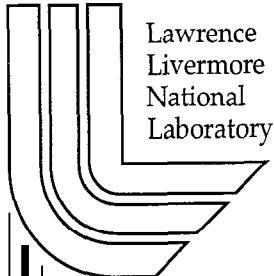


Grain-scale Dynamics in Explosives

J.E. Reaugh

September 30, 2002

U.S. Department of Energy



Lawrence
Livermore
National
Laboratory

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This work was performed under the auspices of the U. S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.

This report has been reproduced directly from the best available copy.

Available electronically at <http://www.doc.gov/bridge>

Available for a processing fee to U.S. Department of Energy
And its contractors in paper from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-mail: reports@adonis.osti.gov

Available for sale to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-mail: orders@ntis.fedworld.gov
Online ordering: <http://www.ntis.gov/ordering.htm>

OR

Lawrence Livermore National Laboratory
Technical Information Department's Digital Library
<http://www.llnl.gov/tid/Library.html>

Grain-scale Dynamics in Explosives

John E. Reaugh

September 30, 2002

Contents

Executive Summary	1
Project Objective.....	1
Milepost Objective	1
Report Objective.....	2
Milepost Achievements	2
Future Work	3
Background	4
Fundamental aspects of reactive flow models.....	4
Prior Research.....	4
Focus of this project	5
Project History	6
Structure of Plastic Bonded Explosives (PBXs)	7
Summary.....	7
Monomodal size distribution.....	7
Grain-conforming Mesh.....	7
Shaping-in the crystals.....	8
Arbitrary size distribution.....	10
Introduction of defects.....	15
Shock waves in PBX.....	18
Summary.....	18
Simulations of shocks in an assembly of crystals	18
Collapse of a void.....	20
Ignition after void collapse.....	21
Deflagration.....	26
Summary.....	26
Chemical kinetics and equations of state	26
Equation of state	27
Transport properties	27
Simulations of deflagration.....	28
Measurements of deflagration	30
Kinetic rates at high pressure.....	33
Resolution of the differences.....	35
Future work	36
Growth of hot spots in PBX.....	37
Summary.....	37
Simplified model of hot spot growth	37
Corrections to the simplified model	38
Growth of a single hot spot in ALE-3D	39
Growth of an assembly of hot spots	40
Application of initial continuum model	43
Summary.....	43
Baseline results	43

Results with parameters developed from grain-scale simulations	43
Significance of r_0	45
Summary and Conclusions	47
Future Directions	49
Acknowledgements	50
Project Documentation	52
References	53
Appendix: Improvements to ALE-3D that have resulted from this project	54
Level set Issues	54
General Material Issues:	54
Chemical Material/Reaction Issues	55
Thermal Solver Issues	55
IO issues	55
General stuff	55
Problem Generation	56

Figures

Figure 1. The particle size distribution is narrow, centered at 76- μm diameter. Binder is 15% by volume, and 3 μm thick between crystals. Slice dimensions are 300 μm by 300 μm	9
Figure 2. Particle size distribution for LX-04 (solid) and two other distributions introduced later in the text – the monomodal coarse distribution (dot) and the bimodal distribution (dash).....	10
Figure 3. Coarse monomodal distribution with particle size centered at 76- μm diameter. Slice is 300 μm by 300 μm	12
Figure 4. Bimodal distribution with particle sizes 76- μm and 6- μm diameter. Slice is 300 μm by 300 μm	13
Figure 5. Bimodal distribution with particle sizes 76- μm and 6- μm diameter. Modified method with limited number of planes to describe large particles. Slice is 300 μm by 300 μm	14
Figure 6. Truncated bimodal distribution on a 300- μm cube. Cube contains 0.1 M particles defined by 10 M planes on 200 M zones.	16
Figure 7. Coarse monomodal distribution in a 300- μm cube with defects in the larger particles (top) and microphotograph of defects in a single 150 μm HMX crystal (bottom).	17
Figure 8. Shock wave in a 150 μm by 150 μm brick, 300 μm long. Binder is HTPB, which is a poor density match for the crystalline HMX. The shock front is seen to be	

sharp within an explosive crystal. Other calculations with Viton binder, an excellent density match, show similar behavior.....	19
Figure 9. Spherical hole collapsed by a sharp shock before the arrival of the front (left) and during passage of the front (right).	21
Figure 10. Region occupied by final product in the wake of a 5 GPa shock wave in HMX. Original 14 μm diameter void was located between $x=0$ and $x=14 \mu\text{m}$	22
Figure 11. Region occupied by final product in the wake of a 20 GPa shock wave in HMX. Same initial geometry as in Fig. 10.....	23
Figure 12. Mass of final product scaled to mass of reactant that would fill the original defect volume. The solvent-filled cavity has a much higher pressure cut-off for reaction than the air-filled cavity does.	24
Figure 13. Direct numerical simulation of a flame at 1 GPa. The mass fractions of the four species are scaled so that the value 2000 represents 100 %.	29
Figure 14. Direct numerical simulation of a flame in HMX at 12 GPa. Other conditions the same as in Fig. 13.....	29
Figure 15. Schematic of the diamond anvil cell experiments to measure the high-pressure deflagration rate. The ignition pulse starts a cylindrically symmetric deflagration.....	31
Figure 16. Velocity correction as a function of scaled radius of the flame location.	32
Figure 17. Experimental and corrected computational flame speeds as a function of pressure. The uniform grain size data are incompletely compacted below 6 GPa, so the observed values are the speed of gas products diffusing ahead in a porous medium. 33	
Figure 18. Appearance of nitrogen (top) and water (bottom) in the QMD simulations of Manaa (symbols) and calculated with the kinetic rates using our simplified chemistry for HMX decomposition (lines). In both figures the slower chemistry is the original kinetics from low-pressure experiments.....	34
Figure 19. Species diffusivity (solid) and thermal diffusivity (dot) on the 3000K isotherm for HMX final products.	35
Figure 20. Model result for apparent velocity of deflagration front and front velocity from ALE-3D For both 0.5 and 0.67 μm zones, the agreement is good after the initial transient in ALE-3D. The input deflagration velocity is 60 m/s.....	40
Figure 21. Growth of reaction from a number of 14 μm diameter hot spots in a 150- μm cube at early time (left) and at later time (right). Red spheres mark the flame fronts and enclose final product gas.....	41

Figure 22. Calculated form factors for arrays of hot spots with differing radii and volume factor. Curves have been scaled by $3vf_v^{1/3}/r_0$ 42

Figure 23. Detonation speed as a function of reciprocal diameter for Tarver's parameters (solid dot) and the new parameters (other symbols) compared with the experimenters' fit to their data (solid) 45

Tables

Table 1. Parameter values for the growth of reaction. Time is in μs 41

Table 2. Parameter values for reactive flow models for PBX-9404. Units are g, cm, μs , Mbar. 44

Executive Summary

Project Objective

High explosives can have reactions to external stimuli that range from mild pressure bursts to full detonation. The ability to predict these responses is important for understanding the performance as well as the safety and reliability of these important materials. At present, we have only relatively simple phenomenological computational models for the behavior of high explosives under these conditions. These models are limited by the assumption that the explosive can be treated as homogeneous. In reality the explosive is a highly heterogeneous composite of irregular crystallites and plastic binder. The heterogeneous nature of explosives is responsible for many of their unique mechanical and chemical properties.

We use computational models to simulate the response of explosives to external mechanical stimuli at the grain-scale level. The ultimate goal of this work is to understand the detailed processes involved with the material response, so that we can develop realistic material models, which can be used in a hydrodynamics/multi-physics code to model real systems. The new material models will provide a more realistic description of the explosive system during the most critical period of ignition and initiation.

The focus of this work is to use the results of grain-scale simulations to develop an advanced macroscopic reactive flow model that is consistent with our understanding of the grain-scale details, and that can incorporate such information quantitatively. The objective is to connect changes to observed properties of the explosive (grain size distribution, binder thickness distribution, void shape, size, and separation distribution, binder mechanical properties, *etc.*) with predictions of the resulting sensitivity and performance.

Milepost Objective

Arriving at this milepost does not attain the long-term objective of the project. Instead, we will demonstrate for the first time the direct connection between grain-scale dynamics and the parameters of a reactive flow model. Our deliverable is an initial model, not the final one. The milepost model requirements are:

1. The mixture rule for the model will be derived from analysis of numerical simulations of the growth of reaction from hot spots.
2. The functional form and parameter values for the growth of reaction in the model will be determined by results of grain-scale simulations of the growth of reaction from hot spots.

3. The pressure- and temperature-dependent deflagration rate in the model will be taken from numerical and/or analytic and/or experimental results.

4. The model as described will be incorporated in numerical simulations (at the continuum scale) in which non-ideal behavior of the explosive is observed.

Report Objective

The purpose of this report is to document our attainment of the milepost, and to summarize the status of the project by showing how the various aspects of the project fit together to achieve the project objective. To that end, the sections of the report do not have the completeness that would be expected of a journal article, but rather are intended to explain our approach and to summarize the project achievements. We anticipate a sequence of technical articles to be submitted over the next year. Those articles that have already been submitted are included in the section "Project Documentation."

Milepost Achievements

1. We derived a mixture rule [Reaugh and Lee, in the section Project Documentation] based on numerical simulation of the growth of reaction by flame spread from ignition sites. The rule is to treat computational zones containing both unreacted explosive and product gases as a mixture at pressure equilibrium, but not at thermal equilibrium. As part of that work we demonstrated that for a JWL product equation of state, the choice of mixture rule does not affect the result. As a result, we concentrated our resources on the other aspects of the model.
2. We performed simulations of the growth of reaction from hot spots at the grain scale. From these simulations we inferred the form factor used in our initial continuum model.
3. We used the measured deflagration rate in our initial continuum model. Our direct numerical simulations of the deflagration rate at high pressure under-predict the measured values. We determined the additional physics modules needed in ALE-3D to resolve the pressure dependence of deflagration.
4. We applied the initial continuum model (reactive flow model) to the calculation of detonation in small diameter, unconfined explosive cylinders, for which the reaction zone is a significant fraction of the diameter, so that the non-ideality is an important feature of the results.

In addition to meeting the milepost objectives, we have achieved a number of technical firsts and additional accomplishments:

1. We developed methods [Bastea] by which the transport properties of non-ideal gas mixtures can be determined that are consistent with the equation of state at high temperature and density.
2. We developed methods to produce realistic particle size distributions for assemblies of plastic bonded explosives. In so doing, we are near the limit imposed by present computer software and hardware to using the full richness of these assemblies of particles in our simulations.
3. We have begun a computational study of the ignition process around intra-crystalline defects, both air-filled and solvent filled, and the pressure dependence of the ignition process. These features have been incorporated into the continuum model in a preliminary way.
4. We have begun an experimental project funded under the joint DoD/DOE Munitions Technology Development Program. The experimental results from that work [Farber, *et al.*], in conjunction with the computational and theoretical work under direct ASCI sponsorship, have been crucial in identifying weaknesses in our numerical simulations, and those additional physics models needed to advance our work.

Future Work

We are well-started on unfolding the events leading from an initial mechanical stimulus to a detonation as an orderly process that depends in a demonstrable way on the thermodynamic and transport properties of the explosive material, and on the heterogeneous nature of a polycrystalline explosive including its imperfections. Nevertheless, during the course of our research, we have identified additional work that needs to be done.

The most important areas for future research are where our understanding is weakest and the consequences are large. At present, we feel these areas are: the morphology of defects and the ignition process that results from mechanical stress, both high-pressure shock waves and low intensity mechanical stimulus; resolution of the differences between our numerical simulations and experiments for the propagation of deflagration at high and low pressure; and the effect of binder material in the deflagration process. These three research areas will all be advanced more quickly and more thoroughly by the continued tightly coupled experimental, theoretical, computational, and numerical methodsefforts, and on the simultaneous ALE-3D development efforts.

Background

Fundamental aspects of reactive flow models

The fundamental reactive flow model is a rate law that describes the chemical change from the reactant, which is the unreacted solid explosive, to product, which is a mixture of gases with perhaps some solid particles such as graphite and/or diamond and/or metal oxide. The rate law incorporates separate terms for ignition and growth to completion of the reaction. In some models the ignition term is consistent with the idea of a shock wave collapsing the defects and creating a hot spot. In other models it is consistent with the idea of localized shear bands creating hot surfaces. The growth to completion of reaction is consistent with the idea of very subsonic deflagration of the explosive moving from the hot spot outward, with the surface area at the flame front increasing as the deflagration proceeds outward from a localized hot spot. Eventually the flame fronts arising from multiple hot spots intersect, and the available surface area at the overall flame front decreases as the reaction completes. The changing surface area, which relates the mass rate of burning to the flame velocity, is represented by a form factor that depends on the extent of reaction. Although the flame velocity is very slow relative to the detonation velocity, it can result in significant mass transformation in the reaction zone (order mm) if the individual hot spots are very close together (order μm)

In addition, there is a mixture rule to determine the equation of state of a mixture of partially reacted material that is present in the reaction zone. The mixture rule is, for many reactive flow models, a construction of additive volumes with pressure and temperature equilibrium enforced. Other reactive flow models may use a partial pressure construction or adiabatic pressure equilibrium.

Prior Research

Starting in the 1940's explosive researchers became aware of the safety hazards that resulted from the presence of 0.1 to 10 μm gas bubbles or voids in solid and liquid explosives. These defects made the explosive much more sensitive to shock. In the engineering design of explosives, these defects were eliminated to make explosive handling safer. In commercial slurry explosives, such defects were deliberately introduced to make them detonable in appropriate geometric dimensions. These advances in explosive design and safety were all achieved using empirical methods.

Reactive flow models have been developed and used since the 1980's to describe the response of explosives to mechanical shock under accident conditions. They have also been applied to configurations with a small geometric dimension, which results in both reduced detonation velocity and reduced explosive energy delivery. Although such models are in general accord with the transformation of bubbles and voids into hot

spots, and the subsequent ignition of the explosive, the parameter values are fit to specific families of experiments. Previous model development has been focused on fitting experimental families of experiments, using functional forms suggested by the continuum-scale experiments, and parameter fits to those experiments. Historically, the number of parameters has increased as each research group sought to include more and more families of experiments. This increase in complexity is not intrinsically bad, but requires each variation of explosive manufacture, age, and temperature of the experiment to be treated as a new material, subject to the same battery of experiments and parameter fitting.

