gas and to the electron pressure. This is possible since these corrections are controlled by the electron temperature. The ion internal energy density, ion pressure, and ion specific heat in an isochoric process per unit of volume are simply

$$p_i = n_a k_B T_i, \qquad E_i = \frac{p_i}{\gamma - 1}, \qquad C_{Vi} = \left(\frac{\partial E_i}{\partial T_i} \right)_\rho = \frac{n_a k_B}{\gamma - 1}, \tag{22}$$

which are due to the contributions from ion translational motions, for which $\gamma = 5/3$.

The relations among electron internal energy density, pressure, density, and temperature are known as the EOS. To solve these relations is usually complex and time consuming. We therefore store these relations in invertible lookup tables. For each material m, our EOS tables have logarithmic lookup arguments $(\log T_e, \log n_a)$. The list of quantities stored in these tables is indicated in Table 10. These lookup tables are populated with quantities that are needed for both single- and two-temperature simulations. For two-temperature plasma simulations, we need p_e, E_e, the electron specific heat C_{Ve}, and the electron-speed-of-sound gamma γ_{S_e}. For convenience, we add the total pressure $p = p_e + p_i$,

quantity	stored quantity	units
total pressure p	p/n_a	eV
total internal energy density E	E/n_a	eV
electron pressure p_e	p_e/n_a	eV
electron internal energy density E_e	E_e/n_a	eV
specific heat C_V	$C_V/(n_a k_B)$	
electron specific heat C_{Ve}	$C_{Ve}/(n_a k_B)$	
speed of sound gamma	γ_S	
electron speed of sound gamma	γ_{S_e}	
inverse of ion-electron interaction time	$1/\tau_{ie}$	s^{-1}
electron conductivity	C_e	$\mathrm{J\,m^{-1}\,s^{-1}\,K^{-1}}$
mean ionization	$\overline{Z}$	
mean square ionization	$\overline{Z^2}$	

Table 10: Quantities stored in the EOS tables as a function of $\log T_e\,[\mathrm{eV}]$ and $\log n_a\,\left[\mathrm{m}^{-3}\right]$

total internal energy density $E = E_i + E_e$, single temperature specific heat C_V, and the single-temperature-speed-of-sound gamma γ_S, which can be used in single temperature simulations. We use high enough table resolutions so that it is sufficient to use a bilinear interpolation in the lookup arguments. If p_e or E_e (or p and E in single temperature mode) are known on entry to the lookup instead of T_e, we do a binary search in the table to find the appropriate electron temperature. The latter works only if the necessary thermodynamic derivatives are sign definite, i.e. the table is invertible. Other thermodynamic quantities that are needed, but not stored in these lookup tables, can be derived. For example, the electron density can be obtained from the mean ionization $n_e = n_a\overline{Z}$.

quantity	symbol	units
specific group Rosseland mean opacities	κ_{Rg}/ρ	$\mathrm{kg^{-1}\,m^2}$
specific group Planck mean opacities	κ_{Pg}/ρ	$\mathrm{kg^{-1}\,m^2}$

Table 11: Quantities stored in the opacity tables as a function of $\log\rho\,\left[\mathrm{kg\ m^{-3}}\right]$ and $\log T_e\,[\mathrm{eV}]$.

In addition, we have lookup tables for the averaged multigroup opacities. These tables are either constructed internally for a given frequency range, number of groups, and selected materials, or externally supplied. For any material m, the logarithmic lookup arguments are $(\log\rho, \log T_e)$. The stored quantities, see Table 11, are the specific Rosseland mean opacity κ_{Rg}/ρ and specific Planck mean opacity κ_{Pg}/ρ for all groups $g = 1,\ldots,G$ used during a simulation. Planck opacities are assumed to be corrected for stimulated emission, as discussed in Section 3.1.1. The groups are always assumed to be logarithmically distributed in frequency space.

3.1.5 Flux-limited Diffusion

Radiation diffusion, if uncorrected, can transport energy too efficiently in the optically thin free streaming limit. In the diffusion limit, the radiation diffusion flux for each group follows Fick's law $\mathbf{F}_g = -D_g \nabla E_g$, where the diffusion coefficient D_g depends on the Rosseland mean opacity κ_{Rg} for the group g via $D_g = c/(3\kappa_{Rg})$. However, this flux is potentially unbounded. In the optically thin free-streaming limit, the magnitude of the radiation flux can be at most cE_g in order to maintain causality. Various flux limiters exist in the literature, see for instance Minerbo [1978], Lund and J.R. [1980], Levermore and Pomraning [1981], that ensure that the diffusion flux is limited by this free streaming flux. We have implemented the so-called square-root flux limiter to obtain the correct progation speed in the optically thin regime [Morel, 2000]. For this flux limiter, the diffusion coefficient is rewritten as

$$D_g = \frac{c}{\sqrt{(3\kappa_{Rg})^2 + \frac{|\nabla E_g|^2}{E_g^2}}}. \tag{23}$$

In the limit that the radiation length scale $L_R = E_g/|\nabla E_g|$ is large, the diffusive limit is recovered. For a small radiation length scale, $D_g = c|E_g|/|\nabla E_g|$ and the radiation propagates with the speed of light.

Similarly, we have implemented the option to limit the electron thermal heat flux (see Drake [2006] for more details on electron flux limiters). The classical Spitzer-Harm formula for the collisional electron conductivity is proportional to $T_e^{5/2}/\overline{Z^2}$, where $\overline{Z^2}$ is the mean square ionization of the used material. The collisional model is only valid when the temperature scale length $L_T = T_e/|\nabla T_e|$ is much larger than the collisional mean free path of the electrons λ_{mfp}. When the temperature scale length is only a few λ_{mfp} or smaller, this description breaks down. This may for instance happen in laser-irradiated plasmas. In such a case, one could determine the heat flux by solving the Fokker-Planck equation for the electrons, but this is computationally expensive. Instead, we use a simplified model to limit the electron heat flux. A free-streaming heat flux can be defined as the thermal energy density in the plasma transported at some characteristic thermal velocity: $F_{FS} = n_e k_B T_e v_{th}$, where $v_{th} = \sqrt{k_B T_e/m_e}$. For practical applications, the maximum heat transport is usually only a fraction of this free-streaming flux: $\mathbf{F} = -(fF_{FS}/|\nabla T_e|)\nabla Te$, where f is the so-called flux limiter. This heat flux model is the threshold model and is also used in other radiation hydrodynamics packages, such as HYADES [Larsen and Lane, 1994]. The flux-limited heat flux can now be defined as

$$\mathbf{F} = -\min\left(C_e, \frac{fF_{FS}}{|\nabla T_e|}\right)\nabla T_e. \tag{24}$$

The flux limiter f is an adjustable input parameter and can be tuned to let the simulated results better fit reality.

3.2 The numerical method

In this section, we describe the discretization of the set of multi-material, radiation-hydrodynamics equations for the density (15), momentum (16), total energy (20), electron internal energies (18), radiation group energy (12), and material level-set functions (21). The equations are time integrated using an operator-split method to solve the equations in substeps. Formally, we may write this system as

$$\frac{\partial \mathbf{U}}{\partial t} = \mathbf{R}_{\text{hydro}}(\mathbf{U}) + \mathbf{R}_{\text{frequency}}(\mathbf{U}) + \mathbf{R}_{\text{diffusion}}(\mathbf{U}), \tag{25}$$

where $\mathbf{U}$ is the vector of state variables. We have split the right hand sides of the equations into three parts and time advance the equations with an operator splitting method in the following order: (1) The right hand side $\mathbf{R}_{\text{hydro}}$

describes the advection and pressure contributions (Section 3.2.1). This part is essentially the ideal hydrodynamic equations augmented with the advection and compression of the radiation energy, the electron internal energy, and the material level sets. (2) The right hand side $\mathbf{R}_{\text{frequency}}$ is the advection of the radiation field in frequency space (Section 3.2.2). (3) The right hand side $\mathbf{R}_{\text{diffusion}}$ takes care of the diffusion and energy exchange terms, which we will solve with an implicit scheme (Section 3.2.3). This choice of operator splitting is not unique. Instead of splitting the hydrodynamic advection operator and the extra advance operator for the frequency advection, one could attempt to discretize the frequency advection flux as an extra flux for the control volume of the four-dimensional (x,y,z,ν) space. However, since the CRASH code is built around the existing BATS-R-US code in 1D, 2D, and 3D, we opted for splitting the frequency advection from the hydro update. Boundary conditions are treated in Section 3.2.4.

3.2.1 Hydro Solve

The first step of the operator splitting is an update of the hydrodynamic equations, including the advection and compression of the radiation energy density, electron internal energy density and the level sets. We have implemented two variants to solve the hydrodynamic equations: (1) using conservation of the total energy (Section 3.2.1.1) and (2) a non-conservative pressure formulation (Section 3.2.1.2). We can also combine the two discretizations in a hybrid manner within a multi-material simulation [Karni, 1996].

3.2.1.1 Conservative Approach

We have implemented several hydrodynamic shock-capturing schemes in the CRASH code: the HLLE scheme [Harten et al., 1983, Einfeldt et al., 1991], the Rusanov scheme [Yee, 1989], and a Godunov scheme [Godunov, 1959] with an exact Riemann solver. In this section, we will explain how we generalized the HLLE scheme for our system of equations that includes radiation, level sets, and an EOS. The other hydrodynamic schemes can be generalized in a similar fashion.

Typical hydrodynamic solvers in the literature assume constant γ. Our problem is to generalize the constant γ hydro solvers for the case of a spatially varying polytropic index, γ_e, which varies due to ionization, excitation and Coulomb interactions. A method that is applicable to all the aforementioned, constant-γ, hydrodynamic shock-capturing schemes is one of splitting the electron internal energy E_e density into the sum of an ideal (translational) energy part $p_e/(\gamma-1)$ and an extra internal energy density E_X. Similarly, we can define an ideal total energy density

$$e_I = \frac{\rho u^2}{2} + \frac{p_i + p_e}{\gamma - 1} + \sum_{g=1}^{G} E_g, \tag{26}$$

which is related to the total energy density by $e = e_I + E_X$. We time advance p_e with the ideal electron pressure equation and E_X by a conservative advection equation, and then apply a correction step as described below.

The time update with the operator $\mathbf{R}_{\text{hydro}}$ solves the following equations:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0, \tag{27}$$

$$\frac{\partial \rho \mathbf{u}}{\partial t} + \nabla \cdot [\rho \mathbf{u}\mathbf{u} + I(p + p_r)] = 0, \tag{28}$$

$$\frac{\partial e_I}{\partial t} + \nabla \cdot [(e_I + p + p_r)\mathbf{u}] = 0, \tag{29}$$

$$\frac{1}{\gamma - 1}\frac{\partial p_e}{\partial t} + \frac{1}{\gamma - 1}\nabla \cdot (p_e \mathbf{u}) + p_e \nabla \cdot u = 0, \tag{30}$$

$$\frac{\partial E_X}{\partial t} + \nabla \cdot [E_X \mathbf{u}] = 0, \tag{31}$$

$$\frac{\partial E_g}{\partial t} + \nabla \cdot (E_g \mathbf{u}) + (\gamma_r - 1) E_g \nabla \cdot u = 0, \tag{32}$$

$$\frac{\partial d_m}{\partial t} + \nabla \cdot (d_m \mathbf{u}) - d_m \nabla \cdot u = 0, \tag{33}$$

where the frequency advection, diffusion, and energy exchange terms are omitted in this first operator step. After each time advance from time t^n to time t^{n+1}, we have to correct e, e_I, p_e, and E_X. We denote the uncorrected variables with a superscript $*$, then we recover at time level $n+1$ the true electron internal energy E_e^{n+1} and the true total energy density e^{n+1} by

$$E_e^{n+1} = \frac{p_e^*}{\gamma - 1} + E_X^*, \tag{34}$$

$$e^{n+1} = e_I^* + E_X^*. \tag{35}$$

Since both e_I and E_X follow a conservation law, the total energy density e is also conserved. The true electron pressure is recovered from the updated electron internal energy and mass density by means of the EOS:

$$p_e^{n+1} = p_{\text{EOS}}(\rho^{n+1}, E_e^{n+1}, m), \tag{36}$$

where the function p_{EOS} can be either a calculated EOS or an EOS lookup table for material m, determined by the level set functions d_m^{n+1} (Section 3.1.3). The extra internal energy E_X is reset as the difference between the true electron internal energy and the ideal electron internal energy for $\gamma = 5/3$:

$$E_X^{n+1} = E_e^{n+1} - \frac{p_e^{n+1}}{\gamma - 1}. \tag{37}$$

This is positive because the EOS state p_{EOS} satisfies $E_e - p_e/(\gamma - 1) \geq 0$ at all times. The ideal part of the total energy density at time level $n+1$ can now be updated as

$$e_I^{n+1} = e^{n+1} - E_X^{n+1}. \tag{38}$$

We have now recovered e^{n+1}, ${e_I}^{n+1}$, p_e^{n+1}, and E_X^{n+1} at time t^{n+1}.

We time advance the hydrodynamic equations to the time level $*$ with a shock-capturing scheme with a constant $\gamma = 5/3$. For an ideal EOS, the speed of sound of the equations (27)–(32) can be derived as

$$c_s = \sqrt{\frac{\gamma(p_i + p_e) + \gamma_r p_r}{\rho}}, \tag{39}$$

which includes modifications due to the presence of the total radiation pressure. This speed of sound will be used in the hydro scheme below. Since the CRASH EOS solver always satisfies $E_X \geq 0$ and $\gamma_e \leq 5/3$, the speed of sound for the ideal EOS is always an upper bound for the true speed of sound.

We use shock-capturing schemes to advance the equations (27)–(33). In the following, we denote the (near) conservative variables by $U = (\rho, \rho\mathbf{u}, e_I, p_e, E_X, E_g, d_m)$ and let U be grid cell averages in the standard finite volume sense. If we assume for the moment a 1D grid with spacing Δx, cell center index i and cell face between cell i and $i+1$ identified by half indices $i+1/2$, then we can write the two-stage Runge-Kutta hydro update as

$$U_i^{n+1/2} = U_i^n - \frac{\Delta t}{2\Delta x}\left(f_{i+1/2}^n - f_{i-1/2}^n\right), \tag{40}$$

$$U_i^{n+1} = U_i^n - \frac{\Delta t}{\Delta x}\left(f_{i+1/2}^{n+1/2} - f_{i-1/2}^{n+1/2}\right). \tag{41}$$

where f is the numerical flux. In particular, the HLLE flux f equals the physical flux $F(U_{i+1/2}^R)$ when $c_s^+ = u_i + c_s \leq 0$, $F(U_{i+1/2}^L)$ when $c_s^- = u_i - c_s \geq 0$, and in all other cases it uses the weighted flux

$$f_{i+1/2} = \frac{c_s^+ F(U_{i+1/2}^L) - c_s^- F(U_{i+1/2}^R) + c_s^+ c_s^- (U_{i+1/2}^R - U_{i+1/2}^L)}{c_s^+ - c_s^-}. \tag{42}$$

Here, the left and right cell face states are

$$U_{i+1/2}^L = U_i + \frac{1}{2}\bar{\Delta}^L U_i, \tag{43}$$

$$U_{i+1/2}^R = U_{i+1} - \frac{1}{2}\bar{\Delta}^R U_{i+1}. \tag{44}$$

We use the generalized Koren limiter, and define the limited slopes as

$$\bar{\Delta}^L U_i = \text{minmod}\left[\beta(U_{i+1} - U_i), \beta(U_i - U_{i-1}), \frac{2U_{i+1} - U_i - U_{i-1}}{3}\right], \tag{45}$$

$$\bar{\Delta}^R U_i = \text{minmod}\left[\beta(U_{i+1} - U_i), \beta(U_i - U_{i-1}), \frac{U_{i+1} - U_i - 2U_{i-1}}{3}\right], \tag{46}$$

for the extrapolations from the left and right. This reconstruction can be third order in smooth regions away from extrema [Koren, 1993, Tóth et al., 2008]. The parameter β can be changed between 1 and 2, but in simulations with adaptive mesh refinement we have best experience with $\beta = 3/2$. We generally apply the slope limiters on the primitive variables $(\rho, \mathbf{u}, p_i, p_e, E_X/\rho, E_g, d_m)$, instead of the conservative variables. We apply the slope limiter on E_X/ρ, instead of E_X, since E_X/ρ is smoother at shocks and across material interfaces. A multi-dimensional update is obtained by adding the fluxes for each direction in a dimensionally unsplit manner.

After each stage of the two step Runge-Kutta, we correct for the EOS effects via the update procedure outlined in equations (34)–(38).

3.2.1.2 Non-Conservative Pressure Equations

In regions away from shocks, it is sometimes more important to preserve pressure balance than to have a shock capturing scheme that recovers the correct jump conditions. This is especially important at material interfaces. We therefore have implemented the option to solve the hydro part of the pressure equations

$$\frac{\partial p_i}{\partial t} + \nabla\cdot(p_i\mathbf{u}) + (\gamma - 1)p_i\nabla\cdot u = 0, \tag{47}$$

$$\frac{\partial p_e}{\partial t} + \nabla\cdot(p_e\mathbf{u}) + (\gamma_{Se} - 1)p_e\nabla\cdot u = 0, \tag{48}$$

instead of the equations for the total energy (29) and the electron internal energy (30). As long as the speed-of-sound gamma for the electrons

$$\gamma_{S_e} = \frac{\rho}{p_e}\left(\frac{\partial p_e}{\partial \rho}\right)_{S_e} \tag{49}$$

is smaller than $\gamma = 5/3$, the numerical scheme is stable. Contrary to the energy conserving scheme, the pressure-based scheme can directly include the EOS and we no longer need the time evolution of the extra internal energy density (31). The EOS contribution in the electron-pressure equation (48) is implemented as a source term $-(\gamma_{Se} - \gamma)p_e \nabla \cdot u$ added to the ideal electron pressure equation.

To facilitate using both the shock capturing properties and the pressure balance at the material interfaces during CRASH simulations, we have several criteria to switch between them automatically. One of the criteria, for instance, uses a detection of steep pressure gradients as a shock identification. The user can select the magnitude of the pressure gradient above which the scheme switches to the conservative energy equations.

3.2.2 Frequency Advection

The set of multigroup equations (12) contains an integral over the group photon frequencies. Performing this integration, the frequency advection update by the $\mathbf{R}_{\text{frequency}}$ operator can be written as

$$\frac{\partial E_g}{\partial t} - (\gamma_r - 1)(\nabla \cdot \mathbf{u})\left[\nu_{g+1/2} E_\nu(\nu_{g+1/2}) - \nu_{g-1/2} E_\nu(\nu_{g-1/2})\right] = 0. \tag{50}$$

These equations, however, still depend on the, as yet, unassigned photon group frequencies ν_g and the spectral radiation energy density E_ν. We will now restrict the analysis to a frequency grid that is uniformly spaced in the frequency logarithm, i.e.,

$$\ln(\nu_{g+1/2}) - \ln(\nu_{g-1/2}) = \Delta(\ln \nu) = \text{constant}. \tag{51}$$

For a large enough number of frequency groups G, the group energy E_g can then be approximated as the product of the photon frequency, spectral radiation energy E_ν, and the logarithmic group spacing $\Delta(\ln \nu)$:

$$E_g = \int_{\nu_{g-1/2}}^{\nu_{g+1/2}} E_\nu d\nu = \int_{\ln \nu_{g-1/2}}^{\ln \nu_{g+1/2}} E_\nu \nu d(\ln \nu) \approx E_\nu \nu \Delta(\ln \nu). \tag{52}$$

Using this approximation in equation (50), we obtain our final form of the frequency advection

$$\frac{\partial E_g}{\partial t} + u_\nu \frac{E_{g+1/2} - E_{g-1/2}}{\Delta(\ln \nu)} = 0, \tag{53}$$

where $u_\nu = -(\gamma_r - 1)\nabla \cdot u$ is the frequency advection speed. The values $E_{g\pm 1/2}$ are to be interpolated from the mesh-centered values E_g towards the group boundaries.

The frequency advection is a conservative linear advection in the log-frequency coordinate, for which the physical flux is defined as $F_{g-1/2} = u_\nu E_{g-1/2}$. For the boundary conditions in the frequency domain we assume zero radiation flux so that no radiation can leak at the edges of the frequency domain. Equation (53) can be discretized with the one-stage second order upwind scheme

$$E_g^{n+1} = E_g^* - \Delta t \frac{f_{g+1/2} - f_{g-1/2}}{\Delta(\ln \nu)}, \tag{54}$$

where time level $*$ is now the state after the hydro update and the numerical flux is

$$\begin{aligned} f_{g-1/2} &= u_\nu \left[E_g - \frac{1-\mathrm{C}}{2} \bar{\Delta}(E_{g+1}-E_g, E_g - E_{g-1}) \right], \qquad u_\nu \le 0, \\ f_{g-1/2} &= u_\nu \left[E_{g-1} + \frac{1-\mathrm{C}}{2} \bar{\Delta}(E_g - E_{g-1}, E_{g-1} - E_{g-2}) \right], \qquad u_\nu \ge 0. \end{aligned} \tag{55}$$

and we use the superbee limiter [Roe, 1986] for the limited slope $\bar{\Delta}$. The Courant-Friedrichs-Lewy number $\mathrm{C} = |u_\nu|\Delta t/\Delta(\ln \nu)$ depends on the hydrodynamic time-step Δt. If C is larger than one, the frequency advection is subcycled with the number of steps equal to the smallest integer value larger than C.

3.2.3 Implicit Diffusion and Energy Exchange

The stiff parts of the radhyro equations are solved implicitly in an operator split fashion. These stiff parts are the radiation energy diffusion, electron heat conduction, and the energy exchange between the electrons and each energy group g and between the electrons and ions. In this section, we will describe two implicit schemes that are implemented: (1) solving all radiation groups, together with electron and ion temperatures in a coupled manner (Section 3.2.3.1), and (2) solving each radiation group energy and the electron temperature independently (Section 3.2.3.2). Our strategy for resolution changes is described in section 3.2.6, while the modifications for the rz-geometry are explained in section 3.2.5.

3.2.3.1 Coupled Implicit Scheme

Discretizing the diffusion and energy exchange terms of equations (17)–(18), and (12) implicitly in time leads to

$$\frac{E_i^{n+1} - E_i^*}{\Delta t} = \sigma_{ie}^* (T_e^{n+1} - T_i^{n+1}), \tag{56}$$

$$\frac{E_e^{n+1} - E_e^*}{\Delta t} = \sigma_{ie}^* (T_i^{n+1} - T_e^{n+1}) + \nabla \cdot C_e^* \nabla T_e^{n+1} + \sum_{g=1}^{G} \sigma_g^* (E_g^{n+1} - B_g^{n+1}), \tag{57}$$

$$\frac{E_g^{n+1} - E_g^*}{\Delta t} = \sigma_g^* (B_g^{n+1} - E_g^{n+1}) + \nabla \cdot D_g^* \nabla E_g^{n+1}, \tag{58}$$

where time level $*$ now corresponds to the state after the hydro update and the frequency advection. The coupling coefficients σ_{ie}^* and σ_g^* and the diffusion coefficients C_e^* and D_g^* are taken at time level $*$ (frozen coefficients). One can either (1) solve the coupled system of $G+2$ equations (56)–(58) implicitly or (2) solve equation (56) for the ion internal energy E_i^{n+1}, substitute the solution back into equation (57), and solve the resulting reduced set of $G+1$ equations (57)–(58) implicitly. Here we describe the second scheme, because it is more efficient, especially for a small number of groups, e.g., for gray radiation diffusion. Note that if we had included ion heat conduction in (58), then we would have to solve the entire coupled system of equations.

First, we introduce the ion Planck function $B_i = aT_i^4$ as a new variable similar to the electron Planck function $B = aT_e^4$, and replace E_i and E_e with these variables using the chain rule

$$\frac{\partial E_i}{\partial t} = \frac{\partial E_i}{\partial T_i}\frac{\partial T_i}{\partial B_i}\frac{\partial B_i}{\partial t} = \frac{C_{Vi}}{4aT_i^3}\frac{\partial B_i}{\partial t}, \qquad \frac{\partial E_e}{\partial t} = \frac{C_{Ve}}{4aT_e^3}\frac{\partial B}{\partial t}, \tag{59}$$

where C_{Vi} and C_{Ve} are the specific heats of the ions and electrons, respectively. Now equation (56) can be replaced with

$$B_i^{n+1} = B_i^* + \Delta t \sigma'_{ie}(B^{n+1} - B_i^{n+1}), \tag{60}$$

where

$$\sigma'_{ie} = \sigma^*_{ie} \frac{4aT_i^3}{C_{Vi}} \frac{1}{a(T_e+T_i)(T_e^2+T_i^2)}, \tag{61}$$

is again taken at time level $*$. The numerator comes from $(T_e^4 - T_i^4)/(T_e - T_i)$. Equation (60) can be solved for B_i^{n+1}. This result can be substituted into the electron internal energy equation (57) to obtain

$$\frac{C'_{Ve}}{\Delta t}(B^{n+1} - B^*) = \sigma''_{ie}(B_i^* - B^{n+1}) + \nabla \cdot C'_e \nabla B^{n+1} + \sum_{g=1}^{G} \sigma_g^*(E_g^{n+1} - w_g^* B^{n+1}), \tag{62}$$

where we have introduced new coefficients at time level $*$:

$$\sigma''_{ie} = \frac{C_{Vi}}{4aT_i^3} \frac{\sigma'_{ie}}{1+\Delta t \sigma'_{ie}}, \qquad C'_e = \frac{C_e^*}{4aT_e^3}, \qquad C'_{Ve} = \frac{C_{Ve}^*}{4aT_e^3}. \tag{63}$$

The Planck weights $w_g^* = B_g^*/B^*$ satisfy $\sum_g w_g = 1$. It is convenient to introduce the changes $\Delta B = B^{n+1} - B^*$ and $\Delta E_g = E_g^{n+1} - E_g^*$ to arrive at

$$\begin{aligned} \left[\frac{C'_{Ve}}{\Delta t} + \sigma''_{ie} - \nabla \cdot C'_e \nabla\right] \Delta B \quad &- \quad \sum_{g=1}^{G} \sigma_g^*(\Delta E_g - w_g^* \Delta B) = \sigma''_{ie}(B_i^* - B^*) \\ &+ \quad \nabla \cdot C'_e \nabla B^* + \sum_{g=1}^{G} \sigma_g^*(E_g^* - w_g^* B^*), \end{aligned} \tag{64}$$

$$\left[\frac{1}{\Delta t} - \nabla \cdot D_g^* \nabla\right] \Delta E_g \quad - \quad \sigma_g^*(w_g^* \Delta B - \Delta E_g) = \sigma_g^*(w_g^* B^* - E_g^*) + \nabla \cdot D_g^* \nabla E_g^*. \tag{65}$$

This is a coupled system of $G+1$ linearized equations for the changes ΔB and ΔE_g. The right hand sides are all at time level *.

A discrete set of equations is obtained by applying the standard finite volume method to the equations (64) and (65) and partitioning the domain in a set of control volumes V_i, enumerated by a single index $i = 1, \ldots, I$. As an example, the fluxes F_{gij} associated with the radiation diffusion operator may be obtained by approximating the gradient of the group energy density with a simple central difference in the uniform part of the mesh:

$$-\int_{V_i} \nabla \cdot (D_g \nabla E_g) dV = \sum_j F_{gij} = \sum_j S_{ij} D_{gij} \frac{E_{gi} - E_{gj}}{|\mathbf{x}_i - \mathbf{x}_j|}, \tag{66}$$

where the index j enumerates the control volumes which have a common face with the control volume i, the face area being S_{ij}, and the distance between cell centers is $|\mathbf{x}_i - \mathbf{x}_j|$. Note that we assumed here an orthogonal mesh. Generalization to curvilinear grids can be done as shown in Tóth et al. [2008]. The diffusion coefficients at a face are obtained by simple averaging of the cell centered diffusion coefficient: $D_{gij} = (D_{gi} + D_{gj})/2$. The discretization of the diffusion operator at resolution changes is described in section 3.2.6.

The linear system (64)–(65) can be written in a more compact form as the linearized implicit backward Euler scheme

$$\left(\mathbf{I} - \Delta t \frac{\partial \mathbf{R}}{\partial \mathbf{U}}\right) \Delta \mathbf{U} = \Delta t \mathbf{R}(\mathbf{U}^*), \tag{67}$$

where $\mathbf{U}$ are the $I \times (G+1)$ state variables B and E_g for all I control volumes, and $\Delta\mathbf{U} = \mathbf{U}^{n+1} - \mathbf{U}^*$. $\mathbf{R}$ is defined by the spatially discretized version of the right hand side of equations (64) and (65). The matrix $\mathbf{A} = \mathbf{I} - \Delta t \partial\mathbf{R}/\partial\mathbf{U}$ is a $I \times I$ block matrix consisting of $(G+1) \times (G+1)$ sub-matrices. This matrix $\mathbf{A}$ is in general non-symmetric due to the Planck weight w_g^* in the energy exchange between the radiation and electrons. To solve this system we use Krylov sub-space type iterative solvers, like GMRES [Saad and Schultz, 1986] or Bi-CGSTAB [Van der Vorst, 1992]. To accelerate the convergence of the iterative scheme, we use a preconditioner. In the current implementation of CRASH, we use the Block Incomplete Lower-Upper decomposition (BILU) preconditioner, which is applied for each adaptive mesh refinement block independently. For gray radiation diffusion the Planck weight is one, and the matrix $\mathbf{A}$ can be proven to be symmetric positive definite (SPD) for commonly used boundary conditions (see for example Edwards [1996b]). In that case, we can use a preconditioned conjugate gradient (PCG) scheme (see for instance Eisenstat [1981]).

For some verification tests, we can attempt to go second order in time under the assumption of temporally constant coefficients using the Crank-Nicolson scheme

$$\frac{\mathbf{U}^{n+1} - \mathbf{U}^*}{\Delta t} = (1-\alpha)\mathbf{R}(\mathbf{U}^*) + \alpha\mathbf{R}(\mathbf{U}^{n+1}), \tag{68}$$

with $\alpha = 1/2$. The implicit residual can again be linearized $\mathbf{R}(\mathbf{U}^{n+1}) = \mathbf{R}(\mathbf{U}^*) + (\partial\mathbf{R}/\partial\mathbf{U})^*\Delta\mathbf{U}$ to obtain the linear system of equations

$$\left(\mathbf{I} - \alpha\Delta t \frac{\partial\mathbf{R}}{\partial\mathbf{U}}\right)\Delta\mathbf{U} = \Delta t\mathbf{R}(\mathbf{U}^*). \tag{69}$$

We use the same iterative solvers as for the backward Euler scheme.

Finally, we show how we use the solution ΔB and ΔE_g for $g = 1, \ldots, G$ from the non-conservative equations (64) and (65) to advance the solution of the original equations (56)–(58) and still conserve the total energy. One needs to express the fluxes and energies on the right hand side in the latter equations in terms of B^{n+1} and E_g^{n+1} while still keeping the coefficients frozen. After some algebra we obtain

$$E_i^{n+1} = E_i^* + \Delta t\sigma_{ie}''(B^{n+1} - B_i^*), \tag{70}$$

$$E_e^{n+1} = E_e^* + C_{Ve}'(B^{n+1} - B^*), \tag{71}$$

$$E_g^{n+1} = E_g^* + \Delta E_g. \tag{72}$$

This update conserves the total energy to round-off error. Note that at this final stage, taking too large time step may result in negative ion internal energy E_i^{n+1} if $B^{n+1} \ll B_i^*$ and negative electron internal energy E_e^{n+1} if $B^{n+1} \ll B^*$. If this happens, the advance might be redone with a smaller time step, to limit the drop in B, or by some other timestep control scheme. A generalization of the conservative update to the Crank-Nicolson scheme is also implemented for verification tests with time constant coefficients.

For completeness, we mention that in the absence of radiation we solve during the implicit step for the temperatures T_e and T_i instead of the radiation energy-like variables aT_e^4 and aT_i^4. In that case the corresponding matrix $\mathbf{A}$ is always SPD. In principle, the formulation in temperatures can be generalized to radiation as well. In Landau and Lifshitz [1980], a spectral temperature $T_\nu(E_\nu, \nu)$ is defined, such that the spectral energy density is locally equal to the spectral Planckian energy density at the temperature T_ν: $E_\nu = B_\nu(T_\nu, \nu)$. This relationship is a one-to-one map. A group temperature, T_g, can also be introduced as the discrete analog such that the group energy density can be obtained by

$$E_g(T_g) = \int_{\nu_{g-1/2}}^{\nu_{g+1/2}} B_\nu(T_g, \nu)d\nu. \tag{73}$$

The equation (58) can be recast as equation for the group temperature T_g. This introduces the group specific heat of the radiation $C_g = dE_g/dT_g$. The set of equations (56)–(58) reformulated as an implicit backward Euler scheme for the temperatures T_i, T_e, and T_g can in a similar way as in Edwards [1996b] be proven to be SPD. While this scheme has the advantage of being SPD, the conservative update of the group energy density $E_g^{n+1} = E_g^* + C_g^* \Delta T_g$ might result in negative energy density E_g^{n+1} for time steps that are too large.

3.2.3.2 Decoupled Implicit Scheme

The coupled implicit scheme of Section 3.2.3.1 requires solution of a large system of equations ($G+1$ variables per mesh cell). The preconditioning of such a system can be computationally expensive and requires overall a lot of memory. We therefore also implemented a decoupled implicit scheme that solves each equation independently.

For some applications, the electron temperature does not change much in exchanging energy with the radiation. This is typically so if the electrons have a much larger energy density than the radiation, so that T_e changes little due to interaction with the radiation in a single time step. In that case, we first solve for the electron and ion temperatures without the contributions from the radiation-electron energy exchange. Let again time level $*$ indicate the state after the hydro update and frequency advection, and again freeze C_e^*, D_g^*, σ_{ie}^*, σ_g^* at time level $*$. Discretization in time now leads to

$$\frac{E_i^{n+1} - E_i^*}{\Delta t} = \sigma_{ie}^*(T_e^{**} - T_i^{n+1}), \tag{74}$$

$$\frac{E_e^{**} - E_e^*}{\Delta t} = \sigma_{ie}^*(T_i^{n+1} - T_e^{**}) + \nabla \cdot C_e^* \nabla T_e^{**}, \tag{75}$$

where the time level $**$ of E_e indicates that we still have to do an extra update to time level $n+1$ with the radiation-electron energy exchange. Each radiation group energy density is solved independently using time level $*$ for the electron temperature in B_g^*:

$$\frac{E_g^{n+1} - E_g^*}{\Delta t} = \sigma_g^*(B_g^* - E_g^{n+1}) + \nabla \cdot D_g^* \nabla E_g^{n+1}, \tag{76}$$

where we have exploited the assumption that B_g^* is not stiff.

Equations (74)–(76) can be recast in equations for the $G+1$ independent changes $\Delta B = B^{**} - B^*$ and $\Delta E_g = E_g^{n+1} - E_g^*$:

$$\left[\frac{C'_{Ve}}{\Delta t} + \sigma''_{ie} - \nabla \cdot C'_e \nabla\right] \Delta B = \sigma''_{ie}(B_i^* - B^*) + \nabla \cdot C'_e \nabla B^*, \tag{77}$$

$$\left[\frac{1}{\Delta t} + \sigma_g^* - \nabla \cdot D_g^* \nabla\right] \Delta E_g = \sigma_g^*(w_g^* B^* - E_g^*) + \nabla \cdot D_g^* \nabla E_g^*. \tag{78}$$

where we have used the definitions (59), (61), and (63) of the coefficients, frozen at time level $*$. Each equation for the changes is in the form of the linearized implicit backward Euler scheme (67) and can be solved independently with iterative solvers like GMRES and Bi-CGSTAB using a BILU preconditioner. As long as the boundary conditions are such that the matrices are symmetric and positive definite, a preconditioned conjugate gradient method might also be used.

In a manner similar to the coupled implicit scheme, a conservative update for the energy densities can be derived

as

$$E_g^{n+1} = E_g^* + \Delta E_g, \tag{79}$$

$$E_i^{n+1} = E_i^* + \Delta t \sigma_{ie}''(B^{**} - B_i^*), \tag{80}$$

$$E_e^{n+1} = E_e^* + C_{Ve}'(B^{**} - B^*) + \Delta t \sum_{g=1}^{G} \sigma_g^*(E_g^{n+1} - w_g^* B^*), \tag{81}$$

which preserves the total energy to round-off errors. The main difference between the conservative update in the coupled and decoupled schemes is that here the energy exchange between the radiation and electrons is added afterwards as the last term in equation (81).

This scheme requires less computational time for preconditioning and for the Krylov solver than the coupled implicit algorithm. However it generally needs more message passing in parallel computations. It is therefore not always guaranteed that the decoupled scheme is faster. The memory usage is always smaller.

3.2.4 Boundary Conditions

The CRASH code allows for any user specified type of boundary conditions. Several commonly used boundary conditions are readily available in the main code for convenience, e.g., fixed, extrapolation with zero gradient, periodic, and reflective.

For the radiation field, we have implemented a zero or fixed incoming flux boundary condition that is used instead of the extrapolation with zero gradient. This type of boundary condition is useful if there are no sources of radiation outside the computational domain and we assume that outflowing radiation does not return back into the computational domain (zero albedo). Note that simple extrapolation with zero gradient can make the radiation diffusion problem ill-posed. The boundary condition is derived as follows: Radiation diffusion approximation corresponds to a linear-in-angle intensity distribution

$$I_g = \frac{c}{4\pi} E_g + \frac{3}{4\pi} \mathbf{F}_g \cdot \mathbf{n}, \tag{82}$$

so we can calculate the radiation flux through a boundary surface. If we define the outward pointing normal vector of the boundary as $\mathbf{n}_b$, the net flux of radiation energy inward through this boundary is

$$F_g^{\text{in}} = -\int_{\mathbf{n}\cdot\mathbf{n}_b<0} \mathbf{n}_b \cdot \mathbf{n} I_g d\Omega = \frac{cE_g}{4} - \frac{1}{2} \mathbf{n}_b \cdot \mathbf{F}_g, \tag{83}$$

where the closure (82) is used. In the radiation diffusion model, the flux is written as $\mathbf{F}_g = -D_g \nabla E_g$, where the diffusion coefficient D_g is a nonlinear function of E_g and ∇E_g in a flux limited diffusion model. The boundary condition satisfies

$$E_g + \frac{2D_g}{c} \mathbf{n}_b \cdot \nabla E_g = \frac{4}{c} F_g^{in}. \tag{84}$$

For the left boundary in the x-direction, for instance, this can be discretized as

$$\frac{E_{g0} + E_{g1}}{2} - \frac{2D_g}{c} \frac{E_{g1} - E_{g0}}{\Delta x} = \frac{4}{c} F_g^{in}, \tag{85}$$

where the index 1 corresponds to the last physical cell and 0 to the ghost cell. This equation can be solved for the ghost cell value. For zero incoming radiation flux boundary conditions we set $F_g^{in} = 0$.

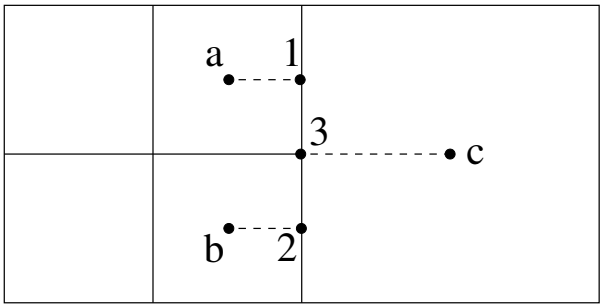

Figure 23: Cell and face centers at the adaptive interface in 2D.

3.2.5 R-Z-Geometry

Incorporating rz-geometry in a finite volume formulation is as follows: the radial cell face area and the cell volume must be made proportional to the distance r from the symmetry axis. In addition, the r component of the momentum equation (16) is modified as

$$\frac{\partial \rho u_r}{\partial t} + \nabla \cdot [\rho \mathbf{u} u_r + \hat{\mathbf{r}}(p + p_r)] = \frac{p + p_r}{r}, \tag{86}$$

where $\hat{\mathbf{r}}$ is the unit vector in the r direction and $u_r = \mathbf{u} \cdot \hat{\mathbf{r}}$. This correction reflects that the pressure term is a gradient, not a divergence.

3.2.6 Discretization of the diffusion operator at resolution changes

In Sections 3.2.3.1 and 3.2.3.2, the diffusion operator is discretized on a uniform mesh with a standard finite volume method in combination with a central difference approximation for the gradient in the flux calculation as in equation (66). The diffusion coefficient needed at the face is obtained by simple arithmetic averaging of the left and right cell center diffusion coefficients. The generalization to resolution changes as in Figure 23 is less straightforward. In the following, we denote the fine cell centers by a and b, the coarse cell center by c. The flux densities at resolution changes in the direction orthogonal to the interface are denoted by F_1 and F_2 at the fine faces, and F_3 at the coarse face.

In Edwards [1996b], a strategy was developed to discretize the diffusion operator on an adaptive mesh in the context of reservoir simulations. The main ingredients of the method are (1) require the continuity of the flux at the resolution change in the strong sense, i.e. $F_1 = F_2 = F_3$, and (2) discretize the gradient in the diffusion flux by a one-sided difference. An expression was found for the diffusion flux $\mathbf{F} = -D\nabla E$ in which the diffusion coefficient is replaced by a weighted harmonic average of the cell centered values D_a, D_b, D_c. In Gittings [2008], it was argued that this discretization does not properly propagate the self-similar Marshak waves of the radiation diffusion model, unless the cell centered diffusion coefficients are calculated on a common face temperature.

In the version of the code discussed here, we follow a different approach that replaces the harmonic average of the diffusion coefficient in Edwards [1996b] by an arithmetic average and obtain for the flux densities normal to the resolution change interface

$$F_1 = F_2 = F_3 = -\frac{2D}{3\Delta x}\left[E_c - (E_a + E_b)/2\right], \tag{87}$$

where Δx is the fine cell size and the diffusion coefficient at the face is averaged as

$$D_3 = [D_c + (D_a + D_b)]/3. \tag{88}$$

We have demonstrated with verification tests including those discussed above that this change produces properly propagating radiative precursor and shock fronts. Generalizations to 1D and 3D are straightforward.

3.2.6.1 Improved diffusion operator at resolution changes

In this section we present an improved conservative and spatially second-order implicit scheme for the radiation diffusion and heat conduction. For convenience, we only derive this scheme for the heat conduction. The generalization for radiation diffusion is straightforward.

Discretizing the electron thermal heat conduction implicitly in time leads to the linearized backward Euler equation for the electron temperature T

$$C_{Ve}^* \frac{T^{n+1} - T^*}{\Delta t} = \nabla \cdot C_e^* \nabla T^{n+1}, \tag{89}$$

where C_{Ve} is the electron specific heat and C_e is the heat conduction coeffcient. The time level $*$ corresponds to the state before the implicit update. During the implicit advance with time step Δt, the coefficients are frozen in at time level $*$, resulting in a temporally first order scheme in general. This equation can be recast in a linearized equation for the change $\Delta T = T^{n+1} - T^*$:

$$\left[\frac{C_{Ve}^*}{\Delta t} - \nabla \cdot C_e^* \nabla\right] \Delta T = \nabla \cdot C_e^* \nabla T^*. \tag{90}$$

The right-hand-side depends only on time level $*$. Once Eq. 90 is solved using a linear solver, the electron energy density can be updated using

$$E_e^{n+1} = E_e^* + C_{Ve}^* (T^{n+1} - T^*), \tag{91}$$

to conserve the energy.

A discrete set of equations is obtained by applying a finite volume method to Eq. 90. To make this scheme spatially second-order accurate on a uniform mesh, we need a second order accurate thermal heat flux at the face centers. This is achieved by approximating the gradient of the electron temperature with a central difference using the cell-centered values

$$-\int_{V_i} \nabla \cdot (C_e \nabla T) dV = \sum_j S_{ij} C_{eij} \frac{T_i - T_j}{|\mathbf{x}_i - \mathbf{x}_j|}, \tag{92}$$

where the control volumes are indexed by i, each having a volume V_i. The index j is for the neighboring cells having a common interface with area S_{ij} and a distance between the cell centers is $|\mathbf{x}_i - \mathbf{x}_j|$. The heat conduction coefficient at the face is the arithmetic average of the coefficient at the two neighboring cell centers, $C_{eij} = (C_{ei} + C_{ej})/2$.

To obtain a second-order heat flux at the resolution change, we need a third-order interpolation of the temperature in the ghost cells. Such an interpolation was previously used in the context of Hall magnetohydrodynamics (MHD) [Tóth et al., 2008]. This interpolation is only needed for the fine cells, since the flux at the coarse side will be obtained as the sum of the fluxes at the neighboring fine cells to preserve conservation of the scheme [Berger and Colella, 1989]. The implementation for the heat conduction is different from the Hall MHD, since for the heat conduction we also have to maintain positivity of the temperature and avoid spurious oscillations. For convenience we will restrict the analysis to two-dimensional domains. The third-order interpolation for the temperature value at the fine ghost cell

$(0, j)$, indicated by the dashed circle, is first performed along the coarse cell values in the transverse direction to obtain the temperature at $(-1/2, j)$ as depicted in Fig. 24:

$$T_{-1/2,j} = \frac{5T_{-1/2,j-3/2} + 30T_{-1/2,j+1/2} - 3T_{-1/2,j+5/2}}{32}. \tag{93}$$

To guarantee the positivity of the interpolated temperature, the value is clipped by the maximum and minimum values of the surrounding points

$$T_{-1/2,j} = \max\{\min[T_{-1/2,j}, \quad \max(T_{-1/2,j-3/2}, T_{-1/2,j+1/2})], \tag{94}$$

$$\min(T_{-1/2,j-3/2}, T_{-1/2,j+1/2})\}. \tag{95}$$

The value in the fine ghost cell can now be obtained by a parabolic interpolation along the fine cells in the direction normal to the refinement interface

$$T_{0,j} = \frac{8T_{-1/2,j} + 10T_{1,j} - 3T_{2,j}}{15}. \tag{96}$$

For this interpolation we need again to clip the obtained value with the surrounding temperatures

$$T_{0,j} = \max\{\min[T_{0,j}, \max(T_{-1/2,j}, T_{1,j})], \min(T_{-1/2,j}, T_{1,j})\}. \tag{97}$$

This ghost cell value is used in the heat conduction formula 92. We still need a second order accurate heat conduction coefficient C_{eij} at the fine face center. This amounts to a second-order prolongation of the heat conduction coefficient to obtain the ghost-cell values at the fine side, followed by an averaging of the cell centered coefficients to the face center. Once the thermal heat fluxes at the fine side are obtained, copying the fine fluxes to the coarse side at the resolution changes restores conservation. Note that this scheme is different from Edwards [1996b] where conservation of the flux at the resolution changes is enforced in the strong sense by enforcing the flux on the coarse side to be equal to each of the fluxes on the fine side.

In the analysis we assumed a Cartesian mesh. Generalization to curvilinear grids is presented in Tóth et al. [2008]. The same third-order interpolation procedure with clipping is also used for heat conduction along the magnetic field lines in the context of solar wind modeling [van der Holst et al., 2010].

We verified the improved implicit heat conduction and radiation solver with the tests described in van der Holst et al. [2011] and find that the solver converges to the analytical solutions and does not under- or over-shoot near discontinuities.

3.3 Parallel Performance of the Baseline CRASH code

We present parallel scaling studies on the Pleiades supercomputer at NASA Ames. This computer is an SGI ICE cluster connected with InfiniBand. Figure 25 shows strong scaling for a problem size that is independent of the number of processors. This 3D simulation is a circular tube version of the full system test described in Section 4.1.4. It uses five materials, 30 radiation groups, and separate electron and ion temperatures. The grid contains $80 \times 8 \times 8$ blocks of $4 \times 4 \times 4$ cells each at the base level and in addition two time-dependent refinement levels. There are overall approximately 2.6 million cells in this problem. We use lookup tables for the EOS and opacities, so that the computational time for that is negligible. For the hydrodynamic equations, we use the HLLE scheme together with the generalized Koren limiter with $\beta = 3/2$. The radiation diffusion, electron heat conduction and energy exchange terms

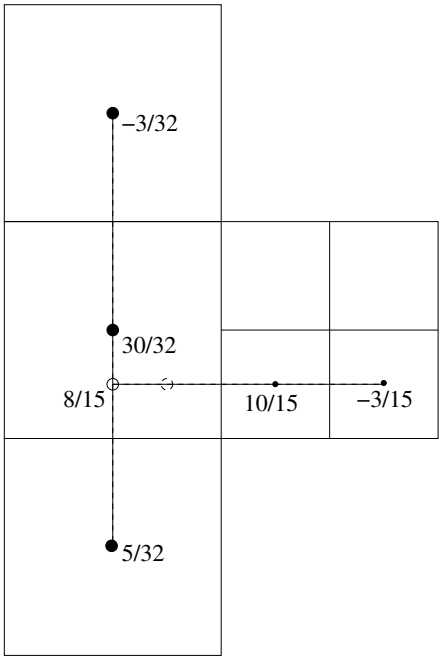

Figure 24: The coefficients of the third-order interpolation near the resolution changes (similar to Tóth et al. [2008]). The order of the interpolation is first along the coarse grid cells tangential to the resolution change, followed by an interpolation along the fine cells to obtain a third-order temperature value in the fine ghost cell indicated by the dashed circle.

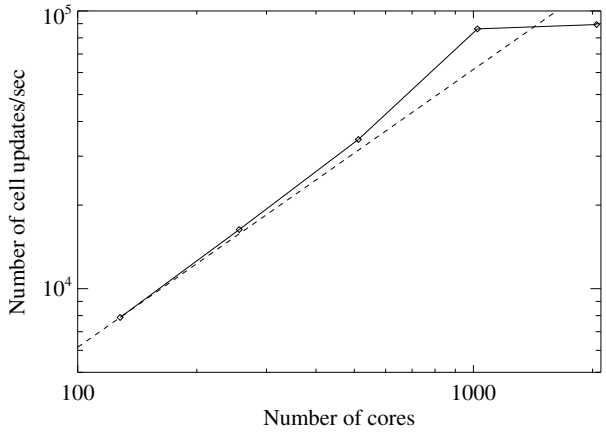

Figure 25: Strong scaling of the CRASH code, running a 3D CRASH application with 5 material level sets, electron and ion temperature, 30 radiation groups, and two levels of time dependent mesh refinement.

are solved implicitly with the decoupled scheme, using the Bi-CGSTAB iterative solver. This simulation is performed for 20 time steps for the number of cores varying from 128 to 2048, but excludes file I/O to measure the performance of the implicit solver. Up to 1024 cores, we get good scaling. However, for more cores we observe saturation in the performance.

3.4 Simulating radiative shocks in nozzle shock tubes using the baseline CRASH code

Before we developed our own laser package to model the laser-plasma physics and laser energy deposition in the radiative shock tube experiments, we used H2D, the 2D version of the Hyades code [Larsen and Lane, 1994], which contains a built-in laser package. H2D is a Lagrangian radiation-hydrodynamics code that utilizes the axisymmetry. Hyades is capable of tracing rays in 3D; the runs shown below used 2D ray tracing for the laser energy deposition. We used the H2D code to simulate our experiments for the first 1.1 ns (sometimes up to 1.3 ns), the time of the full width half maximum (FWHM) 1 ns laser pulse including ramp-up and ramp-down time. The spatial profile of the beam is determined using a model for the irradiance pattern of the laser beams in typical experiments. At 1.1 ns the experiment is in a regime that is well described by radiation-hydrodynamics so that the simulation can be continued with the CRASH code [van der Holst et al., 2011].

Section 3.4.1 describes the radiative shocks produced in nozzles with circular cross-section. We contrast the results with order-of-magnitude estimates based on physical arguments. Section 3.4.2 demonstrates the radiative shocks in nozzles with an elliptical neck. This simulation is fully three-dimensional.

3.4.1 Circular nozzle

A laser pulse irradiates a $20\,\mu$m thick beryllium disk with $0.35\,\mu$m wavelength light with a FWHM duration of 1 ns and with a laser energy deposition of 4 kJ. For the laser spot size we use a FWHM $800\,\mu$m diameter which is smaller than the $1200\,\mu$m diameter of the tube. Tthe runs reported here use the Lagrangian radiation-hydrodynamics code Hyades 2D (H2D) [Larsen and Lane, 1994] to evaluate the laser energy deposition during the first 1.1 ns which includes the laser ramp-up and ramp-down time. While our aim is to simulate radiative shocks in nozzles, the first 1.1 ns is however simulated in a straight tube for convenience. Calculating the first 1.1 ns with a straight tube and then transforming this tube into a nozzle is physically justified as long as the taper and shaft of the nozzle only alters the radiative precursor. The justification originates from the observation that the radiative transport from the precursor back through the shock front is negligible [Drake et al., 2011]. The straight tube does have a cylindrical polyimide wall of $100\,\mu$m thickness and inner radius of $600\,\mu$m filled with xenon with mass density $\rho = 0.0065\,\mathrm{g/cm^3}$, while there is vacuum outside. The beryllium disk is immediately to the left of $x = 0$, where x is the coordinate along the tube. The laser light will come in from the negative x direction. We place a gold washer next to the beryllium disk to protect the outside of the plastic tube from the laser light and use acrylic in between the gold and polyimide tube.

After 1.1 ns of simulation time the output of H2D is used to initialize the Eulerian CRASH code as described in Sec. 3.5.3. The straight tube of H2D is transformed into a nozzle, see the left panel of Fig. 26 for one quarter of the nozzle domain. This nozzle changes cross-section in the following way. For $x < 500\,\mu$m we do not modify the tube as defined in H2D. For $x > 750\,\mu$m we shrink the tube diameter from $1200\,\mu$m to $600\,\mu$m and the polyimide wall thickness correspondingly reduces to $50\,\mu$m. Between $500\,\mu$m and $750\,\mu$m the tube diameter and wall decrease linearly (In the notation of Sec. 3.5.3 we use $x_0 = 500\,\mu$m, $x_1 = 750\,\mu$m, $\varepsilon_y = 1/2$ and $\varepsilon_z = 1/2$). The vacuum outside

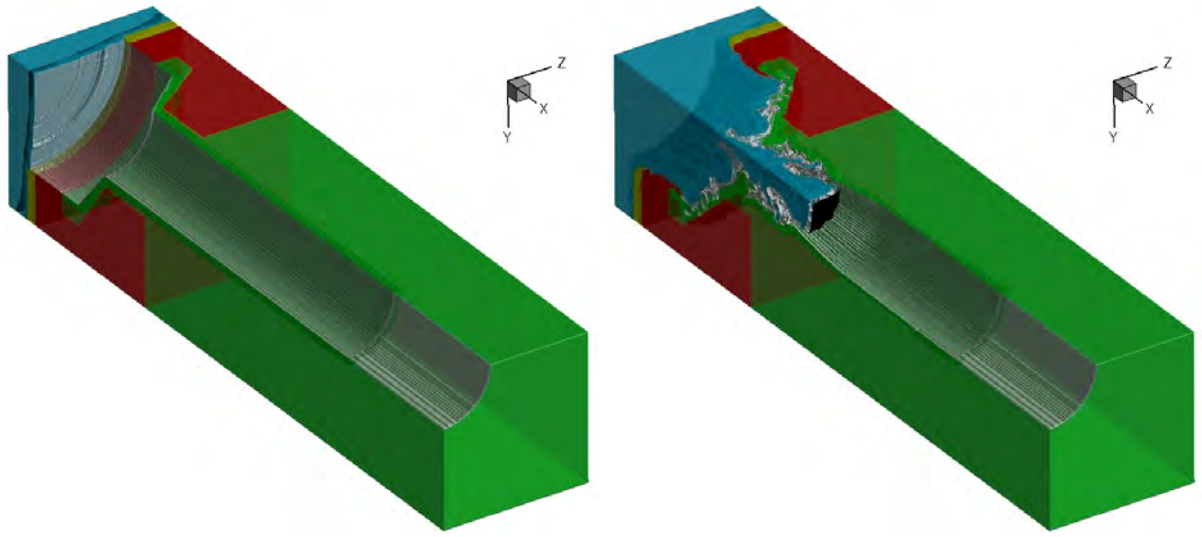

Figure 26: The materials in the 3D circular nozzle experiment at 1.1 ns (left panel) and 13 ns (right panels). Only one quarter of the shock tube is shown. The materials are beryllium (blue), polyimide (green), acrylic (red) and gold (yellow). The xenon inside the tube is not colored, but the border of the xenon volume is indicated by the white surface. The primary shock at 13 ns is depicted by a black iso-surface.

the nozzle is replaced with low density polyimide to avoid interactions with a zero mass density. The left panel of Fig. 26 shows the materials that are present in the simulation in color: beryllium (blue), polyimide (green), acrylic (red), and gold (yellow). The xenon inside the nozzle is, for convenience, not shown in this figure.

The computational domain size is $-150 < x < 3900$, $0 < y < 900$ and $0 < z < 900$ in microns, where y and z are the two directions transverse to the nozzle. The simulation is performed with an effective resolution of $2560 \times 512 \times 512$ grid cells using two levels of refinement. The effective cell sizes are therefore approximately $1.6\,\mu$m along the tube and $1.8\,\mu$m in the two transverse directions. The domain is decomposed in $4 \times 4 \times 4$ grid blocks. The mesh is refined at all interfaces that involve xenon or gold. To capture the shock front and the cooling layer, the mesh is also refined where the xenon density exceeds $0.02\,\mathrm{g/cm^3}$. Grid blocks are also refined if these criteria are satisfied in the ghost cells.

We used for the hydrodynamic part of the equations the HLLE scheme with a Courant-Friedrichs-Lewy number of 0.8 and the generalized Koren limiter with $\beta = 3/2$. For the radiation, we used the multi-group flux-limited diffusion model with 30 groups. The photon energy range is 0.1 eV to 20 keV that is logarithmically distributed over the groups. The radiation diffusion, heat conduction and energy exchanges are solved with the split (decoupled) implicit solver of van der Holst et al. [2011] using the conjugate gradient method with a Schwarz-type Incomplete Upper-Lower (ILU) preconditioner.

Due to the symmetry in the problem we only need to simulate one quarter of the nozzle domain ($y > 0$ and $z > 0$) and use reflective boundary conditions at $y = 0$ and $z = 0$. For all other boundaries of the domain we use extrapolation with zero gradient. For the radiation we use zero albedo boundary conditions (radiation propagating out of the computational domain does not return back).

We simulated the shock evolution from 1.1 ns to 13 ns physical time. It took a little over three days to compute on 1000 cores of the HERA supercomputer at the Lawrence Livermore National Laboratory. The number of cells in the computational domain increased from 26.5 million at the beginning to about 38 million near the end. The 3D material identity at 13 ns is shown in the right panel of Fig. 26. The beryllium has moved into the nozzle and is like a piston driving a shock in the xenon. This shock is indicated by a black iso-surface (of high ion temperature values). The xenon itself is not shown but the edge of the xenon is emphasized by a white color to make the xenon entrainment between the beryllium and polyimide more clear. We can also see the inward moving polyimide which will lead to a wall shock. In the following we will analyze the shock structure in more detail and check if the results are in agreement with back-of-the-envelope estimates presented in Drake [2006], Drake et al. [2011].

The shock structure in the *xy*-plane at time 13 ns is shown in Fig. 27. The top left panel is for the material identification of beryllium (blue), xenon (black), polyimide (green), gold (yellow) and acrylic (red). The nozzle with inner radius of 600 μm, taper, and shaft with inner radius of 300 μm are visible. The black lines indicate the locations where the mesh is dynamically refined. The top right panel shows the mass density. Part of the polyimide is of very low density and represents vacuum. The dense polyimide tube thickness ranges from 50 to 100 μm. The xenon is compressed by the beryllium piston flow resulting in a primary shock that is located at $x \approx 1700\,\mu$m. For convenience we have indicated with black lines where the material interfaces are.

To check the properties of the primary shock we first determine the beryllium piston velocity. For a 1 ns laser pulse of 4 kJ energy, the irradiance is 8×10^{14} W/cm^2 for a 800 μm diameter spot. Most of this light will be absorbed in the beryllium, so that the absorbed energy per unit of area during this 1 ns is 8×10^5 J/cm^2. About 20% of the beryllium mass, corresponding to 4 μm of the 20 μm, will be ablated by the laser [Drake et al., 2011]. The areal mass density of 16 μm beryllium at 1.8 g/cm^3 is approximately 3×10^{-3}g/cm^2. The conversion from laser energy to kinetic energy of the remaining 16 μm beryllium has a hydrodynamic efficiency of around 10%. The areal kinetic energy density $0.1\times8\times10^5$J/cm^2 corresponds therefore to an initial velocity of the beryllium of a little more than $v = 200$ km/s. This beryllium will launch a shock through the xenon-filled tube. This high velocity is only achieved at early times. The simulated shock velocity gradually reduces to a value between 110 km/s and 120 km/s at time 13 ns as shown in the X velocity plot of Fig. 27.

The initial xenon gas pressure is 1.1 atm while the density is $\rho = 6.5\times10^{-3}$ g/cm^3. With the above mentioned shock velocity of $v \approx 110$ km/s we obtain a xenon post-shock pressure at 13 ns on the order of $\rho v^2 \approx 80$ GPa. The post-shock pressure at $x \approx 1700$ obtained in the bottom right panel of Fig. 27 is in agreement with this estimate. As the shock wave heats the ions, the ion temperature in the postshock region of a strong shock wave with compression ratio equal to $\kappa = (\gamma+1)/(\gamma-1)$ is approximately [Drake, 2006]:

$$k_B T_i = A m_p v^2 \frac{1-1/\kappa}{\kappa}, \tag{98}$$

where k_B, m_p and A are the Boltzmann constant, the proton mass and atomic mass, respectively. For xenon with atomic mass $A = 131$, a shock velocity of 110 km/s, and an adiabatic index of $\gamma = 5/3$ for ions gives a postshock temperature of $T_i \approx 3$ keV. We find with the numerical simulation an ion temperature of 1.2 keV in Fig. 28 that shows the region around the shock. With an effective resolution of 1.6 μm we do not fully resolve this narrow ion temperature peak, resulting in a lower than expected maximum temperature. We find two other temperature peaks as well. One peak is near $x = 1600\,\mu$m behind the shock and another is near $y = 250\,\mu$m. They overlap with the beryllium-xenon and the polyimide-xenon material interfaces, respectively. These spikes are probably the result of a low-order convergence

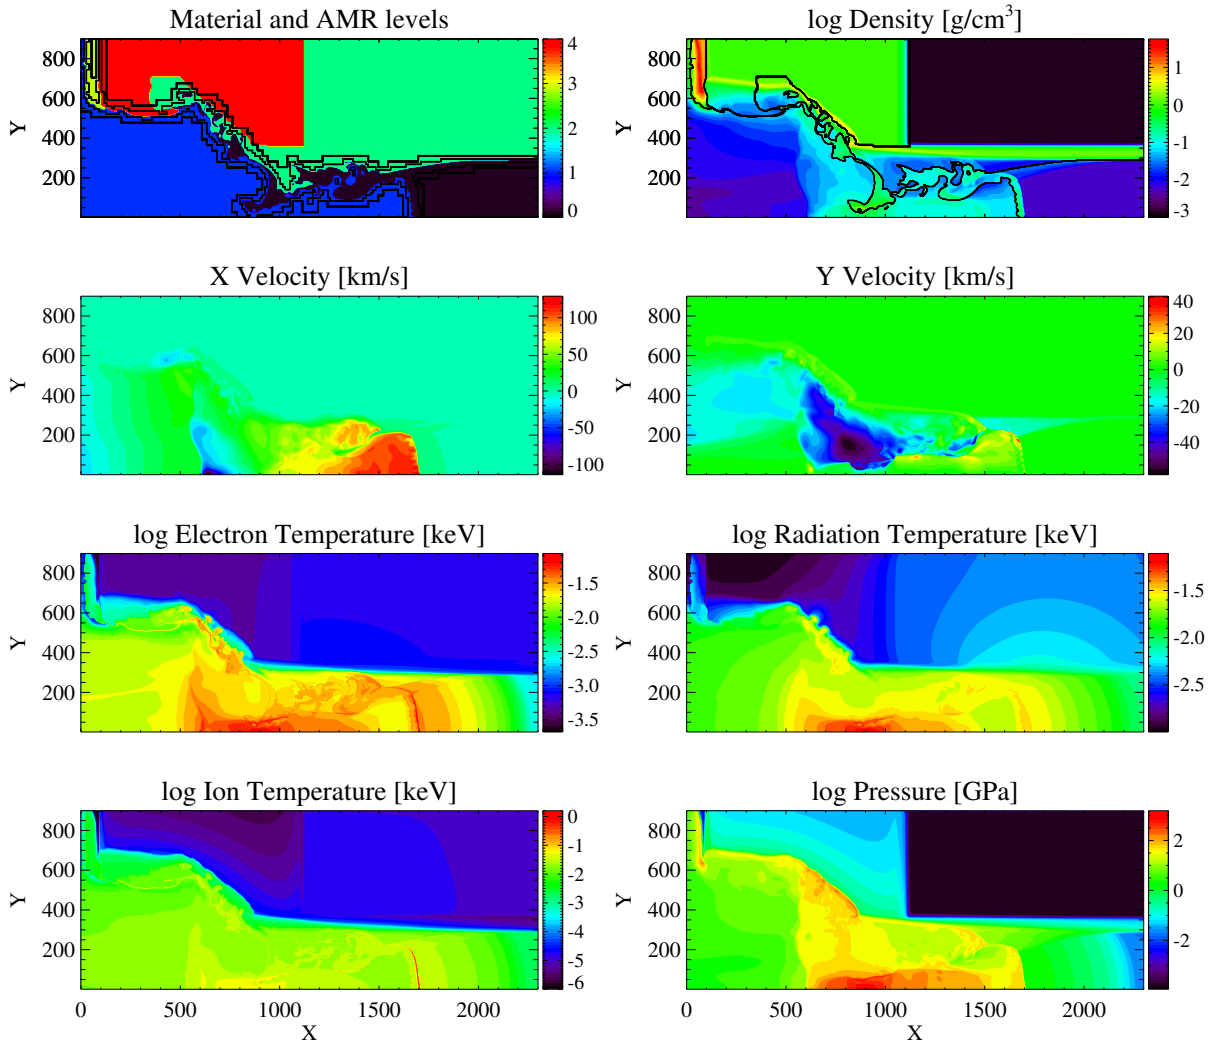

Figure 27: The radiative shock structure in the 3D circular nozzle simulation at 13 ns. The plots show in color contour the plasma and radiation state indicated in the plot titles as a function of the x and y positions in microns. The colors in the top left panel indicate beryllium (blue), xenon (black), polyimide (green), gold (yellow) and acryllic (red). The black lines in this panel show resolution changes, while in the top right panel the lines indicate the material interfaces.

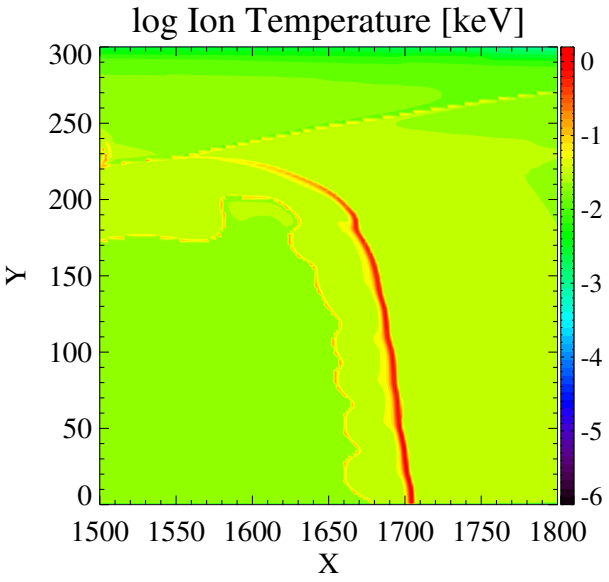

Figure 28: Zoom-in of the ion temperature near the shock with the x and y coordinates in microns.

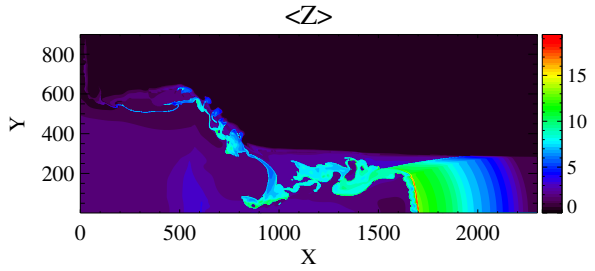

Figure 29: The mean ionization in the circular nozzle simulation as a function of the x and y coordinates in microns.

rate in the interface treatment.

For a high density plasma the collision frequency between electrons and ions is large, so that their temperature will equilibrate. During the shock, however, the ion temperature jumps so that the ions and electrons are out of equilibrium. The ions will heat the electrons via Coulomb collisions and form an equilibration zone directly behind the shock. The electron heating will also increase the ionization in the equilibration zone. Figure 29 shows the average ionization in the xy-plane. We find from the simulation that the ionization is elevated to about $\langle Z\rangle = 17$ at the electron temperature peak in the equilibration zone. At this peak the energy of each ion is shared with $\langle Z\rangle$ electrons. The equilibration temperature $T_{\rm eq}$ can be approximated as the postshock temperature for a strong shock under the assumption that the electrons and ions are in temperature equilibrium:

$$k_B T_{\rm eq} = \frac{1}{\langle Z\rangle + 1} A m_p v^2 \frac{1 - 1/\kappa}{\kappa}. \tag{99}$$

For representative values of the polytropic index between $\gamma = 1.2$ and 1.3 for single-temperature xenon, the estimated equilibration temperature is between $T_{\rm eq} = 76\,{\rm eV}$ and $104\,{\rm eV}$. In reality this temperature will be somewhat lower due to radiative cooling. The simulated value is $T_e \approx 73\,{\rm eV}$ in Fig. 27. This electron temperature is not the final state

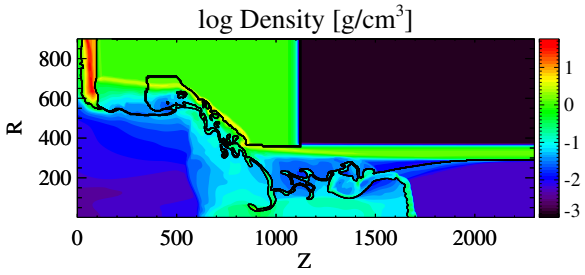

Figure 30: The density at 13 ns in color contour. This simulation is performed in the rz-geometry. The axes are in microns and the black line indicates the change in material level.

in the postshock region since $\sigma T_e^4 \approx 2.8 \times 10^{12}\,\mathrm{W/cm^2}$, where σ is the Stefan–Boltzmann constant. The incoming kinetic energy flux is however $\rho v^3/2 \approx 4.3 \times 10^{11}\,\mathrm{W/cm^2}$ based on the velocity at time $t = 13\,\mathrm{ns}$ of $v = 110\,\mathrm{km/s}$ and xenon density $\rho = 0.0065\,\mathrm{g/cm^3}$. There is therefore not enough incoming energy. There must be a cooling layer [Reighard et al., 2007] through which the electron temperature falls to the final temperature T_f estimated by $2\sigma T_f^4 = \rho v^3/2 \approx 4.3 \times 10^{11}\,\mathrm{W/cm^2}$. The factor 2 is because the radiative cooling layer emits in both directions equally. The final temperature is thus $T_f \approx 38\,\mathrm{eV}$.

The heated electrons are the main energy source for radiation. In Fig. 27 we show the radiation temperature as a measure for the total radiation energy density. The photons travel upstream of the shock where they preheat and ionize the unshocked xenon in the precursor as depicted in the electron temperature panel in Figs. 27 and 29. The radiation transport in the unshocked xenon is not diffusive and we rely on the flux-limited diffusion to recover the optically thin free-streaming limit. This free-streaming approximation is accurate enough as long as the radiation transport from the precursor back to the shock is negligible, in contrast with the almost omnidirectional photon distribution function as assumed in the diffusive limit. A fraction of the upstream radiation expands sideways and heats the polyimide wall ahead of the primary shock. This will ablate the polyimide of the wall. The resulting inward polyimide flow is visible in the Y velocity panel of Fig. 27. This inflow extends at 13 ns to $x \approx 2100\,\mu\mathrm{m}$ and does have a magnitude of about 10 km/s. The exact magnitude of this flow will depend on the radiation transport fidelity used in simulations. The ablated polyimide compresses the xenon as can be seen in the density panel in Fig. 27 by the faint tilted feature between $x \approx 1700\,\mu\mathrm{m}$ and $x \approx 2100\,\mu\mathrm{m}$. The resulting wall shock has an angle with the primary shock and their shock properties were analyzed in Doss et al. [2009, 2011a].

The material identity in the top left panel of Fig. 27 also demonstrates the entrainment of xenon in between the beryllium and polyimide. In Doss et al. [2011a] the entrained flow was shown to first get shocked by the wall shock and then again shocked near the triple point of the wall shock and the primary shock. We also mention that in our simulations the entrained shear flows produce Kelvin–Helmholtz roll-ups at, for instance, $x \approx 1400\,\mu\mathrm{m}$.

The simulation presented in this section was performed in 3D Cartesian geometry. The axi-symmetry in the problem does however allow one to perform this simulation in 2D cylindrical rz-geometry as well. We used the same settings for the numerical radiation-hydrodynamics solvers. The computational domain is $-150 < z < 3900$ and $0 < r < 900$ in microns, where r is the radial coordinate and z is now the coordinate along the tube. The effective resolution is 2560×512 using two levels of refinement, so that the cell sizes in the rz-geometry correspond to the cell sizes in the xy-plane in the 3D Cartesian simulation. The 2D mesh is decomposed in 4×4 grid blocks. The mesh

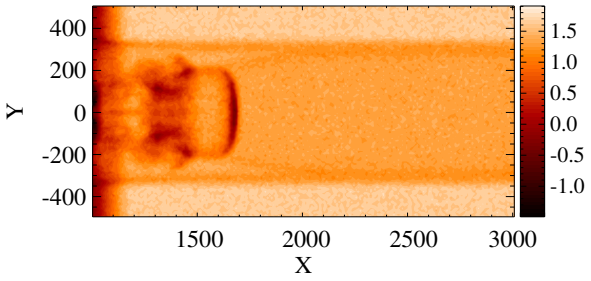

Figure 31: Synthetic radiograph image at 13 ns for the circular nozzle simulation. The X-ray source is located at $(x,y,z) = (2,-12,0)$ in mm.

refinement criteria are also the same as for the 3D case. In Fig. 30 the density is shown at the final time 13 ns. The results are quite similar to those of the previous simulation. The main difference is in the entrained xenon, which requires higher resolution to be fully resolved.

From the CRASH shock tube experiments we typically obtain radiographs by transmitting 5.18 keV (mainly V He-alpha line) X-rays through material whose temperature is of the order of 50 eV. These X-rays interact strongly with the inner-shell electrons. We can therefore use cold opacities for the materials in the experiment to evaluate the transmission. X-ray photons at this energy experience negligible refraction in the experimental target. So it is legitimate to create line-of-sight plots that calculate the optical depth along the straight lines (rays), connecting the X-ray source to the center of the image pixel. To make the simulated images appear more similar to the experimental backlit pinhole radiographs, we need to take into account the finite pinhole size of about 20 mm diameter, the finite exposure time of about 0.2 ns and the effect of a finite number of collected photons (typically 50 photons per 100 nm^2 image pixel). The first two effects are approximated by smoothing the synthetic radiograph over a few pixels. The finite photon count can be taken into account by adding Poisson noise.

The radiograph for the circular nozzle simulation is shown in Fig. 31. We locate the X-ray source at $(x,y,z) = (2000,-12000,0)\,\mu$m. These images show where the dense xenon is. The dense xenon behind the primary shock at $x = 1700\,\mu$m and the wall shock is clearly visible. We can also see the entrained xenon and the Kelvin-Helmholtz roll-up to the left of the primary shock.

3.4.2 Elliptical nozzle

The setup of the elliptical nozzle is very similar to that of the circular nozzle. The main difference is that the wide tube of $1200\,\mu$m inner radius changes the cross-section down the tube into an ellipse with a major axis of $1200\,\mu$m and a minor axis of $600\,\mu$m. The first 1.1 ns of the simulation is, as for the circular nozzle, performed with the H2D radiation-hydrodynamics code to determine the laser energy deposition. For the tube geometry in H2D a straight tube is used with a diameter of $1200\,\mu$m. The output of H2D is used to initialize the CRASH code following the recipe outlined in 3.5.3. If we use the notation that x is the coordinate along the tube and y and z are the two directions transverse to the tube, then we remap the straight tube of H2D to CRASH using the coordinate transformations shown by the Eqs. 113 and 114 in which $x_0 = 500\,\mu$m, $x_1 = 750\,\mu$m, $\varepsilon_y = 1/2$ and $\varepsilon_z = 1$. That means that the circular

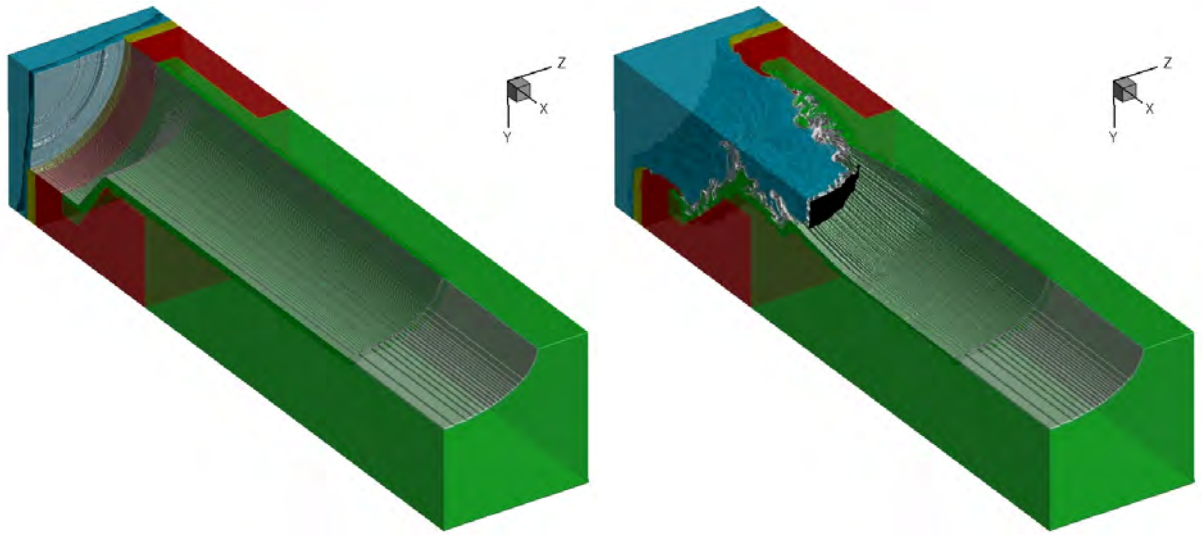

Figure 32: The materials in the 3D elliptical nozzle simulation at 1.1 ns (left panel) and 13 ns (right panel). One quarter of the nozzle is shown. The color code for the materials is the same as in Fig. 26. The black iso-surface indicate the primary shock.

tube is tapered into an elliptical shaft between $x_0 = 500\,\mu$m and $x_1 = 750\,\mu$m. The domain size, effective resolution, refinement criteria, boundary conditions, and the used numerical scheme are the same as for the circular nozzle.

The simulation is performed from 1.1 ns to 13 ns. The number of finite volume cells increased from 29.6 million initially to about 45.3 million at the end of the simulation. The computational time was 3.5 days on 1000 processor cores of the HERA supercomputer. The 3D material identification is shown in Fig. 32. The left panel is for 1.1 ns and the right panel for 13 ns. The color code is blue for beryllium, green for polyimide, red for acrylic and yellow for gold. For convenience xenon is not colored to reveal the elliptical shaft. At 13 ns the beryllium has moved into the shaft and drives a shock in the xenon like a piston. The shock front is indicated with a black surface. The edge of the volume occupied by the xenon is colored in white to visualize the entrainment of xenon between the polyimide and beryllium.

The left panels of Fig. 33 show the material location, mass density and electron temperature, respectively, in the xz-plane at 13 ns. The basic ingredients of the compound radiative shock structure are seen in these panels. The primary shock near $x \approx 1700\,\mu$m in this plane is curved, since the diameter of the laser spot of $800\,\mu$m is smaller than the major axis, $1200\,\mu$m, of the elliptical shaft. The ripple in the compressed xenon region behind the primary shock was analyzed for similar experimental conditions [Doss, 2011]. We also find again the tilted wall shock in front of the primary shock. The electron temperature panel shows the temperature peak in the equilibration zone behind the primary shock and a decreasing temperature behind that in the radiative cooling zone. The right panels of Fig. 33 show the same physical quantities in the vertical yz-plane. The shock structure in this plane is quite similar to the results found for the circular nozzle.

The difference in the shock structure and compressed xenon region behind the shock should also be visible in the synthetic radiographs. Indeed, in Fig. 34 we show in the left panel the image produced by an X-ray source at $(x,y,z) = (2000,-12000,0)\,\mu$m directed at the xz-plane and in the right panel the image produced by an X-ray source

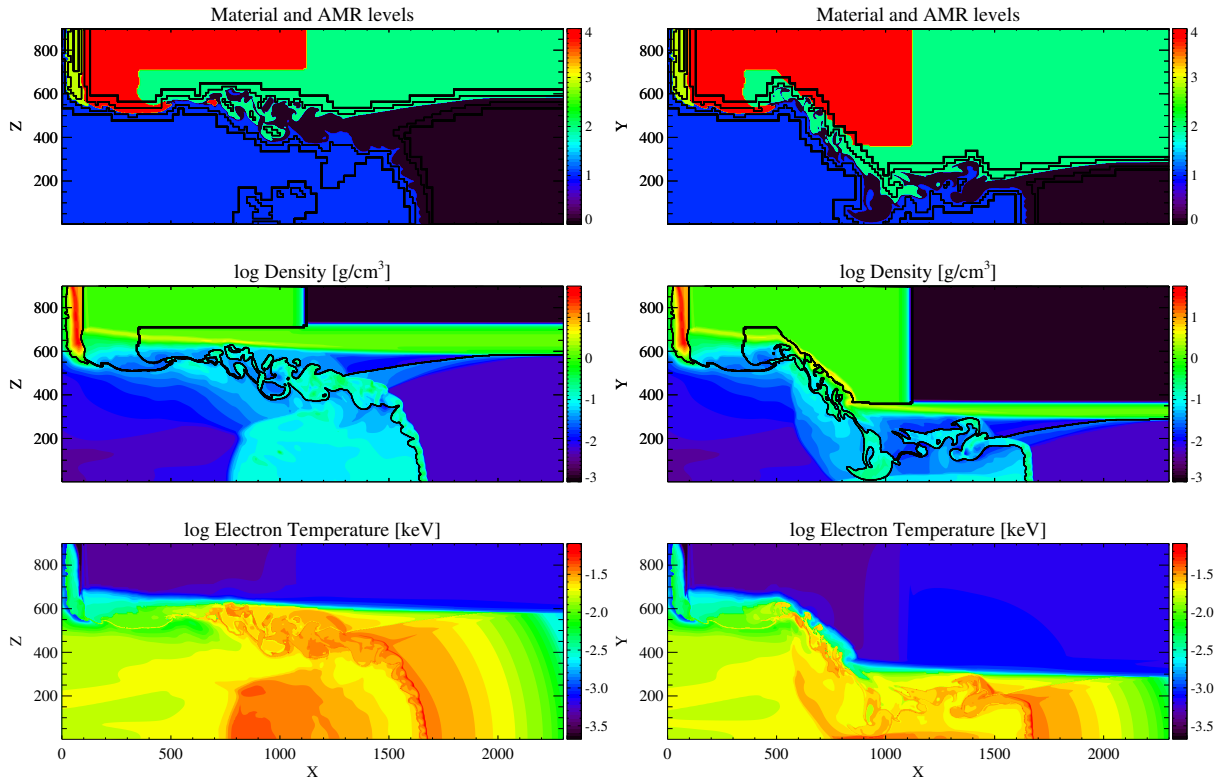

Figure 33: The output of the 3D elliptical nozzle simulation at 13 ns. The state variables indicated in the plot titles are shown in color as a function of the (x,z) coordinates given in microns in the $y = 0$ plane on the left and (x,y) coordinates in the $z = 0$ plane on the right. The color bars provide the scales. The color code of the material levels in the top panels indicates the same materials as in Fig. 27. The black lines in the top panels are for the resolution changes, while the black lines in the middle panels indicate the change in material identity.