Focus of this project

The computing speed and memory of ASCI computer hardware and the incorporation of chemical kinetics and heat transfer in ASCI computer software make it possible to study the behavior of voids and gas bubbles in explosives by direct numerical simulation. At present, and in the near future, it is practical to study only isolated pieces of explosive with dimension less than one cubic mm when cubic μm must be resolved. To use the insight and understanding that result from direct numerical simulation of the dynamics of such defects, we must use models that are consistent with the grain-scale dynamics and that can be applied to the behavior of engineering-scale configurations of explosives.

By introducing results of grain-scale simulations into the framework of such a model, the parameter values can be directly connected with grain-scale characteristics of the specific explosive. The long-term objective is to connect changes to observed properties of the explosive (grain size distribution, binder thickness distribution, void shape, size, and separation distribution, binder mechanical properties, *etc.*) with predictions of sensitivity and performance. This long-term objective will not be reached with this milepost. Instead, this milepost will demonstrate for the first time the direct connection between grain-scale dynamics and the parameters of a reactive flow model. Our deliverable is an initial model, not the final one.

A model must be demonstrated by its use in numerical simulations. It would, however, detract from the long-term objective if the criterion for successfully achieving this milepost were to be specified by a required percent deviation from experiment. That level of accuracy has already been demonstrated by the success of reactive flow models. The parameters in such a model can be successfully tuned to a specific family of experiments. The achievement of this milepost is that the parameters and functional forms can come from enforcing consistency with the grain scale dynamics, and so are obtained independently from the family of experiments.

To date the project has used HMX (1,3,5,7-tetranitro-1,3,5,7-tetraazacyclooctane) as the prototypical explosive molecule, and LX-04, (85% HMX, 15% by weight Viton binder) as

the prototypical engineered plastic bonded explosive (PBX). We anticipate that future work of this project will include the TATB explosive molecule as well.

Project History

This project was begun in January 1999 as a part of the ASCI project to form a bridge between the micro-scale atomic and molecular modeling, which has been a part of the Lawrence Livermore National Laboratory's explosive research for many years, and the continuum models, which have been used to describe the behavior of ideal or nearly ideal explosives. The idea was that information would flow from the research at the micro- and meso-scale to develop and improve continuum models for macro-scale applications.

In FY00 the computational work was expanded to include an experimental effort under the Munitions Technology Development Program to measure the deflagration rate of explosives at high pressure, using diamond anvil cell techniques developed at this laboratory in the 1990s. Knowledge of these rates was identified as a key to validating the computational research.

The first level 1 milepost for the materials effort in ASCI was assigned to this project, due in the 4th quarter of FY02. During CY02, it was officially downgraded to a level 2 milestone. This laboratory has chosen to continue to treat the project and the milestone as though it were still a level-1 milepost.

Structure of Plastic Bonded Explosives (PBXs)

Summary

To study the grain-scale dynamics of PBX requires a geometric description of the explosive crystals and binder. We developed methods to describe the structure of explosives as an assembly of crystals held together with binder at 100% crystal density. We introduce intra- and inter-crystalline imperfections into this assembly. Each crystal may have the same material properties, or may be assigned a unique crystal orientation that is independent of its neighbors. The method may also be used to describe metals and unconsolidated assemblies of crystals (loose powder) with substantial (50%) pore volume.

Monomodal size distribution

Our early description of PBX used a distribution of particles developed by randomly selecting crystal centroids within a specified volume, but rejecting those trial centroids that were too close to their nearest neighbor. With the modest number of crystals we were then using (a few hundred) efficient algorithm programming was preferred over efficient algorithm execution. A trial centroid was simply tested against all previously accepted centroids before deciding to retain it.

We have used the QHULL computer program, which we downloaded from the University of Minnesota Geometry Center site, at present found at (<http://www.rtporal.umn.edu>). Source code (the program is written in C), documentation, and sample problems are contained in the downloaded file, which can be accessed with links provided at that site. The Geometry Center is now closed, but former researchers from the center have provided a recently updated version (August 20, 2002). At present, Brad Barber, Cambridge MA is apparently keeping the activity updated. We have used QHULL throughout this project to find the Delaunay triangulation of an assembly of points, and the Voroni region construction for that assembly. In this construction, a region contains all points that are closer to the region's center point than they are to any other center point. This construction is then used to describe the assembly of crystals with a uniform coating of binder on each crystal facet.

Grain-conforming Mesh

For our initial work we constructed grain-conforming meshes. These meshes were constructed by subdividing the crystal particles into tetrahedrons. The surface of each crystal is the family of planes that bisect the line segments joining the centroids to its nearest neighbors. The Delaunay construction is used to form the list of neighbors for each centroid. One node of each tetrahedron used to describe a grain is the crystal centroid. The opposite base of the tetrahedron lies on a surface facet of the crystal. That

facet is subdivided into triangles by first ordering the vertices that define the facet counterclockwise when viewed from the crystal centroids. Then adjacent pairs of points form one side of the tetrahedral face, and the third point is the center of the facet. As long as each tetrahedron has the same number of zones along each edge, the nodes will line up and can be merged by the mesh-generating program TRUEGRID used by ALE-3D and other finite-difference and finite-element computer programs. Each crystal is described by a regular array of hexahedral elements. The method is much better suited to describing assemblies of crystals that are similarly sized, rather than assemblies with disparate grain sizes. This method was abandoned after we observed that the resulting irregular mesh imprinted the shock structure, even when the mesh only contained a single material.

Shaping-in the crystals

We reverted, instead, to painted-in shapes on an underlying regular mesh. The advantage to this method is that the mesh boundaries need not have any special relationship to material boundaries. This method, which is common to Eulerian-frame simulation programs, has the disadvantage of creating many mixed cells, finite difference volumes that comprise multiple materials. Such mixed cells are intrinsically less accurate than clean cells in either Eulerian or Lagrange formulations.

The method we use to describe the crystal shapes is one of the standard input methods for ALE-3D, namely a family of planes that describe each facet of the crystal. The original implementation of the shaping-in algorithm in ALE-3D was suitable for a small handful of shapes. Our description of even a few hundred crystals, which require 20 to 50 facets each, pointed out significant inefficiencies. Gary Friedman has removed the inefficiencies, so that an assembly of 10^4 crystals, which requires nearly 10^6 planes, is routine. As of September 2002, the creation of an assembly of 10^5 crystals, using 10^7 planes described on a mesh with 10^8 finite-difference zones is just within the apparent limit of ALE-3D software and ASCI White hardware.

The use of a uniform distribution of centroids, and the use of bisectors of the lines joining nearest neighbors results in a monomodal size distribution of particles that is quite narrow. In Fig. 1 we show a cross section of such a distribution. It has 15% binder by volume, with nominal binder thickness $3 \mu\text{m}$. In each of our assemblies, the binder thickness is constant. In Fig. 1 and the rest of our illustrations, the crystals are assigned region numbers at random. The visualization program MESHTV cycles through a limited palate of 30 different colors before repeating, so that some adjacent crystals with different region numbers can have the same color. In ALE-3D, an arbitrarily chosen crystal orientation can be assigned to each region number.

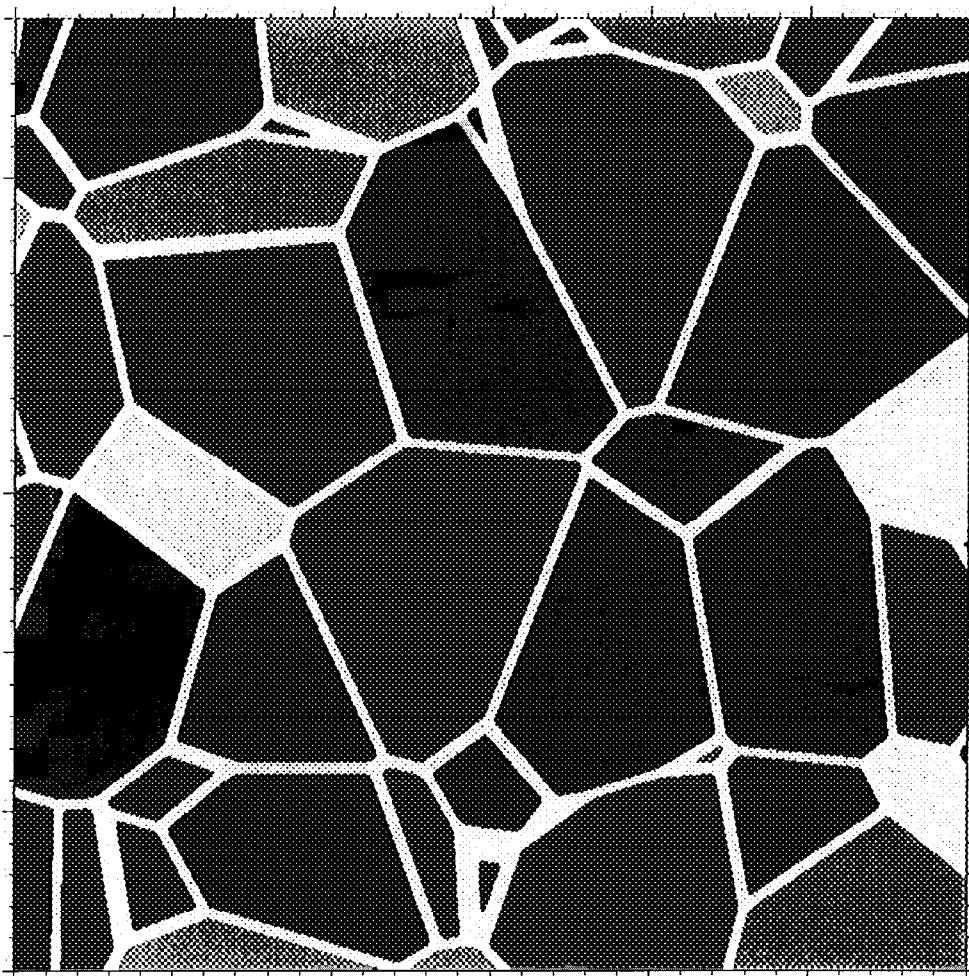


Figure 1. The particle size distribution is narrow, centered at 76- μm diameter. Binder is 15% by volume, and 3 μm thick between crystals. Slice dimensions are 300 μm by 300 μm .

Although a single crystal size is accurate and appropriate for polycrystalline metals and for some biological assemblies of cells, it is inaccurate and inappropriate for PBXs. For such materials, the crystal size distribution is deliberately made bimodal or even trimodal, with the peaks of the distribution separated by an order of magnitude. In

addition, each of the modes has a broad mix of sizes. The size distributions are typically displayed as the volume (or mass) fraction retained on a sieve, which were passed by the next larger sieve but retained on the sieve in question. In such distributions, plotted against the logarithm of the equivalent diameter, the diameters at the half-maximum of the distribution are typically a factor of 2 or 3 larger or smaller than the diameter at the maximum. The particle size distribution shown in Figure 2 is the measured distribution for LX-04. Variations of that distribution, which we have used to describe the structure of PBX and which are illustrated below, are also indicated in Fig. 2.

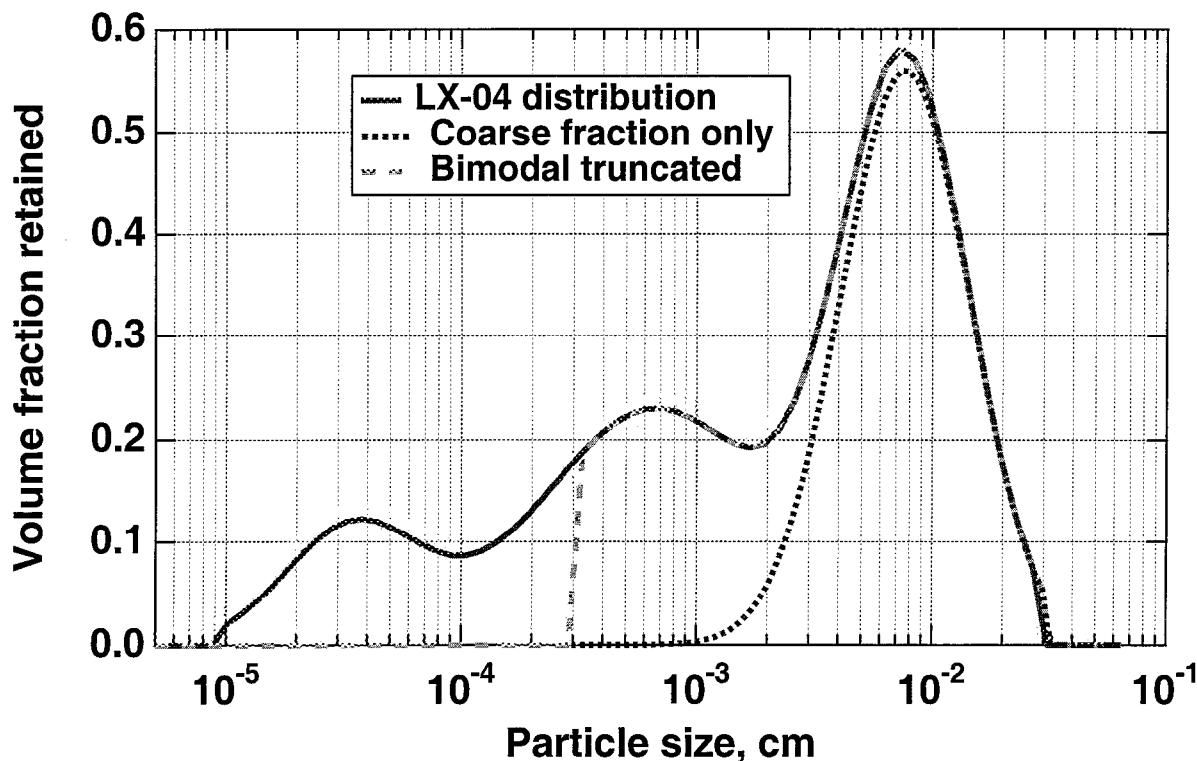


Figure 2. Particle size distribution for LX-04 (solid) and two other distributions introduced later in the text – the monomodal coarse distribution (dot) and the bimodal distribution (dash).