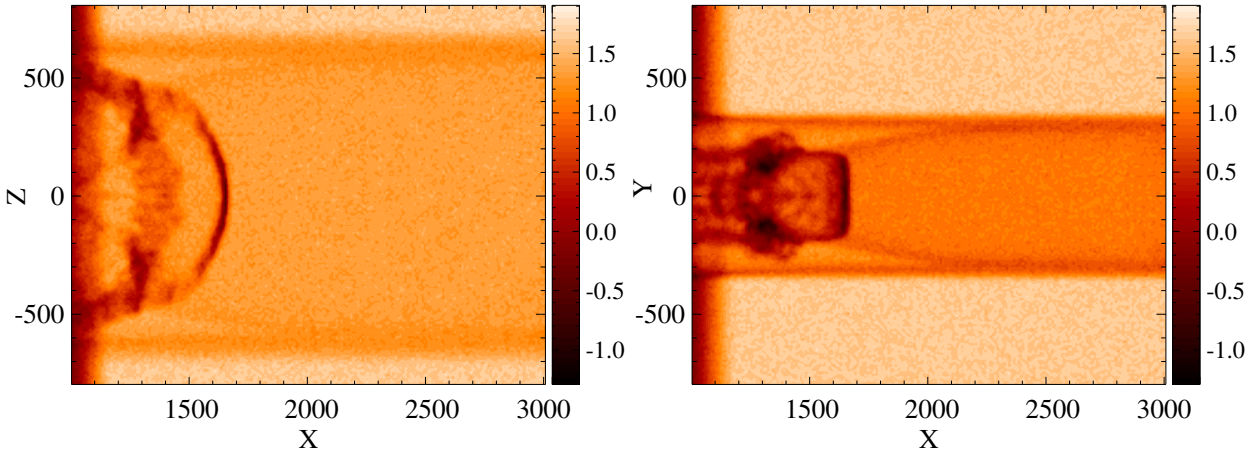

Figure 34: Synthetic radiographs created along the two transverse axes of the 3D elliptic nozzle at 13 ns. The X-ray sources are at $(x,y,z) = (2,-12,0)$ mm directed at the $y = 0$ (left panel) and at $(x,y,z) = (2,0,12)$ mm directed at the $z = 0$ plane (right panel).

at $(x,y,z) = (2000,0,12000)\,\mu$m directed at the xy-plane. The compressed xenon behind the primary and wall shock are found as dark features in these images. Note that the image in the right panel is darker than the image in the left panel, since the rays are going through twice as much xenon (the major axis is twice the minor axis). Also note that the rippled structure of the compressed xenon layer behind the primary shock is somewhat smoothed out in these images.

3.4.3 Improved geometrical transformation for the elliptical nozzle

When we start a 3D simulation from a 2D output file, we take all grid points in the 3D grid, calculate the R-Z coordinates, and interpolate to this coordinate pair from the triangulated 2D grid. This procedure results in an axially symmetric initial condition. For elliptical nozzles, however, we apply a distortion by replacing the actual Y coordinate of the 3D grid cell with A*Y, where the coefficient A is a function of the axial coordinate. This transforms the circular cross section of the tube into an elliptical cross section. Note that A=1 near the Beryllium disk where the laser heating has an affect after 1 nanosecond. The distortion is only applied in the unperturbed part of the domain where the velocities are still zero, and the various material states are still uniform. Although this simple procedure results in an elliptical cross section, the thickness of the plastic tube will vary together with the tube diameter. This means that the tube thickness will be twice as large near the ends of the short axis of the ellipse than at the end points of the long axis for an axis ratio of 2. In the real experimental target the plastic tube has a roughly uniform thickness even after it is squeezed into an ellipse. To fix this problem, we modified the algorithm by making A a function of Y as well as the radial distance R. Points inside the inner side of the plastic tube are stretched according to the original algorithm, but points outside this radius are shifted rather than stretched so that the thickness of the tube is preserved.

3.5 Improving on the baseline CRASH code

In this section, we describe several improvements to the baseline CRASH code, made between years 3 and 5 of the project. First of all we have implemented a new parallel laser energy deposition library as an integral part of our code. This allows the code to simulate the laser heating and the subsequent radiation-hydrodynamic response in a self-consistent and efficient way with one single model. Another improvement was needed for the xenon opacities, since the atomic data provided to our opacity solver were inaccurate. We now use high quality xenon opacities calculated with the super-transition-arrays (STA) model [Bar-Shalom et al., 1989, 1999, Busquet et al., 2010] as an alternative. Both the introduction of the laser package and improved xenon opacities turned out to make the radiation transport much stiffer in some regions. The algebraic multigrid preconditioner using the BoomerAMG solver from the HYPRE library [Falgout and Yang, 2002] resulted in more accurate solutions. It is the purpose of this section to demonstrate that these code changes result in improvement in the fidelity of the simulation results. The reported distortion of the compressed xenon layer on axis [Drake et al., 2011] is now significantly reduced and the wall shock is now more realistic.

Section 3.5.7 describes how we setup the shock tube and laser pulse with our new laser package that is consistent with the experiments performed with the Omega laser facility. This is followed in Section 3.5.8 by a discussion of the simulation results.

3.5.1 CRASH initialization

At the beginning of the CRASH project, we had relatively simple and few initial conditions for the CRASH experiments. The basic initial geometry and physical state of the experiment was hard coded with a few adjustable parameters, such as tube radius, beryllium disk thickness, density and pressure of xenon, etc. Later we switched to initialize CRASH with HYADES output, so the initial state was taken care of by setting up the HYADES run. Starting in year 3, we implemented the laser package into CRASH, and found superior results as compared to those obtained by initializing CRASH with HYADES. This brought up again the need to initialize runs from the very beginning in CRASH. The number of applications using CRASH also increased during the project, which meant a large variety of initial conditions. To match this need, we have implemented a general 2D initialization algorithm, which is controlled from the input parameter file. The user can specify an arbitrary number of material states defined by the material (e.g. xenon or acrylic) and physical properties (e.g. density and pressure). The geometry can be described by an arbitrary number of line segments separating two different material states. The code reads in the material states and the segment positions, and initializes the physical state and the levelset functions for every grid cell in the simulation domain based on the closest segment (see van der Holst et al. [2012]) , Simulating the long-term evolution of radiative shocks in shock tubes.)

3.5.1.1 Initializing CRASH with the laser package

In this section, we present the new package that models the laser energy transport and deposition. In previous reported radiative shock tube modeling efforts [van der Holst et al., 2011, 2012], we used the Lagrangian radiation-hydrodynamics code H2D [Larsen and Lane, 1994] for the first $\sim 1.1\,\mathrm{ns}$. This initializing of the CRASH simulations with H2D turned out to be problematic in that this produced significantly different shock structures when compared

to the observations [Drake et al., 2011]. We therefore opted to implement our own laser package directly into the CRASH code. This also allows us now to simulate the radiative shock experiment in a single self-consistent model. In addition, the new laser package is parallel, while H2D is a serial code, resulting in improved computational speed.

This laser package model decomposes the laser pulse into many rays. We use a ray-tracing algorithm based on a geometric optics approximation with laser absorption via inverse Bremsstrahlung along the trajectory of the rays. Geometric optics is acceptable as long as the electron density does not vary significantly over one wavelength of the laser pulse. This is satisfied most of the time for our applications of interest, with the exception of the startup phase of the laser heating. The inverse Bremsstrahlung absorption is the most important absorption mechanism for the CRASH laser applications [Drake, 2006].

The laser package works both in the 2-D axi-symmetric geometry as well as in 3-D cartesian. For the axi-symmetric geometry we have implemented two versions of the ray tracing: (1) rays confined to the axi-symmetric plane and (2) ray tracing in 3-D. In the 2-D ray tracing case we experienced the problem that all rays that are not parallel to the cylindrical axis will eventually also heat the plasma near the cylindrical axis, resulting in an excessive increase of the electron temperature near the axis. Using 3-D ray tracing in the axi-symmetric geometry mitigates this problem, resulting in improved simulations of shock break-out time and evolution of the radiative shocks compared to the 2-D rays.

Our computationally parallel ray-tracing algorithm is based on previous work on tracing radio rays in the solar corona [Benkevitch et al., 2010]. Here, we briefly summarize the implementation as needed for the laser heating. At each time step and for each ray we trace the trajectory with a ray equation that can be derived from Fermat's principle: a ray connecting two points $\mathbf{r}_1$ and $\mathbf{r}_2$ will follow a path which minimizes the integral of refractive index n, i.e. the variation of the integral

$$\delta \int_{\mathbf{r}_1}^{\mathbf{r}_2} n(\mathbf{r}) ds = 0, \tag{100}$$

where the independent variable s is the arc-length of the ray. This can be shown, see Ref. [Benkevitch et al., 2010], to be equivalent to

$$\frac{d}{ds}\left(n\frac{d\mathbf{r}}{ds}\right) - \nabla n = 0, \quad \text{or,} \quad \frac{d\mathbf{r}^2}{ds^2} = \frac{d\mathbf{r}}{ds} \times \left(\frac{\nabla n}{n} \times \frac{d\mathbf{r}}{ds}\right). \tag{101}$$

By introducing the ray direction, $\mathbf{v} = d\mathbf{r}/ds$, the system (101) of three second-order differential equations is transformed in a set of six first-order equations

$$\frac{d\mathbf{r}}{ds} = \mathbf{v}, \tag{102}$$

$$\frac{d\mathbf{v}}{ds} = \mathbf{v} \times \left(\frac{\nabla n}{n} \times \mathbf{v}\right). \tag{103}$$

For isotropic collisionless plasmas the refractive index is

$$n^2 = \varepsilon = 1 - \frac{\omega_p^2}{\omega^2}, \tag{104}$$

where ε is the dielectric permittivity of the plasma, ω is the angular frequency of the laser light, and the plasma frequency $\omega_p = \sqrt{e^2 n_e / m_e \varepsilon_0}$ depends on the electron density n_e, electron mass m_e, electron charge e and the permittivity of vacuum ε_0. The refraction index is therefore determined from the mass density ρ via $n^2 = 1 - Z\rho/\rho_c$, in which Z is the ion charge and the critical mass density is defined as

$$\rho_c = \frac{\varepsilon_0 A m_p m_e \omega^2}{e^2}, \tag{105}$$

and where A is the mean atomic weight and m_p is the proton mass. This provides us with a final expression for the relative gradient of the refractive index

$$\frac{\nabla n}{n} = -\frac{\nabla(\rho Z)}{2(\rho_c - \rho Z)}. \tag{106}$$

Once this gradient is known, the integration of Eqs. (102)–(103) is performed with the CYLRAD algorithm [Boris, 1971]. The ray trace algorithm of [Benkevitch et al., 2010] is implemented with adaptive step size to handle the steep gradients in the plasma density. For each integration step, every ray is checked for accuracy and correctness, and we ensure that rays do not penetrate in regions where $Z\rho > \rho_c$.

Electron-ion collisions modify the refractive index to a complex value, where the imaginary part corresponds to absorption

$$n^2 = \varepsilon = 1 - \frac{\omega_p^2}{\omega(\omega + i\nu_{\text{eff}})}. \tag{107}$$

The effective electron-ion collision frequency is defined as [Ginzburg, 1964]

$$\nu_{\text{eff}} = \frac{2\pi}{3}\left(\frac{e^2}{4\pi\varepsilon_0 k_B T_e}\right)^2 \sqrt{\frac{8k_B T_e}{\pi m_e}} \left\langle n_i Z^2 \right\rangle \ln\Lambda. \tag{108}$$

Here n_i is the ion number density, k_B is the Boltzmann constant, T_e is the electron temperature, and $\ln\Lambda$ is the Coulomb logarithm. It is due to these electron-ion collisions that the laser energy is absorbed into the plasma. The absorption coefficient (in units of 1/m) is then found as [Drake, 2006]

$$\alpha = \frac{\nu_{\text{eff}}}{c}\frac{Z\rho/\rho_c}{1 - Z\rho/\rho_c}. \tag{109}$$

While performing the integration along each ray, energy is gradually deposited in the plasma.

We have added code infrastructure in CRASH to facilitate the setup of a laser pulse using a set of rays. For the sake of brevity we will only describe the 3-D laser pulse implementation in a 2-D axi-symmetric simulation. In general a 3-D laser pule will generate a 3-D laser heating, and hence 3-D simulations are required. We have code infrastructure to perform such 3-D simulations, but this is computationally expensive. We assume that the departure from axi-symmetry is small. The laser pulse is defined with an irradiance (in units of J/s) and a time profile with a given linear ramp-up time, decay time, and total pulse duration. The laser pulse is further decomposed in a number of laser beams with a circular cross-section. Each beam is defined by a slope θ_b with respect to the x-axis as in Fig. 35 and an initial starting point of the beam defined by x_{start} and a beam offset r_b from the x-axis. For the beams we select a spatial profile of irradiance that is super-Gaussian as a function of the beam radius r': $\propto \exp(-(r'/a)^b)$, where a is the beam radius and b is the super-Gaussian order usually chosen to be 4.2. We use a beam margin of $r' = 1.5a$, beyond which the laser energy deposition is assumed to be negligible. Each beam is discretized with the same number of rays $n_{r'}$ for the radial beam direction and $n_{\phi'}$ rays for the angular direction ϕ'. Due to symmetry properties for 3-D beams in axi-symmetric geometry, we only need to consider half of the angular direction of the beam, i.e. we limit to the angular beam range $[0, \pi]$. The number of radial rays should be high enough to obtain a smooth laser heating profile, but as low as possible for computational speed. This number $n_{r'}$ is typically chosen such that every computational cell at the critical density surface, $Z\rho = \rho_c$, is crossed by at least one ray, preferably more. If the slopes θ_b of the laser beams are not too large, the number of angular rays $n_{\phi'}$ can be selected to be small for the sake of computational speed, while still achieving sufficient accuracy.

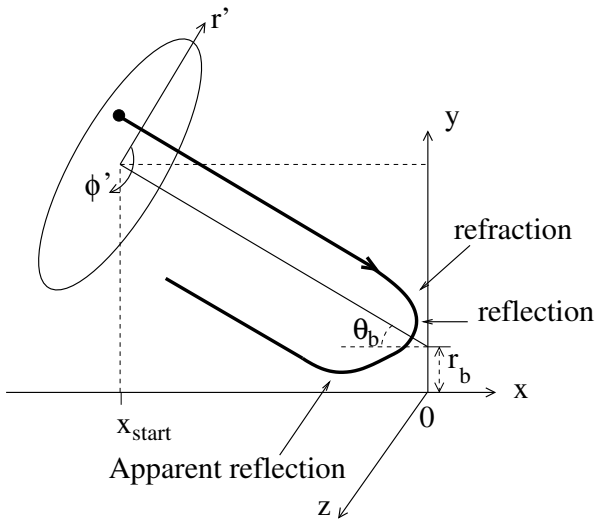

Figure 35: Each laser beam has a different direction θ_b relative to the x-axis. The beam rays start at x_{start} and r_b defines the initial y-position. The beam irradiance profile and the initial beam ray locations are defined as a function of the polar beam coordinates r' and ϕ'. The thick drawn line is a 3-D ray projected onto the axi-symmetric plane, here shown as the xy-plane.

For the laser energy deposition, we evaluate the energy emitted by the laser at each time step. The beam energy is distributed over the number of rays. The local intensity for each ray is obtained from the super-Gaussian beam profile. The laser energy transport in the CRASH code uses 3-D ray-tracing based on geometric optics. During the 3-D ray tracing, we use the x and $r = \sqrt{y^2 + z^2}$ positions in the axi-symmetric plane to determine where in the solution plane the laser energy deposition via inverse bremsstrahlung occurs as well as the further evaluation of the trajectory in 3-D. For these evaluations we need to map the density and density gradient into the 3-D space. A drawing of a 3-D ray projected onto the axi-symmetric plane is shown in Fig. 35. The projected ray shows not only the reflection before approaching the critical density surface and the refraction, but also the apparent reflection at a finite distance from the x-axis instead of a reflection on axis. This apparent reflection is because the 3-D ray is in general not in the axi-symmetric plane and hence has a minimum distance with respect to the x-axis. It is this effect that avoids the excessive laser heating on axis as is the case for 2-D rays in which the rays are confined to the axi-symmetric plane. The deposited laser energy at the current ray location in the axi-symmetric plane is distributed to the nearest computational zone volumes with the sum of the interpolation coefficients equal to one and subsequently added as an explicit source term to the right-hand-side of the electron energy density equation. The scheme is fully conservative since the total energy that is deposited by the laser pulse equals the total laser energy that is absorbed by the plasma.

3.5.1.2 Laser Package Extension to 3D rays with Verification

We have generalized the CRASH laser package to 3D rays in both 2D axial symmetric geometry and 3D Cartesian geometry. In the previous version of the laser package, the rays were all confined to the same axial symmetric plane, resulting in an excessive increase of the electron temperature near the axis, since all rays not parallel to the cylindrical axis would eventually also heat the plasma near the axis. The new 3D laser package mitigates this problem, resulting

in improved simulations, as measured by validation with experimental results for shock break-out time and evolution of the radiative shocks.

For the laser energy deposition, we evaluate the energy emitted by the laser at each time step. The energy is distributed over a finite number of discrete rays conforming to the beam cross-section; the local intensity for each ray is determined by a selected super-gaussian beam profile that depends on the radial beam coordinate. The laser energy transport in the CRASH code uses 3D ray-tracing based on geometric optics. The ray-tracing now uses a new efficient parallel adaptive mesh refinement implementation using the block adaptive tree library (BATL). During the 3D ray tracing the *rz*-plane is rotated to the current position of the ray trajectory to determine laser energy deposition via inverse bremsstrahlung as well as the further evaluation of the trajectory. The deposited laser energy at the current ray location is distributed to the nearest zone volumes at each time step and subsequently added as an explicit source term to the right-hand-side of the electron energy density equation.

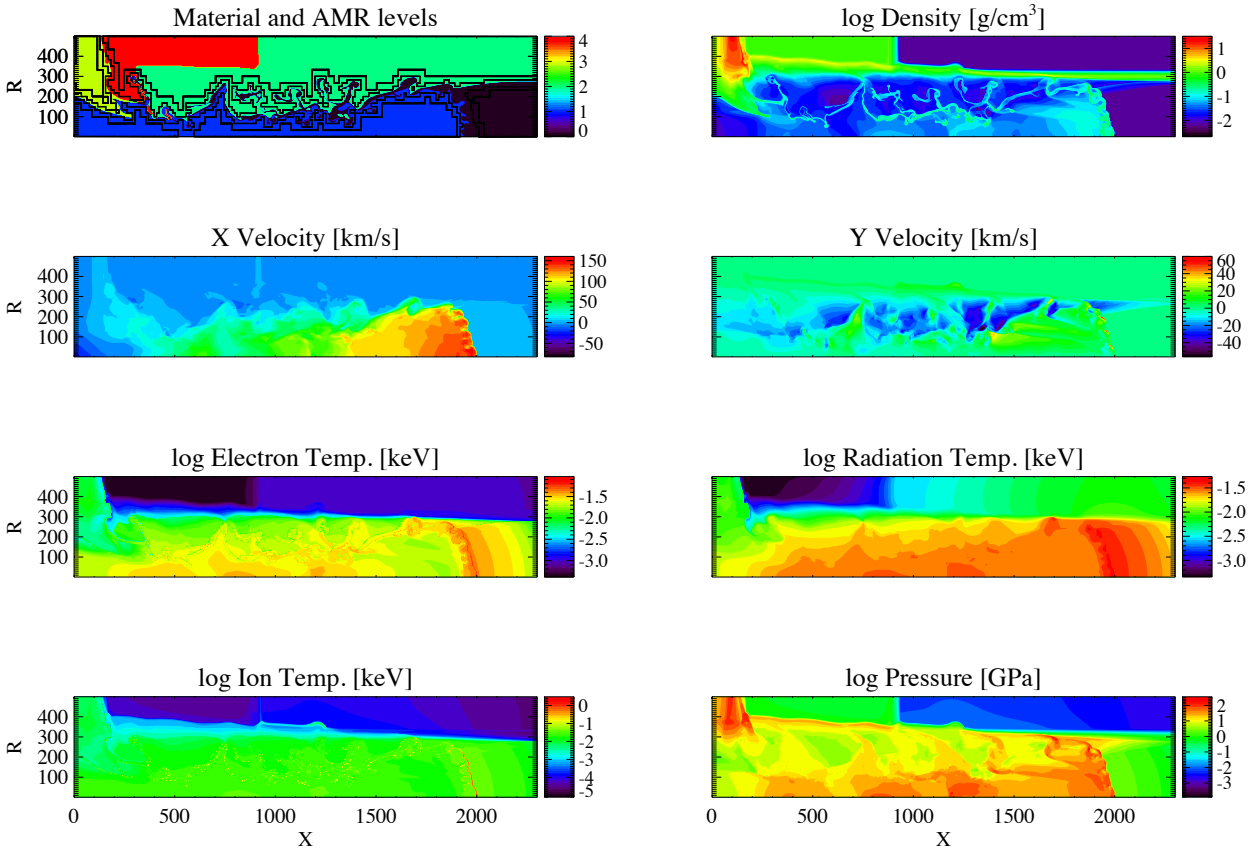

Figure 36: The radiative shock structure at 13 ns of a 3D laser simulation. The plots show in color contour the plasma and radiation state indicated in the plot titles as a function of the x and y positions in microns. The colors in the top left panel indicate beryllium (blue), xenon (black), polyimide (green), gold (yellow) and acrylic (red). The black lines in this panel show resolution changes, while in the top right panel the lines indicate the material interfaces.

We have performed simulations with the new 3D laser package, with results shown in Figure 36. The numerical

set-up is as follows: An Omega laser pulse irradiates a 21 micron beryllium disk with 0.35 micron wavelength light for the full width half maximum (FWHM) duration of 1 ns and with a laser energy deposition of 3.8 kJ. The laser spot size is 820 micron FWHM diameter. The laser pulse consists of 10 beams and are modeled as 4 beams with the power per beam weighted by the number of beams approximately at that angle with respect to the normal: 50.6° (times 1), 42° (times 3), 31.7° (times 4), and 10.2° (times 2). The intensity of each beam is modeled as a super-gaussian profile of order 4.2 in the radial beam coordinates. In the simulation below we use 900 by 4 rays for the radial and angular direction to discretize the laser energy of each beam and multiply the total laser energy deposition by a 0.7 laser scale factor. For this simulation we used the ARTEP opacities (see section 3.5.6.5) for xenon and CRASH calculated opacities for the other materials. In addition we use the new HYPRE preconditioner, which improves the convergence behavior in the flux-limited multi-group radiation diffusion solver. The primary and wall shock are recovered with the updated CRASH code.

The laser package has been tested in three significant ways. First, the 3D capability of the ray-tracing algorithm is tested as part of our nightly code verification tests. In the test, a laser ray transports and deposits energy while propagating through a plasma that has a linear density profile. The refracted ray is tested for angle of incidence, accuracy of the closest approach to the critical density surface (ray turning point), and total energy deposition along the ray's path against the known analytical results (Kruer 2003). Figure 37 shows the result of a 3D laser ray test in an RZ geometry. (Labeled RX by the code for technical reasons.) Note that the ray does not intersect the X axis, because it is not in the RX plane at that point.

Second, the nightly tests verify by comparison to the exact solutions that the code maintains the 2nd order convergence of the numerical solution for the turning point and the 1st order convergence of the energy deposition of the ray.

Lastly, we are validating the full CRASH code, but particularly the laser package, against our experimental results of shock breakout of a 20 μm Be foil. In the experiments, 10 beams at OMEGA are used to irradiate a Be foil with at a laser intensity of 7.0×10^{14} W/cm^2, for a 1.1 ns pulse. The shock breakout is measured by two different experimental techniques to be 450 ps with a systematic error of $\pm$ 50 ps. Although our final simulations with the 3-D laser package are still underway, the preliminary results for shock breakout of a 20 um Be foil with the full experimental intensity are approximately 400 ps.

3.5.2 Reading 2D CRASH output as an initialization for 3D CRASH runs

Although the laser package works for 3D CRASH runs, it is very expensive to simulate the first nanosecond in a full 3D simulation. This is due to the very small time step required by the laser package, and also the number of rays required for a highly resolved 3D grid is huge. Since the first nanosecond of the simulation affects the cylindrical part of the tube only for the PY5 experiment, it is practical to do this simulation in 2D axial symmetry (R-Z geometry). After the laser heating is finished, we save the full 2D information and start a new 3D simulation. The 3D simulation reads in the 2D data file and interpolates onto the 3D grid using Delaunay triangulation. The procedure is similar to the one used for reading the HYADES files, but the list of state variables and their units are fully compatible when the 2D CRASH file is read into the 3D CRASH run. On the other hand the 2D AMR grid in the CRASH file is more complicated than the distorted Cartesian grid produced by HYADES, so it needs a general triangulation.

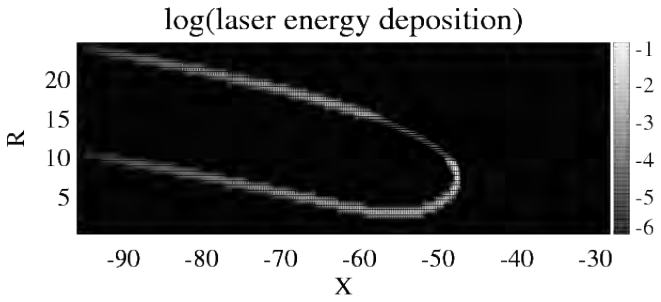

Figure 37: 3D laser ray turning point test for a ray with a 10° angle of incidence in an RZ geometry. Note that the ray does not intersect the X axis, because it is not in the RZ plane at that point.

3.5.3 Initializing CRASH with Hyades

The laser energy deposition during the first 1.1 ns is evaluated with the Lagrangian radiation hydrodynamics code Hyades 2D (H2D) [Larsen and Lane, 1994]. The output of H2D is on a distorted logically Cartesian mesh of $n_x \times n_y$ cells. When we initialize the CRASH code with H2D, this mesh is first triangulated by splitting each quadrilateral along the shorter diagonal. The interpolation from the triangulated Hyades grid requires finding the triangle that surrounds the center of a given grid cell of the CRASH code. A simple linear search can become very inefficient when we have many grid cells (order of a hundred thousand) per processor. To accelerate this method, first we create a uniform grid with about 200 by 200 resolution that covers the whole domain. For each rectangular cell in the uniform grid we find and store the list of triangles that intersect it, which can be done rapidly. Then we perform the interpolation onto the CRASH grid by first finding the rectangular cell that surrounds it, and then we only check the triangles that intersect this cell to see which one contains the CRASH grid cell center.

The H2D simulations are performed on a Lagrangian grid where all cells correspond to a unique material. These materials are identified by a material index in the H2D output. Our code uses an Eulerian grid and we track the material by means of level set functions. These level set functions are smooth and signed distance functions, initialized with the following algorithm. For each cell i of the H2D grid having material index α_i, the level set function for material α_i is set to the distance to the closest H2D cell that contains a material different from α_i:

$$d_{\alpha_i}(i) = + \min_{j,\alpha_j \neq \alpha_i} |\mathbf{r}_i - \mathbf{r}_j|, \tag{110}$$

where $\mathbf{r}_i$ is the location of cell i. The level set functions for the other materials, $\beta \neq \alpha_i$, are set to the negative distance to the closest cell containing material β:

$$d_\beta(i) = - \min_{j,\alpha_j = \beta} |\mathbf{r}_i - \mathbf{r}_j|. \tag{111}$$

We interpolate these level set functions to the cells on the CRASH grid. In our first-order level set scheme the material

in the cell is indicated by the largest level set function. For the simulations with H2D of our shock tube experiments we use xenon, beryllium, polyimide, gold, acrylic and "vacuum". In CRASH, vacuum is reassigned as polyimide with low mass density. At later times during the evolution, the location of material m follows the simple advection equation for the level set function

$$\frac{\partial d_m}{\partial t} + \nabla \cdot (d_m \mathbf{u}) = d_m \nabla \cdot \mathbf{u}. \tag{112}$$

Again, the material having the largest d_m is assigned to be the material of the cell.

A 3D nozzle is created from 2D Hyades as follows. We first read the H2D shock tube output and triangulate the data. After that we transform the coordinates in the transverse directions y and z. The circular cross section of the tube shrinks into an ellipse or smaller circle by means of

$$y' = y\left[1+(\varepsilon_y-1)\max\left(0,\min(1,\frac{x-x_0}{x_1-x_0})\right)\right], \tag{113}$$

$$z' = z\left[1+(\varepsilon_z-1)\max\left(0,\min(1,\frac{x-x_0}{x_1-x_0})\right)\right], \tag{114}$$

where ε_y and ε_z are the factors with which the y and z coordinates contracts for $x > x_1$ along the tube. This shrinking varies linearly between x_0 and x_1. For $x < x_0$ the tube is not modified. This means that the plasma at the far end is relocated closer to the axis. This will change the output of H2D to an unphysical state (the solution of the straight tube is not the same as the solution of the nozzle). The impact of this change in geometry is however minimal since at 1.1 ns the shock dynamics did not yet reach the far end of the shock tube.

3.5.3.1 Limitations of Hyades

A rate-limiting element for our predictive-science studies has been the Hyades code, a serial, Lagrangian, rad-hydro code that we have used to model the laser-energy deposition phase of the experiment. With the arrival of a rezoner for this code, we proved able to accomplish the sets of order 100 runs needed for our UQ studies. In particular, we completed during 2010 a run set that we are using to initiate the current multi-dimensional UQ runs for CRASH. However, we encountered a number of issues that led us to decide that it was necessary for us to implement a laser package in CRASH. Figures 38 and 39 illustrate two of the issues that support this decision. Figure 38 illustrates that rezoning impacts the region where Be, Xe, and Au interact, which is where the wall shock-primary shock interactions will later become important. It is worth noting that manual rezoning will produce similar variations though in a less controlled and less quantifiable fashion. Figure 39 shows that running the same physical problem, driven by thermal radiation, produces very different results when Hyades is used for the first ns than it does when only CRASH is used throughout. We do not know why this occurs; it might reflect some inadequacy in our mapping of the parameters between codes, although we have examined this closely, or the lack of conservation of vorticity in Hyades, or something else. We believe that Hyades is a very useful code for experiment design and for other things Lagrangian codes are generally used for. However, doing UQ studies of systems that generate substantial vorticity, as we need to do, is not an area of strength for this code.

3.5.4 Boundary conditions for radiation diffusion

We have also been studying effects of boundary-condtion (BC) treatments for radiation problems. Ambiguity of BCs arises because angle averaging, inherent in diffusion approximations, discards information; physics on the boundary

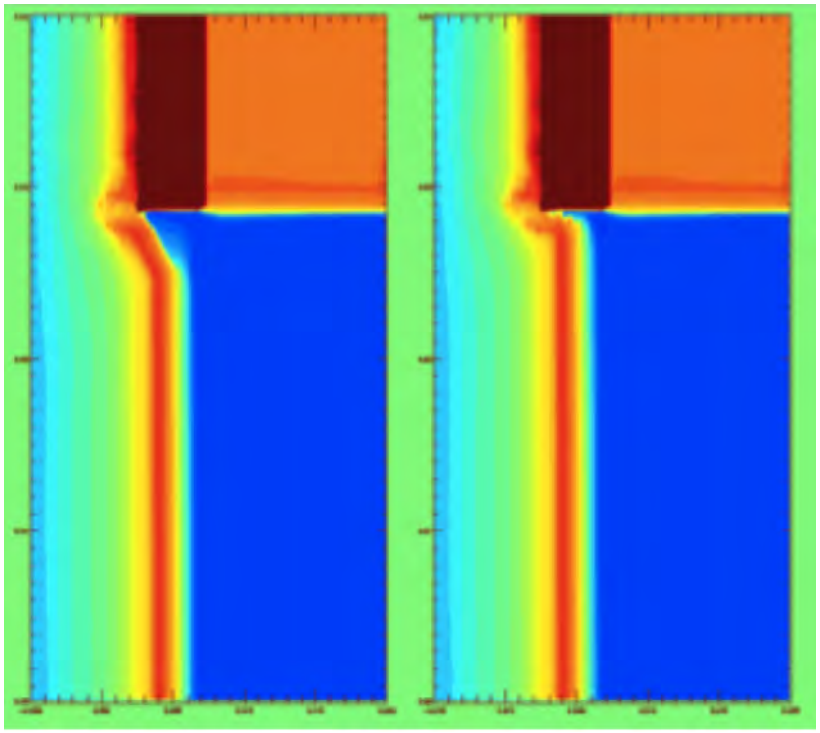

Figure 38: Effect of rezoner settings on Hyades simulation outputs. Log density is shown on an RZ plot. The rezoner is automatically adjusting a specified number of zones near the intersection of Be, Xe, and gold on the upper left. The only difference in these runs is that 6 zones are adjusted for the left case while 3 are for the right.

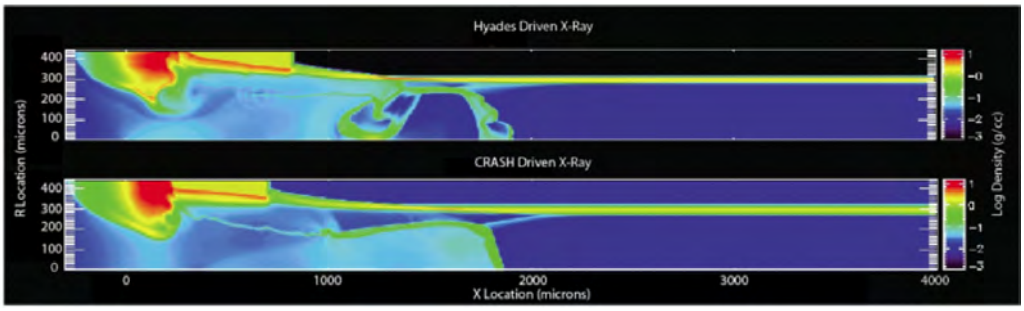

Figure 39: Results from two simulations that should give identical results. In both cases, a 167 eV thermal x-ray flux irradiates a 34 m thick Be disk for 1 ns and results are shown at 13 ns. Top: run in Hyades for 1 ns then CRASH. Bottom: run only in CRASH.

can be only incompletely described. Thus, an approximation is required, which can take the form of an integral constraint that makes assumptions about the discarded angular information [Pomraning, 1973, p.51] [Castor, 2004, p.242].

Applying an integral constraint, and specifying that no incident radiation can originate outside the domain, gives a suitable boundary condition, which applies to both gray transport or to individual energies of multigroup transport,

$$-\frac{1}{2}\langle\mu\rangle cE_\varepsilon + \frac{1}{2}F_{\text{out}} = 0, \tag{115}$$

where E_ε is the radiation energy density, F_{out} is the radiation energy flux density emerging from the boundary, $\langle\mu\rangle$ is the mean cosine of the propagation angle of the emergent radiation, and c is the speed of light. Until now, in the CRASH project, we have used a hard-coded value of $\langle\mu\rangle = 1/2$, which assumes radiation is in the diffusion limit at the boundary. We now allow arbitrary values of $\langle\mu\rangle$ and, in particular, have investigated $\langle\mu\rangle = 1$, which corresponds to free-streaming at the boundary [Pomraning, 1982], a situation which often obtains in HEDP problems.

Following Myra and Hawkins [2013], the flux at the boundary interface is

$$F_{n+\frac{1}{2}} = \langle\mu\rangle c\left(\frac{2D_{\text{eff.}}E_n}{2D_{\text{eff.}} + c\Delta x\langle\mu\rangle}\right), \tag{116}$$

where $D_{\text{eff.}}$ is the diffusion coefficient at the interface and Δx is the grid spacing. Having the flux at the boundary, we can calculate, self-consistently, the energy flow off the computational mesh. Comparisons of this flow, using PDT and CRASH, are currently in progress.

We have also introduced a scheme for calculating an extrapolation length as a function of energy group, which is also a function of $\langle\mu\rangle$. Using this scheme, we can then solve for the ghost-cell value of E_ε needed by the CRASH code,

$$E_{n+1} = \left(\frac{1-\frac{3}{2}\kappa\Delta x\langle\mu\rangle}{1+\frac{3}{2}\kappa\Delta x\langle\mu\rangle}\right)E_n, \tag{117}$$

where E_{n+1} is the value of E_ε in the first ghost zone (here at a right boundary), E_n is E_ε at the last computed cell, and κ is the absorption coefficient. With this result, we can compute BCs in CRASH using a variety of integral schemes.

Figure 40 shows the effects of setting $\langle\mu\rangle$ at the boundary for FLD. Especially notable: Discrepancies in T_{rad} are cumulative with time. By 200 ps, there is $\sim 7\%$ discrepancy in E_ε at the right boundary. In addition, changing the value of $\langle\mu\rangle$ at the right boundary—which is the only difference between the two simulations—significantly influences evolution of the whole domain (note the resulting discrepancy at the *left* boundary). In part, these counter-intuitive results follow from the temperature-dependent, real opacities and multiple materials used in these runs.

3.5.5 Improvements to the Modeling of Equations of State in Dense Plasmas

The development of theoretical models to simulate the Equation Of State (EOS) for dense plasmas had been mostly finished in 40s-50s (see Zeldovich and Raizer [1966]). One can solve numerically the thermodynamic parameters and radiation transport in high-temperature plasmas, if the atomic and ion "parameters"are known such as all ionization potentials, an excitation energy for all levels, as well as the degeneracy for both ground and energy levels. The researcher may be tempted to solve the coupled system of *the Saha equations*, which relate the populations of ions with different cherge numbers, however, the latter approach appears to be inoptimal, if plausible at all.

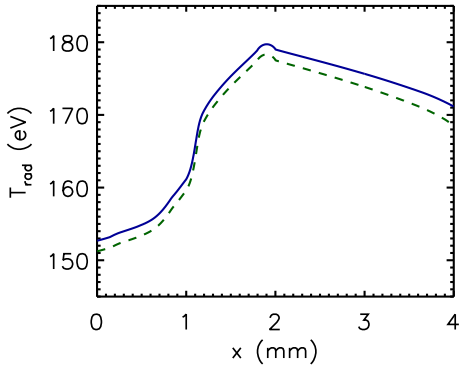

Figure 40: Result from a 10-group calculation of the 1D CRASH-based problem described in Ref. [Myra and Hawkins, 2013] using the CRASH code. Equation 117 is used to set the BCs at both the left and right boundaries. The solid curve shows the solution for $\langle\mu\rangle = 1$ at the right boundary, while the dashed curve shows $\langle\mu\rangle = 1/2$ at the right boundary. In both cases, the opaque left boundary uses $\langle\mu\rangle = 1/2$.

Here, we employ the assumption of a Local Thermodynamic Equilibrium (LTE). This means that we assume the population of all atomic and ionic states to be exactly the same as those calculated using the methods of statistical physics [Landau and Lifshitz, 1980]. The Section is split into two parts, the first of them being entirely devoted to the thermodynamic analisys applied to LTE dense plasmas. It is more focused on the EOS function calculation, however, the obtained populations and partition function are also used in the second part, to find opacities. This makes the overall aproach to the radiation hydrodynamics to be consistent. The second part describes the way to apply the results of the first part to a problem of the radiation transport in plasmas. This transport is treated in the framework of Multi-Group Diffusion (MGD). The assumtions of LTE and MGD are key points of the presented work.

3.5.5.1 Effect of the Fermi statistics on Thermal Ionization

While simulating the ionization equilibrium in partially ionized electron-ion plasmas using Saha equations, the electrons are usually assumed to be an ideal Boltzmann gas. However, the electron is a Fermion with spin of 1/2. In practical calculations the conditions for applicability of the Boltzmann gas model for electrons are often not satisfied, resulting in very low accuracy. Therefore, it is worth while checking whether or not one needs to assume Boltzmann statistics for the electrons when solving the ionization equlibrium.

For realistic quantitative simulations this assumption is completly unnecessary, as solving the ionization equilibrium under the incorrect assumption of Boltzmann statistics is not any easier than that with the correct Fermi statistics for electrons.

In a partially ionized plasma the free electron density, corrected by the effects of the Fermi statistics ('the exchange interaction'), should be used in solving the Equation-Of-State (EOS), which is also directly affected by the exchange interactions. There is no need to remind that *at the given electron density* the exchange interaction increases the electron pressure and the internal energy density (see Landau and Lifshitz [1980]). On the other hand, in partially

ionized plasmas this effect may be partially or even fully balanced by the electron density decrease due to the effect of the exchange interaction on the ionization equilibrium.

Helmholtz free energy. Consider an ionized monatomic gas with positive non-complex ions. The Helmholtz free energy, $F = F_{ion} + F_e$, is assumed to be the total of contributions from each of the ion charge states, $i = 0$ to $i_{\max}$ (we apply Eq.(42.3) from Landau and Lifshitz [1980] to account for these contributions), as well as the contribution from electrons:

$$F = -T \sum_{i=0}^{i_{max}} N_i \log \left[g_i \frac{eV}{N_i} \left(\frac{MT}{2\pi\hbar^2} \right)^{3/2} \exp \left(- \sum_{j=0}^{i-1} \frac{I_j}{T} \right) \right] + F_e, \tag{118}$$

where $N_i = n_i V$ is the total number of ions in the charge state i in the volume V and I_j ($j = 0, 1, 2, \dots$) is the energy needed to ionize an atom or ion from the charge state j to the charge state $j+1$ (the ionization potential), and g_i is a statistical weight of an ion (atom) in a given charge state (see below for more detail).

Ionization equilibrium: formulation of the problem. Now we formulate the requirement for the ionization equilibrium with respect to the reaction $(i) \leftrightarrow (i+1) + e$ for each ion charge state, i. The Helmholtz free energy is a minimum at the equilibrium set of N_i and N_e. Therefore, the total derivative of F with respect to N_e should be zero:

$$\frac{\partial F}{\partial N_e} + \frac{dN_i}{dN_e} \frac{\partial F}{\partial N_i} + \frac{dN_{i+1}}{dN_e} \frac{\partial F}{\partial N_{i+1}} = 0. \tag{119}$$

For the reaction under consideration, the increments in the particle numbers should be related as follows: $dN_i = -dN_{i+1} = -dN_e$. Therefore the requirement, $dF/dN_i - dF/dN_{i+1} - dF/dN_e = 0$, gives:

$$-T \log \left[\frac{g_i}{N_i} e^{-\Sigma_{j=0}^{i-1} I_j/T} \right] + T \log \left[\frac{g_{i+1}}{N_{i+1}} e^{-\Sigma_{j=0}^{i} I_j/T} \right] - \mu_e = 0, \tag{120}$$

where we applied the definition of the chemical potential, $\mu = (\partial F/\partial N)_{T,V}$, to the electron gas. The solution of the ionization equilibrium, therefore, reads:

$$N_{i+1}/g_{i+1} = (N_i/g_i) e^{(-\mu_e - I_i)/T}, \tag{121}$$

or, applying this recursively:

$$N_i/g_i = (N_0/g_0) e^{(-i\mu_e - \Sigma_{j=0}^{i-1} I_j)/T} = (N_0/g_0)(g_e)^i e^{-\Sigma_{j=0}^{i-1} I_j/T}, \; g_e = e^{-\mu_e/T}, \tag{122}$$

where g_e is the effective statistical weight of a free electron. Indeed, the effective statistical weight of i electrons combined with an ion in the charge state, i, is the product of statistical weights for each of the particles under consideration, $(g_e)^i g_i$, in accordance with Eq.(122).

In the limiting case of $g_e \gg 1$ (a Boltzmann gas of electrons with large negative value of μ_e), g_e might be interpreted as the large number of elementary quantum states the detached electron can occupy, which facilitates the ionization by resulting in a higher total probability for the ionized state. In the opposite limiting case of a degenerate Fermi gas of electrons, the positive chemical potential, $\mu_e > 0$, tends to the Fermi energy E_F, which in this limiting case is much greater than the temperature. Accordingly, the exponentially low value of $g_e = e^{-E_F/T}$ in this case means the low probability for an electron to jump from a bound state with negative energy to a free state above the threshold of the positive Fermi energy.

Partition function and electron density. We now introduce the ion partition function, $p_i = N_i/N_a$, N_a, being the total number of atoms. Since the partition function is normalized by unity, we have:

$$p_i = \frac{g_i (g_e)^i e^{-E_i/T}}{S}, \tag{123}$$

where we introduced the statistical sum:

$$S = \sum_{i=0}^{i_{max}} \left[g_i (g_e)^i e^{-E_i/T} \right], \tag{124}$$

as well as the total ionization energy spent to ionize the atom to the state i:

$$E_i = \sum_{j=0}^{i-1} I_j. \tag{125}$$

Introducing the averaging operator acting on an arbitrary function of the ion charge number, $\langle f_i \rangle = \sum p_i f_i$, and assuming quasi-neutrality, $N_e = \sum i N_i$, we obtain the expression for the electron density:

$$Z = N_e/N_a = \langle i \rangle. \tag{126}$$

On the other hand, for given T and $n_a = N_a/V$ the electron concentration may be found as a function of $g_e (= e^{-\mu_e/T})$. Now we assume electrons form an ideal Fermi gas. This assumption immediately gives us another relationship between the electron density and g_e (see Eq.(56.5) from Landau and Lifshitz [1980]):

$$Z = g_{e1} \mathrm{Fe}_{1/2}(g_e). \tag{127}$$

The coupled equations (126) and (127) are used below to solve Z and g_e. Here

$$g_{e1}(T, N_a/V) = \frac{2V}{N_a} \left(\frac{m_e T}{2\pi\hbar^2} \right)^{3/2} \tag{128}$$

is a value such that in the *Boltzmann* electron gas $g_e = g_{e1}/Z$ would hold. $\mathrm{Fe}_\nu(g_e)$ is the Fermi function:

$$\mathrm{Fe}_\nu(g_e) = \frac{1}{\Gamma(\nu+1)} \int \frac{x^\nu dx}{g_e e^x + 1}, \tag{129}$$

where Γ-function is introduced as usually: $\Gamma(\nu+1) = \nu\Gamma(\nu), \Gamma(1/2) = \pi^{1/2}$. Below we use the following auxiliary functions:

$$R^-(g_e) = \frac{\mathrm{Fe}_{-1/2}(g_e)}{\mathrm{Fe}_{1/2}(g_e)}, \qquad R^+(g_e) = \frac{\mathrm{Fe}_{3/2}(g_e)}{\mathrm{Fe}_{1/2}(g_e)}. \tag{130}$$

Derivatives along the ionization curve. Taking differential of the equation of ionization equilibrium, $G = 0$, one gets an equation relating the differentials of different variables along the curve of ionization equilibrium:

$$\frac{dg_e}{g_e} \left(\langle i^2 \rangle - Z^2 + Z R^-(g_e) \right) + \frac{dT}{T^2} \left(\langle i E_i \rangle - \langle E_i \rangle Z \right) \Big) = \left(\frac{3}{2} \frac{dT}{T} + \frac{dV}{V} \right) Z, \tag{131}$$

here we substitute wherever possible Z for $\langle i \rangle$ and $g_{e1}\mathrm{Fe}_{1/2}(g_e)$. Accordingly, the differential of Z is:

$$dZ = \left(\frac{3}{2} \frac{dT}{T} + \frac{dV}{V} - \frac{R^-(g_e) dg_e}{g_e} \right) Z. \tag{132}$$

To solve the ionization equilibrium for the given T and $n_a = N_a/V$, one needs to solve g_e from equation $G(g_e) = 0$, $G(g_e) = \langle i \rangle - g_{e1}\mathrm{Fe}_{1/2}(g_e)$, which is obtained by means of excluding Z from Eqs.(126,127). It may be solved using the Newton-Rapson iterations with any trial value, $\log(g_e)_{old}$, the improved value, $\log(g_e)_{new}$, is obtained from the equation as follows:

$$\log(g_e)_{new} = \log(g_e)_{old} - \frac{\langle i \rangle - g_{e1}\mathrm{Fe}_{1/2}((g_e)_{old})}{\langle i^2 \rangle - \langle i \rangle^2 + Z R^-((g_e)_{old})}, \tag{133}$$

where the derivative $g_e G'$, which should stand in the denominator of Eq.(133), is derived using the easy-to-check equations as follows:

$$g_e \frac{d\mathrm{Fe}_\nu(g_e)}{dg_e} = -\mathrm{Fe}_{\nu-1}(g_e), \tag{134}$$

and for any set of values in the charge states, f_i:

$$g_e \frac{\partial \langle f_i \rangle}{\partial g_e} = \langle i f_i \rangle - \langle f_i \rangle \langle i \rangle. \tag{135}$$

We see that to solve iterations in Eq.(133) for Fermi gas of electrons is not a more computationally intense problem comparing with the same problem in assuming electrons to be a Boltzmann gas. In the latter case, $g_e \to +\infty$, we have $\mathrm{Fe}_\nu(g_e) \approx 1/g_e$ (see Eq.(146) below) and $g_e \approx g_{e1}/Z$, which allows us to iterate Eq.(133) as a somewhat simpler equation for Z: $\log(Z_{new}) = \log(Z_{old}) + (\langle i \rangle - Z_{old})/(Z_{old} + \langle i^2 \rangle - \langle i \rangle^2)$. However in any case, the most cumbersome computations while solving Eq.(133) for a Fermi gas, or the equation for Z for a Boltzmann gas, is in explicitly calculating the numerous partition functions for many charge states and excitation levels. Compared with these bulk computations, the presence of the Fermi functions in Eq.(133), which may be tabulated for all interesting cases of $\nu = -1/2, 1/2, 3/2$, does not matter at all.

Plasma thermodynamics and Equation-Of-State. Now we substitute the ion partition function into Eq.(118). After some algebra we obtain:

$$F = -TN_a \log\left[\frac{eV}{N_a}\left(\frac{MT}{2\pi\hbar^2}\right)^{3/2}\right] - TN_a \log S + \Omega_e, \tag{136}$$

the thermodynamic potential $\Omega_e = F_e - \mu_e Z N_a$ for Fermi gas of electrons is given by Eq.(56.6) from Landau and Lifshitz [1980]:

$$\Omega_e = -[g_{e1} N_a] T \mathrm{Fe}_{3/2}(g_e), \tag{137}$$

the product in the square brackets, being independent of N_a because $g_{e1} \sim N_a^{-1}$. Eq.(136) provides the free energy in the case of local thermodynamic equilibrium with the first term being the contribution from the ion translational energy. This term may be written as the function of the ion temperature, in the case the latter differs from the electron temperature. Unless the ion-ion interaction is taken into account, this first term gives the contributions of $n_a T_i$ and $3n_a T_i/2$ to the total plasma pressure and total energy density correspondingly. The second term is the Boltzmann distribution of ions over the ionization and excitation states, expressed in terms of the statistical sum. Finally, the electron gas with the variable particle number gives the contribution of Ω_e instead of F_e.

While differentiating Eq.(136) with respect to T and V, it is important that the derivatives by g_e from the second and third terms cancel each other: $g_e(\partial \log S/\partial g_e) = \langle i \rangle = Z$ and $-g_e g_{e1} \mathrm{Fe}'_{3/2}(g_e) = g_{e1}\mathrm{Fe}_{1/2} = Z$. That is why for the internal energy density, $\mathscr{E}$, and for the pressure we find:

$$\mathscr{E} = -\frac{T^2}{V}\left(\frac{\partial}{\partial T}\left(\frac{F}{T}\right)\right) = \mathscr{E}_i + \mathscr{E}_e, \qquad \mathscr{E}_i = \frac{3}{2} T n_a, \qquad \mathscr{E}_e = n_a\left[\frac{3}{2} TZR^+(g_e) + \langle E_i \rangle\right], \tag{138}$$

$$P = -\frac{\partial F}{\partial V} = P_i + P_e, \qquad P_i = n_a T, \qquad P_e = n_a TZR^+(g_e). \tag{139}$$

However, while calculating the second order thermodynamic derivatives, like the specific heat, the derivatives of g_e essentially sophisticate the calculations. The result may be expressed in terms of covariances: $\langle \delta^2 i \rangle = \langle (i - Z)^2 \rangle$,

$\langle\delta^2 E_i\rangle = \langle(E_i - \langle E_i\rangle)^2\rangle$ and $\langle\delta i\delta E_i\rangle = \langle(E_i - \langle E_i\rangle)(i-Z)\rangle$. In a similar way one can find the specific heat in an isochoric process, per the unit of volume:

$$C_{Ve} = \frac{\partial \mathscr{E}_e}{\partial T} = n_a \left[\frac{\langle\delta^2 E_i\rangle}{T^2} + \frac{15}{4} Z R^+ - \frac{\left(\frac{3}{2}Z - \frac{\langle\delta E_i \delta i\rangle}{T}\right)^2}{\langle\delta^2 i\rangle + Z R^-} \right], \tag{140}$$

the temperature derivative of pressure:

$$\frac{\partial P_e}{\partial T} = n_a Z \left[\frac{5}{2} R^+ - \frac{\frac{3}{2}Z - \frac{\langle\delta E_i \delta i\rangle}{T}}{\langle\delta^2 i\rangle + Z R^-} \right], \tag{141}$$

as well as the isothermal compressibility:

$$V \frac{\partial P_e}{\partial V} = -\frac{Z^2 n_a T_e}{\langle\delta^2 i\rangle + Z R^-}. \tag{142}$$

For simplicity in the above equations, the contributions due to ion translational motions,

$$C_{Vi} = \frac{3}{2} n_a, \qquad \frac{\partial P_i}{\partial T} = n_a, \qquad V \frac{\partial P_i}{\partial V} = -n_a T, \tag{143}$$

are omitted.

The speed of sound, C_s, is defined in terms of the adiabatic comressibility (at constant entropy), $C_s^2 = \left(\frac{\partial P}{\partial \rho}\right)_{\rm ad}$, which may be parametrized in terms of effective adiabatic index, γ, such that $\gamma \frac{P}{\rho} = \left(\frac{\partial P}{\partial \rho}\right)_{\rm ad}$, herewith ρ is the mass density. To calculate this, one can take Eq.(3.72) from Drake [2006],

$$\left(\frac{\partial P}{\partial \rho}\right)_{\rm ad} = \left(\frac{\partial P}{\partial \rho}\right)_T - \frac{\rho}{C_V}\left[\left(\frac{\partial(\mathscr{E}/\rho)}{\partial \rho}\right)_T - \frac{P}{\rho^2}\right]\left(\frac{\partial P}{\partial \rho}\right)_T.$$

Note that in Drake [2006] both the internal energy and the specific heat are related per a unit of mass, while we relate them to the unit of volume. Now we apply the thermodynamic identity as follows:

$$\left(\frac{\partial(\mathscr{E}/\rho)}{\partial \rho}\right)_T - \frac{P}{\rho^2} = -\frac{T}{\rho^2}\left(\frac{\partial P}{\partial T}\right)_\rho, \tag{144}$$

which gives:

$$\gamma = \frac{\rho}{P}\left(\frac{\partial P}{\partial \rho}\right)_T + \left(\frac{\partial P}{\partial T}\right)_\rho^2 \frac{T}{C_V P}. \tag{145}$$

Note that in the last three equations we denote the derivative at constant V as that at constant ρ and used the derivatives over ρ instead of those over V: $V\frac{\partial}{\partial V} = -\rho\frac{\partial}{\partial \rho}$.

Discussion: Estimating the effect of the Fermi statistics on the ionization degree. Eqs.(131,132) allow us to evaluate the effect of electron Fermi statistics on ionization. From Eq.(146) one can see that for large g_e (Boltzmann gas) the equation, $G(g_e) = 0$, reduces to $\langle i\rangle - g_{e1}/g_e = \delta_{\rm Fe}$, where $\delta_{\rm Fe}$ is a small negative correction to the Fermi function at $g_e \to \infty$:

$$\mathrm{Fe}_\nu(g_e) = \frac{1}{\Gamma(\nu+1)}\int \frac{x^\nu dx}{g_e e^x + 1} = \frac{1}{\Gamma(\nu+1)}\int \frac{x^\nu dx}{g_e e^x} + \delta_{\rm Fe} = \frac{1}{g_e} + \delta_{\rm Fe}. \tag{146}$$

Assuming $\delta_{\rm Fe}$ to be a small increment in the right hand side of Eqs.(131,132), and by finding $dg_e/\delta_{\rm Fe}$ from Eq.(131) at $dV = dT = 0$ and then finding $dZ/\delta_{\rm Fe}$ from Eq.(132) gives:

$$\delta Z = Z\delta_{\rm Fe}\frac{\langle i^2\rangle - Z^2}{\langle i^2\rangle - Z^2 + Z}. \tag{147}$$

The correction is negative as long as $\delta_{\rm Fe}$ is negative. Thus, due to the Fermi gas effects for detached electrons, the ionization degree is always lower than that predicted by the Saha equilibrium equations under the assumption of a Boltzmann electron gas.

This effect should be accounted for while treating the effect of the Fermi statistics for electrons in the equation of state. Specifically, *at constant electron density* the exchange interactions between the electrons increases the electron pressure, but in the partially ionized plasma the magnitude (if not the sign!) of this effect can be compromised by the pressure reduction due to the decrease in electron density.

3.5.5.2 Madelung approximation of electrostatic energy

Coulomb interactions may be accounted for within the Madelung approximation. According to this model, we assume that the net charge of an ion and all of the electrons bound to it is concentrated in one point. For an ion in the charge state i, i free electrons are considered to be coupled with it, being uniformly distributed over its *ion sphere,* which is a sphere of volume $1/n_a$, with the ion at its center. While calculating the electrostatic energy we neglect the Coulomb interactions between charges related to different ion spheres. As a result, we only need to calculate the electrostatic energy of an ion with free electrons coupled to it, which is equal to the energy of electron-electron interactions and the energy required to put the ion at the center of the ion sphere:

$$E_E = \frac{1}{2}\int_{r=0}^{r_{iono}} q_e n_e \varphi_e(r) dV - q_e i \varphi_e(0) = \frac{3}{5}\frac{q_e^2 i^2}{r_{iono}} - \frac{3}{2}\frac{q_e^2 i^2}{r_{iono}} = -\frac{9}{10}\frac{q_e^2 i^2}{r_{iono}}, \tag{148}$$

where r_{iono} stands for the radius of the ion sphere, and the potential of the electrostatic field of electrons is as follows:

$$\varphi_e(r) = ei\left(-\frac{3}{2}\frac{1}{r_{iono}} + \frac{1}{2}\frac{r^2}{r_{iono}^3}\right). \tag{149}$$

Helmholtz free energy. The extra term in the free energy, which accounts for the electrostatic energy, is as follows (see also Eq.(3.50) in Drake [2006]):

$$F_M = -E_M \sum_{i=0}^{i_{max}} i^2 N_i, \qquad E_M = \frac{9}{10}\frac{q_e^2}{r_{iono}}, \qquad r_{iono} = \left(\frac{4\pi}{3} n_a\right)^{-\frac{1}{3}}. \tag{150}$$

Here the Madelung energy,

$$E_M = \frac{9}{10}\frac{q_e^2}{r_{iono}} = \frac{1.8Ry}{(r_{iono}/a)}, \tag{151}$$

characterizes the electrostatic energy related per atomic cell. It is conveniently expressed in terms of the Rydberg constant, $Ry = \frac{q_e^2}{2a} \approx 13.60$ eV, as long as the ion sphere radius, r_{iono}, is related to the Bohr radius, $a = \hbar^2/m_e q_e^2 \approx 0.5 \cdot 10^{-10}$ m.

Ionization equilibrium. Accordingly, in the requirement for ionization equilibrium, $\partial F/\partial N_i - \partial F/\partial N_{i+1} - \partial F/\partial N_e = 0$ (with respect to the reaction $(i) \leftrightarrow (i+1)+e$), the term $\partial F_M/\partial N_i - \partial F_M/\partial N_{i+1}$ will give the contribution of $(2i+1)E_M$ to the left side:

$$-T\log\left[\frac{g_i}{N_i} e^{\Sigma_{j=0}^{i-1} I_j/T}\right] - i^2 E_M + T\log\left[\frac{g_{i+1}}{N_{i+1}} e^{\Sigma_{j=0}^{i} I_j/T}\right] + (i+1)^2 E_M - \mu_e = 0. \tag{152}$$

The solution of the ionization equilibrium, hence, reads:

$$\frac{N_{i+1}}{g_{i+1}} = \frac{N_i}{g_i}\exp\left(-\frac{1}{T}\left(I_i - (2i+1)E_M + \mu_e\right)\right), \tag{153}$$

or, applying recursively and reducing the sum $\sum_{j=0}^{i-1}(2j+1)=i^2$,

$$\frac{N_i}{g_i}=\frac{N_0}{g_0}(g_e)^i\exp\left(i^2\frac{E_M}{T}-\frac{\sum_{j=0}^{i-1}I_j}{T}\right). \tag{154}$$

This may be interpreted as the ionization potential lowering caused by the Coulomb interaction. Each of the potentials, I_i, is reduced by $(2i+1)E_M$. Energy of the ion of the charge state i is:

$$E_i^*=\sum_{j=0}^{i-1}I_j-i^2E_M. \tag{155}$$

This effect shifts the ionization equilibrium towards higher ionization degrees for a given temperature and atomic density.

Partition function. The common multiplier, $\frac{N_0}{g_0}$, in each of Eqs.(154) may be also represented as $\frac{N_a}{S}$. From the normalization condition, $\sum N_i=N_a$, we find that S is a statistical sum:

$$S=\sum_{i=0}^{i_{max}}g_i(g_e)^i\exp\left(-\frac{E_i^*}{T}\right), \tag{156}$$

so that:

$$p_i=\frac{N_i}{N_a}=\frac{1}{S}g_i(g_e)^i\exp\left(-\frac{E_i^*}{T}\right). \tag{157}$$

Differentials along the curve of ionization equilibrium obey the following equations:

$$A_{g_e}\frac{dg_e}{g_e}=A_V\frac{dV}{V}+A_T\frac{dT}{T}, \tag{158}$$

$$A_{g_e}=\langle\delta^2 i\rangle+ZR^-(g_e),\qquad A_T=\frac{3}{2}Z-\frac{\langle\delta i\delta E_i^*\rangle}{T},\qquad A_V=Z+L\langle\delta(i^2)\delta i\rangle, \tag{159}$$

where

$$L=\frac{E_M}{3T}=\frac{3}{5}\frac{Ry[eV]}{T[eV]r_{iono}[a]}. \tag{160}$$

Again, we express the result in terms of covariances, $\langle\delta a\delta b\rangle=\langle(a-\langle a\rangle)(b-\langle b\rangle)\rangle$, and mean values, which are now being calculated using the modified partition functions. Differentiation of mean values, which is necessary for derivation of the above equation on differentials, is not a complicated problem with the following formula:

$$d\langle f_i\rangle=\left\langle\delta f_i\delta\left(\frac{dp_i}{p_i}\right)\right\rangle, \tag{161}$$

where f_i is a function of the only argument i, for example, iE_i or i^2+i.

The ionization equilibrium can be solved using the old technique, i.e. the Newton-Rapson iterations, defined in Eq.(133).

The full Helmholtz free energy now includes the contribution of the electrostatic field energy as in Eq.(150):

$$F=-T\sum_{i=0}^{i_{max}}N_i\log\left[g_i\frac{eV}{N_i}\left(\frac{MT}{2\pi\hbar^2}\right)^{3/2}\exp\left(-\sum_{j=0}^{i-1}\frac{I_j}{T}\right)\right]+F_e-\sum_{i=0}^{i_{max}}N_ii^2E_M. \tag{162}$$

With the ion partition functions in Eq.(157) one can rewrite Eq.(162) in the following form:

$$F=-TN_a\log\left[\frac{eV}{N_a}\left(\frac{MT}{2\pi\hbar^2}\right)^{3/2}\right]-TN_a\log S+\Omega_e, \tag{163}$$

where, again, $\Omega_e = F_e - \mu_e N_a \sum i p_i = F_e - \mu_e N_a \langle i \rangle$.

Plasma thermodynamics and Equation-Of-State. While differentiating Eq.(163) with respect to T and V, again we see that the derivatives by g_e from the second and third terms cancel each other: $g_e(\partial \log S/\partial g_e) = \langle i \rangle = Z$, which is evident from $d \log S = \langle d(\log p_i) \rangle$, and $-g_e g_{e1} \mathrm{Fe}'_{3/2}(g_e) = g_{e1} \mathrm{Fe}_{1/2} = Z$. Accordingly, for the internal energy density, $\mathscr{E}$, and for the pressure, P, we find the following general expressions:

$$\mathscr{E} = -\frac{T^2}{V}\left(\frac{\partial}{\partial T}\left(\frac{F}{T}\right)\right) = \mathscr{E}_i + \mathscr{E}_e, \qquad \mathscr{E}_i = \frac{3}{2} T n_a, \qquad \mathscr{E}_e = n_a \left[\frac{3}{2} T Z R^+ + \langle E_i^* \rangle\right], \tag{164}$$

$$P = -\frac{\partial F}{\partial V} = P_i + P_e, \quad P_i = n_a T, \quad P_e = n_a \left[T Z R^+ - V \langle \frac{\partial E_i^*}{\partial V} \rangle\right], \tag{165}$$

where we assume a general dependence $E_i^* = E_i^*(V)$, such as that in Eq.(155).

In the above equations we add the Madelung corrections to the energy of the electron gas, $\mathscr{E}_e$, and to the electron pressure, P_e, because those corrections are controlled by the electron temperature.

The thermodynamic derivatives can also be expressed in a general form for $E_i^* = E_i^*(V)$ in such a way that one can find the specific heat in an isochoric process, per the unit volume:

$$C_{Ve} = \frac{\partial \mathscr{E}_e}{\partial T} = n_a \left[\frac{\langle \delta^2 E_i^* \rangle}{T^2} + \frac{15}{4} Z R^+ - \frac{\left(\frac{3}{2} Z - \frac{\langle \delta E_i^* \delta i \rangle}{T}\right)^2}{\langle \delta^2 i \rangle + Z R^-}\right], \tag{166}$$

the temperature derivative of pressure:

$$\frac{\partial P_e}{\partial T} = n_a \left[\frac{5}{2} Z R^+ - \left(Z + \frac{V}{T} \langle \delta \frac{\partial E_i^*}{\partial V} \delta i \rangle\right) \frac{\frac{3}{2} Z - \frac{\langle \delta E_i^* \delta i \rangle}{T}}{\langle \delta^2 i \rangle + Z R^-} - \frac{V}{T^2} \langle \delta \frac{\partial E_i^*}{\partial V} \delta E_i^* \rangle\right], \tag{167}$$

as well as the isothermal compressibility:

$$V \frac{\partial P_e}{\partial V} = n_a T \left[-\frac{\left(Z + \frac{V}{T} \langle \delta i \delta \frac{\partial E_i^*}{\partial V} \rangle\right)^2}{\langle \delta^2 i \rangle + Z R^-} + \frac{V^2}{T^2} \langle \delta^2 \frac{\partial E_i^*}{\partial V} \rangle - \frac{V^2}{T} \langle \frac{\partial^2 E_i^*}{\partial V^2} \rangle\right]. \tag{168}$$

Again, for simplicity in the above equations the contributions due to ion translational motions,

$$C_{Vi} = \frac{3}{2} n_a, \qquad \frac{\partial P_i}{\partial T} = n_a, \qquad V \frac{\partial P_i}{\partial V} = -n_a T, \tag{169}$$

are omitted.

To apply the Madelung theory we calculate the first and the second partial derivatives of the energy levels over volume using Eq.(155):

$$\frac{V}{T} \frac{\partial E_i^*}{\partial V} = L i^2, \qquad \frac{V^2}{T} \frac{\partial^2 E_i^*}{\partial V^2} = -\frac{4}{3} L i^2. \tag{170}$$

Eqs.(155,170) allow for specification of all averages and covariances in the expressions for thermodynamic variables and derivatives.

3.5.5.3 Excited states of atoms and ions

To account for excitation of atoms and ions we need not only to involve the distribution over the ionization states, i, but also to quantify the ion distribution over the ground and excited levels.

For simplicity, multiple excitation and autoionization are neglected. We count as the separate excited levels only those for which the principal quantum number, n, of the outermost electron, exceeds that for the atom or ion in its ground state, n_{gr}, by at least one. The excited states, therefore, can be enumerated using two indexes, namely, the ion charge, $i = 0, 1, 2...$, and the principal quantum number $n(i) = n_{gr}(i), n_{gr}(i)+1, ...$

The partition function, p_i, describing the ion distribution over charge state i (recall that $\sum_i p_i = 1$), is now split into smaller populations, $p_{i,n}$, with each relating to a particular excitation level, n. Note that:

$$\sum_n p_{i,n} = p_i, \qquad \sum \frac{p_{i,n}}{p_i} = 1. \tag{171}$$

Which is why the statistical weights, $w_{i,n}$, and, hence, the statistical sum and the partition functions become more complex:

$$w_{i,n} = g_{i,n} \cdot g_e^i \exp\left(-\frac{E^*_{i,n}}{T}\right), \qquad S = \sum_{i=0}^{i_{max}} \sum_{n=n_{gr}}^{\infty} w_{i,n}, \qquad p_{i,n} = \frac{w_{i,n}}{S}, \tag{172}$$

where $g_{i,n}$ stands for the excited level degeneracy, and the energy of the excited state,

$$E^*_{i,n} = E^*_i + E^{exc}_{i,n}, \tag{173}$$

now includes the contribution from the ionization energy, the excitation energy and, generally speaking, the contribution from the electrostatic energy.

We take the values of the excitation energy, $E^{exc}_{i,n}$, from several sources: [spe, 2009], [nis, 2009], [Martin and Zalubas, 1979], [Saloman, 2004]. When some values are missing in those sources, we calculate them using the following formula:

$$E^{exc}_{i,n} = I_i - Ry\frac{(i+1)^2}{n^2}. \tag{174}$$

For the degeneracies of excited levels we assume $g_{i,n} = 2n^2$. To calculate the ground level degeneracy we use the electron configurations from Carlson et al. [1970]. Some of the values of degeneracies have been modified to match the data from Cox [2000].

Up to this point mean values and covariances were being calculated over $i_{max}+1$ possible values of i, but now i is not the only variable that the energy levels depend on. Accounting for excitation levels in the way described above leads to energy levels being a function of i and n. This only increases the number of terms in the statistical sum, but does not change anything in principle. To obtain the thermodynamic variables and the thermodynamic functions we can just replace the mean values over i with mean values over i and n in the Eqs.(164), (165), (166), (167), (168) and substitute the new expression, as in Eq.(174), for energy levels.

3.5.6 Modeling of Radiation Transport and Opacities in Dense Plasmas

3.5.6.1 Multi-group diffusion: governing equations and general ralationships

Governing equation to describe the radiation transport in the multigroup diffusion approximation may be written as the partial differential equation for the spectral energy density, E_ε, which is related to the unit ov volume and the unit

interval of the photon energy, ε. The energy is assumed to be integrated over the solid angle of directions of the photon propagation. Once the spectral energy is integrated density over photon energies, the total radiation energy density is obtained:

$$E = \int_0^\infty E_\varepsilon d\varepsilon. \tag{175}$$

The governing equation for the spectral energy density is as follows:

$$\frac{\partial E_\varepsilon}{\partial t} + \nabla\cdot(\mathbf{u}E_\varepsilon) - (\gamma_R - 1)(\nabla\cdot\mathbf{u})\varepsilon\frac{\partial E_\varepsilon}{\partial\varepsilon} = \text{diffusion} + \text{emission} - \text{absorption}. \tag{176}$$

The second and third terms in the left hand side of Eq.(176) express the time evolution of the spectral energy density resulting from: (1) the radiation advection and comression with the background, which moves with the velocity, $\mathbf{u}$; as well as (2) the photon systematic blue (red) shift in the convergent (divergent) motions, which is analogous to the first order Fermi acceleration of chanrged particles in a moving plasma with the frozen in magnetic field. Herwith $\gamma_R = 4/3$ is the adiabat index of a relativistic (photon) gas. The processes described by the symbolic terms in the right hand side of Eq.(176) are described below.

A set of multi-group equations may be indroduced when we choose a set of frequency groups. Here we enumerate groups with the index, $g = 1, G$. The interval of the photon energies, relating to the gth group is denoted as $[\varepsilon_{g-1/2}, \varepsilon_{g+1/2}]$. The discrete set of unknowns, E_g, is introduced in terms of the integrals of the spectral energy density of the frequency group interval:

$$E_g = \int_{\varepsilon_{g-1/2}}^{\varepsilon_{g+1/2}} E_\varepsilon d\varepsilon. \tag{177}$$

Note that according to Eqs.(175,177)

$$E = \sum_g E_g. \tag{178}$$

Although some of the formulae below are not sensitive to the choice of the group set, here we specify the boundaries of the frequency groups to be such that *the frequency logarithm* is equally spaced (rather than the frequency itself):

$$\log(\varepsilon_{g+1/2}) - \log(\varepsilon_{g-1/2}) = \Delta(\log\varepsilon) = const. \tag{179}$$

Note that as long as the number of groups, G, tends to infinity, the ratio $E_g/\Delta(\log\varepsilon)$ tends to the local value of $\varepsilon E_\varepsilon$, rather than to E_ε. Therefore, the grequency integrals on the equally spaced logarithmic frequency grid allow us to approximate not a spectral energy density, but its product by the photon energy:

$$\frac{E_g}{\Delta(\log\varepsilon)} \approx \varepsilon E_\varepsilon. \tag{180}$$

Now we can integrate Eq.(176) to arrive at the desired set of the multigroup equations:

$$\begin{aligned}\frac{\partial E_q}{\partial t} &+ \nabla\cdot(\mathbf{u}E_g) + (\gamma_R - 1)(\nabla\cdot\mathbf{u})E_g + \\ &+ [-(\gamma_R - 1)(\nabla\cdot\mathbf{u})]\left[\varepsilon_{g+1/2}E_\varepsilon(\varepsilon_{g+1/2}) - \varepsilon_{g-1/2}E_\varepsilon(\varepsilon_{g-1/2})\right] = \\ &= \int_{\varepsilon_{g-1/2}}^{\varepsilon_{g+1/2}} (\text{diffusion} + \text{emission} - \text{absorption})d\varepsilon.\end{aligned} \tag{181}$$

If the number of frequency groups, G is sufficiently large, we employ the approximation as in Eq.(180), which allows us to close Eq.(181) in the following form:

$$\frac{\partial E_q}{\partial t} + \nabla\cdot(\mathbf{u}E_g) + (\gamma_R - 1)(\nabla\cdot\mathbf{u})E_g + \frac{-(\gamma_R-1)(\nabla\cdot\mathbf{u})}{\Delta(\log\varepsilon)} \times \\ \times \left[E_{g+1/2} - E_{g-1/2})\right] = \int_{\varepsilon_{g-1/2}}^{\varepsilon_{g+1/2}} (\text{diffusion} + \text{emission} - \text{absorption}) d\varepsilon. \tag{182}$$

where the values $E_{g\pm1/2}$ should be interpolated from the mesh-centered values E_g towards the frequency values corresponding to the inter-group boundary.

Note that we arrived to difference-differential equation, with the left-hand side including: (1) the conservative advection of the radiation energy density with the velocity $\mathbf{u}$; the work done by the radiation pressure $P_g = (\gamma_R - 1)E_g$; (3) and, as a new element, a linear conservative advection with respect to the log-frequency coordinate. The flux-to-control-volume ratio for the latter effect equals

$$F_{g-1/2} = -(\gamma_R - 1)(\nabla\cdot\mathbf{u})E_{g-1/2}/\Delta(\log\varepsilon).$$

3.5.6.2 Absorption, emission and stimulated emission

An account of the stimuated emission is not less important in the context of the multi-group radiation diffusion than the use of the locat termodynamic equailibrium assumption. As long as the stimulated emission is not often discussed in books regarding the diffisive radiation transport, the code developer may meet a problem while bridging from the absorption coefficient, a_ε, which is calculated from single-photon diagrams of Quantum ElectroDynamics (QED), describing an absorption of a single photon, to the absorption coefficient which is to be used in the radiation trasport simulations. Below we follow Zeldovich and Raizer [1966], in which this subject is presented quite transparently.

Consider a part of the radiation spectrum, of a small width of $\Delta\varepsilon$, about the photon energy, ε, which is resonant with some bound-bound transition, that is

$$\varepsilon = \mathscr{E}_E - \mathscr{E}_A, \tag{183}$$

where $\mathscr{E}_A$ and $\mathscr{E}_E$ are some energy states of an atom (or ion, below we refer the system in these states to as an Emitter or an Absorber). For the sake of simplicity, assume for a while that the degeneracy (multiplicity) is equal to one for both upper and lower level:

$$g_E = 1, \qquad g_A = 1,$$

the use of the same denotation for the degeneracy as for the group index should not confuse the reader.

Assume there is no photon in the initial state. The only radiation process which can occur in this case is the spontaneous emission. Introduce the probability, $dw_p/d\mathbf{o}$, for the spontaneous emission from a single emitter, into the element of a solid andgle, with a given polarization of the photon, p, per a unit of time. The total spontaneous emission from the unit of volume is, hence,

$$N_E \frac{dw_p}{d\mathbf{o}} d\mathbf{o},$$

where $N_{E,A}$ are the abundancies of the emitters and the absorbers correspondingly. The contribution from the spontaneous emission to Eq.(177) should have an extra factor of 4π (due to the integration over the photon directions) times 2 (summation by polarizations) times ε (an emitted energy per a photon) divided by $\Delta\varepsilon$.

Now assume that there are N_p photons in the considered unit volume of plasma, which are in the same as the emitted photon. The rules to express the probability of stimulated emission and absorption may be found in any textbook on QED, particularly, we follow Berestetskii et al. [1982]. The probability of a total emission (spontaneous plus stimilated) equals

$$(1+N_p)\frac{dw_p}{d\mathbf{o}}$$

while the number of the absorbed photons from the same photon state, per an absorber, equals

$$(1+N_p)\frac{dw_p}{d\mathbf{o}}$$

On integrating the total absorbtion and emission by $d\mathbf{o}$ and suming the result up by polarizations, we arrive at the following expression for the emission and absoption terms in Eq.(177):

$$\text{emission} - \text{absortion} = \frac{8\pi\varepsilon}{\Delta\varepsilon}\frac{dw_p}{d\mathbf{o}}(N_E(1+N_p) - N_A N_p)$$

Now we may abandon the assumption about non-degeneracy of the emeitter and absorber states. The extra freedom of the ssytem in its *final* state to occupy any of g_f final states simply multiplies the transition probability by a factor of g_f. Properly applying the factors $g_f = g_E$ and $g_f = g_A$, we find:

$$\text{emission} - \text{absortion} = \frac{8\pi\varepsilon}{\Delta\varepsilon}\frac{dw_p}{d\mathbf{o}}(N_E g_A(1+N_p) - N_A g_E N_p).$$

The population of the photon state, N_p, may be related to the spectral energy density. A radiation energy in the unit volume within the interval of the photon energies, $d\varepsilon$, may be written as $E_\varepsilon d\varepsilon$. The same energy may be obtained if we multiply $2N_p\varepsilon$ (the factor of two accounts for two polarizations of the photon) by the number of photon states within $d\varepsilon$. In the unit volume the number of photon states per the pase volume of $dk_x dk_y dk_z$ equals $dk_x dk_y dk_z/(2\pi)^3$, with k_x, k_y, k_z being the three components of the wave vector. The number of photon states per the interval of dk is, hence, $4\pi k^2 dk/(2\pi)^3 = 4\pi\varepsilon^2 d\varepsilon/(hc)^3$. We find:

$$E_\varepsilon = \frac{8\pi N_p \varepsilon^3}{h^3 c^3},$$

and

$$\text{emission} - \text{absorption} = c a_\varepsilon \left[\frac{N_E g_A}{N_A g_E}\left(\frac{8\pi\varepsilon^3}{h^3c^3} + E_\varepsilon\right) - E_\varepsilon\right]. \tag{184}$$

where the absorption coefficient is introduced (the radiation energy dissipation per a unit of length), which is related to the emission probability as follows:

$$a_\varepsilon = \frac{h^3 c^2 N_A g_E}{8\pi\varepsilon^2}\frac{dw_p}{d\mathbf{o}\Delta\varepsilon}. \tag{185}$$

Under the condition of a local thermodynamic equilibrium, the partition function of atoms and ions in different charge and energy states is governed by the Boltzmann statistics, so that the abundance of each excited level is proportianal to its mutiplicity and to the Boltzbann factor, $\exp(-E/(k_B T))$. Therefore, under these circumstances, using Eq.(183), we find

$$N_{E,A} \propto g_{E,A}\exp[-E_{E,A}/(k_B T)], \quad \frac{N_E g_A}{N_A g_E} = \exp\left[\frac{E_A - E_E}{k_B T}\right] = \exp[-\frac{\varepsilon}{k_B T}]$$

and

$$\text{emission} - \text{absorption} = ca'_\varepsilon \left(\frac{8\pi}{h^3c^3} \frac{\varepsilon^3}{\exp[\varepsilon/(k_BT)]-1} - E_\varepsilon \right). \tag{186}$$

where the effective absorption coefficient, a'_ε, is indroduced, which is corrected to account for stimulated emission:

$$a'_\varepsilon = a_\varepsilon \left(1 - \exp\left[-\frac{\varepsilon}{k_BT} \right] \right) \tag{187}$$

Following Zeldovich and Raizer [1966] we can make now two important conclusions. First, the assoption coefficient to be used in simulationg the radiation transport should be corrected for the stimulated emission, by means of applying a simple correction factor. Second, the spontaneous emission from a plasma with the equilibrium distribution over the energy states is related to the corrected absorption coefficient in sunch a way that their ratio,

$$E_\varepsilon^{(Pl)}(T,\varepsilon) = \frac{8\pi}{h^3c^3} \frac{\varepsilon^3}{\exp[\varepsilon/(k_BT)]-1}, \tag{188}$$

is the spectral energy density distribution of the black body radiation (the Planckian). We will also apply it in a normalized form as follows:

$$E_\varepsilon^{(Pl)}(T,\varepsilon) = \alpha T^4 \frac{15}{\pi^4} \frac{x^3}{\exp(x)-1} \frac{1}{k_BT}, \qquad \alpha = \frac{8\pi^5 k_B^4}{15h^3c^3}, \qquad x = \frac{\varepsilon}{k_BT}. \tag{189}$$

The spectral function in (188) is normalized by a unity:

$$\int_0^\infty \frac{15}{\pi^4} \frac{x^3}{\exp(x)-1} \frac{d\varepsilon}{k_BT} = \frac{15}{\pi^4} \int_0^\infty \frac{x^3 dx}{\exp(x)-1} = 1,$$

about the intregrals like this see §58 in Landau and Lifshitz [1980]. Therefore the total energy density in the Planck spectrum equals αT^4, as it should.

3.5.6.3 Contributions to Opacity: Photoionization and Photorecombination, Effect of the Fermi Statistics.

The formula for the photoionization cross-section from atom (ion) in the charge state, ζ, and with the electron configuration, i, reads:

$$\sigma^{ion}_{\zeta,i\to\zeta+1,j} = \sigma_0 \frac{w_\eta n_\eta}{(\zeta+1)^2} \left(\frac{E_{ij}}{h\nu} \right)^3. \tag{190}$$

Here

$$\sigma_0 = \frac{64\pi}{3\sqrt{3}} \alpha a_0^2 \approx 7.9 \cdot 10^{-22} [\text{m}^2], \tag{191}$$

is the near-threshold ($E_{ij} = h\nu$) semi-classical cross-section of the photoionization of the hydrogen atom from the ground state ($\zeta = 0$, $w_\eta = 1$, $n_\eta = 1$), which may be found in Zeldovich and Raizer [1966]

$$\alpha = \frac{e^2}{\hbar c} \approx 1/137.04 \tag{192}$$

is a fine structure constant, and

$$a_0 = \frac{\hbar^2}{m_e e^2} \approx 0.5 \cdot 10^{-10} [\text{m}] \tag{193}$$

is the Bohr radius. Eqs.(192,193) are only valid in CGSE system of units, however, with the known numerical values for them all the other formulae become insensitive to the choice of units. $h\nu$ is the energy of the photon, which is absorbed in the course of the photoionization. It should exceed the transition energy,

$$h\nu \geq E_{ij}. \tag{194}$$

The transition energy equals:

$$E_{ij} = I_\zeta - E_{\zeta,i} + E_{\zeta+1,j}, \tag{195}$$

where the excitation energy, $E_{\zeta,i} \geq 0$ is introduced, which is the non-negative difference between the energy of the given electron configuration with respect to the ground state energy for the atom (ion) with the charge number equal to ζ.

We consider only the transition to states j, such that the electron confiruration in these states are obtained from the initial configuration, i, by removing a single electron from ηth electron orbit. w_η and n_η in Eq.(190) mean the number of electron on this orbit and its principal quantum number correspondingly.

However, if the photoionization occurs in a plasma and the effects of the Fermi- statistics in this plasma are not negligible, Eq.(190) should be revisited. Recall, that the probabilities of the photon absorption processes in Quantum ElectroDynamics (QED) are calculated as follows (see Berestetskii et al. [1982]): (1) first, the matrix element of perturbation is calculated using the wave functions of bound and free electrons in initial and final states; (2) and then the matrix element should be integrated over all possible value of the free electron momentum, $\mathbf{p}$, as long as there is a free electron in the final state. In this integration the number of free electron states per the unit of space volume is introduced as follows:

$$N_\mathbf{p} = 2\frac{dVd^3\mathbf{p}}{(2\pi\hbar)^3}. \tag{196}$$

To account for the Fermi statistics effect on the electron gas we should note that: (1) the calculation of the matrix element does not change, because in any case the wave function for a free electron is a plane wave; (2) however, the way to calculate the number of states of free electrons change specifically. Specifically, we find that among the states as in Eq.(196) some places are already occupied, with the occupation numbers equal to $1/(\exp[(\varepsilon_e-\mu)/(k_BT)]+1)$ (see Landau and Lifshitz [1980]). Here μ is a chemical potential of a free electron gas and $\varepsilon(\mathbf{p})$ is its non-relativistic kinetic energy of the electron. With this account, the actual number of free electron states is

$$N_\mathbf{p}^{corr} = 2\frac{dVd^3\mathbf{p}_e}{(2\pi\hbar)^3}\left(1-\frac{1}{\exp[\frac{\varepsilon_e-\mu}{k_BT}]+1}\right) = N_\mathbf{p}\frac{\exp[\frac{\varepsilon_e-\mu}{k_BT}]}{\exp[\frac{\varepsilon_e-\mu}{k_BT}]+1}. \tag{197}$$

Therefore, the corrected cross-section becomes:

$$\sigma_{\zeta,i\to\zeta+1,j}^{corr,ion} = \sigma_{\zeta,i\to\zeta+1,j}^{ion}\frac{\exp[\frac{\varepsilon_e-\mu}{k_BT}]}{\exp[\frac{\varepsilon_e-\mu}{k_BT}]+1}. \tag{198}$$

The electron energy here is related to the photon energy, $h\nu$, via the conservation law:

$$\varepsilon_e = h\nu - E_{ij} \tag{199}$$

The contribution from the bound-free transition to the absorption coefficient (not yet corrected for the stimulated emission) may be now written in terms of the corrected cross-section:

$$\kappa_\nu^{bf} = \sigma_0\sum_{\zeta,i,j} N_{\zeta,i}\frac{w_\eta n_\eta}{(\zeta+1)^2}\left(\frac{E_{ij}}{h\nu}\right)^3\frac{\exp[\frac{h\nu-E_{ij}-\mu}{k_BT}]}{\exp[\frac{h\nu-E_{ij}-\mu}{k_BT}]+1}, \tag{200}$$

where $N_{\zeta,i}$ is the concentration of ions with the given electron configuration.

Now, the cross-section of a photorecombination of an electron on a collision with the ion in the state, $\zeta+1,j$ is related to the cross-section as in Eq.(190) in the following way:

$$\sigma^{rec}_{\zeta+1,j\to\zeta,i}=\frac{g_{\zeta,i}}{g_{\zeta+1.j}}\frac{(h\nu)^2}{2m_ec^2}\frac{1}{\varepsilon_e}\sigma^{ion}_{\zeta,i\to\zeta+1,j}=\frac{g_{\zeta,i}p_p^2}{g_{\zeta+1.j}p_e^2}\sigma^{ion}_{\zeta,i\to\zeta+1,j}, \tag{201}$$

where $p_p=\varepsilon_p/c$ and $\varepsilon_p=h\nu$ are the momentum and the energy of a photon correspondingly.

The rule as in Eq.(201) expresses the so-called *cross-invariance* property of all QED processes, which claims the possibility to relate the probability of the direct and inverse processes if we accordingly account for the weights of the initial and final states. This property comes from the point that the matrix elements for the direct and reverse processes are exactly the same, the difference in the cross-sesctions occurs when the matrix element is integrated over the final states, which are different, because the final and initial states in the cross-invariant processes are changed by places. We emphasize, that the relationship as in Eq.(201) holds with the *uncorrected* cross-section $\sigma^{ion}_{\zeta,i\to\zeta+1,j}$, because it is based on the *uncorrected* statistical weight of free electron as in Eq.(196), which results in the multiplier, p_e^2, in the denominator of (201). If we account the correcting factors both in $\sigma^{ion}_{\zeta,i\to\zeta+1,j}$ and in (196) these factors cancel each other and the expression for $\sigma^{rec}_{\zeta+1,j\to\zeta,i}$ keeps unchanged.

Integrate $(p_e/m_e)\sigma^{rec}_{\zeta+1,j\to\zeta,i}$ with the distribution function, $f_{\mathbf{p}}$, of an ideal Fermi gas of electrons,

$$f_{\mathbf{p}}d^3\mathbf{p}_e=\frac{2}{\exp[\frac{\varepsilon_e-\mu}{k_BT}]+1}\frac{d^3\mathbf{p}_e}{(2\pi\hbar)^3}, \tag{202}$$

and represent the result of integration by the electron momentum directions in the following form.

$$\int\frac{8\pi}{c^2h^3}\sigma_0\frac{g_{\zeta,i}}{g_{\zeta+1.j}}\frac{w_\eta n_\eta}{(\zeta+1)^2}\frac{1}{h\nu}\frac{E_{ij}^3d\varepsilon_e}{\exp[\frac{\varepsilon_e-\mu}{k_BT}]+1}. \tag{203}$$

Now discuss the physical meaning of the integrand as in (203). In accordance with Eq.(199), $d\varepsilon_e=d\varepsilon_p=d(h\nu)$. Hence the integrand in (203) may be interpreted as the probability of emission, per a unit time interval, from an ion in the given state, with the emitted photon being within the energy interval $d(h\nu)=d(E_{ij}+\varepsilon_e)$. Once been multiplied by $(h\nu)$, this becomes the emitted energy and once multipled by the emitters density, $N_{\zeta+1,j}$, it becomes the volumetric emissivity due to recombination, per a unity interval of the photon energies, $\mathscr{E}_\varepsilon$:

$$\mathscr{E}^{bf}_\varepsilon=\frac{8\pi}{c^2h^3}\sigma_0\sum_{\zeta,i,j}N_{\zeta+1,j}\frac{g_{\zeta,i}}{g_{\zeta+1.j}}\frac{w_\eta n_\eta}{(\zeta+1)^2}\frac{E_{ij}^3}{\exp[\frac{h\nu-E_{ij}-\mu}{k_BT}]+1}. \tag{204}$$

Discuss briefly the interaction of plasma with the euqilibrium black-body radiation. The absroption of the black-body radiation (again, our Eq.(200) is not yet corrected for a stimulated emission) balances the emissivity, if the following condition is satisfied:

$$\mathscr{E}^{bf}_\varepsilon=\frac{8\pi(h\nu)^3\exp[-(h\nu)/(k_BT)]}{c^2h^3}\kappa^{bf}. \tag{205}$$

Comparing this with Eqs.(200,204), we find that the condition is satisfied, if:

$$\frac{g_{\zeta,i}N_{\zeta+1,j}}{g_{\zeta+1.j}N_{\zeta.i}}=\exp[\frac{-E_{ij}-\mu}{k_BT}]=\exp[\frac{-I_\zeta+E_{\zeta,i}-E_{\zeta+1,j}-\mu}{k_BT}]. \tag{206}$$

This is exactly the expression relating the populations of the electron states in ions, which are in a thermodynamic equilibrium with a Fermi gas of electrons (see Part 1).

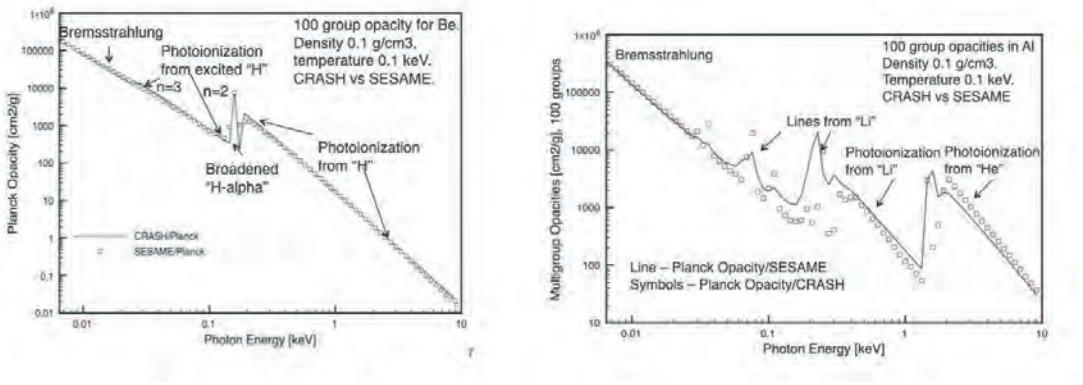

Figure 41: The results of simple opacity model, analogous to that used in the IONMIX code, for low-Z elements: beryllium (the left panel) and aluminum (the right panel). For comparison the SESAME table data are also provided

3.5.6.4 Contributions to Opacity: Free-Free Transitions

The analisys of free-free absorption is mostly based on the same considerations as that for bound-free absorption. The distinction is in the use of the Fermi-Dirac distribution function not only for a free electron in the final state (what we did in the previous section) but also for the electron in the initial state.

The emission from an electron having the velocity, v_E (again the subscript E means "emitter", subscript A means "absorber"), is given well known formula:

$$\frac{dE_\varepsilon}{dt} = \frac{32\pi^2 e^6}{3\sqrt{3}c^3 h m_e^2 v_E} N_e \sum_\zeta N_\zeta \zeta^2 G_f f \tag{207}$$

The way to account for the effect of the Fermi statistics on the population of free electron states and the absorption-emission probability is analogous to that discussed above for the bound-free transitions

3.5.6.5 The main Contribution to the Opacity: Bound-Bound Transitions.

The main contribution to the opacity of plasmas of mid-Z and high-Z ions comes from "bound-bound" transitions between different excited levels of atoms and ions. For low-Z elements for which the number of such levels is moderate and the emission lines can be resolved in the emission spectrum, we successfully used the model analogous to that used in the IONMIX code [J.J.MacFarlane, 1989]. In Fig. 41 we compare the results obtained with this (over-simplified) model and compare the results with SESAME tables. The agreement is good for Beryllium (the left panel) but not so good for Aluminum. For Xenon with Z=54, which radiative property was of a major importance for the project, the opacity data obtained within the framework of simplified model with sparse atomic data present in literature cannot be relied upon. Surprisingly, the attempts to use the opacity data from other sources, such as PROPACEOS, did not provide us with any better data.

The opacity data for Xenon were provided by the ARTEP group. They are calculated within the framework of the STA model (see A.Bar-Shalom et al. [1989]). In Fig. 42, the old CRASH are compared to the STA opacities. The STA opacities are close to the INFERNO [Liberman, 1979] opacities. In the relevant ranges of xenon densities

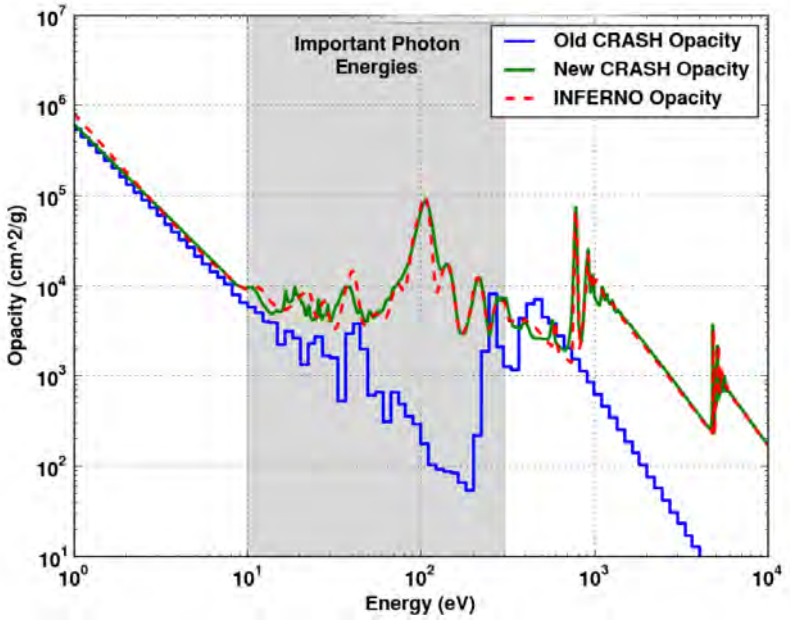

Figure 42: Comparison of the old and new (STA) CRASH opacities for xenon at a density of $\rho = 0.011\,\mathrm{g/cm^3}$ and temperature of 49.99 eV. The STA opacities are close to the INFERNO opacities. In the CRASH relevant photon energy range indicated by the gray box, the STA opacities are significantly higher than the old CRASH xenon opacities.

and temperatures, STA produces higher opacity values than we find by running the CRASH opacity model using the limited available atomic data, e.g. for xenon at a temperature of 49.99 eV and a density $\rho = 0.011\,\mathrm{g/cm^3}$ the STA opacities around photon energies of 100 eV are three orders of magnitude larger. The STA Fe and Ni opacities have recently been compared with other models in Gilles and et al [2011].

3.5.7 Numerical setup of the shock tube experiment

In the baseline experiment of CRASH, a radiative shock is created by means of ten laser beams from the Omega laser facility. The resulting laser pulse irradiates a 20 μm beryllium target with approximately 3.8 kJ laser light of 0.35 μm wavelength for the duration of 1 ns. This first ablates the beryllium, generates a shock, and then accelerates the plasma to over 100 km/s. The front of this plasma drives a shock through a xenon-filled polyimide tube with an initial shock velocity of 200 km/s, see also Drake et al. [2011]. In the shocked xenon region, the shock-heated ions exchange energy with the electrons so that they are also heated, after which both temperatures drop through a radiative cooling layer [Reighard et al., 2007]. The emitted photons from this layer can propagate ahead of the shock and preheat the unshocked xenon. A fraction of this radiation also expands sideways and heats the tube wall, leading to ablation of the polyimide, which in turn drives a wall shock into the xenon [Doss et al., 2009]. In this section, we will describe in more detail how we numerically setup this experiment with our new laser package.

The simulations in this section have been performed with the new 3-D laser package described in 3.5.1.1. The numerical set-up is as follows: A laser pulse of 0.35 μm wavelength irradiates a 20 μm thick beryllium disk for 1 ns. The corresponding critical electron number density, below which all light is absorbed, is $8.9 \times 10^{21}\,\mathrm{cm^{-3}}$. In the radiative

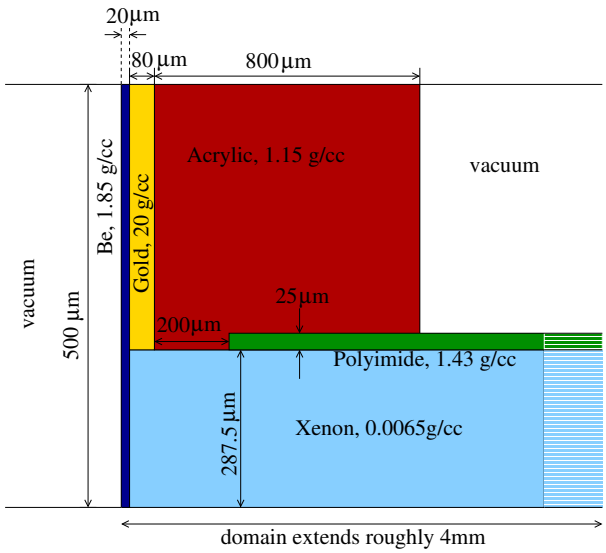

Figure 43: Details of the axi-symmetric radiative shock tube target.

shock experiments, the total laser energy deposition is typically 3.8 kJ, but in our simulations we have to scale this down to arrive at similar results in the shock position. This laser scale factor accounts for that part of the energy for which the laser-plasma interactions cause reflection or absorption into distributions of particles that do not effectively generate ablation. These laser-plasma processes include wave-wave instabilities and related phenomena [Kruer, 2001]. In these runs, we use an energy of 2.7 kJ. The laser spot size is 820 μm full width half maximum (FWHM) diameter. In our application, the Omega laser pulse can be represented by 10 beams with a circular cross-section and the following angles with respect to the shock tube axis: $10.13^\circ, 10.49^\circ, 31.37^\circ, 31.6^\circ, 31.71^\circ, 31.94^\circ, 41.96^\circ, 42.04^\circ, 42.37^\circ, 50.62^\circ$. For computational efficiency, this is modeled using 4 beams with the power per beam weighted by the number of beams approximately at that angle: one beam at 50.6°, three beams at 42.0°, four beams at 31.7°, and two beams at 10.2°. The laser profiles are spatially chosen as super-Gaussian of the order 4.2 and the time profile is split in a 100 ps linear ramp-up phase, 0.8 ns with constant power, and a 100 ps linear decay time. Each beam is discretized with 900 by 4 rays for the radial and angular coordinates of the beam cross-section. The radial beam domain size is up to 1.5 times the FWHM beam radius of 410 μm, while the angular direction is limited to half the domain $[0, \pi]$ due to symmetry considerations. The resulting beam resolution is sufficiently high to obtain a smooth laser heating profile, but also as coarse as possible for computational speed.

3.5.8 Radiative shocks in straight tubes

We use the CRASH code [van der Holst et al., 2011] to simulate both the laser energy deposition and the radiative shock propagation. This code solves the multi-material radiation-hydrodynamic equations in an operator split fashion. For each time step, we split the dynamical equations in the following way: (1) The hydrodynamic equations, level sets, and advection of radiation groups are explicitly solved with a shock capturing scheme. We typically use the HLLE scheme with a Courant–Friedrichs–Lewy (CFL) number of 0.8 and the generalized Koren limiter with $\beta = 3/2$. (2) Optionally we add a frequency advection in the radiation group energies caused by fluid compression. In these runs,

this is switched off. (3) The contribution of the laser heating is explicitly added to the electron internal energy. In our code this energy is split between electron pressure and an extra internal energy to account for EOS corrections like the ionization, excitation, and Coulomb interactions of partially ionized ion-electron plasma. (4) The radiation diffusion, heat conduction, and energy exchanges are solved implicitly. We use a multigroup flux-limited radiation diffusion method in which the flux limiter is the square-root flux limiter [Morel, 2000]. For the electron thermal heat conduction, we use the so-called threshold model [Drake, 2006, van der Holst et al., 2011], for which the heat flux limiter has a value of 0.06. We have made several improvements to the implicit solver including the implementation of the HYPRE preconditioner, to make the solutions more accurate. The results demonstrated below have been produced with these new code changes.

The photon energy range in our multigroup radiation model is 0.1 eV – 20 keV, which is divided in 30 groups. These groups are non-logarithmically distributed to improve the accuracy for absorption edges in the used materials. In our code, the frequency-dependent absorption coefficients are calculated internally and include the effects of Bremsstrahlung, photo-ionization of the outermost electrons, and bound-bound transitions with spectral line broadening. Multigroup opacities are then determined by averaging the absorption coefficients over the photon energy groups. The resulting specific Rosseland and Planck mean opacities for all groups are stored in lookup tables. For xenon opacities, our initially available atomic model was too incomplete, so that the methods used by our code produced substantially inaccurate opacities. We therefore use for xenon high quality opacity tables calculated with the super-transition-arrays (STA) model [Bar-Shalom et al., 1989, 1999, Busquet et al., 2010]. In Fig. 42, the old CRASH are compared to the STA opacities. The STA opacities are close to the INFERNO [Liberman, 1979] opacities. In the relevant ranges of xenon densities and temperatures, STA produces higher opacity values than we find by running the CRASH opacity model using the limited available atomic data, e.g. for xenon at a temperature of 49.99 eV and a density $\rho = 0.011\,\mathrm{g/cm^3}$ the STA opacities around photon energies of 100 eV are three orders of magnitude larger. The STA Fe and Ni opacities have recently been compared with other models in Gilles and et al [2011].

The shock tube is defined on a 2-D axial symmetric computational domain. The size of the domain is $-550 < x < 3850$ along the tube and the radius is limited to $0 < r < 500$, with all distances measured in microns. The base level grid is decomposed of 165×15 grid blocks of 8×8 mesh cells for the x and r directions, respectively. Two levels of adaptive mesh refinement are applied, so that the effective grid resolution is 5280×480 grid cells of approximately $0.8\,\mu$m by $1\,\mu$m. The grid refinement is applied at all interfaces involving xenon or gold. We also apply grid refinement when the xenon mass density exceeds $0.02\,\mathrm{g/cm^3}$ in order to resolve the xenon shock front, the electron-ion equilibration zone, and the radiative cooling layer in the shocked xenon. In addition, all beryllium to the right of $x = -5\,\mu$m is mesh refined during the laser heating. All mentioned grid refinements are applied when any of the mentioned criteria is met in the mesh as well as ghost cells. Note that the effective cell size of $0.8\,\mu$m means that there are 25 cells in the x-direction to span the beryllium disk thickness. This turns out to be sufficient to accurately describe the beryllium ablation and shock breakout.

The boundary conditions at the symmetry axis $r = 0$ are reflective, while we use at all other boundaries extrapolation with zero gradient. However, for the radiation groups, we use a zero incoming flux boundary condition at the outer boundaries, i.e. all radiation leaving the computational domain will not return back.

The simulated laser energy deposition and radiative shock evolution is modeled from 0 to 18 ns. During the first 200 ps the time step is reduced from about 3×10^{-17} s at the very beginning of the simulation and gradually increases towards the end of the 200 ps to a time step based on a CFL of 0.8. The increase (or decrease) in time step is controlled

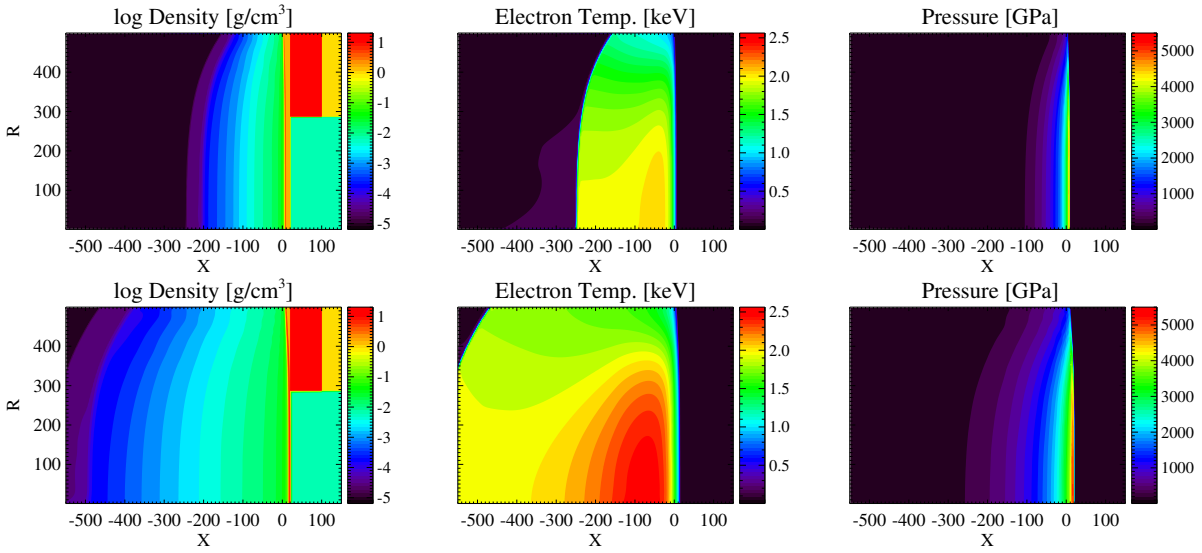

Figure 44: The mass density (left panels), electron temperature (middle panels) and total plasma pressure (right panels) as a function of the x and r coordinates in microns at time $t = 200\,\mathrm{ps}$ (top row) and near the shock breakout time $t = 400\,\mathrm{ps}$ (bottom row).

by the change in the extra internal energy that accounts for the ionization, excitation, and Coulomb interactions. From 200 ps to 18 ns, the time step is set by the default CFL number. This computation was performed on 100 processors of the FLUX supercomputer at the University of Michigan using dual socket six-core Intel Core I7 CPU nodes connected with infiniband and took 39.5 hours. This includes 10 hours and 50 minutes for the laser heating, 21 hours for the Krylov solver, and 6 hours for the setup time of the HYPRE BoomerAMG preconditioner.

In Fig. 44 we show the early time response to the laser heating. The top row is for the density, electron temperature and total plasma pressure at time $t = 0.2\,\mathrm{ns}$. Here one can see the early ablation of the beryllium disk, initially located between $x = 0$ and $x = 20\,\mu\mathrm{m}$. The bottom row shows the same but at time $t = 0.4\,\mathrm{ns}$. The region to the left of the beryllium disk is the laser corona. This region can be split in three main regions [Drake, 2006]: (1) The leftmost low density region is the so-called expansion region in which the plasma expands, but hardly absorbs the laser light. (2) Between the expansion region and the critical density is the absorption region where the laser energy deposition takes place. (3) Between the critical density and the not yet ablated beryllium disk is the transport region in which electron heat conduction transports heat from the low density and hot laser corona to the high density and low temperature beryllium target. It is in this region that the classical Spitzer-Härm (SH) formalism for heat transport overestimates the heat flux for the steep temperature gradients. The artificial heat flux limiter is used to prevent the SH heat conduction from exceeding the free-streaming heat flux (see also Drake [2006]).

The resulting ablation pressure of approximately 5000 GPa drives a shock through the beryllium disk. At $t \approx 400\,\mathrm{ps}$ the shock has reached the right boundary of the beryllium disk. This is the shock breakout time. After this time the high pressure due to the laser heating will further accelerate the shocked plasma. In a forthcoming paper, we will present a more detailed analysis and compare the simulated shock breakout with experiments.

The shock structure at 13 ns is shown in Fig. 45. In the top left panel the materials are displayed: xenon (black),

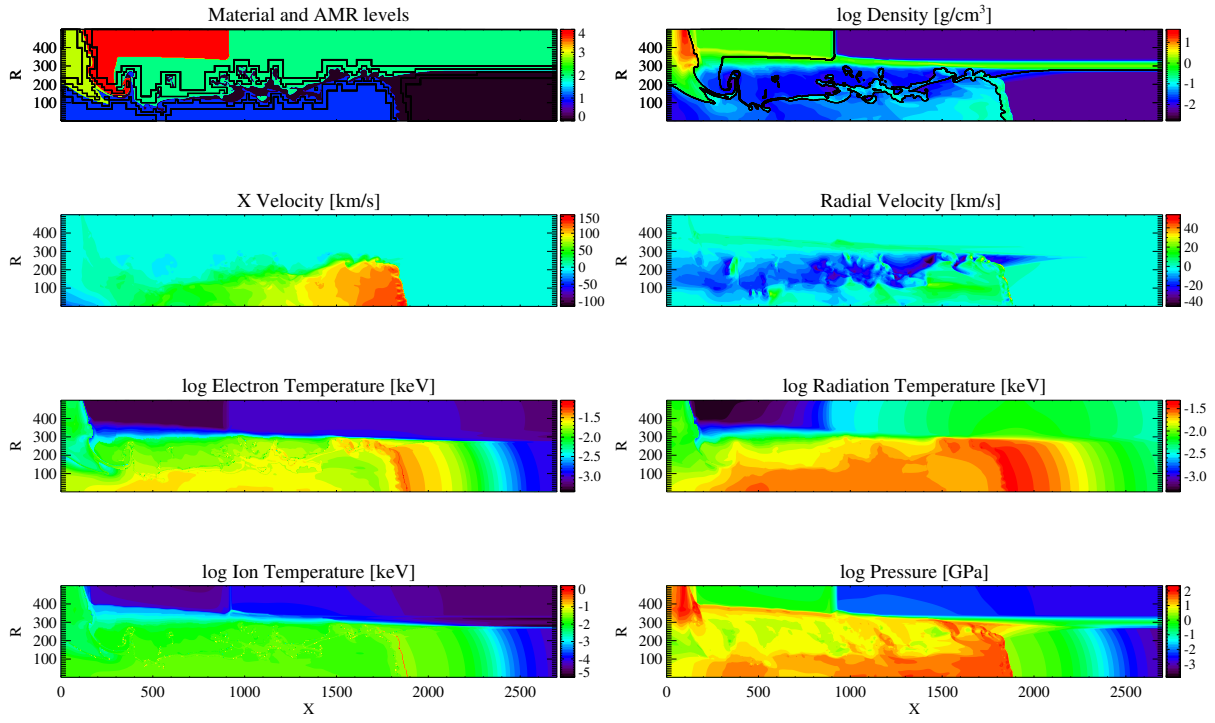

Figure 45: The radiative shock structure at 13 ns. The colors in top-left panel indicate the materials: xenon (black), beryllium (blue), gold (yellow), acrylic (red) and polyimide (green), while the black lines indicate grid resolution changes.

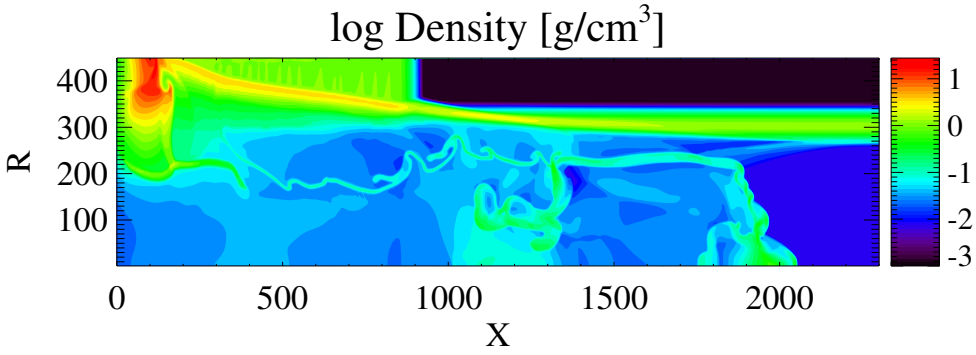

Figure 46: The mass density at 13 ns as a function of the axial distance (x) and radial distance (r) in mircons of the axi-symmetric radiative shock tube target. The simulation shows a shock protusion at $(x,r) \approx (2000, 0)$. These results were produced with an older CRASH code version and reported in Drake et al. [2011].

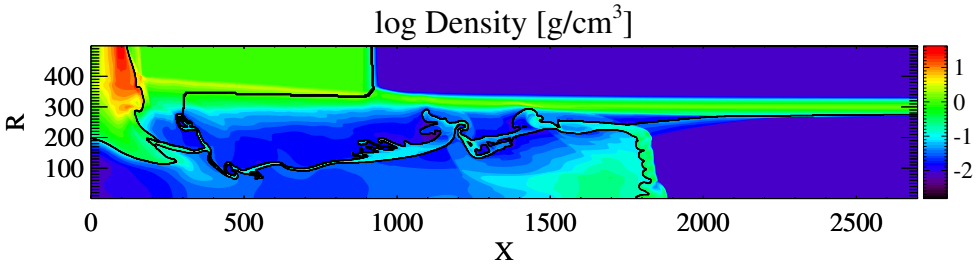

Figure 47: The mass density at 13 ns as a function of x and r in microns. This simulation was performed using the new laser package and the old CRASH xenon opacities instead of the STA opacities in Fig. 45.

beryllium (blue), gold (yellow), acrylic (red), polyimide (green). The two levels of dynamic mesh refinement are indicated by the black lines. The beryllium moves through the polyimide tube and drives like a piston a shock into the xenon. The compressed xenon between this shock front and the beryllium is found around $x = 1850\,\mu$m in the mass density plot of the top right panel. The plot also shows, via a black line, where the material interfaces are. The physics described in the remaining panels is similar to that in previous studies, see van der Holst et al. [2012]. We repeat here only the main results for completeness. The shock velocity at 13 ns has gradually reduced from the early velocity of 200 km/s to about 150 km/s. The bottom right panel shows the pressure jump at the shock, while the bottom left panel shows that the ions are shock heated. In the compressed xenon region this leads first to an electron-ion temperature equilibration due to Coulomb collisions, resulting in a cooling of the ions and heating of the electrons. Further to the left in the compressed xenon region, the electrons cool down by emitting photons. This is called the radiative cooling layer [Reighard et al., 2007]. The emitted photons can propagate ahead of the shock and produce a radiative precursor as depicted in the radiation temperature panel. The sideways propagation of the radiation heats the polyimide tube. The ablation pressure of the polyimide at $x \approx 2000\,\mu$m in the pressure plot then drives a wall shock radially inward as is visible at the same x location in the density and radial velocity plots.

The main goal of this section is to demonstrate that with the new improved physics fidelity and numerical schemes since the CRASH code release of van der Holst et al. [2011], the radiative shock simulations in straight tubes are no

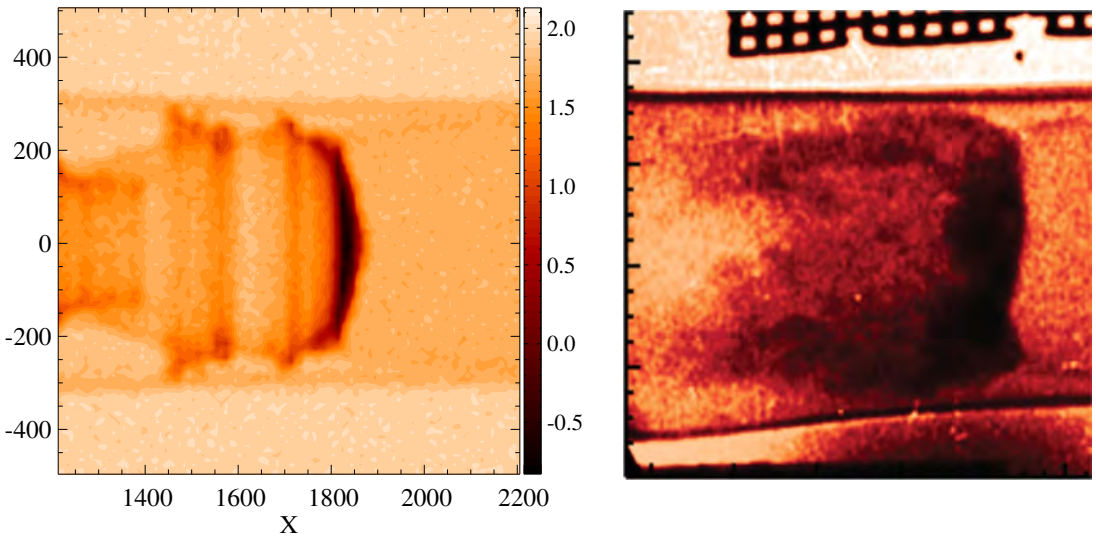

Figure 48: (left panel) Simulated radiograph image at 13 ns. (Right panel) Experimental X-ray radiograph from Doss et al. [2010]. The experimental and simulated set-up is not identical and hence shock positions are different.

longer susceptible to distortion of the dense xenon layer on axis. In Fig. 46, we have reproduced the axi-symmetric and straight shock tube simulation results reported in Drake et al. [2011]. This figure shows the mass density as a function x and r in microns. A significant shock distortion can be seen at $(x,r) \approx (1950,0)$ with the old CRASH code, while in the new code the primary shock front in Fig. 45 is nearly straight and only slightly slanted, but does not display the unwanted protrusion of the shock as shown in Drake et al. [2011]. To demonstrate the improvement due to the STA opacities, we have performed a similar simulation as in Fig. 45 with the CRASH laser package, but using the old xenon opacities. The dense xenon layer is still at $x \approx 1850\,\mu$m. Note that there is in this case still a small shock protusion present on axis. This demonstrates that besides a new laser package also improved xenon opacities were required to improve the quality of the shock tube simulations.

From the radiative shock tube experiments we obtain backlit-pinhole radiograph images [Doss et al., 2009]. These images are produced by transmitting 5.18 keV through the CRASH target and in essence show regions of dense xenon. From these images, we can then deduce the location of the primary and wall shock. With our code, we can produce simulated X-ray radiographs, see the left panel of Fig. 48, and use such images in future code validation and uncertainty quantification. The importance of the code improvements is that there is no longer dense xenon in front of the center of the primary shock at $(x,y) \approx (1850,0)\,\mu$m and hence there is no longer a dark feature ahead of the dense xenon layer in the radiograph as in the experimentally obtained radiograph in the right panel of Fig. 48.

To demonstrate that the shock is also correct at later times, the density and radial velocity at time 18 ns is shown in Fig. 49. The primary shock has reached $x \approx 2500\,\mu$m and is somewhat more slanted. The radially inward moving wall shock is at the far right in these panels.

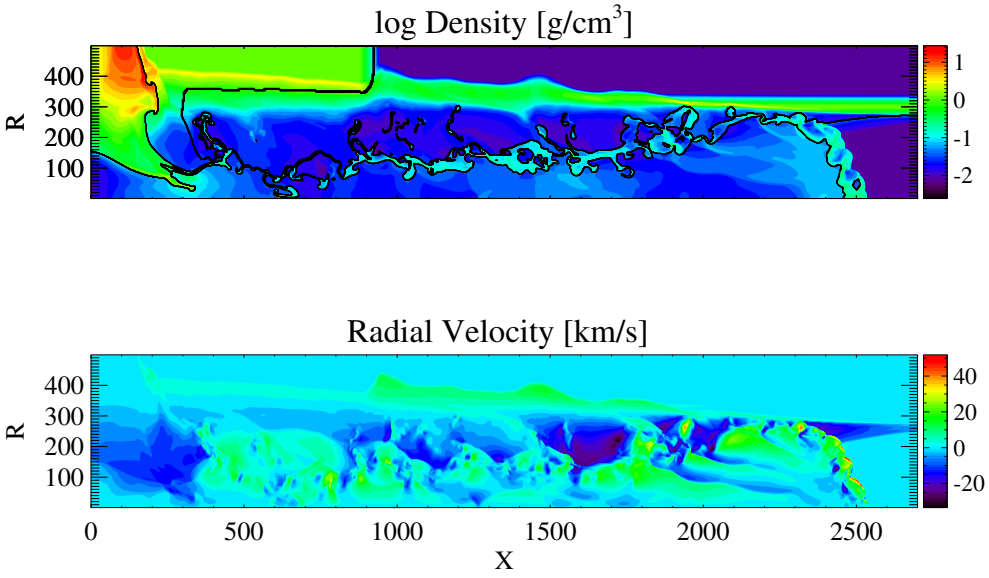

Figure 49: The shock structure at 18 ns.

3.5.8.1 Switching the AMR code in BATSRUS/CRASH completely to BATL

The Block Adaptive Grid Library (BATL) was developed in the CRASH project because the original adaptive mesh refinement (AMR) code in BATS-R-US did not meet the demands of the new project: it could not do AMR in 2D and the dynamic AMR algorithm was rather inefficient. We developed BATL from scratch in about a year to the point that it became fully usable for the CRASH project. Unlike the original AMR code, BATL is a fully object oriented, fully verified and self-contained library. In the last year we completed the transition to BATL, and completely replaced the original AMR code with it. This required implementation of generalized coordinates into BATL, new AMR algorithms, and modification of a large fraction of the BATS-R-US/CRASH code. While this work does not have a significant affect on the current CRASH project, it has far reaching consequences for future development of our code. Bart van der Holst has recently demonstrated good parallel scaling of CRASH with the BATL library from 1 all the way to 260,000 cores on the Jaguar CRAY system at ORNL. Scaling results are shown in Figure 50.

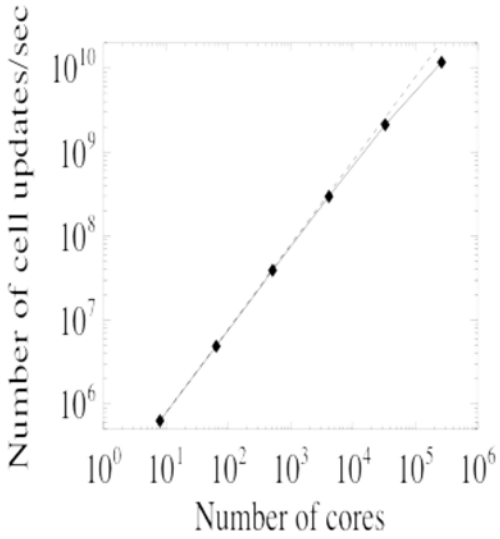

Figure 50: Scaling of CRASH on Jaguar

4 Code Quality Assurance

4.1 Code Verification

To test the CRASH as well as the BATS-R-US and SWMF codes, we have implemented numerous tests. These tests are subdivided in two categories: functionality tests and verification tests. Both test suites are performed automatically and return pass or fail messages depending on whether or not certain predefined tolerance criteria are met. This automated testing process provides software quality confidence especially when used in combination with a software version control system like CVS (Concurrent Versions System) to recover previous correctly performing code.

The functionality tests are performed nightly on several computer platforms with different compilers and numbers of processors. They consist of unit tests and full system tests. Unit tests are designed to test a particular unit, for example a linear equation solver. The full system tests on the other hand, exercise the code in the way end-users will use it for their research applications. We always try to cover as much code as possible with these tests so that we can discover bugs and other unwanted side effects early on.

To test the correctness of the implemented algorithms we have also constructed a suite of verification tests. This suite is executed daily on a dedicated parallel computer and runs specific simulations to quantify against analytic and semi-analytic solutions, whenever possible. The CRASH test repository currently covers a wide range of tests for hydrodynamics, multi-material advection methods, gray and multigroup radiation diffusion, heat conduction, to mention a few. These are performed to test for grid and/or time convergence, as deemed necessary. We also simulate full system laboratory experiment configurations in various geometries, dimensionality, and physics fidelity. The results are either validated against laboratory experiments or simply used to check that the code keeps performing these simulations as expected. Once a week, we also perform a parallel scalability test on a large parallel computer to

verify that the code does not degrade in performance during further development of the software.

In the following sub-sections, we highlight some specific verification tests related to the implicit radiation (Section 4.1.2) and heat conduction (Section 4.1.3) solver. These tests cover both Cartesian and rz-geometry, and some of them also involve the hydrodynamic solver. We demonstrate a 3D full system test in Section 4.1.4 and describe the parallel scalability in Section 3.3.

4.1.1 Error Assessment

For assessment of the accuracy of the solutions in the test suites, an appropriate definition of the numerical errors has to be defined. We use two types of errors to quantify the verification analysis: The relative L1 error is defined as

$$E_{\mathrm{L1}} = \sum_{\alpha=1}^{N} \frac{\sum_{i=1}^{I} |\mathbf{U}_{\alpha i} - \mathbf{V}_{\alpha i}|}{\sum_{i=1}^{I} |\mathbf{V}_{\alpha i}|}, \tag{208}$$

where $\alpha = 1,\ldots,N$ indexes the state variables of numerical solution vector $\mathbf{U}$ and the reference solution $\mathbf{V}$, and $i = 1,\ldots,I$ indexes the grid cells of the entire computational domain. For test problems with smooth solutions, we will also use the relative maximum error defined by

$$E_{\mathrm{L}\infty} = \sum_{\alpha=1}^{N} \frac{\max_{i=1,\ldots,I} |\mathbf{U}_{\alpha i} - \mathbf{V}_{\alpha i}|}{\max_{i=1,\ldots,I} |\mathbf{V}_{\alpha i}|}. \tag{209}$$

Quite often, the reference solution is defined on a grid with higher resolution than that of the numerical solution. In that case, we first coarsen the reference solution to the resolution of the numerical solution.

4.1.2 Radiation Tests

4.1.2.1 Su-Olson Test

Su and Olson [1996] developed a one-dimensional Marshak wave test, to check the accuracy of the scheme and the correctness of the implementation of the time-dependent non-equilibrium gray radiation diffusion model. In this test, radiation propagates through a cold medium that is initially absent of radiation. The equations are linearized by the choice of the specific heat of the material $C_V = 4aT^3$ as well as by setting the Rosseland and Planck opacities to the same uniform and time-independent constant $\kappa_R = \kappa_P = \kappa$. The cold medium is defined on a half-space of the slab geometry $0 \leq x < \infty$. At the boundary on the left, a radiative source is specified, creating an incident radiation flux of $F^{in} = aT_{in}^4$, where $T_{in} = 1\,\mathrm{keV}$. As time progresses, the radiation diffuses through the initially cold medium and by energy exchange between radiation and matter, the material temperature rises. In Su and Olson [1996], a semi-analytical solution is derived for the time evolution of the radiation energy and material temperature. We use this solution for our verification test.

For convenience, we locate the right boundary at a finite distance $x = 5$ cm and impose a zero incoming radiation flux on that boundary. We decompose the computational domain into 6 grid blocks at the base level with 10 cells per block. Between $x = 5/6$ cm and $x = 5/3$ cm, the domain is refined by one level of AMR. During the time evolution, radiation diffuses to the right through the resolution changes. The system is time evolved with the implicit radiation diffusion solver by using a preconditioned conjugate gradient method until the final time 0.02 ns. The solver steps through a series of fixed time steps of 5×10^{-4} ns and we use a Crank-Nicolson approach to achieve second

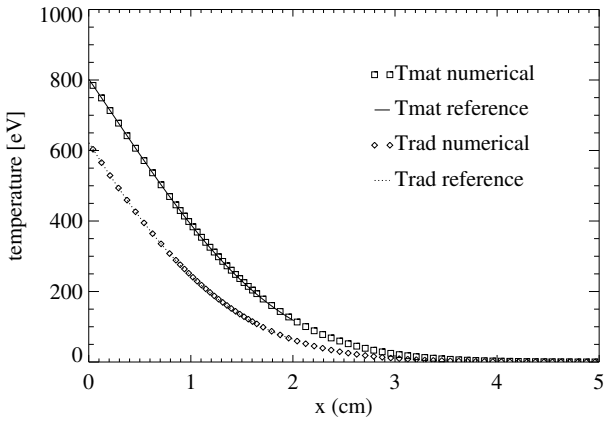

Figure 51: Material (T_{mat}) and radiation (T_{rad}) temperature solution of the [Su and Olson, 1996] non-equilibrium Marshak radiation diffusion problem obtained with the CRASH code on a non-uniform grid. The reference temperatures of the analytical method of Su and Olson [1996] are shown as lines.

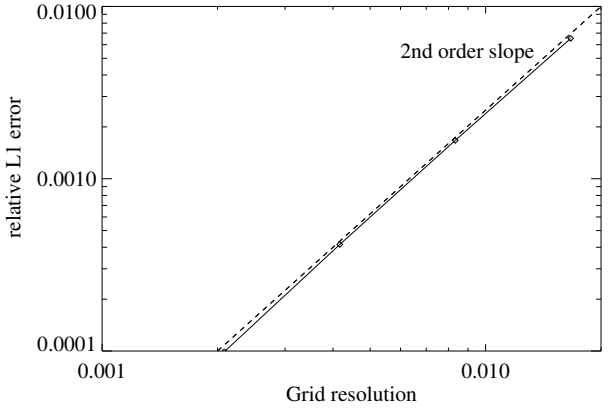

Figure 52: The relative L1 error for the Su-Olson test on a non-uniform grid.

order accurate time-integration. Note that this is possible because coefficients of the matrix to be solved are not time dependent. The computed radiation and material temperatures at the final time are shown in Figure 51 and agree well with the semi-analytical solution. Figure 52 shows the relative L1 error of the radiation and material teperatures versus increasing grid resolution of the base level grid. We did not use the semi-analytical solution as the reference, since it is difficult to get an accurate enough solution with the quadrature method as mentioned by Su and Olson [1996]. Instead, we use a very high resolution (1920 cells) numerical reference solution obtained with the CRASH code. Four different base level resolutions with 60, 120, 240, and 480 cells are used to demonstrate the second order convergence. The time step is proportional to the cell size Δx.

4.1.2.2 Lowrie's Non-equilibrium Radiation Hydrodynamics Solutions

Lowrie and Edwards [2008] designed several shock tube problems for the non-equilibrium gray radiation diffusion

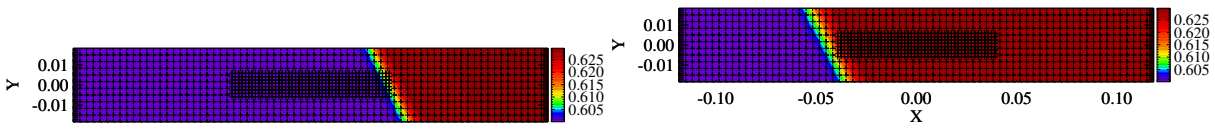

Figure 53: Rotated shock-tube test on a 2D AMR grid based on the Mach number 1.05 non-equilibrium gray radiation hydrodynamic test in Lowrie and Edwards [2008]. Shown is the radiation temperature in color contour at the initial (top panel) and final (bottom panel) times. Black crosses indicate cell centers.

coupled to the hydrodynamic equations that can be used for code verification. These solutions are planar radiative shock waves where the material and radiation temperatures are out of equilibrium near the shock, but are assumed to be in radiative equilibrium far from the shock. Depending on the Mach number of the pre-shock state, a wide range of shock behavior can occur. For the CRASH test suite, we selected a few of the semi-analytic solutions from Lowrie and Edwards [2008]. In this section we will describe the Mach 1.05 flow with uniform opacities as an example. Here the shock is smoothed by energy exchange with the diffusive radiation. Another more challenging Mach 5 problem with non-uniform opacities will be described in Section 4.1.3.3.

The Mach 1.05 test is performed on a 2D non-uniform grid. The initial condition is taken to be the same as the original steady state reference solution. Since the system of equations is Galilean invariant, we can add an additional velocity -1.05 so that the velocity on the left boundary is zero while the smoothed shock will now move to the left. This new initial condition as well as the velocity vector are rotated by $\tan^{-1}(1/2) \approx 26.56^\circ$. This means that there is a translational symmetry in the $(-1,2)$ direction of the xy-plane as shown in Figure 53. The computational domain is $-0.12 < x < 0.12$ by $-0.02 < y < 0.02$ decomposed in 3×3 grid blocks of 24×4 cells each. We apply one level of refinement inside the region $-0.04 < x < 0.04$ by $-0.02/3 < y < 0.02/3$. The initial smoothed shock starts at the right boundary of the refined grid and we time-evolve the solution until it reaches the resolution change on the left as shown in Figure 53. For the boundary conditions in the x direction, we use zero radiation influx conditions for the radiation field, while a zero gradient is applied to the remaining state variables. On the y boundaries, we apply a sheared zero gradient in the $(-1,2)$ direction for all variables.

The hydrodynamic equations are time evolved with the HLLE scheme with a CFL number 0.8. We use the generalized Koren limiter with $\beta = 3/2$ for the slope reconstruction. For the implicit radiation diffusion solver, we use the GMRES iterative solver in combination with a BILU preconditioner. The specific heat is time dependent since it depends on the density, therefore the implicit scheme is only first order accurate in time. To enable second order grid convergence for this smooth test problem, we compensate for this by reducing the CFL number proportional to the grid cell size, in other words $\Delta t \propto \Delta x^2$, so that second order accuracy with respect to Δx can be achieved. We increase the spatial resolution by each time doubling the number of grid blocks at the base level in both the x and y directions.

The convergence of the numerically obtained material and radiation temperatures along the $y = 0$ cut at the final time $t = 0.07$ is shown in Figure 54. The solid, dotted, and dashed lines correspond to the solutions with the 3×3, 6×6, and 12×12 base level grid blocks, respectively. The advected semi-analytical reference solution is shown as a blue line for comparison.

To assess the order of accuracy, the grid convergence is shown in Figure 55 for the three resolutions. The relative L1 error is calculated using the density, velocity components, and both the material and radiation temperatures. We

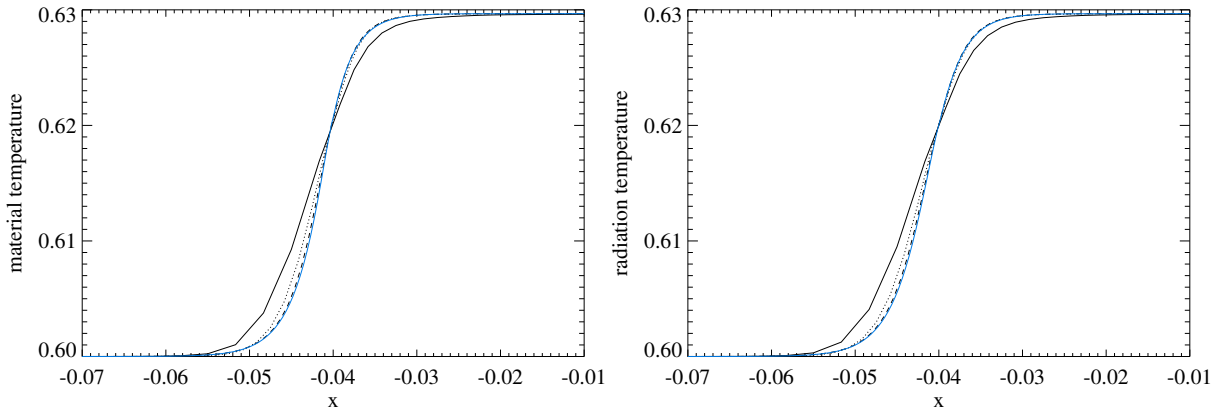

Figure 54: Material (left panel) and radiation (right panel) temperatures for the Mach 1.05 radiative shock tube problem at the final time are shown in the x-direction. The solid, dotted, and dashed lines correspond to three different grid resolutions, respectively. The blue line is the semi-analytical reference solution of Lowrie and Edwards [2008].

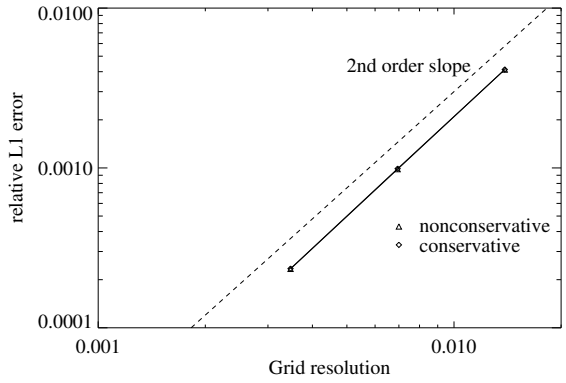

Figure 55: The relative L1 error for the Mach 1.05 non-equilibrium radiation diffusion test on a non-uniform grid. Both non-conservative as well as conservative hydrodynamic schemes are tested.

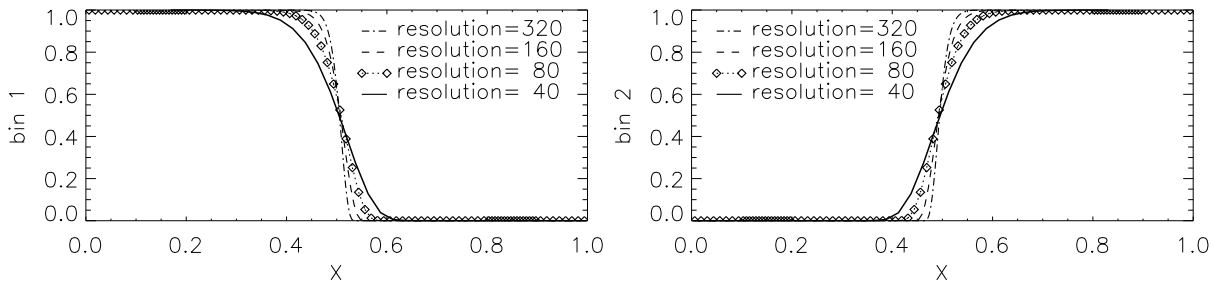

Figure 56: Solutions for the 1D double light front test for 4 different non-uniform grid resolutions. The radiation energy for group 1 (left panel) enters from the left boundary, for group 2 (right panel) it enters from the right boundary. The symbols for base resolution 80 shows one level of grid refinement for $0.1 < x < 0.2$ and $0.8 < x < 0.9$.

obtain second order convergence for both the conservative as well as the non-conservative (using the pressure equation instead of the total energy) hydrodynamic schemes. The latter scheme can be used because in the Mach 1.05 test the hydro shock is smoothed out by the interaction with the radiation.

4.1.2.3 Double Light Front

As a test for the multigroup radiation diffusion model we developed a double light front test problem. This test is used to verify the implementation of both the group diffusion and flux limiters. At the light front, the discontinuity in the radiation field switches on the flux limiter. This limiter is used to correct the radiation propagation speed in the optically thin free streaming regime. With the light front test we can then check that we obtain the speed of light propagation of the front and that the front maintains as much as possible the initial discontinuity.

This test is constructed as follows: We use a 1D computational domain of the size of 1 m in the x-direction. On this domain, we initialize the two radiation group energy densities E_g $(g = 1,2)$ with a very small, positive number to avoid division by zero in the flux limiter. Also the Rosseland mean opacities are set to a small number corresponding to strong radiation diffusion, while the Planck mean opacities are set to zero corresponding to an optically thin medium. The radiation energy density of the first group enters from the left boundary by applying a fixed boundary condition with value one in arbitrary units. On the right boundary this group is extrapolated with zero gradient. Note that these are the proper boundary conditions in the free-streaming limit and not the diffusive flux boundary conditions described in Section 3.2.4. The second radiation group enters from the right boundary with density one, and it is extrapolated with zero gradient at the left boundary. We time evolve both groups for 0.5 m$/c$ seconds. The analytic solution is then two discontinuities that have reached $x = 0.5$ m, since both fronts propagate with the speed of light c.

The computational domain is non-uniform. In the coarsest resolution, there are 10 grid blocks of 4 cells each at the base level. Inside the regions $0.1 < x < 0.2$ and $0.8 < x < 0.9$, we use one level of refinement. The total time evolution is divided into 400 fixed time steps. We use GMRES for the radiation diffusion solver in combination with a BILU preconditioner. For the grid convergence, we reduce the fixed time step quadratically with grid resolution. This time step reduction mimics second order discretization in time. In Figure 56, the two group energy densities are shown for the base level grid resolutions 40, 80, 160, and 320. Clearly, with increasing number of cells, the solution converges towards the reference discontinuous fronts at $x = 0.5$.

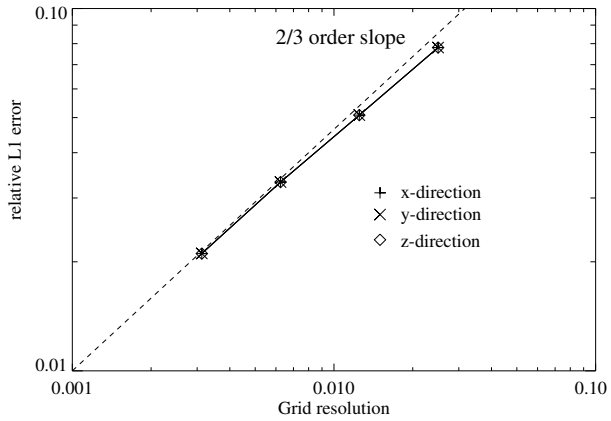

Figure 57: Relative L1 error for the double light front test on a non-uniform grid. The test is performed for the x, y, and z directions.

In Figure 57, the grid convergence is shown for the four resolutions. The relative L1 error is calculated using both radiation group energy densities and compared to the analytical reference solution with the discontinuities at $x = 0.5$. In Gittings [2008], it was stated that for a second order difference scheme the convergence rate for a contact discontinuity is 2/3. Indeed, we find this type of convergence rate, due to numerical diffusion of the discontinuities, for the light front test. We have also performed the tests in the y and z directions to further verify the implementation.

4.1.2.4 Relaxation of Radiation Energy Test

This test is designed to check the relaxation rate between material and radiation. The energy exchange between the material and radiation groups can be written as

$$C_V \frac{\partial T}{\partial t} = \sum_{g=1}^{G} \sigma_g (E_g - B_g), \tag{210}$$

$$\frac{\partial E_g}{\partial t} = \sigma_g (B_g - E_g). \tag{211}$$

For a single radiation group, an analytic expression can be found to describe the relaxation in time. However, for an arbitrary number of groups, a time dependent analytic solution is less obvious, except for some rather artificial cases. Here we make the assumption of an extremely large value of the specific heat C_V to make analysis more tractable. In this case, the material temperature is time independent, so that B_g is likewise time independent. The solution is then $E_g = B_g(1 - e^{-\sigma_g t})$ assuming $E_g(t = 0) = 0$ initially. At time $t = 1/\sigma_g$, the group radiation energy density is $E_g = B_g(1 - 1/e)$. Note that this test only needs one computational mesh cell in the spatial domain. We set $T = 1$ keV and the resulting Planckian spectrum, defined by B_g, is depicted by the dotted line in the left panel of Figure 58. We use 80 groups logarithmically distributed over the photon energy domain in the range of 0.1 eV to 20 keV. The computed E_g at time $t = 1/\sigma_g$ are shown as + points. For the simulation we used the GMRES iterative solver and the Crank-Nicolson scheme. To assess the error, we repeated the test with time steps of $1/20$, $1/40$, and $1/80$ of the simulation time. The second order convergence rate is demonstrated in the right panel of Figure 58.

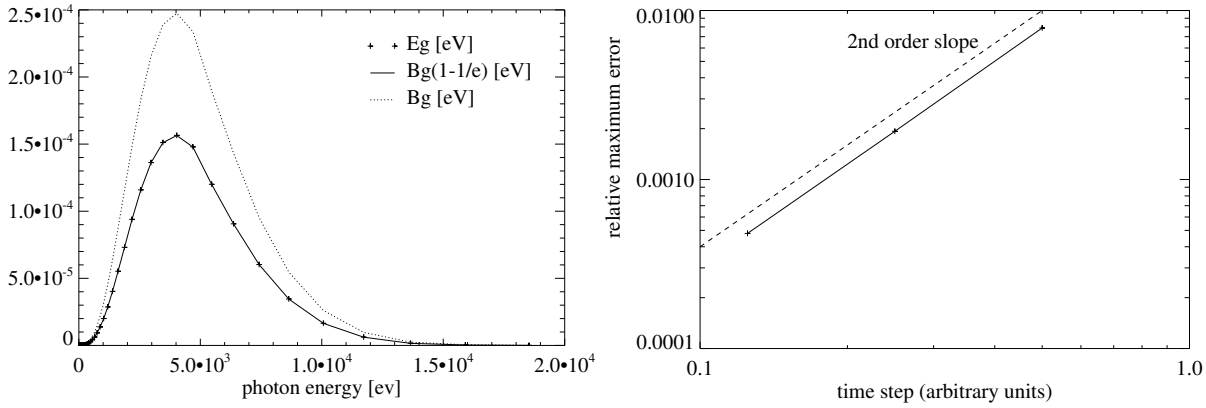

Figure 58: The relaxation of radiation energy test for 80 groups. Left panel is for the time independent spectrum B_g (dotted line) and the group radiation energy solution E_g at time $1/\sigma_g$ (+ points) versus the photon energies after 80 time steps. The analytical reference solution is shown as a solid line. Right panel shows the relative maximum error for 20, 40, and 80 time steps demonstrating second order convergence rate.

4.1.3 Heat Conduction Tests

4.1.3.1 Uniform Heat Conduction in rz-geometry

This test is designed to verify the implicit heat conduction solver in rz-geometry. It tests the time evolution of the temperature profile using uniform and time independent heat conductivity. In rz-geometry, the equation of the electron temperature for purely heat conductive plasma follows

$$C_{Ve}\frac{\partial T_e}{\partial t} = \frac{1}{r}\frac{\partial}{\partial r}\left(rC_e\frac{\partial T_e}{\partial r}\right) + \frac{\partial}{\partial z}\left(C_e\frac{\partial T_e}{\partial z}\right). \tag{212}$$

We set the electron specific heat $C_{Ve} = 1$ and assume electron conductivity C_e to be constant. In this case, a solution can be written as a product of a Gaussian profile in the z-direction and an elevated Bessel function J_0 in the r-direction [Arfken, 1985]:

$$T_e = T_{\min} + T_0\frac{1}{\sqrt{4\pi C_e t}}e^{-\frac{z^2}{4C_e t}}J_0(br)e^{-b^2C_e t}, \tag{213}$$

where $b \approx 3.8317$ is the first root of the derivative of $J_0(r)$. We select the following values for the input parameters: $T_{\min} = 3$, $T_0 = 10$, and $C_e = 0.1$ in dimensionless units.

The computational domain is $-3 < z < 3$ and $0 < r < 1$ discretized with 3×3 grid blocks of 30×30 cells each. In the region $-1 < z < 1$ and $1/3 < r < 2/3$, we apply one level of mesh refinement. We impose a symmetry condition for the electron temperature on the axis. On all other boundaries the electron temperature is fixed to the time dependent reference solution. We time evolve this heat conduction problem with a preconditioned conjugate gradient method from time $t = 1$ to the final time at $t = 1.5$. The Crank-Nicolson approach is used to achieve second order accurate time integration.

The initial and final solutions for the electron temperature are shown in Figure 59 in color contour in the rz-plane. The heat conduction has diffused the temperature in time to a more uniform state. The black line indicates the region in which the mesh refinement was applied. The relative maximum error of the numerically obtained electron temperature

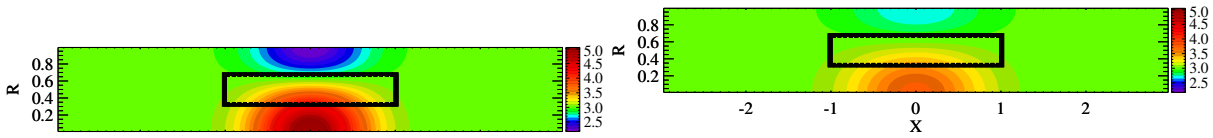

Figure 59: The electron temperature for the uniform heat conduction test on a non-uniform grid in *rz*-geometry. The top panel shows the electron temperature in the initial condition while the bottom panel is the electron temperature at the final time. The black box indicates the region within which the grid is refined by one level.

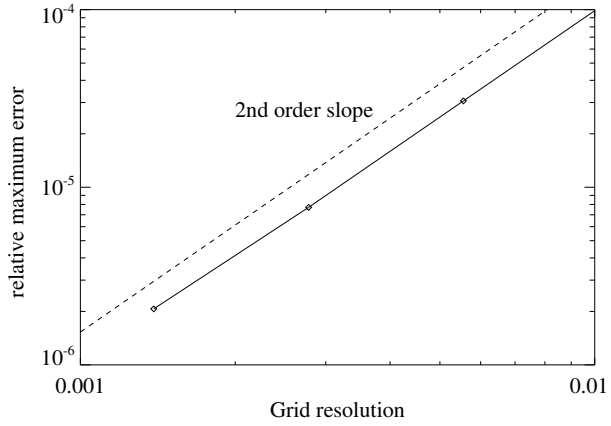

Figure 60: The relative maximum error for the uniform heat conduction test on a non-uniform grid in the *rz*-geometry.

versus the analytical solution is shown in Figure 60. Here we used the non-uniform grid with base resolutions of 90^2, 180^2, 360^2, and 720^2 cells and set the time step proportional to the cell size to demonstrate a second order convergence rate.

4.1.3.2 Reinicke Meyer-ter Vehn Test

The Reinecke and Meyer-ter Vehn [1991] problem tests both the hydrodynamics and heat conduction implementations. This test generalizes the well-known Sedov–Taylor strong point explosion in single temperature hydrodynamics by including the heat conduction. The heat conductivity is parameterized as a non-linear function of the density and material temperature: $C_e = \rho^a T^b$. We select the spherically symmetric self-similar solution of Reinecke and Meyer-ter Vehn [1991] with coefficients $a = -2$ and $b = 13/2$ and the adiabatic index is $\gamma = 5/4$. This solution produces, similar to the Sedov–Taylor blast-wave, an expanding shock front through an ambient medium. However, at very high temperatures, thermal heat conduction dominates the fluid flows, so that a thermal front precedes the shock front. With the selected parameters, the heat front is always at twice the distance from the origin of the explosion as is the shock front.

We perform the test in rz-geometry. The computational domain is divided in 200×200 cells. The boundary conditions along the r and z axes are reflective. The two other boundaries, away from explosion, are prescribed by the self-similar solution. The time evolution is numerically performed as follows: For the hydrodynamics, we use the HLLE scheme with the CFL number set to 0.8. Since this test is performed on a uniform grid without adaptive mesh refinement, we can use the generalized Koren limiter with $\beta = 2$. This is the same as the original Koren [1993] slope limiter. The heat conduction is solved implicitly with the preconditioned conjugate gradient method. The test is initialized with the spherical self-similar solution with the shock front located at the spherical radius 0.225 and the heat front is at 0.45. The simulation is stopped once the shock front has reached 0.45 and the heat front is at 0.9.

A 1D slice along the r-axis of the solution at the final time is shown in Figure 61. We normalize the output similar to Reinecke and Meyer-ter Vehn [1991]: The temperature is normalized by the central temperature, while the density and radial velocity are normalized by their values of the post-shock state at the shock. The numerical solution obtained by the CRASH code is shown as $+$ symbols and is close to the self-similar reference solution, shown as solid lines. Note that the temperature is smooth due to the heat conduction, except for the discontinuous derivative at the heat front. The wiggle at $r = 0.3$ in the density and radial velocity is due to the diffusion of the analytical shock discontinuity in the initial condition during the first few time steps. In the left panel of Figure 62, the spherical expansion of temperature at the final time is shown. Clearly, the Cartesian grid with the rz-geometry does not significantly distort the spherical symmetry of the solution. The spatial distribution of the error in the temperature is shown in the right panel. The errors are largest at the discontinuities of the shock and heat fronts as expected.

A grid convergence study is performed with resolutions of 200^2, 400^2, and 800^2 cells. The relative L1 error in Figure 63 is calculated using the density, velocity components, and the material temperature. The convergence rate is first order due to the shock and heat front.

4.1.3.3 Heat Conduction Version of Lowrie's Test

Any of the verification tests for non-equilibrium gray-diffusion coupled to the single temperature hydrodynamics can be reworked as a test for the hydrodynamic equations for the ions coupled to the electron pressure equation with

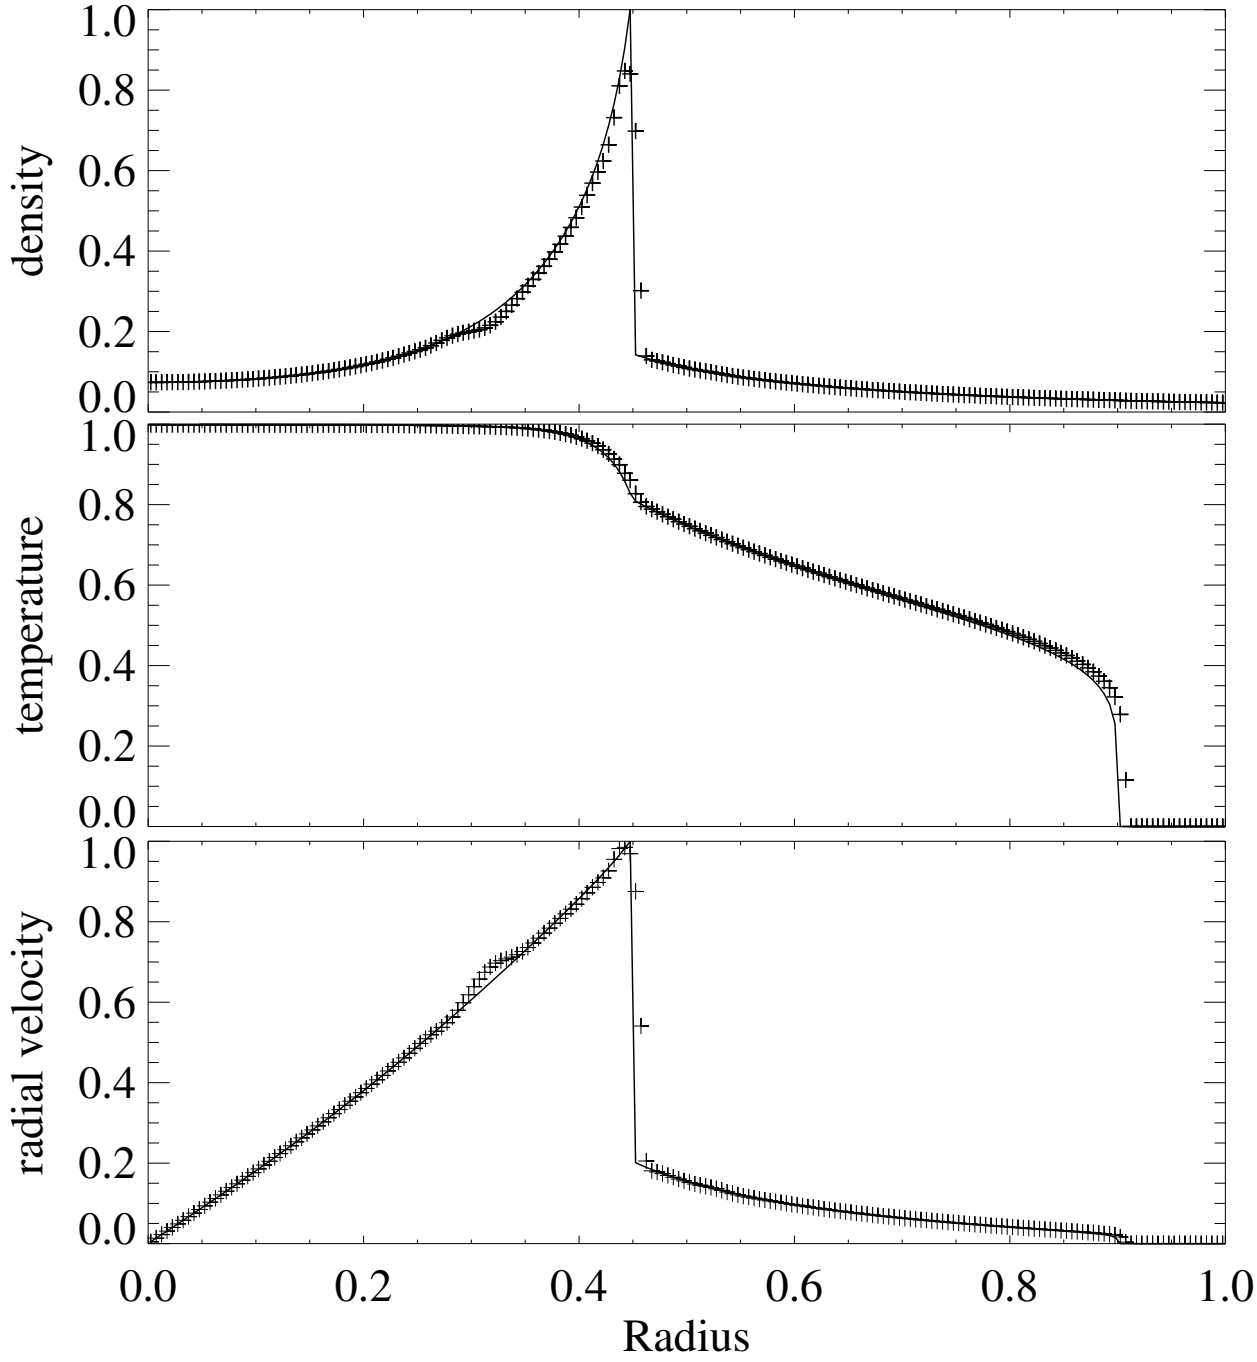

Figure 61: Density (top panel), temperature (middle panel), and radial velocity (bottom panel) along the $z = 0$ cut for the Reinicke Meyer-ter Vehn test in *rz*-geometry. The numerical solution (+ symbols) is at the final time compared to the self-similar analytical reference solution (solid lines).

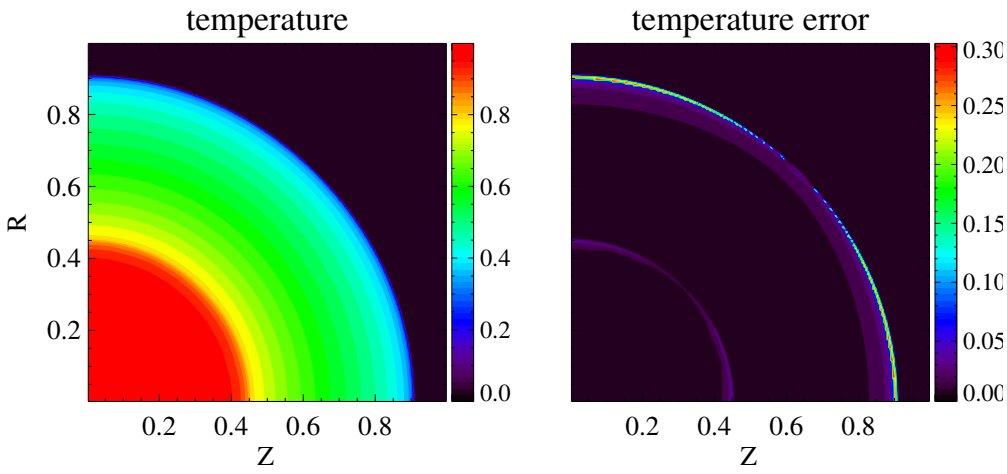

Figure 62: The temperature (left panel) and temperature error compared to the reference solution (right panel) for the Reinicke Meyer-ter Vehn test in *rz*-geometry.

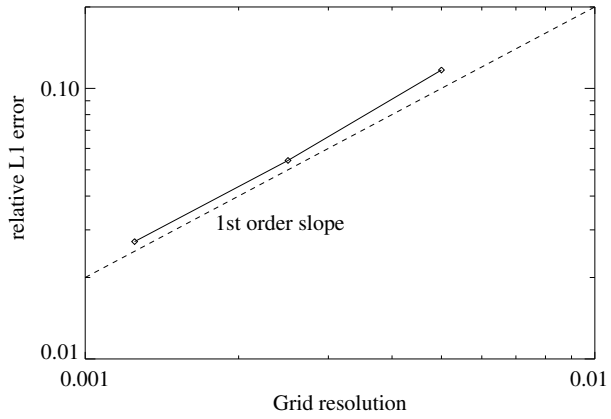

Figure 63: The relative L1 error for the Reinicke Meyer-ter Vehn test in *rz*-geometry.

electron heat conduction and energy exchange between the electrons and ions. As an example, we transform one of the non-equilibrium gray diffusion tests of Lowrie and Edwards [2008] to verify the heat conduction implementation.

The electron energy density equation (18) without the radiation interaction can be written as

$$\frac{\partial E_e}{\partial t} + \nabla \cdot [E_e \mathbf{u}] + p_e \nabla \cdot u = \nabla \cdot [C_e \nabla T_e] + \sigma_{ie}(T_i - T_e), \tag{214}$$

where the heat conduction and energy exchange terms on the right hand side depend on the gradients and differences of the temperatures. The equation for the gray radiation energy density (14) on the other hand depends on the gradients and differences of energy densities. By defining the radiation temperature T_r by $E_r = aT_r^4$ and using the definition of the Planckian $B = aT_e^4$, we can rewrite the energy density equation for the radiation as

$$\begin{aligned} \frac{\partial E_r}{\partial t} + \nabla \cdot [E_r \mathbf{u}] + \frac{1}{3} E_r \nabla \cdot u &= \nabla \cdot [D_r \nabla E_r] + c\kappa_P (aT^4 - E_r) \\ &= \nabla \cdot \left[\overline{D}_r \nabla T_r\right] + c\overline{\kappa}_P (T - T_r), \end{aligned} \tag{215}$$

where $\overline{D}_r = D_r 4aT_r^3$ and $c\overline{\kappa}_P = c\kappa_P a(T^2 + T_r^2)(T + T_r)$ are the new coefficients that appear due to this transformation. The equations (214) and (215) are now of the same form. To translate a gray diffusion test to a heat conduction test, we reinterpret $\overline{D}_r$ as the heat conductivity C_e and $c\overline{\kappa}_P$ as the relaxation coefficient σ_{ie} in the ion-electron energy exchange. In addition, the material temperature T and radiation temperature T_r have to be reinterpreted as the ion temperature T_i and electron temperature T_e, respectively. Note that we also have to relate the electron pressure and internal energy by $p_e = E_e/3$ similar to the radiation field corresponding to $\gamma_e = 4/3$, and let the electron internal energy and electron specific heat depend on the electron temperature as $E_e = aT_e^4$ and $C_{Ve} = 4aT_e^3$, respectively.

As an example, we transform the Mach 5 non-equilibrium gray diffusion shock tube problem of Lowrie and Edwards [2008]. It uses non-uniform opacities that depend on the density and temperature defined by $D_r = 0.0175(\gamma T)^{7/2}/\rho$ and $c\kappa_P = 10^6/D_r$. The above described procedure is used to translate this problem to an electron heat conduction test with energy exchange between the electron and ions. The heat conductivity for this test is $C_e = 4aT_e^3 0.0175(\gamma T_i)^{7/2}/\rho$ and the relaxation coefficient between the electron and ions is $\sigma_{ie} = a(T_i^2 + T_e^2)(T_i + T_e)4aT_e^3 10^6/C_e$.

We perform this Mach 5 heat conduction test on a 2D non-uniform grid. For the initial condition, the 1D semi-analytical steady state reference solution of Lowrie and Edwards [2008] is used. There is a Mach 5 pre-shock flow on the left side of the tube resulting in an embedded hydro shock as well as a steep thermal front (a look at Figure 64 will help to understand this shock tube problem.) We add an additional velocity of Mach -5 so that the pre-shock velocity is zero and the shock is no longer steady, but instead will move to the left with a velocity -5 (in units in which the pre-shock speed of sound is 1). The problem is rotated counter-clockwise on the grid by $\tan^{-1}(1/2)$. The translational symmetry is now in the $(-1,2)$ direction in the xy-plane similar to the Mach 1.05 shock tube problem described in Section 4.1.2.2. The computational domain is $-0.0384 < x < 0.0384$ by $-0.0048 < y < 0.0048$. Inside the area $-0.0128 < x < 0.0128$ and $-0.0016 < y < 0.0016$, we apply one level of refinement. This refinement is set up such that both the heat front as well as the shock front propagate through the resolution change on the left (from fine to coarse) and right (from coarse to fine), respectively. For the boundary conditions in the x-direction, we fix the state on the right side with the semi-analytical solution, but for the left side we use zero gradient. On the y boundaries, we apply a sheared zero gradient in the $(-1,2)$ direction.

For the evolution until the final time $t = 0.0025$, we use the HLLE scheme together with the generalized Koren limiter with $\beta = 3/2$ to solve the hydrodynamic equations. The CFL number is set to 0.8. The heat conduction and

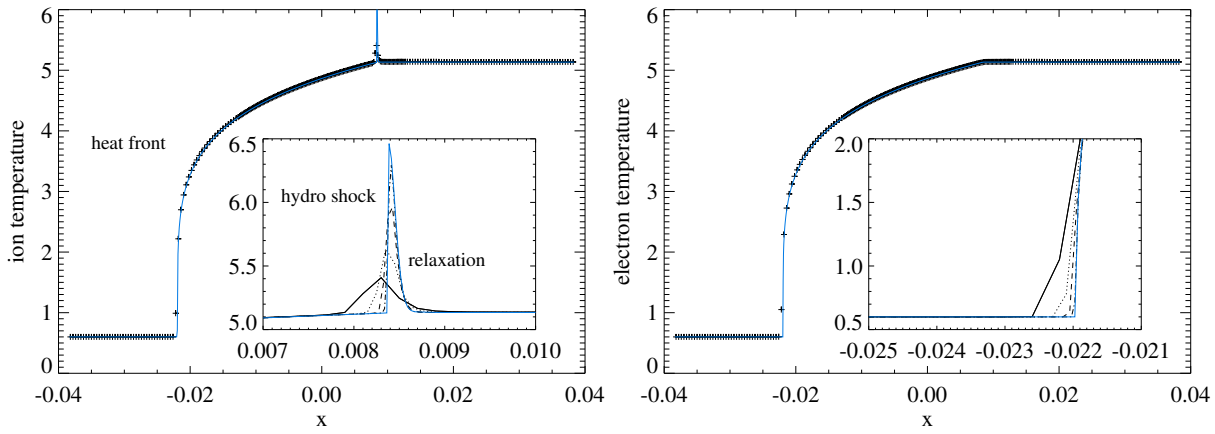

Figure 64: Mach 5 shock tube problem of Lowrie and Edwards [2008] transformed to a non-uniform heat conduction and ion-electron collision frequency test and rotated on a 2D non-uniform grid. Ion (left panel) and electron (right panel) temperatures at the final time are shown in the x-direction. The blue line is the reference solution. In the left panel, the grid convergence near the shock is shown in the inset. In the right panel, a blow-up of the grid convergence to the reference heat front is shown.

energy exchange between electrons and ions are solved implicitly with the backward Euler scheme using the GMRES iterative solver in combination with a BILU preconditioner.

In Figure 64, the electron (right panel) and ion (left panel) temperatures are shown at the final time along the x-axis. The semi-analytical reference solution is shown as a blue line, while the numerical solution is shown with $+$ symbols for a simulation with 192×24 cells at the base level in the x and y direction. The hydro shock is located near $x \approx 0.0085$ and shows up in the ion temperature as a jump in the temperature, followed directly behind the shock by a strong relaxation due to the energy exchange between the ions and electrons. The electron temperature stays smooth due to strong heat conduction. The heat front is seen with a steep foot at $x \approx -0.022$. This front corresponds to the radiative precursor in the non-equilibirum gray diffusion tests of Lowrie and Edwards [2008]. We repeat the test with four different resolutions at the base level: 192×24, 384×48, 768×96, and 1536×192 cells in the x and y direction. The insets in both panels of Figure 64 show the four resolutions as solid, dotted, dashed, dashed-dotted lines, respectively. In the left panel, the zoom-in shows the convergence of the ion temperature towards the embedded hydro shock and the temperature relaxation. In the right panel, the blow-up shows the convergence towards the reference precursor front. Note that no spurious oscillations appear near the shock or near the precursor.

Due to the discontinuity in both the shock and heat precursor, the convergence rate can be at most first order. Indeed, in Figure 65 the relative L1 error shows first order accuracy. The error is calculated using all the density, velocity components and both temperatures. Note that the spike in the ion temperature is spatially so small that a huge number of grid cells are needed to get a fully resolved shock and relaxation state.