Arbitrary size distribution

We have modified the method described above to obtain arbitrary size-distributions of crystals. The method incorporates two ideas used by Edward Garboczi at the National Institute of Standards and Technology (NIST) to describe concrete. We select spheres with diameter determined by an arbitrary grain-size distribution to fill a cubical box. We effectively choose the number of such spheres by preselecting an unfilled (void) volume, typically 5 to 15% of the total. The particles are sorted by size, and using one of the ideas of Garboczi, we begin by selecting the locations of the particles from largest to smallest. Each successive trial location is tested for overlap with those already chosen. If

overlap is found, a new potential site is randomly chosen. For some combinations of box dimension, particle dimension, and particle size distribution, the method gets stuck, and no solution is found. We have always been able to find a solution by increasing the void volume a few percent. The result of this artifice is to modify the actual size distribution slightly. We have not found it necessary to iterate by modifying the basis size-distribution function.

With broad and multi-modal size distributions, there are significantly more particles to check. We incorporated a second feature of Garboczi's method, which is to gain efficiency by subdividing the box into smaller boxes. In this way, a test particle need only be checked against nearby particles. Although we have not further optimized the selection process, even 10^5 particles have been placed with only a few minutes elapsed computer time.

After the centroids are placed, QHULL is used to select neighbors as before. However, we have found that the direct application of QHULL, using bisector planes for the crystal faces, does a poor job at recovering the basis grain-size distribution. When the bisector is used, a small crystal cannot lie adjacent to a large one. Instead the algorithm creates two middle-sized particles. Rather than use a bisector, we modified the location of the plane so that it is closer to the small particle. We do this by creating a companion file that contains the particle diameter for each centroid. This method, however, causes some crystals to overlap others. We accommodate this overlap by painting-in the crystals from smallest to largest, and surrounding each crystal with binder, which requires two shapes per crystal. We show the results of using only the coarse fraction of LX-04 in cross section in Fig. 3. The presence of smaller crystals increases the surface area, so that to maintain the 15% binder requires a somewhat thinner binder layer, 2.4 μm .

We apply the same method, but include some of the middle distribution of LX-04 centered at 6.6- μm diameter. The result (Fig. 4) shows an order of magnitude increase in the number of crystals. The binder thickness for the same volume fraction drops to less than 0.5 μm . Experimental measurement of the binder thickness of LX-17, which is only 7% binder by volume, ranges from 0.1 to 0.3 μm . The resulting size-distribution recovers the basis one well, but the relatively rare large crystals are unnaturally spherical, because each large one has many smaller crystal neighbors. Our present modification limits the number of facets for each grain to about 40, which is comparable to the number of facets for each crystal in our original (Fig. 1) distribution. The resulting large particles are much more realistically shaped, (Fig. 5) although the overlap of similarly sized particles is not yet totally satisfactory. We had no difficulty generating the cross sectional slices of Figs. 4 and 5. Gary Friedman was recently able to successfully generate a cube with the same bimodal particle size distribution (Fig. 6) on ASCI White. He used 40 nodes, but only half the possible number of CPUs at each node

to allow each processor to have enough memory. The cube contains 216 M zones, and 93 K particles defined by 2.8 M planes.

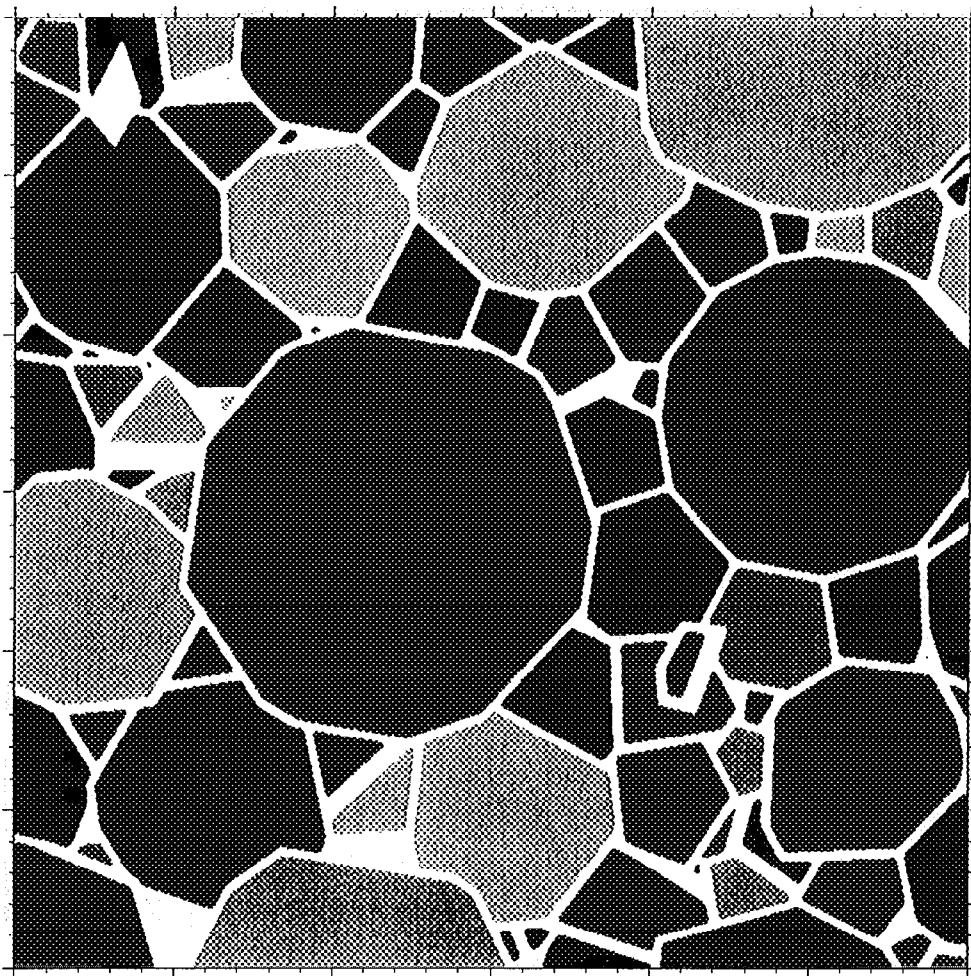


Figure 3. Coarse monomodal distribution with particle size centered at 76- μm diameter. Slice is 300 μm by 300 μm .

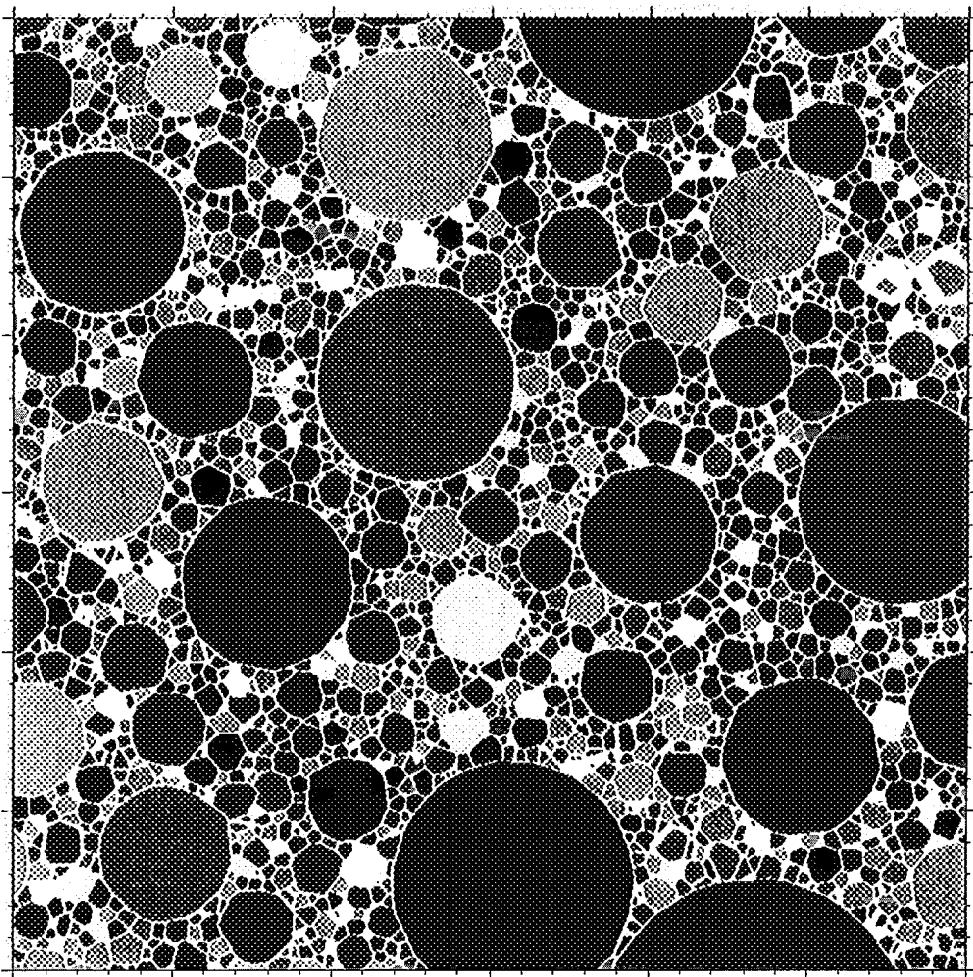


Figure 4. Bimodal distribution with particle sizes 76- μm and 6- μm diameter. Slice is 300 μm by 300 μm .

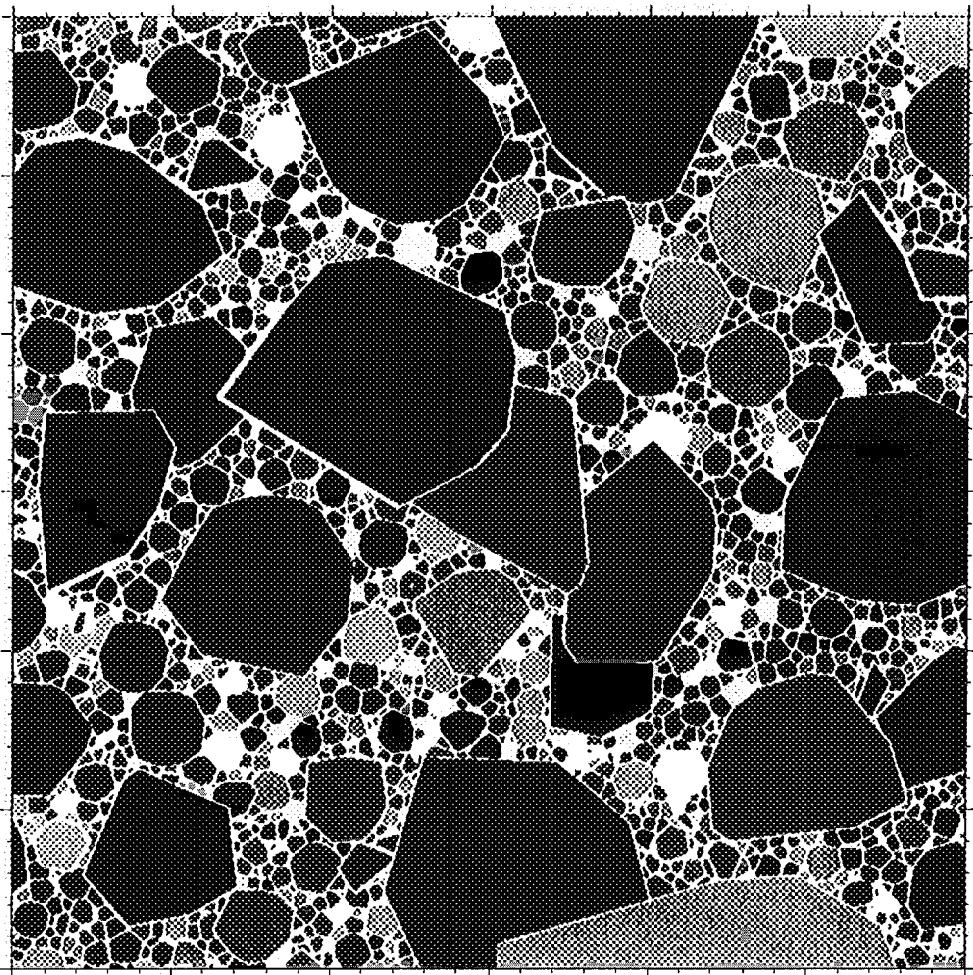


Figure 5. Bimodal distribution with particle sizes 76- μm and 6- μm diameter. Modified method with limited number of planes to describe large particles. Slice is 300 μm by 300 μm .

Introduction of defects

We must introduce defects in our assembly of HMX crystals in order to create the hot spots that are the initiation sites that result from a shock wave. Our original method was to introduce inter-crystalline voids by omitting some binder where three or more crystals meet, leaving voids whose shape resembles deflated playground balls. When we were using assemblies of rather coarse particles, the binder was thick enough that the resulting defects had minimum dimensions of a few μm . This size hotspot has been shown to be sufficiently large (have a sufficiently small surface to volume ratio that heat conduction across the surface into cold material is unable to remove enough heat to prevent thermal runaway). When fine particles are included in our assembly, then the defects caused by absence of binder must be much smaller as well. Hot spots that size are so small that in some circumstances they are unable to grow.