4.1.4 Full System Tests

The CRASH test repository contains a range of full system configurations to be used for validation with future laboratory experiments. In Figure 66, we show the configuration of a 3D elliptical nozzle through which a fast shock of

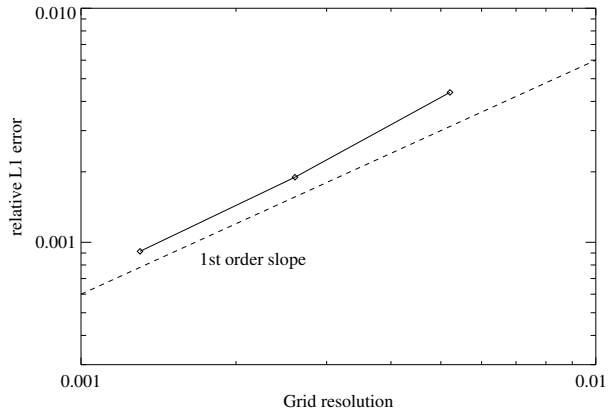

Figure 65: Relative L1 error for the Mach 5 non-equilibrium heat conduction test on a non-uniform grid.

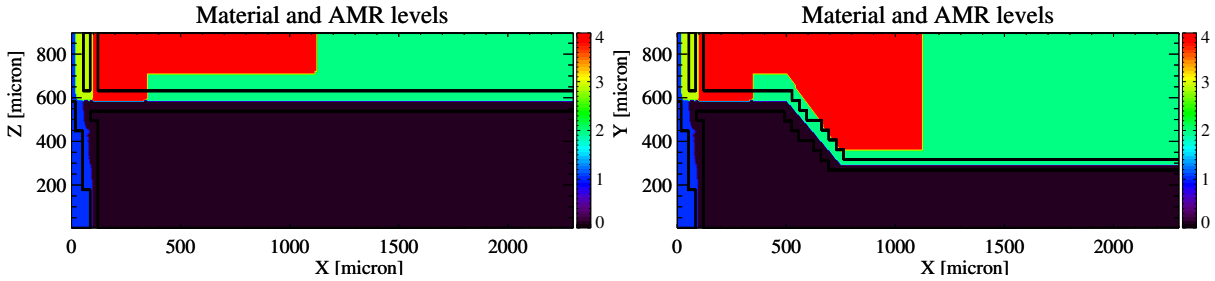

Figure 66: The geometry of the 3D elliptic nozzle experiment after 1.1 ns, consisting of 5 materials: beryllium (blue), xenon (black), polyimide (green), gold (yellow), and acrylic (red) in both panels. The radius of the inside of the polyimide tube is 600 μm in the $y = 0$ plane (left panel). In the $z = 0$ plane (right panel), the radius of the inner tube is 600 μm for $x < 500\,\mu$m, but shrinks to 300 μm beyond $x = 750\,\mu$m. The lines represent the mesh refinement at material interfaces and shock fronts.

the order of 150 km/s will be launched, which is still significantly slower than the speed of light. The shock wave is produced by a 1.1 ns laser pulse from the left with 4 kJ of energy irradiating a 20 μm thick beryllium disk, initially located at $x = 0$. A layer of gold is glued to the plastic tube to protect the outside of the tube from the laser-driven shock. The plastic (polyimide) tube is circular for $x < 500\,\mu$m with a radius of 600 μm. Beyond $x = 750\,\mu$m the tube is made elliptical by flattening the tube in the y-direction by a factor 2.

A laser energy deposition library is currently under construction and the implemented will be reported elsewhere. For this paper, we will instead perform the first part of the simulation with the 2D, Lagrangian, radiation hydrodynamics code HYADES [Larsen and Lane, 1994] to time-advance the heating due to the irradiation by the laser beams and the response of the plasma until 1.1 ns. This laser pulse first shocks and then accelerates the beryllium to the right. After 1.1 ns, the output of HYADES is used as an initial condition of the CRASH code.

This simulation is performed for a two-temperature, electron and ion, plasma. For the radiation, we use the flux limited diffusion approximation with 30 groups. The photon energy is in the range of 0.1 eV to 20 keV, logarithmically distributed over the groups. Due to the symmetry in the problem we only simulate one quadrant ($y > 0$ and $z > 0$), with reflective boundary conditions at $y = 0$ and $z = 0$. At all other boundaries we use an extrapolation with zero gradient for the plasma and a zero incoming flux boundary for the radiation. The domain size is $[-150, 3900] \times [0, 900] \times [0, 900]$ microns for the x, y, z coordinates. The base level grid consists of $120 \times 20 \times 20$ blocks of $4 \times 4 \times 4$ mesh cells. One level of dynamic mesh refinement is used at material interfaces and the shock front. Overall, the effective resolution is $960 \times 160 \times 160$ cells and there are approximately 4.5 million finite volume cells. The hydrodynamic equations are solved with the HLLE scheme with a CFL number 0.8 together with the generalized Koren limiter with $\beta = 3/2$. The diffusion and energy exchange of the radiation groups as well as the heat conduction are solved with the decoupled implicit scheme using a Bi-CGSTAB iterative solver. The simulation from 1.1 ns to 13 ns physical time took 1 hour and 55 minutes on 480 cores of the FLUX supercomputer at the University of Michigan.

In Figure 67, we show the shock structure at 13 ns. The accelerated beryllium compresses the xenon directly to the right of the interface, which is seen as a high density plasma near $x = 1700\,\mu$m in the top right panel of Figure 67. This drives a primary shock and the velocity jump at $x \approx 1700\,\mu$m is seen in the middle left panel. Behind the shock front, the ions are heated as depicted from the middle right panel, followed directly behind the shock by cooling due to the energy exchange between ions and electrons. Early on, electron heating produces ionization and emission of radiation, and the radiation in turn heats and ionizes the material ahead of the primary shock. The radiation temperature, a measured by the total radiation energy density, is shown in the bottom left panel. The photons will interact again with the matter, sometimes after traveling some distance. This is the source of the wall shock seen ahead of the primary shock [Doss et al., 2009, 2011a]: photons traveling ahead of the shock interact with the plastic wall, heat it, and in turn drive a shock off the wall into the xenon. The ablation of plastic is depicted in the top left panel as a radially inward moving polyimide (in green color) near and even ahead of the primary shock. The compressed xenon due to plastic ablation is seen in the top right panel as a faint density feature that is ahead of the primary shock front, between $x = 1700\,\mu$m and $x = 2000\,\mu$m. Interaction between the photons and matter is also seen by the radiative precursor to the right of the radiative shock elevating the electron temperature ahead of the shock in the bottom right panel. This is due to the strong coupling between the electrons and radiation field. The reader is referred to Drake et al. [2011] for more details on radiative effects in radiative shock tubes.

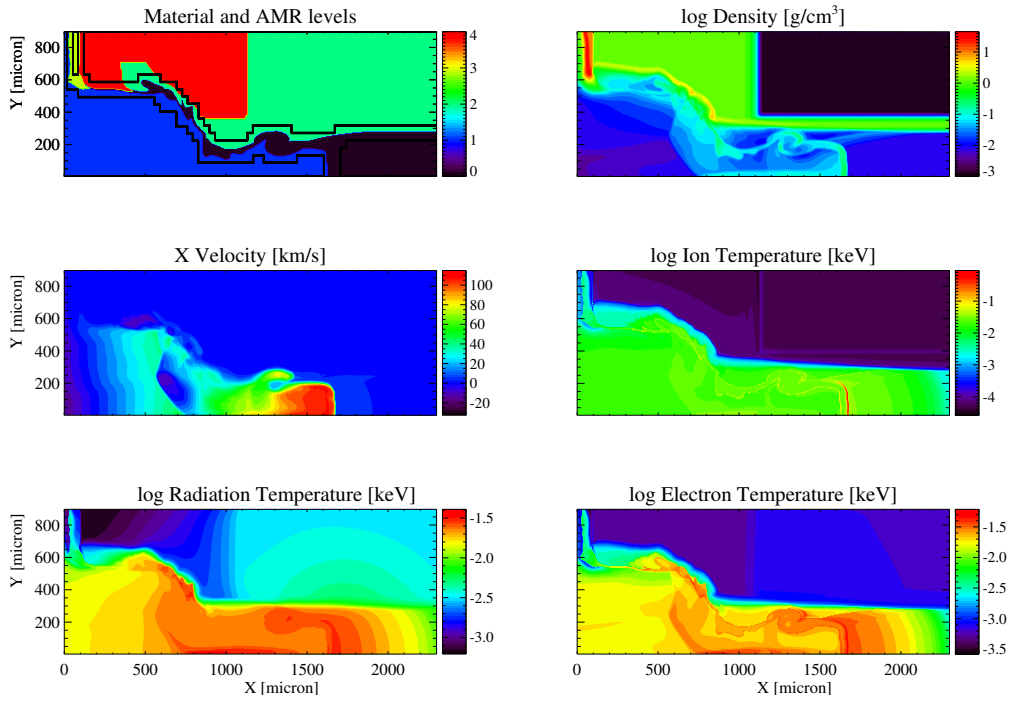

Figure 67: Simulated radiative shock structure at 13ns in a 3D elliptic nozzle consisting of the 5 materials indicated in Figure 66. The plots show in the xy-plane in color contour the variables indicated in the plot title. The primary shock is at $x \approx 1700$.

4.2 Validation

The entire CRASH project can properly be viewed as a validation study, assessing the ability of the CRASH code to correctly predict the observed behavior of our radiative shock experiment. And indeed, during our early years during which we had to rely on a separate Lagrangian code to calculate the laser energy deposition and early radiative-shock behavior, our calculations were *not* validated. This drove the project, in both its focus on certain code improvements that ultimately succeeded and in the selection of experiments, most of which involved obtaining more extensive data regarding the behavior of our primary radiative shock system.

However, we found little time to pursue component validation in a formal sense, by which is usually meant the modeling of experiments that stress one or another aspect of the physics in the code. The shock breakout data we obtained could be used to focus on validation of the laser-energy-deposition calculation, and was used in that way for our initial simulations using Hyades [Kuranz et al., 2013]. However, we did not find time to return to this issue for the CRASH simulations once the CRASH laser package was working. Other programmatic priorities intervened. We did complete a validation study of the hydrodynamic solver, described next.

Validation is often also used in a less rigorous sense, to describe the successful application of a code to model various experiments. Predictive science purists sometimes are disdainful toward such efforts, but they do have merit in the sense that they stress the physics in a code in new ways. One of our annual reviewers opined that the best way to assure that a code was good was to run it for many different problems. While we appreciate that the formal approach is more rigorous, and may matter a lot for precision applications, we also think his point is of some value. For whatever reason, he did not return as a panel member the next year.

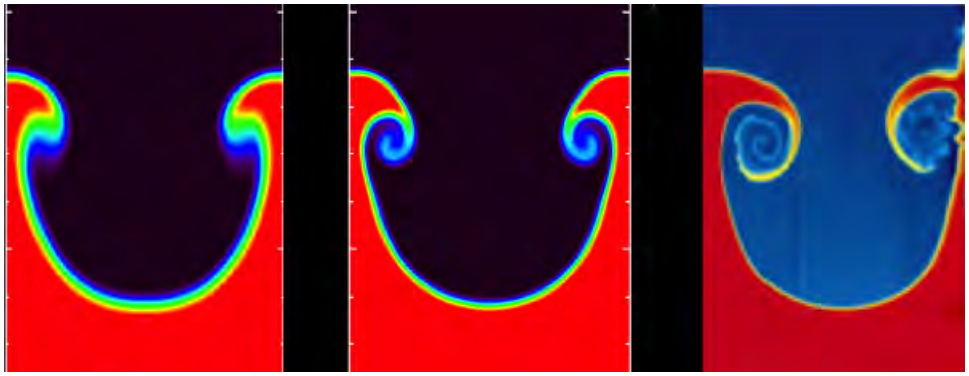

Figure 68: Flow morphology obtained by CRASH simulations with 128 grid points per wavelength (left) and 256 grid points per wavelength (center). The plot on the right shows the result from the experiment. Although the amplitude of the instability is similar in all three plots, the experiment shows much more detail, especially in the rollup.

4.2.1 Richtmyer Meshkov validation study

Component validation, in which one compares a simulation code to data from field experiments that test one of the physics components of the code, are a recognized part of uncertainty quantification. The CRASH project has had little time for such studies, under the pressure of the need to complete integrated simulations and to be ready to predict the

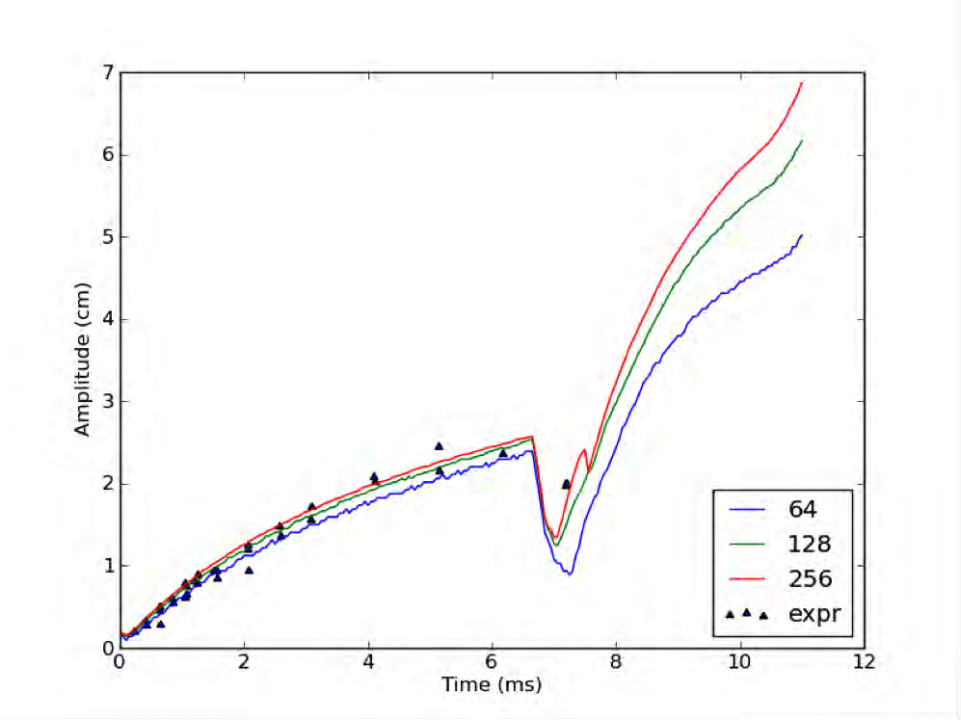

Figure 69: Amplitude of the perturbations vs. time for the three CRASH simulations with resolutions of 64, 128, and 256 grid points per wavelength. The triangles show the data from the experiment. Agreement is very good before re-shock, and the one experimental data point after re-shock is consistent with the highest resolution simulation.

year-5 experiment. Despite the competing priorities, we have completed the component validation study described here.

An effort has been made to validate the hydrodynamic modules in the CRASH code by simulating the vertical shock tube experiments of Richtmyer-Meshkov instabilities carried out by Collins and Jacobs at the University of Arizona. The two materials used were air and SF_6, providing a density ratio of approximately 1:5. A two-dimensional sinusoidal perturbation was applied to the interface by shaking the tube. The shock propagated downward through the air, refracted through the interface into the SF_6, and excited the instability, causing the amplitude of the perturbation to grow. The nonlinear behavior of the instability produced the spike and bubble morphology characteristic of the RM instability, with "mushroom cap" structures at the tips of the spikes. The shock was then allowed to reflect off the bottom of the shock tube, propagate upward, and re-shock the interface. At this point the flow became very complex,turbulent, and fully three dimensional.

The CRASH simulations were carried out in two-dimensional Cartesian geometry, which should be adequate to model at least the early stages of the experiment before re-shock. The grid resolution needed to obtain a converged growth rate and to reproduce the growth rate seen in the experiment was measured by varying the wavelength of the perturbation from 64 to 256 grid points. It is not possible to obtain convergence for the flow morphology, since smaller-scale structures will continue to appear as the grid is refined, given the lack of a physical dissipation mechanism in the code.

The results at resolutions of 128 and 256 grid points per wavelength are shown in Figure 68. Also shown is the result from the experiment. In these plots, the spike is at the edge of the grid and the bubble at the center. The growth rate obtained by the code is consistent with that seen in the experiment, but the details of the flow are slightly different. There is much less roll up in the simulations, especially at a resolution of 128 grid points, indicating that the numerical dissipation in the code at this resolution is larger than the physical dissipation in the experiment. The amplitude of the instability as a function of time at each grid resolution is plotted in Figure 69, along with the data from the experiment. The results before re-shock agree very well with the experiment. After re-shock, it is difficult to tell how good the agreement is, since there is only one experimental data point. However, at this point the highest-resolution simulation is consistent with the experiment. The amplitude has not quite converged at these resolutions, but it appears to be close. The lowest-resolution simulation produced a considerably slower growth rate after re-shock. It would be somewhat surprising if the simulations and experiment agreed well after re-shock, since the three-dimensional behavior of the experiment at this time cannot be reproduced by the two-dimensional simulations.

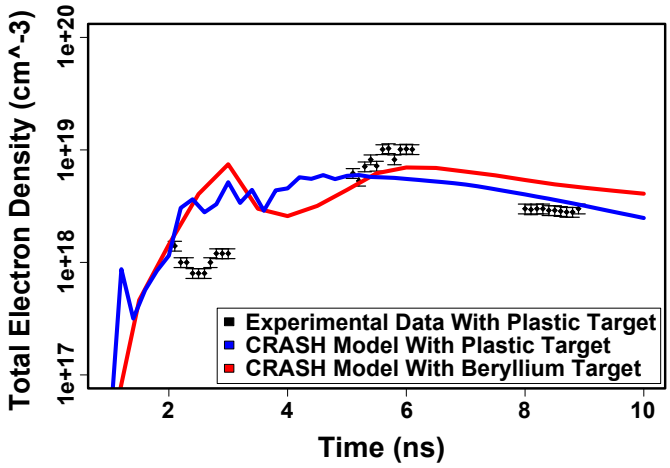

Figure 70: Comparison of CRASH simulation results and electron density data at 4 mm from a target surface, for collisionless-shock development experiments by the ACSEL collaboration.

4.2.2 Other applications

The CRASH code is now in active use for modeling HED experiments. A recent example where multidimensional effects are critical is the collisionless-shock experiments on which we collaborate, as part of the Astrophysical Collisionless Shock Experiments with Lasers (ACSEL) project led by Hye-Sook Park of LLNL. This project sought modeling support from us because it was not possible to obtain sufficient support from LLNL. (This is common in the NNSA labs, where obtaining simulation and design support even for internal basic science experiments is often hampered by their perceived lack of programmatic priority. One of the ways we can increase our interactions with the laboratories is to provide such support in collaborative projects.) Figure 70 shows the density from one of these simulations, used to help refine experimental plans for a late 2011 ACSEL shot day.

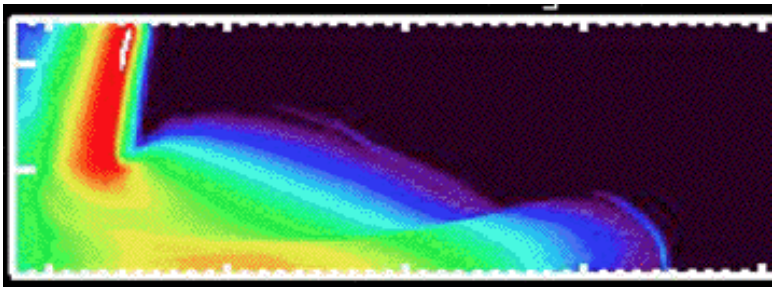

Figure 71: Color contour pot showing log density in simulations of an experimental design for a long-lived jet, several of which can drive a rotating, turbulent plasma.

Other uses of CRASH to model our HED experiments were shown above; we are using it for all of them. CRASH aided in examination of design choices for radiative reverse shock experiments. CRASH is being actively used to evaluate target designs for magnetized flow experiments. (See Fig. 71) CRASH has also modeled radiative-shock experiments using Ar gas rather than Xe, for which the simulations are essential to enable interpretation of transverse optical pyrometry data. CRASH has just begun to be used by our first non-Michigan group, at Florida State University, to model the Diverging Supernova Rayleigh Taylor experiments accepted for NIF. Figure 72 shows results from a preliminary run. The code is also being used to model previous, planar, Supernova Rayleigh Taylor experiments both as a validation study and to explore the potential generation of magnetic fields in these systems [Kuranz et al., 2010]. Figure 73 shows results from a simulation.

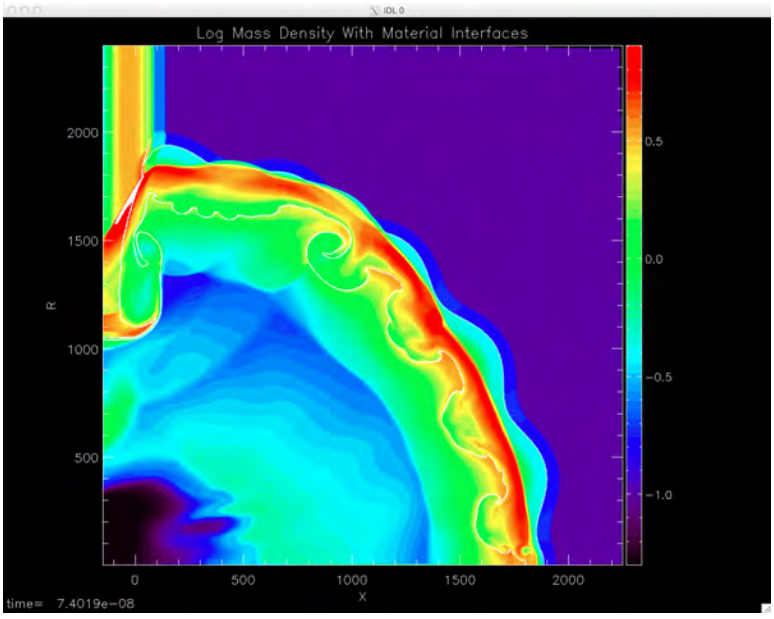

Figure 72: Color contour pot showing log density in simulations of an experiment design to produce a diverging, unstable explosion on NIF relevant to supernovae.

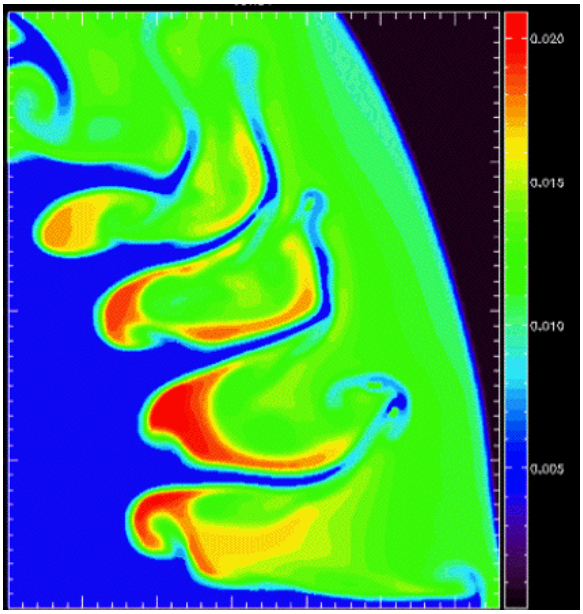

Figure 73: Color contour pot showing log density in simulations of past experiments producing blast-wave-driven instabilities relevant to supernovae.

4.3 Solution Verification for the CRASH Code

It is easy to assess the errors and grid convergence properties for a simulation code in cases where an analytic solution exists. The error is just the difference between the simulated solution and the analytic solution, and the rate at which the simulated solution approaches the analytic solution provides the rate of convergence of the method. For the case where no analytic solution exists, other techniques must be used. An unreliable method would be to estimate the error in the general case by running a number of problems with analytic solutions that are chosen to be as close as possible to the problem of interest. However, when there are extreme non-linearities in the physics and the chance for complex interactions between various physics modules, this method has the potential to severly underestimate the numerical errors.

One method proposed by Roache is to compute a quantity called the Grid Convergence Index (GCI), which measures the relative error between solutions on different grid sizes. This can then be used to construct approximate error bars for the solution. This quantity is defined as

$$\mathrm{GCI} = \frac{F_S}{R^p - 1} \left| \frac{y\left(\Delta x_f\right) - y\left(\Delta x_c\right)}{y\left(\Delta x_f\right)} \right|, \tag{216}$$

where Δx_f and Δx_c are the grid spacings on the fine and coarse grids, y is the simulated solution on those grids, R is the refinement ratio $\Delta x_f / \Delta x_c$, and F_s is a "safety factor". There is some controversy over the appropriate value to use for F_S, but typical values are in the range of 1 to 3. The correct solution is then assumed to lie in the range given by

$$y_{\mathrm{GCI}} \approx y\left(1 \pm \mathrm{GCI}\right). \tag{217}$$

If the exact solution to the problem is known, the order of convergence of the numerical method can be easily calcualted from the following equations

$$\varepsilon\left(\Delta x_c\right) = \left|y_{\text{exact}} - y\left(\Delta x_c\right)\right| = \beta\left(\Delta x_c\right)^p + \text{HOT} \tag{218}$$

$$\varepsilon\left(\Delta x_f\right) = \left|y_{\text{exact}} - y\left(\Delta x_f\right)\right| = \beta\left(\Delta x_f\right)^p + \text{HOT}, \tag{219}$$

where ε is the error between the simulated solution and the exact solution, p is the order of the method, β is an unknown constant, and HOT represents the higher order terms. Equations 218 and 219 represent a system of two equations with two unknowns (p and β) that can be solved to give the order of convergence as

$$p = \frac{1}{\log\left(R\right)} \log\left(\frac{\varepsilon\left(\Delta x_c\right)}{\varepsilon\left(\Delta x_f\right)}\right). \tag{220}$$

If the exact solution is not known, a "converged" solution on a very fine grid can be used in place of the exact solution.

If the solution on three grids can be obtained, it is possible to use Richardson extrapolation to determine the order of convergence as well as an extrapolated estimate of the solution on an infinitely fine grid. This analysis assumes that the solution on the three grids are in the asymptotic regime of convergence, in which all higher order terms are negligble. If this is not the case, the results will not be reliable. It also assumes that convergence is monotonic. For non-monotonic convergence, the analysis fails. The numerical solutions on the three grids (coarse, medium, and fine) are given by

$$y\left(\Delta x_c\right) = y_{\text{RE}} + \beta\left(\Delta x_c\right)^p + \text{HOT} \tag{221}$$

$$y\left(\Delta x_m\right) = y_{\text{RE}} + \beta\left(\Delta x_m\right)^p + \text{HOT} \tag{222}$$

$$y\left(\Delta x_f\right) = y_{\text{RE}} + \beta\left(\Delta x_f\right)^p + \text{HOT}, \tag{223}$$

where y_{RE} is the extrapolated solution. There are now three equations with three unknowns (p, β, and y_{RE}), which can be solved to give

$$p = \frac{1}{\log\left(R\right)} \log\left(\frac{y\left(\Delta x_m\right) - y\left(\Delta x_c\right)}{y\left(\Delta x_f\right) - y\left(\Delta x_m\right)}\right) \tag{224}$$

$$y_{RE} = y\left(\Delta x_f\right) + \frac{y\left(\Delta x_f\right) - y\left(\Delta x_m\right)}{R^p - 1}. \tag{225}$$

The remainder of this section will describe application of these techniques to the CRASH code. The first problem to be discussed will be a radiative shock tube that exercises many of the physics modules relevant to the full CRASH problem. An analysis of a one-dimensional version of the full CRASH problem will then be presented, followed by some results for the two-dimensional CRASH problem.

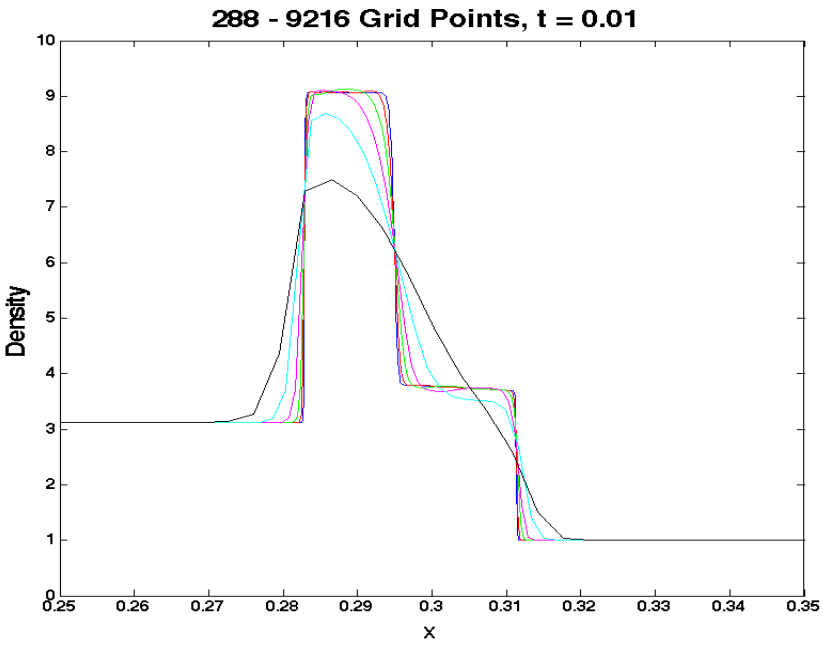

Figure 74: Denisty profiles at t = 0.01 s for the radiative shock tube test problem using six different grid resolutions. The lowest resolution shown is 288 grid points, represented by the black curve. A factor of two was used for each grid refinement, with the highest resolution shown, 9216 points, represented by the blue curve. Two additional simulations were performed using 18,342 and 36,864 points. The results of these were indistinguishable from the 9216 point case on the plot, indicating that grid convergence had been achieved at 9216 points.

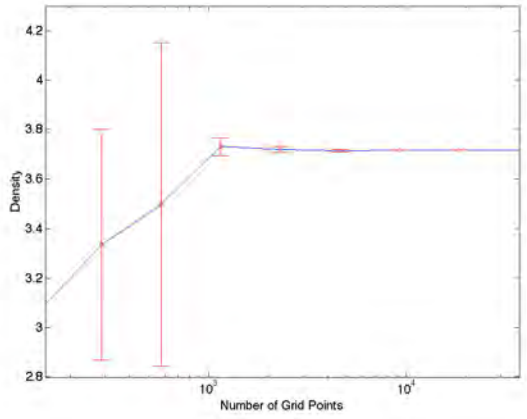

Figure 75: Density at $x = 0.30825$ for the radiative shock tube test problem as a function of grid resolution. The error bars were obtained using the grid convergence index with a safety factor of $F_S = 3$.

4.3.1 Shock Tube with Radiation

A relatively simple problem involving both gray radiation diffusion and hydrodynamics was proposed by Lowrie and Morel. A one-dimensional shock tube is defined on a grid ranging from $x = 0.$ to $x = 1.$ with the discontinuity initially at $x = 0.25$. The initial values of density, fluid velocity, and pressure to the left of the interface are 3.123, 6.798, and 2.874, respectively. To the right of the interface, the initial density, fluid velocity, and pressure are 1.0, 0., and 0.6, respectively. The material is assumed to be a gamma-law gas with $\gamma = 5/3$. The radiation energy and flux are given by

$$E_r = a_r T^4, \quad F_r = 4/3 E_r v, \tag{226}$$

where v is the fluid velocity, $a_r = 44.93$ is the radiation constant, and the speed of light is $c = 100$. A constant opacity is assumed with a value $\sigma_t = 100$.

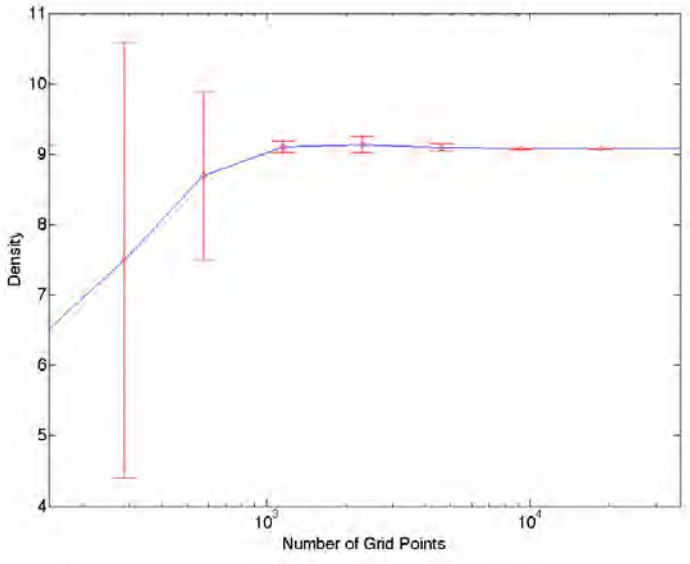

Figure 76: Maximum density in the radiative shock tube test problem as a function of grid resolution. The error bars were obtained using the grid convergence index with a safety factor of $F_S = 3$.

Grid convergence was tested for this problem using nine grid sizes, ranging from 144 points to 36,864 points. A constant refinement ratio of $R = 2$ was used. Fig. 74 shows the density profile at t=0.01 for six of these grids. Grid sizes represented are 288 (black), 576 (cyan), 1152(magenta), 2304(green), 4608 (red), and 9216 (blue). For the two higher-resolution grids, the results were virtually indistinguishable from the 9216-point case, suggesting that a converged solution has been reached. The solution on the 144-point grid was so inaccurate that it was not included on the plot.

The density at a fixed point ($x = 0.30825$) as a function of grid size is plotted in Fig. 75. The error bars were obtained using the GCI with a safety factor $F_S = 3$. It appears that these error bars represent a reasonable estimate of

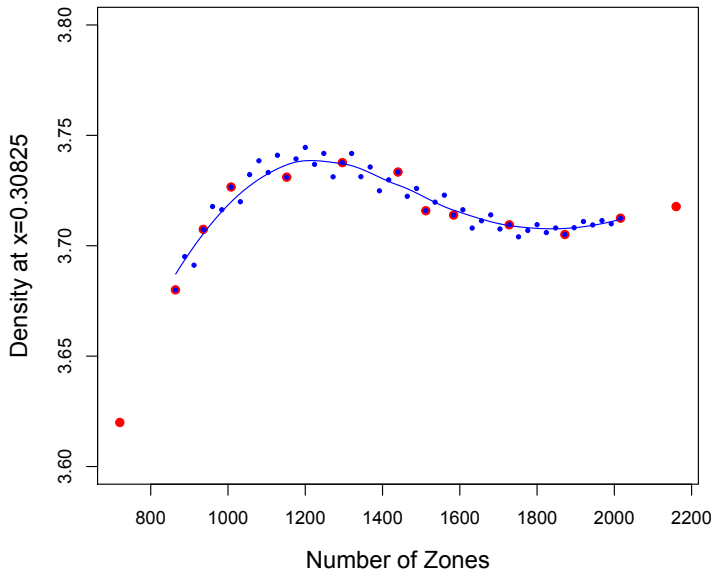

Figure 77: Closeup of the density at $x = 0.30825$ as a function of grid size for the radiative shock tube problem. The convergence is oscillatory, preventing use of Richardson extrapolation. The scatter in the points is due to errors in the feature extraction algorithm. The curve was obtained using a smoothing algorithm.

the error in the solution. Fig. 76 shows a similar plot for maxium density on the grid. Again, the error bars shown appear to provide a reasonable estimate of the numerical errors.

Unfortunately, it is impossible to use Richardson extrapolation for this problem, since the convergence is non-monotonic. This is illustrated in Fig. 77, which shows the density at $x = 0.30825$ for a large number of grids between 576 points and 2100 points. The scatter in the points is caused by errors in the feature extraction algorithm. The red points indicate a first sample of grids. Gaps between the red points were later filled in with the blue points. The curve was obtained using a smoothing algorithm.

Even though Richardson extrapolation cannot be used for these results, it is still possible to estimate the convergence order of the code from Eq. 220. The convergence rates for both the density at a fixed point and the maximum density between each two adjacent grid sizes are listed in Table 12. The wide variation in these numbers results from the non-monotonic grid convergence of the code. For the convergence of the maximum density, some grid sizes actually show a negative order of convergence, indicating that the solution on the finer grid is farther from the reference solution than the solution on the coarser grid.

4.3.2 One-dimensional CRASH Problem with Hyades Initialization

Given the limited success in applying solution verification techiniques to the radiative shock tube problem, one might expect even more difficulties in applying them to the full CRASH problem, and indeed, this turned out to be true.

Table 12: Order of Convergence

Points	ρ(x=0.30825)	Max ρ
144	0.304	0.554
288	0.741	0.726
576	0.807	2.304
1152	3.845	3.778
2304	2.287	-0.992
4608	1.265	1.672
9216	1.882	5.730
18432	0.737	-2.402

Nevertheless, performing the analysis provided some very valuable information about the CRASH code. For this section, a one-dimensional version of the CRASH problem was used. The code was initialized using a solution from the Hyades code at 1.1 ns.

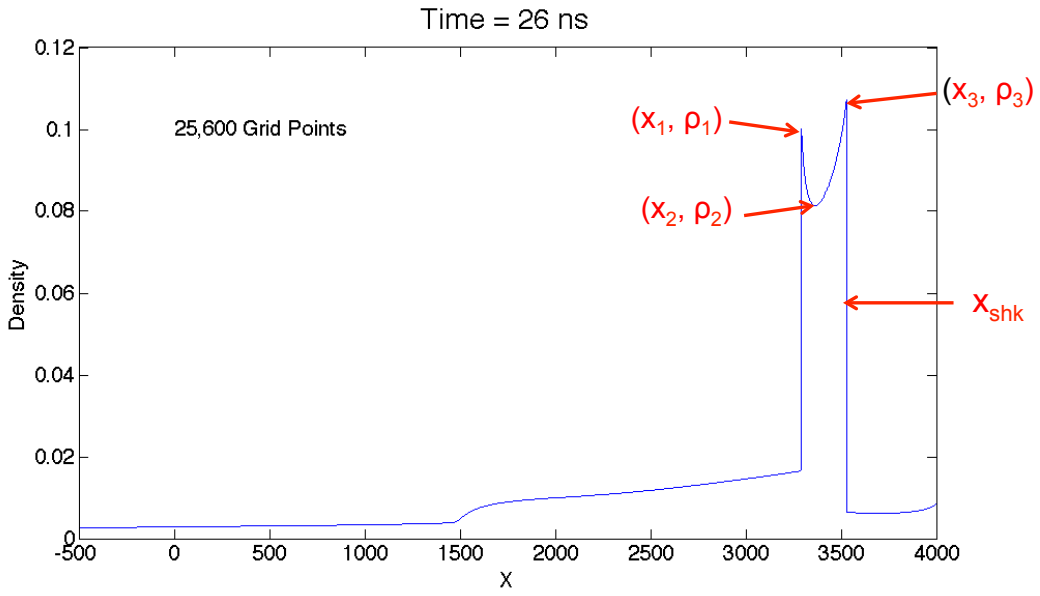

Figure 78: Density profile at 26 ns for a one-dimensional version of the CRASH problem. The simulation was initiated using output from Hyades at 1.1 ns. The positions and densities shown in red indicate the output quantities of interest that were used to test grid convergence.

The density profile at 26 ns on a grid of 25,600 points is shown in Fig. 78. There are two very sharp peaks, which require extremely fine grids to resolve. Grid convergence was tested using the shock location (x_{shk}), the location of the peak on the left (x_1), the location of the density minimum (x_2), and the density at the two peaks and at the minmum

(ρ_1, ρ_2, and ρ_3). These values are shown in red in Fig. 78.

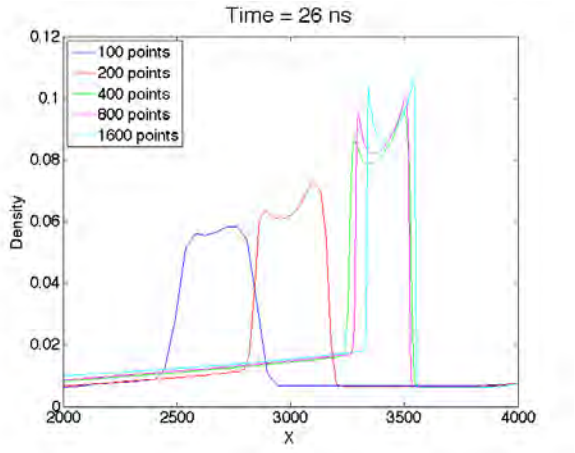

Figure 79: Density profiles at 26 ns for a one-dimensional version of the CRASH problem with Hyades initialization obtained using the five lowest-resolution grids. The density peak moves to the right and increases in height with increasing resolution for these relatively small grids.

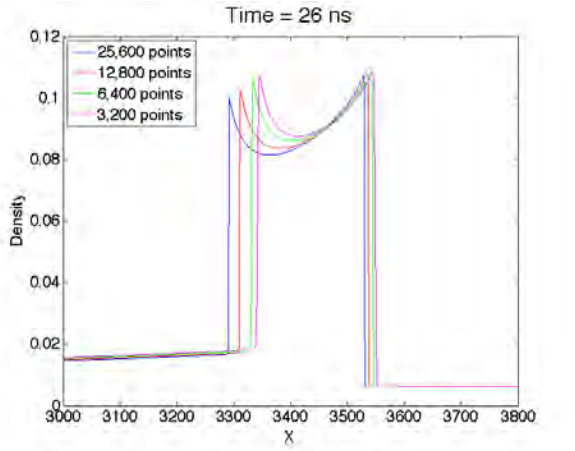

Figure 80: Density profiles at 26 ns for a one-dimensional version of the CRASH problem with Hyades initialization obtained using the four highest-resolution grids. The density peak moves to the left and descreases in height with increasing grid size at these resolutions, which is opposite to the behavior seen for low-resolution grids. The code is showing no sign of convergence at these resolutions, since the differences between the solutions do not seem to be decreasing.

Nine different grids were used to test convergence, ranging from 100 points to 25,600 points, using a constant refinement ratio of $R = 2$. The density profile at 26 ns on the five coarsest grids is plotted in Fig. 79. The locations of the two density peaks and the density minimum move to the right as the resolution is increased. The value of the

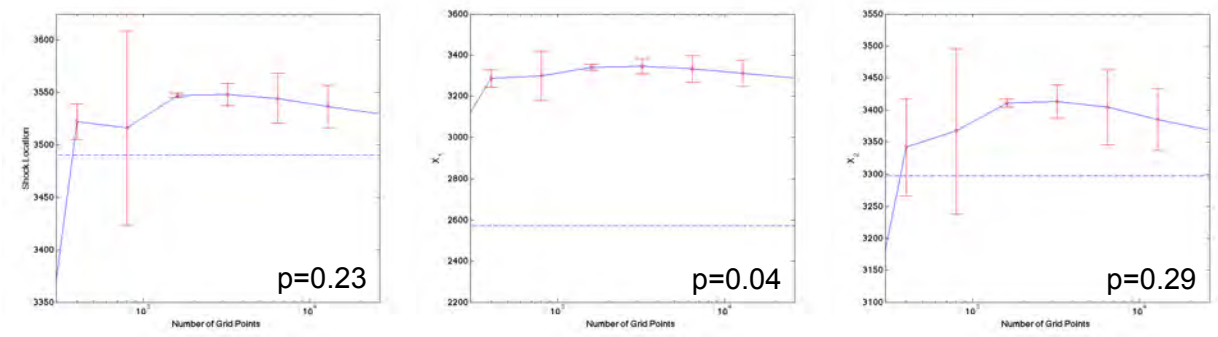

Figure 81: Location of three features in a one-dimensional version of the CRASH problem with Hyades initialization as a function of grid resolution. The plots represent, from left to right, the shock location x_{shk}, x_1, and x_2 as shown on Fig. 78. The error bars were obtained using the grid convergence index with a safety factor of $F_S = 3$. Due to the non-convergence of the code for this problem, these error bars do not give a very meaningful estimate of the convergence error. The horizontal line represents the value obtained using Richardson extrapolation for the three highest-resolution values. This value is also very unreliable due to the convergence problems. The value of p shown is the order of convergence obtained using the three highest-resolution simulations.

density at all three locations increases with increasing resolution. Judging from these five grids, it appears that the code is converging properly. However, the situation changes for higher-resolution grids, as shown in Fig. 80. For these grids, the location of all three features moves to the left with increasing resolution and the density values decrease.

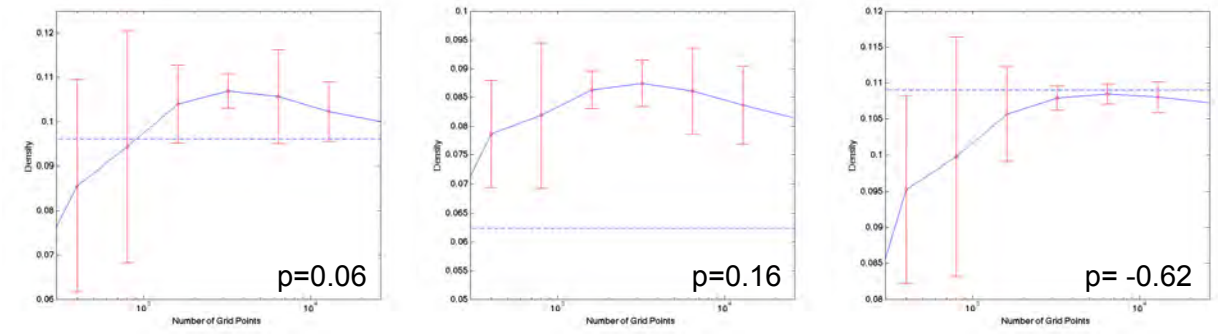

Figure 82: From left to right, values of ρ_1, ρ_2, and ρ_3 (see Fig. 78) as a function of grid resolution for a one-dimensional version of the CRASH problem with Hyades initialization. The error bars were obtained using the grid convergence index with a safety factor of $F_S = 3$. Due to the non-convergence of the code for this problem, these error bars do not give a very meaningful estimate of the convergence error. The horizontal line represents the value obtained using Richardson extrapolation for the three highest-resolution values. This value is also very unreliable due to convergence problems. The value of p shown is the order of convergence obtained using the three highest-resolution simulations. For ρ_3, the code is actually diverging with a negative order of convergence.

Moreover, the difference between adjacent curves does not seem to be decreasing, and there is no sign that the locations of the three features or the densities at those points are converging toward any particular values.

Fig. 81 shows (from left to right) the locations x_{shk}, x_1, and x_2 as a function of grid resolution. The error bars were obtained from the GCI with a safety factor $F_S = 3$. The three highest-resolution points for each position were monotonic, so it is possible to use Richardson extrapolation. The extrapolated solutions are given by the horizontal dashed lines. In most cases, the extrapolated solution lies outside of the error bars. The order of convergence p obtained for each case using the three highest-resolution points is shown in the lower right hand corner of each plot. Fig. 82 shows a similar analysis for ρ_1, ρ_2, and ρ_3. For ρ_3, the code is diverging, with a negative order of convergence. This type of behavior is an example of what can go wrong when solution verification techniques are applied to a code that is not converging properly.

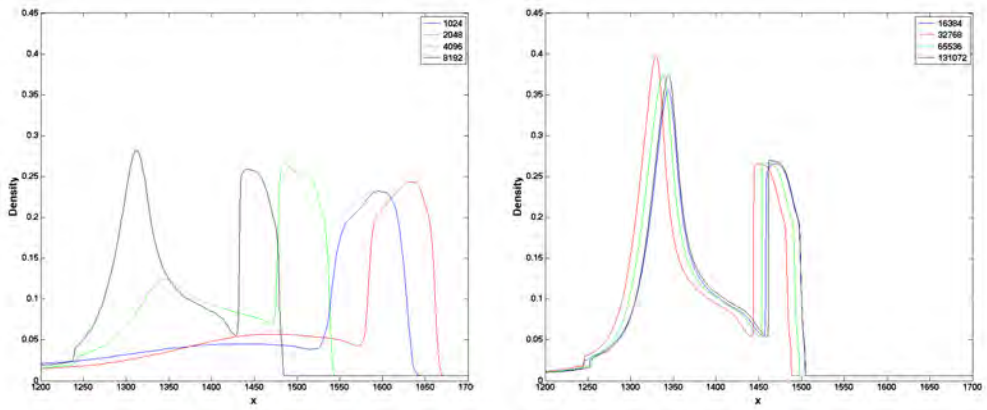

Figure 83: Density profiles obtained at 13 ns for a one-dimensional version of the CRASH problem with laser package initialization at eight different grid resolutions. At the four lowest resolutions, the shock position varies dramatically, increasing as the grid resolution is changed from 1024 to 2048 points and then decreasing as the resolution is increased further to 8192. At higher resolution, the variation is much smaller, but the location of the shock remained erratic, showing no sign of convergence. The simulations at 16,384 and 131,072 grid points agree quite well, but this is just coincidence, since the intermediate resolutions are rather different.

4.3.3 One-dimensional CRASH Problem with Laser Package Initialization

A number of simulations were attempted in an effort to understand why the CRASH code was not converging for the one-dimensional CRASH problem. The first was to use the CRASH laser package instead of the Hyades initialization. Fig. 83 shows density profiles at 13 ns obtained using eight grid sizes ranging from 1024 points to 131,072 points. A constant refinement ration of $R = 2$ was used. There is a large variation in location of the primary shock for the four lowest resolutions. At higher resolution, the variation is much smaller, but the behavior is erratic, showing no clear trend. The results for 16,384 points and 131,072 points are nearly identical, but this is just a coincidence and does not indicate convergence, since the intermediate results are significantly different. Similar behavior is seen at 26 ns (Fig. 84).

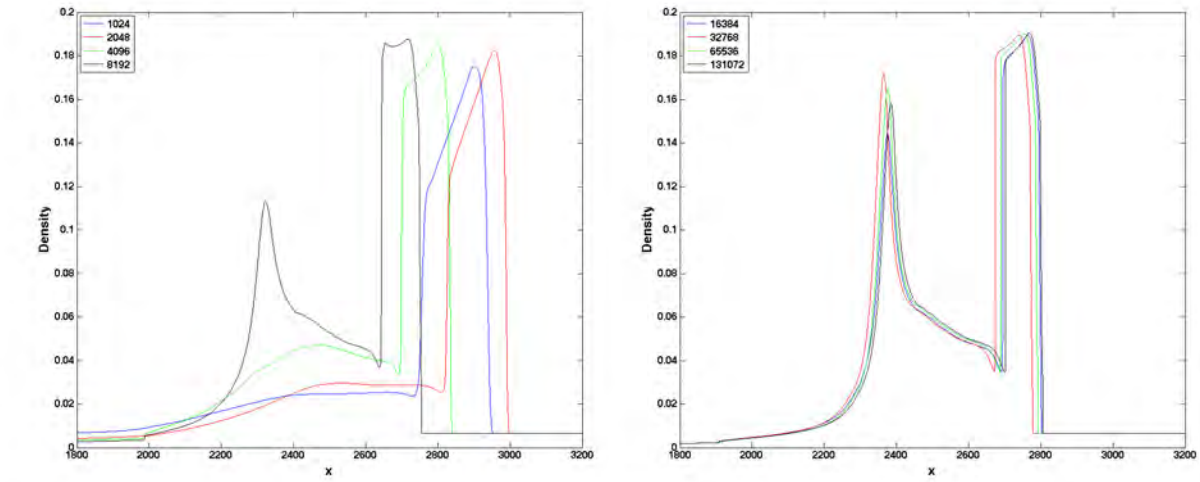

Figure 84: Density profiles obtained at 26 ns for a one-dimensional version of the CRASH problem with laser package initialization at eight different grid resolutions. The behavior is similar to that at 13 ns (Fig. 83). Although the location of the primary shock does not seem to be converging, the range of variation is quite small compared to experimental uncertainty at high resolution.

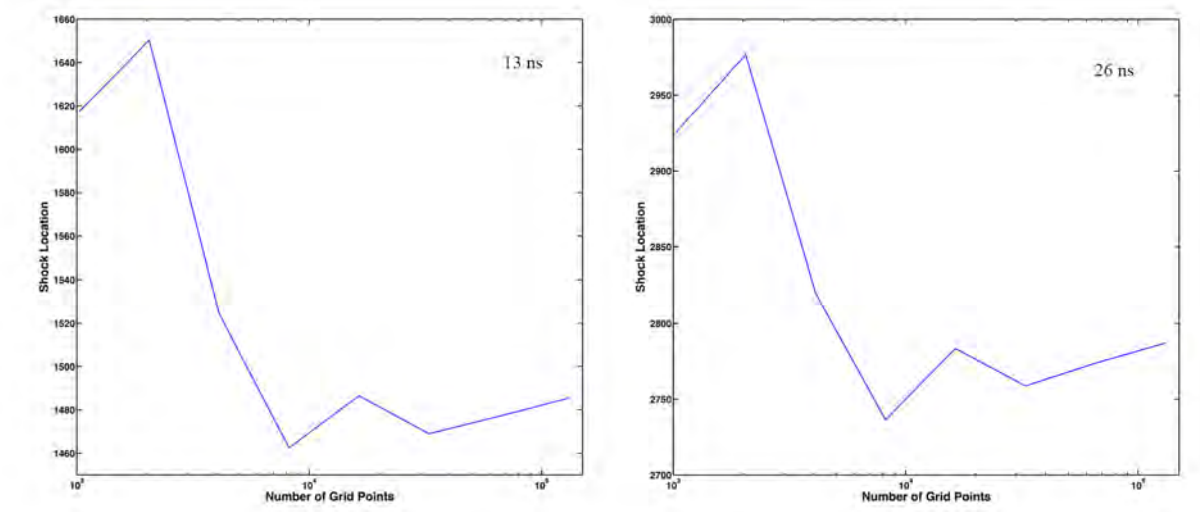

Figure 85: Location of primary shock as a function of grid size for a one-dimensional version of the CRASH problem with laser package initialization. Results are shown at both 13 ns and 26 ns. There is no indication that the location is converging toward a particular value.

The location of the primary shock as a function of grid resolution is plotted in Fig. 85 at both 13 ns (left) and 26 ns (right). No clear convergent behavior is detectable. At the highest resolutions, the variation in shock location is small compared to experimental uncertainties. However, there is no way to determine how much the shock location will change at even higher resolutions. It is also interesting to observe that the resolution required to get close to a converged solution is considerably higher when the laser package is used for initialization than when the result from Hyades is used.

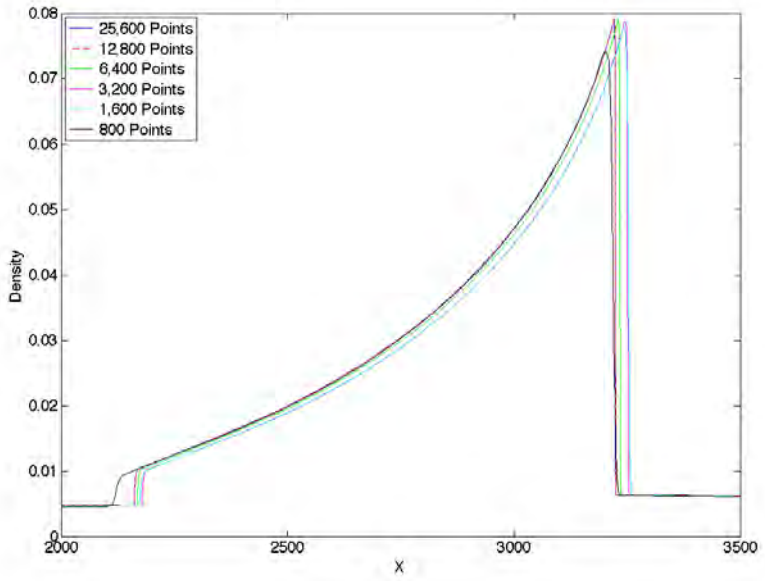

Figure 86: Density profiles obtained for a single-fluid version of the one-dimensional CRASH problem in which all of the material is assumed to be pure xenon. Results are shown for six different grid resolutions. With a single material, there is only one peak in density.

4.3.4 Single Material Test

The next set of tests performed investigated the effects on the convergence properties of using level sets to track material interfaces in the CRASH code. The level set algorithm used in CRASH does not permit grid cells with a mixture of materials. Each cell contains only one pure material. The advantage of this approach is that it avoids errors having to do with computing equations of state and opacities for a mixture of materials. The disadvantage is that it does not conserve the mass of each material, although the total mass within the grid is conserved. Since a non-conservative method is not guaranteed to converge on the grid, the use of level sets seems a likely source of the odd behaviors described in the previous sections.

For the test described here, all the material in the grid was assumed to be pure xenon. Since there are no longer any material interfaces in the problem, this disables any effect of the level sets on the solution. The density profiles obtained on six grids ranging from 800 points to 25,600 points are plotted in Fig. 86. There is now only a single density

peak in the solution. Fig. 87 shows a closeup of the region in the vicinity of the shock. Although the location of the shock does not vary monotonically, at the highest resolutions, the solution appears to be in the asymptotic regime of convergence. The two highest-resolution results are nearly identical, suggesting that a converged solution has been achieved. This result tends to support the idea that the non-convergence in the code is caused by the use of level sets to track material interfaces.

4.3.5 Mixed Cell Test

An additional test of the level set algorithm involved a multi-material test with no interface tracking. The various materials were allowed to mix with each other. This created errors in the equation of state and opacity calculations, but the amount of each material in the grid was conserved to roundoff error. Additional error was created due to increased diffusion across the material interface.

Fig. 88 shows the density profiles obtained on the same grids as for the single-material case. Unlike the single-material case, the density again shows two seaparate peaks between the shock and the material interface. The increased diffusion error can be seen by the spreading of the material interface, which for low resolutions is extreme. There is no sign of non-convergent behavior, and the highest-resolution solutions are nearly identical. Given this result and the result from the previous section, it appears nearly certain that the non-convergent behavior described above was due to the use of level sets for tracking material interfaces.

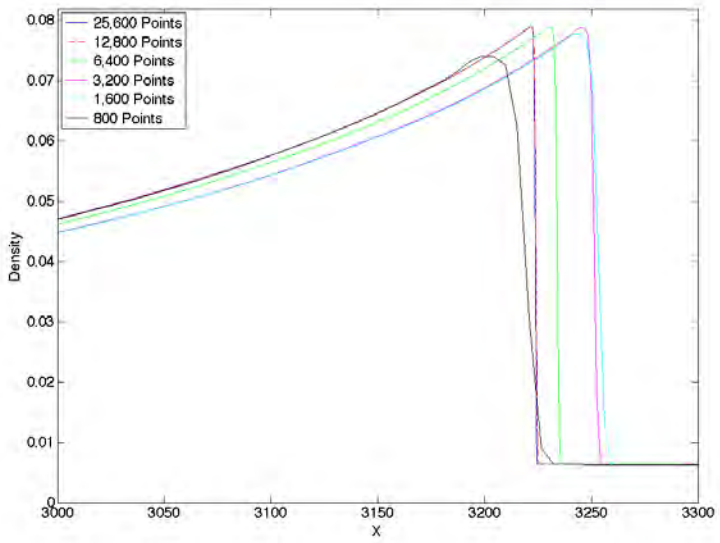

Figure 87: Closeup of the region near the primary shock for the simulations shown in Fig. 86. The two highest-resolution results are nearly identical, suggesting that a converged solution has been obtained.

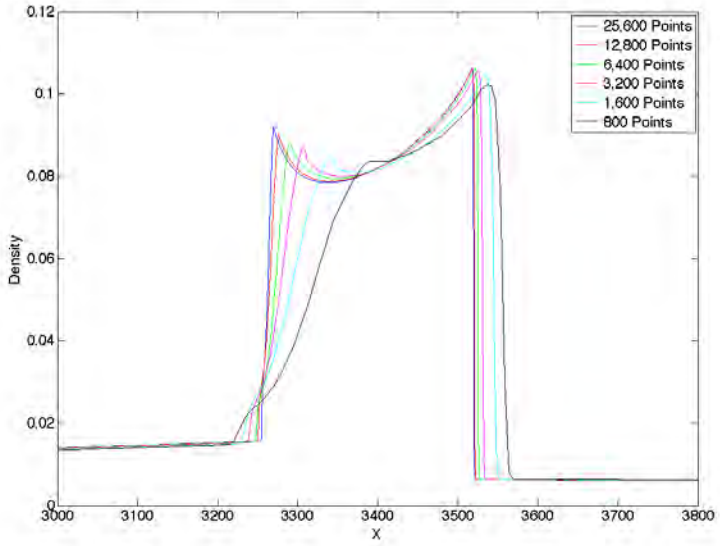

Figure 88: Density profiles obtained on six different grids for a one-dimensional version of the CRASH problem. For these simulations, the level sets used to track material interfaces were turned off, and the fluids were allowed to mix. Although there is considerably more diffusion at the material interface between the xenon and beryllium, the solution appears to be nearly converged at the highest resolution.

4.3.6 Two-dimensional CRASH Problem with Laser Package Initialization

Uncertainties in the location of the primary shock are more complicated in multidimensional simulations. The errors due to non-convergence described above are still present. To make matters worse, the huge grids used for some of the one-dimensional tests are not feasible in two dimensions. The highest-resolution two-dimensional simulation described in this section required months to complete. In addition to this error, the two-dimensional simulations are subject to fluid instabilities. Due to the lack of a real physical viscosity in CRASH, short-wavelenth instabilities have the fastest growth rates. As the grid is refined, the effective Reynolds number of the simulation increases, and even shorter wavelengths grow. Without a dissipation mechansim to suppress these short-wavelength instabilities, there is no way to obtain a converged solution. If these instabilites have an effect on the location of the primary shock, this additional uncertainty should be included in the CRASH uncertainty quantification analysis. Finally, for reasons that are not completely clear, CRASH produces an anomolous protruding feature on the symmetry axis that is not seen in the experimental results. This leads to considerable uncertainty concerning the actualy location of the primary shock, and different feature detection algorithms can produce quite different results.

The density distributions on five different grids obtained for the two-dimensional CRASH problem with laser package initialization are shown in Fig. 89. Results are shown at both 13 ns (top) and 26 ns (bottom). It was not feasible to use a fixed grid due to the enormous computational resources that would have been required. Instead, higher resolution was obtained by increasing the number of levels of Adaptive Mesh Refinement (AMR). The effective resolution for each refinement increased by a factor of two, ranging from 1320×120 to $21,120 \times 1920$. The resolution

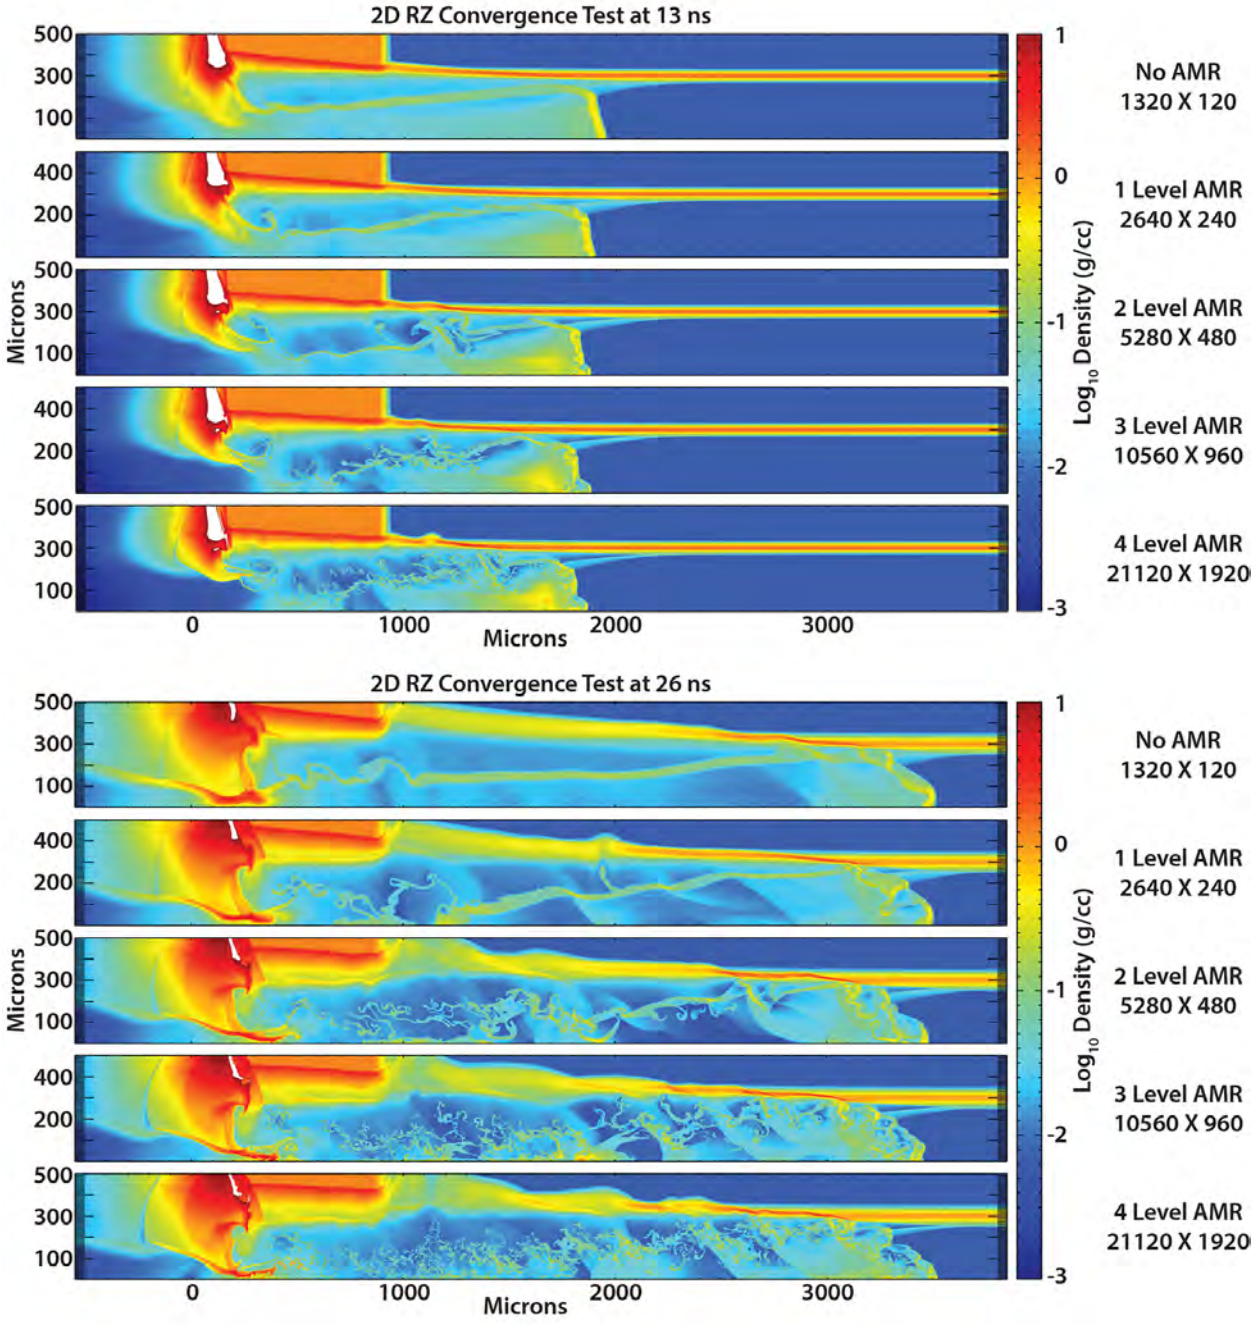

Figure 89: Density distribution at 13 ns (top) and 26 ns (bottom) for the two-dimensional CRASH problem with laser package initialization. Results are shown at five different grid resolutions. Instead of using a fixed grid, the resolution was changed by increasing the number of levels of Adaptive Mesh Refinement (AMR). This was necessary to achieve high resolution without using an unreasonable amount of computer resources.

typically used for CRASH uncertainty quantifaction studies typically corresponding to somewhat less that the two-level results in two dimensions and somewhat less than the one-level results in three dimensions.

The location of the primary shock at both 13 ns (left) and 26 ns (right) are plotted in Fig. 90 as a function of the number of AMR levels (grid resolution). At 13 ns, it appears that the shock position has converged, but given the results described above, it seems likely that additional variation in shock location might occur at higher resolutions. At 26 ns, there is no sign of a converged value for the shock location.

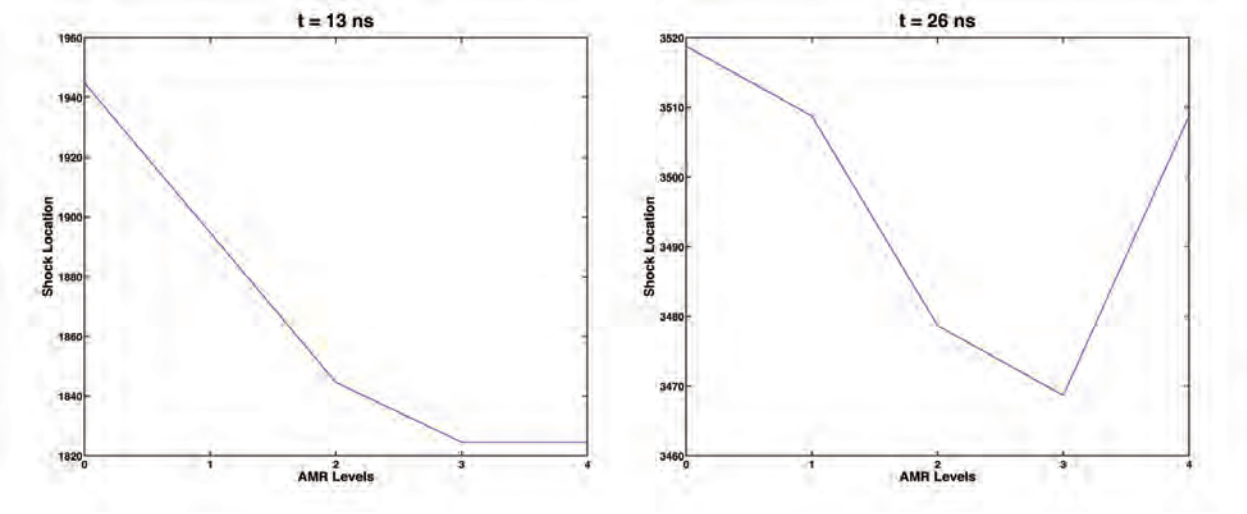

Figure 90: Shock location for the two-dimensional CRASH problem with laser package initialization as a function of the number of AMR levels (grid resolution). Results are shown at times of 13 ns and 26 ns. At $t = 13$ ns, it appears that the solution is converging. However, given the one-dimensional results described above, it is quite likely that the shock location will show additional variation at higher resolutions.

4.3.7 Summary of Solution Verification Studies

Solution verification techniques have been applied to the CRASH code with only limited success. For a relatively simple radiative shock tube problem, good grid convergence was achieved. However, the oscillatory nature of the convergence prevented use of Richardson extrapolation and made it impossible to determine the order of convergence of the code.

For a one-dimensional version of the full CRASH problem with initialization from a Hyades data file at 1.1 ns, the situation was much worse. The code appeared to be converging properly for low-resolution grids. However, for grids larger than about 3200 points, the direction of variation of the quantities of interest with increasing resolution reversed, and it appeared unlikely that a converged solution could be found at any resolution. When the problem was initialized using the CRASH laser package, convergence was even more erratic, and the number of grid points required to get a "nearly converged solution" was even larger.

Two additional tests were performed to determine if the use of level sets for tracking material interfaces was the cause of the poor convergence properties of the code. Since the level set algorithm used in CRASH does not conserve

the mass of individual materials, convergence of the code is not guaranteed. For the first test, all of the material in the grid was assumed to be pure xenon. Since there are no longer any material interfaces to track, the level set alogrithm has no effect on the results. In this case, good grid convergence was achieved, although the convergence was still not monotonic, making it impossible to use some solution verification techniques. The second test involved multi-material simulations in which the level set tracking was turned off and the materials were allowed to mix. Here also, good grid convergence was observed. It appears fairly certain that the use of level sets in CRASH to track material interfaces prevents the code from converging on the grid.

The multidimensional CRASH problem suffers from the same convergence issues as the one-dimensional simplification. In addition, the presence of fluid instabilities further complicates the situation. Since there are no physical dissipation terms in the CRASH code to limit the growth of short-wavelength instabilities, it is impossible to obtain a converged solution. As the grid size is increased and the effective Reynolds number of the simulation increases, smaller-wavelength instabilities will continually appear. It appears that these instabilities can affect the morphology and location of the shocks in the problem, introducing additional uncertainties into the simulation.

In spite of the limited success of solution verification techniques, it is still possible to obtain a rough (although not mathematically rigorous) estimate of the error in the shock location due to the use of non-converged simulations. Even though the shock position is not converged even at high resolution, the range of variation is relatively small. The difference between the shock location at the grid resolution typically used and the location at the highest resolutions is approximately 100 μ for two-dimensional simulations and 150 μ for three-dimensional simulations. The variation in shock location with resolution for the two-dimensional simulations is only about 50 μ, but this number would likely increase if higher-resolution simulatons could be performed. These errors due to lack of grid convergence were not yet included in the overall CRASH uncertainty quantification analysis at project end.

4.4 Code-to-code Comparisons

Contrary to the view of some who write about uncertainty quantification, we believe that code-to-code comparison can be a valuable tool for checking the accuracy and reliability of numerical simulation codes. Agreement between the codes does not guarantee that they are producing the correct answer. They could all be getting the same wrong result. However, if there is disagreement, it is clear that at least one of the codes is wrong. Understanding the cause of these discrepancies can point to programming errors or differences in physics modules that deserve further investigation. We undertook a collaboration to compare several radiation-hydrodynamics codes. The codes currently in the test suite are CRASH (University of Michigan), FLASH (University of Chicago), and xRAGE (LANL).

4.4.1 Description of Simulation Codes

All three codes contain sufficient physics for simulating a wide range of High-Energy-Density Physics (HEDP) phenomena, including radiative shocks. They solve equations for plasmas containing electron, ions, and thermal radiation. Each of these components is allowed to have a different temperature and thus the plasma is described by the "3T" or "three-temperaturee" approximation.

The gas dynamics modules for compressible gas dynamics in each code solve Euler's equations in conservative form using finite-volume methods on a structured grid with adaptive mesh refinement capability. CRASH and FLASH use a block-structured adaptive mesh, while xRAGE uses cell-by-cell refinement. The effects of viscosity in the prob-

lems of interest here are typically much less than numerical diffusion and are, therefore, omitted from the equations. Each code uses a different variation on the high-order Godunov formalism, and are accurate to second order in both space and time. Radiative transfer is calculated using a flux-limited diffusion approximation with both single-group and multigoup capability, although the details of the radiation solvers differ from code to code. Single-group calculations with FLASH are run in multigroup mode using a single energy group. The codes also compute energy diffusion by flux-limited electron heat conduction and include terms to account for the transfer of energy between electrons and ions through collisions, as well as the absorption and emission of radiation by electrons. The effects of scattering, magnetic fields, and relativity are not included.

A major difference between the codes is in the way they treat material interfaces. CRASH determines interface locations using a level set method. An initial level set function is defined for each material. The function for a given material is positive in each zone which contains that material and negative in all other zones. The initial magnitude of the function is equal to the distance between the interface and that zone. This function is advected along with the fluid at each time step. At the end of each time step, the level set function with the largest value in each zone determines the identity of the material in that zone. In this approach there are no mixed zones − each zone contains only a single material. The advantage of this approach is that it eliminates the need to determine equations of state and opacities for a fluid containing a mixture of materials. The disadvantage is that it does not conserve the mass of each material, although the total mass within the grid remains fixed, except for material flowing through the boundaries. FLASH and xRAGE solve a separate advection equation for each material and allow the materials to mix within each zone. FLASH uses a steepener that prevents diffusive spreading of the interface beyond two zones. The opacities in mixed cells are then weighted by the number density of each material. For the equation of state, FLASH assumes common electron and ion temperatures. xRage produces sharp interfaces using either an interface preserver or a volume of fluid method. It also weights the opacities by number density. The equation of state in mixed cells assumes temperature and pressure equilibrium.

There are also differences in the 3T approximations used by the three codes. FLASH and xRAGE solve separate equations for the total energy, electron energy, and radiation energy. CRASH uses an electron pressure equation instead of the equation for the electron energy. All three codes divide shock heating among the electrons, ions, and radiation in proportion to their pressure ratios. FLASH has an additional option to solve a separate electron entropy equation to assign shock heating to only the ions.

4.4.2 Results

4.4.2.1 Reverse Radiative Shocks

The first comparison attempted was a one-dimensional shock tube that produces a reverse radiative shock. The plasma is fully ionized helium at an initial density of 0.0065 $\mathrm{gm/cm^3}$. The boundary condition at the left of the grid was set to be a reflecting wall, and material was allowed to flow into the grid through the right boundary at a velocity of 100 km/s. A radiative shock developed and propagated through the grid from left to right. A schematic of the initial setup is shown in Figure 91.

Figure 92 shows the results obtained for the density and for the electron and ion temperatures using both FLASH and xRAGE. This problem was not attempted with CRASH. The plot shows a close up of just the region in the vicinity of the shock. The shock structure differs substantially between the two codes. The shock location is different, as

are the peak ion temperature and the extent of the pre-heated region to the right of the shock. The two codes show good agreement in the rest of the grid. At this point, it was decided that there were too many possible sources for this discrepancy to easily diagnose. For this reason, we designed a set of simpler test problems that would test each term in the equations individually. A subset of these tests will be discussed in the following sections.

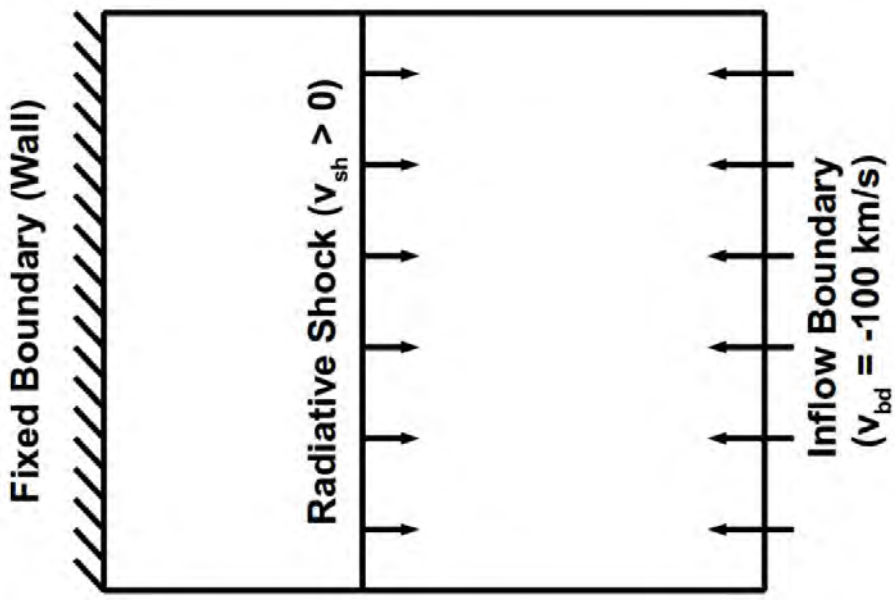

Figure 91: Schematic of the reverse radiative shock test problem, in which fully ionized helium flows into the grid from the right at a velocity of 100 km/s with a solid wall boundary at the left. The result is a radiative shock propagating to the right.

4.4.2.2 Temperature Relaxation Tests

The simplest set of problems in our comparison suite tested the temperature relaxation terms in the equations. All other terms were turned off. The grid was filled with an infinite medium of fully ionized helium at a density of $0.0065 \mathrm{gm/cm}^3$ with no spatial gradients. The helium was assumed to be an ideal gas with adiabatic index of 5/3. The electron, ion, and radiation temperatures were set to different initial values and allowed to evolve into equilibrium.

The first problem in this set tested just the ion-electron equilibration. The initial ion temperature was 500 eV, while the electrons were cooler at 250 eV. The time evolution of the two temperatures is shown in Figure 93. The electron temperature is plotted in blue, with the ion temperature shown in green. The CRASH and xRAGE results are indicated by the solid lines. The two results agree so closely that only a single line is visible. The FLASH results, shown as individual points, also agree with the other two codes to several significant figures. An equilibrium temperature of approximately 420 eV is achieved in an evolution time of less than 1 ns.

Temperature equilibration of the matter with the radiation field was also tested using mediums with constant opacity, an opacity that depended on the electron temperature, and multigroup calculations, with different (but constant) opacity in each group. The multigroup simulations used either four or eight energy groups. The initial radiation tem-

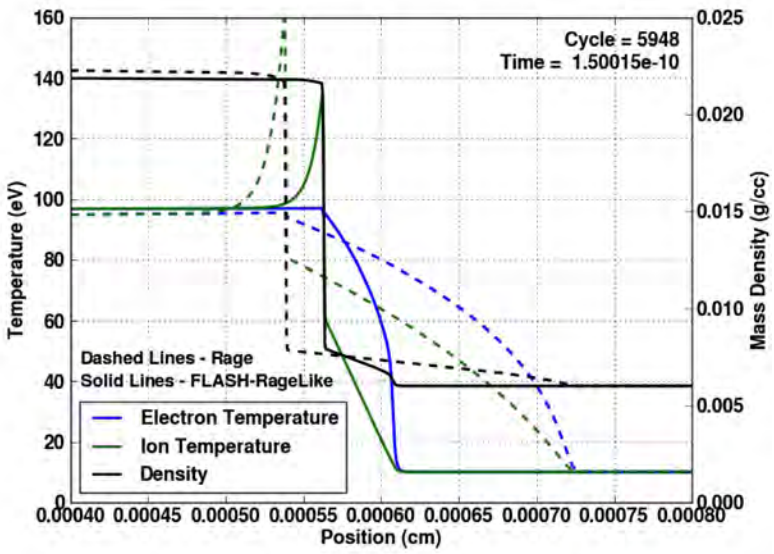

Figure 92: Results obtained from FLASH and xRAGE for the reverse radiative shock test problem. The plot shows only a small portion of the grid in the vicinity of the shock. The two codes produce different shock locations and structures.

perature was set to 100 eV. Only two of these tests are described here. The time evolution of the three temperatures for the constant opacity case is plotted in Figure 94. The opacity of the plasma was assumed to be 0.03 cm^{-1}. The initial temperature of the radiation was set to 100 eV. As in Figure 93, the electrons are plotted in blue and the ions in green. The radiation temperature is shown in red. The results from CRASH and xRAGE are represented by the solid lines, with the FLASH results indicated by the points. In this case, equilibrium had not quite been achieved by 1 ns. However, it appears that the final equilibrium temperature will be approximately 280 eV. As with the previous tests, all three codes gave nearly identical results.

The final test from this set to be described is the multigroup case with four energy groups. The boundaries of the energy groups were located at 1 eV, 400 eV, 800 eV, 1.5 keV, and 100 keV. The opacity in each group was given by

$$\sigma_i = 1.0 \times 10^8 \left(\sqrt{E_{i-1/2} E_{i+1/2}} \right)^{-3} \text{cm}^{-1},$$

with i ranging from 1 to 4. The time evolution of the energy density in each group is shown in Figure 95 for each of the three codes. The CRASH results are represented by the red line, the xRAGE results by the green line, and the FLASH results by the individual points. In this case, CRASH and FLASH produce nearly identical results. However, xRAGE uses different algorithms and disagrees slightly from the other two codes. This difference is unlikely to have any real consequence for practical problems. Initially, significant differences existed between the results from the three codes for the multigroup cases. However, after some programming errors were fixed, the codes came into acceptable agreement for all the tests in this set.

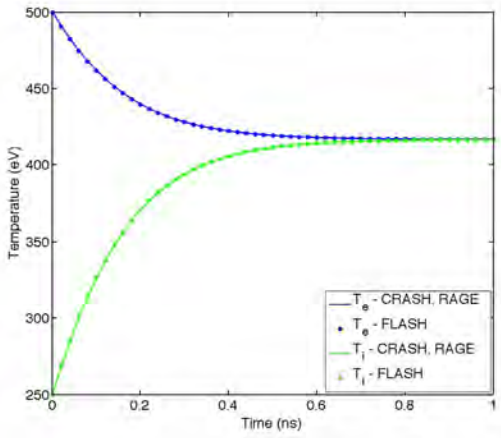

Figure 93: Results of the ion-electron relaxation test for all three codes. The initial condition is a uniform medium of ionized helium with different ion and electron temperatures. The plot shows the time evolution of the ion (green) and electron (blue) temperatures. An equilibrium temperature is reached after about 1 ns. The results from CRASH and xRAGE are represented by the solid lines, while the FLASH results are plotted as individual points. All three results are indistinguishable on this scale and agree to several significant figures.

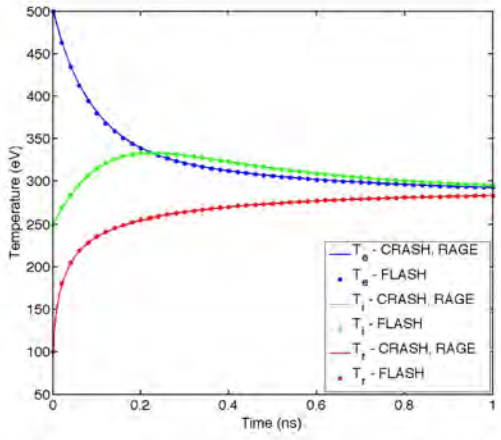

Figure 94: . Results of the ion-electron-radiation equilibration test using a single radiation energy group with constant opacity. The initial condition is a uniform medium of ionized helium with different ion, electron, and radiation temperatures. The plot shows the time evolution of the ion (green), electron (blue), and radiatio (red) temperatures. An equilibrium temperature is not quite achieved by 1 ns. The results from CRASH and xRAGE are represented by the solid lines, while the FLASH results are plotted as individual points. All three results are indistinguishable on this scale and agree to several significant figures.

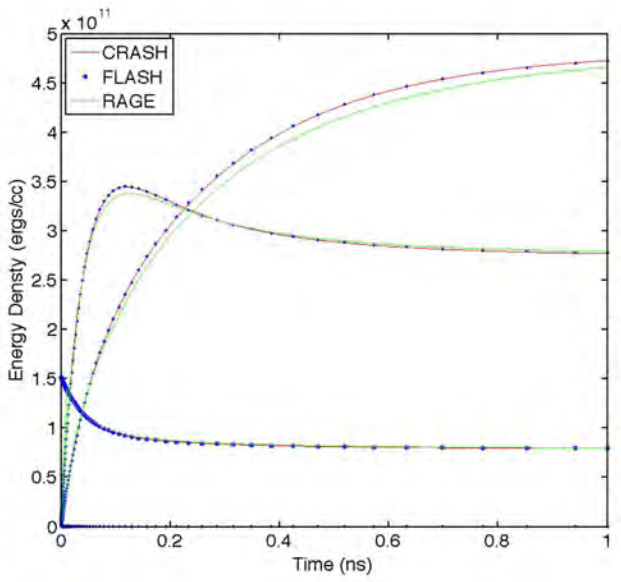

Figure 95: Results from the ion-electron-radiation equilibration test using four radiation energy groups with a constant, but different opacity in each group. The plot shows the time evolution of the energy density in each energy group. Only three of the energy groups contain a significant energy density. The fourth group is represented by a flat line on the horizontal axis. The results from CRASH are shown in red, those from xRage in green, and those from FLASH as individual points. CRASH and FLASH show good agreement, but the xRage results differ slightly, due to the use of different algorithms.

4.4.2.3 Diffusion Tests

The next set of tests to be discussed investigates the accuracy of the diffusion modules in the codes, including both the electron conduction and radiation diffusion modules. Tests were performed using electron conduction alone, electron conduction with ion-electron equilibration, gray radiation alone, and a full-physics test including both electron conduction and single-group radiation diffusion along with ion-electron equilibration, emission, and absorption. For these tests, the plasma is fully ionized helium at a density of 0.0018 $\mathrm{gm/cm^3}$. Again, the helium is assumed to be an ideal gas. The domain extends from 0.0 cm to 0.1 cm with zero-flux boundary conditions at both edges of the grid. The electron thermal conductivity is taken from Atzeni and Meyer-ter-Vehn. Initial temperature profiles for both the ions and electrons are given by

$$T = 450 \tanh\left[-1000\left(\mathrm{x} - 0.2\right)\right] + 550 \text{ eV},$$

where x is the spatial coordinated in cm. This profile is shown in Figure 96. For tests involving radiation, the initial radiation temperature profile is the same as for the ions and electrons but at one tenth the value. The opacities for emission, absorption, and diffusion were all set to the same value, given by

$$\sigma = 1000\left(0.43\frac{z^3}{A^2}\right)\rho^2\left(\frac{T_e}{1000}\right)^{-7/2}\ \mathrm{cm}^{-1},$$

where z is the ionization level, A is the atomic mass, ρ is the mass density, and T_e is the electron temperature.

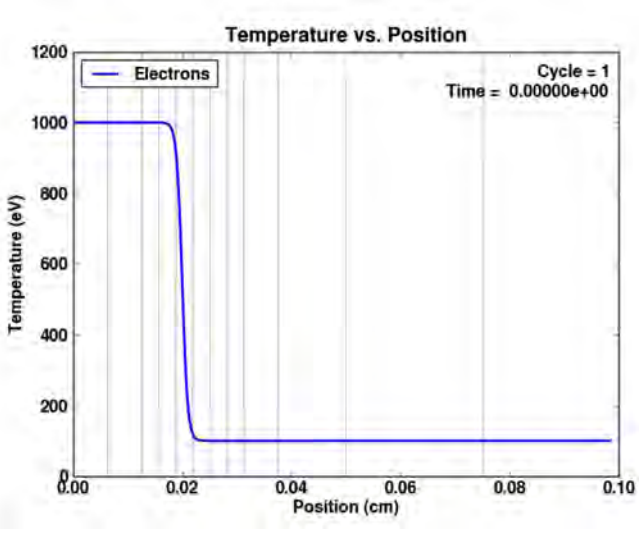

Figure 96: Initilal electron temperature profile for the diffusion tests. The vertical lines indicate the non-unifrom adaptive grid that was used by FLASH. The initial ion temperature profile was identical to that of the electrons. The initial radiation profile was also the same, but at one tenth the value.

The first test involves only electron heat conduction. The temperature profile obtained by the three codes at time 1.5 ns is plotted in Figure 97. The initial results were very different from each other due to an error in calculating specific heats in one of the codes. This error was quickly found and corrected, and the results now show excellent agreement. The results from CRASH and xRAGE are represented by the solid lines in the graph, while the results from FLASH are indicated by the individual points.

This same test was then repeated with ion-electron equilibration turned on. The initial temperature of the electron and ion temperature assumed the same tanh profile described above. Of course, the diffusion operated on only the electrons. The results at 1.5 ns for FLASH and xRAGE are plotted in Figure 98. The xRAGE results are shown as solid lines, and the FLASH results are represented by dashed lines. Again, the results from these two codes are virtually identical. Although the CRASH results are not shown on the plot, they are also indistinguishable from the results from the other two codes.

The next test is for single-group radiation diffusion, without emission, absorption, electron heat conduction, or ion-electron equilibration. FLASH does not have a gray radiation solver, so for this problem, it was run in multigroup mode using a single energy group with T_e ranging from 0 to 1 MeV. The three temperature profiles from FLASH and xRAGE at 2×10^{-14} seconds are shown in Figure 99. This problem was not attempted with CRASH. The electron and ion temperature profiles are shown in blue and green (only the green shows since the values are exactly equal and the green curve overlies the blue curve). The radiation temperature is plotted in red. Since only the radiation diffusion

terms are turned on, the electrons and ions maintain their original profiles. Only the radiation temperature gradient spreads due to diffusion. The two codes are in excellent agreement.

The final diffusion problem tests all the terms in the equations, except for hydrodynamics. It includes electron heat conduction, single-group radiation diffusion, ion-electron equilibration, emission, and absorption. Initial conditions were set as described above. The temperature profiles obtained by the three codes at 0.2 ns are plotted in Fig. 10. The electron temperature profile is plotted in blue, the ion temperature profile in green, and the radiation temperature profile in red. The results from CRASH and xRAGE are represented as solid lines. The CRASH results overly the xRAGE results, so only one line appears on the plot. The FLASH results are represented by individual points, which lie exactly on top of the lines from the other two codes. The radiation evolves on a much shorter time scale than the electrons and ions in this problem, so by the time shown in Fig. 10, the radiation temperature gradient has almost completely diffused away.

4.4.2.4 Hydrodynamics Test

The final comparison described in this paper adds hydrodynamic effects and consists of a steady shock in a two-temperature plasma, with $T_e \neq T_i$. An important feature of this test problem is the existence of an analytic solution, which was derived by Shafranov. This provides an additional sanity check on the reliability of the codes, since the correct solution is known. Radiation does not play a role in this test, but electron heat conduction is included. The conductivity varies as $T_e^{2.5}$ and the ion-electron equilibration varies as $T_e^{1.5}$. Two constant states are initially separated by a simple discontinuity and allowed to evolve until a steady state shock profile is achieved.

The initial upstream conditions are given by a fluid density of 0.0018 gm/cm^3 and electron and ion temperatures

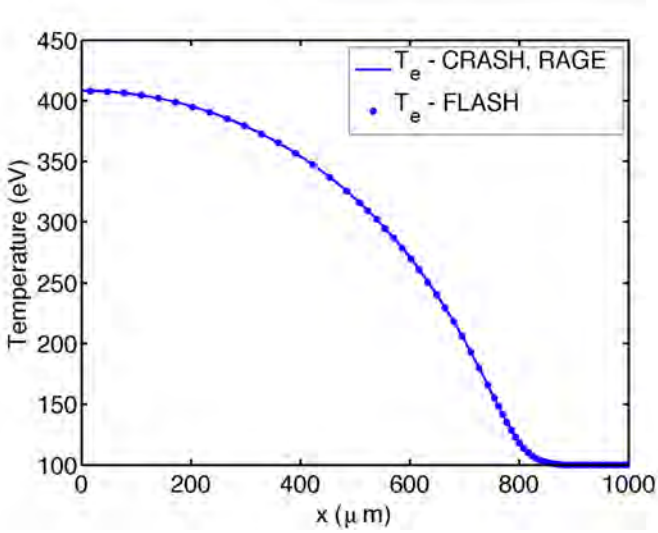

Figure 97: Electron temperature profile obtained by applying electron heat conduction to the temperature profile shown in Figure 96. The results are plotted at time 1.5 ns. The spatial profiles obtained by CRASH and xRAGE are represented by the solid line, while the FLASH results are plotted as individual ponits. The three codes give indistinguishible results for this problem.

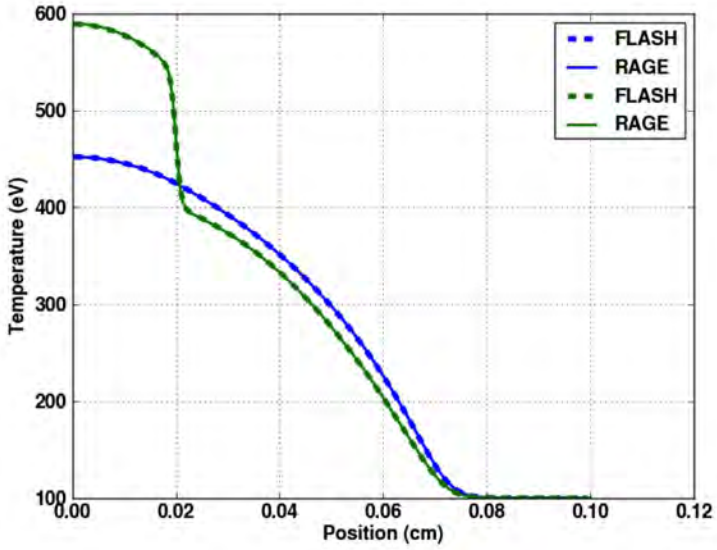

Figure 98: Results at 1.5 ns for the electron conduction test problem when ion-electron equilibration is included. xRAGE results are shown as solid lines, while the FLASH results aree represented by dashed lines. The code are again in excellent agreement. Although not shown here, the results obtained by CRASH were virtually identical to those obtained by the other two codes.

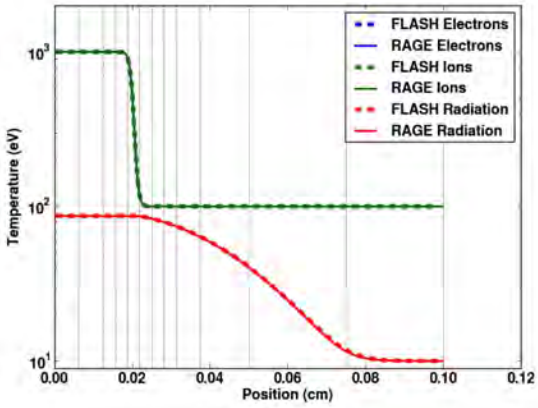

Figure 99: Single-group radiation diffusion test.. All three temperatures are initially set to tanh profiles, as described in the text. The ion temperature profile is plotted in green and overlies the blue curves for the electron temperature. The radiation temperature is shown in red. Since only the radiation diffusion terms are turned on, the electron and ion temperatures maintain their original profiles. Only the radiation temperature shows the effects of the diffusion operator. FLASH and xRAGE give good agreement for this problem, which was not attempted with CRASH.

of 40 eV, with the fluid at rest. For the initial downstream conditions, the density is 0.004466 gm/cm^3, the electron and ion temperatures are 102.4393 eV, and the fluid velocity is 9.9635×10^6 cm/s. The plasma is assumed to be fully ionized helium with an adiabatic index of 5/3. For these conditions, the steady shock speed is 1.6234×10^7 cm/s.

The steady profiles of the ion and electron temperatures near the shock obtained by the three codes are shown in Figures 101 and 102. The FLASH results are represented by the blue line, the xRAGE results by the red line, and the CRASH results by the dashed green line. FLASH was run using its entropy advection algorithm for this test. For the sake of comparison, the shock locations were adjusted to coincide, since they achieved their steady profiles at different times. The most difficult feature to get right in this problem is the peak ion temperature. The analytic value is 124.2 eV. FLASH achieved the closest value, reaching a peak ion temperature of 122.8 eV. CRASH was not far behind, with a value of 122.5 eV. Initially, xRAGE reached a peak value of only 117 eV, as shown in Fig. 11, but after a small error in the setup of the problem was uncovered, the xRAGE value improved to 121.6 eV. These small differences are not likely to be significant for practical problems, and the agreement between the three codes and the analytic solution is completely acceptable.

4.4.3 Summary

Three High-Energy-Density Physics (HEDP) codes (CRASH, FLASH, and xRAGE) have been compared on a variety of simple one-dimensional test problems. The subset of problems described in this paper tested the ion/electron equilibration terms, the radiation emission and absorption terms, and the diffusion modules in the codes. A final test combined hydrodynamic effects, including shock propagation, with electron heat conduction and ion/electron equilibration. At the start of the collaboration, significant differences existed between the codes for several of the

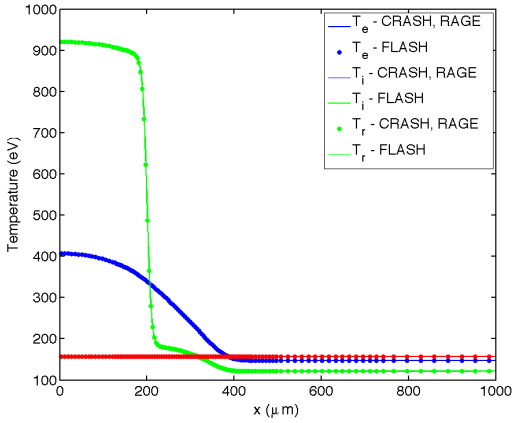

Figure 100: Diffusion test including electron heat conduction, single-group radiation diffusion, ion-electron equilibration, emission and absorption. The electron temperature is plotted in blue, the ion temperature in green, and the radiation temperature in red. The CRASH and xRAGE results are both plotted as solid lines and lie on top of each other. The FLASH results are represented by individual points and agree very well with the results from the other two codes. The radiation evolves on a much faster time scale than the electrons and ions, so at this time (0.2 ns) the radiation temperature gradient has almost completely diffused away

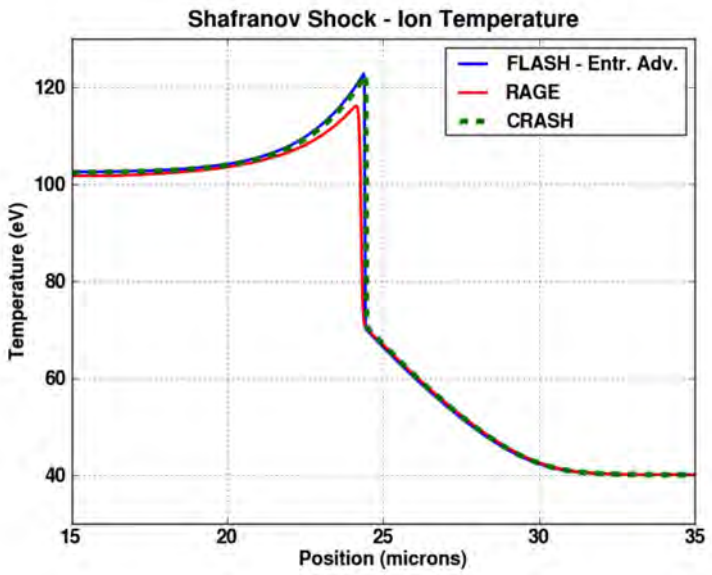

Figure 101: Steady-state ion temperature profiles for the Shafranov problem obtained by the three codes. The problem consists of a steady shock with unequal electron and ion temperatures. Electron heat conduction and ion-electron equilibration are both included, but radiation diffusion is not. The FLASH results, using the entropy advection algorithm, are represented by the blue line, the xRAGE results are represented by the red line, and the CRASH results are given by the dashed green line. The codes obtain slightly different peak value for the ion temperature. The low value of the xRAGE value was due to a slight error in the setup of the problem, which has since been corrected. All three codes are now in good agreement with each other and with the analytic solution.

comparisons. These differences were due to a number of factors. In some cases, the codes used different formulas for the various transport coefficients being tested. In other cases, programming errors were discovered or errors were made in the setup of the test problem. In still other cases, there were differences caused by using time steps that were too large. Once all these issues were resolved, the codes produced nearly identical results. Future comparisons will also include tests using flux limiters for the conduction and radiation diffusion, multidimensional problems, and perhaps tests of additional physics modules, such as MHD. The ultimate goal of this collaboration is to compare the codes on simulations of full HEDP experiments and to either get good agreement or at least to understand the causes of any remaining discrepancies.

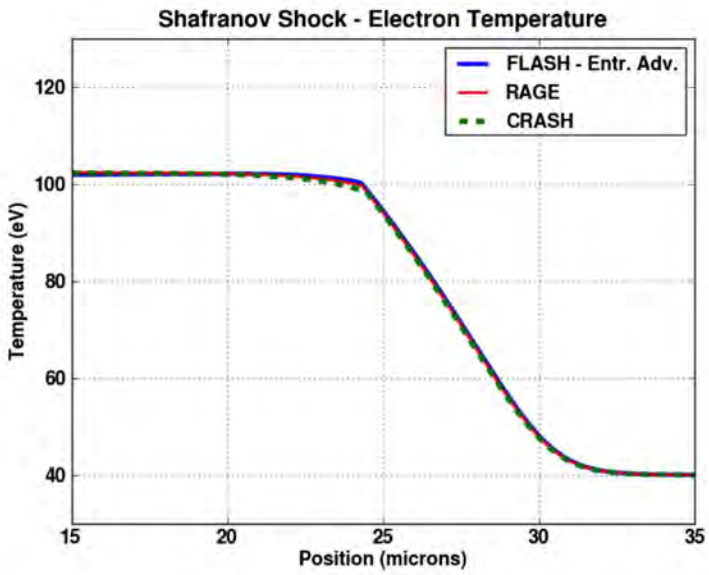

Figure 102: Steady-state electron temperature profiles for the Shafranov problem obtained by the three codes. The FLASH results are shown by the solid blue line, the xRAGE results by the solid red line, and the CRASH results by the dashed green line. Only very small differences between the three results can be seen.

5 Evaluation of Model Form Uncertainties

In any complex, multiphysics simulation approximate models must be used. This is a source of uncertainty in the results, often referred to as "model form uncertainty". During this project we performed several studies aimed at evaluating to some extent the errors that might be associated with various model assumptions in the CRASH code. This was not a primary area of emphasis in the project, but we did make some progress, discussed in this section.

5.1 Transport–diffusion comparison using PDT and CRASH

We began this project with the intention of coupling the CRASH code with PDT, a highly scalable, radiation-transport code from Texas A&M University, that solves that full Boltzmann Transport Equation. This would have provided a model with higher fidelity than we accomplished. In pursuit of this goal we developed a version of PDT that calculates the transport of thermal radiation. Coupling this to CRASH would have given us a code with capabilities similar to those of a number of European codes used in solar-physics studies. However, combining these codes proved impractical once we became familiar with some governmentally imposed limitations on the use of such a combination. In addition, the limited computational resources that proved to be available limited the utility of even 3D calculations with multigroup, flux-limited diffusion (FLD), and so we would have been unlikely to have accomplished predictive science studies with such a combined code anyway. Facing these limitations, we instead focused the use of PDT on the quantification of the errors in the diffusion method implemented in CRASH. The research reported here is a summary a portion of the project that aimed to quantify the error associated with the use of radiation-diffusion to model radiation

in the CRASH project.

The most important computational tools in carrying out this study are (i) PDT, a highly scalable, radiation-transport code from Texas A&M University, that solves that full Boltzmann Transport Equation, but which lacks modules to calculate material hydrodynamics, and (ii) the CRASH code, developed during the course of the CRASH project, which is also highly scalable. It is a flux-limited-diffusion-based radiation-hydrodynamics code with the ability to simulate the evolution of a radiation-hydrodynamic system relevant to the CRASH project.

To gain insight into the error associated with flux-limited diffusion, we constructed a suite of code-to-code-comparison test problems. This suite consists of analytical and near-analytical tests in the diffusion regime (where transport and diffusion are expected to produce nearly identical solutions) as well as tests that are more reflective of the physical conditions found within the CRASH experiment.

In performing a code-to-code comparison that involves codes based on different underlying equations, several concerns arise. Because of the dimensional reduction that is fundamental to radiation-diffusion approximations, diffusion and transport problems generate results in different solution spaces. The question then arises as to how one performs a meaningful comparison. Of additional concern, and also inherent in dimensional reduction, is the need to add, what are essentially, ad hoc features to a diffusion model (e.g., limiters, diffusion coefficients). Another concern involves consistent treatment of matter–radiation energy exchange: Except for simple problems, collision integrals cannot be made fully consistent between models.

Aside from differences already listed, there are also algorithmic and code-implementation differences between the codes that we needed to consider before drawing meaningful conclusions. In general terms, these include issues such as discretization, time-step selection, and solution methods. The mode in which a code is used for a given test problem can vary significantly from the mode that is used for production calculations.

Finally, there are operational considerations. A test problem that is easy to set up and run for one code can be difficult for another. These differences can be extreme enough to preclude any ability to run a comparison test for a meaningful length of time. An example of this, relevant to radiation transport, is diffusion-regime problems. Diffusion-based algorithms are set up to run diffusion-regime problems efficiently. However, even allowing for the much larger number of variables solved in transport calculations, diffusion-regime problems are often computationally taxing on a transport code.

5.1.1 Testbed for the comparison study

In most conceivable regimes encountered in high-energy-density physics, radiation is properly described by the Boltzmann Transport Equation (hereafter, transport equation), which describes the evolution of radiation through six-dimensional phase space—three spatial and three momentum dimensions. For the purpose of ease and speed of computation, flux-limited diffusion (FLD) has been developed as an approximation to the full transport equation. In a flux-limited formulation, two of the angular dimensions of momentum space have undergone averaging, which results in dimensional reduction. As an approximation, it sacrifices solution accuracy as well as the ability to describe radiation in complete generality.

The transport equation, as used by PDT, can be written in mixed-frame form as

$$\begin{aligned}&\frac{1}{c}\frac{\partial I(\mathbf{r},\hat{\Omega},\varepsilon,t)}{\partial t}+\hat{\Omega}\cdot\nabla I(\mathbf{r},\hat{\Omega},\varepsilon,t)\\&=\frac{1}{4\pi}\rho(\mathbf{r},t)j(\mathbf{r},\varepsilon,t)-\kappa(\mathbf{r},\varepsilon,t)I(\mathbf{r},\hat{\Omega},\varepsilon,t)\\&+\rho(\mathbf{r},t)\frac{\mathbf{u}(\mathbf{r},t)}{c}\cdot\mathbf{C}(\mathbf{r},\hat{\Omega},\varepsilon,t)\,,\end{aligned}\tag{227}$$

where I is the specific intensity, expressed as a function of position $\mathbf{r}$, time t, $\hat{\Omega}$ the vector solid angle, indicating the direction of photon propagation, and ε the photon energy. The material density is given by ρ, and j and κ are the material emissivity and opacity, respectively. In the event that the background material is moving relative to the laboratory frame, the quantity $\mathbf{C}$ contains the information needed to convert terms containing comoving emissivities and opacities for use in inertial-frame equations. For the problems we have considered in this study, the material velocity $\mathbf{u}$ is always zero. Thus, we ignore the distinction between laboratory and comoving frames.

The radiation module in the CRASH code is formulated using multigroup flux-limited diffusion. The formal derivation is constructed in the comoving frame, where the zeroth angular moment is taken of the boosted transport equation. (Although physically equivalent, the boosted equation on which this FLD derivation is based has a different form from Eq. 227.) Upon dropping terms of higher order than u/c, and substituting Fick's Law diffusion for the radiative energy flux density $\mathbf{F}$, we obtain the monochromatic radiation energy equation

$$\frac{\partial E_\varepsilon}{\partial t}+\nabla\cdot(\mathbf{u}E_\varepsilon)-\nabla\cdot(D_\varepsilon\nabla E_\varepsilon)-\varepsilon\frac{\partial}{\partial\varepsilon}\left(\mathsf{P}_\varepsilon:\nabla\mathbf{u}\right)=4\pi j_\varepsilon-c\kappa_\varepsilon^{\text{Planck}}E_\varepsilon.\tag{228}$$

Here, E_ε is the monochromatic radiation energy density (in units of energy per unit volume per unit spectral energy interval), which is the zeroth angular moment of I, P is the radiation pressure tensor, which is the second angular moment of I, ε is the spectral photon energy, j_ε is the radiative emissivity of the matter and $\kappa_\varepsilon^{\text{Planck}}$ is its Planck mean opacity, each at photon energy ε. We note, again, that for problems of interest in this study, the material velocity, $\mathbf{u}$, is always zero and, hence, the term in $\nabla\mathbf{u}$ vanishes. The monochromatic diffusion coefficient is given by

$$D_\varepsilon=\frac{\lambda_\varepsilon c}{\kappa_\varepsilon^{\text{Ross}}},\tag{229}$$

where $\kappa_\varepsilon^{\text{Ross}}$ is the Rosseland mean opacity at photon energy ε, and λ_ε is the flux limiter used in the CRASH code. For this limiter, we have used the well-known prescription of Larsen,

$$\lambda_\varepsilon^{\text{Larsen}}=\frac{1}{\left[3^n+\left(\frac{1}{\kappa_\varepsilon^{\text{Ross}}}\frac{|\nabla E_\varepsilon|}{E_\varepsilon}\right)^n\right]^{1/n}},\tag{230}$$

in its $n=2$. form. As part of the study, we also investigated heterogeneous diffusion and the effects of the treatment of diffusion-coefficient interpolation across cell interfaces in high-energy-density-physics (HEDP) regimes.

In the foregoing, three opacities appear, the Planck mean, the Rosseland mean and, in Eq. 227, an opacity that is, as yet, unspecified. In diffusion theory, use of σ_{Ross} yields the correct spatial energy transport in the diffusion limit, while σ_{Planck} yields the correct local balance of matter–radiation energy exchange. This introduces a choice as to which opacity to assign to Eq. 227. In the results we present here, we follow a frequently used convention and employ the Rosseland mean. Because the Rosseland formulation is based on a harmonic average of opacities over a spectral

range, it tends to make matter appear more transparent to radiation than the Planck mean, especially if the radiation spectrum is insufficiently resolved.

Details of the CRASH code are presented elsewhere in this report. PDT, the transport code used in this study, is a deterministic radiation transport code, where the specific intensity can be computed in either a multigroup or spectrally integrated gray version. In the code comparisons we consider here, both multigroup and gray transport modules are used. PDT uses a discrete-ordinates discretization to describe the direction of radiation propagation. In all the simulations reported here, 256 angles have been used. Spatially, PDT uses a discontinuous finite element method and is fully implicit in time.

For this study, all problems in which the CRASH code is used are run in the hydrostatic limit. Similar to PDT, the radiation energy can expressed in either a multigroup or gray formulation. Both are used here. As outlined above, the flux-limited diffusion approximation sacrifices the ability to resolve the angular dependence of radiation. Instead, radiation is constrained to propagate in the direction of the negative gradient of E_ε.

In all results reported here, both codes use vacuum boundary conditions, where the specific intensity (or an approximation to it, in the case of FLD) drop to zero outside the solution domain. Vacuum boundary conditions introduce special issues in FLD codes, which we examined in the context of our current test suite.

To help better understand discrepancies between the results of transport and diffusion calculations, both PDT and CRASH have been outfitted with the capability of tracking energy flow and its conservation throughout the computational mesh. In addition, bookkeeping is available to track radiative energy flow off the mesh. This is especially important for gaining insight into the results some test problems based on the CRASH experimental setup. Since the CRASH-experimental domain contains many regions that are in the translucent and free-streaming regimes, a great deal of radiation flows off the mesh—both in test problems as well as in realistic radiation-hydrodynamic simulations of the full CRASH experiment.

5.1.2 Code-comparison problems and results

The starting point for our code-to-code comparison were infinite-medium test problems. Transport and diffusion should agree almost exactly for these problems, since none of the physical quantities has spatial dependence. Quantities in infinite-medium problems evolve as a result of non-zero source terms, which take similar forms in both transport and diffusion when spatial dependence is absent. Since such tests are relatively simple to construct, any discrepancy between codes can be used to uncover possible code implementation issues and to check on the equality of local rates for relevant physical processes. As part our testing, we ran a number of infinite-medium code-comparison problems and found precise agreement between PDT and CRASH, as expected.

The next step in adding complexity to our code comparison was incorporating spatial dependence into the tests. Radiation fronts were used to test wave propagation and the implementation of transport coefficients. Marshak waves were used to test both wave propagation and local rates. Both codes also have the option of adding heat sources to the matter. This mechanism can be used as a proxy for shock heating. Throughout the problems in this class, we used gray transport and simple analytic opacities. However, to reflect the conditions we face in more realistic problems, opacities were allowed to vary sharply across an interface.

Once satisfactory agreement was achieved on simpler tests, the next step was to set up test problems that more accurately reflected the actual CRASH experiment and its computational models. Agreement here was not expected

to be as close as in diffusion-limit problems. This is because much of the computational domain in CRASH is in the intermediate transport regime (i.e., neither diffusive nor free streaming). It is here where diffusion is known to be a poor substitute for full transport. Assessing the importance of any observed discrepancies in these problems has been a key goal of our investigation.

5.1.2.1 Diffusion front

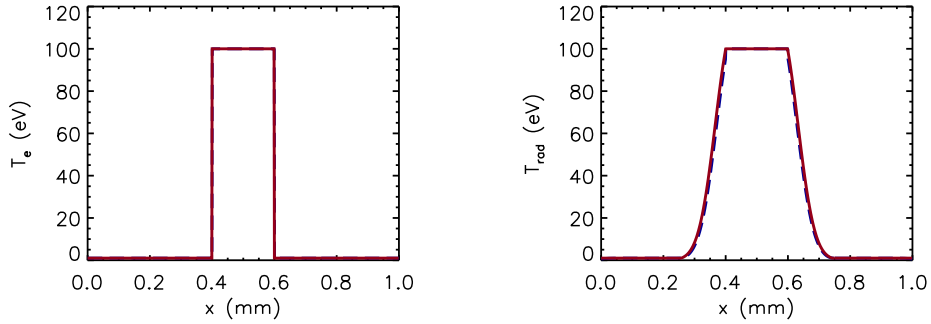

Figure 103: Images of the diffusion front problem described in Section 5.1.2.1, showing the initial electron-temperature ($T_{\rm elec}$) profile (left) and the radiation-temperature ($T_{\rm rad}$) profile (right) after 3.0 ps of evolution. PDT results are shown by the solid maroon line; CRASH results are shown in the dashed blue line. At 3.0 ps, both $T_{\rm rad}$ profiles are virtually identical, and the dashed curve is difficult to discern.

One of the most straightforward tests of radiation propagation is the simulation of a diffusion wave. This test is similar to the classic light front problem except that it is run in the diffusion, rather than free-streaming limit. Our version of the problem is constructed using gray transport in both codes. The test is run in a 1D medium of uniform density 1 g cm^{-3}, which extends from 0–1.0 mm. Initially, the medium is set to have identical electron and radiation temperatures ($T_{\rm elec}$ and $T_{\rm rad}$, respectively) of 1.0 eV, except for a central strip extending from 0.4–0.6 mm, where $T_{\rm elec} = T_{\rm rad} = 100$ eV. (The initial setup for $T_{\rm elec}$ is shown in the left image of Fig. 103.) In this test, we use the hot central strip as the source of radiation that initiates the front. The cooler regions are modeled as a medium in which the matter–radiation interaction is restricted to purely elastic scattering. Since neither CRASH nor PDT includes scattering modules, scattering is simulated by appropriate settings of microphysical parameters.

For CRASH, we set the opacities such that

$$\sigma_{\rm Ross} = \begin{cases} 10^4 \text{ cm}^2 \text{ g}^{-1} & x < 0.4 \text{ mm}; \\ & x > 0.6 \text{ mm} \\ 10^5 \text{ cm}^2 \text{ g}^{-1} & 0.4 \leq x \leq 0.6 \text{ mm} \end{cases} \tag{231}$$

$$\sigma_{\rm Planck} = \begin{cases} 0 & x < 0.4 \text{ mm}; \\ & x > 0.6 \text{ mm} \\ 10^5 \text{ cm}^2 \text{ g}^{-1} & 0.4 \leq x \leq 0.6 \text{ mm}. \end{cases} \tag{232}$$

Setting $\sigma_{\text{Planck}} = 0$ outside the central strip insures no matter–radiation energy exchange occurs in this region, while a finite σ_{Ross} ensures that diffusion takes place in a manner consistent with elastic scattering being the sole source of opacity.

As noted earlier, the equation of radiation transport (Eq. 227) contains only one kind of opacity. For the purposes of the present test, we set its value as follows:

$$\sigma = \begin{cases} 10^4 \text{ cm}^2 \text{ g}^{-1} & x < 0.4 \text{ mm}; \\ & x > 600 \text{ mm} \\ 10^5 \text{ cm}^2 \text{ g}^{-1} & 0.4 \leq x \leq 0.6 \text{ mm}. \end{cases} \tag{233}$$

In PDT, in the cool regions, elastic scattering is modeled differently: matter–radiation energy exchange is effectively turned off by setting the specific heat in the cool regions to zero.

The right-hand image of Fig. 103 shows T_{rad} in the system after 3.0 ps of evolution by both codes. (T_{elec} is essentially unchanged from its initial values.) The two curves, representing PDT and CRASH results, are virtually identical. We note that the diffusion time scale for this problem is given by $t_{\text{diff}} = l^2/D \sim 10$ ns, where $l \approx 0.1$ cm. In contrast, the free-streaming time scale is $t_{\text{fs}} = l/c \approx 3$ ps, which, with our results, confirms that we are indeed computing in the diffusion limit.

5.1.2.2 Marshak wave

As a modification to the diffusion-front problem described above, we next allow emission–absorption in all regions of the computational domain. This allows us to follow the evolution of a Marshak wave. In the test described here, we also include a heat source in the central strip, which adds internal energy to this region at a rate of 4.25×10^{33} eV cm^{-3} s^{-1}. The initial temperatures are identical to those in the diffusion-front problem: $T_{\text{elec}} = T_{\text{rad}} = 100$ eV within the central strip; $T_{\text{elec}} = T_{\text{rad}} = 1$ eV elsewhere.

Since we find the Marshak wave has a much slower propagation speed than the diffusion front, we lower the opacity in the initially cooler region to allow the problem to be calculated in a reasonable time. Thus, for both CRASH and PDT all opacities are set as

$$\sigma = \begin{cases} 10^3 \text{ cm}^2 \text{ g}^{-1} & x < 0.4 \text{ mm}; \\ & x > 0.6 \text{ mm} \\ 10^5 \text{ cm}^2 \text{ g}^{-1} & 0.4 \leq x \leq 0.6 \text{ mm}. \end{cases} \tag{234}$$

Lowering the opacity in the cooler regions reduces the diffusion time scale to ~ 1 ns, which is still large relative to the free-streaming time scale.

Figure 104 shows this system after 100 ps of evolution. Although agreement between diffusion and transport is good, there are small discrepancies apparent. Most visible are differences between both T_{elec} and T_{rad} in the central strip. CRASH shows consistently higher temperatures in this region. At least part of this disagreement may be explained by differences in the energy-transfer mechanism between diffusion and transport. This difference creates a surface vs. volume effect: At the beginning of the calculation, the radiation-energy density is a pure step function. For diffusion, energy can flow only where there is a gradient in radiation energy. This occurs only at the step—the surface of the strip. In contrast, full transport constrains the flow of radiation only via the opacity of the medium. This allows

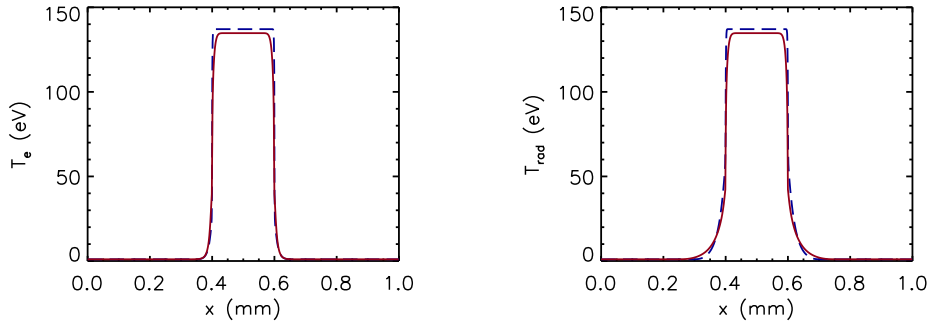

Figure 104: Results for the Marshak wave problem described in Section 5.1.2.2, showing electron-temperature ($T_{\rm elec}$) profiles (left) and the radiation-temperature ($T_{\rm rad}$) profiles (right) after 100 ps of model evolution. PDT results are shown by the solid maroon line; CRASH results are shown by the dashed blue line. At 100 ps, results for both $T_{\rm elec}$ and $T_{\rm rad}$ profiles agree well, but are not identical. Differences in energy transport between diffusion and full transport are the likely cause of the discrepancies.

emergent radiation to have originated directly from within the volume of the strip. This explanation is also consistent with our observation that $T_{\rm rad}$ is slightly higher in the initially cool medium in PDT relative to CRASH.

5.1.2.3 1D CRASH-based problems

A key element of helping determine the error in using FLD in the CRASH project has been the construction of a set of realistic test problems that, in a hydrostatic setting, can compare the results of diffusion and full transport. A schematic of the setup of two classes of such tests is shown in Fig. 105. In the problems based on this setup, the geometry of the CRASH shock tube is approximated by a 2D Cartesian central cross-section of the tube. The domain consists of regions containing the same materials in approximately the same configuration as used in the CRASH experiment (Be, Xe, polyimide, and Au).

Throughout the simulations, realistic opacities, using the CRASH opacity tables, (as developed by CRASH team member I. Sokolov) are employed. As shown in the figure, a heat source acts as a proxy for shock heating, adding 4.25×10^{33} eV cm^{-3} s^{-1} to the internal energy of heated area. Initially, $T_{\rm elec} = T_{\rm rad} = 1.0$ eV.

In the first set of results we describe here, we restrict ourselves to a horizontal 1D cross-section of the domain shown in red in Fig. 105, where a cross-section has been taken halfway through the vertical extent of the tube (as shown by the blue dotted line). In this 1D *longitudinal* version of the problem, there are two materials present—Be and Xe. For this test, we set the specific heat capacities as $c_V^{\rm Be} = 1.1 \times 10^{19}$ eV g^{-1} K^{-1} and $c_V^{\rm Xe} = 9.9 \times 10^{17}$ eV g^{-1} K^{-1}. Figure 106 shows the density profile (which remains constant throughout the simulation) and the initial Rosseland gray opacity (which the changing temperature affects as the system evolves). Electron heat conduction is not included in these calculations. Vacuum boundary conditions are used by both codes.

Agreement between transport and diffusion is considerably better in multigroup calculations than in gray. (See Myra and Hawkins [2013] for details.) In Figs. 107–110, we show calculations using both PDT and CRASH, which have been run using 10 logarithmically spaced energy groups. (The total spectrum is bounded by 1.0 eV and 20.0 keV.)

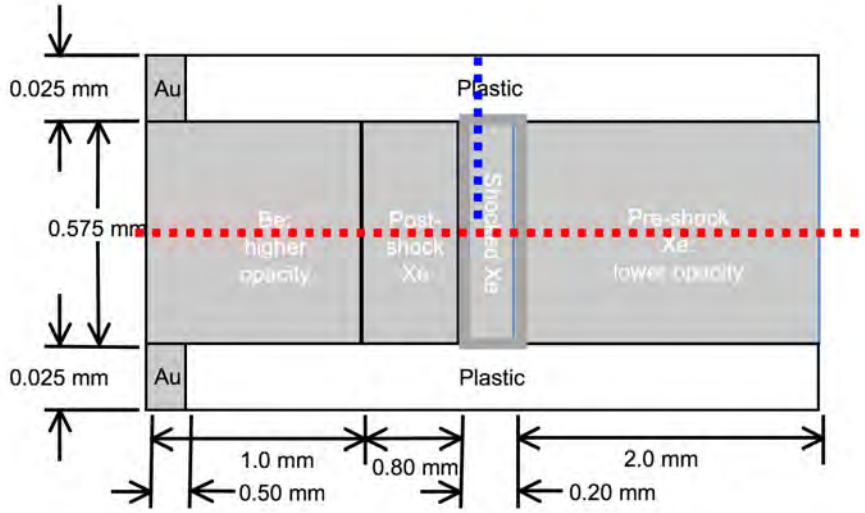

Figure 105: A schematic diagram showing a class of test problems, based on the CRASH experiment, which is used as a method to quantify the error associated with using flux-limited diffusion instead of full transport in CRASH simulations. Shock heating is approximated by a heat-source, which is active in the region marked as "Shocked Xe," near the center of the diagram. Only the results from two types of 1D calculations are presented in this report. In the first type, the single dimension corresponds to the 1D lineout shown as a red dotted line in the figure. In the second class, the 1D computational domain corresponds to the blue dotted.

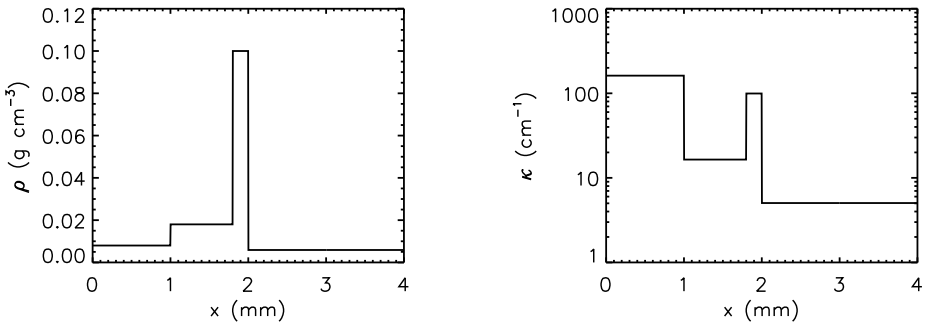

Figure 106: For the 1D CRASH-based longitudinal problem, the density profile (which remains constant throughout a simulation) is shown in the left-hand figure. The initial Rosseland opacity for a gray calculation is shown at the right. Since we are using realistic material opacities in this problem, these values will change as the system evolves, reflecting their dependence on the changing $T_{\rm elec}$.

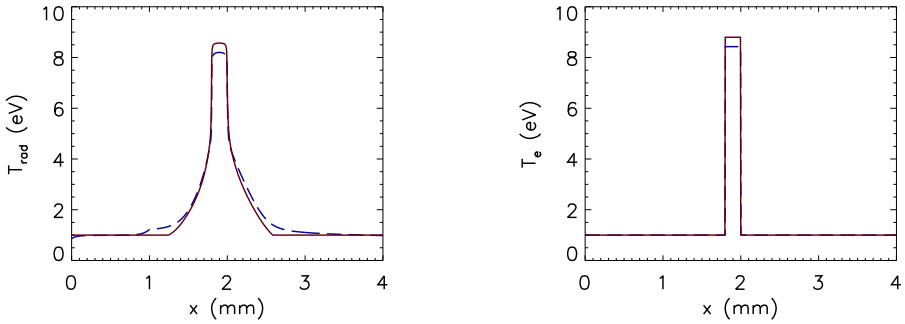

Figure 107: Agreement between transport and FLD on the 1D CRASH-based longitudinal problem improves with number of groups used in multigroup radiation–shown here for 10 groups. The comparison is at $t = 2.0$ ps. Results of radiation transport (maroon, solid curve) using PDT and flux-limited diffusion (blue, dashed curve) using the CRASH code. Vacuum boundary conditions are used for both codes.

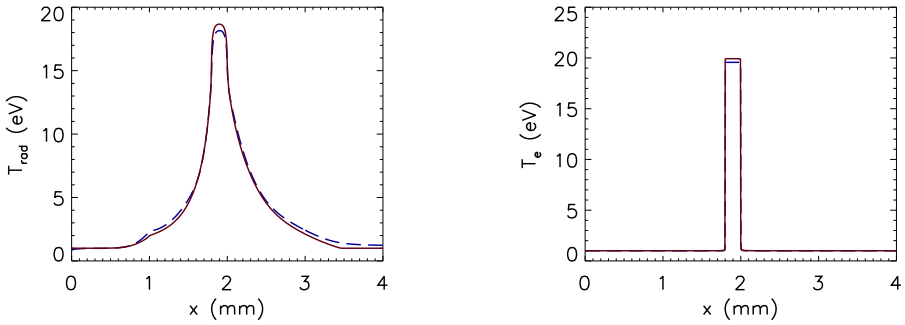

Figure 108: Agreement between transport and diffusion solutions continues to be good for the 10-group version of the 1D CRASH-based longitudinal problem. Here, the comparison is at $t = 5.0$ ps, as advanced from the results shown in Fig. 107. PDT transport results are plotted in the maroon, solid curve and CRASH FLD results are in the blue, dashed curve.

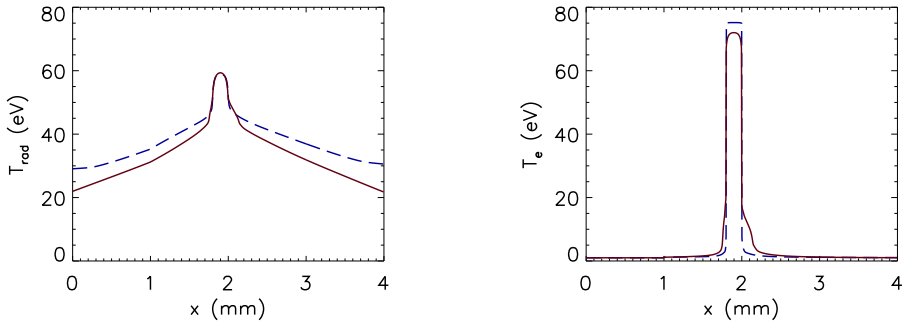

Figure 109: By $t = 20.0$ ps, shown in this figure, transport and diffusion are starting to show some noticeable disagreement for the 10-group version of the 1D CRASH-based longitudinal problem. PDT transport results are plotted in the maroon, solid curve and CRASH FLD results are in the blue, dashed curve.

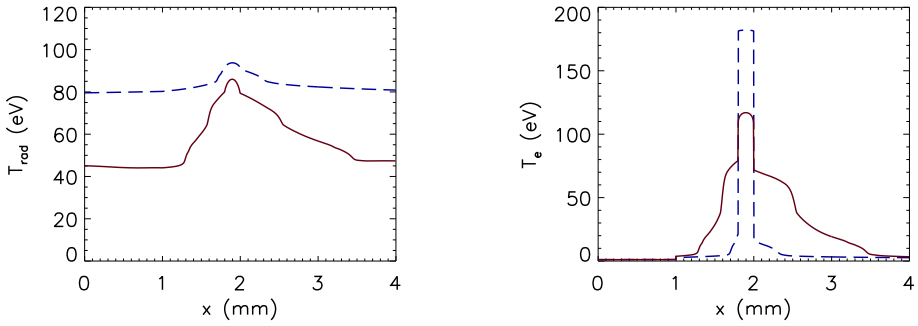

Figure 110: By $t = 50.0$ ps, shown in this figure, transport and diffusion solutions have significantly diverged for the 10-group version of the 1D CRASH-based longitudinal problem. Also noticeable is the significant signature that radiation discrepancies have imparted on the matter, as shown by T_{elec}, on the right. PDT transport results are plotted in the maroon, solid curve and CRASH FLD results are in the blue, dashed curve.

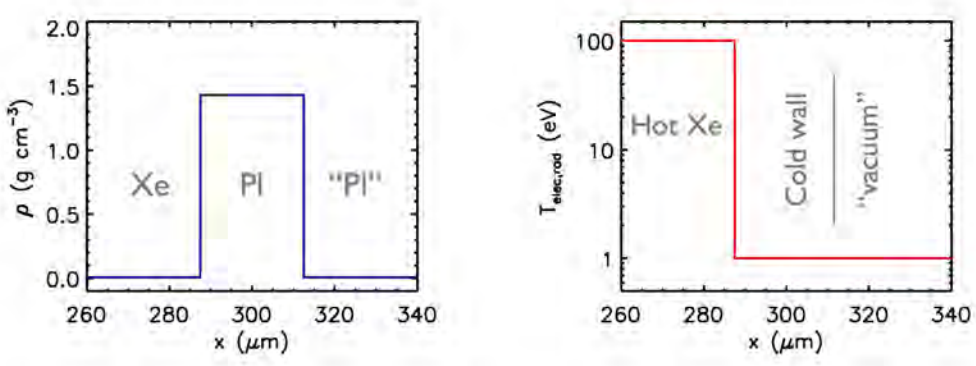

Figure 111: Initial conditions for CRASH-based transverse 1D problem in density (left figure) and electron and radiation temperatures (right figure). The domain for this problem is shown by the dotted blue line near the center of Fig. 105. The diffuse area on the left side of the domain corresponds to the center of the shock tube in the full CRASH problem. Here, the temperature is initially set to 100 eV for both matter and radiation. The denser plastic region, labeled as "Pl" corresponds to the plastic tube in the CRASH experiment. The temperatures here are initially set to 1 eV. Finally the diffuse region on the right is a near-vacuum, which is is approximated by cold diffuse plastic.

Agreement between transport and diffusion is excellent at 2.0 and 5.0 ps; it is good at 20 ps; however, by 50 ps, substantial divergence is evident in both $T_{\rm rad}$ and $T_{\rm elec}$. In analysis made subsequent to the original publication of these results, subtle differences between the codes in how vacuum boundary conditions are implemented may be a contributing factor to the worsening agreement with time. In this set of problems, by 50 ps, a considerable amount of energy has left the grid in both the transport and diffusion versions of the problem. Additionally, since transport and diffusion propagate energy at slightly varying rates within the mesh, by the time the first waves reach the boundaries, they will look different in each code. Hence, even if extrapolation lengths are identical in the transport and radiation codes, the rate of energy flow will not be. Once that situation obtains, the discrepancies between solutions continue to mount.

Another source of discrepancy between transport methods concerns the choice of opacity in the calculations. As mentioned above, use of the Rosseland opacity ensures accurate transport of radiative energy in the diffusion limit, while using the Planck opacity ensures accurate matter–radiation energy exchange locally. In each calculation, we ran the CRASH code in two different modes: (i) using Planck opacities for emission–absorption and Rosseland opacities to calculate diffusion coefficients, as is the standard mode for full-scale CRASH simulations; and (ii) using Rosseland opacities for both quantities. (In this study, all PDT calculations use only Rosseland opacities.) For the results we present in this report, CRASH has been run in mode (ii), which gives us significantly better agreement with PDT. The discrepancy between Planck and Rosseland results likely indicates that the energy-group structure is insufficiently fine to resolve important features in the material opacity profiles.

For the CRASH project, accurate solution for $T_{\rm elec}$ is considered more important than for $T_{\rm rad}$ since it is material properties that are directly measurable by experiment. Obviously, however, $T_{\rm rad}$ and $T_{\rm elec}$ are closely coupled.

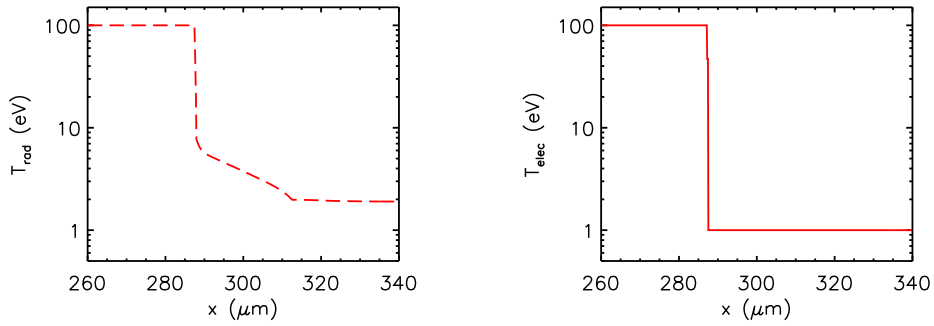

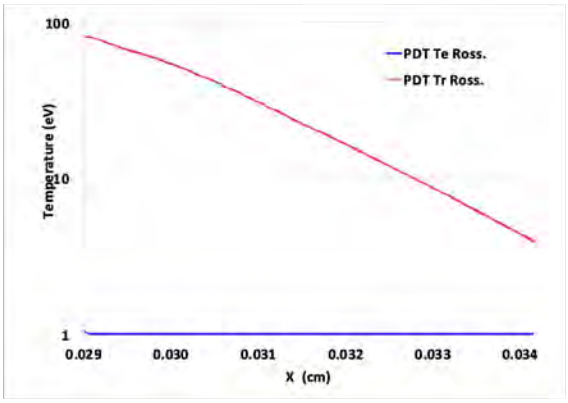

Figure 112: Results showing T_{rad} (top left) and T_{elec} (top right) for the CRASH code running the transverse problem, after 0.05 ps of evolution. On the immediate left are corresponding PDT results. A "prompt" radiation front has propagated into the plastic in both sets of results, though PDT shows greater penetration—to the point of significant propagation into the vacuum region. In both codes, T_{elec} is essentially unchanged.

The second set of CRASH-based test problems we investigated represented a *transverse* lineout within the CRASH shock tube. This region is indicated by the vertical dotted blue line near the center of Fig. 105. It serves to test the effect of hot x-rays incident on an initially cool plastic wall. With this setup, we can compare the differences between radiation transport and diffusion in the response of the wall. This details of the response of the wall is important due to wall ablation and the formation of a wall shock, which strongly influences the primary shock in the CRASH problem. In Fig. 111, the diffuse area on the left side of the domain corresponds to the center of the shock tube in the full CRASH problem. Here, the temperature is initially set to 100 eV for both matter and radiation. The denser plastic region, labeled as "Pl" corresponds to the solid portion of the plastic tube in the CRASH experiment. The temperatures here are initially set to 1 eV. Finally the diffuse region on the right is a near-vacuum region that is approximated by cold diffuse plastic. For the results presented here, multi-group radiation was used, with five energy groups spaced logarithmically between 1 eV and 20 keV.

Results for both CRASH and PDT are shown in Figs. 112 and 113 at 0.05 ps and 200 ps of evolution, respectively. Although there has been considerable evolution of the radiation field by 0.05 ps, there is also considerable discrepancy between the results of the two codes. By this time, a "prompt" radiation front has propagated into the plastic in both calculations, though PDT shows greater penetration of the plastic—to the point of there being significant propagation into the vacuum to the right of the tube. In both codes, T_{elec} is essentially unchanged.

In contrast, by 200 ps, CRASH and PDT results agree well. At this time, within the dense plastic, evolution of T_{rad} and $T_{\text{elec}}$is are seen to be driven by the competing effects of radiation transport through the plastic wall and thermal

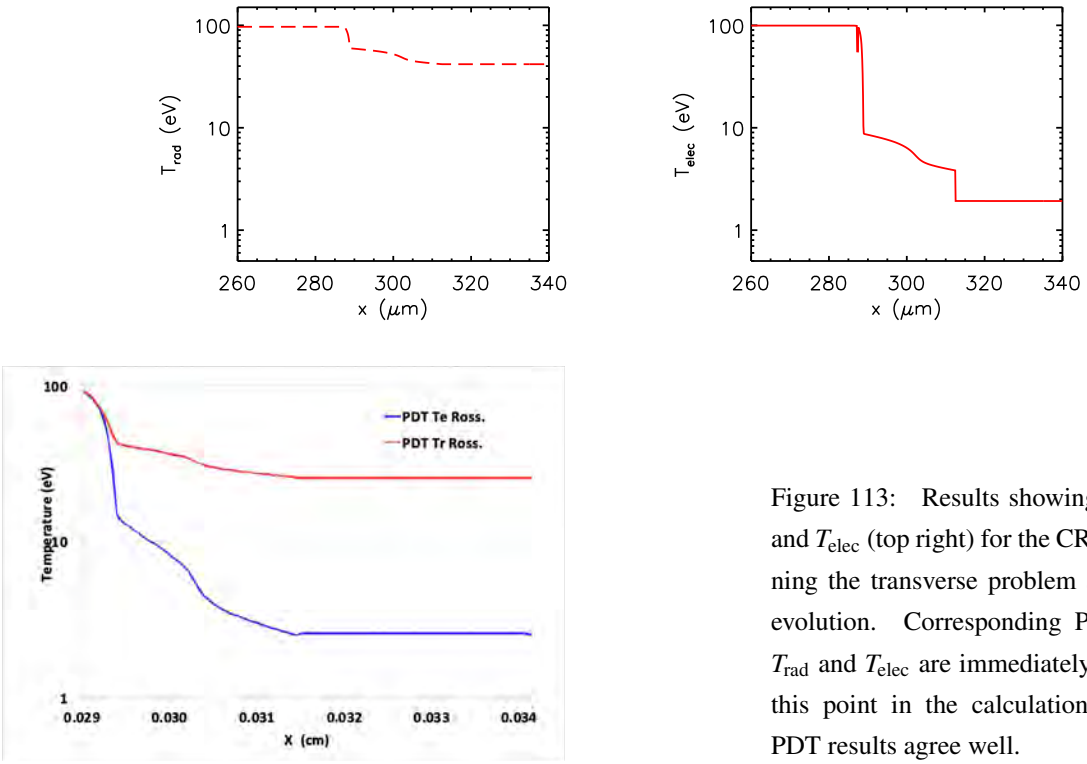

Figure 113: Results showing T_{rad} (top left) and T_{elec} (top right) for the CRASH code running the transverse problem after 200 ps of evolution. Corresponding PDT results for T_{rad} and T_{elec} are immediately to the left. At this point in the calculation, CRASH and PDT results agree well.

equilibration between matter and radiation. With the 100-eV boundary condition on the left-hand edge, and given enough time, a steady-state temperature of $T_{\text{rad}} = T_{\text{elec}} = 100$ eV will eventually obtain.

An interesting contrast between the results of the transverse problem shown in Figs. 111–113 and the longitudinal problem shown earlier in Figs. 106–110 is that in the longitudinal case, the agreement between CRASH and PDT is at its best early in the calculation, while in the transverse problem, agreement improves with time. As mentioned above, an important contribution to late-time divergence in the longitudinal problem likely results for significant energy flow off the mesh as time advances. The setup of the transverse problem setup does not allow this to occur to any great extent over the time scales of interest. As a result, we conclude that the transverse problem constitutes a better code-to-code comparison test. In analogy, it is possible to revise the longitudinal problem to exhibit the same behavior—a process in which we are currently engaged [Myra and Hawkins, 2013].

5.1.3 CRASH-code issues uncovered

5.1.3.1 Heterogeneous Diffusion

In computational problems in high-energy-density physics, opacities may vary by significant amounts from cell to cell. Commonly, these large changes occur due to the presence of material interfaces and/or sharp changes in density between adjacent cells. In these situations, obtaining a physically meaningful radiative flux from diffusion-based radiation codes requires special care.

In finite volume codes, because of where physical quantities are defined on the computation mesh, it is usually most natural to evaluate diffusion coefficients at cell centers. In contrast, calculation of radiative fluxes most naturally

takes place at cell interfaces. The need for a diffusion coefficient in evaluating a flux therefore requires an appropriate interpolation from among the neighboring cell-centered diffusion coefficients for use at a cell interface.

A number of interpolation treatments are in use in these heterogeneous situations The CRASH code, in its final version, evaluates the diffusion coefficient at an interface as an arithmetic average of the two adjacent cells. Although this is probably adequate for many situations, it clearly leads to an incorrect result in any situation where two neighboring cells have opacities that vary by several orders of magnitude. (An arithmetic average vastly overestimates radiation transport within the opaque region and leads to unphysical penetration of the opaque material and, potentially, spurious ablation—a potential issue at the interface between the tube interior and inner surface of the plastic tube.

As an ancillary activity to the PDT-CRASH code comparison study, we investigated the effect of this interface-implementation error. When replaced with a more accurate harmonic-mean of diffusion coefficients, which is a better approach that yields continuity of flux across an interface, we found that use of the harmonic mean eliminated the spurious diffusion and the incorrect front speed seen with arithmetic averaging. Thus, we recommend that in a multi-material HEDP simulation, treatment of interfaces deserves more attention than afforded by the CRASH project. Details of these findings are presented elsewhere [Myra and Hawkins, 2013].

5.1.3.2 Boundary conditions for radiation diffusion

Setting boundary conditions in radiation-diffusion codes presents interesting implementation issues. This is because boundary conditions need to be based on the physics of the problem but, in the case of the radiation-diffusion approximation, it is physics that cannot be adequately described within the formalism of the approximation.

As another ancillary activity to our code comparison study, we investigated the effect of CRASH's implementation of the radiative boundary conditions that permit no incoming radiative flux—the so-called vacuum conditions. The CRASH code, in its final version, evaluates the cosine of mean scattering angle of emergent radiation, $\langle\mu\rangle$, at each boundary in the pure diffusion limit (i.e., $\langle\mu\rangle = 1/2$). This is consistent with the familiar Mashak-Milne condition. We find, however, that in many problems in the HEDP regime—in particular, at the far end of the shock tube in the CRASH setup—radiation behaves in a free-streaming fashion and is better approximated as $\langle\mu\rangle = 1$. Although radiation effects at the far end of the shock tube would seem to play a minor role in the overall evolution of the system, we have demonstrated [Myra and Hawkins, 2013] that this is not necessarily true, and its effect is measurable.

5.1.4 Conclusions of the CRASH/PDT Study

Code-to-code comparison of the radiation-transport code, PDT, and the radiation-hydrodynamic-diffusion code, CRASH, proved challenging. Reasons for this included such issues as (i) the two codes being based on different underlying equations, resulting from different underlying approximations to the physical phenomena they track; (ii) test-problem design issues in yielding tractable problems for both codes using realistic computational resources; (iii) issues in choosing comparison metrics, given that each code evolves different physical quantities that cover solution spaces of differing dimensions, (iv) operational considerations, given that each code is developed and maintained at different institutions, and (v) operational constraints preventing implementation of a hydrodynamics module to incorporate into PDT.

The above difficulties notwithstanding, we found that CRASH and PDT produce good agreement in regimes where they should be expected to agree closely, i.e., in the diffusion and free-streaming limits. In problems that more closely

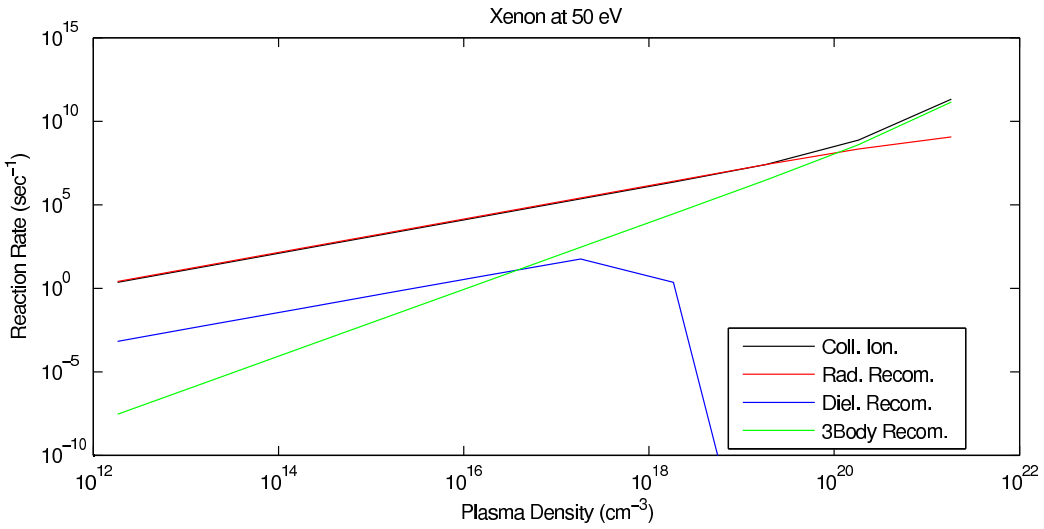

Figure 114: Reaction rates in xenon at $T_e = 50$ eV as functions of the atomic density, N_a. The LTE approximation is applicable if the collisional recombination dominates over photo-recombination, i.e. at $N_a > 10^{20}\text{cm}^{-3}$

represented the physical conditions found in the CRASH problem, we found the level of agreement conditional on the specifics of the test problem. However, we can fairly conclude that the variation in results that one would observe among different "reasonable" implementations of flux-limited diffusion are at least comparable to variations between radiation-diffusion, as implemented, and full transport.

A significant disappointment was our inability, within the resource and managerial constraints of the project, to make a comprehensive set of comparison tests in multiple spatial dimensions. Without such tests, it is difficult to conclude whether use of radiation transport (rather than diffusion) for a complete CRASH simulation would have a measurable effect on some phenomena important to the CRASH measureables (e.g., radiative wall ablation, wall-shock development, and its affect on the primary-shock evolution).

Finally, this portion of the project allowed a "double check" of the implementation of various details of radiation transport in the each code, which allowed us to uncover some shortcomings and bugs. Like the effect of full transport, the magnitude of shortcomings in the CRASH code on a complete CRASH simulation is uncertain. Although obvious in relevant test problems, the effects of these shortcomings may be washed out by other approximations elsewhere in the code. The solution-convergence results presented elsewhere in this document support this observation.

5.2 Assessment of Non-LTE effects

In calculating the equation-of-state functions as well as the opacities, we produced the tables used in the base CRASH code on the assumption that these quantities may be calculated assuming the radiation to be in thermal equilibrium with the plasma. Detailed balance occurs as the part of LTE, so that all processes involving the photons (such as photoionization and photorecombonation) balance each other separately of the electron-ion collisional processes (such as 3-body recombination and collisional ionization). Prof. G. Moses in the course of the annul review observed that the CRASH radiative shock system might not satisfy the validity condition for LTE. He provided Fig.114, which shows that, for Xe at an electron temperature $T_e = 50$ eV, the atomic density N_a must exceed 10^{20}cm^{-3} for LTE to be valid.

This would require only fourfold compression of our Xe gas, whose initial atomic density is $3 \cdot 10^{19}\mathrm{cm}^{-3}$. In addition, the threshold density for the validity of LTE drops rapidly as the temperature decreases. On this basis, we would expect that LTE would ve valid or very nearly so for the CRASH radiative shock.

In contrast, it is well known that laser-heated coronae are not in LTE. These regions are also optically thin, so the resulting radiation rarely has significant consequences. However, it is possible that the resulting difference in electron temperature might create some differences in ablation rate. In the context of the overall modeling by the CRASH code, these effects show up primarily in the Laser Energy Scale Factor necessary to reproduce the data. So long as these effects were consistent and one did not scale an application from one regime of laser irradiance to another, this would not be expected to affect one's ability to assess the predictive capability of the code.

It is also the case that the use of an LTE assumption in a code with a multigroup treatment of radiation is logically inconsistent. One introduces a multi-group description for the radiation energy density, with the energy density of group i being U^i_{rad}, and explicitly assume that these energy densities differ from the black-body-radiation energy densities, $B^i_{\mathrm{rad}}(T_e)$, integrated over the same photon energy groups. The ratios, $U^i_{\mathrm{rad}}/B^i_{\mathrm{rad}}$, may provide a quantitative estimate of the accuracy of the Local-Thermodynamic-Equilibrium (LTE) approximation. Under optically thick conditions, i.e. when

$$U^i_{\mathrm{rad}}/B^i_{\mathrm{rad}} \approx 1 \tag{235}$$

for the energy group carrying significant energy density, the LTE approximation works very well.

In the optically thin media Eq.(235) may not hold, so that the radiation multi-group energy densities may be much less than the Planckian energy density calculated at the electron temperature. Such situation is typical, say, for the solar corona or laser-heated corona. Ahead of the shock wave in the CRASH system the opposite inequality holds and $U^i_{\mathrm{rad}}/B^i_{\mathrm{rad}} > 1$, as has long been known to be expected for radiative shocks. The radiation "temperature" there much exceeds the electron temperature. Even when Eq.(235) does not hold and the system does not have detailed balance, the LTE approximation may be still applicable to equation-of-state calculations. Specifically, this is true at high densities. The reaction rates for the collisional processes are proportional to a second or third (for three-body recombination) power of density, while the NonLTE radiation recombination is linear in density and comparatively negligible at high densities. Therefore, the presence of non-equilibrium radiation at high densities may not affect the accuracy of the LTE approximation.

Motivated by the above considerations, and at the behest of our review committe, we proceeded to implement a model in CRASH to account for the non-LTE effects and assess their impact. Specifically, we merged the RADIOM code by Michel Busquet to the BATSRUS/CRASH code. The model implemented in the RADIOM reduces the non-LTE effects to the calculation of the effective temperature,

$$T_z = T_z(T_e, N_e, U^i_{\mathrm{rad}}/B^i_{\mathrm{rad}}) \tag{236}$$

This temperature is used to characterize the degree of the plasma ionization and electron excitation, hence, it can be used as an input parameter for the tabulated LTE equation of state. Particularly, the internal energy density, $E^{LTE}(T_z, N_a)$, may be calculated once the effective temperature, T_z is known. However, the kinetic energy of free electrons is still assumed to be characterized by the kinetic electron temperature, T_e, using the following correction equation:

$$E^{NLTE} - \frac{3}{2} k_B \rho Z^*(T_z)(T_e - T_z) = E^{LTE}(T_z, N_a) \tag{237}$$

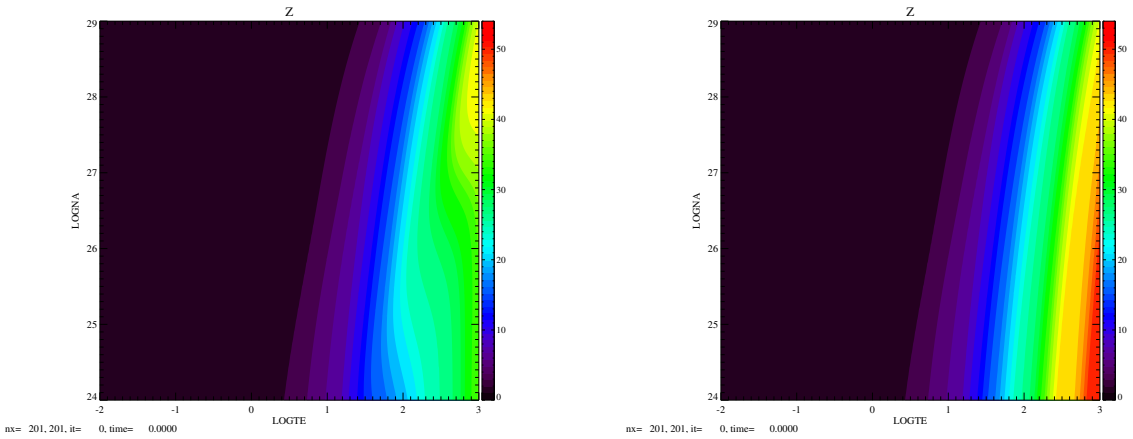

Figure 115: Non-LTE ionization degree, $Z^* = N_e/N_a$, in xenon (left panel). The radiation multi-group energy ratios, $U^i_{\rm rad}/B^i_{\rm rad}$ are assumed to be zero. Right panel presents the same plot for LTE equation of state.

In the left panel of Fig.115 we present, for Xe, the distribution of the ionization degree (the ratio of electron to atomic densities) as a function of $\log_{10}(T_e[\mathrm{eV}])$ and $\log_{10}(N_a[\mathrm{m}^{-3}])$ for the extreme limit of non-LTE effets, which corresponds $U^i_{\rm rad}/B^i_{\rm rad}(T_e) = 0$. The right panel shows the distribution obtained from the LTE EOS. As anticipated, for the conditions of the radiative-shock precursor ($\log_{10}(T_e[\mathrm{eV}]) < 1.8$ and $N_a[\mathrm{m}^{-3}] > 3\times 10^{25}$) there is no perceivable difference. This is also the case for the post-shock conditions with an electron density 10 to 10 times higher and an electron temperature no more than 0.15 keV. For hotter, less-dense Xe plasmas, toward the right edge and lower righr corner of the plot, the differences become larger.

In order to preperly implement the RADIOM model in CRASH, we had to pay special attention to the fact that our semi-implicit multi-group radiation diffusion solver approximates the differences of the internal energies densities in terms of the specific heat coefficients. Therefore, the EOS that includes NLTE effects should provide reliable values for the specific heat coefficient, in a thermodynamically consistent manner. The specific heat is the (partial) thermodynamic derivative of the internal energy density with respect to *temperature*, at constant mass density. However, while introducing an effective *temparature* T_z different from the real electron *temperature* the concept of the specific heat must be revisited. We have done this. As one can see from Eq.(237), the derivative with respect to temperature of the internal energy involves the derivative, $\partial T_z/\partial T_e$, which may be approximated as T_z/T_e with reasonable accuracy.

Although the thermodynamically consistent computation of the specific heat is an extra *requirement*, we can also *benefit* from its implementation. Indeed, in most occurences a code based on a conservative scheme should solve the *inverse* equation of state, that the unknown electron temperature should be calculated implicily from the given internal energy density. Once the calculation of the temparature derivative (the specific heat) is implemented consistently, the temperature can be solved using the reliable and rapidly converging Newton-Rapson procedure, which was impossible until we implemented the specific heat computation.

Enabled by this newly implemented Newton-Rapson-based algorithm our work was first focused on verifying that the coupled CRASH-RADIOM code reproduces the LTE results correctly. To do so, we artificially set $E_{\rm rad}/B = 1$ it the data that CRASH passed to RADIOM. We designate results from this model as 'fake-LTE' model and compared these results with those from the standard, LTE, CRASH code, as shown in Fig. 116.

Figure 116: Electron temperature distributions in the CRASH target at the time 0.9 ns. Top panel: the output from the standard version of the CRASH code (no NLTE effects). Bottom panel: the coupled CRASH-RADIOM code, with $E_{\rm rad}/B$ set to 1 in the data passed to RADIOM. There is a slight difference in color scale, reflecting the behavior in the laser corona discussed in the text.

Doing this comparison revealed an issue that we had not anticipated. To make it feasible to simulate the CRASH PY5 experiment, which is inherently 3D, it is important not to carry any more energy groups than needed. We developed present group structure to optimally resolve the structure of the STA Xe opacities using a computationally viable number of groups (30) and a physically sensible maximum energy (20 keV). The RADIOM code, however, is written to require the presence of energy groups up to 20 times T_e. In zones where this condition is not met, the values of $E_{\rm rad}/B$ are set to zero by RADIOM in groups for which no energy density is available. This approach is sensible for NLTE calculations, but it prevented the 'fake-LTE' test from producing LTE results in zones in the laser corona for which $T_e > 1$ keV.

Even so, the test was a success, as follows. We found that in zones for which $T_e < 1$ keV (where 1 keV=$\left(\frac{1}{20}\right) \cdot 20$ keV, 20 keV being the boundary of the radiation energy grid) the computation of T_z gives $T_z = T_e$ identically, as it should for LTE. This corresponds to the region that is not red (approximately to $x > 0$) in the figure. In the laser corona, the maximum temperature for for the 'fake-LTE' result in 3.3 keV, which is higher than that for the LTE run (2.9 keV). When accounting for this difference in maximum temperatures (accordingly, all contours in the bottom panel corresponds to somewhat higher values than those in the top panel) the results look the esame for the temperature range $T_e \leq 1$ keV.

Then we performed one of our standard tests (in $r-z$ geometry, with 3D distribution of the laser beams) with the time duration reduced to 1.3 ns. In Fig.116 the electron temparture distributions are presented, for the LTE version of the code (standard CRASH) in the top panel and the full NLTE result (coupled CRASH-RADIOM) in the bottom panel. We see that the account of the NLTE effect produces an observable shortening of the length of the precursor. This might potentially reduce the expansion that drives the wall shock. Thus there are two effects that might lead CRASH to overpredict the wall-shock separation from the wall – NLTE and the use of a diffusive radiation transport model. There is some evidence of this, discussed in Sec. 6.2.2. Despite this difference, examination of the shocked

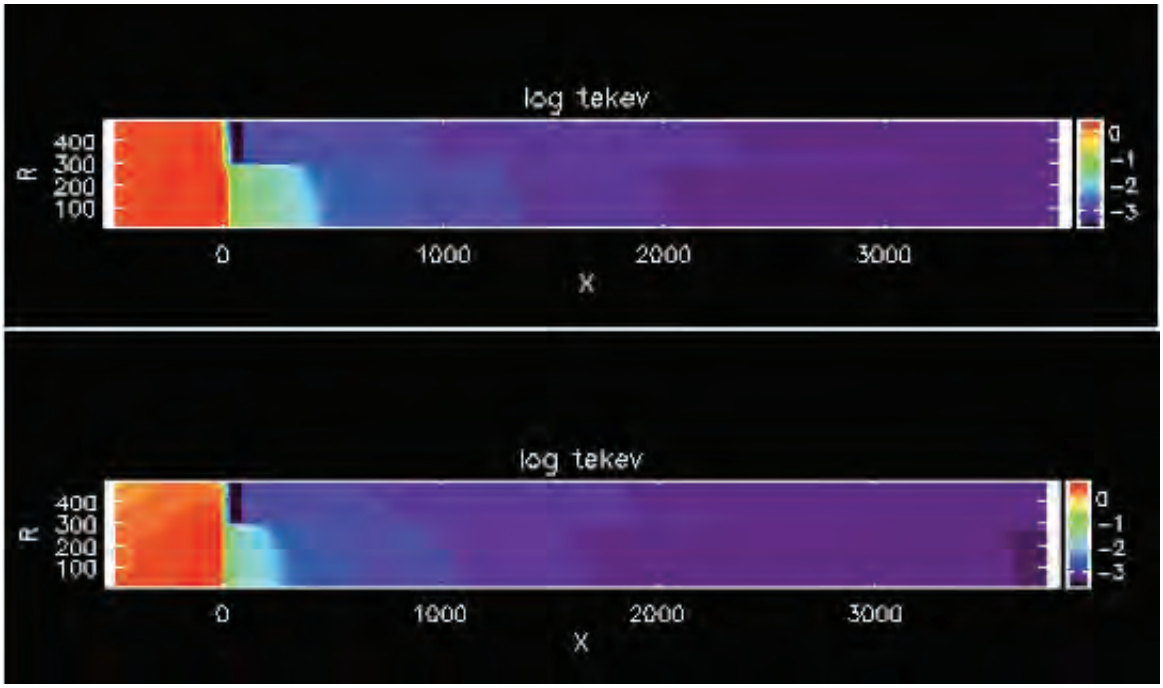

Figure 117: Electron temperature distributions in the CRASH target at the time 0.9 ns. Top panel: the output from the standard version of the CRASH code (no NLTE effects). Bottom panel: the coupled CRASH-RADIOM code (the NLTE effects are included). Note that the maximum temperature for the NLTE result is higher than that for the LTE run, therefore, the scales for the figures are somewhat different.

matter in the two simulations finds no structural differences. Since the simulations must use a Laser Energy Scale Factor to account for other unresolved physics anyway, our opinion is that NLTE effects did not introduce significant errors that would have impacted the predictions of the CRASH project.

Our results and observations from this study are as follows. The NLTE version of the code is slower by a factor ≥ 5. The EOS function necessary to implement NLTE effects within the CRASH structure is computationally intense. Using the NLTE package noticably degraded the efficiency of the preconditioner. NLTE effects appear to produce modest changes in the radiative precursor, but we do not believe that these had a significant effect on our predictions, at least at the level achieved in the course of the project.

6 Predictive Studies

Our predictive studies were performed in the context of the Bayesian framework established by Kennedy and O'Hagen [2001]. This ultimately involves the use of data from code runs, in conjunction with data from physical experiments, to construct a statistical predictive model. The code runs are performed in sets, each of which is designed to sample the probability distributions of some specific combination of experimental parameters, physical parameters, and numerical parameters. Examples of these are gas pressure, electron flux limiter, and number of levels of mesh refinement, respectively. We summarize these run sets next. We then discuss the extraction of metrics, which is necessary in order to use the computational or physical data for prediction. An example is shock location. These activities enable the statistical analysis of the run sets, some of which are sensitivity studies and some of which produce data for prediction.

6.1 Sets of Code Runs

In practice, one cannot sample the probability distributions of all possible variables in a set of code runs that can be accomplished in a reasonable time on available computer resources. One instead uses a combination of expert judgement and sensitivity studies to identify which variables have a significant effect on the metrics of interest. Several of our run sets were sensitivity studies, intended to inform the selection of variables for the run sets to be used for prediction. Beyond the run sets described here, we also did many smaller-scale assessments of sensitivity using a few runs to observe the sensitivity of the output to one or two input parameters. Other run sets were designed to enable analysis via the Kennedy-O'Hagan framework. Some turned out to serve both purposes.

6.1.1 Early Run Sets

The CRASH project organized sets of simulation runs designed for various uncertainty quantification, sensitivity, and scoping studies into groups of runs called run sets (abbreviation RS). Each run set consists of 10's to 100's of runs producing 10's to 1000's of values of quantities of interest (QOI), with each run differing by various input parameters. We did the following run sets:

1. Sensitivity study – 512 runs using 1D Hyades
2. Shock location study – 320 runs using 1D Hyades & CRASH 1.1
3. Shock breakout study – 1024 runs using 1D Hyades & CRASH 1.1
4. 2D shock structure study – 104 runs using H2D and CRASH 2.1
   A. 3D Sensitivity study – 64 runs using CRASH 1.1 from one H2D run
5. 1D convergence study – 512 runs 1D-Multigroup CRASH 2.1 from one H2D run
6. 2D convergence study – 128 2D-Multigroup CRASH 2.1 from one H2D run
7. 2D nozzle study with large tubes – 107 runs H2D & CRASH; Gray and Multigroup
8. 2D study of sensitivity to nozzle properties – 18 runs CRASH 3.0 using laser package
9. 3D study of sensitivity to ellipticity and shape – 10 runs H2D & 3D CRASH 3.0
10. 2D shock structure study for base experiment 120 runs CRASH 3.0 (replacing RS4)
11. 2D shock structure study for base experiment 120 runs CRASH 3.2 (replacing RS10)
12. 2D shock structure study for base experiment 120 runs CRASH 3.2 (replacing RS10)

13. 3D shock structure study for year-5 target – 3D-Gray 64 runs CRASH 3.2
14. 3D shock structure study for year-5 target – 3D-Multigroup (see discussion below)

We did four studies during our first two years, and analyzed the results in some depth as our initial work in prediction. The results are discussed in detail below. Following an initial sensitivity study (RS1), RS2 was integrated with our first predictive study (see Sec. 6.5). RS3 and RS4 were done in support of our combined analysis of shock breakout from the Be disk and later shock location. We undertook to examine shock breakout because of its central importance in determining the totol momentum available to the shock wave. RS 4 was a UQ study of 2D experimental variations for the basic CRASH radiative shock experiment. Five parameters were varied; laser energy, beryllium disk thickness, beryllium gamma, plastic wall opacity multiplier, and electron flux limiter. The 104 runs were initialized using H2D for the laser phase, then run with CRASH out to 26 ns. These results were used for the methodological development of a prediction that combined models, discussed below in Sec. 6.9.

Runset A was a 3D study in year 2 (see Sec. 6.6), intended to assess the sensitivity of the multidimensional behavior to parameters characterizing the equations of state and opacity. The study consisted of 64 simulations at each grid resolution varying four input parameters: the equation of state gamma for Be was varied between 1.4 and 1.66667; the gamma for Xe was varied between 1.1 and 1.4; and the opacity scale factors for both Be and Xe were varied independently between 0.7 and 1.3.