Microscopy by Mark Hoffman has illustrated the presence and size of intra-crystalline defects in HMX crystals. One such microscopic picture, Fig. 7, shows defects on the order of $15\text{-}\mu\text{m}$ diameter in a $150\text{-}\mu\text{m}$ crystal. Hot spots this size are definitely large enough to grow. As a result, we have introduced defects within the largest crystals computationally as shown in Figs. 6 and 7. The distribution of particles in the cube of Fig. 7 is taken from the coarse distribution of HMX and Fig. 6 from the bimodal distribution. Our defects are spherical, but when a crystal is large enough to contain multiple defects, they are permitted to overlap, in accord with Hoffman's observations.

There is, at present, no observational guidance for the morphology of inter-crystalline defects in PBX, which we assume would be voids or vapor-filled pockets created during the pressing process. The intra-crystalline voids may be solvent pockets, or they may be voids. Hoffman's technique does not distinguish, although he states that a defect partially filled with liquid would be identifiable. He has seen no such defects. This leads to the tentative conclusion that the defects he sees are voids, rather than solvent filled pockets. If solvent were present, further dissolution of the explosive crystal would lead to defect growth and incompletely filled defects.

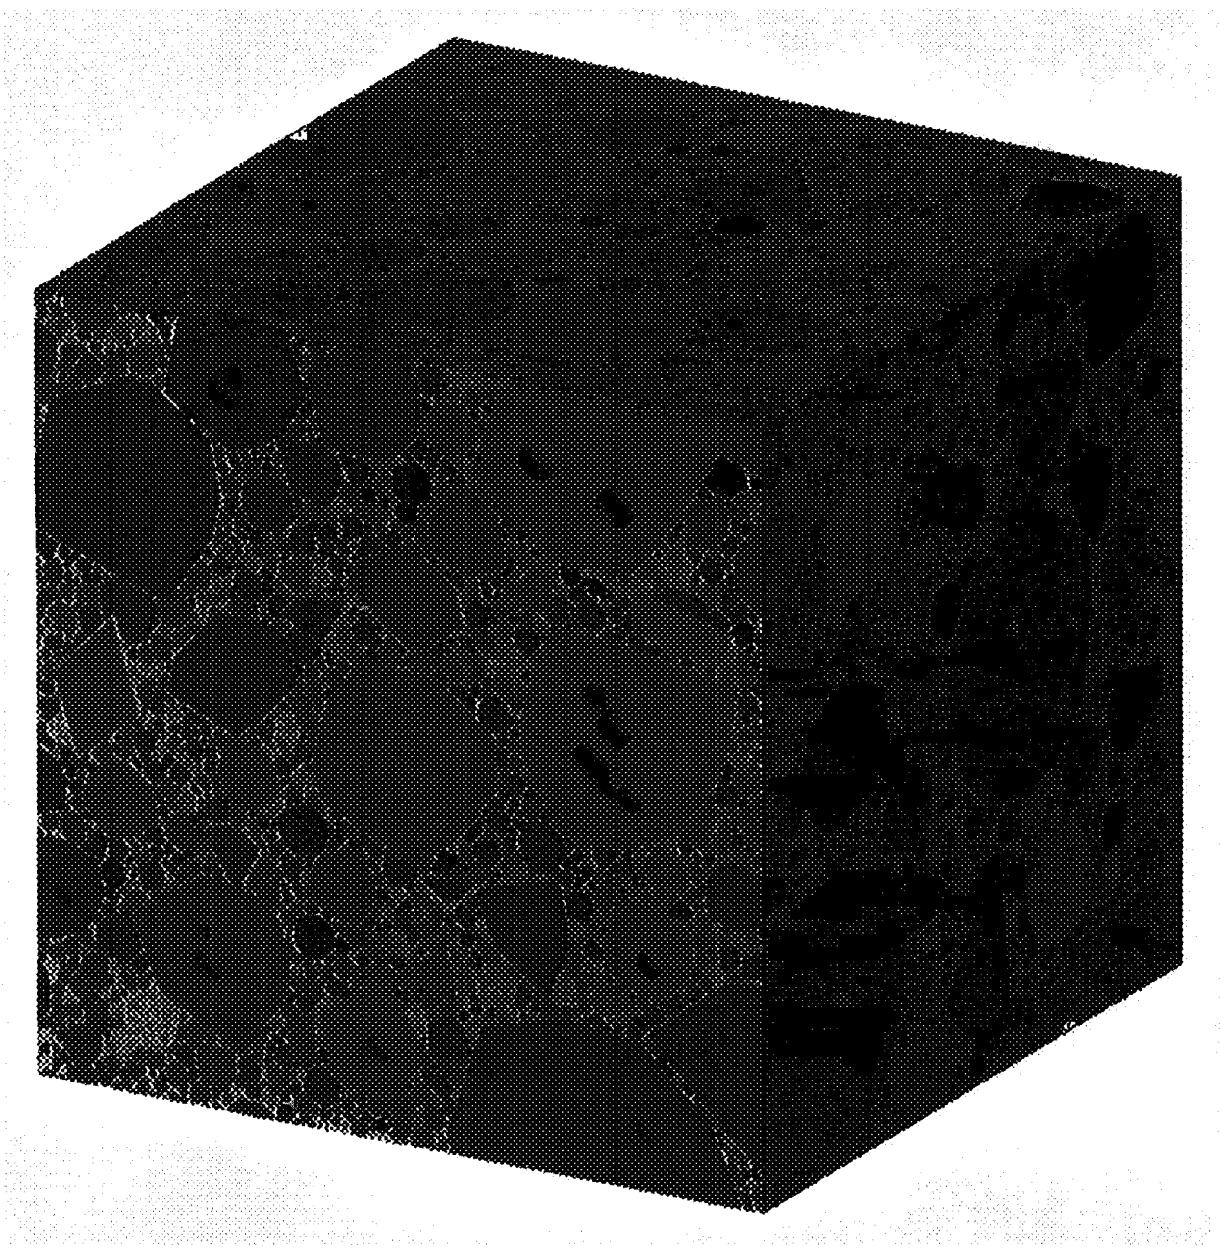


Figure 6. Truncated bimodal distribution on a $300\text{-}\mu\text{m}$ cube. Cube contains 0.1 M particles defined by 10 M planes on 200 M zones.

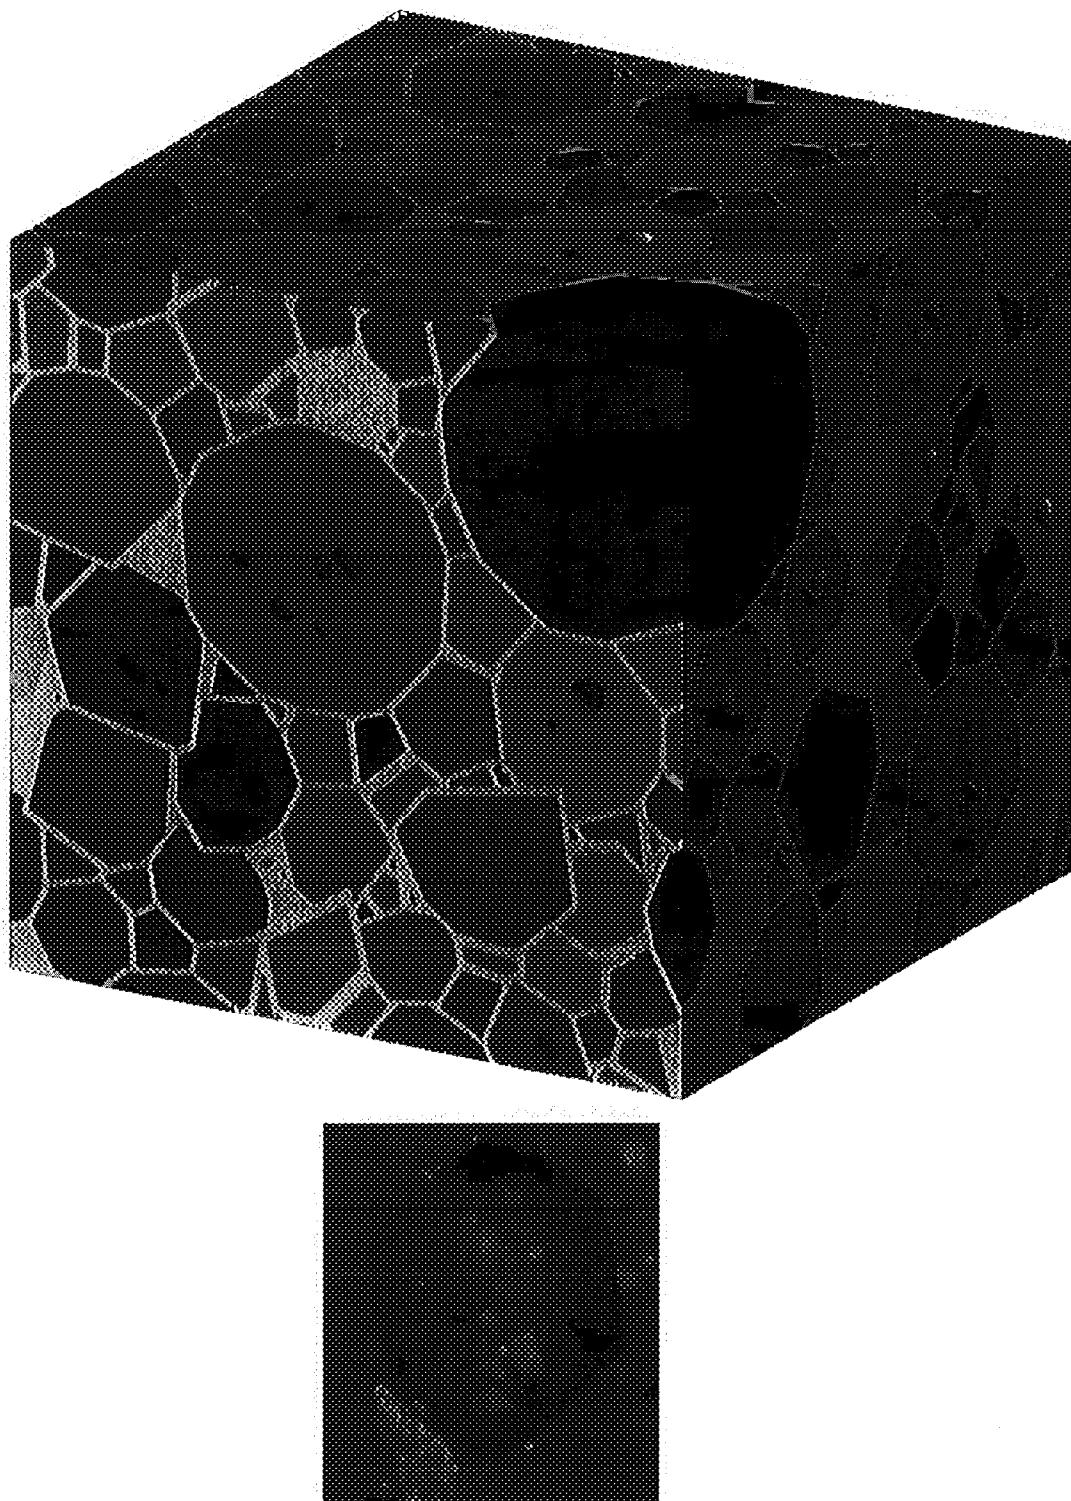


Figure 7. Coarse monomodal distribution in a 300- μm cube with defects in the larger particles (top) and microphotograph of defects in a single 150 μm HMX crystal (bottom).

Shock waves in PBX

Summary

The fundamental idea in the model to describe shock ignition in heterogeneous explosives is that the shock wave collapses a void, leading to temperature localization and the subsequent runaway of an exothermic, thermally activated chemical reaction. The collapse of a void is fundamentally different when the shock front is narrow relative to the void dimension, as compared to collapse by a shock front that is broad relative to the void dimension. We constructed assemblies of explosive particles and binder to examine the structure of a shock after passage over such an assembly.

We determined that on a macroscopic scale the shock front is broad. On the scale of the void, however, the shock is narrow. As a result, void collapse is not well described by a spherically symmetric, quasi-static compression, but by an unsymmetrical collapse from a sharp shock.

After void collapse, a volume of material comparable to the initial void volume is heated above 700 K where prompt reactions that would convert the material to gas products. When chemical reactions are permitted, the completion of our simplified reaction chain to final product spreads rapidly in the heated material not as a deflagration front, but with a phase velocity that tracks the chemical reaction rates in a volume with a temperature gradient.

Simulations of shocks in an assembly of crystals

We created an assembly of crystals from a monomodal distribution in brick-shaped volumes $150\mu\text{m}$ by $150\mu\text{m}$ wide and either $300\mu\text{m}$ or $600\mu\text{m}$ long. We applied a 10 GPa shock that propagates down the long axis of the brick. The lateral faces are roller boundaries. Computational gauge points were located on planes at various distances from the input face, and the normal stress component recorded in these planes at 1 to $3\mu\text{m}$ intervals.

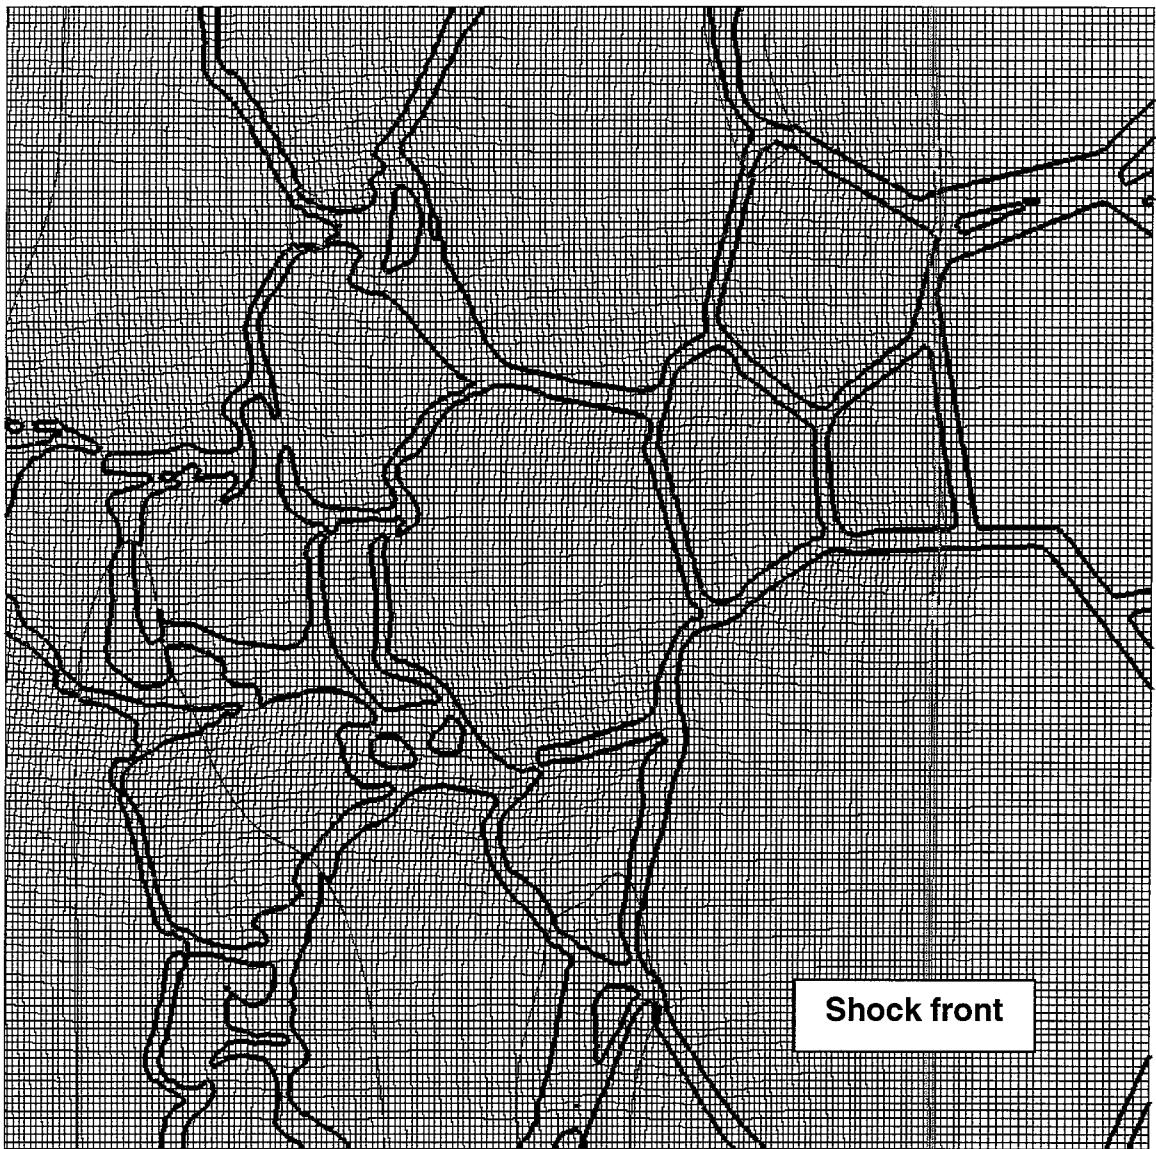


Figure 8. Shock wave in a $150 \mu\text{m}$ by $150 \mu\text{m}$ brick, $300 \mu\text{m}$ long. Binder is HTPB, which is a poor density match for the crystalline HMX. The shock front is seen to be sharp within an explosive crystal. Other calculations with Viton binder, an excellent density match, show similar behavior.

Taking the average response of all the gauge points on a plane simulated the response of a macroscopic stress gauge. In addition, we recorded the average

time-of-arrival of the half-maximum stress, and the standard deviation of the time-of-arrival at each of the gauge planes. We also performed baseline simulations of the propagation of a shock in homogeneous unreacted explosive. In summary, our simulations show that the shock front in each of the crystals of the assembly is sharp. If the stress wave is averaged over a macroscopic plane, the response is a notably longer rise time. This is caused by a statistical distribution of arrival times over the given test plane. At a specific location in that plane, the arrival time depends on the relative amount of binder and crystal that the shock has traversed. When the shock has progressed far enough down the assembly, the statistical distribution becomes steady. We show one cross-section from an assembly that uses a low-density plastic binder (HTPB) to exaggerate the response (Fig. 8). Within an HMX crystal, the shock is sharp, with a rise time that corresponds to the time it takes for the shock to pass over two finite-difference zones, just as it does in homogeneous material.

Collapse of a void

The consequence of these results is that the collapse of a void within such an assembly is the response to a sharp shock, rather than to a slow pressure ramp. It is therefore unrelated to the quasistatic collapse of a spherical hole in a visco-plastic material. Instead, it shows the unsymmetrical response to a sharp shock including the formation of a jet (see Fig. 9). We simulated the collapse of an air-filled void inside an explosive crystal. We use a simplified global reaction scheme that comprises 3 reactions and 4 species for HMX decomposition, in place of the more than 250 reactions and 75 species known to describe low-pressure decomposition of HMX and RDX. In the absence of chemical reactions, the cavity collapses and forms a jet that collides on the downstream wall. The extreme temperature sensitivity of Arrhenius kinetics would result in the chemical reactions effectively stopping at a certain temperature contour. However, in ALE-3D the chemical reactions of each computational zone take place in a well-stirred reactor, so that there is the uncertainty in location of this sharp demarcation by a small number of computational zones, even in the absence of heat conduction.

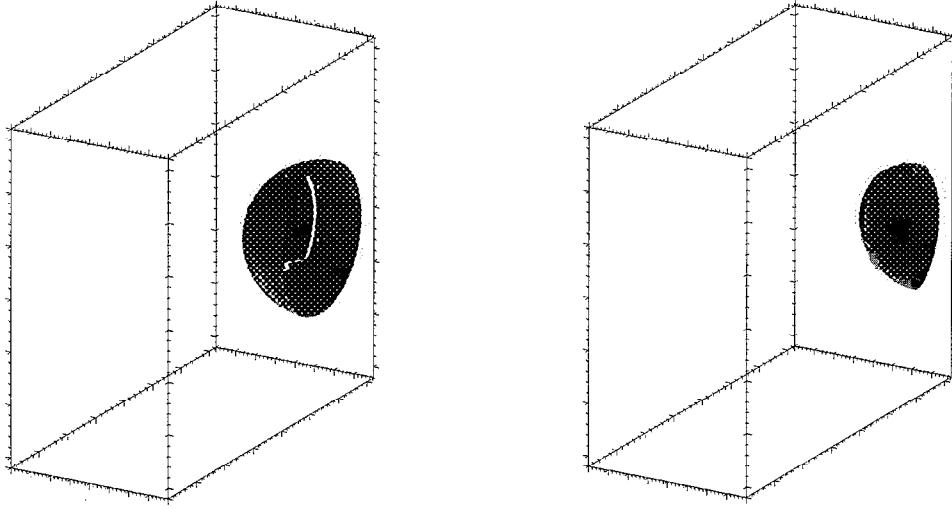


Figure 9. Spherical hole collapsed by a sharp shock before the arrival of the front (left) and during passage of the front (right).

Ignition after void collapse

After passage of a shock front, the void has collapsed, and high temperatures lead to rapid chemical transformation to the product gas mixture. Our simulations at 5 and 20 GPa (Figs. 10 and 11) show a volume of final product comparable to the volume occupied by the 14- μm -diameter sphere. In Fig. 12 we show the results from shock waves of 5, 10, and 20 GPa interacting with the same size imperfection filled with air and filled with water, which represents the mechanical resistance to volume change of a solvent-filled void. For the air filled void, the rates of mass conversion are approximately dependent on the second power of excess compression, μ . These simulations are performed with chemical reactions, but with no heat transfer, so that hydrodynamic scaling laws apply. Our numerical results can be represented by

$$\dot{F} = I \frac{f_v}{r_0} \mu^{2.0}, \text{ where} \quad (1)$$

$$\mu = \frac{\rho}{\rho_0} - 1$$

and where f_v is the volume fraction of void, r_0 is the radius of a typical void, and μ is the excess compression. The coefficient I takes the value 1.04 for time in μs and r_0 in cm.

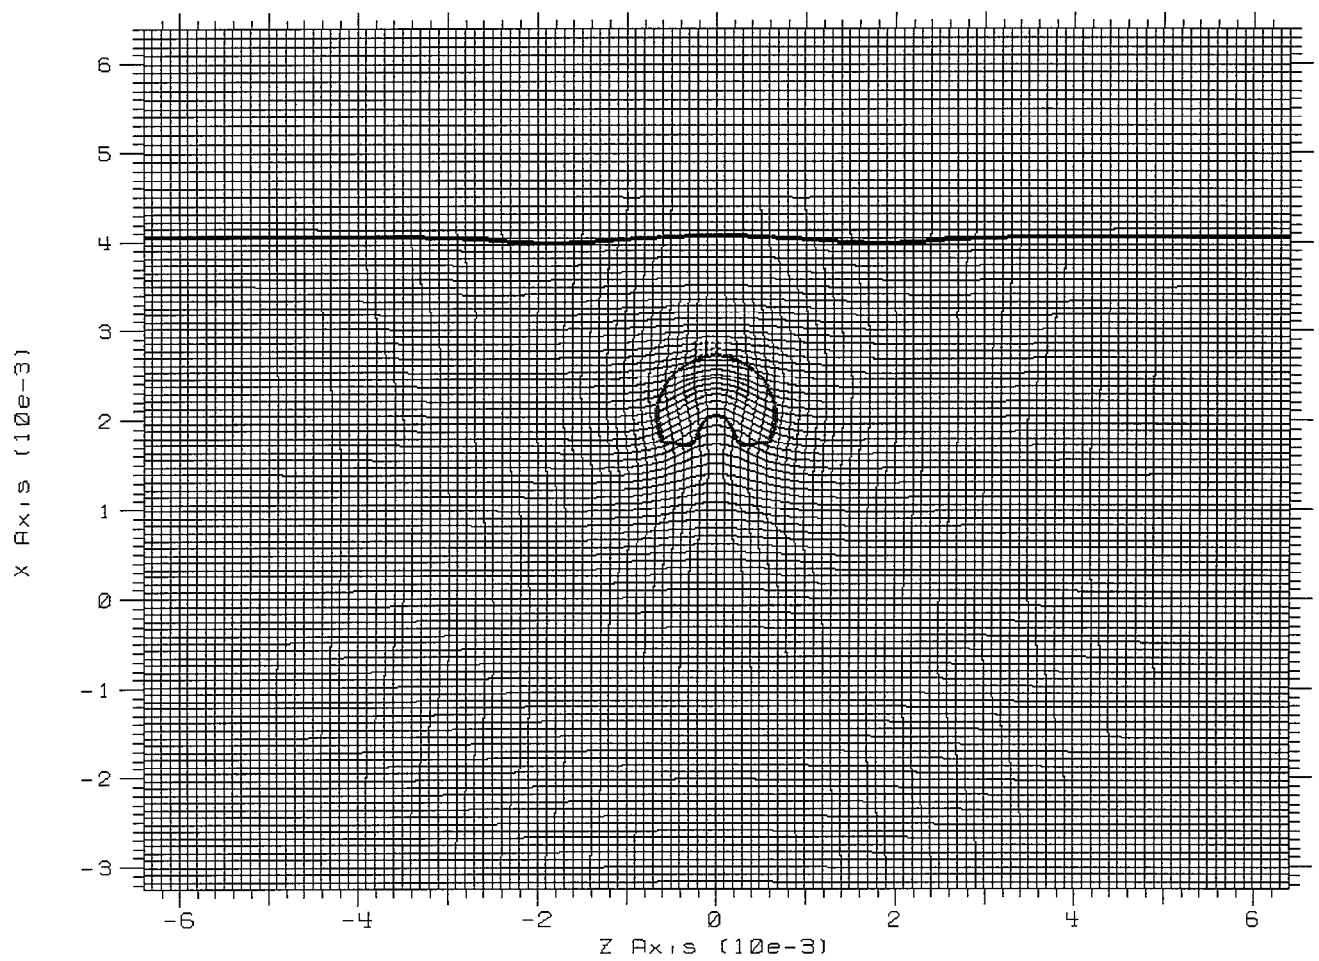


Figure 10. Region occupied by final product in the wake of a 5 GPa shock wave in HMX. Original 14 μm diameter void was located between $x=0$ and $x=14 \mu\text{m}$.