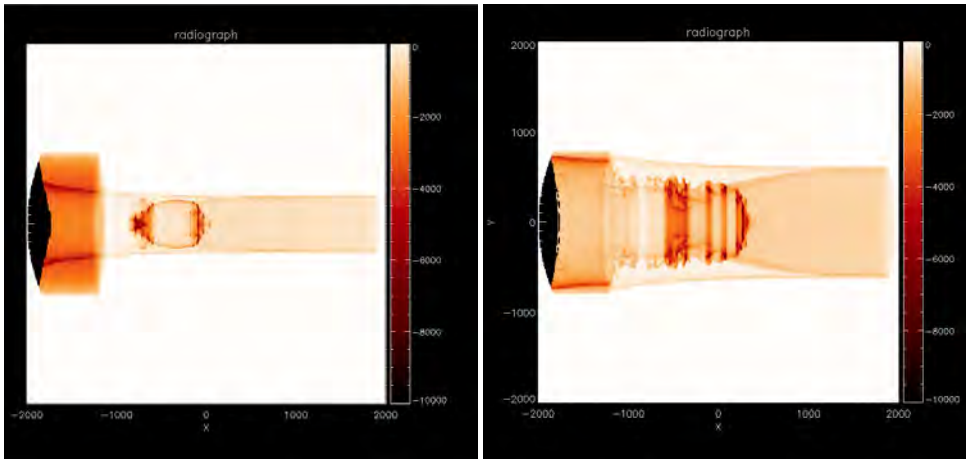

Figure 118: Radiographs of nominal tube radius (right) and wide tube radius (left) at 13 ns

In addition to runsets assessing the sensitivity of the behavior to experimental and theoretical variations, two run sets were performed to quantify the uncertainties of numerical parameters in the CRASH code. RS 5, initialized from output of a single 1D hyades run provided a systematic multi-parameter numerical convergence study by varying the number of photon groups, the minimum and maximum bounds of photon energy, the size of equation of state tables, and mesh resolution, in the subsequent 512 CRASH runs. RS 6, featuring 128 runs, was an extension of RS 5 in 2D and varied the number of photon groups, the minimum and maximum bounds of photon energy, the size of the equation of state tables, mesh resolution in two dimensions, and the convergence criterion for the Krylov solver. The output was analyzed in radiographic form using integrated metrics in 2D, and determining shock position in 1D. The integrated metrics are discussed below; they proved most useful when the fidelity of the simulations was relatively

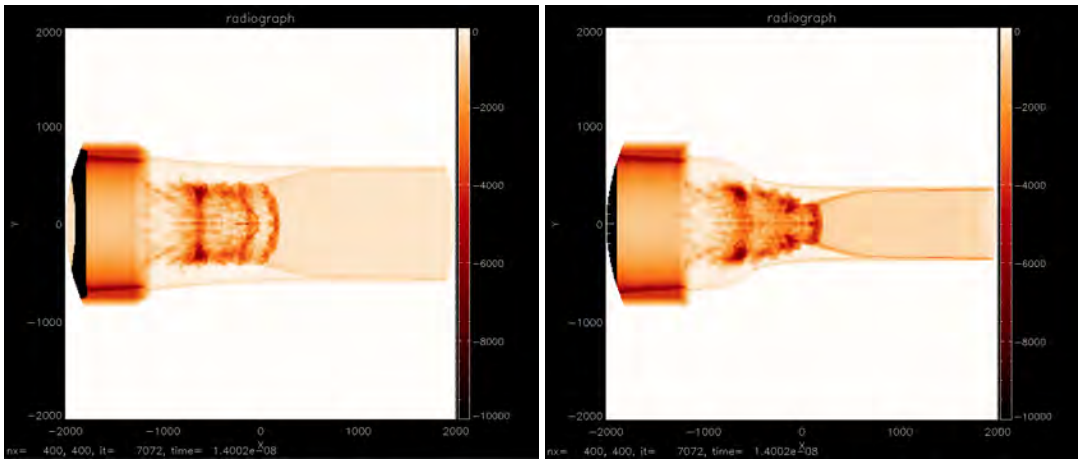

Figure 119: Cassini oval target pictured in orthogonal views at 14 ns.

poor. The UQ analysis based on a flexible regression fitting of the response surface determined the most efficient group structures for use in later run sets.

RS 7 focused on variations in experimental parameters in order to scope the Year 5 capstone experiment. It was 107 2D, H2D-initialized shock tube, performed using gray radiation and multi-group radiation, with a multi-group structure determined from RS 6. The radius of the tube was varied along with six other parameters: laser energy, laser scale factor, beryllium thickness, xenon fill pressure, electron flux limiter, and observation time. Results at 13 and 27 ns were analyzed using feature extractions from radiographs; samples are pictured in Figure 118. These wider H2D runs were used to initialize later run sets featuring cylindrical and elliptical nozzles.

To scope out the optimum nozzle shape for the Year 5 experiments, RS 8 was performed to investigate sensitivities to nozzle parameters. Initialized with the CRASH laser package, these 18 2D CRASH runs were performed with both gray and multi-group radiation. The nozzle length, post-nozzle tube diameter, and distance from the drive surface to the beginning of the nozzle were varied and the results were analyzed using feature extraction metrics of simulated radiographs.

Once the nozzle shape was determined with the help of RS 8, it became necessary to determine the effects of altering the ellipticity of the tube. RS 9 consisted of 10 runs initialized with H2D, and run in 3D CRASH with multi-group radiation. The only parameter varied in this run set was the value of the ellipticity in the range $[0.5, 1]$. Due to the methods of constructing elliptical tubes, both a regular ellipse and a Cassini oval were used in the modeling. Simulated radiographs, such as those shown in Figure 119, are then analyzed from orthogonal views to determine inherently 3D metrics to be compared to experimental data.

After the CRASH laser package was completed, a run set effectively replicating RS 4 but utilizing the CRASH laser package in place of H2D was performed. These 120 runs were run using 2D CRASH with 2D RZ laser ray tracing, in both gray radiation and multi-group radiation. The input parameters varied were laser energy, beryllium disk thickness, xenon fill pressure, electron flux limiter, laser scale factor, xenon opacity scale factor, and observation time. A scale factor on the xenon opacity replaced the scale factor on the plastic opacity after investigation of the sensitivities indicated that the xenon opacity was a more important factor. The simulated radiographs were analyzed, integrated met-

rics extracted, and combined with similar data from experimental radiographs using the Kennedy-O'Hagen predictive model form described above.

We then completed a version of run set 10 (RS 10), composed of 120 input points over a 6D input space. The input space covers Be thickness $[18, 22]\mu$m, laser energy $[3594, 4060]$ J, Xe fill pressures $[0.852, 1.46]$ atm, and observation times from the interval $[0.5, 27]$ ns in steps of 0.5 ns. In addition the inputs cover an electron flux limiter range of $[0.02, 0.1]$, a laser energy scale factor in $[0.2, 1.1]$ and a Xe opacity scale factor range $[0.1, 10]$. These runs were completed using the 2D CRASH laser package. A similar run set, RS 11, with improved Xe opacities and improved (3D) laser ray tracing was designed in order to observe the consequent improvement in the predictions. Comparison of RS 10 and 11 (now 12; see below) could be analyzed to assess the impact of the improvement in physics represented by the improved opacities and ray tracing.

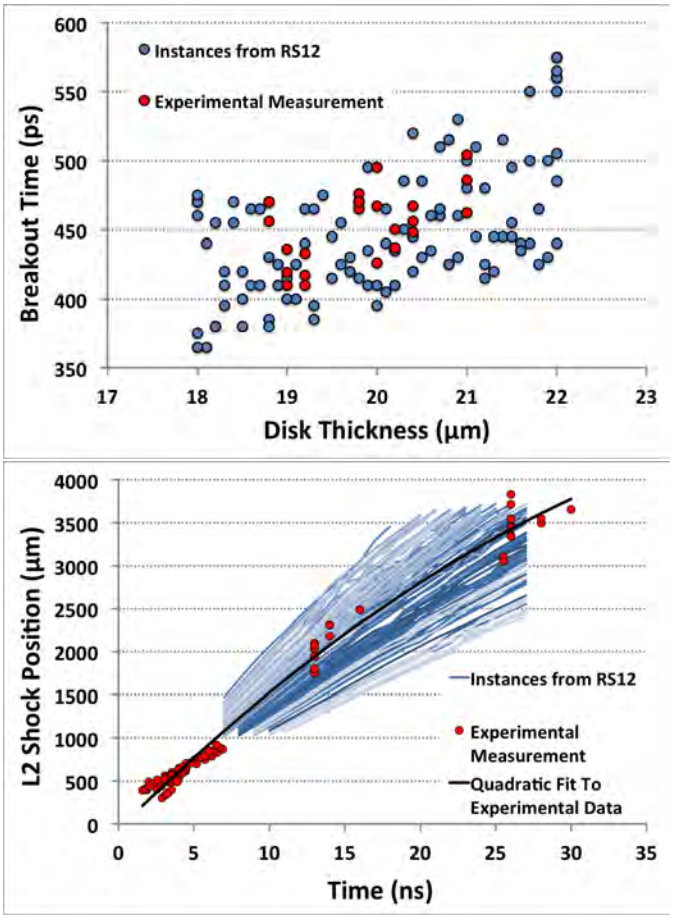

Figure 120: Results showing that RS12 spans the range of the experimental data.

With the further developments of the laser package—including the addition of 3D laser ray tracing as well as the acquisition of high-quality opacity tables for xenon and an improved multi-group structure, RS 11 was designed to replicate run sets 4 and 10 with the highest quality 2D CRASH model (v 3.2), but we soon realized that some issues with table ranges were compromising the results. To avoid confusion, after correcting the table ranges we restarted this case as RS12. Figure 120 shows that the runs in this set are spanning key QOI's from the physical data. The physical parameters for these 120 runs are the same as with RS 10, with tighter ranges on the calibration parameters of laser and opacity scale factors. Similarly to RS 4 and 10, RS 12 was then used to incorporate experimental data to make predictions of the QOI at 26 ns.

To formulate the prediction and prediction uncertainty of the Project Year 5 (PY5) experiment, we combined information from two simulation datasets run set 12 (RS12) and run set 13 (RS13) and data from experiments on the OMEGA laser facility. Experiments from October 2008 on OMEGA provided data of shock location and wall shock parameters at 13 ns for a radiative shock driven down a cylindrical, 575-μm inner diameter (ID) shock tube. Additionally shock location and wall shock parameters were measured in data taken in the same tube geometry at 27 ns in December 2010, as well as for cylindrical, 1500-μm ID tube data near 27 ns in October 2011. The simulation data to integrate into the predictive framework comes from RS12, 128 simulations varying parameters for the 575-μm ID ,

cylindrical tube, and RS13, 64 simulations varying parameters for the PY5 target geometry. Because these simulations involve handling large amounts of data, the extraction of the output data from the results is automated. The data extraction routines, described in Sec. 6.2.2, have been designed to act both on the simulated data and the experimental data, in order to be able to make direct comparisons. This information was combined using the Kennedy-OHagan framework utilizing Bayesian Gaussian process emulators previously discussed, for shock location. Follow-up work using the locations of the intersection of the wall shock and the primary shock would be possible.

The predictive process utilized experimental data and results from RS12 and RS13 to calibrate a model parameter, the laser energy scale factor (LESF), then use the calibrated distribution for the model parameter, as well as the information on experimental variability and uncertainty to predict the PY5 experiment with accompanying uncertainty estimates. The design for both RS12 and RS13 contain the LESF to account for uncertainty in the ability of the simulation to accurately model the coupling of laser energy to the shock system, at minimum because of laser-plasma energy coupling processes that are not modeled. This model parameter was calibrated using both RS12 and RS13 and the data in the geometry relevant to each run set. In RS12 there were 2 model parameters – LESF and electron flux limiter – and 4 experimental parameters – laser energy, drive disk thickness, gas density, and observation time. In RS13 there was 1 model parameter –LESF– and 7 experimental parameters – laser energy, gas density, nozzle length, taper length, aspect ratio, tube major diameter, and observation time. The deliberate lack of experimental data for the system in the (elliptical) RS13 geometry necessitates the inclusion of RS12 model information for calibration. RS12 also contained a second model parameter, the electron flux limiter; however sensitivity analysis revealed that the simulation results were insensitive to the electron flux limiter so it was set to a fixed value in RS13. (This parameter was important in 1D simulations, but the fact that it is not important in 2D or 3D is not surprising, because the laser energy penetrates less far in multi-D simulations with ray tracing and as a result the heat flux toward higher densities is reduced.) For making predictions of the PY5 experiment with quantified uncertainty bounds, the calibrated distribution for the LESF was combined with the distributions for the experimental parameters in the KOH framework to generate the predictive distributions of the system outputs.

Figure 121 shows the radiographs from RS12. Because the output data are obtained by automated routines, bad data from extraction routine error can emerge within the data set. Discussions between the code applications team and the statisticians were critical in identifying the outlier points in order to remove only data where the extraction routines had clearly failed. This also has led to discussions on the best methods for improving the extraction routines to minimize the loss of data. Much more progress in this respect would have been possible had the project been extended.

The 3D simulations for RS13 were challenging with respect to available computing resources, but proved feasible. Figure 122 shows radiographs along the two orthogonal diagnostic views. We evaluated a shock curvature metric, corresponding to the curvature near the radial axis, along the direction of the major axis of the elliptical tube. This was interesting but we did not have time to follow up on its ultimate utility. The runs were spread across two DOE machines, Mapache at LANL and Cab at LLNL, as well as one University of Michigan cluster, Flux.

6.1.2 Run Set 14

In order to directly compare the effect of different fidelity of radiation transport model on the features of interest in the radiative shock simulations, a subset of RS13 was carried out using the multigroup radiation transport. These runs were initialized from 2D, RZ geometry runs to model the laser portion of the experiment, with those results mapped

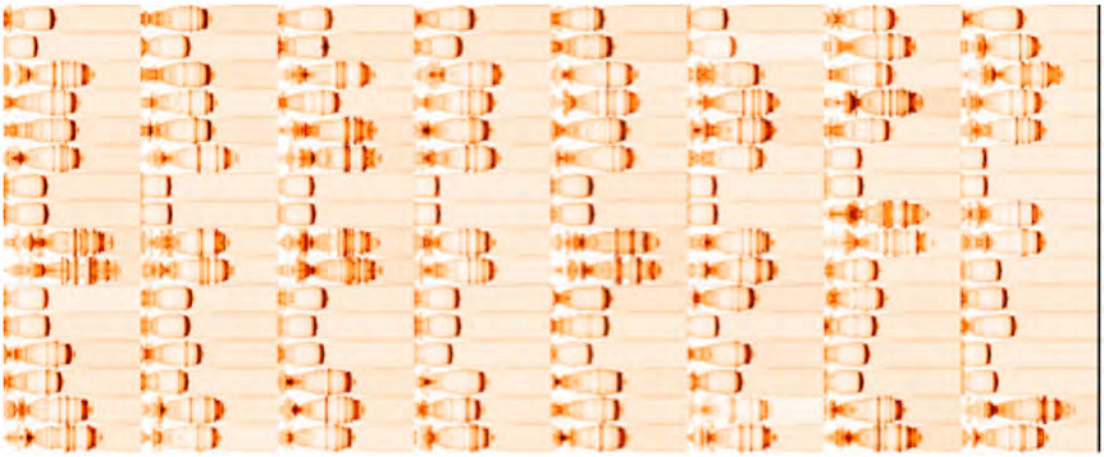

Figure 121: Tiled image of 128 simulated radiographs from the runs in RS12 at 16 ns. The array of images shows the large variations between the many runs.

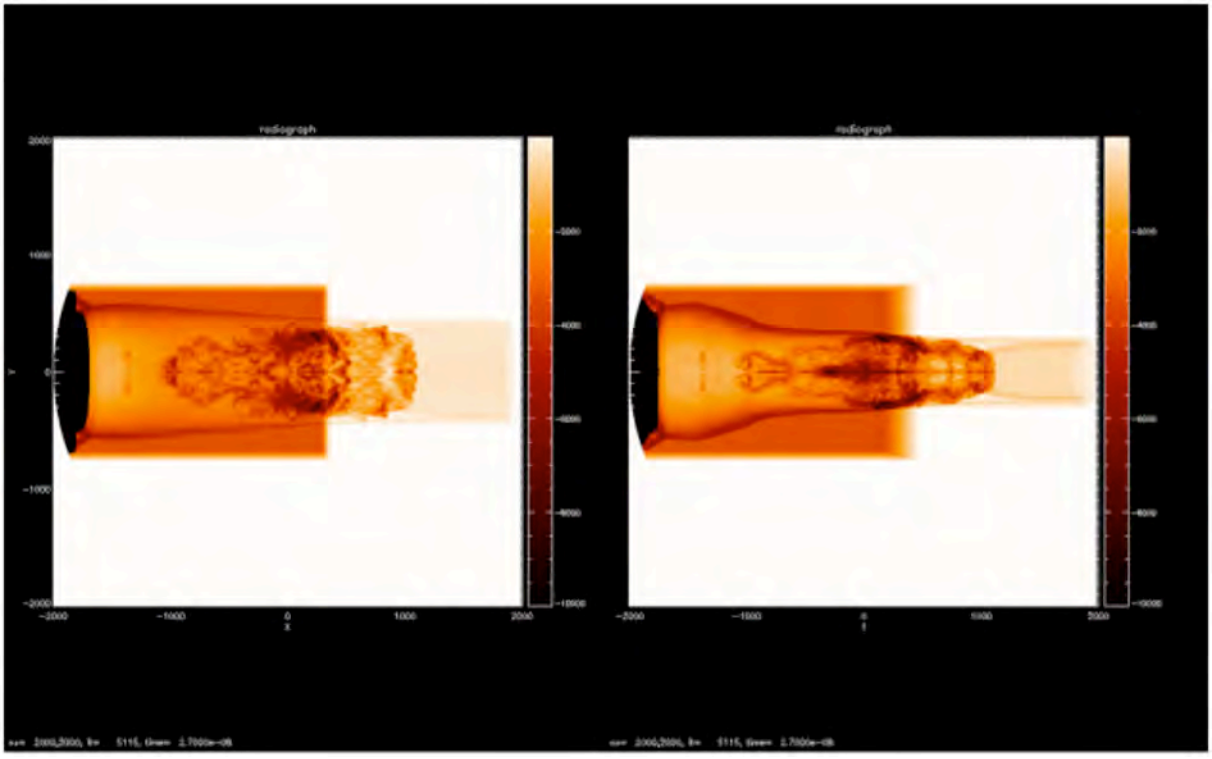

Figure 122: Example of a simulated radiograph from a run in RS13 at 27 ns. The image shows the simulated radiograph from the major and minor axis view of the elliptical tube experiment.

to a full 3D grid with 30 groups for multigroup radiation transport to model the remainder of the simulation.

RS13 runs used 1024 cores and required between 16 and 32 hours of wall time on the Cab supercomputing cluster

at LLNL to complete. This meant that each run needed only one or two passes through the queue in order to be finished. Due to the increased computational cost associated with 30-group radiation transport in RS14, despite expanding to 1536 core per run, the simulations were limited to advancing approximately 0.6-1.4 ns per 16 hour time block on Cab. This caused the time waiting in queue to become a large percentage of the overall time dedicated to the run set. Additionally, the many, large runs quickly used the projects allocation, causing files to wait much longer in the queue as each month went along. Because of these time constraints, only two of the RS14 runs finished and only 3 others getting past 18 nanoseconds (See table below).

Table 13: Output times reached in the most completed runs from RS14.

File Number	Last Output Time
1	18.0
3	27.0
5	25.5
7	27
17	15.5
34	12.5
36	15.0
37	12.0
59	22.5
64	15.0

Figures 123 through 125 below show a comparison between RS14 runs and their corresponding outputs in RS13. One would not expect the radiation transport to have an impact on the basic one-dimensional hydrodynamics, and this is confirmed in Fig. 123. Figures 124 and 125 show the vertical location of the triple point, or of the nearby shock-shock intersection, because the vertical location of the triple point is the output metric most sensitive to the details in the radiation transport. The figures show the vertical location in RS14 versus RS13, with the slope 1 line indicated where the points would fall if the values were equal. Figure 124 shows the values looking at the major axis of the target and Figure 125 shows the values for the view looking at the minor axis. As you can see, the vertical location of the triple point is notably greater in RS14, indicating that the triple point is less separated from the wall and thus that the wall shocks are weaker in the simulations with higher radiation transport fidelity.

6.2 Extraction of Metrics

To implement a statistical analysis that compares simulation output with the results of physical experiments, one must first identify something to compare. This is far from trivial. If the data are obtained in the form of images, a natural step is to implement software that emulates the function of the measuring instrument and produces an analagous image. One then seeks somehow to compare the images. This process is fraught with challenges. Idealistic discussions often

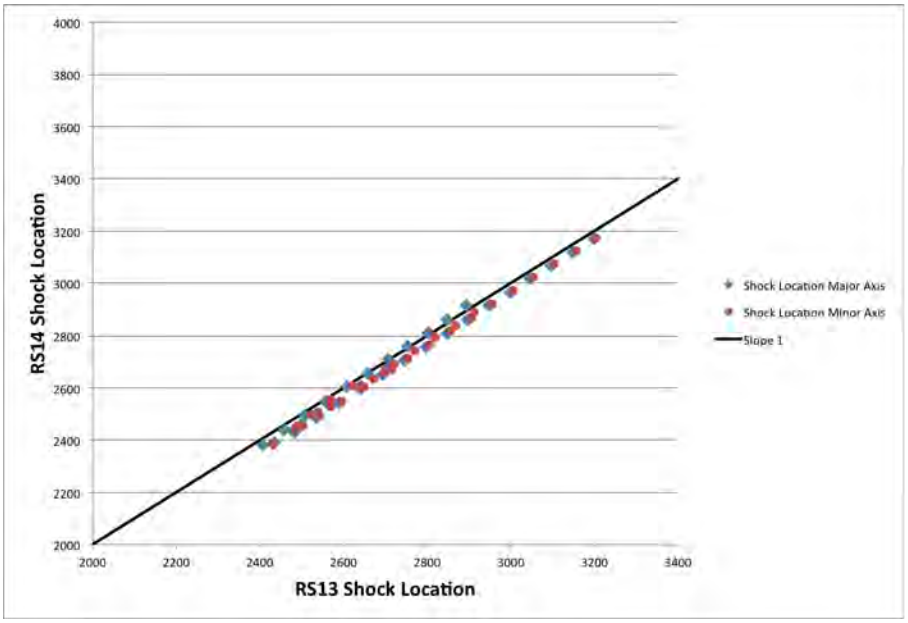

Figure 123: Comparison of axial shock location, in microns, between RS14, which used multigroup radiation transport, and RS13, which used gray radiation transport. Results extracted from the two simulated radiographic views are shown.

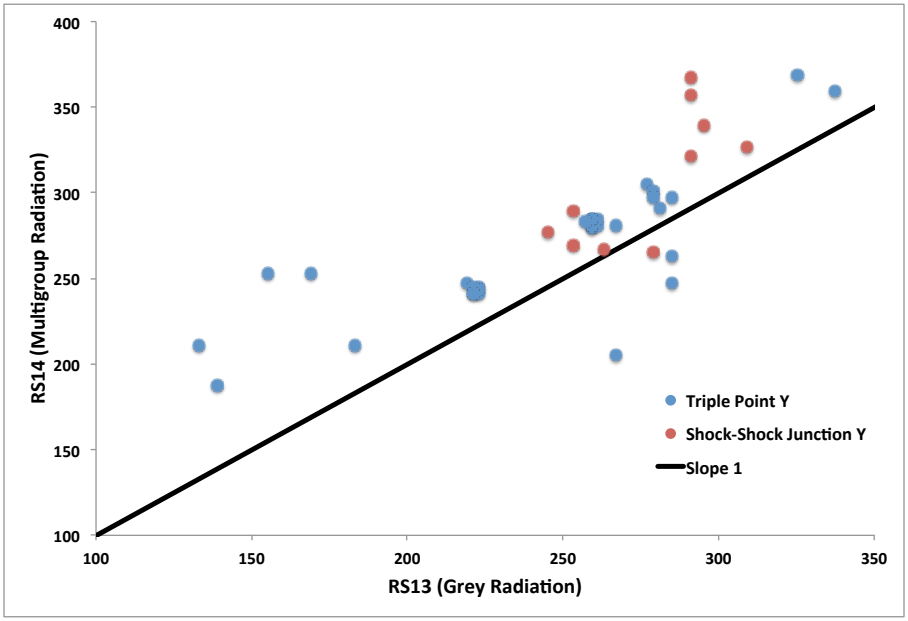

Figure 124: Comparison of triple-point and shock-shock axial positions, in microns, between RS14, which used multigroup radiation transport, and RS13, which used gray radiation transport. This shows results in the major axis (wide tube) view.

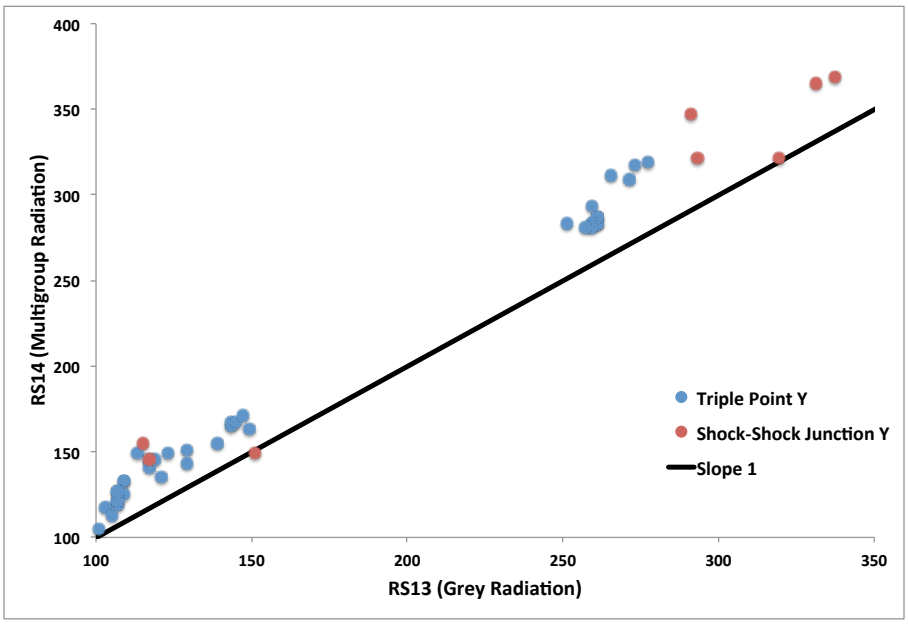

Figure 125: Comparison of triple-point and shock-shock axial positions, in microns, between RS14, which used multigroup radiation transport, and RS13, which used gray radiation transport. This shows results in the minor axis (narrow tube) view.

imagine that one might do some kind of analysis based on a pixel-by-pixel comparison, but our experience would suggest that this may only prove fruitful on systems whose behavior is so well understood and so well modeled that there is little point to doing predictive studies. Similarly, it would be ideal to apply an identical analysis to the raw data from the experiment and from the simulation, but this too may well be unproductive. For sparse physical data, experimental realities like calibration grids and scratches in film often pose few problems for manual extraction of metrics but large problems for automated approaches. On the computational side, developing algorithms that mimic the pattern recognition ability of the human eye and brain is notoriously difficult. This, in turn, increases the difficulty and decreases the accuracy of automated extraction of metrics from simulation data into which realistic noise has been introduced. This tends to lead one to make some compromises away from the idealistic approaches, because their full implementation would be risky, inordinately expensive, and ultimately out of scope for the project.

Beyond these considerations, the fidelity of the code output to the data from the physical experiment, in particular with regard to structure, has a major impact on which metrics will turn out to work. By design, one is varying some input parameters over ranges that exceed those present in reality. This is especially true of uncertain physical constants. If these variations produce variations in the output morphology, then metrics that depend on the morphology are likely to fail. In the early years of the present project, our challenge was even larger, as we were not yet getting any results that were morphologically consistent with the physical data. In such a context, one is drawn toward integrated metrics, and we made progress on using these during that period. Later, when we had consistently good morphology in the simulation output, we were finally able to return to metrics based on aspects of the structure present in the physical data. Ultimately, though, our predictive studies were primarily based on shock location, whose identification is fairly robust. We were in a position when the project ended to go further than this, and would have done so had funding

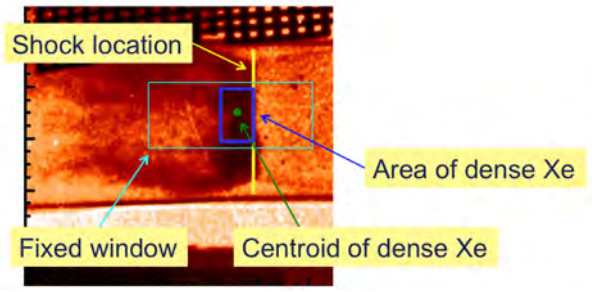

Figure 126: Illustration of some of the integrated metrics adopted in year 3.

continued.

6.2.1 Integrated Metrics

Our predictive work is predicated on extracting a few scalar parameters from both experimental data (in the form of radiographs) and from simulation data (also in the form of simulated radiographs). Because there is considerable variability in the structures in our system, both in physical experiments and in simulations, we developed a robust set of integrated metrics that are less sensitive to interface details. The integrated metrics provide information about the fundamental information given by the radiographs: how much dense xenon there is, where is it located, and how much is flowing near the edges. Within a fixed window (shown in Figure 126) we extract a set of metrics:

1. the (projected) area of dense Xe, defined as the area where the optical depth is larger than a given threshold times the optical depth of the unshocked Xe
2. the axial centroid of the dense Xe
3. the radial rms of the dense Xe over a window extended to the tube walls (not shown in the figure)
4. the breakpoints of a piecewise constant fit with 4 to 7 segments (the first such breakpoint is diagnostic of the shock location)

Shown in Figure 127 are the simulated radiographs from the first 64 runs of UQ Run Set 6, each paired with an image showing the location of pixels whose intensity exceeds a threshold value - a multiple of the absorption of the upstream xenon. The window extends from 1.3 mm to 2.5 mm in the axial direction and 0.1 mm from the centerline radially. The area of PATs is calculated by using the known pixel size in microns. The axial centroid is then calculated by: nx*x/N where nx is the number of PAT at a location x and N is the total PAT in the window. The process of finding PATs is repeated over a window spanning the entire tube diameter to calculate the radial variance of the PATs.

We use a separate program to fit a plot of the optical depth with a piecewise constant function and return the breakpoints of the best L1 (or L2) fit. The breakpoints give information about the shock location as well and thickness of the shocked layer can be extracted using this information. These fits are done using the same axial window as the area and centroid metrics. Robustness of the piecewise fit is established by varying the number of segments in the fit.

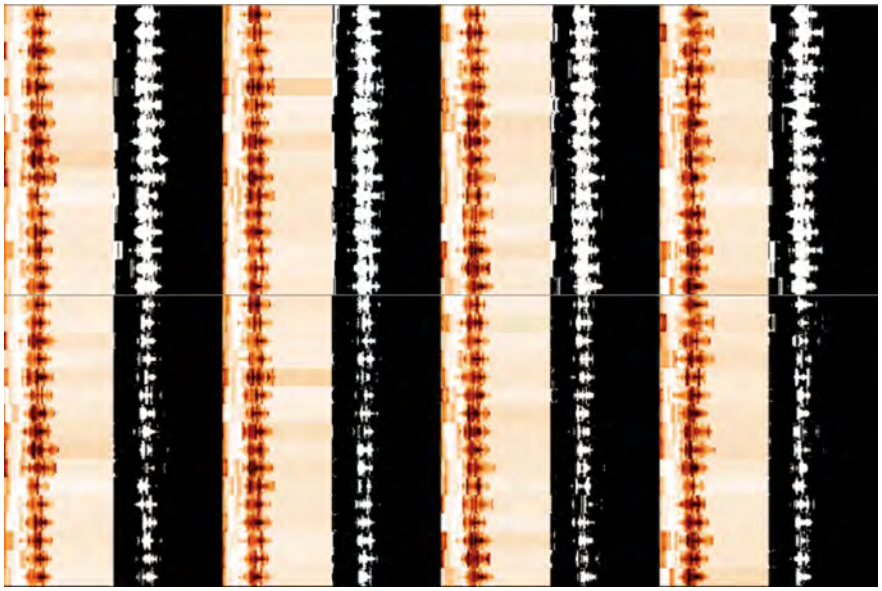

Figure 127: Output from the first 64 runs in RS6. The colored images show the simulated radiographs. The white images to their right show the location of the pixcels whose intensity exceeds a threshold value.

6.2.2 Improved Wall Shock Metrics

The main metrics used for assessment of the computational model and prediction in the CRASH radiative shock experiments have continued to be focused on the location of the main shock and properties of the radiation-driven wall shock upstream of the main shock. Methods for identifying the location of the shock in a robust, automated fashion were relatively straightforward to develop. Because the wall shock is sensitive to different aspects of the physical system than the main shock, such as the optical properties of the upstream materials and emission properties of the shocked xenon, quantification of the effects of its properties provides an important additional constraint on model physics and parametric uncertainty. Persistent effort has gone into designing and testing approaches to quantifying physically relevant, robustly identifiable aspects of the wall shock. We will review the early stages of these metrics, the issues that arose, and discuss the current wall shock metrics and methodology for extracting them. The speed of the wall shock as an indication of shock strength is a sensible measurable reflective of the important physical processes. Assuming the radiation flux from the main shock is adequately steady, the distance that the wall shock has separated from the wall when interacting with the main shock is a surrogate for the speed. Early results indicated that developing an automated algorithm for identifying the location that the wall shock met the main shock was not feasible. This led to the decision to define the first wall shock metric to be the separation of the wall shock from the wall at a distance of 100 m upstream of the main shock. Because this metric was a measurement in relation to the main shock location and the simulations showed large variations in morphology of the main shock inconsistent with the experimental systems, this metric varied wildly on ostensibly similar wall shocks and was determined to be a poor measure of the underlying dynamics. The decision was made to move to a wall shock metric that was independent of the shock location, with the focus on identifying two quantities, the slope of the wall shock on the image plane and the intersection of the wall shock with the wall of the target. These names are became somewhat misnomers, as both are defined by identifying

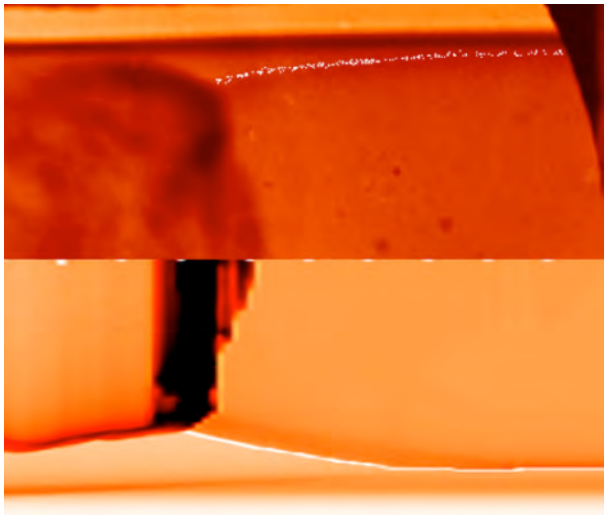

Figure 128: Experimental (top) and simulated (bottom) radiographs of radiative shock system. In both the top and bottom image, the locations of the interface between the plastic ablated from the walls and the xenon gas shocked by this ablated material are denoted in white. In the experimental image, the 200 points are identified manually, while the interface in the simulated radiograph is found by identifying the minimum absorption in the subset of a column of pixels. A piecewise linear fit to these points is used to define the wall shock slope and intersect point. The X-coordinates in the two images have been shifted to align the main shocks in order to best illustrate the wall shock properties.

the discontinuity in absorption marked by the interface between the ablated plastic from the wall material and the xenon that has been shocked by this ablation. The wall shock slope and intercept are then defined as parameters of a piecewise linear fit to the set of points marking this transition at multiple locations axially between the end of the tube and the location that the plastic-xenon interface meets the main shock. Figure 128 shows an example of the set of identified points on an experimental and a simulated radiograph, labeled in white. On the simulated image, the points appear as a nearly continuous line because one point is identified for each column of pixels. The details of the identification procedure for each type of image will be described below. The piecewise fit allows one breakpoint, such that one segment is fit to the points that have little to no separation from the tube wall, and one segment to the data that separate from the wall to intersect with the main shock. The parameters of this fit are chosen to minimize the residual sum of squares of the fit. The wall shock slope is then defined as the slope of the segment closest to the main shock and the wall shock intersect is defined as the location of the breakpoint. Because of the volume of images from simulated data, the identification of the two metrics is automated. The radial location of the interface near the end of the domain is found by first detecting the location of the wall by a parametric peak-fitting algorithm. (The wall is found by fitting a Gaussian peak added to an approximation to the background absorption profile to the radial absorption data.) The point of minimum absorption within 25 pixels (50 m) radially inward of the wall is then classified as the interface at that axial location. The next step in the algorithm is to step one pixel axially nearer to the shock and identify the minimum absorption within nine pixels radially of the previous interface location. This assumes that the wall shock in the simulations demonstrates sufficient smoothness across the image. The algorithm continues to step in the negative

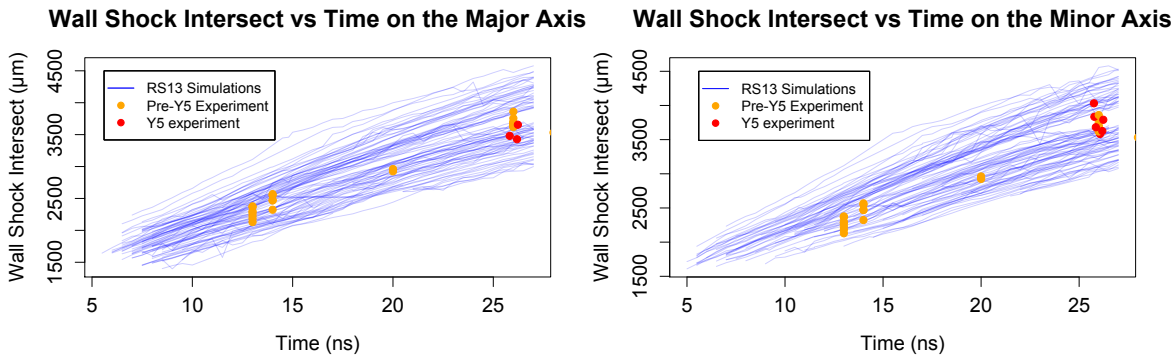

Figure 129: The values for the wall shock intersect location in RS13 on both the major and minor axis views of the elliptical tube. The experimental data from both pre-year 5 experiments and the year 5 experiments are shown. While the simulations seem to provide coverage of the experimental data, there is a gap in the values at late times due to problems fitting data near the end of the computational domain.

x-direction until the interface meets the shock. These points, with an example in the lower half of Figure 128, are then used to determine the wall shock slope and intercept. For the experimental data, the interface between the ablated

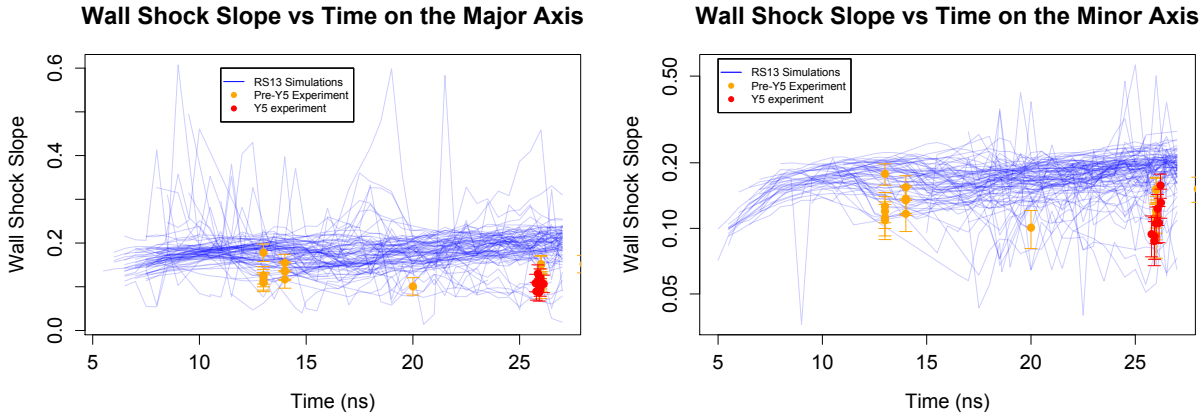

Figure 130: The values for the wall shock slope in RS13 on both the major and minor axis views of the elliptical tube. The experimental data from both pre-year 5 experiments and the year 5 experiments are shown. The simulation output shows significant noise due to inconsistencies in the fitting procedure. Additionally, late time slopes were biased toward higher values due to limitations in the computational domain. Work to improve the fitting algorithm led to the decision to change paths and go back to using the triple point as the wall shock metric.

plastic and shocked xenon is determined manually. The procedure for the simulated radiographs is not feasible due to the lower contrast and Poisson noise. Instead, a set of 200 points for each image was used to define the interface. The number of points chosen to minimize the impact of measurement error of any particular point and to ensure a stable fit of the piecewise linear model to determine the wall shock metrics. One example of the set of points used for the fit is

shown on the top of Figures128. In cases where the wall shock intercept extends beyond the viewable area of the film, the slope of a simple linear fit to the interface data is used for the wall shock slope and the interface value is omitted.

This algorithm worked adequately for the wide range of simulation results, but ran into a few issues related to the details of the fitting procedure. Especially at late times, the wall shock metrics became strongly sensitive to the length of the computational domain. As the number of points available for the flat portion of the piecewise fit were available on the domain became small, their weight on the piecewise fit became less, leading the wall shock intersection point to drift toward the main shock, biasing the output metric toward smaller values. Figure 129 shows the effect of this fitting problem as a gap in the wall shock intersect values at late times, with the values above the gap being from a subset of runs redone with a larger domain. Computational resource limitations prevented us from redoing all of the runs. As the intersection point moved toward the main shock, it shortened the axial component of the sloped portion of the piecewise fit, leading to steeper values for the slope. Additionally, the slope values showed a great deal of noise, especially at late times. The values for wall shock slope versus time can be seen in Figure 130. Despite these issues, the fitting mechanism worked well overall and efforts to improve the fit were promising; however, in testing improvements to those methods, it became clear that the algorithm being used to identify the points along the wall shock was capable of finding the triple point, the location where the wall shock meets the main shock. As this had been the original feature of interest related to the wall shocks, the decision was made to refocus on extracting it instead of extracting the slope and intercept values of the wall shock. Triple point locations extracted from the experimental data where possible are shown in Table 14

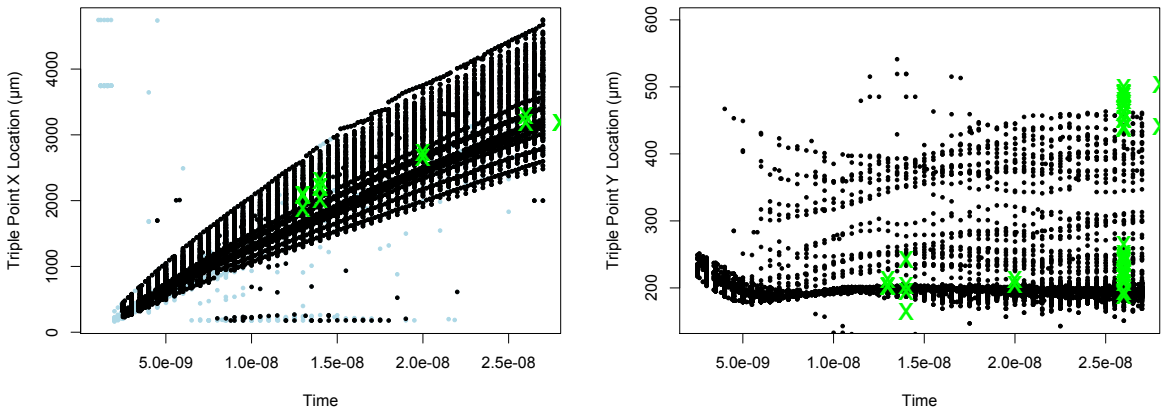

Figure 131: The values for the triple point locations in RS12 and RS13. The axial (x) location is shown on the right while the vertical displacement from the tube axis (y) is shown on the left. The experimental data from both pre-year 5 experiments and the year 5 experiments are shown with green crosses. The wall location for the baseline tube and narrow-view tube in year 5 is 287 μm while that for views of the wide tube is 575 μm. The simulations show good coverage in the horizontal values, though are biased to somewhat smaller vertical locations, indicating stronger wall shocks in the simulations.

The routine used to trace the points of the interface between the ablated plastic wall and the xenon in the wall shock in order to find the wall shock slope and intercept, as defined previously, needed only slight changes to be able

Shot number	Wall position (μm)	Horizontal position (μm)	Vertical position (μm)	Horizontal position (μm)	Vertical position (μm)
53665	$\pm$ 287	2192	195	2209	-203
56667	$\pm$ 287	2092	203	2096	-201
53668	$\pm$ 287			2308	-164
53671	$\pm$ 287	1882	212		
59027	$\pm$ 287	2654	205	2732	-213
59029	$\pm$ 287	3429	195	3454	-214
59029	$\pm$ 287	3487	216	3509	-223
63768	$\pm$ 287			3471	-208
63768	$\pm$ 287	3559	189		
63769	$\pm$ 575	3190	473		
63769	$\pm$ 575			3226	-514
63773	$\pm$ 575	3179	441	3185	-503
67701	$\pm$ 287	3802	227	3786	-239
67702	$\pm$ 287	3785	227	3675	-239
67702	$\pm$ 575	3593	480	3638	-461
67703	$\pm$ 287	3273	216		
67703	$\pm$ 575	2998	438	3011	-491
67704	$\pm$ 287	3532	223	3564	-246
67704	$\pm$ 575	3526	461	3600	-489
67706	$\pm$ 287	3512	227	3628	-264
67706	$\pm$ 575	3441	490	3426	-453
67706	$\pm$ 287	3553	246	3442	-236
67707	$\pm$ 575	3376	461	3376	-488
67710	$\pm$ 287	3622	221	3594	-248
67710	$\pm$ 575	3347	451	3429	-499
67711	$\pm$ 287	3655	243	3532	-250
67711	$\pm$ 575	3312	441	3345	-485

Table 14: Shot number and triple point location from multiple experiments. In some cases 2 triple point locations were measured from the top wall and the bottom wall. The uncertainty for the horizontal triple point location is $\pm$ 100 μm and $\pm$ 20 μm for vertical triple point location.

to trace back to the point where the wall shock intersected with the main shock. The procedure outlined earlier for finding the wall near the end of the tube was used, but at each step backward axially along the tube, the point of maximum absorption was found. While this did not perfectly trace along the wall shock front far from the main shock, the absorption profile of the wall shocked xenon near the main shock allowed the triple point to be identified with greater reliably. Figure 131 shows the horizontal and vertical locations of the triple point for run sets 12 and 13. The vertical location is the distance from the center of the target, for the same tube diameter, larger values correspond to wall shocks that have less separation from the wall. Because run set 13 varied the tube diameter, triple point vertical (Y) locations have a larger spread. The simulations show a bias toward smaller values for the vertical wall shock location, indicating stronger wall shocks than in the experimental system. This was consistent with expectation due to the near-grey radiation transport model. The simulations did show a wider spread in value than the simulations for similar tube diameter.

6.3 1D HYADES sensitivity study

Because CRASH was being initialized based on output from runs using two-dimensional HYADES (H2D), in order to have a calibrated initial condition modeling the laser irradiation phase of the experiment, it was important to assess the uncertainties associated with H2D to gain a full understanding of the uncertainties in CRASH. Lagrangian simulations in 2D can be very time and effort intensive due to mesh-tangling issues, which is one of several reasons why the first uncertainty study was done using 1D HYADES. The results of this study provided evidence of the importance of different parameters within the 1D code and were used to direct the future 2D study. A 15-D parameter space-filling Latin hypercube distribution was designed by Derek Bingham at Simon Fraser University to define a 512-run dataset for HYADES (RS1). The 15 parameters are as follows:

- Drive Laser Energy
- Drive Disk Thickness
- Gas Density
- Drive Pulse Duration
- Tube Length
- Laser Rise Time
- Slope of Laser Pulse
- Mesh Resolution
- Photon Group Resolution
- Electron Flux Limiter
- Time Step Multiplier
- Beryllium Opacity scale factor
- Beryllium Gamma
- Xenon Gamma
- Xenon Opacity scale factor

The parameter list encompasses both experimental parameters, such as the laser energy and the gas density, as well as physical or numerical code parameters, such as the xenon gamma or the mesh resolution. The experimental

Figure 132: An influence plot based on the study with 1D HYADES

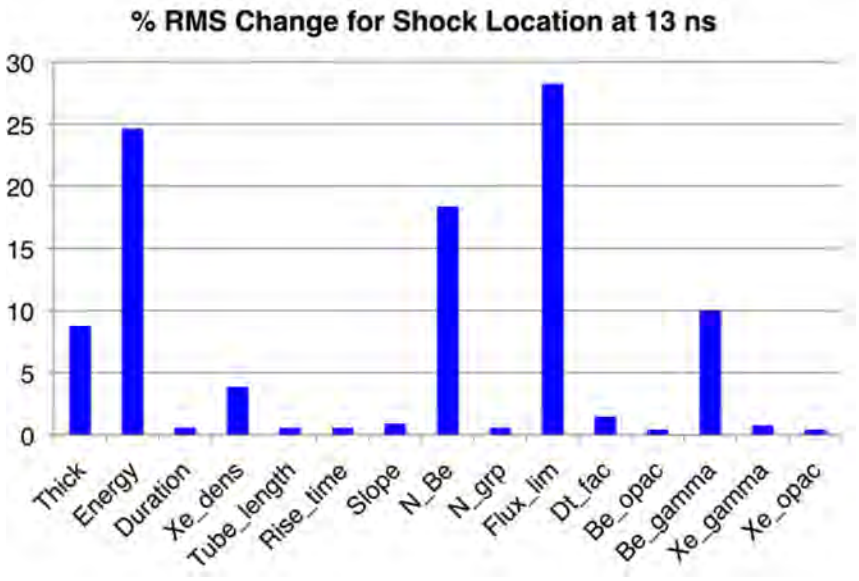

parameters were varied over a range defined by estimates of the variances from the experiments carried out at the Omega laser facility. The ranges of the code parameters were determined by careful analysis of sensible ranges for each variable. The results from the 1D uncertainty study were then used to further refine what we consider to be a sensible range of parameters for the H2D simulations. Before undertaking the full study, test runs were done to confirm the exclusion of some other parameters on the grounds that we did not believe they could have significant effects.

We evaluated global sensitivity by functional fitting with flexible regression methods (e.g. MARS and MART) followed by random permutations of each input and computation of average RMS change over such permutations. Using this technique we can determine which inputs have the most significant effect on the response surface. These results are plotted in influence plots to show those inputs that have the largest global influence on the outputs. Figure 132 shows an influence plot based on using 1D HYADES.

This identifies the Be disk thickness, gamma, and laser energy as important physical variables to the output, over the ranges investigated. Also notable is the number of zones (N_{Be}) used in the code; this set of data is based on a 512 point input design over a 15 dimensional input space. In using this set of data to construct the initial state for CRASH, we marginalize over only those values of N_{Be} that have no influence on the outputs (that is, we use sufficiently large N_{Be} that this mesh parameter has no influence on the results). The heat conduction flux limiter also stands tall as having a large influence. This then tells us that we should more closely investigate this parameter; subsequent review of the literature indicates that we should have used a more restricted range of values for the heat conduction flux limiter. In this way the sensitivity study makes apparent parameters that require more attention, and the UQ process drives the physics modeling and code development.

Another sensitivity metric comes from the length parameters from Gaussian Process (GP) fits of response surfaces. The covariance models in the GP model have the form

$$\prod_k \exp\left[-\frac{(x_{ik}-x_{jk})^2}{2\ell_k^2}\right], \tag{238}$$

Figure 133: Normalized significance of the 15 input parameters used in the HYADES run set.

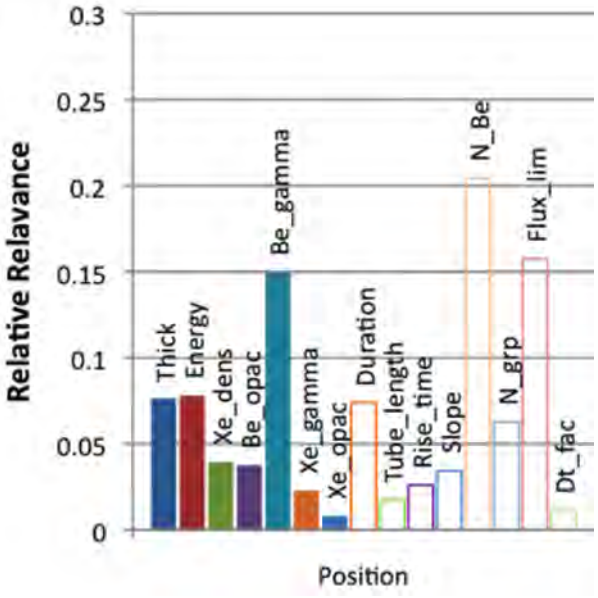

where ℓ_k is the length parameter along coordinate k. A large length scale implies that distant points are highly correlated. The relative relevance for input k is defined as $r_k = 1/\ell_k$; a small value for r_k means that large changes in variable x_k have little relevance to the output, while large values of r_k imply that a small change in x_k has a significant relevance to the output. While the influence plots describe large-scale influences of inputs on outputs, they operate over the whole input range and are not sensitive to more local structure. The relative relevance describes the scale over which a variable operates, and so provides information about variations that are more localized than the entire input range investigated.

Figure 133 shows the relative relevance of the 15 input parameters for shock location at 1.3 ns (at the time when CRASH was initialized). The input variables have been standardized over their ranges, so the relative relevance is normalized to the width of the input space along each dimension. In the system response at this early time more input variables are in play, besides those that have influence at 13 ns. In particular, besides the influential 5 variables seen before, pulse duration and number of groups (N_{grp}), all produce relatively rapid variation in output compared to the other 8 variables. These then become candidates for closer study. A study of the variation in N_{Be} marginalized over all other variables reveals a curve asymptoting to a shock location independent of N_{Be} (number of Be zones), as would be expected.

6.4 Development of a Physics Informed Emulator

The task of quantifying the uncertainty in the output from computer codes due to input uncertainties is important element of making informed predictions based on simulation. This task is especially challenging when there are a large number of uncertain input parameters or number of output quanitities of interest. Both of these characteristics are present when functional data from one computer code are used to initialize another computer code: the number of uncertain inputs to the second computer code is the large number of outputs from the first computer code. During

the first three years of this project we used results from the Hyades code to initialize our runs with CRASH. Here and in McClarren et al. [2011] we detail a particular approach to this aspect of uncertainty quantification that we believe could have wide applicability.

At CRASH we endeavor to simulate and predict the behavior of laser-driven shock waves traveling at high Mach number down a gas-filled tube. The temperatures reached in the experiment are high enough that energy carried by x-ray radiation affects the dynamics of the shock evolution. To simulate these shocks an Eulerian radiation hydrodynamics code, the CRASH code, has been developed based on a 3D, adaptive, massively parallel magnetohydrodynamics code [van der Holst et al., 2011]. Using the CRASH code we will predict experimentally observable quantities such as shock location as a function of time. To initialize the CRASH code we use a Lagrangian radiation-hydrodynamics code, Hyades [Larsen and Lane, 1994], to compute the laser energy deposition and early time shock formation. Therefore, we take the field data for the hydrodynamics and radiation variables from Hyades (hundreds of outputs) to initialize the CRASH simulation. As a result, to compute the sensitivity/uncertainty of the CRASH predictions for the result of an experiment, we need to known the sensitivity of the Hyades outputs to its uncertain inputs. If we had to compute sensitivities at each of the hundreds of Hyades outputs this task would be nearly hopeless.

Nevertheless, we have been able to reduce the number of outputs from a Hyades runs significantly using physical insight. The CRASH simulation is not sensitive to every detail of the Hyades output, and we have leveraged this fact to characterize the Hyades output using only 40 parameters. We can compute the sensitivity of these 40 parameters to quantify how much the uncertainties in the input to Hyades affect its output. As a result we have identified the areas of input space where we need to focus effort in reducing uncertainties.

6.4.1 Description of the Simulation

Using computer simulation we seek to predict the several features of the shock such as its position down the tube as a function of time and its speed. The paradigm we have adopted uses a Lagrangian radiation hydrodynamics code, Hyades, to compute the laser energy deposition and system evolution for the first 1.3 ns (the laser pulse width is 1 nanosecond at full width half maximum). The result of the Hyades computation is an initial condition for the CRASH code, an adapative mesh refinement (AMR) Eulerian radiation hydrodynamics code that computes the shock behavior in the xenon. This handoff from Hyades to CRASH is necessitated by the fact that, when modeling the system in multiple dimensions, the time a Lagrangian code can model is limited by tangling of the mesh. This mesh-tangling problem is severe enough that Hyades cannot be used to model the system of interest at later times. However, because CRASH cannot model the laser absorption, modeling at early times with Hyades is necessary.

As a result of this arrangement, we need to understand the sensitivity of the Hyades results as function of 1) experimental conditions such as laser irradiance, beryllium disk thickness, xenon gas pressure, etc. 2) numerical parameters such as the number of mesh points or the number of energy groups. Moreover, because we are interested in using the output from Hyades to initialize the main simulation code, we are presented with a large number of parameters (hundreds of mesh points times the number of hydrodynamic variables) for which we desire sensitivity information. In the next section we discuss our approach for reducing the number of parameters needed to characterized the Hyades output. Later we will describe the process of generating a regression model for the Hyades output.

6.4.2 Physics-Informed Partitioning of Hyades Output

One might hope to develop an analytic model of the structure of the Hyades output at a given time. For our application this time would be the time at which the Hyades output is used to initialize the CRASH code. This would offer the merit of having physically based parameters that could be adjusted to fit the Hyades results. Unfortunately, this turns out to be quite complex because the effects of the laser-heated electrons continue for several hundred picoseconds after the end of the laser irradiation. Alternatively, one might hope to do fitting of the Hyades output to identify a simple, parameterized, description of the structure. Our initial attempts to do this by purely statistical methods did not go well, in part because the statistical model has no physical judgement. Here we present and discuss a physically motivated partitioning of the Hyades 1D output. This makes the development of an emulator much more efficient because we need to predict the simulation output of many fewer degrees of freedom.

Figure 134 shows the velocity, density, and pressure profiles from a 1D Hyades run, at a time of 1.25 ns (the initialization is standardized at 1.3 ns). This run is tuned to match the observed location of the shock in the Xe gas at 13 ns. The smooth, dot-dash curves show the Hyades profiles while the solid lines show a linearly interpolated fit. First we discuss the origin of the observed structures. The laser irradiates the Be disk, first driving a shock wave through it. The shock wave breaks out of the rear of the Be disk in about 500 picoseconds, after which two things happen. First, the rarefaction ahead of the rear surface drives a shock into the Xe gas at just over 100 km/s. Second, a pressure gradient develops from the released rear material toward the denser material heated by electron heat transport from the laser absorption region. This pressure gradient accelerates the bulk of the Be, a process that can be simply modeled as rocket acceleration. By the decline of the laser irradiance from 1 to 1.1 ns, the bulk of the Be has reached the same velocity as the initial shock, about 120 km/s (this equals 120 μm/ns and 1.2 $\times 10^7$ cm/s).

Nevertheless, the end of the laser irradiation is not the end of the Be acceleration, because the pressure gradient that accelerates the Be remains present until the electrons cool by expansion and heat conduction. This cooling will in actual fact be much stronger in 2D, so that the 2D temporal structure will not reproduce in detail the 1D structure. Primarily as a result of this, we expect that less tuning will be required in 2D. As the electrons cool, the pressure profile develops the peaked structure seen in Figure 134. The Be to the right of the peak continues to be accelerated, the Be to the left of the peak is decelerated. Correspondingly the region to the left of the peak has very little impact on the subsequent dynamics. One would expect that there is no need to model this precisely, or even to include the material far enough from the peak, in order to accurately initiate the CRASH calculation. The pressure drops with time, so that by 2 ns the pressure profile has Flattened out and no longer accelerates the Be. Much later on, the structure evolves toward that of a blast wave in which the pressure accelerates material gradually away from the shock.

In the figure, one can see a region of reduced pressure gradient to the right of the peak. This is a remnant of the initial launching of the shock in the Xe at shock breakout. If one follows the line upward from the left boundary of this Flat region, one can see that this corresponds to the location of maximum velocity. What has happened is that the further acceleration of the Be after 1 ns has launched a velocity impulse forward through the leading edge of the Be. The maximum velocity is about 180 km/s at 1.25 ns, and in this calculation reaches 220 km/s by 1.5 ns when the velocity impulse has overtaken the shock in the Xe. This corresponds to the maximum post-shock ion temperature found by Hyades, which is about 2 keV. At 1.25 ns, the shock is established in the Xe but is evident in the figure primarily in the density. At this time, the immediate postshock ion temperature is about 700 eV. The structure in the Xe is not well resolved at this time, with both the pressure and the velocity showing gradual transitions. This is

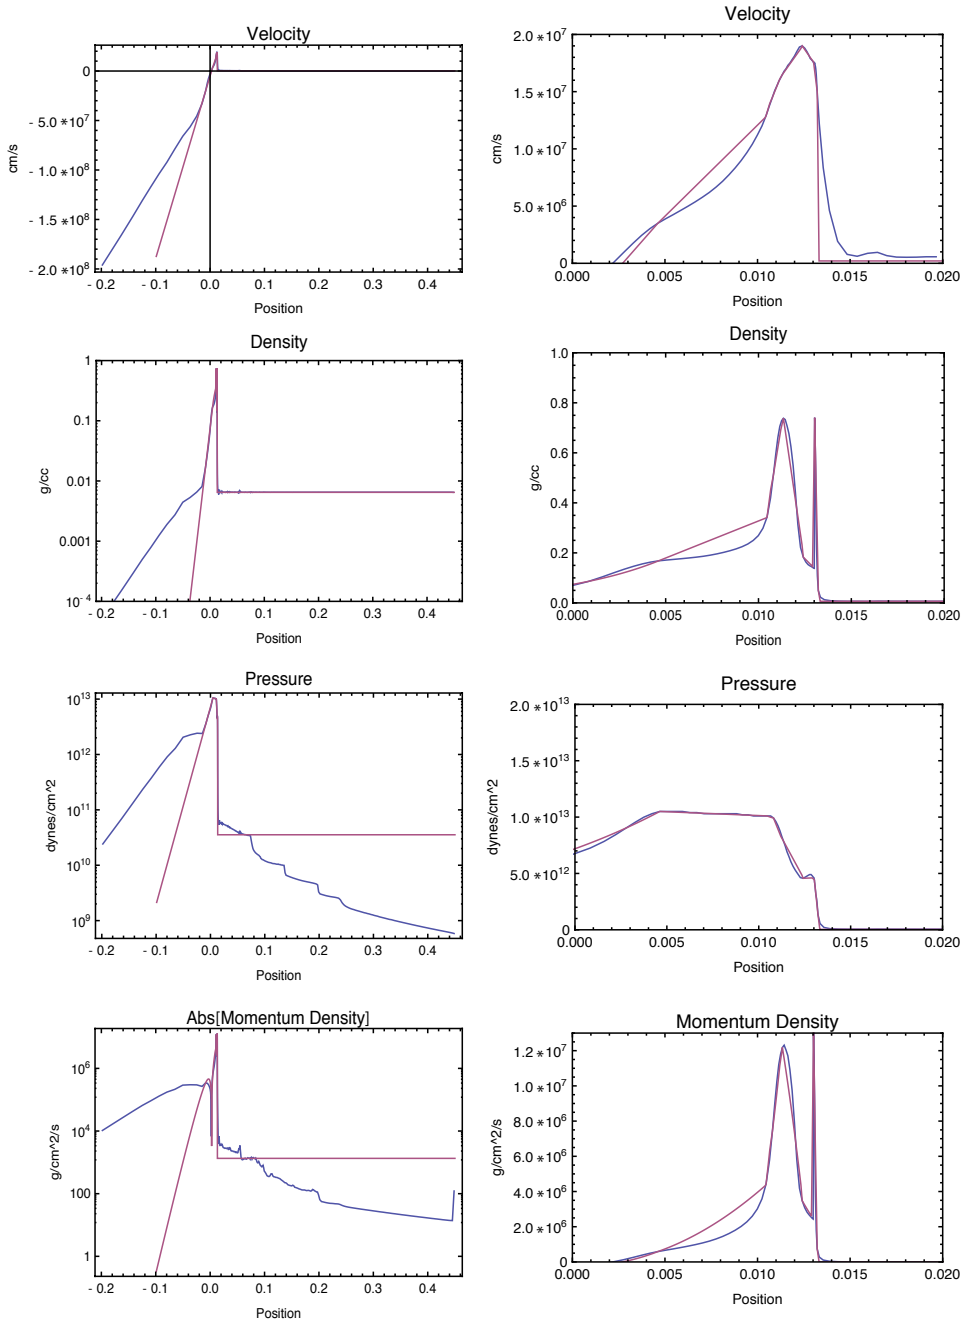

(a) Hyades Results

(b) Detail

Figure 134: Hyades output at 1.25 ns compared with piecewise linear fit from the physics informed emulator.

unavoidable in the context of doing a viable 1D Lagrangian model.

The density maximum is located in between the maxima of pressure and velocity, and is the natural result of the ablation of Be to the left and the expansion of the rear Be surface to the right. Once the velocity maximum has overtaken the shock, the velocity profile becomes and stays quite linear. This is typical of freely evolving hydrodynamic systems, and corresponds to steady expansion with time. As the system expands, the density decreases and the density maximum eventually disappears.

Beyond these physical considerations, in the process of making this fit work we were led to include some additional parameters further to the left, with an exponential fit to the density and pressure there. In the long run, after we had sorted out all the details, this turned out not to be essential.

With the above context, the following assumptions seem reasonable for construction of a fit to the Hyades output. 1) Material having significant negative velocity can be approximated by very simple and inaccurate profiles, because it just continues to slowly accelerate to the left and cannot impact the dynamics of interest. The fit shown in the figure and described here, though, ignores the laser-heated corona at low density and approximately captures the exponential density profile from the material that was heated by electron heat transport. 2) The initial state of the precursor does not need to be modeled in detail, because the energy through the shock by 1.5 ns is less than 10% of the energy by 13 ns. We demonstrated this in test runs with CRASH that removed the initial precursor heating. 3) The following parameters will suffice to define the profiles.

First, the minimum position of the Hyades output is needed. The fit shown uses half this value as its limit, because as one can see the density and pressure profiles fall off much more steeply (due to ignoring the laser heated corona). Call this value $r_{\min}$. Second, one needs the position, velocity, density, and pressure at the locations where (from left to right):

- the velocity first exceeds -3×10^7;
- the velocity is 1/2 the maximum value;
- the derivative of the pressure abruptly decreases;
- the derivative of the pressure abruptly decreases again (becoming negative or more negative);
- the density is maximum;
- the velocity is maximum;
- the interface is, from the Be side;
- the interface is, from the Xe side;
- the shock is located;
- the precursor properities are steady.

The data just described need to be ordered so that the position monotonically increases. The fit then is piecewise linear in all regions, beginning at $r_{\min}$, except for the density and pressure, which are fit as linear exponentials between the first two locations from this list and left from there to rmin . There is no need to fit the details of the shock

Table 15: Parameters from 1D Hyades (cgs units)

	Position	Velocity	Density	Pressure	Material
$u > -3 \times 10^7$	-0.008363001	-2.04E+07	0.016829163	3.62E+12	Be
p left corner	0.004615684	3.50E+06	0.168068666	1.05E+13	Be
ρ half max	0.010460472	1.29E+07	0.342857445	1.01E+13	Be
p rt corner	0.010801793	1.45E+07	0.510666015	9.92E+12	Be
ρ max	0.011342457	1.65E+07	0.738250317	7.88E+12	Be
u max	0.012412547	1.90E+07	0.182946	4.57E+12	Be
Be at interface	0.013004106	1.76E+07	0.138409625	4.61E+12	Be
Xe at interface	0.01302491	1.76E+07	0.740472601	4.43E+12	Xe
shock	0.013222783	1.50E+07	0.049271551	1.35E+12	Xe
Precursor	0.063511755	2.07E+05	0.006476877	3.55E+10	Xe

structure, both because they are unimportant for the long-term dynamics and because they are incorrect as represented by Hyades. The piecewise linear fit shown in Fig. 1 used the list of values shown in Table 2. One could calibrate the fit to preserve some defined quantity of mass, momentum, or energy, but at the present level of detail this would be overkill.

It is hard to see how fewer locations might adequately represent the physical system, except that one might drop the one with the most negative velocity. The upshot here is that the profiles are minimally represented by four parameters (plus material identification) at 10 locations. The total of 40 parameters include some that might be inferred from correlations, but not many.

To verify that these 40 parameters are adequate to describe the Hyades data used to initialize CRASH, we performed simulations of the shock evolution using CRASH with the nominal configuration (see Section 6.4.4) using the full Hyades output and the 40 parameters of the physics motivated partitioning of the data. We colloquially refer to this partitioning, especially when in the context of an emulator, as the physics informed emulator (PIE). In Figure 135 we show the results of these simulations. Here we see that the different initializations does affect the simulation output. Nevertheless, the change in shock position is less than our experimental uncertainty. Other features in the solution have changed slightly: the behind the shock the velocity, pressure, and density have a slightly different shape between 0.001 and 0.002 cm.

6.4.3 Models

We consider a regression model for the responses of the outputs Y_i, $i = 1, \ldots, n$, from n runs of simulations or experiments on the predictors or the inputs $\mathbf{X}_i = (X_{i1}, \ldots, X_{im})$ such that

$$Y_i = f(\mathbf{X}_i) + \varepsilon_i, \quad \mathbf{X}_i \in D \subset \mathbb{R}^m \tag{239}$$

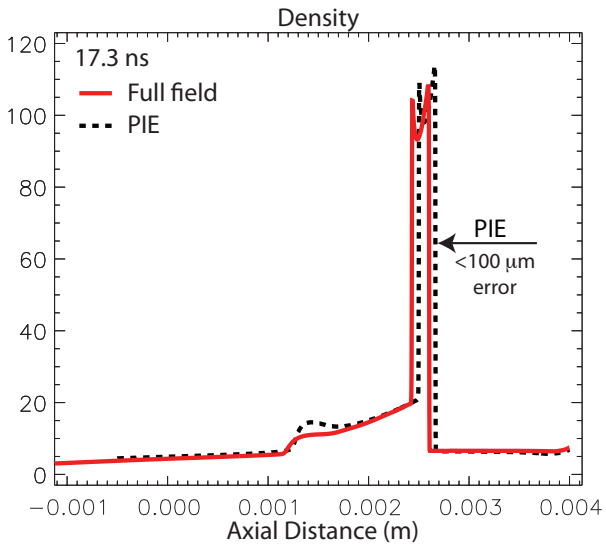

Figure 135: The density at 17.3 ns after initiation of the laser pulse as computed with two different initializations: using the 1D Hyades output directly (full field) and using the physical fields as estimated using the PIE. At this late time the difference in shock location is only about 75 μm off the prediction based on the full Hyades output field; this is comparable to the experimental uncertainties in locating the shock.

where f is an unknown regression function that we wish to estimate, ε_i is a random error with zero mean, mostly assumed to be from a Gaussian distribution, and D is the domain of interest, e.g., convex hull defined by the predictors.

In our case, because we are dealing with computer simulation and not performing a measurement, the value of ε is zero. Also, to predict the 40 PIE parameters for a given run of Hyades, each of the 40 parameters is a Y_i that we will model independent of the other 39 parameters. Modeling these parameters as being related is the topic of ongoing research, including the use of seemingly unrelated regression models [Holmes et al., 2002]. In priniciple we should be able make use of the known relations between the parameters. For instance, the parameters are ordered so that the position is an increasing function, and the other values of density are all less than the density maximum.

The two approaches used to construct a regression model are detailed in the remainder of this section.

6.4.3.1 Multivariate adaptive regression splines (MARS)

Multivariate adaptive regression splines (MARS) as proposed by Friedman [1991] is a nonparametric regression. With k basis functions B_i and suitable coefficients a_i, $i = 1,\ldots,k$, the MARS estimates f such that

$$\widehat{f}(\mathbf{x}) = \sum_{i=1}^{k} a_i B_i(\mathbf{x}), \tag{240}$$

where $\mathbf{x} \in D$. The basis function B_i with the degree of the interaction J_i consists of the sign indicators $s_{ij} = \pm 1$, knot points t_{ij}, and the order r_i

$$B_i(\mathbf{x}) = \begin{cases} 1, & i = 1, \\ \prod_{j=1}^{J_i} [s_{ij}(x_{v(ij)} - t_{ij})^{r_i}]_+, & i = 2, 3, \ldots, \end{cases} \tag{241}$$

where $(\cdot)_+ = max(0, \cdot)$ and $v(i, j)$ gives the index of the predictor variable split on t_{ij}. The optimum basis functions including knot points can be achieved by the generalized cross-validation criterion. The Bayesian MARS (BMARS) suggested by Denison et al. [1998] assigns prior to every unknown parameter in the model. The sign indicators s_{ij} and order r_i are assumed from uniform on the set $\{-1, 1\}$ and $\{0, 1, \ldots, R\}$ respective, for the maximum order R. The interaction terms J_i and the components of interaction effects are also uniformly selected. For example, if two-way interaction $J_i = 2$ is selected with predictors $\mathbf{X}_i = (X_{i1}, X_{i2}, X_{i3})$, then the interaction effect is equally likely to be one of $\{X_{i1}X_{i2}, X_{i1}X_{i3}, X_{i2}X_{i3}\}$. The number of knots and their locations are assumed to be uniform, respectively, on positive integers and on the data points. Conjugate priors are assigned to other nuisance parameters. By using a class of reversible jump Metropolis–Hastings algorithms for Markov chain Monte Carlo (MCMC), the BMARS collects samples of parameters from their joint posterior distribution. The advantage of BMARS is to identify significant main effects and interaction effects. In addition, the distribution of knots of each predictor reveals the complexity of the relationship between each predictor and the response.

6.4.3.2 Gaussian Process Regression

Gaussian process regression generates a Gaussian distribution of functions that attempts to interpolate the output data. Specifically, the Gaussian process is a collection of random variables where any finite subset of the random variables has a joint Gaussian distribution. The random variables for Gaussian process regression is the value of $f(\mathbf{X}_i)$ at a given point $\mathbf{X}_i$. In our notation, we write the value of f at all the points $\mathbf{X}_i$ as $\mathbf{f}(\mathbf{X})$. Like a Gaussian distribution, a Gaussian process is entirely determined by its mean and covariance. The data are normally standardized to have a mean of zero, and the covariance is determined by a covariance function that is chosen to have certain properties. In our case we use the squared exponential covariance function which assures that the function $f(\mathbf{X}_i)$ is smooth. This covariance function has the form

$$k(\mathbf{X}_i, \mathbf{X}_j) = \sigma_f^2 \exp\left\{ -\sum_{k=1}^{m} \frac{(x_{ik} - x_{jk})^2}{2l_k^2} \right\}.$$

The covariance function has $m+1$ parameters: the maximum allowable covariance (σ_f^2), and the m length parameters. These parameters are estimated using the empirical Bayes procedure. Using the covariance function we construct a covariance matrix $\mathbf{K}$ of size $n \times n$ with elements

$$K_{ij} = k(\mathbf{X}_i, \mathbf{X}_j),$$

and then write the Gaussian process regression distribution as

$$\mathbf{f}(\mathbf{X}) \sim \mathscr{GP}(0, \mathbf{K}).$$

Then to predict $\mathbf{f}(\mathbf{X}^*)$ at some number n^* input points, we use the fact that $f(\mathbf{X})$ and $f(\mathbf{X}^*)$ are distributed by a joint Gaussian. Then the conditional expected value of $\mathbf{f}(\mathbf{X}^*)$ is

$$\mathrm{E}[\mathbf{f}(\mathbf{X}^*)|\mathbf{f}(\mathbf{X}), \mathbf{X}^*, \mathbf{X}] = \mathbf{K}(\mathbf{X}^*, \mathbf{X})\mathbf{K}(\mathbf{X}, \mathbf{X})^{-1}\mathbf{f}(\mathbf{X}).$$

The covariance for $\mathbf{f}(\mathbf{X}^*)$ is given by

$$\text{var}[\mathbf{f}(\mathbf{X}^*)|\mathbf{f}(\mathbf{X}),\mathbf{X}^*,\mathbf{X}] = \mathbf{K}(\mathbf{X}^*,\mathbf{X}^*) - \mathbf{K}(\mathbf{X}^*,\mathbf{X})\mathbf{K}(\mathbf{X},\mathbf{X})^{-1}\mathbf{K}(\mathbf{X},\mathbf{X}^*).$$

Therefore, in Gaussian process regression we have a value for the mean and covariance at each point we wish to evaluate the regression model.

6.4.3.3 Comparison of Models

The two approaches, Gaussian process regression and Bayesian MARS take two different approaches to the regression problem. Gaussian process regression takes a holistic view of the data: it builds a regression model by evaluating a covariance function at every input point. This contrasts with the MARS approach that uses knot points to segregate the data into snapshots that can be described using different basis functions. As result of these different perspectives, Gaussian process regression can take a look at the overall effects of the data on the regression model, whereas MARS is formulated to look at both the main effects of the input data as well as their interactions. The approach of MARS is superior when the true function has different regimes. For example if it has a rapidly varying region and an asymptotic regime as is the case with the function $f(x) = (\log x)^2/\sqrt{x}$ which varies rapidly near $x = 0$ and slowly as $x \to \infty$, MARS can seperate $f(x)$ into a rapidly varying piece and a slowly varying piece. Gaussian process regression in looking at the totality of the data cannot discrimmate between these two different regimes of the function. We show results for this regression problem in Figure 136. The GPR results demonstrate the above point that GPR cannot capture the change in behavior of the underlying function.

In terms of the number of parameters required to describe the regression model, Gaussian process regression requires only $m+1$ parameters. This contrasts with MARS where there are parameters that describe the knot points, the interactions, the sign indicators, and basis function order. Of course, for the extra parameters one receives a more flexible regression model as discussed above.

6.4.4 Simulations

To map the relevant input space for our radiating shock simulation we have run 512 simulations using 1D Hyades in run set RS1. The input space we consider is 15 dimensional; see Table 16 for a list of the different input parameters and how each was varied. Of these 15 different inputs 7 inputs describe the experimental configuration being simulated (the first seven entries in Table 16), and 8 are parameters that relate to numerical accuracy (e.g. number of zones in the Be) and model calibration (e.g. Be gamma). The 512 simulations used a Latin hypercube design to partition the 15 dimensional input space.

In Figure 137 the values for the shock location and the density at the shock measured at 1.3 ns are shown as a function of six of the input parameters. In these scatter plots a discernible trend in the data as a function of an input parameter indicates that output quantity of interest is strongly influenced by that parameter. Figure 137 indicates that the shock location is dependent on the electron flux limiter and the laser energy: as each of these parameter values increases the shock value appears to increase. The density at the shock does not have such obvious influence by one of these six parameters. Later, we will return to the question of which input parameters are most important by performing statistical analyses on the emulator results.

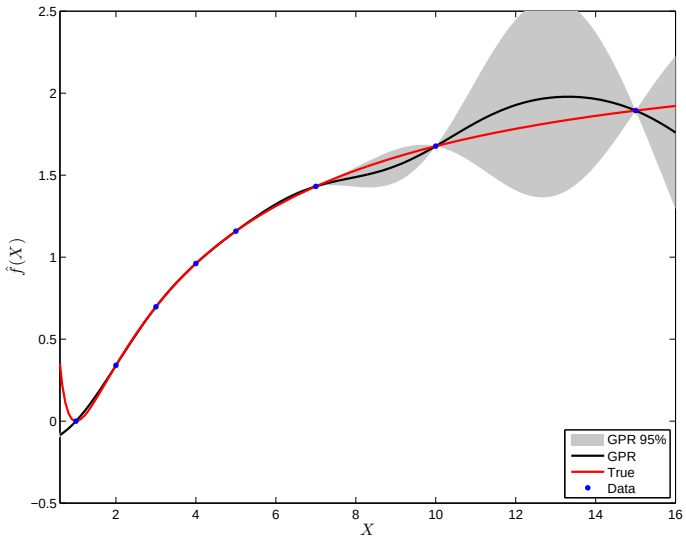

(a) GPR

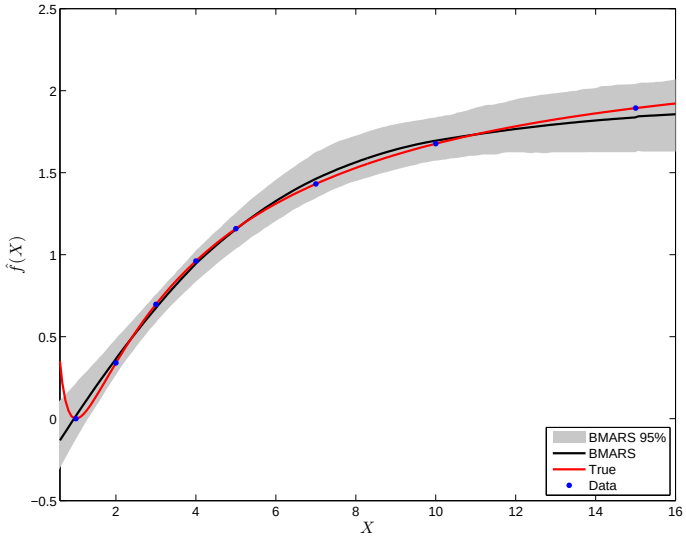

(b) BMARS

Figure 136: Comparison of the two regression models under consideration, GPR and BMARS, for $f(x) = (\log x)^2/\sqrt{x}$.

Table 16: Hyades Simulation Input Parameters and Ranges

Parameter	Nominal	Range	Min	Max
Beryllium Thickness	0.020 mm	10%	0.018	0.022
Laser Energy	3.8 kJ	15%	3.23	4.37
Pulse Duration FWHM	1 ns	10%	0.9	1.1
Xenon Density	.0065 g/cc	10%	0.00585	0.00715
Tube Length	5 mm	-20%	4	5
Laser Rise Time	100 ps	50%	50	150
Slope of Laser Pulse		30%	0.85	1.15
Mesh Resolution (number of zones in Be)	65 zones		20	200
Number of Photon Groups			20	100
Electron Flux Limiter	0.05		0.03	0.1
Time Step Control Multiplier	1		0.25	1
Beryllium Opacity Multiplier	1		0.7	1.3
Beryllium Gamma			1.4	1.667
Xenon Gamma			1.2	1.4
Xenon Opacity Multiplier	1		0.7	1.3

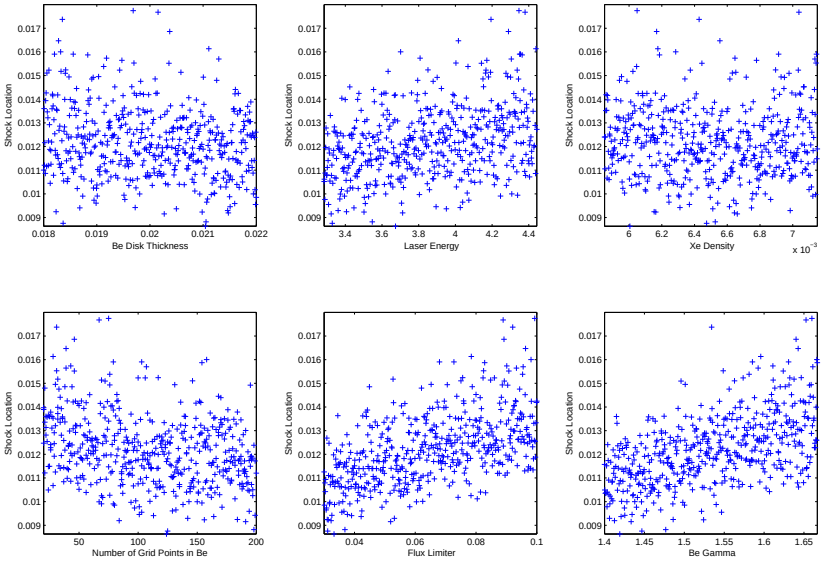

(a) Shock Location

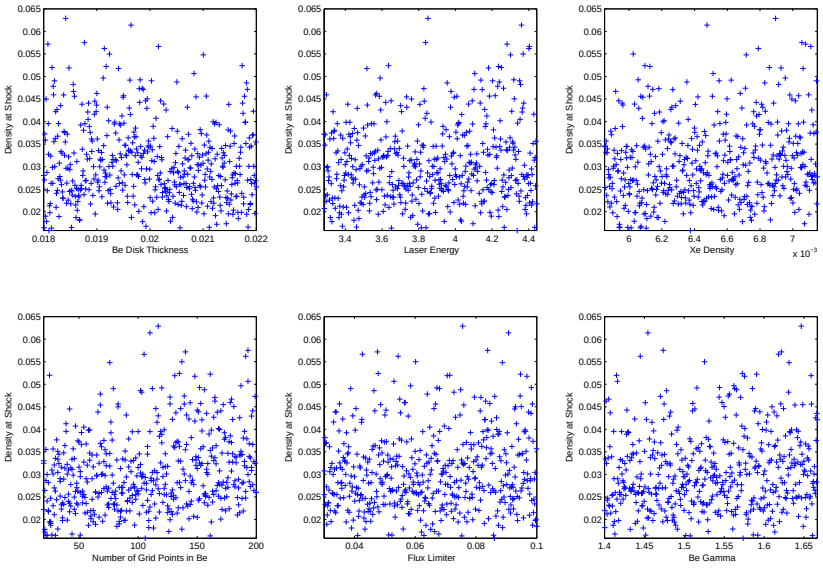

(b) Density at Shock

Figure 137: Scatter plots for the shock location and density at the shock from the 512 Hyades simulations at 1.3 ns as a function of the six most important input parameters.