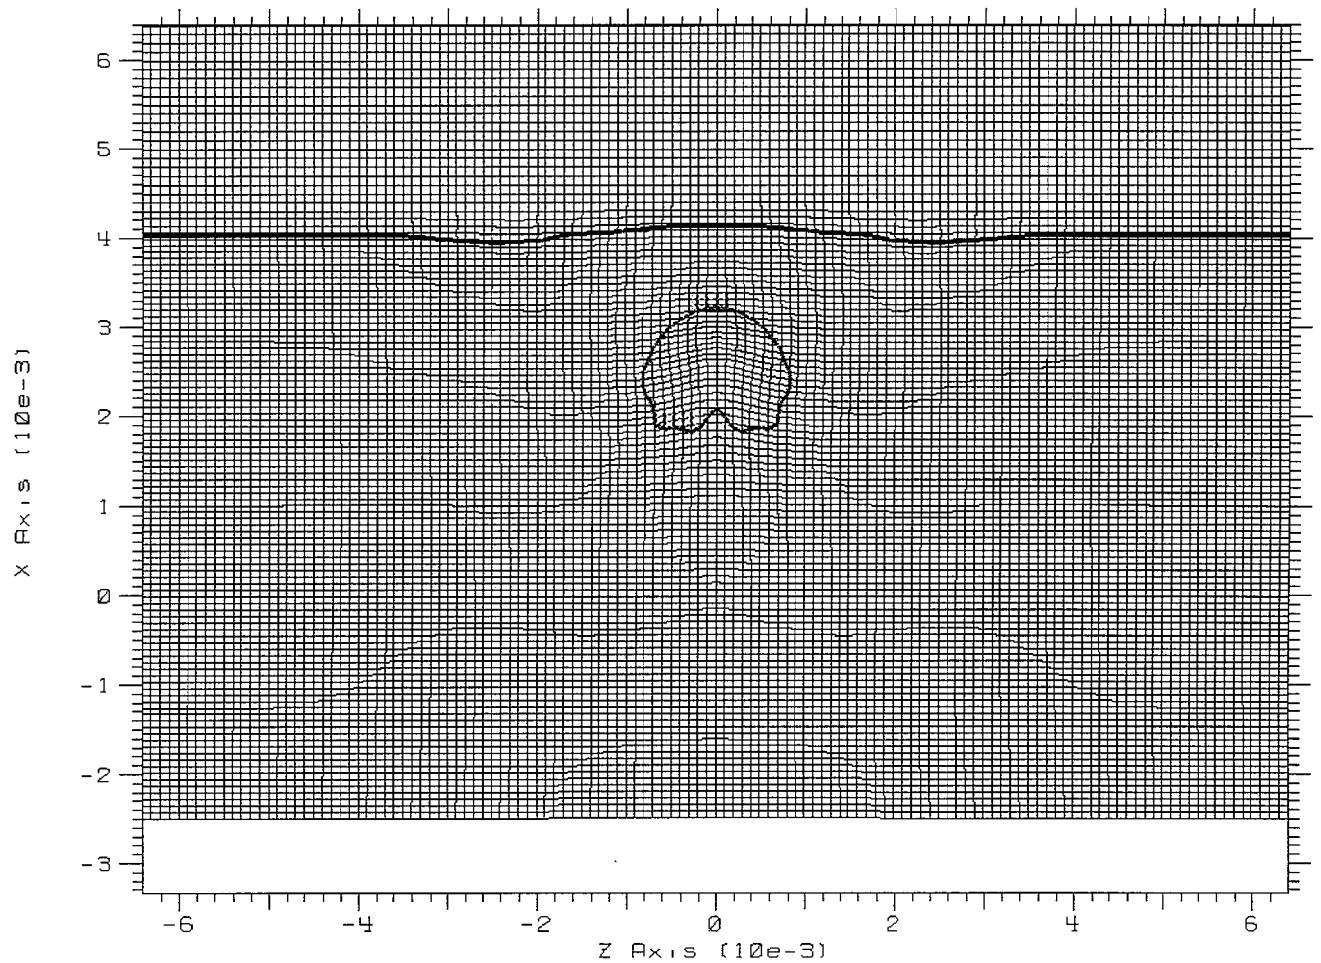


Figure 11. Region occupied by final product in the wake of a 20 GPa shock wave in HMX. Same initial geometry as in Fig. 10.

The results of our simulations overestimate the reaction that occurs in the wake of a shock. Advection is used in these calculations to prevent zone tangling. However, the consequence of advection is that hot final product is mixed with

neighboring cool reactant. At present, the reactant and all its products, both intermediate and final, are at pressure and temperature equilibrium within a given finite-difference zone. As a result, the reactant heats up by the addition of hot final product, and the temperature sensitive reactions occur to make still more product. Our expectation is that, instead, the reaction would slow down and stop. At first the reaction front propagates quickly in the volume of material that is heated strongly. Later, the reaction front apparently spreads to adjacent cooler material, because the reaction rates are slower, and so the appearance of final product is delayed. In effect, the reaction front propagates with a phase velocity that depends on the temperature gradient caused by the shock collapse of the void.

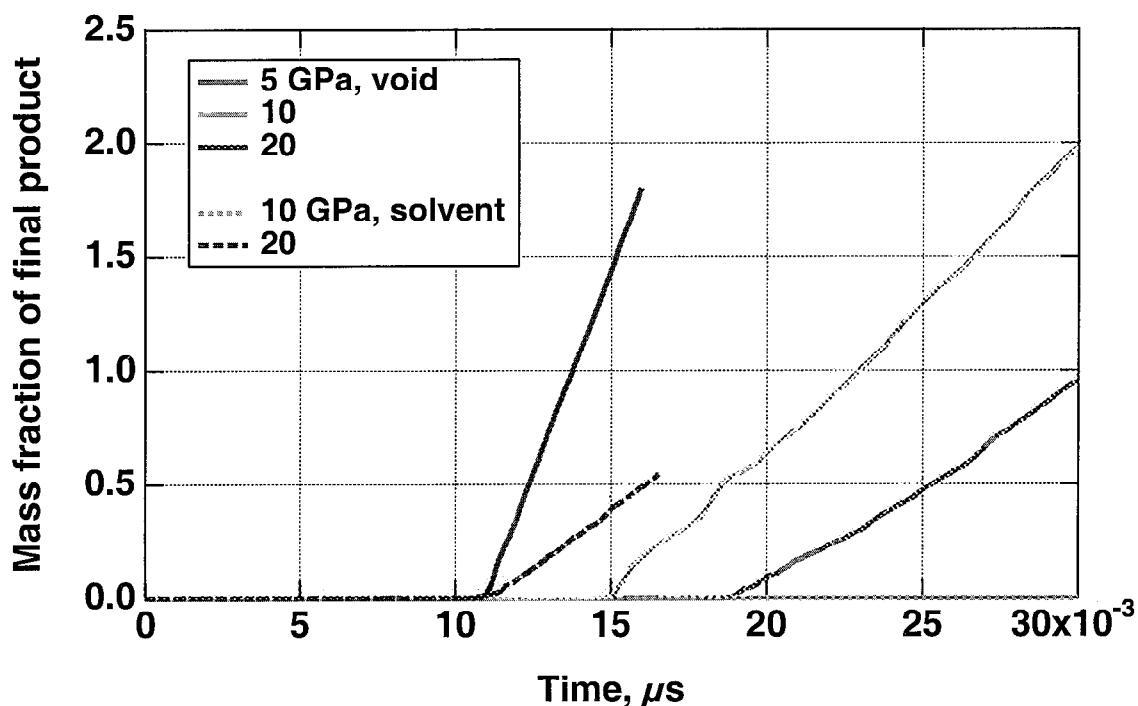


Figure 12. Mass of final product scaled to mass of reactant that would fill the original defect volume. The solvent-filled cavity has a much higher pressure cut-off for reaction than the air-filled cavity does.

We tentatively use the results of our simulations, nevertheless, by the expedient of cutting them off when the mass reacted reaches a mass comparable to that which would fit into the void volume. Further numerical simulations await a detailed study of alternative methods of keeping track of a chemical reaction front in mixed cells.

For the water-filled or solvent-filled imperfection, the localized heating is reduced and the reaction is likewise reduced substantially. In Fig. 12 it is

apparent that solvent-filled holes show a much higher pressure threshold for ignition than is present for air-filled holes or voids. If our preliminary results are correct, then solvent pockets cannot be effective in providing ignition sources for low-pressure events (safety-related ignition). There remains the possibility that a decomposition pathway (which does not exist in our reduced-manifold kinetic rates) is catalyzed by the presence of hot solvent. In such an case, our evaluation of the importance of solvent-filled defects might change.

Deflagration

Summary

In the standard view of reactive flow modeling, isolated hot spots connect with each other by deflagration to form a volume of hot products. This deflagration front moves through the unreacted solid with a velocity that depends on pressure and temperature. We use direct numerical simulation to calculate the velocity of the deflagration. Experiments at high pressure using a diamond anvil cell, and theoretical QMD simulations of HMX decomposition at high pressure and temperature demonstrated that our reaction rates at more than 10 GPa are significantly wrong. Calculated values of the species diffusivity over a broad pressure range have helped us understand our results and point to additional capabilities needed in ALE-3D.

In order to simulate deflagrations at high pressure, we need complete equations of state of the various species involved, thermal transport properties of the species, and the chemical reaction kinetic rates embedded in a computer simulation program.

Chemical kinetics and equations of state

Craig Tarver has used one-dimensional time-to-explosion experiments to obtain a reduced-manifold chemical reaction scheme for the decomposition of explosives. We have used his published 3-reaction, 4-species model for HMX. As with all explosives that are thermally stable enough to be useable, the first reaction step is endothermic, which we tentatively identify with the solid-solid β to δ phase transformation of HMX. Since only one-directional kinetic rates are used, the phase transformation in our simplified scheme does not exhibit the equilibrium pressure-temperature dependence demanded by the Clausius-Clapeyron phase rule. Nevertheless, we give the second species in this step the properties of δ -HMX. We assume that the compressibility is the same as β -HMX, but that the reference density is reduced. We use Bridgman's model for the thermal conductivity of both solid species.

The second reaction Tarver identified as slightly exothermic, producing the high-molecular weight gases CH_2O , N_2O , HCN, and HNO_2 . Sorin Bastea has recently (September, 2002) produced an equation of state with transport properties for an equilibrium mixture of these species. Christine Wu obtained the molecular potentials for these molecules from recent molecular dynamics simulations. All our previous work, and all work described here, has been based on a much less accurate estimate of the equation of state using the older BKWC equation of state form, and the Bridgman model for thermal conductivity. The final step in this

global reaction scheme is the gas-phase reaction to final products. The equation of state for the final products is from CHEQ, although either CHEETAH or CHEQ could be used interchangeably. The transport properties are the more accurate estimates of Bastea.

Equation of state

We use the thermochemical equilibrium computer programs CHEETAH and CHEQ to calculate enough points on the equation of state surface to construct a tabular equation of state in one of this laboratory's standard formats (LEOS). Historically, these thermochemical codes were unreliable estimates of states far removed from the Chapman-Jouget (CJ) adiabat. They were tuned to obtain reliable estimates of detonation properties for CHNO explosives using equation of state forms with few adjustable parameters, and with a concomitantly limited range of validity. More recent developments have led to more realistic exponential-6 molecular potentials and to improved estimates of the equation of state surface for a mixture of molecules with that potential. Although CHEQ and CHEETAH use the same molecular potential, they use different estimates of the equation of state surface for an assembly of molecules with that potential. It is reassuring that these two independent estimates of the equation of state surface are as close to each other as they are.

Both programs are focused on the behavior of HE gas products at high pressure. In particular, they only treat states above the liquid-vapor coexistence region (steam dome). Although states in and near the steam dome are not sampled in simulations of explosives, the algorithm used by ALE-3D to achieve pressure equilibrium occasionally samples portions of the equation of state surface far removed from the equilibrium solution. As a result, a necessary additional chore is to ensure that reasonable values for the equation of state are defined within the steam dome.

Transport properties

Calculation of flame advance also requires knowledge of the transport properties. Recently, Bastea has modified Enskog theory to obtain thermal conductivity consistent with the mixture of spheres represented by the (radial) exponential-6 potentials. The values he calculates at high pressure are typically within a factor of two or so of earlier estimates we made using Bridgman's model for high-density fluids. For future work in this and other programs, he has also calculated the viscosity and species diffusivity.

Simulations of deflagration

To the extent that the deflagration rates are subsonic, the formulation of the numerical problem could be cast as quasistatic. However at the small zone size needed to resolve the deflagration front, the stable time step for explicit finite-difference calculation of heat transfer is comparable to the stable time step for explicit finite-difference calculation of compressible hydrodynamics. As a result, it is numerically efficient to calculate the compressible flow dynamically.

With the properties as outlined above, we performed a series of one-dimensional planar calculations of deflagration in HMX at various constant pressures from 0.1 to 30 GPa using ALE-3D. The results of our simulations show a deflagration rate that increases with both pressure and the temperature of the unreacted solid. See Figs. 13 and 14. We observe that the flames are quite thin, ranging from 300 to 100 nm over the pressure range. To resolve the flame fronts in these calculations requires a mesh size on the order of 1 to 10 nm. Imposing a hot temperature and constant pressure boundary condition on one end, and a constant position, adiabatic boundary on the other specifies the boundary conditions for these transient calculations. We carry out the calculations long enough that the flame propagation has a steady velocity and the flame profile is unchanging. The results of our calculations at room temperature are summarized by

$$v = 2p^{0.5} \text{ m/s}$$

where the pressure, p , is in GPa. The pressure exponent is relatively small compared to experience at low pressure (<1 GPa) with a variety of nitramine-containing propellants, although the value at low pressure is within the range commonly reported with both strand burner and closed bomb tests.

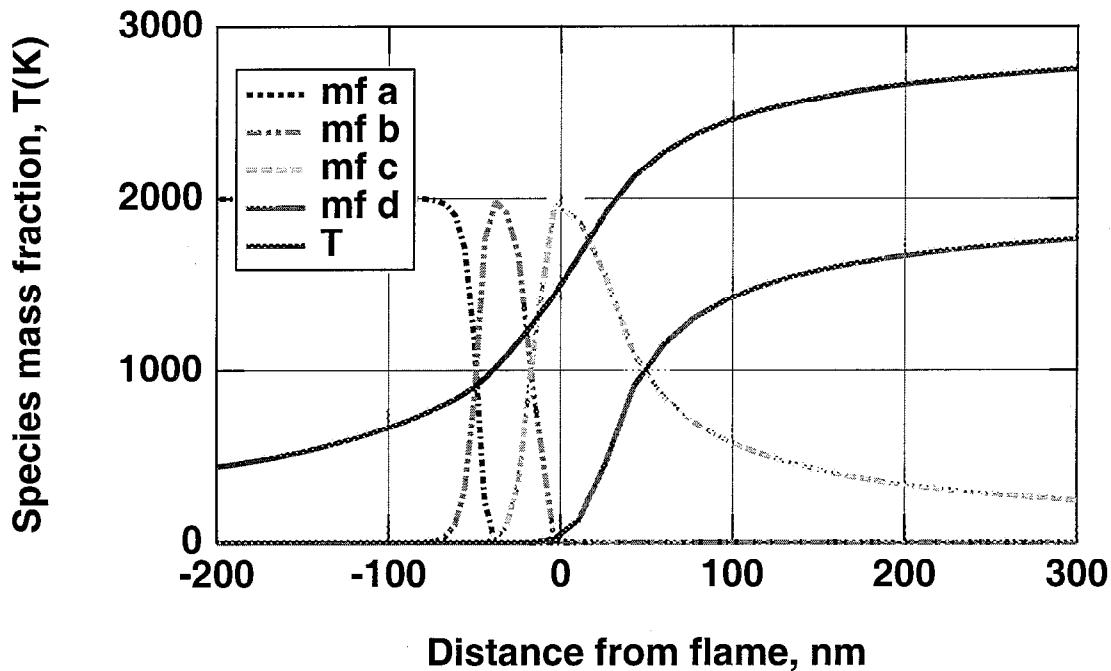


Figure 13. Direct numerical simulation of a flame at 1 GPa. The mass fractions of the four species are scaled so that the value 2000 represents 100 %.

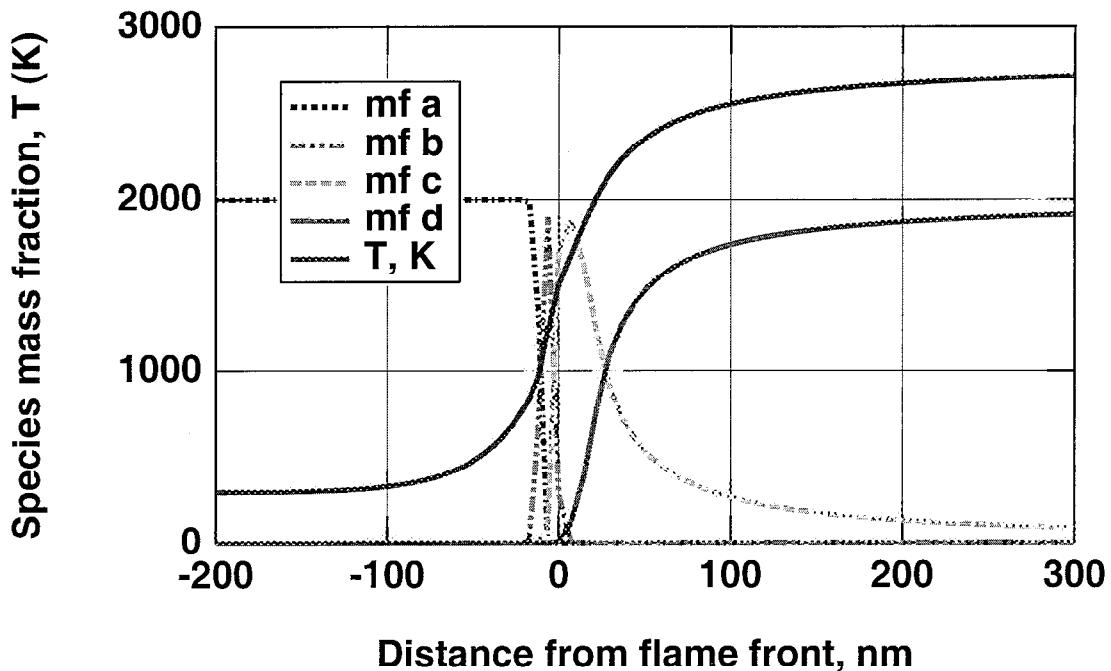


Figure 14. Direct numerical simulation of a flame in HMX at 12 GPa. Other conditions the same as in Fig. 13

Measurements of deflagration

Joe Zaug and his coworkers have performed diamond anvil cell (DAC) measurements with the support of the Munitions Technology Development Program to extend the range of the measured flame speed for HMX to over 30 GPa. In the DAC the sample disk is nominally 140- μm in diameter by 50 μm thick. It is ignited at the center with a focused pencil of laser light approximately 2- μm diameter. A schematic of the experiment is shown in Fig. 15. The flame spreads cylindrically, and the location of the flame front in time is recorded with a streak camera. The observed advance of the flame front is a combination of the motion of unreacted explosive away from the burned volume as the pressure increases and the unreacted solid compresses, and the motion of the flame front through the solid. For a constant volume system,

$$\rho_s(R^2 - r^2) = \rho_{s0}R^2(1 - F) \quad (2)$$

where ρ_s is the changing density of the unreacted solid, R is the radius of the hole in the DAC gasket, r is the location of the flame front, ρ_{s0} is the density of the unreacted solid at pressure, but before ignition, and F is the mass fraction transformed to product gas.

Taking the derivative of Eq. 2 with respect to time,

$$\dot{\rho}_s(R^2 - r^2) - 2\rho_s r \dot{r} = -\rho_{s0}R^2 \dot{F} \quad (3)$$

where \dot{r} is the measured flame speed. The value of \dot{F} is given by

$$\dot{F} = \frac{\rho_s v (2\pi r)}{\pi \rho_{s0} R^2} \quad (4)$$

where v is the velocity that the flame advances in stationary solid reactant, corresponding to our calculated flame speed. The correction to translate from one measure to the other is

$$\dot{r} = v + \frac{1}{2} \frac{\rho_{s0}}{\rho_s} \frac{\dot{\rho}_s}{\rho_s} \frac{R^2}{r} (1 - F). \quad (5)$$

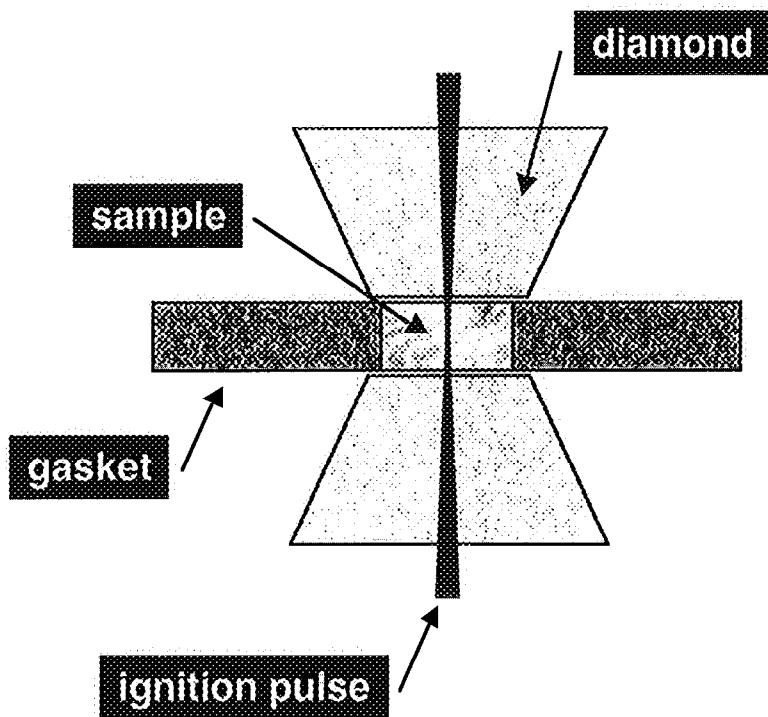


Figure 15. Schematic of the diamond anvil cell experiments to measure the high-pressure deflagration rate. The ignition pulse starts a cylindrically symmetric deflagration.

The density of the solid is pressure dependent, because the solid is compressed adiabatically. The velocity is also pressure dependent, and in our simulations the pressure is a smooth and nearly linear function of the mass fraction reacted, F .

The correction to our calculated flame speed, v , is seen to be 20% when the measurements are taken near the beginning of the burn, before the reaction front has progressed through half the radius. See Fig. 16. Nevertheless even our corrected estimate of the DAC flame speed remains significantly less than the measured value at high pressure. See Fig. 17.

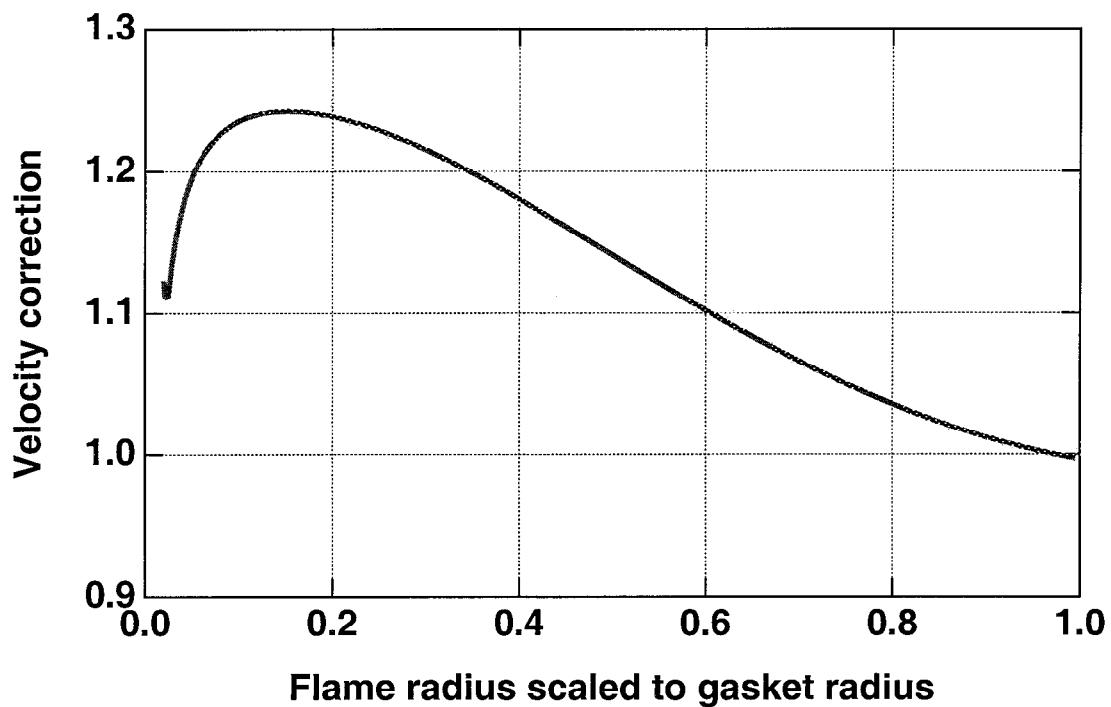


Figure 16. Velocity correction as a function of scaled radius of the flame location.

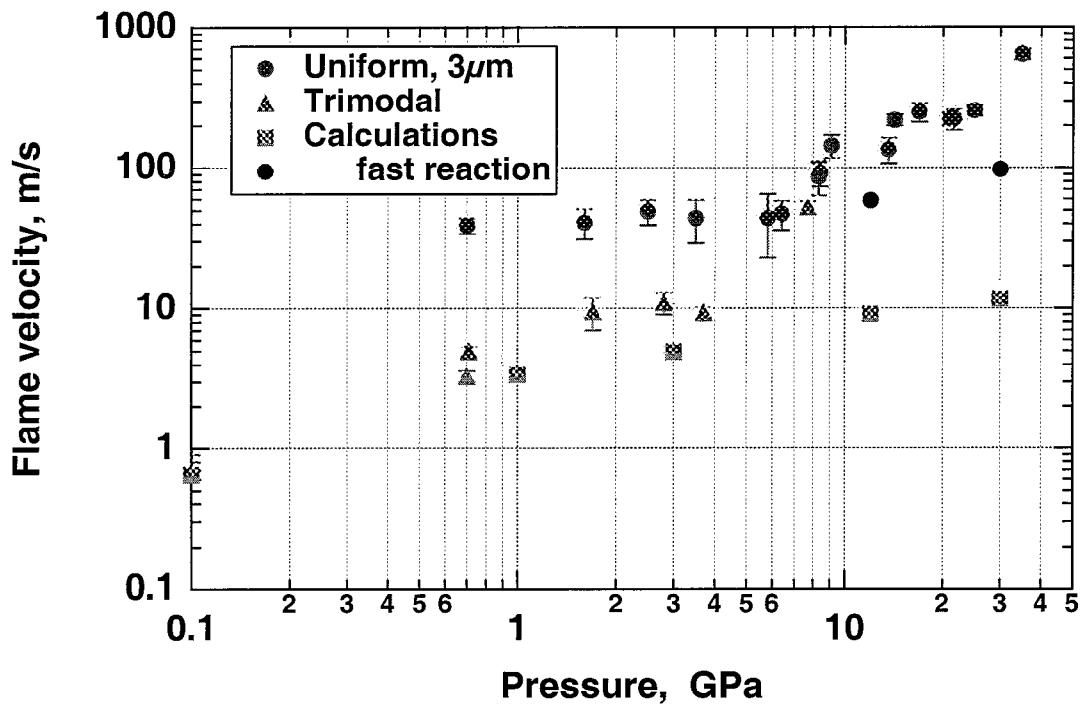


Figure 17. Experimental and corrected computational flame speeds as a function of pressure. The uniform grain size data are incompletely compacted below 6 GPa, so the observed values are the speed of gas products diffusing ahead in a porous medium.

Kinetic rates at high pressure

Riad Manaa has performed quantum molecular dynamics simulations of the decomposition of HMX at high pressure and temperature (3500 K, density 1.9 g/cc). When we compare the results of his simulations for the concentrations of typical species in the final products to those that result from the kinetic rates we have used in our simulations, it is apparent that our rates are much too slow and that a single rate does not represent both the appearance of N_2 and of H_2O (Fig. 18). We performed a few direct numerical simulations of the flame speed using rates that approximate those found by Manaa by the expedient of assuming the same activation energy that we had been using for the final reaction step, and fitting the rate at 3500 K. Upon doing so, our high-pressure deflagration speed is much closer to the measured values (Fig. 17). However, the low-pressure results are also affected, so that we no longer predict values of the deflagration speed in accord with experiment there. Further, the use of these higher rates in simulations of the ODTX experiments gives a much more thermally sensitive result than is measured.