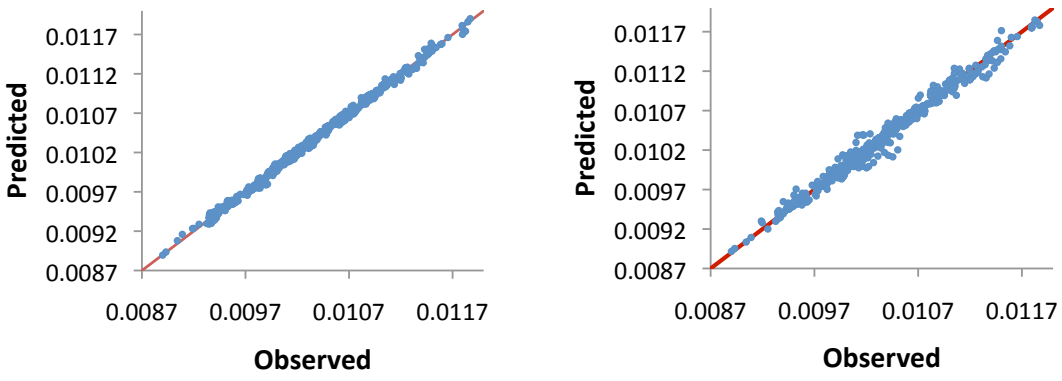

(a) BMARS

(b) Gaussian Process Regression

Figure 138: Comparison of the shock position as predicted by the regression models with the observed Hyades value from the 363 test data results.

6.4.5 Hyades Outputs Emulation and Discussion

The obvious litmus test for an emulator is how well it can predict the output of the code. To test the ability of GPR and BMARS to predict the output of Hyades we used 363 Hyades runs as test data. The results for shock position from BMARS and GPR emulators compared with the actual Hyades shock position on the test data are shown in Figure 138. Shock position is one of the most important output parameters because it is experimentally measurable and the location of the shock in the initial conditions for CRASH should have a large effect on the CRASH output. In Figure 138 perfect emulation would have the data fall on the red line given by $y = x$. A cursory glance at the figure shows that BMARS did a better job of predicting the shock prediction than GPR. Although, it should be said that both regression methods did predict the shock position within 3%. The emulator results for output parameters other than shock position demonstrated similar performance.

6.4.5.1 Analysis of GPR results

From the emulator we constructed using Gaussian process regression we can use the values of l_k found via the empirical Bayes method to get information on which input parameters affect the outputs the most. Specifically, the value of $1/l_k$, called the relative relevance, tells us how important changes in variable k are to the output.

In Figure 139 we show the relative relevance for each input parameter on the shock position. From this figure we see that numerical and model calibration parameters, the mesh resolution and the electron flux limiter, have the largest effect on the shock position. The fact that these parameters are important for the shock position is not surprising. Changing the number of mesh zones in the beryllium can change the shock position because the error in the numerical results is related to the mesh resolution as well as the fact that adding a Be zone might cause a step change in the shock position. In general, one would like to know the output when the number of zones is large to make the discretization error as small as possible. This consideration will be made when initializing CRASH.

Similarly, the electron flux limiter is model parameter that attempts to account for the fact that a diffusion model

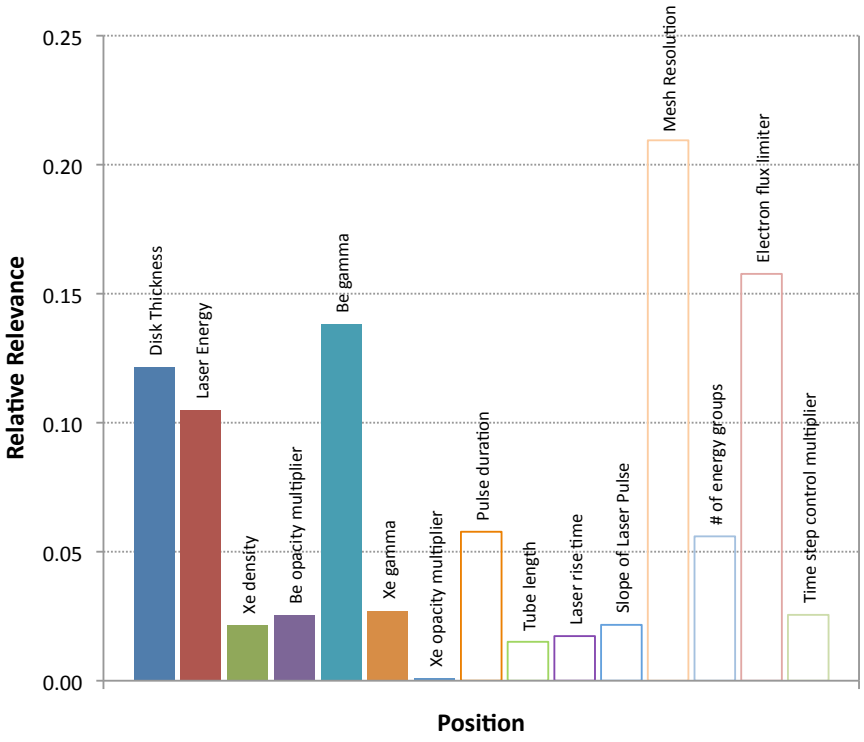

Figure 139: Relative relevance ($1/l_k$) for each input parameter in the GPR emulator for the shock position.

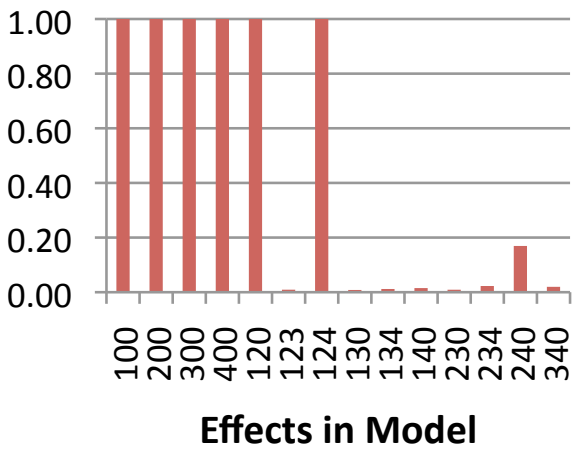

Figure 140: Significance of effects for shock position. The effects are numbered 1 = laser energy, 2 = laser pulse duration, 3 = Xe density, 4 = Be disk thickness, 0 = no effect.

cannot properly capture the streaming of hot electrons. By changing this parameter the bulk speed at which heat is conducted by electrons is changed. In turn changing how heat moves through the problem also changes how the shock moves. The large importance of the electron flux limiter has spurred us to further investigate the germane literature to properly constrain the range of this parameter.

Next in importance is the gamma of the beryllium material, a model calibration parameter that relates to the compressibility of the Be plasma. Experimental parameters, specifically the Be disk thickness and laser energy, are the fourth and fifth most important parameters. In turn adjusting the disk thickness and laser energy would have a larger effect effect on the shock position at 1.3 ns than for instance adjusting the laser rise time. The fact that parameters that describe the experiment are not the leading parameters in terms of relative relevance, points to the fact that the numerics and model calibration aspects of a Hyades simulation are the dominant mechanism for changing shock speed.

6.4.5.2 Analysis of BMARS results

Using the BMARS results we have estimated the interactions between experimental parameters in the emulator. We do this by looking at the MCMC samples from the posterior distribution for the interaction parameters for the emulator and calculating the probability that a sample has that interaction term. In Figure 140 we plot the probability that a particular interaction is in a sample of the emulator model for the shock position. The figure uses a 1 for laser energy, 2 for laser pulse duration, 3 for Xe density, and 4 for Be disk thickness; 0 denotes no interaction. For example, 100 denotes the effect of laser energy only, and 124 is the interaction between laser energy, laser pulse duration, and Be disk thickness. From the figure we can see that each of these four parameters is significant because all samples from the posterior distribution has these effects, i.e., we could not leave out these parameters and have an emulator of similar accuracy. Similarly, changing one of these parameters would change the shock position. Also, the two way interaction between laser energy and pulse duration and the three way interaction between laser energy, pulse duration, and Be

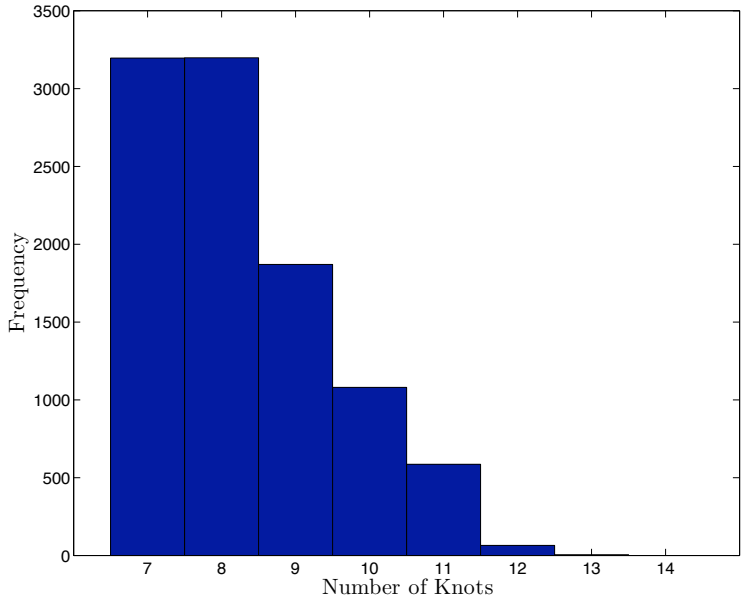

Figure 141: Histogram for the number of knots in the BMARS emulator from samples of the posterior distribution.

disk thickness are important. These interactions are not completely unexpected. For instance, if the Be disk thickness is increased, then the laser energy and pulse duration could be changed to compensate for the greater mass of Be to accelerate.

From the BMARS emulator we can also estimate how many different "regimes" or snapshots of the input/output pairs are needed to predict the shock position. A histogram for the number of knots in the samples from the posterior distribution of the emulator for the shock position is shown in Figure 141. From this figure we can see that the emulator needs at least 7 knots to describe the training data with 7 and 8 being the most common number of knots.

6.5 Predictive study from the 1D simulations

In this section and in [Holloway et al., 2011] we describe our first steps in creating and testing a predictive model that combines a set of data from 1D CRASH code runs (RS2, with CRASH 1.1) with field data from a campaign at the Omega facility [Boehly et al., 1995] from October 2008. As part of this activity a model discrepancy function was created to provide a quantitative picture of the quality of the 1D CRASH model compared to experiments; this discrepancy informs us regarding the extent to which the simulation matches the experiments, modulo the effect of random experimental error. If the discrepancy is significant over some region of input space then this is worthy of investigation to understand why, and assess if reduction in the discrepancy is warranted, and if so to determine what changes in the CRASH code physics will reduce it.

A second aspect of the predictive model is its use to calibrate some uncertain physical model constants. By starting with a non-informative prior distribution of these physical parameters (described below) over an initial range determined by expert judgement, we can learn which of these parameters can be calibrated by the experimental data, and

test the predictive power of the model through a set of leave-one-out exercises in which we predict one experimental result while using the others for calibration.

The topic of quantification of margins and uncertainty (QMU) has received considerable attention of late, particularly by the National Nuclear Security Agency, as reflected for example in the recent National Research Council study [Ahearne et al., 2008]. A core component of QMU is uncertainty quantification, and the NRC report emphasizes the need to further develop the methodology for this work, especially in using computational models to predict, with uncertainty, the threshold or design values for successful operation of an engineered system. The report also emphasizes the importance of developing a diversity of approaches to the methods of establishing margins and quantifying uncertainty. There is also value in applying these diverse approaches to a variety of problems. This section applies one such methodology to a problem in radiative shock hydrodynamics.

6.5.1 The physical system

In the CRASH experiments of interest here, from October 2008, a laser irradiates a Be disk, driving a shock into a cylindrical Xe-filled tube. The shock is sufficiently fast that energy balance requires that it radiate energy away, and so forms a radiative cooling layer immediately behind the shock front. The radiation travels out in front of the shock and heats the wall of the tube, leading to ablation of the plastic and the generation of a second shock (the "wall shock") that travels into the tube radially and interacts with the primary shock [Doss et al., 2009]. The physics is this dominated by the complex interaction among a laser-driven radiative shock, the ablation-driven wall shock, and the Xe-Be interface behind the primary shock. This physics is relevant to astrophysics and fundamental high-energy-density physics research. The experiments are discussed in Sec. 2.1. Their purpose was to explore the variability of the experimental data [Doss et al., 2010]. These were used as our first set of data to build a calibrated predictive model. The shot numbers and input data (nominal x values) are shown in Table 1, and the measured shock locations and shocked Xe layer thicknesses (the outputs y) are shown in Table 2. Because there are two views of the experiment, some of the shots provide two sets of input that differ by observation time. We will use these as independent measurements.

Figure 142 summarizes the shock location from these field data. View 2 in shot 52661 is at 16 ns, and is not shown on this plot. For the predictive model construction in this section we will use 9 experimental points (8 as prior experiments, and the 9th to predict); we eliminate the data marked 52661^* because it is at a late time (although we will explore prediction of this observation below). Two additional shots (52670 & 52671) from the same experimental campaign are not used because the target had alignment issues that produced a shock that was not perpendicular to the tube.

6.5.2 The simulation tools

The CRASH code used for these simulations (version 1.1) had the following properties: a finite-volume MUSCL (monotone upstream centered scheme for conservation laws) [van Leer, 1974, 1979] approach to solve the radiation-hydrodynamics in the gray diffusion approximation, so that the dependent variables are the mass, momentum and total energy densities of the flow, and the radiation energy density. An exact Riemann solver for a sophisticated EOS based on an artificial relaxation method is used to upwind the solution scheme in a mixed implicit/explicit update with heat conduction, radiative transfer and energy exchange between radiation and matter treated implicitly, while the advection terms are explicit. The implicit scheme uses a preconditioned Newton-Krylov-Schwarz method. Multiple materials

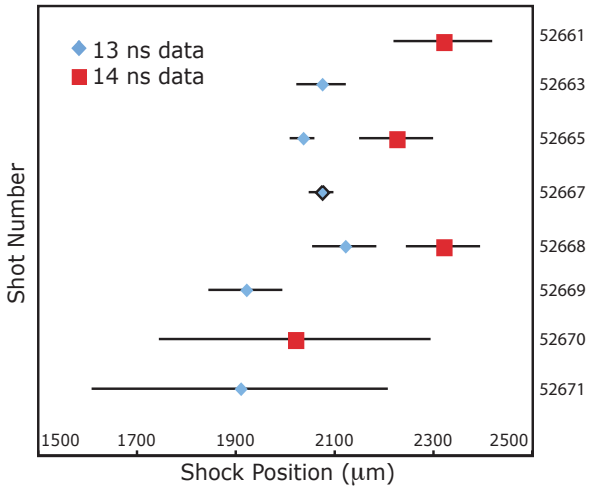

Figure 142: Initial shock location data from the October shots. Note that some shots provide data at two values of t_o, as shown in Table 2. Shot numbers 52670 and 52671 are not used in this section because the targets had alignment problems.

are treated using level set functions initialized based on the signed distance from material interfaces. The level-set function is advected with the fluid to keep track of material location, with one level set function per material. Cells are treated as a single material, assigned to that of the largest level-set function. The material interfaces can be located by adaptive mesh refinement, so that we do not treat mixed material cells; this is important for radiation transport as the different materials can have very different opacities. In the 1D simulations used for this section the Be-Xe interface is well resolved.

The 1D CRASH simulations used for this section are based on gray flux-limited radiation diffusion [Morel, 2000, Drake, 2006]. The full system is updated at each time by first doing a full update on the hydrodynamic variables, with radiation/electron momentum and energy change ignored in this fractional step. The code then passes hydrodynamic variables at cell centers to the radiation solver, which performs an implicit update of the radiation field. The radiation/electron energy change source terms are passed back to the hydro solver, which completes the hydro step, incorporating these source terms [Edwards, 1996a]. The main radiation transport effect on shock position is to remove the correct amount of energy from the shocked Xe; since this shocked Xe is optically quite thick, the gray model can be expected to do well. The use of gray radiation diffusion is seen to be adequate post hoc for the predictions being made here; as we later move into 3D simulations including the plastic tube walls the multigroup effects are expected to be more important in order to compute the correct energy deposition in the wall of the tube.

Because the CRASH code lacks a package to model laser deposition within the Be, it is not able to model the first nanosecond of the CRASH experiments. This modeling is provided by a preprocessor. The Hyades code [Larsen and Lane, 1994] can perform this initial calculation, but it cannot in general perform the full calculation to some 13 to 20 ns after the initiation of the laser pulse because it is a Lagrangian code and in 2D the Lagrangian mesh becomes hopelessly tangled within a few nanoseconds; even with automated remap it cannot follow the CRASH hydrodynamics. It also lacks the parallel performance of CRASH, and does not have a 3D capability, so it is not suited to our long term goals.

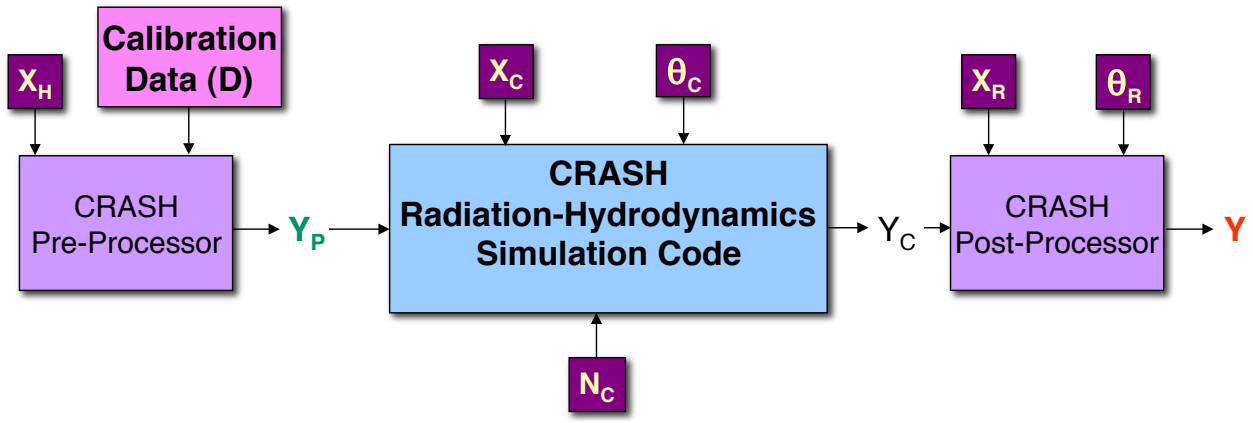

Figure 143: The major components of the CRASH code system prior to implementation of the CRASH laser package.

We therefore use Hyades only as a component in building a preprocessor that can initialize CRASH. The handoff between CRASH and Hyades has been performed at 1.3 ns; we have not examined the sensitivity of our results to variations to this value of 1.3 ns.

The CRASH code system used in this phase of the project then consists of 3 major blocks. A preprocessor that uses a description of the initial system geometry and experimental parameters to estimate the state of the system at about 1.3 ns after the laser pulse is initiated (and shortly after the laser irradiance has returned to zero). This initial state provides matter and radiation fields (ρ, u, p and T_r) to initialize the CRASH radiation hydrodynamics code. In the present study, it turned out that T_r was not needed. The radiative energy fluxes are large, but the radiative energy densities are very small. As a result, the matter recreated the radiation field almost immediately and with negligible energy loss. The code then advances these fields to some later observation time, and which point the fields are provided to a postprocessor which does feature extraction of key system outputs y_s. A block diagram of the system is shown in Fig. 143.

In this section CRASH was run in 1-D with gamma-law EOS and gray diffusion and with opacity scale factors. In addition, in the 1D model there are only two materials, Be and Xe. The use of a 1D simulation to model the CRASH experiment was expected to be adequate for modeling the primary shock location, which is the only experimental output considered here. The shock position is governed largely by the Be slug driving the shock through the Xe, and the position is determined largely by the resulting momentum balance. The 3D aspects of the system come into play when the details of the wall ablation and associated ablation driven shocks are being predicted. The spatial mesh in the 1D simulations was refined in a mesh refinement study so that the shock location was converged to less than 10 μm, less than half of the experimental uncertainties.

6.5.2.1 Preprocessor construction

In this 1D model we must initialize density, velocity, pressure and radiation temperature (ρ, u, p and T_r). We will use a tuned set of Hyades 1D runs to help provide these data; in particular, we must tune the laser energy down in the simulations because the 1D geometry lacks the radial loses of energy and so will result in larger shock speeds than

observed in the actual field experiments. The result of this is that the laser energy must be scaled by a factor f in order to properly launch the shock as observed in experiments. Therefore as we vary the laser energy x_2 to explore the input space, the laser energy input to the code system (and Hyades in particular) is actually fx_2.

In previous radiative shock experiments [Reighard et al., 2007], it was found that in order to match with the shock position, the irradiance needed to be lowered by a factor of $f = 0.46$. Using the new experimental data and new Hyades runs this could be refined. In particular, using $f = 0.42$ and all other inputs (x and θ) kept at their nominal values, the shock is predicted to have traveled 1.983 mm in 13 ns. This compares well with the average shock location of 1.971 mm, and falls well within the error in the data. The need to change the factor is likely mainly due to variances in the resolution of the Hyades runs between the two radiative shock simulations; for the Hyades meshing we have conducted mesh convergence study to converge the shock location to within less than a few percent of the experimental uncertainty. In addition, this calibration exercise assumes that the timing of the measurement at 13 ns is precisely known. The relative error in the timing from one experiment to the next is quite small, but the absolute error is near 0.5 ns. One could explore the impact of the uncertainty in this timing by repeating the study described here using a distribution of measurement times. A better approach, which we are now pursuing, is to obtain independent calibration data based on the early behavior of the experimental system.

In order to construct a fast running preprocessor a set of 512 Hyades runs was conducted to cover the input space using a space-filling Latin Hypercube design (LHD). From each of this set of Hyades runs we constructed a simple profile of piecewise linear and exponential elements that fit the Hyades outputs based on 10 landmark locations in the system state at 1.3 ns. We then build an emulator that models the response function mapping the inputs (x, θ) to the parameters of this fit; this response function is constructed using a Gaussian process model. This provides a physics-based dimension reduction; we therefore call this proprocessor for constructing the CRASH initial state the PIE (physics informed emulator). The PIE is described more fully in McClarren et al. [2011].

The long term value of such a preprocessor for our work is that we describe the intial state with relatively few parameters compared to the 1000's required to define the fields at 1.3 ns. The use of physical insight to determine the important features in system state as modeled by 1D Hyades allowed us to identify only 40 parameters that could adequately describe the system state at 1.3 ns. These 40 parameters are based on identifying 10 landmark locations in the system state at 1.1 ns and using the 3 field values at each of these 10 points (ρ, u, p); the radiation temperature T_r is set to a small value because there is little energy in the radiation field at 1.3 ns. These relatively few parameters are functions of the system input parameters, and this relationship could be captured by a Gaussian process model.

The 512 Hyades runs used to construct this emulator covered a 15 dimensional input space, and provided a sample of 512, 40-dimensional outputs over the 15 dimensional input space. This response was fitted using a Gaussian process model, so that we could compute the response (40 numbers) at any point of the 15 dimensional input space. The additional input parameters included Hyades mesh parameters in space and energy, laser pulse shape parameters, etc. The resulting model was marginalized over input parameters that are not relevant to CRASH. The resulting response surface is a function of 7 of the input data, $x_1, x_2, x_3, \theta_1, \theta_2, \theta_3, \theta_4$, but not x_4 because x_4 is the observation time (between 12.5 and 14.5 ns). This input is not relevant to Hyades, so the PIE is constant as a function of x_4. The PIE can therefore map the 8 dimensional input space to 40 numbers, and given these 40 numbers and the interpolation rules of PIE, we can initialize CRASH 1D at any point in an 8-dimensional input space.

The ability of the PIE to describe the CRASH initial conditions was confirmed by computational studies comparing results where CRASH was initialized using 1) the full Hyades output and 2) the piecewise linear reconstruction of the

Description	Average/Typical Values	Range for Study	Symbolic Name
Be drive disk thickness	21 μm	18 – 22 μm	x_1
Laser total energy	3870 J	3.23 – 4.37 kJ	x_2
Xe gas pressure	6.5 mg/cc (1.15 atm)	5.85 – 7.15 mg/cc	x_3
Observation time t_o	13, 14 and 16 ns	12.5 – 14.5 ns	x_4

Table 17: Experimentally controlled inputs x.

fluid variables using the 10 feature locations as given by the PIE. A sample result is shown in Fig. 135 in the previous section, which shows that the shock position at 17.3 ns is modified by about 75 μm by using the PIE in place of the Hyades full field data to initialize CRASH at 1.1 ns. Note that this error in the shock location is about 20% larger than the average experimental uncertainty in this particular case.

We used the PIE to initialize this uncertainty quantification run set of the CRASH code (RS3). The run set was based on a Latin hypercube design and had 320 prescribed CRASH runs, described further below. To do this initialization we chose the GPR emulator. This decision was based on the fact that we were planning on constructing a GPR-based Kennedy-O'Hagan type model for the final simulation output relative to experiment [Kennedy and O'Hagen, 2001].

To construct the initial conditions we took advantage of the fact that 8 of the input parameters were numerical or model calibration parameters. Because we are primarily interested in the nominal result of a Hyades simulation as the experimental configuration is changed, with the exception of number of Be zones and number of groups, our emulator evaluates each of the 8 numerical or calibration parameters at the median value of its input range. For the number of groups and the number of Be zones we evaluate at the 75th percentile of the input range.

The results of generating the CRASH UQ run set using the GPR emulator for Hyades is shown in Figure 144. Here we show the 40 PIE parameters as a function of input set number. Using the GPR emulator we have information about the distribution about the mean of the parameters. Later analysis of the CRASH outputs, perhaps by sampling from the distributions of the PIE parameters, will be able to discern if the uncertainty in the PIE parameters will impact the CRASH output.

6.5.2.2 System inputs

We will focus here on a specific set of input data to the CRASH code system, and explore the uncertainty in output due to these inputs. The codes have other inputs, including initial physical fields for hydro and radiation, but in this first study we will not treat those other inputs as variable from the UQ perspective. We separate inputs into experimental parameters x that can be controlled and measured to some degree and precision, and uncertain physical parameters θ that are calibrated by physical knowledge and measurement. Tables 17 and 18 list the inputs. For x_1, x_2 and x_3 the ranges selected for the study are about twice the experimental uncertainty in the parameter. Note that the observation times, t_o, of the Omega data are at 13 ns, 14 ns, or 16 ns. We treat this observation time as an input to CRASH because our measured data are at different values of t_o.

The physics parameters are listed in Table 18. While the CRASH 1.1 code—which was used for the work described

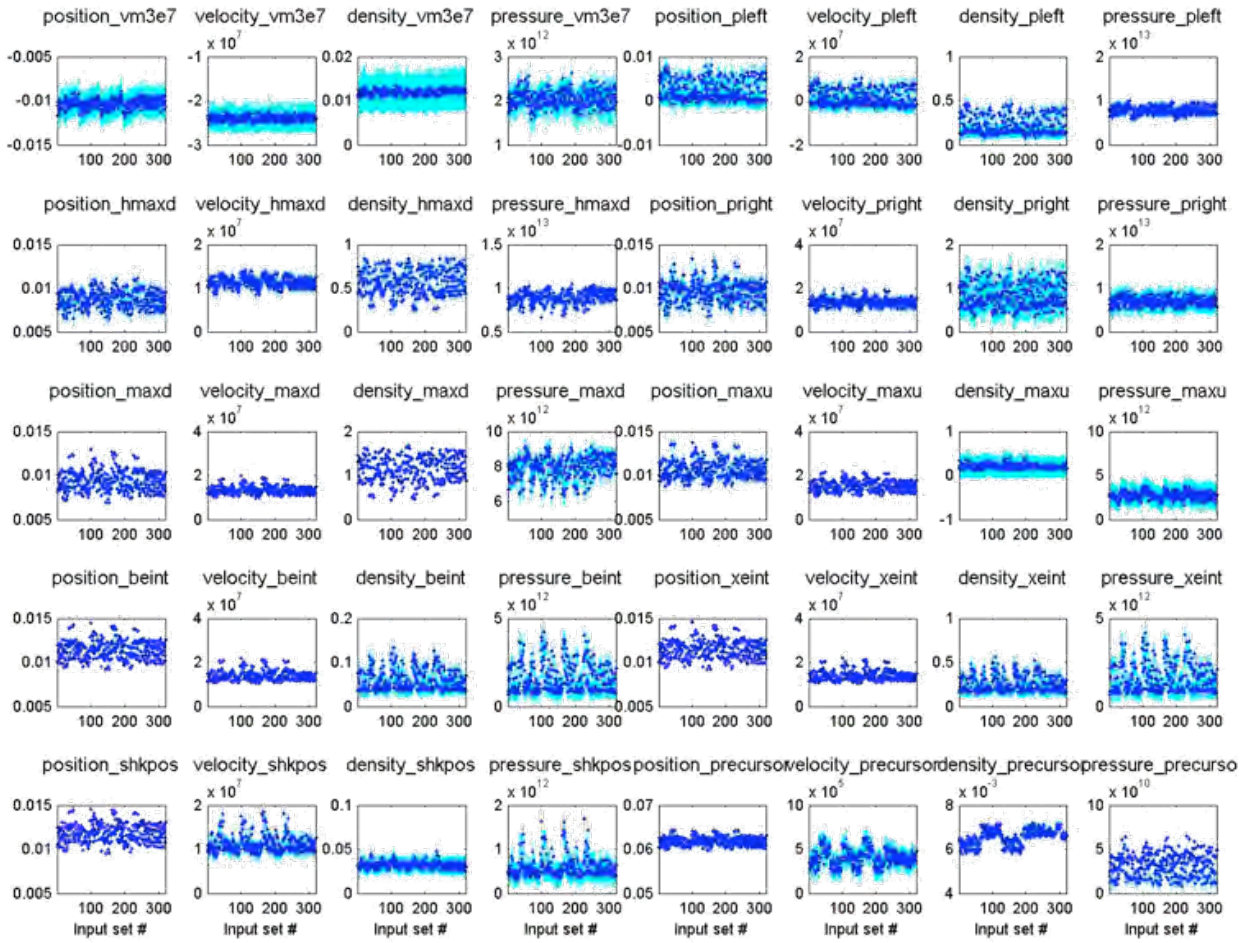

Figure 144: Means and one standard deviation bounds of Hyades outputs to give 320 different initial conditions to CRASH.

Description	Nominal Value	Range for Study	Symbolic Name
γ_{Be} Be gamma	5/3	1.4 – 5/3	θ_1
γ_{Xe} Xe gamma	1.2	1.1 – 1.4	θ_2
Be opacity scale factor	1	0.7 – 1.3	θ_3
Xe opacity scale factor	1	0.7 – 1.3	θ_4

Table 18: Physical/calibration inputs θ.

here—has self-consistent equation of state and opacity models, in this first exercise we will use γ-law equations of state for Xe and Be, and will introduce opacity scale factors on the nominal opacity computed from from SESAME tables [Magee et al., 1995]. Because this is a one-dimensional model of the CRASH experiment, only two materials appear, Be and Xe. The other materials that appear in the actual experiment, most significantly the plastic wall, do not appear in these simulations. For Be we expect the $\gamma_{\text{Be}} \lesssim 5/3$ because the Be is a fully ionized gas, while for Xe, which is only partially ionized, the SESAME tables correspond to $\gamma_{\text{Xe}} \approx 1.2$ to 1.3 over the relevant parameter range. We set the upper limit for the $\gamma_{\text{Be}} = 5/3$ because there cannot be fewer than 3 degrees of freedom for the fully ionized Be plasma; this effectively represents prior knowledge of the physics of the system. We use these simpler models of the physics for this first predictive model exercise so that we can directly explore the sensitivity of the results due to the EOS and opacities, and could explore the computation of posterior calibrating distributions of these simpler models by using the experimental data.

In total, then, we have an 8 dimensional input space to explore, including 4 parameters x that describe the experimental configuration, and 4 parameters θ that describe some aspects of the physics of the system. These parameters are individually referred to as x_i, $i = 1,\ldots 4$ and θ_i, $i = 1,\ldots 4$, as enumerated in Tables 17 and 18.

6.5.3 Predictive model construction

At a fundamental level, we are interested in the following problem: We have a set of inputs, x, which for this section comprise those listed in Table 17. These are controlled, but only within some uncertainty. We also have an output quantity, y, of interest, which in this section is the location of the shock at the observation time. We want to predict y from x.

This problem might be approached by constructing a mapping from x to y, but this straightforward mapping ignores the existence of nuisance variables, ω. Even if we were able to exactly reproduce inputs x between two experiments, $x_1 = x_2$, we could nevertheless *not* expect $y_1 = y_2$ because the nuisance variables would generally be different. Instead we might hypothesize that there is a mapping $y = f(x,\omega)$ from the inputs (both "controlled" and "nuisance") to the output, and there is uncertainty in y both because of the uncertainty in x and the uncertainty in ω.

Furthermore, our ability to construct the function f is limited by our imperfect physical knowledge, and by the need to make approximations both physical and numerical in realizing a model of practical utility. We therefore construct a function $\eta(x,\theta,h)$ that is a function of both the controlled physical inputs x, phenomenological parameters θ that describe physics, and numerical parameters h. In this section the phenomenological parameters are opacity scale factors that measure the strength of interactions of photons with matter and the ratio of specific heats γ in a gamma-

law equation of state for a gas. Generally speaking, the function η is evaluated by solving differential equations in computer codes, and can be quite expensive to evaluate. However we can generate some set $\eta(x_i, \theta_i, h)$, $i = 1, \ldots N_s$ of simulation values to sample the input space of the problem (and at present we fix h sufficiently small via convergence studies, such that the outputs y are insensitive to these numerical parameters) and use these sampled values to construct a model of the output, conditioned also on experimental data as described below. The use of such sampling assumes that the function is sufficiently smooth, and the sampling sufficiently refined to capture its behavior; in our problem we have no reason to expect that the output would not be a smooth function of the inputs, and the success (or not) of the predictions of the predictive model construction will address the extent to which we have sufficiently sampled the input space. In particular, we will use the model to predict shock location at 16 ns after calibrating at 13 and 14 ns, and the shock location is expected to be a smooth function of this observation time.

We can improve our model by using experiments, which result in observations y_j of the output of interest, measured with some finite precision and hence having some uncertainty. These experiments are made with the controlled variables set at some values x_j. However, in reality x_j is known only imperfectly, and there are still nuisance variables ω about which we can say little. We can however select a set of experimental inputs x_j, $j = 1, \ldots N_e$ and measure the corresponding measurements $y_j = f(x_j, \omega)$ with the understanding that y_j is not a single-valued function of x_j because of the nuisance inputs ω that differ from measurement to measurement. Finally, we assume that we have some prior distribution $\pi(\theta)$ that captures our initial belief about the values of the phenomenological parameters, and, if the uncertainty in x_j is to be accounted for, distributions $\pi_j(x)$ that describe the uncertainty in the experimental inputs for each experiment j.

Our task is then this: construct a single combined model that is consistent with both the measurements $Y_e = \{y_j\}_{j=1,\ldots N_e}$ and the computer runs $Y_s = \{y_i\}_{i=1,\ldots N_s}$ and which also provides a distribution of the physics parameters $\pi(\theta|Y_e, Y_s)$ conditioned upon the experimental data. From this model and for a new experiment described by a distribution of input parameters $\pi_0(x)$ we could then construct a distribution of outputs $\pi(y)$. It is this construction of $\pi(y)$ that we call a predictive model.

At present we will simplify this program. In particular, in our work so far we have set aside uncertainty in x. In addition, treating h as a parameter to calibrate is a misguided strategy—at best we will learn that the mesh should be fine and the convergence criteria tight. At worst, the model will suggest the opposite, imply that a coarse mesh is better than fine one and attempting to compensate for some missing physics through a discretization error. We have therefore preceded this predictive model construction with a mesh convergence study to establish a computational mesh sufficiently fine to converge the shock location to within less than 10 μm, which is less than half of the experimental uncertainty.

We have 10 field measurements. However, shot 52661* is left out out because it is a measurement at 16 ns; we will later predict this value as an extrapolation beyond the range of parameters used in calibration. Of the 9 remaining, we leave one out of the model construction process, so $N_e = 8$, and we then predict the 9th. Every one of the 9 will be left out in turn, so that we can construct the model 9 different times to see if all of the outputs can be predicted. We also will have 320 simulation runs, so $N_s = 320$.

We use the modeling framework proposed by Kennedy-O'Hagen [Kennedy and O'Hagen, 2001, Higdon et al., 2004] to describe our predictive model. In particular, the statistical model for the observations is written as

$$y = \eta(x, \theta) + \delta(x) + \varepsilon \tag{242}$$

where η is the simulation code prediction, $\delta(x)$ is the discrepancy that describes the degree to which the simulator does not match the real output y, and ε is random, zero-mean noise that captures experimental variability (and hence, to some degree, captures the impact of the nuisance variables ω). The observations and simulations are combined using the joint model

$$y_j = \eta(x_j, \theta) + \delta(x_j) + \varepsilon_j \qquad j = 1, \ldots, 8 \tag{243}$$

$$y_i = \eta(x_i, \theta_i) \qquad i = 1, \ldots, 320. \tag{244}$$

The first equation represents the modeling of 8 of the field measurements (we always leave one out to later predict), and the second from modeling the output of the 320 CRASH code runs. Notice that the simulations are indexed by θ_i. Recall that the true value of θ is unknown, but some value, θ_i, must be inserted for the simulation code to run. In addition to building a predictive model with associated measures of uncertainty, we also aim to estimate the calibration parameter(s) θ.

Following Kennedy-O'Hagen [Kennedy and O'Hagen, 2001, Higdon et al., 2004], each term, η, δ, ε, in the above model is estimated using a Gaussian process model. Specifically, the output (e.g., shock location y) is specified by a Gaussian process with constant mean μ and covariance structure

$$\Sigma = \Sigma_\eta + \begin{pmatrix} \Sigma_\delta & 0 \\ 0 & 0 \end{pmatrix} + \begin{pmatrix} \Sigma_\varepsilon & 0 \\ 0 & 0 \end{pmatrix}, \tag{245}$$

where Σ is a 328×328 matrix of covariances.

The covariance of the η values between two points (x, θ) and (x', θ') are modeled in the form

$$\mathrm{Cov}_\eta((x,\theta),(x',\theta')) = \frac{1}{\lambda^\eta} \exp\left[-\sum_{k=1}^{4} \beta_k^\eta (x_k - x_k')^2 - \sum_{q=1}^{4} \beta_{4+q}^\eta (\theta_q - \theta_q')^2 \right] \tag{246}$$

where Σ_η is the 328×328 matrix of covariances between all the pairs of points in the set representing both field measurements and simulation data $\{[(x_j, \theta)]_{j=1\ldots8}, [(x_i, \theta_i)]_{i=1\ldots320}\}$. This model represents an assumption that the computer model output, y, is smoothly varying as (x, θ) varies; we know of no physical effects that would invalidate this assumption of smooth variation in shock location with the inputs specified in Tables 17 and 18. Similarly,

$$\mathrm{Cov}_\delta(x, x') = \frac{1}{\lambda^\delta} \exp\left[-\sum_{k=1}^{4} \beta_k^\delta (x_k - x_k')^2 \right] \tag{247}$$

defines the 8×8 matrix Σ_δ of covariances between all pairs of the field measurement inputs $\{x_j\}_{j=1\ldots8}$. Finally,

$$\Sigma_\varepsilon = \frac{1}{\lambda^\varepsilon} I \tag{248}$$

is the 8×8 diagonal matrix representing the measurement uncertainty in the observations, considered as independent. Notice that only the observations inform the random error or discrepancy components directly.

The 16 hyperparameters in this model, μ, λ^ε, λ^δ, λ^η, β_k^η $(k = 1, \ldots, 8)$, β_k^δ $(k = 1, \ldots, 4)$, are estimated by jointly sampling these along with samples of the calibration parameters θ from the posterior distribution $\pi(\theta, \mu, \lambda^\varepsilon, \beta^\eta, \lambda^\eta, \beta^\delta, \lambda^\delta | Y_e, Y_s)$ using a Metropolis algorithm. The likelihood model for this posterior distribution

$$L(z|\theta, \mu, \lambda^\varepsilon, \beta^\eta, \lambda^\eta, \beta^\delta, \lambda^\delta) \propto |\Sigma^{1/2}| \exp\left(-\frac{1}{2}(z - \mu)\Sigma^{-1}(z - \mu) \right) \tag{249}$$

where z is the 328-dimensional vector representing any possible set of 8 field measurement values and 320 computer simulation values. The approach is described more fully by Higdon et al. [Higdon et al., 2004].

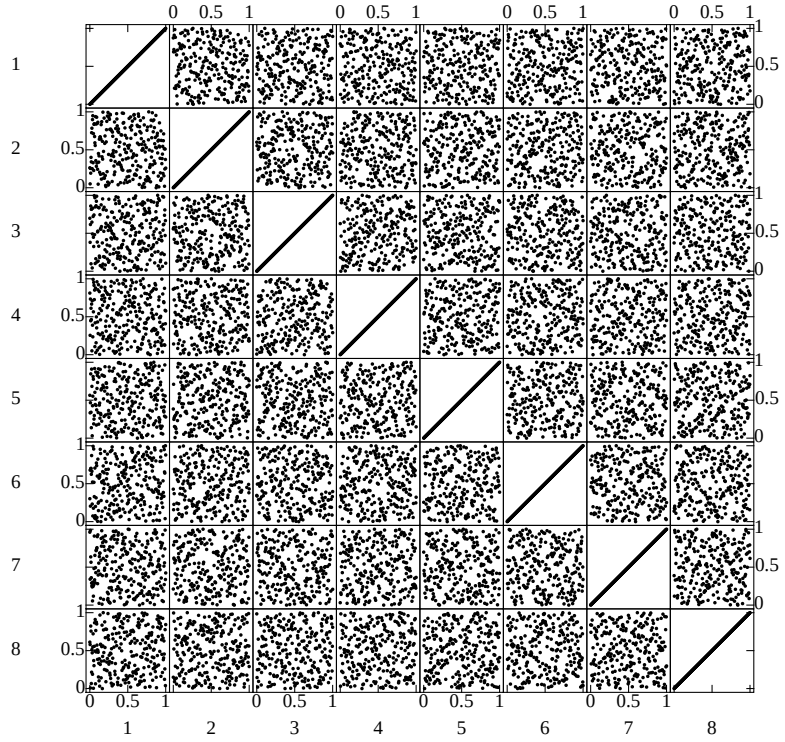

Figure 145: The 256 point part of the input design over 8 scaled parameters (θ and x).

6.5.3.1 Input design

The input design consists of two parts. A 256-point orthogonal array LHD [Tang, 1993] with a space-filling criterion over 8 input parameters (4 xs and 4 θs) added to spread out the points. To this LHD design another 64 points were added. These were constructed by doing a space-filling design (not a LHD) over only the 4 θ's, and each group of 8 points from this design was augmented with the x values from the field measurements 52661, 52663, 52665, 52665*, 52667, 52667*, 52668 and 52668*. Thus, in our set of simulator runs we will have 8 of the field measurement x values nominally matched, with 8 different θ values for each of these x values. The 9th field measurement of interest, 52669, was not given this treatment. These designs are shown in Figs. 145 and 146 in the scaled data (in the range $[0, 1]$).

6.5.3.2 Sensitivity and calibration

Inference for the Gaussian process model is done by sampling from the posterior distribution $\pi(\theta|Y_e, Y_s)$; this provides a distribution for the calibration parameters θ. The posterior distributions of the model parameters are not available in closed form. Instead they are explored numerically using Markov Chain Monte Carlo (MCMC) (e.g., [Besag et al., 1995]). For this model, examination of the trace plots indicated that convergence occurred within about 500 MCMC steps. Indeed, we ran the MCMC for 3000 steps and conservatively discarded the first 1000 points, keeping the last 2000 as a sample from the posterior distribution. Figure 147 shows the marginal posterior distributions of each of the calibration parameters. For each, the prior distribution was uniform (i.e., flat) between the minimum and maximum

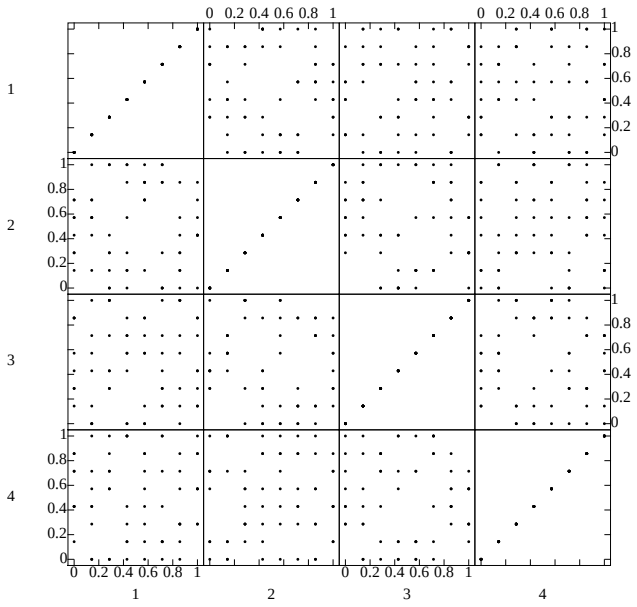

Figure 146: The 64 point part of the input design over 4 scaled parameters (θ).

values. Considering θ_1, physically we expect the fully ionized Be to behave as a perfect gas with only 3 translational degrees of freedom, and hence to have $\gamma_{\text{Be}} = 5/3 = 1.67$; the posterior distribution is consistent with this expectation. For the opacity scale factors the physically anticipated value is 1, as these are multiplicative scale factors multiplying the nominal opacity values, used to explore sensitivity. The posterior distribution for θ_3 shows some preference for the value of 1, while that for θ_4 is very similar to the flat prior distribution. We conclude that there is no information in the measurements to calibrate θ_4. The Xe gamma, θ_2, shows a preference for the high end of the range investigated. It was expected that it might be lower because the Xe is partially ionized, opening many degrees of freedom among the various ionization states; however, over the range investigated the Xe gamma does not have a strong impact on the shock location (compared to experimental uncertainty), and so the experimental data does not provide much information to calibrate the Xe gamma.

Using the model in Eq. 242 we can evaluate the variation of the shock location as a function of any of the individual inputs. Specifically, we look at the expected response as a function of any one of the simulator inputs [Oakley and O'Hagan, 2004]. This gives us a picture of the univariate sensitivity of the output to the input. Figure 148 shows the variation in shock location (in meters) as the inputs vary over their input ranges (see Table 17 and 18). Considering first the θ values, we see that only θ_1 (specific heat ratio γ_{Be}) and θ_3 (the opacity scale factor for Be) show much influence on the primary shock location. The experimental uncertainties range from 1% to 4%, and shock location changes by about 4% as θ_2 varies over the input range; the shock location varies by more, about 15%, as θ_1 varies.

The sensitivity to the experimental inputs x is greater. These show the expected decrease in shock location as the Be disk thickness (x_1) or Xe fill gas pressure (x_3) increases. In contrast, as the laser energy increases (x_2) the shock location also increases, as a higher ablation pressure will be created. These give us a visual picture of the uncertainty in shock location that will result from our uncertainty in these input values. The plot of shock location versus observation time (x_4) shows the linear behavior expected of a shock with nearly constant speed over the range of observation times

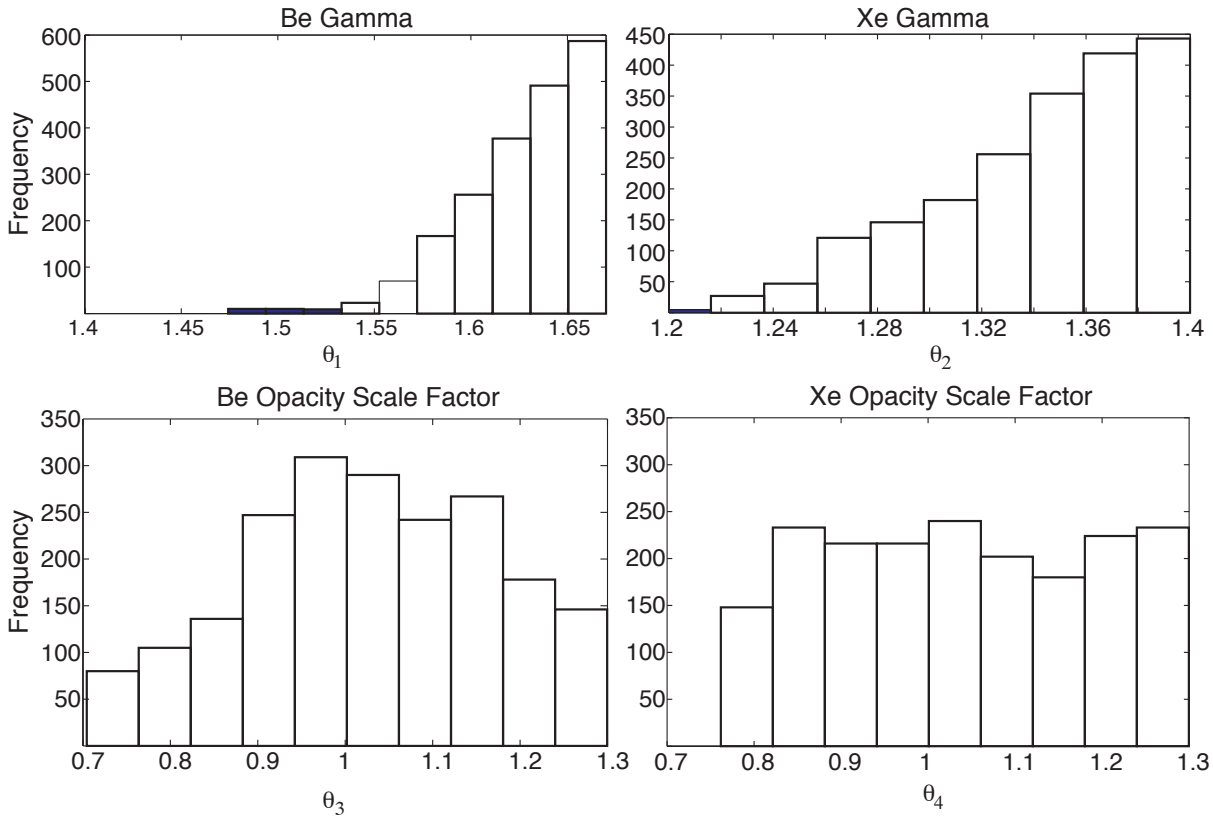

Figure 147: The marginal posterior distributions of calibration (θ) values, based on uniform priors. The calibration shows a preference for the high values of θ_1 and θ_2; this is consistent with expectations for θ_1, while for θ_2 we had expected midrange values to be more probable. The posterior distribution for θ_3 shows some preference for the values around 1, as expected, while the posterior for θ_4 is very similar to the flat prior distribution, reflecting the lack of sensitivity to calibrate this value.

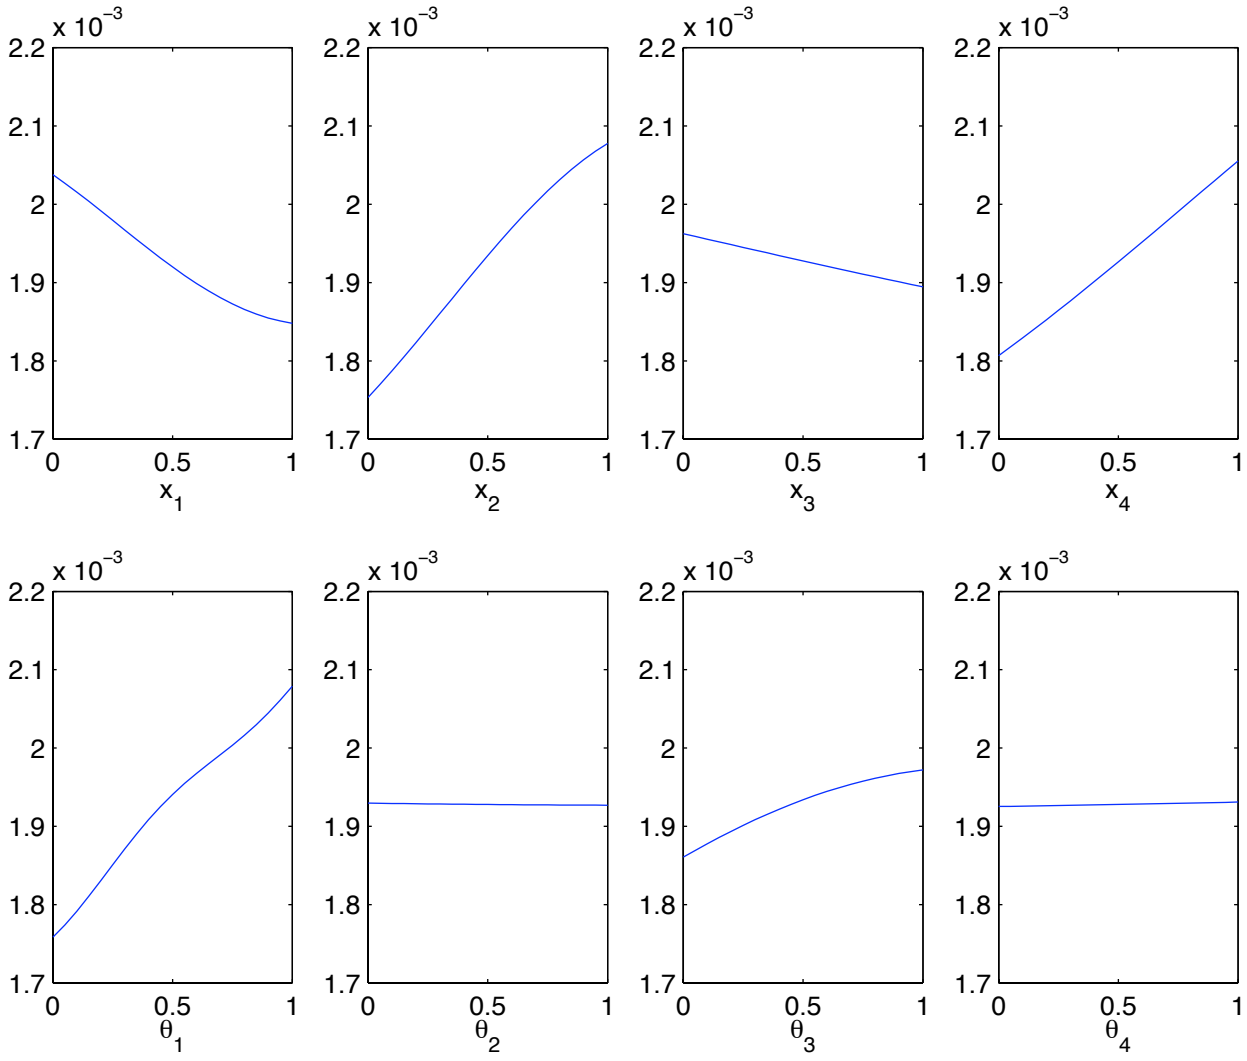

Figure 148: Variation in shock location (in meters) as the inputs vary over their input ranges (scaled to $[0,1]$); each plot shows the variation as a function of one input and averaged over all other inputs. For the calibration parameters θ_1 (specific heat ratio γ_{Be}) and θ_3 (the opacity scale factor for Be) show the most significant impact on shock location. The controlled experimental inputs produce the expected decrease in shock location as the Be disk thickness (x_1) or Xe fill gas pressure (x_3) increases, while as the laser energy increases (x_2) the shock location increases because of the higher ablation pressure created. Shock location versus observation time (x_4) shows the expected linear behavior.

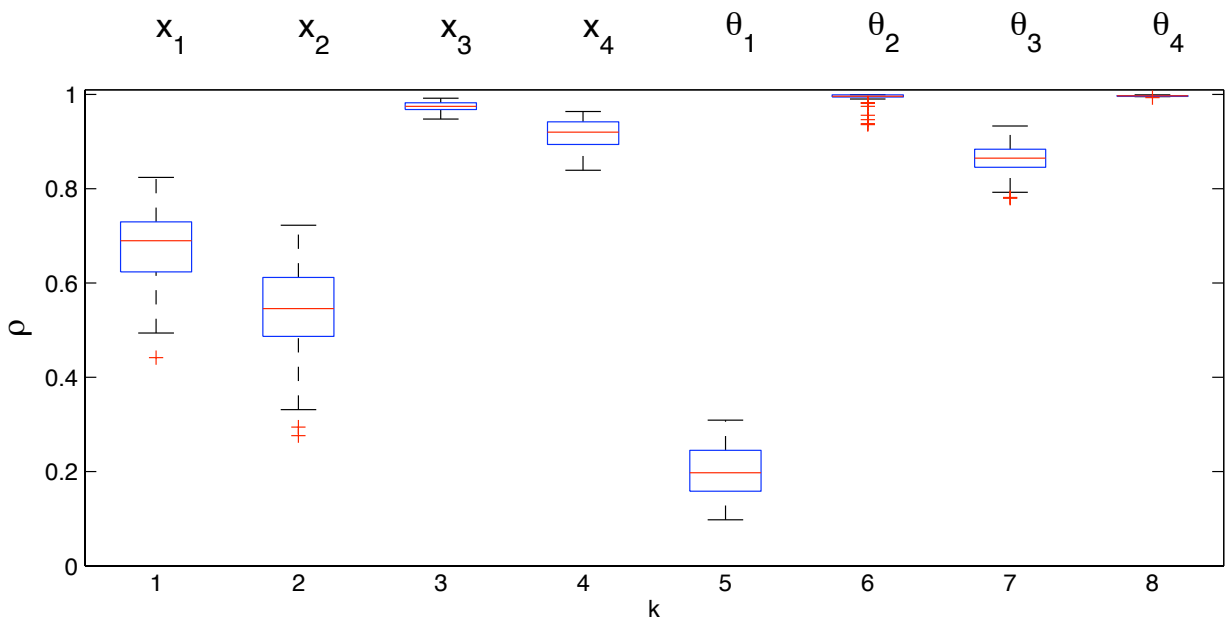

Figure 149: The distribution of correlations ρ_k^η. The box plots show the median (central red line), and 1st and 3rd quartiles (edges of box). The whiskers on each box show the location of the most extreme data in each direction that is within 1.5 times the interquartile range beyond the edge of the box. Outliers are then shown by red crosses. These show very little sensitivity to θ_4 and θ_2, although some outliers on ρ_6^η are more significant and lead to the non-uniform posterior shown in Fig. 147. Input x_3, the Xe fill pressure, is the next next in order of sensitivity.

explored in the model (12.5 to 14.5 ns.)

Sensitivity can be measured in a variety of ways. We are first primarily interested in the impact of each input factor on the response as Fig. 148 aims to visually demonstrate. One can also measure sensitivity by the degree of correlation along each input dimension: the faster the correlation dies off, the more variable the response in that dimension. The correlation in the direction of each input is governed by β_k^η. Let ρ_k^η, denote the correlation in the direction of the k^{th} input at one-half the distance input range; since we have scaled the input to $[0,1]$ this distance is $1/2$ and

$$\rho_k^\eta = \exp(-\beta_k^\eta (1/2)^2) \qquad k = 1,\ldots,8 \tag{250}$$

with β_k drawn from its posterior distribution (hence yielding a distribution of ρ_k values). A small value of β_k results in a value of $\rho_k \approx 1$, and indicates that there is strong "spatial" correlation along this input. In contrast, a large value of β_k results in $\rho_k \approx 0$, and indicates relatively short range correlations. Note that small correlation is evidence of a large sensitivity. Figure 149 gives a box-plot of the posterior distribution of the ρ_k's. A quick glance at the plot indicates that inputs x_1, x_2 and x_4, as well as θ_1 and θ_3, are most important. This is consistent with the conclusions reached when viewing the main effects plots in Fig. 148.

The main effects plots of Fig. 148 are averaged over all input variables except one. So, for example, the effect of θ_2 looks very small. The average correlation corresponding to this input, ρ_6^η, is close to 1, again indicating that is it an unimportant parameter. However, there are a considerable number of outlier values of ρ_6^η at values below

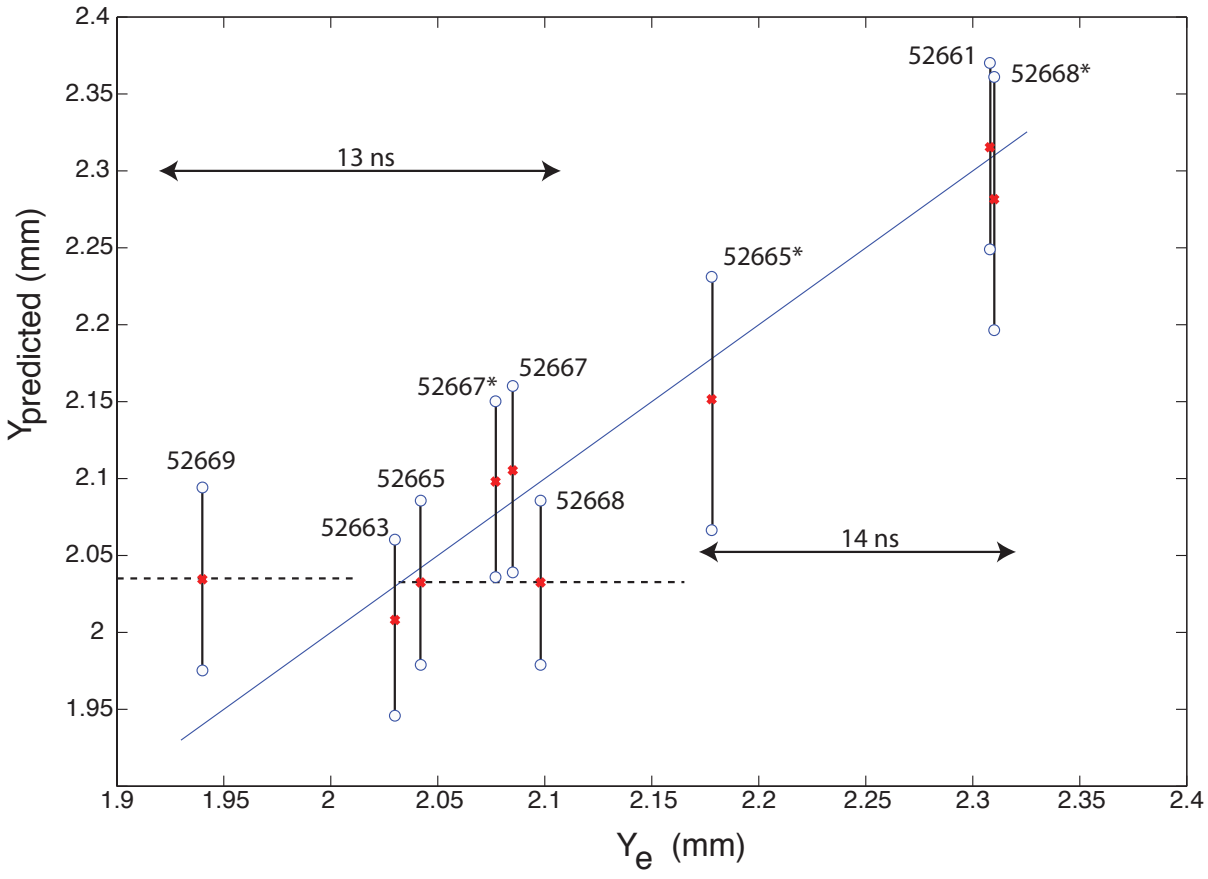

Figure 150: Predictions (on the vertical axis) vs. field measurements (on the horizontal axis) of shock location. Each prediction uses 8 of the field measurements to calibrate the model and predicts the 9th (labeled by shot number). A 90% prediction interval is shown. The data cluster into those from the 13 ns observation time and those from 14 ns. On two of the measurements (52668 and 52669) the experimental uncertainty is shown as a dashed line (others are not shown to keep the plot simple, but are of similar magnitude, and are given in Table 2).

1; this suggest the presence of some input values for which θ_2 is more important, and it is for this reason that the posterior distribution of θ_2 is no longer uniform, even though at first sight the field measurements would seem to lack the information to calibrate θ_2.

6.5.3.3 Predictions

Figure 150 shows the results of a series of leave-one-out predictions. Each prediction uses 8 of the field measurements to predict a 9th (whose shot number labels the prediction). Each prediction is made multiple times and a 90% prediction interval is computed.

The vertical axis shows the predicted mean and 90% prediction interval, while the horizontal axis shows the field measurement. Exact agreement between the predicted mean and the field measurement would be indicated by the

mean falling on the diagonal; of course, this does not happen. However, with only two exceptions (52669 and 52668), the field measurement falls within the 90% confidence band of the predictions. On these two predictions the figure also shows the measurement uncertainty band on the experimental value. Taking this into account the measurement can easily be within the 90% prediction interval on the prediction.

Because the predictive model included a discrepancy $\delta(x)$, it is important to explore that discrepancy and see what it implies. It could be that the simulation is doing a terrible job, and the prediction is totally based on the correction provided by the discrepancy. To explore this, the discrepancy $\delta(x)$ was sampled over the x input space. The values ranged from -5×10^{-5} m to 7×10^{-5} m. This is quite small compared to the shock location of 2×10^{-3} m, and the same as the measurement uncertainty. This indicates that the discrepancy is not dominating the predictions, and indeed, is not needed in this model. The CRASH code is correctly predicting shock location to within experimental error. In fact, refitting the model without the discrepancy resulted in predictions that were again within experimental uncertainty. This does not mean that the CRASH code has full physics (we know it does not), but that when tuned and initialized at 1.3 ns as we described above, and when the θ values are calibrated to the experiments, there is no systematic discrepancy in predicted shock locations over the input range explored. Any difference between the predictions and the field measurements can be explained as zero mean noise of a magnitude similar to the experimental uncertainty.

One of the field measurements, 52661* (see Table 1), provides data with an observation time of 16 ns. This is outside the set of input values that was studied to calibrate the model, so using the model to predict the shock location at 16 ns is an exercise in extrapolation. Performing this extrapolation yields a predicted mean shock location of 2383 μm, to be compared to the measured value 2485 ± 70 μm. The difference between the prediction and nominal measurement is 101 μm. Given this difference we must then question if there is sufficient evidence to suggest that this difference is significant. The 90% prediction intervals are roughly 100 μm in total width and very nearly symmetric, so the difference of just over $100\,\mu$m puts the nominal measurement outside of the 90% prediction interval. But on the experimental side the measurement uncertainty in shock location is $\pm 70\,\mu$m, and in addition there is uncertainty in shock location due to uncertainty in the observation time of approximately ± 0.5 ns which in turn corresponds to an additional $\pm 80\mu$m uncertainty in the shock location at 16 ns. The 100μm difference between prediction and measurement can certainly be understood as uncertainty in the measurement (which can be as much as $\pm 150\mu$m), even without accounting for the predictive uncertainty.

6.5.4 Conclusions of the 1D predictive study

We have constructed a predictive model that can compute the distribution of shock locations in a radiative shock experiment. The uncertainty in shock location in the model is due to uncertainty in the phenomenological parameters θ and uncertainty in the hyper-parameters of the Gaussian process model. Additional uncertainty in the shock location due to uncertainty in the inputs x has not yet been accounted for, but the main effects plots in Fig. 148 can give us some insight into the influence of these parameters. The range of x studied (and represented in Fig 148) is approximately twice that observed as experimental input uncertainty, so looking at half the output variation due to the various x's suggests an uncertainty of about 200 μm due to the uncertainty in x. Any set of eight of the field measurements provide sufficient information to calibrate the parameters θ; the most important of these is the Be-gamma, which is well calibrated. Next most important is the Be opacity scale factor, and the posterior distribution of this parameter

Figure 151: Posterior distribution of shock position.

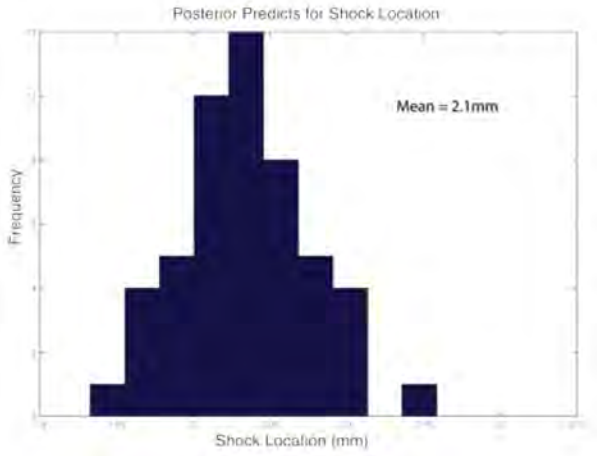

shows the expected peaking around 1, but there is not enough information in the field measurements to calibrate this parameter very well. These results are not unexpected, given that we are using ideal gas law equations of state and a 1D simulation to model the 3D field experiments.

Overall the predictions for shock location are quite accurate; when the field measurement uncertainty is accounted for the measurements are within the 90% prediction interval of the predictions. The discrepancy is small, and predicting without discrepancy is just as good as predicting with discrepancy for shock location. The noise, ε, in the model fully captures the experimental scatter. The shock location is most strongly dependent on observation time, and, within experimental uncertainty, not on the other input variables.

In summary, we have used 320 runs of CRASH 1D and 9 field experiments for our analysis. We focus on one output variable: shock location. To evaluate predictive ability we hold one experiment out, use only the 8 remaining field experiments to construct the model, and predict the ninth. This is repeated for each experiment. As an example of the model output, Figure 151 shows the posterior distribution of the location of a single shock

6.6 Sensitivity analysis from early 3D simulations

One-dimensional studies can provide useful information about the sensitivities of output quantities of interest to variations in the input parameters. However, a number of important output features, such as the wall shock, cannot appear in one-dimensional geometry. As a result, we performed a preliminary three-dimensional sensitivity study of the baseline CRASH experiment. Each simulation was performed on a uniform Cartesian 1200 x 240 x 240 mesh using the CRASH code. A second set of runs was performed on a 600 x 120 x 120 mesh to test grid convergence. For this preliminary analysis, we used simplified physics, treating each material as a gamma-law gas and computed radiation transport with gray flux-limited diffusion. All simulations were initialized from a two-dimensional HYADES output file at 1.3 ns using nominal values of the input parameters. Each simulation required approximately 5 hours on 1024 cores of hera at LLNL. The entire set of simulations was completed in about 2 weeks.

The study consisted of 64 simulations at each grid resolution varying four input parameters: the equation of state gamma for Be was varied between 1.4 and 1.66667; the gamma for Xe was varied between 1.1 and 1.4; and the opacity scale factors for both Be and Xe were varied independently between 0.7 and 1.3. The parameter combinations were

Figure 152: Mass density at 13 ns showing the three output quantities of interest from runs using an early version of the CRASH code, initialized using output from a Lagrangian code that evaluated the behavior for the first 1.1 ns.

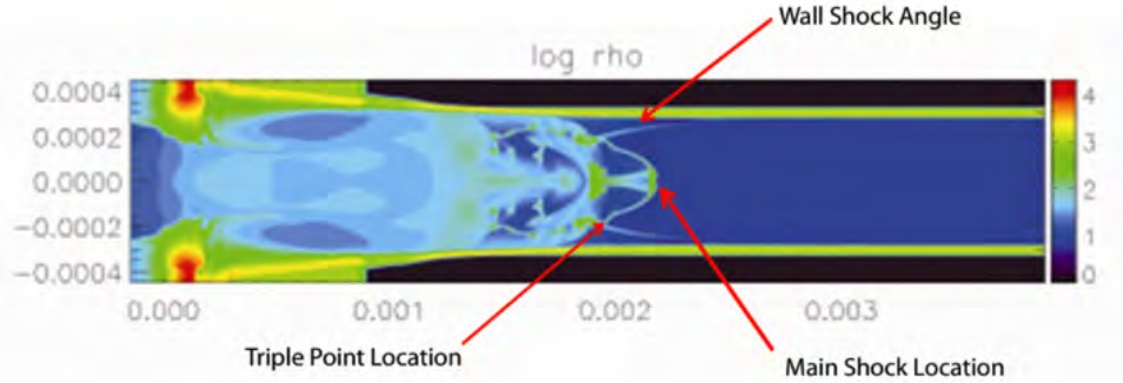

determined using a Latin hypercube design. All other parameters were fixed at their nominal values. We examined three output quantities of interest the location of the main shock, the angle between the wall shock and the plastic tube, and the distance of the triple point from the wall. A typical result showing these output quantities is plotted in Figure 152. All three of these quantities showed surprisingly good agreement with the experiments, even though the overall morphology of the flow in the experiments shows significant differences from the simulations. Figure 153 gives an indication of the wide variety of flow morphologies that are possible with different combinations of the input parameters.

The results indicate that the location of the main shock, defined here as the forward most location of a significant density jump in the xenon, is quite insensitive to the variations in these input parameters. The variation in location was much smaller than the range of values observed in the experiment. However, the shock location may still be sensitive to variations in these input parameters during the first 1.3 ns, before the initialization of the three-dimensional CRASH simulations. We also observed that the location of the main shock is not converged at these grid resolutions, although the error is still less than the experimental range. The angle of the wall shock shows a strong linear correlation with Xe opacity, but no correlation with the other three input variables. This makes physical sense, since the Xe opacity determines how far the radiation can penetrate ahead of the shock. For lower opacities, the wall shock begins further down the tube so that the angle is reduced. The triple point location shows a weak correlation with the Xe gamma, but no noticeable correlations with the other three input variables.

We also constructed plots showing relative importance of the input parameter variations using both MARS and MART. The two methods produced nearly identical results. The most important source of variation in the shock location is the Xe opacity scale factor, although as stated above, the variation in position is very small (at least after 1.3 ns). As expected, the variation in wall shock angle is determined entirely by the Xe opacity scale factor. On the other hand, the Xe gamma is almost entirely responsible for the variation in the triple point location. However, as mentioned above, input parameters that did not prove to be important in this study may still be important during the first 1.3 ns.

Figure 153: Pressure at 13 ns for various combinations of the input parameters, showing the wide variety of flow morphologies that are possible.

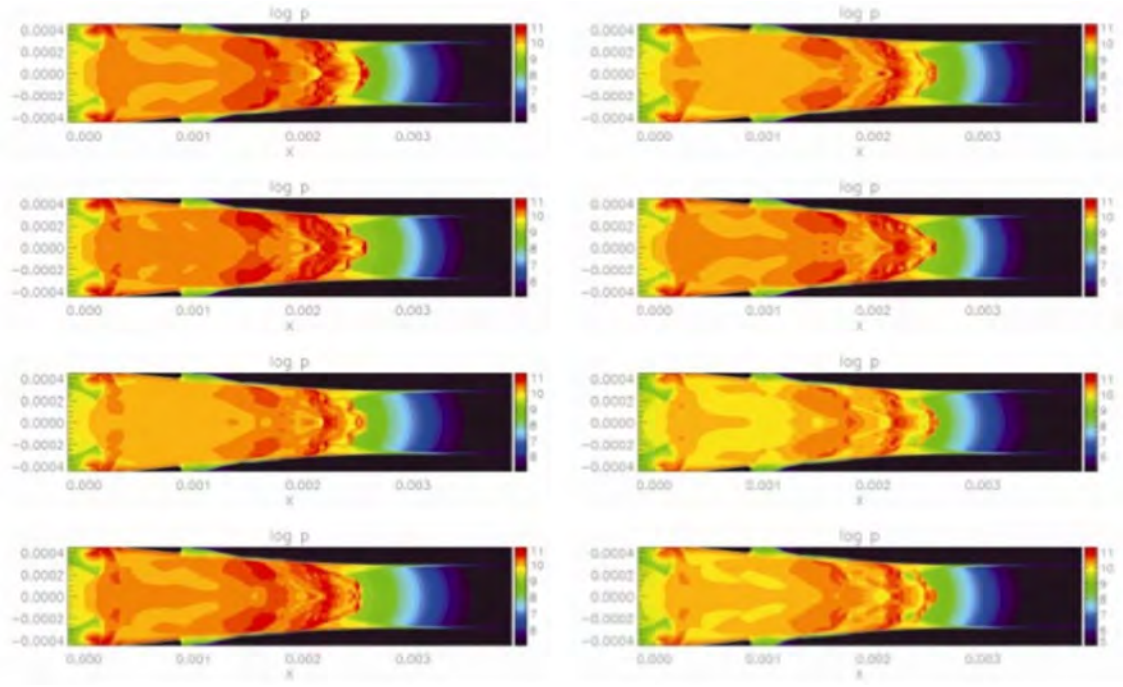

6.7 Verification of uncertainty quantification software

Verification and validation of simulation codes has been a major topic of research for many years. However, little attention has been devoted to verification of software used for uncertainty quantification analysis. As a first step in this process, we have performed a UQ analysis using a simplified problem with an analytic solution to determine if our UQ software is producing sensible results. The analytic solution in this case can be used as a substitute for experimental data and compared to the simulation results. As part of the verification process, we compared the results from four different UQ methodologies Gaussian process, MARS, Bayesian MARS, and MART. We also tested the ability of the UQ software to distinguish between active and inert input parameters. Finally, we performed a blind calibration of an input parameter whose correct value was known.

The problem we studied was a simple shock tube, which exercised only the hydrodynamics solver in the code. The initial conditions consisted of gas at high density and pressure separated from a second gas with a lower density and pressure by a membrane. Both gases were initially at rest. The size of the jumps in density and pressure across the membrane were chosen to match those encountered in the CRASH experiments. Three input parameters were varied the pressure and density in the high-density gas and the value of gamma in the equation of state. The pressure and density in the low-density gas were held fixed. Five inert input parameters that had no effect on the solution were also varied. The study consisted of 62 parameter combinations using a Latin hypercube design. The same parameter combinations were used for both the simulations and the analytic solutions. Eight output quantities were examined. These included the locations of the shock front xshk, the contact discontinuity xcd, the head of the rarefaction xhead, and the tail of the rarefaction, xtail. The other four quantities were the values of density rhshk, pressure Pshk, and

Figure 154: Sensitivity study using Bayesian MARS showing the probability of importance of each primary effect as well as the most important interactions for the four feature locations in the shock tube solution.

x_{shk}		x_{cd}		x_{tail}		x_{head}	
Effect	Prob.	Effect	Prob.	Effect	Prob.	Effect	Prob.
ρ	1.00	ρ	1.00	ρ	1.00	ρ	1.00
P	1.00	P	1.00	P	1.00	P	1.00
γ	1.00	γ	1.00	γ	1.00	γ	1.00
(ρ, P)	0.64	(ρ, P)	0.02	(ρ, P, γ)	0.03	(ρ, P)	0.04
(ρ, P, r_4)	0.05	r_5	0.02	(P, γ)	0.02	r_5	0.02
(ρ, P, γ)	0.02	r_4	0.02	(ρ, P)	0.02	r_3	0.02
(ρ, γ)	0.02	(ρ, P, γ)	0.01	r_4	0.01	(ρ, r_4)	0.01
(ρ, P, r_3)	0.01	(ρ, r_4)	0.01	r_5	0.01	r_2	0.01
(P, γ)	0.01	r_2	0.01	(γ, r_1, r_2)	0.01	(ρ, P, r_1)	0.01
(P, γ, r_1)	0.01	(P, r_1)	0.01	(ρ, P, r_2)	0.01	r_4	0.01

Figure 155: Relative importance plots obtained using MART. The first three columns are density, pressure, and gamma. The remaining five columns are inert parameters. Each plot shows results for one of the characteristic output variables.

ρ_{shk}		ρ_{cd}		P_{shk}		u_{shk}	
Effect	Prob.	Effect	Prob.	Effect	Prob.	Effect	Prob.
ρ	1.00	ρ	1.00	ρ	1.00	ρ	1.00
P	1.00	P	1.00	P	1.00	P	1.00
γ	1.00	γ	1.00	γ	1.00	γ	1.00
(P, γ)	0.59	(ρ, P)	1.00	(ρ, P)	1.00	r_4	0.01
(ρ, γ, r_5)	0.07	(ρ, γ)	1.00	(P, γ)	0.52	(ρ, P)	0.01
(P, γ, r_3)	0.06	(P, γ)	1.00	(ρ, γ)	0.04	r_1	0.01
(ρ, P)	0.04	(P, γ, r_4)	0.04	(P, γ, r_3)	0.02	r_5	0.01
(P, γ, r_5)	0.03	(ρ, r_1)	0.03	r_2	0.02	(ρ, r_4)	0.01
r_2	0.02	(ρ, r_3, r_4)	0.02	(P, γ, r_5)	0.02	(ρ, γ)	0.01
(ρ, γ)	0.02	(γ, r_4, r_5)	0.02	(ρ, r_3, r_5)	0.01	r_2	0.01

velocity ushk behind the shock front and the value of density rhocd between the contact discontinuity and rarefaction. These four quantities are constant both in space and in time.

The simulations and analytic solutions were virtually identical, except for a small bias in detecting the locations of discontinuities on the finite difference grid. The sensitivity analysis performed using all four methods produced consistent, but not identical, results. Since the four methods use different measures of relative importance of the input parameters, perfect agreement was not expected. In addition, all four methods successfully distinguished between the active and inert variables.

Figures 154 and 155 contain the results obtained using Bayesian Mars. The first three lines show the probability that each of the three active input parameters is important in producing variations in each of the eight output parameters. The remaining lines show the probability that various two-way and three-way interactions are important. Variables r1 through r5 are the inert input parameters. As expected, none of these parameters are important either by themselves or in interactions with other variables.

Figure 156: Sensitivity study using Bayesian MARS showing the probability of importance of each primary effect as well as the most important interactions for the other four output parameters of interest.

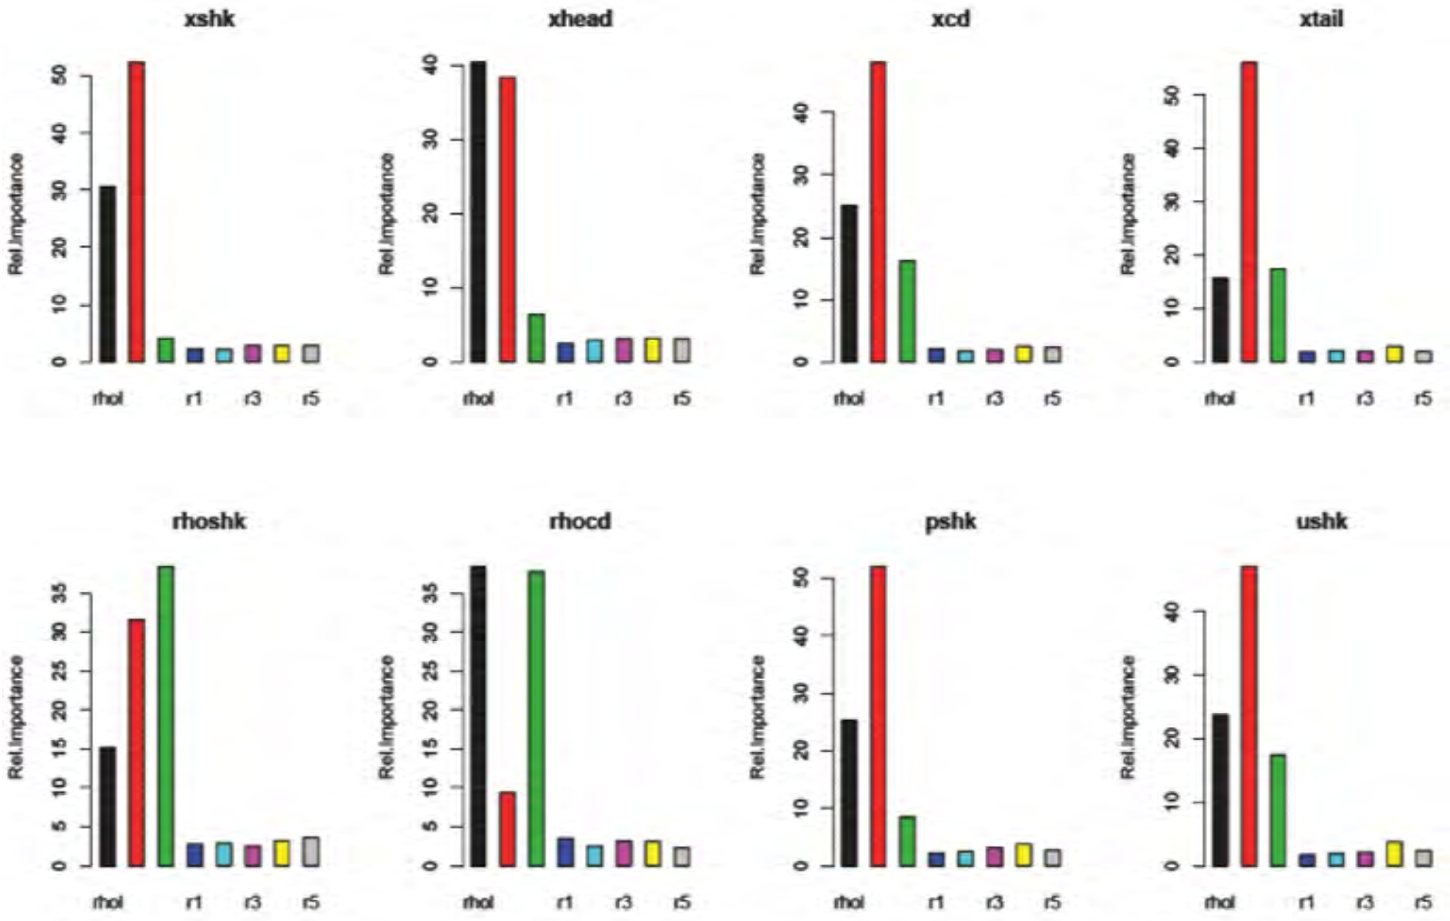

Figure 156 shows relative importance plots of the eight input variables for each of the output parameters. The first three columns represent the input density, pressure, and gamma followed by the five inert parameters. As expected, none of the five inert variables produced a significant signal in the analysis. The feature locations, shown in the top row of plots, are most sensitive to the initial density and pressure. However, the value of gamma is very important in determining the variability in the two density values.

The final test was to calibrate the value of gamma. A set of ten analytic solutions was computed varying only the density and pressure with gamma fixed at 1.4. The number of analytic solutions was reduced for this study to represent the normal case where there are many fewer experiments than simulations. These results were compared with the full set of numerical simulations in which all three input parameters were varied. The posterior distribution for gamma, shown in Figure 12, had a mean of 1.41, in excellent agreement with the expected value. As the number of analytic solutions is increased, the mean of the posterior distribution of gamma converges to the correct value and the standard deviation decreases.

In summary, all four of our UQ methods work well for this problem and provide believable results that are consistent with each other and with the physics of the problem. They reliably differentiated between active and inert input parameters and produced reasonable posterior distributions for calibrating the value of gamma. From this study, it appears that all four of our UQ methods should provide reliable results when applied to simulations and experiments of the complete CRASH problem.

6.8 2D Predictive Study for Shock Location at 20 and 26 ns

In the course of the thesis research of Forrest Doss, we acquired some data at later times. Later, we realized that constraints imposed by target fabrication would mean that the year-5 experiment would also need to be at later times. This provided an opportunity for a predictive study that would extrapolate in time. This study was based on run set RS4, which used CRASH 1.1.

The CRASH experiment creates a high energy density radiative shock in a Xe-filled tube, with a shock velocity on the order of 100 km/s. The shock is first created in a Be metal disk 20 microns thick by a 1 ns laser pulse of 360 J. The shock breaks out of this disk some 400 ps after the initiation of the laser pulse, and continues down the Xe-filled tube, compressing and heating the Xe sufficiently to radiate, and this radiation in turn preheats the Xe ahead of the shock and ablates the plastic wall of the tube, creating additional radial shocks traveling inwards from the wall. To predict the location of the primary shock we use two radiative hydrodynamics codes, Hyades and CRASH. Hyades models the laser-plasma interaction and can predict the shock breakout time and the state of the system at 1.1 ns after the initiation of the laser pulse. The CRASH code, when initialized with this state at 1.1 ns, can predict the shock location at later times when shock location can also be observed in experiments at observation times from 13 to 26 ns.

Our general interest is in using the simulation tools together with experiments conducted in one region of input space, to make predictions in a new region of input space in which no prior experiments have been made. We are generally interested in extrapolation from one region of input space to another, with this extrapolation accomplished by a simulation code that contains the necessary physics. In particular, we have two data sets on which to base predictions: shock break time data, and shock location data at 13, 14 and 16 ns, and wish to predict shock locations at 20 and 26 ns (which are then compared to subsequent field measurements). We use two models of the Kennedy-OHagen form to combine field measurements with simulations, using one to inform the other, and we interpret the discrepancy in these

Figure 157: Posterior distribution for electron flux limiter parameter calibrated using shock breakout times. This is the marginal distribution, but the analysis produces samples from the full joint distribution of calibration parameters.

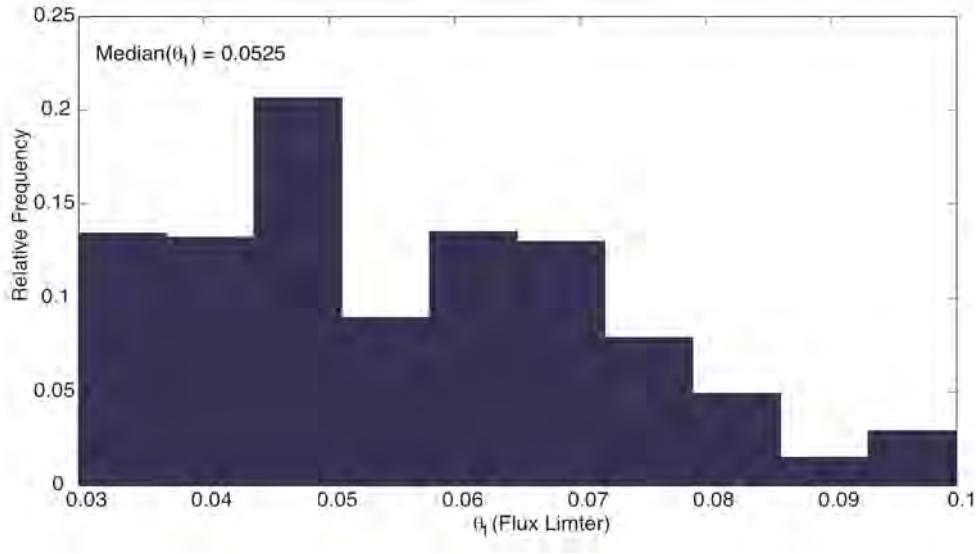

models in a way that allows us to gain some understanding of model error separately from parameter tuning.

To model shock breakout times we construct a model of the form $t = \eta BO(x,\theta) + \delta_{BO} + \varepsilon_{BO}$ that jointly fits the field measurements, T, of shock breakout time t along with a set of 1024 Hyades simulations over a 6 dimensional input space with 4 experimental variables x and 2 calibration parameters θ. This model provides posterior distributions for various modeling parameters, including a posterior $\pi(\theta|T)$ for the calibration parameters, as well as for the parameters in Gaussian process models of the emulator $\eta BO(x,\theta)$, the discrepancy function δ_{BO} and the replication error ε_{BO}. A sample of such a posterior for the electron flux limiter parameter (marginalized over the other calibration parameter) appears in Figure 157.

If the discrepancy function is significant compared to measurement uncertainty we would call this process tuning, but if, as is in our case, the discrepancy is small, then we refer to this as calibration (for shock breakout time). Figure 158 shows a leave-one-out predictions of shock location and of shock breakout time, showing the discrepancy compared to the measurement uncertainty. The discrepancy for shock location is significant, while for breakout time it is insignificant. We therefore can calibrate using the breakout data, and will then use shock location time data from 13 to 16 ns to estimate discrepancy for better understanding of predictions at 20 and 26 ns.

Once we have posterior distributions for calibration parameters we then use the shock location field data (at times of 16 ns and less) along with 1024 simulations of shock locations from CRASH (v 1.1) to construct a model of the form , but in this model is no longer treated as a calibration parameter, but instead is treated as an experimental parameter and is drawn from the posterior constructed in the previous step . The are also drawn from distributions that represent the understood uncertainties in the experimental parameters. This second model is used to construct the emulator and its discrepancy , as well as a best estimate of the replication error , all for shock location. The discrepancy from this model can be studied to understand the defects of the physics model; note that because we have separated calibration from the construction of this discrepancy, the calibration of the is not simply masking errors in predicted

Figure 158: The full set of simulation results (green), a single measurement to predict (open circle), calibrated code predictions (blue), discrepancy (pink) compared to measurement error (black), and finally the prediction of the measurement. All ranges shown are 95% prediction intervals.

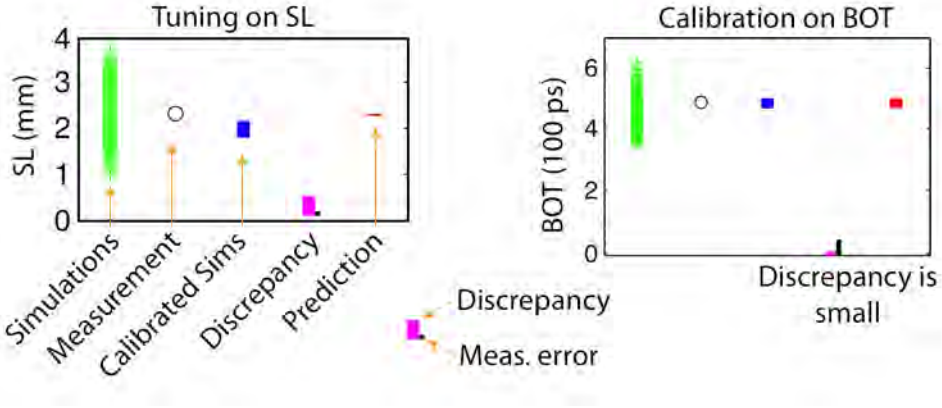

shock location. The result shows that our model tends to under predict shock location.

Finally we can use to predict shock location at 20 and 26 ns, a region of phase space in which we had simulations but no previous measurements. This produced the results shown in Figure 159. In doing this analysis we can separate the code prediction and the uncertainty due to this prediction (caused by uncertainty in x, , and in the Gaussian process modeling parameters) from the uncertainty due to discrepancy . The uncertainty in discrepancy is of course large, because we are extrapolating the discrepancy to a new region of input space. The uncertainty in the emulator is significantly smaller because there were simulation data in this region.

Comparison of the predictions with field measurements at 20 and 26 ns show that even the smaller predictive interval from the emulator alone contains the actual field measurements. The results, shown here as 95% predictive intervals show a median shock location of 2750 microns at 20 ns, and 3200 microns at 26 ns. These compare well with experimental measurements of microns and microns, respectively.

6.9 Combining Different Models

We have developed an approach to combine and simultaneously calibrate multiple models. This is important in support of our plan to base our predictions of the Year-5 experiment on combinations of models of varying fidelity, such as 3D Gray and 2D MG. We began by working with results from 1D and 2D run sets, as a way to develop and demonstrate the methodology. The structure of the combined model for this purpose is

$$y_{f,k} = \eta(x_k, \theta_1, \theta_f) + \delta_1(x_k, \theta_2, \theta_f) + \delta_2(x_k) + \varepsilon \qquad j = 1 \ldots N_f \tag{251}$$

$$y_{s2,j} = \eta(x_j, \theta_1, p_{f,j}) + \delta_1(x_j, p_{2,j}, p_{f,j}) \qquad j = 1 \ldots N_2 \tag{252}$$

$$y_{s1,i} = \eta(x_i, p_{1,i}, p_{f,i}) \qquad i = 1 \ldots N_1 \tag{253}$$

where $y_{s1,i}$ is a set of N_1 data values of a QOI from a set of 1D CRASH runs (RS3) using experimental parameter values x_i, phenomenological/physical parameter values $p_{1,i}$ that are unique to the 1D code, and phenomenological parameter values $p_{f,i}$ that are available for tuning in both 1D and 2D CRASH. In this η is seen to be a code emulator that is

Figure 159: Results of the Kennedy-OHagan type analysis with calibration, predicting new observations at 20 ns and 26 ns in advance of knowing the data. The red predictive interval is that due to propagating uncertainties in θ and x through the simulator, and the uncertainties in the emulator $\eta_{SL}(x, \theta)$. In contrast, the blue predictive interval includes the uncertainty due to the discrepancy, which is large because it has been extrapolated from the 13 16 ns range out to 20 and 26 ns.

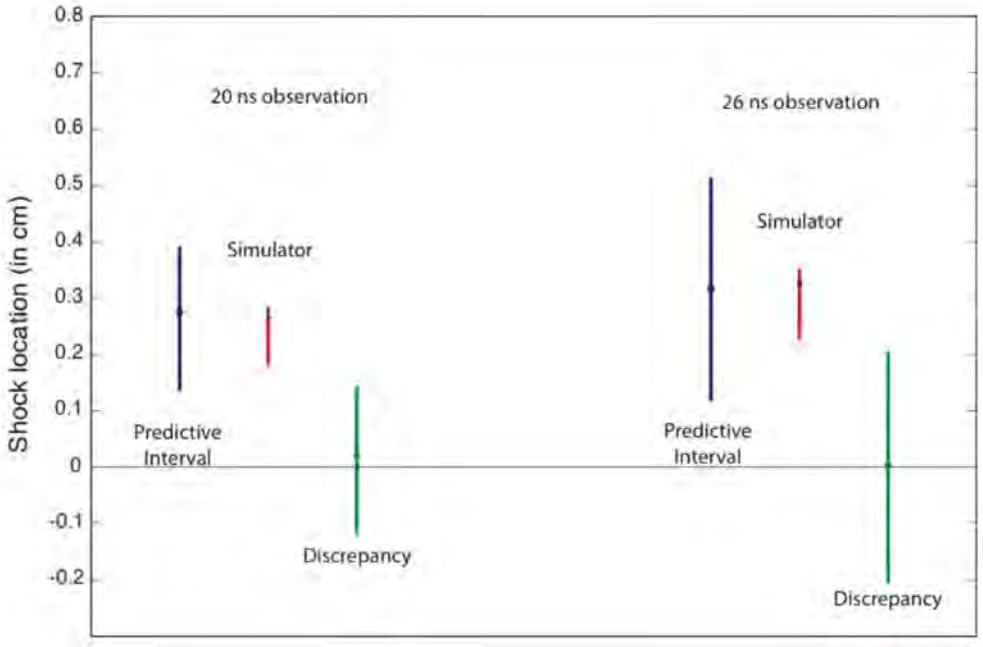

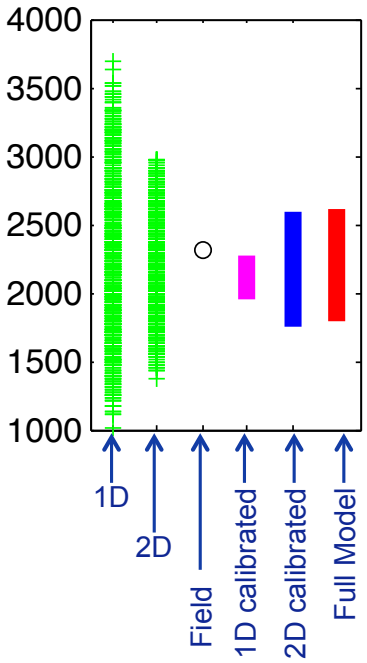

Figure 160: Prediction of shock location based on combined 1D and 2D CRASH simulations.

informed largely (but not entirely) by the 1D code. Similarly $y_{s2,j}$ are a set of N_2 values from runs of 2D CRASH (RS4), using experimental parameter values x_j, phenomenological parameters values that are unique to the 2D code $p_{2,j}$, and values of parameters that should be shared across the two codes of $p_{f,j}$. The value θ_1 represents the calibrated values of the parameters that are unique to the 1D code, and δ_1 is a discrepancy function that captures the difference between the 1D and 2D codes. While θ_1 is a parameter value, it is uncertain (being determined by a Bayesian inverse process) and so in fact is sampled from a (joint) posterior of various parameters. Finally, $y_{f,k}$ are N_f values of the QOI from field measurements. These are predicted by the emulator and discrepancy $\eta + \delta$ with calibrated parameter values θ_f, θ_2 and θ_1, along with a discrepancy δ_2 that is informed largely by the systematic difference between the model and realty, as well as a noise term ε that captures replication error in the experimental values (due to uncontrolled and even unknown experimental conditions).

Note that in this model structure a non-zero discrepancy δ_1 is almost guaranteed, because the 1D and 2D code cannot give the same results (unless the problem really is 1D).

Figure 160 shows a prediction from this combined 1D and 2D model, for shock location. The green crosses show the simulation values, while the circle is a field measurement of shock location (from the 2008 experimental campaign). This specific measurement was left out of the statistical analysis; the other measurements from that campaign were included. The pink bar shows the 95% prediction interval from the calibrated 1D model ($\eta(x, \theta_1, \theta_f)$), while the blue bar is the 95% prediction interval from calibrated the 2D model. Finally, the red band is the full prediction $\eta + \delta_1 + \delta + \varepsilon$.

The joint posterior distribution is also informative. The goal of this analysis combine experimental data on shock breakout time from 2009 experiments and on shock location at 13 ns from 2008 experiments with model results of

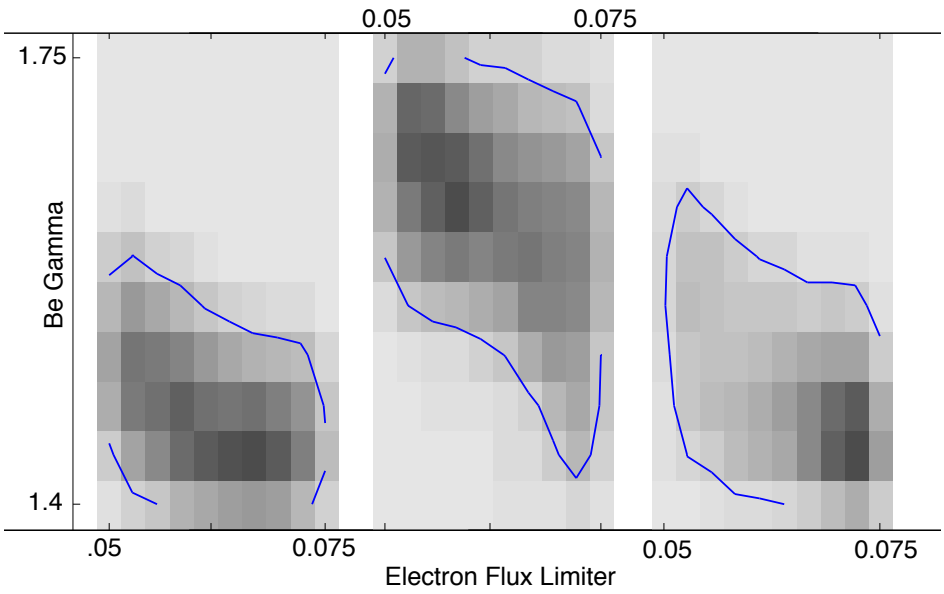

Figure 161: Joint posterior distributions of electron flux limiter and Be gamma. The calibration is based on (left) shock breakout time data only; (center) shock location data only; and (right) both breakout time and shock location.

those quantities, as well as model results at 26 ns, to do predictions of shock location at 26 ns of a future experiment. Figure 161 shows the joint posterior distribution for the electron flux limiter and Be gamma, based on which data are used for calibration. Note that the breakout time provides the most compelling evidence for these parameters (which is expected), but that the joint calibration still refines and tightens the posterior distribution for calibration.

This ability to combine multiple simulation models can be directly applied to combine gray and multi group simulations, so we can, for example, use many gray 2D simulations together with fewer multi-group 2D simulations. We are also extending this model to allow the combination of multiple models where none of the models can be considered the high fidelity model, in other words where there is no hierarchy or ranking of models. This is necessary so that we can, for example, combine 3D gray with 2D multi group models.

6.10 Prediction of fifth year shock location

6.10.1 Introduction

An emerging problem in many areas of scientific study is making, or improving, inferences on physical systems using deterministic computer simulators (Sacks et al., 1989; Santner et al., 2003). An important application of computer models is that of statistical model calibration (or simply *calibration*) where the aim is to combine simulator outputs with physical observations to build a predictive model and also estimate unknown parameters of interest given the real process (e.g., Kennedy and O'Hagan, 2001).

The central application of interest for the Center for Radiative Shock Hydrodynamics (CRASH) is radiative shocks. These are shock waves where radiation from shocked matter dominates the energy transport and results in a complex evolutionary structure. Radiative shocks involve many subprocesses, making them computationally challenging to

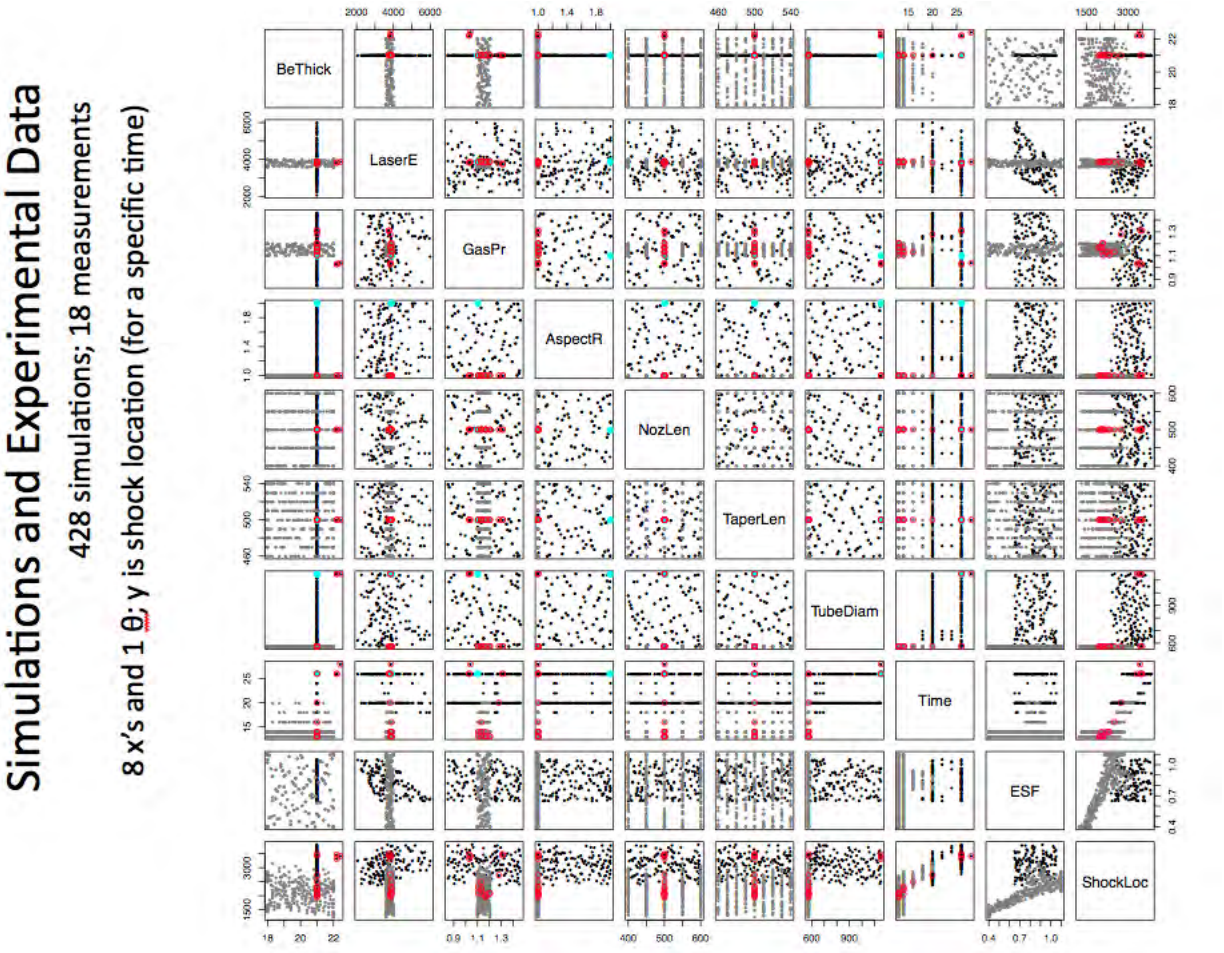

Figure 162: Scatterplot matrix of inputs for observations, simulations and extrapolations. Grey dots are the RS12 simulations, black dots are the RS13 simulations, red circles are the 2008 and 2010 experiments and the light blue circles are the locations of the fifth year experiments.

simulate. In addition, some inputs that govern the behaviour of the shocks are unknown - or at least imprecisely known. A main goal of CRASH is to use a limited number of complex computer simulations and physical observations to predict features of radiative shock experiments in an untested regime. That is, the aim is to calibrate the CRASH code and use the calibrated model to make predictions of features of radiative shocks that require extrapolation in several of the inputs.

6.10.2 Data

The statistical model calibration task utilizes simulations of radiative shock and experiments conducted at the Omega laser facility in 2008 and 2010. This section describes the inputs to the observations and CRASH simulations and the inputs.

The experiments that are used for calibration and prediction are those from the 2008 and 2010 experiments. All

of the 2008 experiments, and all but three of the 2010 experiments, were performed on circular shock tubes with a small diameter (in the area of 575 microns). Three of the 2010 experiments used circular tubes with a diameter of 1150 microns. The 2012, fifth year, experiments to be predicted have oval shaped tubes and nozzle at the front end of the tube. We had two observations with a nozzle on circular tubes, but none with an oval tube. Thus, there are no observations in the regime (oval tube) where the predictions are to be made. The defining feature that captures the shape of the tube is the *aspect ratio*.

The code runs arise from two separate simulation suites - run-sets 12 and 13 (RS12 and RS13), respectively. The inputs for each code and the ranges explored are shown in Table 1. Notice that several of the inputs are held constant in RS12 and RS13, respectively. Briefly, RS12 explores the input region for small, circular tubes used in the 2008 experiments. RS13, on the other hand, explores a similar input region, but also varies the tube geometry and nozzle geometry. The simulation region for RS13 contains the 2010 experiments with large diameter, circular tubes and also the fifth year oval tube experiments to be predicted. The energy flux limiter was held constant in RS13 because the QOIs were found to be relatively insensitive to this input in earlier studies, thus the only calibration parameter to be estimated is the energy scale factor. In addition, the thickness of the beryllium disk is also held constant in RS13 because improvements in manufacturing meant that the disk thickness had very little variability in the experiments.

Design variables		
Input	**RS12**	**RS13**
Be thickness (microns)	[18,22]	21
Laser energy (J)	[3600,3990]	
Effective laser energy (J)*		[2156.4,4060]
Xe fill pressure (atm)	[1.100,1.2032]	[0.852,1.46]
Tube diameter (microns)	575	[575,1150]
Taper length (microns)	500	[460,540]
Nozzle length (microns)	500	[400,600]
Aspect ratio (microns)	1	[1,2]
Calibration parameters		
Input	**RS12**	**RS13**
Electron flux limiter	[0.04, 0.10]	0.06
Energy scale factor	[0.40,1.10]	[0.60,1.00]

Table 19: Input ranges for RS12 and RS13. A single value means that the variable was constant for all simulation runs. * The effective laser energy is the laser energy $\times$ energy scale factor.

The inputs for a sample of 428 simulations from RS12 and RS13, the 2008 and 2010 experiments and the fifth year experiments are shown in Figure 162. There are a few observations that one should take away from this plot. The different input regions RS12 and RS13 are clearly shown. In addition, it is quite noticeable that the 2008 and 2010

experiments (red circles) fall within the input region covered by the simulations. However, the fifth year experiment (light blue circles) are clearly an extrapolation in the aspect ratio. That is, the aspect ratio of the experiments is equal to 1, but the fifth year experiments have an aspect ratio of two.

6.10.3 Prediction Strategy

A Bayesian hierarchical model is used to combine observations and simulations to make predictions with associated uncertainty. The basis for this approach is found in Kennedy and O'Hagan (2001) and Higdon et al. [2004]. The components of the prediction strategy are outlined in the previous quarterly report.

The approach that we use writes the observations as the sum of the simulator output, a discrepancy term and observation error. That is,

$$y(\mathbf{x}) = \eta(\mathbf{x}, \boldsymbol{\theta}) + \delta(\mathbf{x}) + \varepsilon, \tag{254}$$

where $y(\mathbf{x})$ denotes the quantity of interest (QOI) at design variable setting $\mathbf{x}$, $\eta(\mathbf{x}, \boldsymbol{\theta})$ is the simulator output at input settings $(\mathbf{x}, \boldsymbol{\theta})$, $\delta(\mathbf{x})$ is the discrepancy that accounts for systematic differences between the code output and the system mean, and ε is observation error. Note that a CRASH code run takes in inputs $\mathbf{t}$ for the calibration parameter. However, the value of the calibration parameter in the field, $\boldsymbol{\theta}$, is unknown and is estimated as part of fitting the Bayesian hierarchical model. Furthermore, η, δ, $\boldsymbol{\theta}$ and the observation error are estimated and the resulting uncertainties included in the predictions.

The nominal settings for the radiative shock experiments to be predicted are listed in Table 20. From past experiments, we have found that the nominal values of the design variables are not always achieved in practice. Consequently, Table 20 also includes the distribution of several of the inputs. We will propagate these distributions through the predictive model to account for the uncertainty in the inputs.

The strategy used to make predictions attempts to account for all known sources of uncertainty. These include variability in the nominal settings of the system, variability in unknown parameters, uncertainty in predicting the code output at unsampled inputs, the systematic discrepancy between the model response and the physical system, and also observation error. To do so, there are four main components to the approach used to obtain a sample from the posterior predictive distribution of the QOIs: (i) sample the input settings, $\mathbf{x}$, from the distributions specified in Table 20 and also the parameters for the statistical model; (ii) use the sampled parameters to estimate the simulator output, $\eta(\mathbf{x}, \boldsymbol{\theta})$; (iii) use the sampled parameters to estimate the discrepancy, $\delta(\mathbf{x})$; and adjust the prediction; and (iii) sample a value for the observation error add it to the prediction in step (iii). These steps are repeated many times to obtain a sample from the posterior predictive distribution of $y(\mathbf{x})$ at the nominal settings.

6.10.3.1 Prediction of shock location in the training experiments

Before moving on to the predictions of the fifth year experiments, we detour slightly to try to gain some insight from the calibration model. A main component of the statistical model is estimating the code output at un-sampled inputs. That is, an emulator of η is found. While it is not possible to visualize the emulator as a function of all of the inputs simultaneously, we can integrate out all but one of the inputs and plot the averaged response over the remaining input as a form of sensitivity analysis.

Figure 163 shows the sensitivity plots for each of the inputs. Interestingly, the output appears to be fairly insensitive to several of the inputs representing the tube geometry (i.e., nozzle length, taper length and, to a lesser degree, aspect

Design variables		
Input	**Nominal value**	**Distribution**
Be thickness (microns)	21	Uniform (20.5, 21.5)
Laser energy (J)	3800	Normal (3800, 81.64)
Xe fill pressure (atm)	1.15	Normal(1.15, 0.10)
Tube diameter (microns)	1150	
Taper length (microns)	500	
Nozzle length (microns)	500	
Aspect ratio (microns)	2	
Time (ns)	26	

Table 20: Settings and distributions for the design variables in the 2012 experiments. The Be thickness is uniform over the specified range and the Laser energy and Xe fill pressure are both normal with the specific mean and standard deviation.

ratio). The final geometric input, tube diameter, seems to be quite active. A quick glance at Figure 1 (seventh row) shows that there are observations from previous experiments with a large tube diameter. Thus when making the predictions for the fifth year experiment, we will not be extrapolating in this input. The model output appears quite sensitive to the remaining outputs where no extrapolation is required.

The statistical model will be used in the next section to make predictions of the fifth year experiment. We use different formulations for the discrepancy model in our predictions:

(i) Predictions are made using the Kennedy and O'Hagan (2001) model using simulations (RS12 and RS13) and observations (2008 and 2010 experiments).

(ii) Predictions are made using simulations (RS12 and RS13) and observations (2008 and 2010 experiments), but the discrepancy is specified as a random effect with a variance that grows linearly in time.

To evaluate how the statistical model performs, the 2008 and 2010 experiments are predicted. Figure 164 shows a sample of predicted trajectories at a random set of inputs (grey curves), 95% prediction intervals with and without discrepancies and observational error (black and green curves respectively) and the observations of shock location (black dots) from the 2008 and 2010 experiments. Notice that each of the observations lies within the 95% prediction intervals, lending some confidence to the predictive ability of the model.

6.10.3.2 Prediction of shock location in the fifth year experiments

The predictive models ((i) and (ii)) discussed in the previous subsection were used to predict shock location for the fifth year experiments. The predictions and associated uncertainty intervals, as well as the observations, are shown in Figures 4 and 5 for shock location versus time. Both of the figures show a sample of predicted trajectories at a random

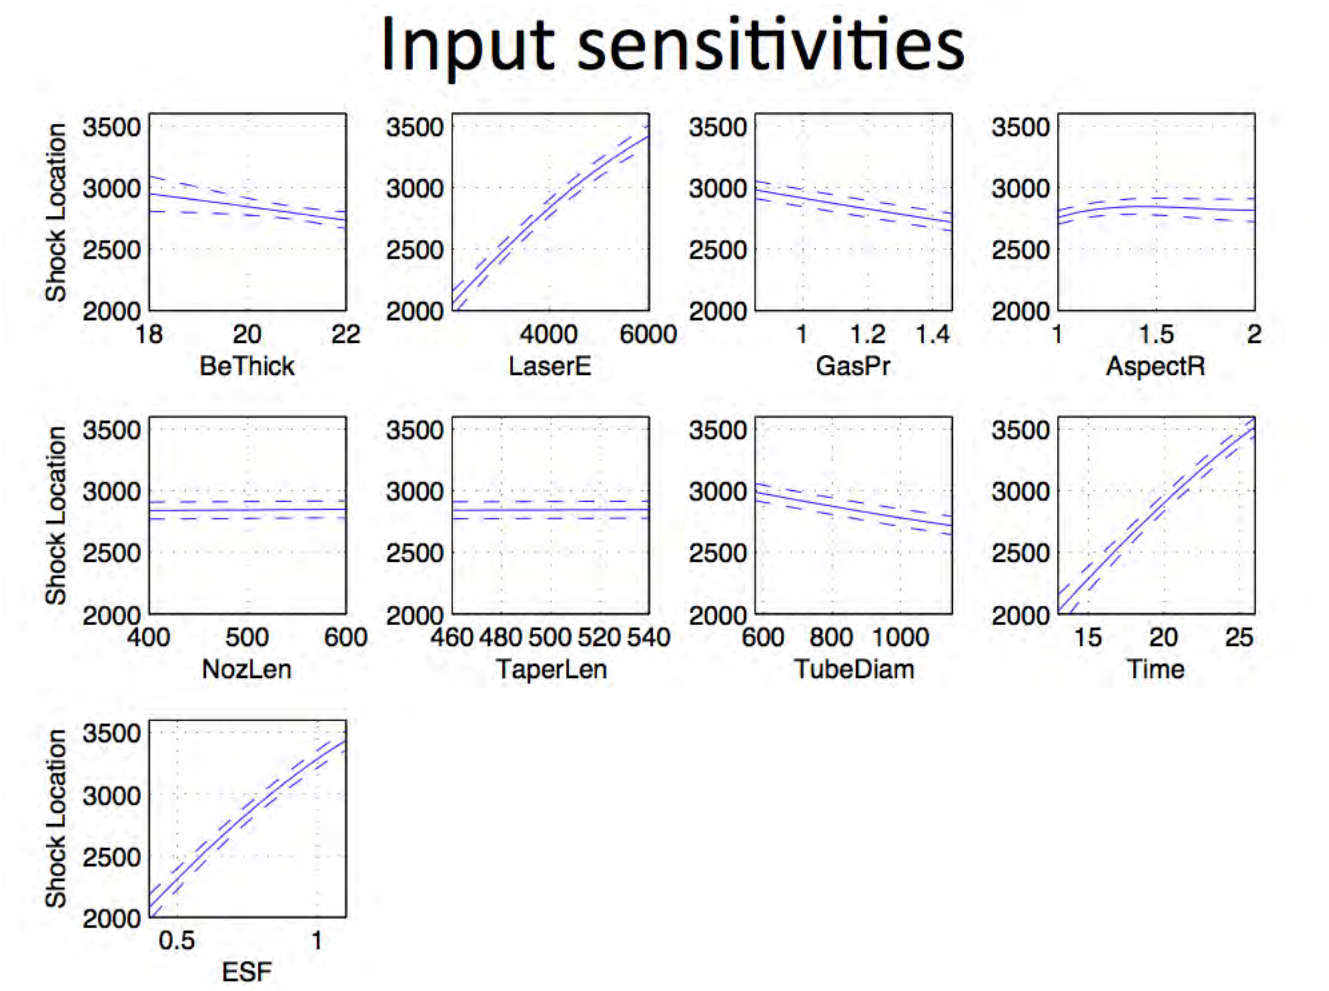

Figure 163: Sensitivity plots for each of the inputs to the computer model. Dashed lines shows 95% prediction intervals and the solid line is the median prediction.

set of inputs (grey curves), 95% prediction intervals with and without discrepancies and observational error (black and green curves respectively) and the observations of shock location from the fifth year experiments.

A quick glance at the figures reveals that the two prediction strategies yield qualitatively similar results, though the uncertainty is slightly larger for latter approach (ii). Furthermore, in both cases, the fifth year shock locations are slightly under-predicted. That is, the location of the the observed shocks are further down the tube than is predicted by the statistical model. Put another way, the majority of the shocks locations lie above the prediction intervals specified by the statistical models. Note, however, that the distribution of predicted locations does overlap with the distribution of observed location, thus achieving the stated goal for the project.

At this point, the reason for the under-prediction of the shock is not known. Looking at Figure 162, we see that the variable with the largest degree of extrapolation in the prediction is the aspect ratio. Furthermore, Figure 163 indicated that the impact of aspect ratio on the system output was limited. From this vantage point, nothing in the experiments initially indicated that the model was likely to under-predict the shock location. One possibility is that this difference

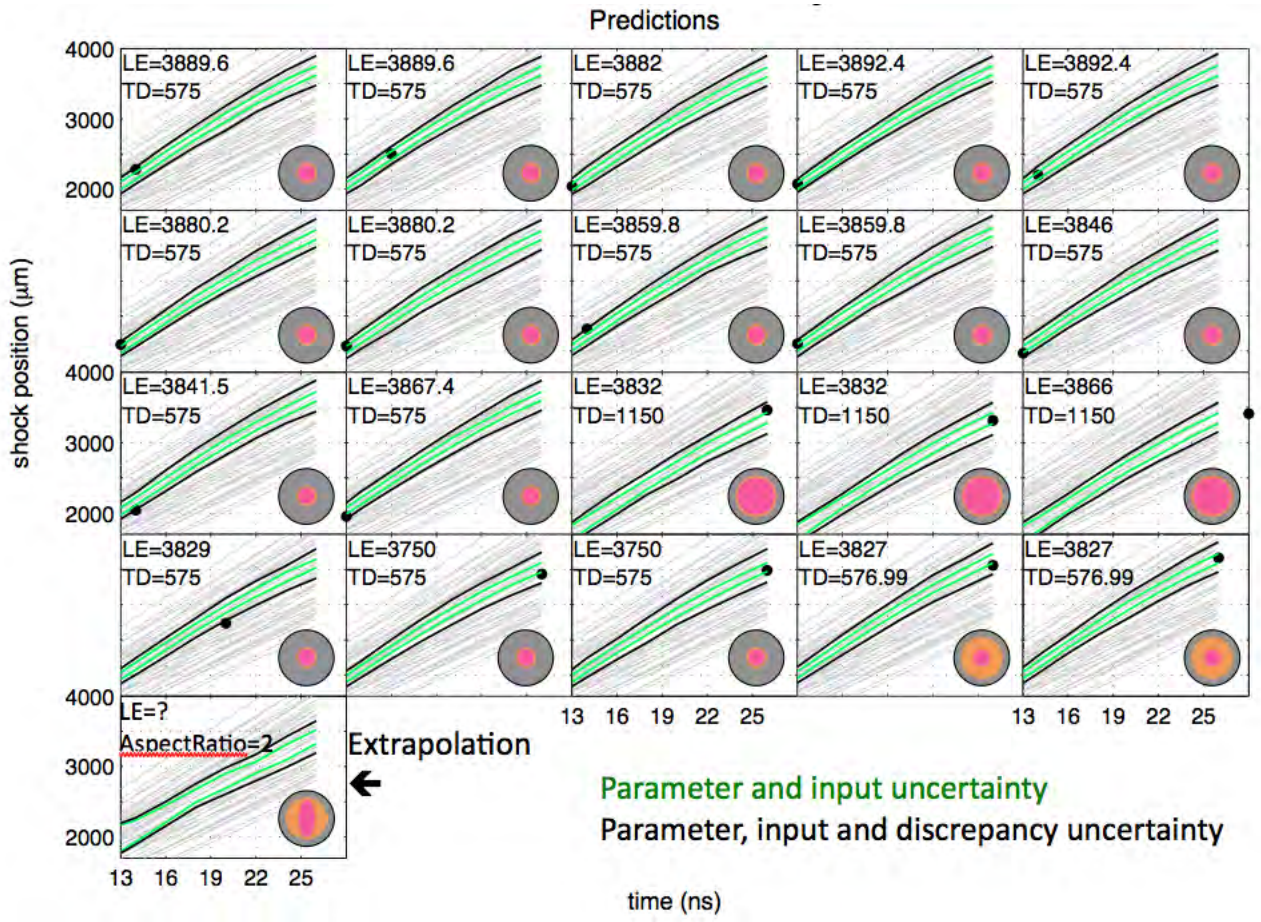

Figure 164: Posterior prediction intervals at the nominal settings for shock location versus time viewed along the minor axis using prediction strategy (ii). Grey curves show a sample of predicted trajectories at different input settings, green curves are 95% prediction intervals incorporating parameter and input uncertainty, black curves form 95% prediction intervals incorporating parameter and input uncertainty, discrepancy variation and observational uncertainty and the circles are the observations from the 2008 and 2010 experiments.

arises from lack of fidelity in the physics in the model, as follows. The diffusion model for radiation transport used in CRASH is well known to transport too much energy laterally. In the present context this would be expected to produce more wall blowoff and a greater displacement of the wall shock in the simulation results than in the physical data. Indeed there are indications of this (see Sec. 6.2.2). The more-extensive wall shock might reduce the amount of momentum that can be delivered through the nozzle, thus reducing the momentum that carries the shock forward at late times.

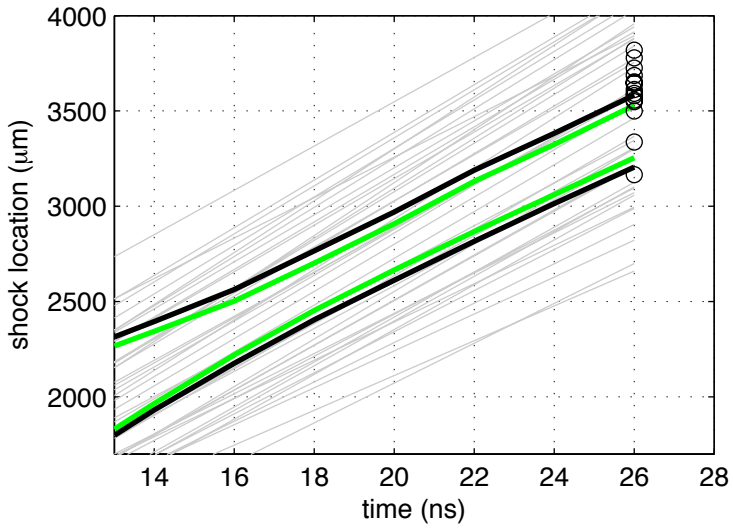

Figure 165: Posterior prediction intervals at the nominal settings for shock location versus time viewed along the minor axis using prediction strategy (i). Grey curves show a sample of predicted trajectories at different input settings, green curves are 95% prediction intervals incorporating parameter and input uncertainty, black curves form 95% prediction intervals incorporating parameter and input uncertainty, discrepancy variation and observational uncertainty and the circles are the observations from the fifth year experiments.

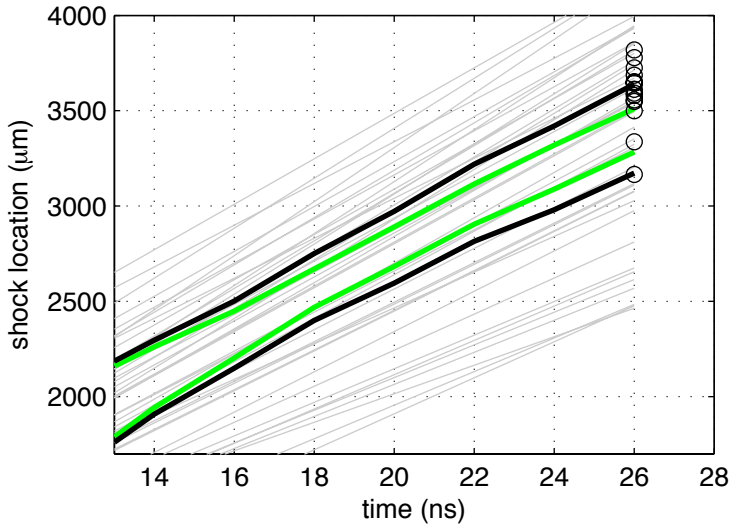

Figure 166: Posterior prediction intervals at the nominal settings for shock location versus time viewed along the minor axis using prediction strategy (ii). Grey curves show a sample of predicted trajectories at different input settings, green curves are 95% prediction intervals incorporating parameter and input uncertainty, black curves form 95% prediction intervals incorporating parameter and input uncertainty, discrepancy variation and observational uncertainty and the circles are the observations from the fifth year experiments.

References

Nist atomic spectra database. http://physics.nist.gov/PhysRefData/ASD/index.html, 2009.

Spectr-w3 online database on spectral properties of atoms and ions. http://spectr-w3.snz.ru, 2009.

A.Bar-Shalom, J.Oreg, W.H.Goldstein, D.Shvarts, and A.Zigler. Super-transition-arrays: A model for the spectral analysis of hot, dense plasma. *Physical Review A*, 40(6), 1989.

John F. Ahearne, Marvin Adams, John Cornwall, Douglass Eardley, B. John Garrick, Richard Garwin, Sydell P. Gold, Yogendra Gupta, David Hammer, Theodore Hardebeck, John Kammerdiener, Sallie Keller-McNulty, Ernest J. Moniz, Michael Ortiz, Jerry Paul, Robert Rosner, and Robert Selden. *Evaluation of Quantification of Margins and Uncertainties Methodology for Assessing and Certifying the Nuclear Stockpile*. National Academies Press, 2008.

G. Arfken. *Mathematical Methods for Physicists, 3rd edition*. Academic Press, San Diego, CA, 1985.

A. Bar-Shalom, J. Oreg, W.H. Goldstein, D. Shvarts, and A.‘ Zigler. Super-transition-arrays: A model for the spectral analysis of hot dense plasma. *Phys. Rev. A.*, 40:3183–3193, 1989.

A. Bar-Shalom, J. Oreg, M. Klapisch, and T. Lehecka. Effect of configuration interaction on shift, widths, and intensity redistribution of transition arrays. *Phys. Rev. E*, 59:3512–3525, 1999.

L.M. Barker and R.E. Hollenback. Laser interferometer for measuring high velocities of any reflecting surface. *J. Appl. Phys.*, 43(11):4669, 1972.

L. Benkevitch, I. Sokolov, D. Oberoi, and T. Zurbuchen. Algorithm for tracing radio rays in solar corona and chromosphere. arXiv:1006.5635v3, 2010.

V. B. Berestetskii, E. M. Lifshitz, and L. P. Pitaevskii. *Quantum Electrodynamics*. Pergamon, 1982.

M.J. Berger and P. Colella. Local adaptive mesh refinement for shock hydrodynamics. *Journal of Computational Physics*, 82, 1989.

Julian Besag, Peter Green, David Higdon, and Kerrie Mengersen. Bayesian computation and stochastic systems. *Statistical Science*, 10:3, 1995.

T. Boehly. personal communication.

T.R. Boehly, R.S. Craxton, T.H. Hinterman, J.H. Kelly, T.J. Kessler, S.A. Kumpman, S.A. Letzring, R.L. McCrory, S.F.B. Morse, W. Seka, S. Skupsky, J.M. Soures, and C.P. Verdon. The upgrade to the OMEGA laser system. *Rev. Sci. Intsr.*, 66(1):508–510, 1995.

J.P. Boris. Relativistic plasma simulation optimization of a hybrid code. In *Proc. 4th Conf. Num. Sim. Plasmas*, 1971.

J.R. Buchler. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 30, 1983.

M. Busquet, M. Klapisch amd A. Bar-Shalom, and J. Oreg. Spectroscopic databases for warm xenon and iron in astrophysics and laboratory astrophysics conditions. *Bull. Am. Phys. Soc.*, 55, 2010.

T.A. Carlson, C.W. Nestor, N. Wasserman Jr., and J.D. McDowell. Calculated ionization potentials for multiply charged ions. *At. Data*, 2, 1970.

J.I. Castor. *Radiation Hydrodynamics*. Cambridge University Press, 2004.

Arthur N. Cox, editor. *Allen's Astrophysical Quantities*. Springer, 2000.

D. G. T. Denison, B. K. Mallick, and A. F. M. Smith. Bayesian mars. *Statistics and Computing*, 8(4):337–346, 1998. ISSN 0960-3174. doi: 10.1023/a:1008824606259. URL <GotoISI>://WOS:000080385000004http://link.springer.com/article/10.1023%2FA%3A1008824606259. Times Cited: 40 Denison, DGT Mallick, BK Smith, AFM 40.

F.W. Doss. *Structure in Radiative Shock Experiemnts*. Thesis, University of Michigan, 2011.

F.W. Doss, H.F. Robey, R.P. Drake, and C.C. Kuranz. Wall shocks in high-energy-density shock tube experiments. *Physics Of Plasmas*, 16:112705, 2009.

F.W. Doss, R.P. Drake, and C.C. Kuranz. Repeatability in radiative shock tube experiments. *High Energy Density Physics*, 6(2):157–161, 2010. doi: doi:10.1016/j.hedp.2009.12.007.

F.W. Doss, R.P. Drake, and E.S. Myra. Oblique radiative shocks, including their interactions with non-radiative polytropic shocks. *Physics of Plasmas*, 18:056901, 2011a.

F.W. Doss, C. C. Kuranz, and R. P. Drake. Statistical inference in the presence of an inclination effect in laboratory radiative shock experiments. *Astrophysics and Space Science*, 336(1):219–224, 2011b. doi: 10.1007/s10509-010-0579-4.

R. P. Drake. *High-Energy-Density Physics: Fundamentals, Inertial Fusion, and Experimental Astrophysics*. Springer-Verlag, Berlin, 2006.

R.P. Drake, F.W. Doss, R.G. McClarren, M.L. Adams, N. Amato, D. Bingham, C.C. Chou, C. DiStefano, K. Fidkowsky, B. Fryxell, T.I. Gombosi, M.J. Grosskopf, J.P. Holloway, B. van der Holst, C.M. Huntington, S. Karni, C.M. Krauland, C.C. Kuranz, E. Larsen, B. van Leer, B. Mallick, D. Marion, W. Martin, J.E. Morel, E.S. Myra, V. Nair, K.G. Powell, L. Raushberger, P. Roe, E. Rutter, I.V. Sokolov, Q. Stout, B.R. Torralva, G. Toth, K. Thornton, and A.J. Visco. Radiative effects in radiative shocks in shock tubes. *High Energy Density Physics*, 7:130–140, 2011.

M. G. Edwards. Elimination of adaptive grid interface errors in the discrete cell centered pressure equation. *Journal of Computational Physics*, 126:356, 1996a.

M.G. Edwards. *Journal of Computational Physics*, 126, 1996b.

B. Einfeldt, C.D. Munz, P.L. Roe, and B. & Sjögreen. *Journal of Computational Physics*, 92, 1991.

S.C. Eisenstat. *SIAM Journal of Scientific and Statistical Computing*, 2, 1981.

R.D. Falgout and U.M. Yang. Hypre: a library of high performance preconditioners. In P.M.A. Sloot, C.J.K. Tan, J.J. Dongarra, and A.G. Hoekstra, editors, *Computational Science - ICCS 2002 Part III*. Springer Verlag, 2002.

J. H. Friedman. Multivariate adaptive regression splines. *Annals of Statistics*, 19(1):1–67, 1991. ISSN 0090-5364. doi: 10.1214/aos/1176347963. URL <GotoISI>://WOS:A1991FF04700001http://projecteuclid.org/euclid.aos/1176347963. Times Cited: 1835 FRIEDMAN, JH 1873.

D. Gilles and et al. Comparison of fe and ni opacity calculations for a better understanding of pulsating stellar envelopes. *High Energy Density Phys.*, 7, 2011.

V.L. Ginzburg. *The Propagation of Electromagnetic Waves in Plasmas*. Pergamon Press, 1964.

M. et al. Gittings. The rage radiation-hydrodynamic code. *Computational Science and Discovery*, 1, 2008.

S.K. Godunov. A difference method for numerical calculation of discontinuous solutions of the equations of hydrodynamics. *Mat. Sbornik*, 47, 1959.

A. Harten, P.D. Lax, and B. Van Leer. *SIAM Review*, 25(1), 1983.

Dave Higdon, Marc Kennedy, James C. Cavendish, John A. Cafeo, and Robert D. Ryne. Combining field data and computer simulations for calibration and prediction. *SIAM Journal of Scientific Computing*, 26(2):448, 2004.

J.P. Holloway, D. Bingham, C.C. Chou, F.W. Doss, R.P. Drake, B. Fryxell, M.J. Grosskopf, B. van der Holst, B. Mallick, R.G. McClarren, A. Mukherjee, V. Nair, K.G. Powell, D. Ryu, I. Sokolov, G. Toth, and Z. Zhang. Predictive modeling of a radiative shock system. *Reliability Engineering and System Safety*, 96(9):1184–1193, 2011. doi: doi:10.1016/j.ress.2010.08.011.

C. C. Holmes, D. G. T. Denison, and B. K. Mallick. Accounting for model uncertainty in seemingly unrelated regressions. *Journal of Computational and Graphical Statistics*, 11(3):533–551, 2002. ISSN 1061-8600. doi: 10.1198/106186002475. URL <GotoISI>://WOS:000178313500003. Times Cited: 10 Holmes, CC Denison, DGT Mallick, BK 10.

J.J.MacFarlane. Ionmix – a code for computing the equation of state and radiative properties of lte and non-lte plasmas. *Computer Physics Communications*, 56, 1989.

S. Karni. *SIAM Journal on Scientific Computing*, 17, 1996.

Marc C. Kennedy and Anthony O'Hagen. Bayesian calibration of computer codes. *Journal of the Royal Statistical Society B*, 63:425, 2001.

B. Koren. *A robust upwind discretization method for advection, diffusion and source terms*. Vieweg, 1993.

G. Kreiss and E. Olsson. *Journal of Computational Physics*, 210, 2005.

W.L. Kruer. *The physics of Laser-Plasma Interactions*. Westview Press, 2001.

C. C. Kuranz, R. P. Drake, M. J. Grosskopf, B. Fryxell, A. Budde, A. R. Miles, J. Knauer, T. Plewa, and N. Hearn. Spike morphology in Rayleigh-Taylor, decelerating interface experiments. *Phys. Plasmas*, 17(052709), 2010.

C. C. Kuranz, R. P. Drake, C. M. Krauland, D. C. Marion, M. J. Grosskopf, E. Rutter, B. Torralva, J. P. Holloway, D. Bingham, J. Goh, T. R. Boehly, and A. T. Sorce. Initial conditions of radiative shock experiments. *Physics Of Plasmas*, 20(5), 2013. ISSN 1070-664X. doi: 05632110.1063/1.4805021.

L.D. Landau and E.M. Lifshitz. *Theoretical Physics, Vol. 5. Statistical Physics, Part 1.* Pergamon Press, 3rd edition edition, 1980.

J. T. Larsen and S. M. Lane. Hyades: a plasma hydrodynamics code for dense plasma studies. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 51:179, 1994.

C.D. Levermore and G.C. Pomraning. *Astrophysical Journal*, 248, 1981.

D.A. Liberman. Self-consistent field model for condensed matter. *Phys. Rev. B*, 20, 1979.

R. B. Lowrie and J. D. Edwards. Radiative shock solutions with grey nonequilibrium diffusion. *Shock Waves*, 18(2): 129–143, 2008. ISSN 0938-1287. doi: 10.1007/s00193-008-0143-0. URL <GotoISI>://000257383600006. Lowrie, Robert B. Edwards, Jarrod D.

C.M. Lund and Wilson J.R. Lawrence Livermore National Laboratory Report UCRL-84678, 1980.

N. H. Magee, Jr. Abdallah, J., R. E. H. Clark, J. S. Cohen, L. A. Collins, G. Csanak, C. J. Fontes, A. Gauger, J. J. Keady, D. P. Kilcrease, and A. L. Merts. Atomic structure calculations and new los alamos astrophysical opacities. In S.J. Adelman and W.L. Wiese, editors, *Astrophysical Applications of Powerful New Databases*, volume 78, page 51, San Francisco, California, 1995. Astronomical Society of the Pacific.

W. C. Martin and R. Zalubas. Energy levels of aluminum, al i through al xiii. *J. Phys. Chem. Ref. Data*, 8, 1979.

Ryan G. McClarren, D. Ryu, R. Paul Drake, Michael Grosskopf, Chuan-Chih Chou, Bruce Fryxell, and James Paul Holloway. A physics informed emulator for laser-driven radiating shock simulations. *Reliability Engineering and System Safety*, 96:1194–1207, 2011.

D. Mihalas and B. Weibel-Mihalas. *Foundations of Radiation Hydrodynamics*. Oxford Univ. Press, 1984.

J. E. Miller, T. R. Boehly, A. Melchior, D. D. Meyerhofer, P. M. Celliers, J. H. Eggert, D. G. Hicks, C. M. Sorce, J. A. Oertel, and P. M. Emmel. Streaked optical pyrometer system for laser-driven shock-wave experiments on OMEGA. *Review of Scientific Instruments*, 78(3), MAR 2007. ISSN 0034-6748. doi: {10.1063/1.2712189}.

G.N. Minerbo. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 20, 1978.

J. E. Morel. Diffusion-limit asymptotics of the transport equation, the $p_{1/3}$ equations, and two flux-limited diffusion theories,. *Journal of Quantative Spectroscopy and Radiative Transfer*, 65:769, 2000.

E. S. Myra and W. D. Hawkins. A comparison study of discrete-ordinates and flux-limited diffusion methods for modeling radiation transport. *High Energy Density Physics*, 9(1):91–102, 2013. ISSN 1574-1818. doi: 10.1016/j.hedp.2012.10.007.

J. E. Oakley and A. O'Hagan. Probabilistic sensitivity analysis of complex models: a bayesian approach. *Journal of the Royal Statistical Society*, Series B 66:751, 2004.

E. Olsson, G. Kreiss, and S. Zahedi. *Journal of Computational Physics*, 225, 2007.

G.C. Pomraming. *The equations of Radiation Hydrodynamics*. Dover, 2005.

G.C. Pomraning. *The Equations of Radiation Hydrodynamics*. Pergamon Press, 1973.

G.C. Pomraning. Flux limiters and eddington factors. *J. Quant. Spectrosc. Radiat. Transfer*, 9:517–530, 1982.

K.G. Powell, P.L. Roe, T.J. Linde, T.I. Gombosi, and D.L. DeZeeuw. *Journal of Computational Physics*, 154(2), 1999.

A. B. Reighard, R. P. Drake, J. E. Mucino, J. P. Knauer, and M. Busquet. Planar radiative shock experiments and their comparison to simulations. *Physics Of Plasmas*, 14:056504, 2007.

P. Reinecke and J. Meyer-ter Vehn. *Phys. Fluids A*, 3:1807, 1991.

P.L. Roe. Characteristic-based schemes for the euler equations. *Annual Review of Fluid Mechanics*, 18, 1986.

Y. Saad and M.H. Schultz. *SIAM Journal of Scientific and Statistical Computing*, 7(3), 1986.

E.B. Saloman. Energy levels and observed spectral lines of xenon, xe i through xe liv. *J. Phys. Chem. Ref. Data*, 33, 2004.

B. Su and G.L. Olson. *J. Quant. Spectr. Radiative Transfer*, 56:337, 1996.

M. Sussman and E.G. Puckett. *Journal of Computational Physics*, 162, 2000.

B. Tang. Orthogonal array-based latin hypercubes. *Journal of the American Statistical Association*, 88:1392, 1993.

G. Toth, I. V. Sokolov, T. I. Gombosi, D. R. Chesney, C. R. Clauer, D. L. De Zeeuw, K. C. Hansen, K. J. Kane, W. B. Manchester, R. C. Oehmke, K. G. Powell, A. J. Ridley, II Roussev, Q. F. Stout, O. Volberg, R. A. Wolf, S. Sazykin, A. Chan, B. Yu, and J. Kota. Space weather modeling framework: A new tool for the space science community. *Journal Of Geophysical Research-Space Physics*, 110(A12), 2005. URL <GotoISI>://000234507600001. 0148-0227 A12226.

G. Tóth, Y. Ma, and T.I. Gombosi. Hall magnetohydrodynamics on block-adaptive grids. *Journal of Computational Physics*, 227, 2008.

G. Toth, B. van der Holst, I. V. Sokolov, D. L. De Zeeuw, T. I. Gombosi, F. Fang, W. B. Manchester, X. Meng, D. Najib, K. G. Powell, Q. F. Stout, A. Glocer, Y. J. Ma, and M. Opher. Adaptive numerical algorithms in space weather modeling. *J. Comp. Physics*, 231(3), 2012.

B. van der Holst, W.B. Manchester IV, R.A. Frazin, A.M. Vásquez, G. Tóth, and T.I. Gombosi. A data-driven, two-temperature solar wind model with alfvén waves. *The Astrophysical Journal*, 725(1), 2010.

B. van der Holst, G. Toth, I.V. Sokolov, K. G. Powell, J. P. Holloway, E.S. Myra, Q. Stout, M.L. Adams, J. E. Morel, and R. P. Drake. A block-adaptive-mesh code for radiative shock hydrodynamics: Implementation and verification. *Astrophysical Journal Supplement Series*, 194:23, 2011.

B. van der Holst, G. Toth, I.V. Sokolov, L.K.S. Daldorff, K.G. Powell, and R.P. Drake. Simulating radiative shocks in nozzle shock tubes. *High Energy Density Physics*, 8:161–169, 2012.

B. van der Holst, G. Toth, I. V. Sokolov, B. R. Torralva, K. G. Powell, R. P. Drake, M. Klapisch, M. Busquet, B. Fryxell, and E. S. Myra. Simulating radiative shocks with the crash laser package. *High Energy Density Physics*, 9(1):8–16, 2013. ISSN 1574-1818. doi: 10.1016/j.hedp.2012.09.011.

H.A. Van der Vorst. Bi-cgstab: A fast and smoothly converging variant of bi-cg for the solution of nonsymmetric linear systems. *SIAM Journal of Scientific and Statistical Computing*, 13, 1992.

B. van Leer. Towards the ultimate conservative difference scheme. II. Monoticity and conservation combined in a second-order scheme. *Journal of Computational Physics*, 14:361, 1974.

B. van Leer. Towards the ultimate conservative difference scheme. V. A second-order sequel to Godunov's method. *Journal of Computational Physics*, 32:101, 1979.

H.C. Yee. NASA TM-101088, 1989.

Y.B. Zeldovich and Y.P. Raizer. *Physics of Shock Waves and High-Temperature Phenomena*. Academic Press, 1966.

Appendix A: CRASH Publications

A. Budde, R. P. Drake, C. C. Kuranz, M.J. Grosskopf, T. Plewa, and N.C. Hearn. Simulation of fabrication variations in supernova hydrodynamics experiments. *High Energy Density Physics*, 6(2):135–142, 2010. doi: doi:10.1016/j.hedp.2010.01.011.

A. Chakraborty, B.K. Mallick, R.G. McClarren, C.C. Kuranz, M.J. Grosskopf, E. Rutter, H.F. Stripling, and R.P. Drake. Spline-based emulators for radiative shock experiments with measurement error. *J. Amer. Stat. Assoc.*, 108(502):411–428, 2013. doi: 10.1080/01621459.2013.770688.

M.T.H. de Frahan and E. Johnsen. Interface capturing for multifluid simulations with shocks using discontinuous galerkin approaches. *J. Comp. Phys.*, page submitted, 2013.

C. A. Di Stefano, C. C. Kuranz, P. A. Keiter, S. R. Klein, D. C. Marion, and R. P. Drake. Late-time breakup of laser-driven hydrodynamics experiments. *High Energy Density Physics*, 8(4):360–365, 2012. ISSN 1574-1818. doi: 10.1016/j.hedp.2012.09.006. URL `<Go to ISI>://WOS:000314395600008`.

F.W. Doss. *Structure in Radiative Shock Experiemnts*. Thesis, University of Michigan, 2011.

F.W. Doss, H.F. Robey, R.P. Drake, and C.C. Kuranz. Wall shocks in high-energy-density shock tube experiments. *Physics Of Plasmas*, 16:112705, 2009.

F.W. Doss, R.P. Drake, and C.C. Kuranz. Repeatability in radiative shock tube experiments. *High Energy Density Physics*, 6(2):157–161, 2010. doi: doi:10.1016/j.hedp.2009.12.007.

F.W. Doss, R.P. Drake, and E.S. Myra. Oblique radiative shocks, including their interactions with non-radiative polytropic shocks. *Physics of Plasmas*, 18:056901, 2011a.

F.W. Doss, C. C. Kuranz, and R. P. Drake. Statistical inference in the presence of an inclination effect in laboratory radiative shock experiments. *Astrophysics and Space Science*, 336(1):219–224, 2011b. doi: 10.1007/s10509-010-0579-4.

R. P. Drake. Isothermal, mass-limited rarefactions in planar and spherical geometry. *Phys. Plasmas*, 18 (September):104506, 2011. doi: 10.1063/1.3642612.

R.P. Drake. Perspective on high-energy-density physics. *Physics Of Plasmas*, 16:055501, 2009. doi: 10.1063/1.3078101.

R.P. Drake. High-energy-density physics. *Physics Today*, June:28–33, 2010.

R.P. Drake. Spike penetration in blast-wave-driven instabilities. *Astrophysical Journal*, 744(2):184–195, 2012.

R.P. Drake and G. Gregori. Design considerations for unmagnetized collisionless shock experiments in homologous flows. *Astrophysical Journal*, 749:171, 2012.

R.P. Drake, F.W. Doss, R.G. McClarren, M.L. Adams, N. Amato, D. Bingham, C.C. Chou, C. DiStefano, K. Fidkowsky, B. Fryxell, T.I. Gombosi, M.J. Grosskopf, J.P. Holloway, B. van der Holst, C.M. Huntington, S. Karni, C.M. Krauland, C.C. Kuranz, E. Larsen, B. van Leer, B. Mallick, D. Marion, W. Martin, J.E. Morel, E.S. Myra, V. Nair, K.G. Powell, L. Raushberger, P. Roe, E. Rutter, I.V. Sokolov, Q. Stout, B.R. Torralva, G. Toth, K. Thornton, and A.J. Visco. Radiative effects in radiative shocks in shock tubes. *High Energy Density Physics*, 7:130–140, 2011.

M. Fatenejad, A. R. Bell, A. Benuzzi-Mounaix, R. Crowston, R. P. Drake, N. Flocke, G. Gregori, M. Koenig, C. Krauland, D. Lamb, D. Lee, J. R. Marques, J. Meinecke, F. Miniati, C. D. Murphy, H. S. Park, A. Pelka, A. Ravasio, B. Remington, B. Reville, A. Scopatz, P. Tzeferacos, K. Weide, N. Woolsey, R. Young, and R. Yurchak. Modeling hedla magnetic field generation experiments on laser facilities. *High Energy Density Physics*, 9(1):172–177, 2013. ISSN 1574-1818. doi: 10.1016/j.hedp.2012.11.002.

B. Fryxell, C. C. Kuranz, R. P. Drake, M.J. Grosskopf, A. Budde, T. Plewa, J. F. Hansen, A. R. Miles, and J. Knauer. The possible effects of magnetic fields on laser experiments of rayleigh-taylor instabilities. *High Energy Density Physics*, 6:162–165, 2010. doi: doi:10.1016/j.hedp.2010.01.008.

B. Fryxell, E. Rutter, and E.S. Myra. Simulations of laser experiments of radiative and non-radiative shocks. *High Energy Density Physics*, 8:141–149, 2012. doi: 10.1016/j.hedp.2011.12.002.

B. Fryxell, F.W. Doss, R.P. Drake, M.J. Grosskopf, C.C. Kuranz, D. Marion, E. Myra, K.G. Powell, E. Rutter, I.V. Sokolov, G. Toth, and B. van der Holst. *Experiments and Simulations of Radiative Shocks*. IAEA, 2013a.

B. Fryxell, R. P. Drake, C. C. Kuranz, and E. S. Myra. Experiments and simulations of radiative shocks. *Physica Scripta*, T155, 2013b. ISSN 0031-8949. doi: 014017 10.1088/0031-8949/2013/t155/014017.

E. J. Gamboa, C. M. Huntington, M.A. Trantham, P.A. Keiter, and R.P. Drake. Imaging x-ray thomson scattering spectrometer design and demonstration. *Review Of Scientific Instruments*, 83:10E108, 2012.

E. J. Gamboa, D. S. Montgomery, I. M. Hall, and R. P. Drake. Imaging X-ray crystal spectrometer for laser-produced plasmas. *JOURNAL OF INSTRUMENTATION*, 6, APR 2011. ISSN 1748-0221. doi: 10.1088/1748-0221/6/04/P04004.

J. Goh, D. Bingham, J.P. Holloway, F. Doss, C.C. Kuranz, M.J. Grosskopf, and E. Rutter. Computer model calibration and prediction using outputs from multi fidelity simulators. *Technometrics*, page in press, 2013.

M. J. Grosskopf, R. P. Drake, C. C. Kuranz, E. M. Rutter, J. S. Ross, N. L. Kugland, C. Plechaty, B. A. Remington, A. Spitkovsky, L. Gargate, G. Gregori, A. Bell, C. D. Murphy, J. Meinecke, B. Reville, Y. Sakawa, Y. Kuramitsu, H. Takabe, D. H. Froula, G. Fiksel, F. Miniati, M. Koenig, A. Ravasio, E. Liang, W. Fu, N. Woolsey, and H. S. Park. Simulation of laser-driven, ablated plasma flows in collisionless shock experiments on omega and the nif. *High Energy Density Physics*, 9(1):192–197, 2013a. ISSN 1574-1818. doi: 10.1016/j.hedp.2012.11.004.

M. J. Grosskopf, R. P. Drake, A. R. Miles, T. Plewa, and C. C. Kuranz. Modeling of aspheric, diverging hydrodynamic instability experiments on the national ignition facility. *High Energy Density Physics*, 9: 439–447, 2013b.

M. J. Grosskopf, D. C. Marion, R. P. Drake, C. C. Kuranz, F. W. Doss, A. J. Visco, C. M. Huntington, C. M. Krauland, C. A. Di Stefano, and E. C. Harding. TARGET FABRICATION AT THE UNIVERSITY OF MICHIGAN. *FUSION SCIENCE AND TECHNOLOGY*, 59(1):250–256, JAN 2011. ISSN 1536-1055. 19th Target Fabrication Meeting, Orlando, FL, FEB 21-26, 2010.

J.P. Holloway, D. Bingham, C.C. Chou, F.W. Doss, R.P. Drake, B. Fryxell, M.J. Grosskopf, B. van der Holst, B. Mallick, R.G. McClarren, A. Mukherjee, V. Nair, K.G. Powell, D. Ryu, I. Sokolov, G. Toth, and Z. Zhang. Predictive modeling of a radiative shock system. *Reliability Engineering and System Safety*, 96 (9):1184–1193, 2011. doi: doi:10.1016/j.ress.2010.08.011.

C.M. Huntington, C.C. Krauland, C. C. Kuranz, S.H. Glenzer, and R. P. Drake. Imaging scattered x-ray radiation for measurement of local electron density in high-energy-density experiments. *High Energy Density Physics*, 6:194–199, 2010a. doi: doi:10.1016/j.hedp.2010.01.009.

C.M. Huntington, C.M. Krauland, C.C. Kuranz, R.P. Drake, H.-S. Park, D.H. Kalantar, A.G. MacPhee, B.A. Remington, and J. Kline. Short-duration backlit pinhole radiography diagnostics on the national ignition facility. *Rev. Sci. Inst.*, 81(10):10E536, 2010b.

P. A. Keiter, E. J. Gamboa, C. M. Huntington, and C. C. Kuranz. Concept to diagnose mix with imaging x-ray thomson scattering. *Review Of Scientific Instruments*, 83(10), 2012. ISSN 0034-6748. doi: 10e534 10.1063/1.4732185.

S. R. Klein, E. J. Gamboa, C. M. Huntington, C. C. Kuranz, P. Susalla, S. Chadwick, B. Lairson, D. E. Hoover, F. Elsner, G. Malamud, C. Di Stefano, R. S. Gillespie, and R. P. Drake. Innovations in target fabrication techniques at the university of michigan. *Fusion Science and Technology*, 63(2):305–312, 2013. ISSN 1536-1055.

C. M. Krauland, R. P. Drake, C. C. Kuranz, R. Sweeney, M. Grosskopf, S. Klein, R. Gillespie, P. A. Keiter, B. Loupias, and E. Falize. Radiative reverse shock laser experiments relevant to accretion processes in cataclysmic variables. *Physics Of Plasmas*, 20(5):056502, 2013a. ISSN 1070-664X. doi: 056502 10.1063/1.4805023.

C.M. Krauland, R.P. Drake, C.C. Kuranz, B. Loupias, T. Plewa, C.M. Huntington, D.N. Kaczala, S. Klein, R. Sweeney, R.P. Young, E. Falize, B. Villette, and P.A. Keiter. Reverse radiative shock laser experiments relevant to accreting stream-disk impact in interacting binaries. *Astrophysical Journal Letters*, 762:1–4, 2013b.

N. L. Kugland, J. S. Ross, P. Y. Chang, R. P. Drake, G. Fiksel, D. H. Froula, S. H. Glenzer, G. Gregori, M. Grosskopf, C. Huntington, M. Koenig, Y. Kuramitsu, C. Kuranz, M. C. Levy, E. Liang, D. Martinez, J. Meinecke, F. Miniati, T. Morita, A. Pelka, C. Plechaty, R. Presura, A. Ravasio, B. A. Remington, B. Reville, D. D. Ryutov, Y. Sakawa, A. Spitkovsky, H. Takabe, and H. S. Park. Visualizing electromagnetic fields in laser-produced counter-streaming plasma experiments for collisionless shock laboratory astrophysics. *Physics Of Plasmas*, 20(5), 2013. ISSN 1070-664X. doi: 056313 10.1063/1.4804548.

C. C. Kuranz, R. P. Drake, C. M. Huntington, C. M. Krauland, C. A. Di Stefano, M. Trantham, M. J. Grosskopf, S. R. Klein, and D. C. Marion. Early-time evolution of a radiative shock. *High Energy Density Physics*, 9(2):315–318, 2013a. ISSN 1574-1818. doi: 10.1016/j.hedp.2012.12.012.

C. C. Kuranz, R. P. Drake, C. M. Krauland, D. C. Marion, M. J. Grosskopf, E. Rutter, B. Torralva, J. P. Holloway, D. Bingham, J. Goh, T. R. Boehly, and A. T. Sorce. Initial conditions of radiative shock experiments. *Physics Of Plasmas*, 20(5), 2013b. ISSN 1070-664X. doi: 056321 10.1063/1.4805021.

C.C. Kuranz, F.W. Doss, R.P. Drake, M.J. Grosskopf, and H.F. Robey. Using wall shocks to measure preheat in laser-irradiated, high-energy-density, hydrodynamics experiments. *High Energy Density Physics*, 6:135–142, 2010a. doi: 10.1016/j.hedp.2010.01.002.

C.C. Kuranz, R.P. Drake, M.J. Grosskopf, B. Fryxell, A. Budde, J.F. Hansen, A.R. Miles, T. Plewa, N.C. Hearn, and J.P. Knauer. Spike morphology in blast-wave-driven instability experiments. *Physics of Plasmas*, 17:052709, 2010b.

C.C. Kuranz, H.-S. Park, B.A. Remington, R.P. Drake, A.R. Miles, H.F. Robey, J.D. Kilkenny, C.J. Keane, D.H. Kalantar, C.M. Huntington, C.M. Krauland, E.C. Harding, M.J. Grosskopf, D.C. Marion, F.W. Doss, E. Myra, B. Maddox, B. Young, J.L. Kline, G. Kyrala, T. Plewa, J.C. Wheeler, W.D. Arnett, R.J. Wallace, E. Giraldez, and A. Nikroo. Astrophysically relevant radiation hydrodynamics experiment at the national ignition facility. *Astrophysics and Space Science*, 336(1):207–211, 2011.

G. Malamud, C. A. Di Stefano, Y. Elbaz, C. M. Huntington, C. Kuranz, P. A. Keiter, and R. P. Drake. A design of a two-dimensional, multimode rm experiment on omega-ep. *High Energy Density Physics*, 9(1): 122–131, 2013a. ISSN 1574-1818. doi: 10.1016/j.hedp.2012.11.008.

G. Malamud, A. Shimony, W.C. Wan, C.A. Di Stefano, Y. Elbaz, C.C. Kuranz, P.A. Keiter, R.P. Drake, and D. Shvarts. A design of a two-dimensional, supersonic kh experiment on omega-ep. *High Energy Density Physics*, 9(1):672–686, 2013b. doi: 10.1016/j.hedp.2013.06.002.

R. G. McClarren and D. Holladay. Benchmarks for verification of hedp/ife codes. *Fusion Science and Technology*, 60(2):600–604, 2011. ISSN 1536-1055. URL `<Go to ISI>://WOS:000293420300031`. McClarren, Ryan G. Holladay, Daniel 19th Topical Meeting on the Technology of Fusion Energy (TOFE-19) Nov 08-11, 2011 Las Vegas, NV Amer Nucl Soc Fus Energy Div, US Dept Energy, Atom Energy Soc Japan, Gen Atom, Univ California.

R.G. McClarren and R. P. Drake. Radiative transfer in the cooling layer of a radiating shock. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 111(14):2095–2105, 2010.

R.G. McClarren and J.G. Wohlbier. Analytic solutions for ion-electron- radiation coupling with radiation and electron diffusion. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 112:119–130, 2010.

R.G. McClarren, R. P. Drake, M.L. Adams, J. E. Morel, and J. P. Holloway. Theory of radiative shocks in the mixed, optically thick-thin case. *Physics of Plasmas*, 17:093301, 2010.

Ryan G. McClarren, D. Ryu, R. Paul Drake, Michael Grosskopf, Chuan-Chih Chou, Bruce Fryxell, and James Paul Holloway. A physics informed emulator for laser-driven radiating shock simulations. *Reliability Engineering and System Safety*, 96:1194–1207, 2011.

J. E. Morel and R. McClarren. Stability of explicit radiation-material coupling in radiative transfer calculations. *Journal Of Quantitative Spectroscopy & Radiative Transfer*, 112(10):1518–1524, 2011. ISSN 0022-4073. doi: 10.1016/j.jqsrt.2011.03.013.

P. Movahed and E. Johnsen. A solution-adaptive method for efficient compressible multifluid simulations, with application to the richtmyer-meshkov instability. *Journal of Computational Physics*, 239:166–186, 2013. ISSN 0021-9991. doi: 10.1016/j.jcp.2013.01.016.

E. S. Myra and W. D. Hawkins. A comparison study of discrete-ordinates and flux-limited diffusion methods for modeling radiation transport. *High Energy Density Physics*, 9(1):91–102, 2013. ISSN 1574-1818. doi: 10.1016/j.hedp.2012.10.007.

K. Raman, O. A. Hurricane, H.-S. Park, B. A. Remington, H. F. Robey, V. A. Smalyuk, R. P. Drake, C. M. Krauland, C. C. Kuranz, J. F. Hansen, and E.C. Harding. Three-dimensional modeling and analysis of a high energy density kelvin-helmholtz experiment. *Physics Of Plasmas*, 19:092112, 2012.

Pritam Ranjan, Derek Bingham Wilson Lu, Shane Reese, Brian J. Williams, Chuan-Chih Chou, Forrest Doss, Michael Grosskopf, and James Paul Holloway. Follow- up experimental designs for computer models and physical processes. *Journal of Statistical Theory and Practice*, 5(1):119–136, 2011.

E.M. Rutter, R.P. Drake, M.J. Grosskopf, G. Malamud, E.C. Harding, C.C. Kuranz, and P.A. Keiter. Comparison between kelvin-helmholtz instability experiments on omega and simulation results using the crash code. *High Energy Density Physics*, 9:148–151, 2013.

D. Ryu, E. N. Li, and B. K. Mallick. Bayesian nonparametric regression analysis of data with random effects covariates from longitudinal measurements. *Biometrics*, 67(2):454–466, 2011. ISSN 0006-341X. doi: 10.1111/j.1541-0420.2010.01489.x. URL `<Go to ISI>://WOS:000292504000013`.

D.W. Savin, N.S. Brickhouse, J.J Cowan, R.P. Drake, S.R. Federman, G.J. Ferland, A. Frank, M.S. Gudipati, W.C. Haxton, E. Herbst, S. Profumo, F. Salama, L.M. Ziurys, and E.G. Zweibel. The impact of recent advances in laboratory astrophysics on our understanding of the cosmos. *Reports on Progress in Physics*, 75:036901, 2012.

V. A. Smalyuk, J. F. Hansen, O. A. Hurricane, G. Langstaff, D. Martinez, H.-S. Park, K. Raman, B. A. Remington, H. F. Robey, O. Schilling, R. Wallace, Y. Elbaz, A. Shimony, D. Shvarts, C. Di Stefano, R. P. Drake, D. Marion, C. M. Krauland, and C. C. Kuranz. Experimental observations of turbulent mixing due to kelvin–helmholtz instability on the omega laser facility. *Physics Of Plasmas*, 19(9):092702, 2012.

V. A. Smalyuk, O. A. Hurricane, J. F. Hansen, G. Langstaff, D. Martinez, H. S. Park, K. Raman, B. A. Remington, H. F. Robey, O. Schilling, R. Wallace, Y. Elbaz, A. Shimony, D. Shvarts, C. Di Stefano, R. P. Drake, D. Marion, C. M. Krauland, and C. C. Kuranz. Measurements of turbulent mixing due to kelvin-helmholtz instability in high-energy-density plasmas. *High Energy Density Physics*, 9(1):47–51, 2013. ISSN 1574-1818. doi: 10.1016/j.hedp.2012.10.001. URL `<Go to ISI>://WOS:000317324500009`.

D.P. Starinshak, S. Karni, and P. Roe. A new level set method. *J. Comp. Phys.*, page submitted, 2013.

H. F. Stripling, M. L. Adams, R. G. McClarren, and B. K. Mallick. The method of manufactured universes for validating uncertainty quantification methods. *Reliability Engineering & System Safety*, 96(9):1242–1256, 2011. ISSN 0951-8320. doi: 10.1016/j.ress.2010.11.012. URL <Go to ISI>://WOS:000293107900022. Times Cited: 0 Stripling, H. F. Adams, M. L. McClarren, R. G. Mallick, B. K. Si.

H. F. Stripling, R. G. McClarren, C. C. Kuranz, M. J. Grosskopf, E. Rutter, and B. R. Torralva. A calibration and data assimilation method using the bayesian mars emulator. *Annals Of Nuclear Energy*, 52:103–112, 2013. ISSN 0306-4549. doi: 10.1016/j.anucene.2012.08.025. URL <Go to ISI>://WOS:000313997700012. Times Cited: 0 Stripling, H. F. McClarren, R. G. Kuranz, C. C. Grosskopf, M. J. Rutter, E. Torralva, B. R. Si.

M. R. Trantham, C. C. Kuranz, G. Malamud, M. J. Grosskopf, E. S. Myra, R. P. Drake, A. R. Miles, H. S. Park, and B. A. Remington. Simulations of radiative effects on the rayleigh-taylor instability using the crash code. *High Energy Density Physics*, 9(2):303–308, 2013. ISSN 1574-1818. doi: 10.1016/j.hedp.2012.12.016. URL <Go to ISI>://WOS:000319952300012.

B. van der Holst, G. Toth, I.V. Sokolov, K. G. Powell, J. P. Holloway, E.S. Myra, Q. Stout, M.L. Adams, J. E. Morel, and R. P. Drake. A block-adaptive-mesh code for radiative shock hydrodynamics: Implementation and verification. *Astrophysical Journal Supplement Series*, 194:23, 2011.

B. van der Holst, G. Toth, I.V. Sokolov, L.K.S. Daldorff, K.G. Powell, and R.P. Drake. Simulating radiative shocks in nozzle shock tubes. *High Energy Density Physics*, 8:161–169, 2012.

B. van der Holst, G. Toth, I. V. Sokolov, B. R. Torralva, K. G. Powell, R. P. Drake, M. Klapisch, M. Busquet, B. Fryxell, and E. S. Myra. Simulating radiative shocks with the crash laser package. *High Energy Density Physics*, 9(1):8–16, 2013. ISSN 1574-1818. doi: 10.1016/j.hedp.2012.09.011.