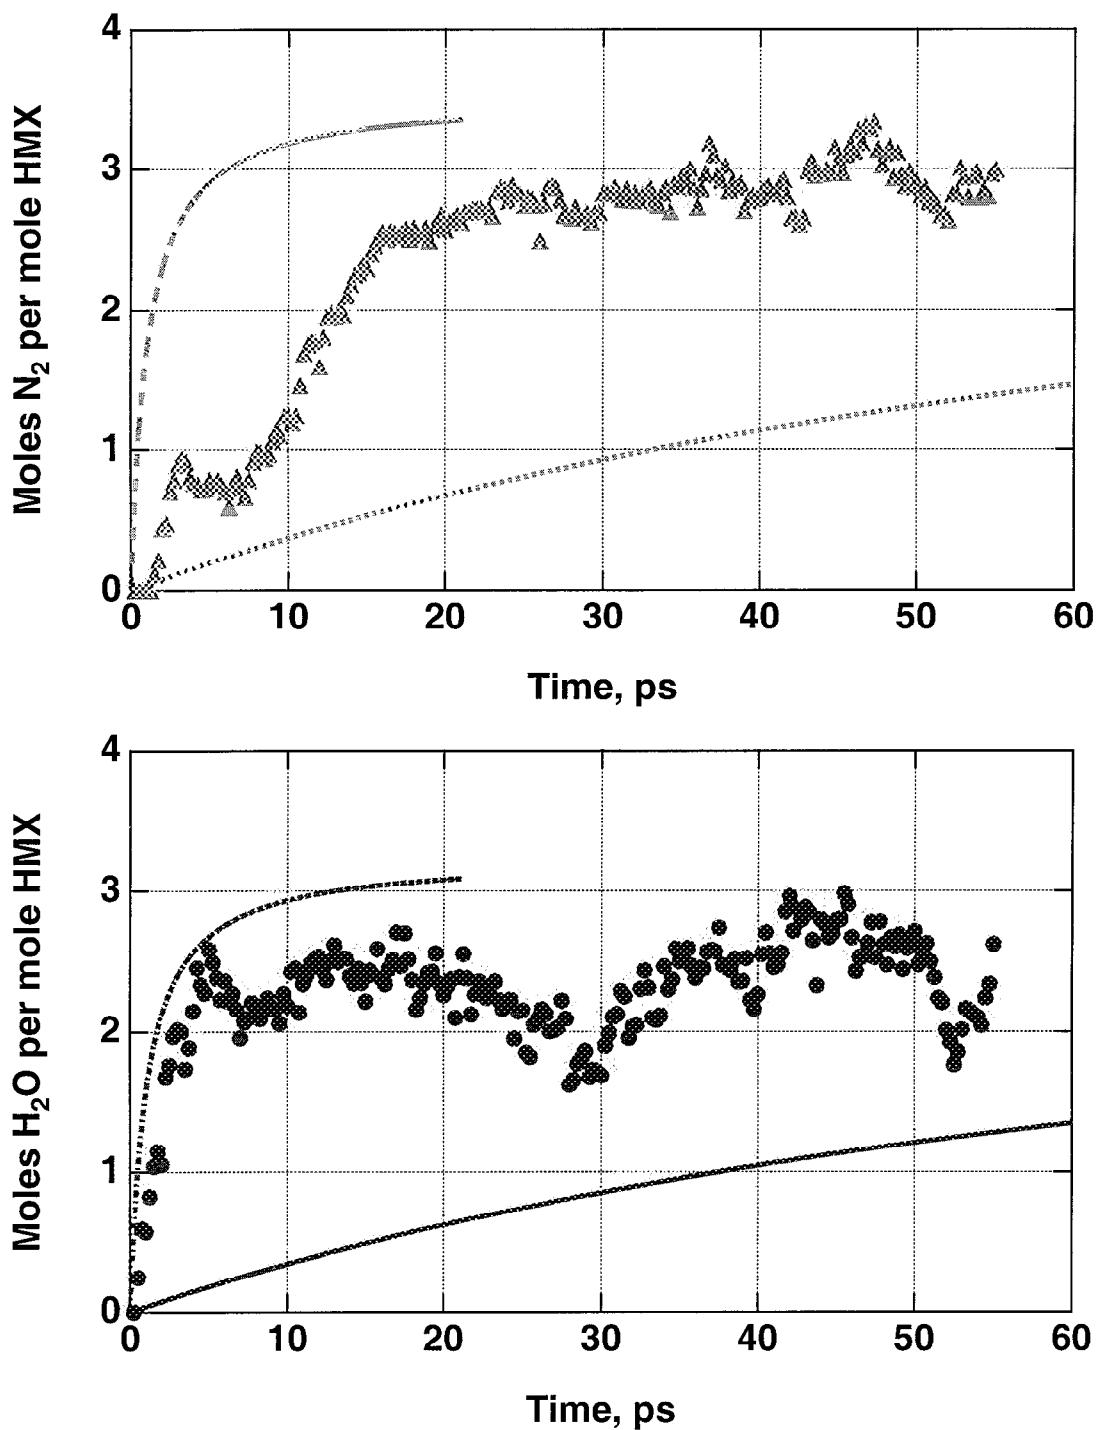


Figure 18. Appearance of nitrogen (top) and water (bottom) in the QMD simulations of Manaa (symbols) and calculated with the kinetic rates using our simplified chemistry for HMX decomposition (lines). In both figures the slower chemistry is the original kinetics from low-pressure experiments.

Resolution of the differences

The transport properties of the products as calculated by Bastea furnish us with a clue as to how the differences between the kinetic rates at low pressure and high pressure may be resolved. The thermal diffusivity and the molecular diffusivity are strongly coupled when statistical mechanics is used to obtain their values.

Comparison of the two diffusivities on the 3000 K isotherm as a function of pressure (Fig. 19) illustrates that there is a significant difference between the two properties at high pressure, and no essential difference at low pressure. None of our previous simulations incorporates species diffusion, even those at low pressure. Zel'dovich and Barenblat formulated and numerically solved the equations of flame propagation for a constant density flame for which the ratio of species diffusion to thermal diffusion is a parameter. In one dimension it is a simple matter (nowadays) to solve their equations numerically. The result is that there is a large difference in the calculated flame speed when thermal and species diffusion have ratio 1 (low-pressure) and when the have a large ratio (high-pressure).

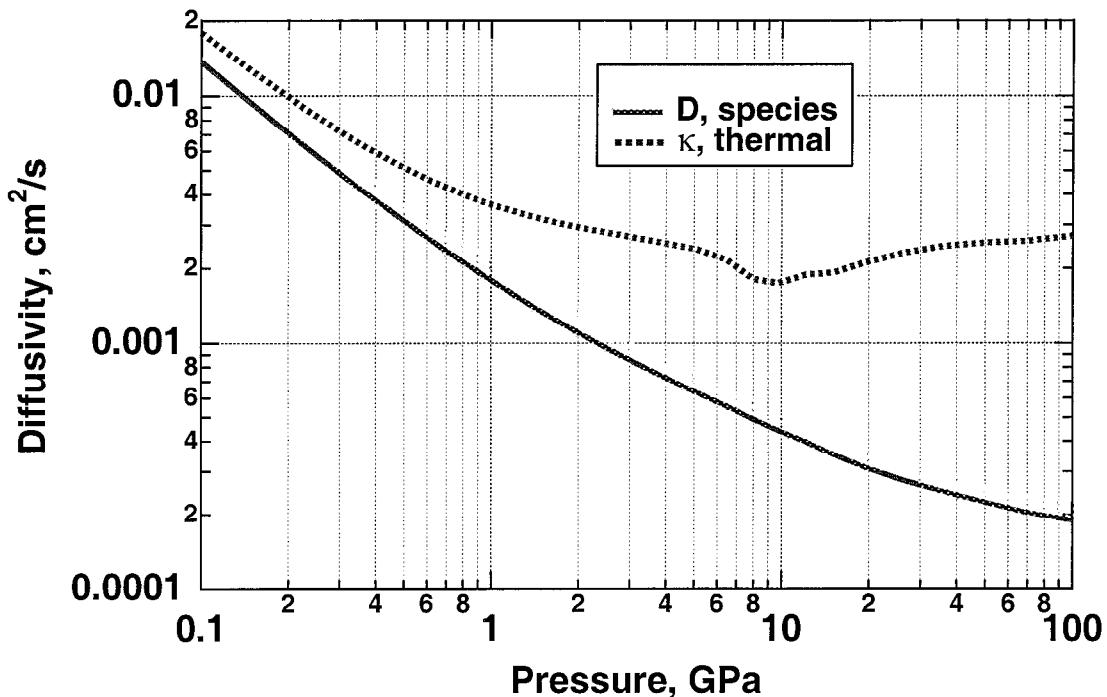


Figure 19. Species diffusivity (solid) and thermal diffusivity (dot) on the 3000K isotherm for HMX final products.

The low-pressure kinetic rates were fitted to low-pressure experiments by assuming species diffusion does not take place. As a result, the kinetic rates are artificially low to compensate for the high concentration of reactant that would

be dispersed by diffusion. So long as all the calculations are performed at low pressure, where the ratio of species and thermal diffusion is constant, the results would track experiment. It is only when we try to apply those rates at high pressure, where the species are actually inhibited from diffusing, that we get into trouble.

Future work

To demonstrate our inference that the pressure dependence of species diffusion is a significant influence on our calculated flame speeds, we must perform calculations at low pressure to refit the kinetic rates, and then perform calculations at high pressure with those rates. This awaits the addition of species diffusion at large concentrations into ALE-3D. At present, the physics is there for tracer species diffusion, but not for significant concentration changes that influence other properties.

Growth of hot spots in PBX

Summary

Once a hot spot is formed, it can either grow or decay, depending on its size. Small spots have sufficient surface area that more heat is lost by conduction to the surrounding cool, unreacted explosive, than is gained by conversion to product gases. The deflagration is a very subsonic process, so that hot spots must be relatively closely spaced in order that the reaction can complete within the observed reaction zone. When hot spots are isolated, a simple model represents the main features, and a modification to that model is in excellent agreement with detailed computer simulations. We rely on computer simulations of a constellation of randomly placed and interacting hot spots to develop the form factor needed for a continuum reactive flow model.

Simplified model of hot spot growth

The advance of a thin flame front through the explosive is accompanied by the conversion of mass from reactant to final product. This same idea is used to describe the mass conversion in the propellant beds of guns and in solid propellant rockets. These latter two fields deal with modest pressures, less than 0.01 GPa for rockets and less than 0.3 GPa for guns. There the change in density of the solid is typically ignored, and the products behave as nearly ideal gases. In all cases, the mass-burning rate is given by

$$\dot{m} = \rho_{s0} v \sum_{i=1}^n A_i(t) \quad (6)$$

where ρ_{s0} is the unreacted (solid) explosive density, v is the pressure- and temperature-dependent deflagration velocity, and A_i is the time-dependent surface area for one of the n isolated hot spots.

Early in the deflagration process, when the hot spots are isolated from each other, a sphere can approximate the surface area of each spot. If we further assume that an effective radius, r_0 , can replace the distribution of hot spot sizes then

$$\dot{m} = 4\pi n \rho_{s0} v (r_0 + vt)^3 \quad (7)$$

where t is the time after the ignition event, n is the number of hot spots and r_0 the original radius of the hot spot. The total mass of the explosive is given by

$$M = \frac{4}{3}\pi n \rho_{s0} f_v^{-1} r_0^3 \quad (8)$$

where f_v is the volume fraction of hot spots, which is approximately the same as the volume fraction of voids.

The time rate of change of F , the mass fraction reacted, which is the progress variable of ignition and growth models, is given by

$$\dot{F} = 3f_v^{1/3} \frac{v}{r_0} \left(1 + \frac{v}{r_0} t \right)^2. \quad (9)$$

Integration of the rate equation at constant flame speed, v , results in

$$F = f_v \left(1 + \frac{v}{r_0} t \right)^3, \quad (10)$$

so that

$$\dot{F} = 3f_v^{1/3} \frac{v}{r_0} F^{2/3}. \quad (11)$$

In reactive flow models, the form factor is given by

$$F^a (1 - F)^b$$

so that as the reaction nears completion, the rate decreases.

Corrections to the simplified model

If we assume that the deflagration takes place at constant superficial density (behind a shock front) the pressure increases as the burn front advances, as so the density of the unreacted material increases. This results in a net motion away from the zone containing the final product. The motion of the burn front, then, is a combination of the motion of the burn front in stationary material, and the motion of the material due to densification.

For all geometries,

$$\dot{F} = \frac{\rho_s v A}{M}. \quad (12)$$

In spherical geometry,

$$M = \frac{4}{3} \pi \rho_{s0} R^3 \quad (13)$$

where R is the dimension of the unit cell, and

$$r_0^3 = f_v R^3. \quad (14)$$

By conservation of mass,

$$\rho_s (R^3 - r^3) = \rho_{s0} R^3 (1 - F) \quad (15)$$

where r is the location of the burn front. Substituting,

$$\dot{F} = 3f_v^{1/3} \frac{\rho_s}{\rho_{s0}} \frac{v}{r_0} \left[1 - \frac{\rho_{s0}}{\rho_s} (1 - F) \right]^{2/3} \quad (16)$$

Similarly for cylindrical and planar geometry,

$$\dot{F} = 2f_v^{1/2} \frac{\rho_s}{\rho_{s0}} \frac{v}{r_0} \left[1 - \frac{\rho_{s0}}{\rho_s} (1 - F) \right]^{1/2}, \text{(cylinder)} \quad (17,18)$$

$$\dot{F} = f_v \frac{\rho_s}{\rho_{s0}} \frac{v}{r_0}, \text{(plane)}$$

Growth of a single hot spot in ALE-3D

We first examine the growth of reaction for a single isolated hot spot. We use the deflagration model that describes the location of the burn front with a level-set method. For a single hotspot, the ALE-3D result and the corrected model are in excellent agreement (Fig. 20). The correction is seen to be approximately a 15% effect in the calculated value of the rate of reaction progress. In the model, we find the density history by calculating the adiabatic compression of the unreacted explosive from the initial pressure and temperature using the equation of state of the solid reactant. The pressure history is fit from the ALE-3D results as a quadratic in the progress variable, F .

$$p = p_0 + F[1 + \delta(1 - \delta F)](p_f - p_0) \quad (19)$$

where p_0 is the initial pressure and p_f the final pressure. The final pressure is just the pressure of the final products in a constant volume (constant density) explosion. We found the pressure increase from initial to final to be about 27 GPa for initial pressures of both 10 and 20 GPa. The value of δ was 0.3 and 0.47 for those same two starting pressures.

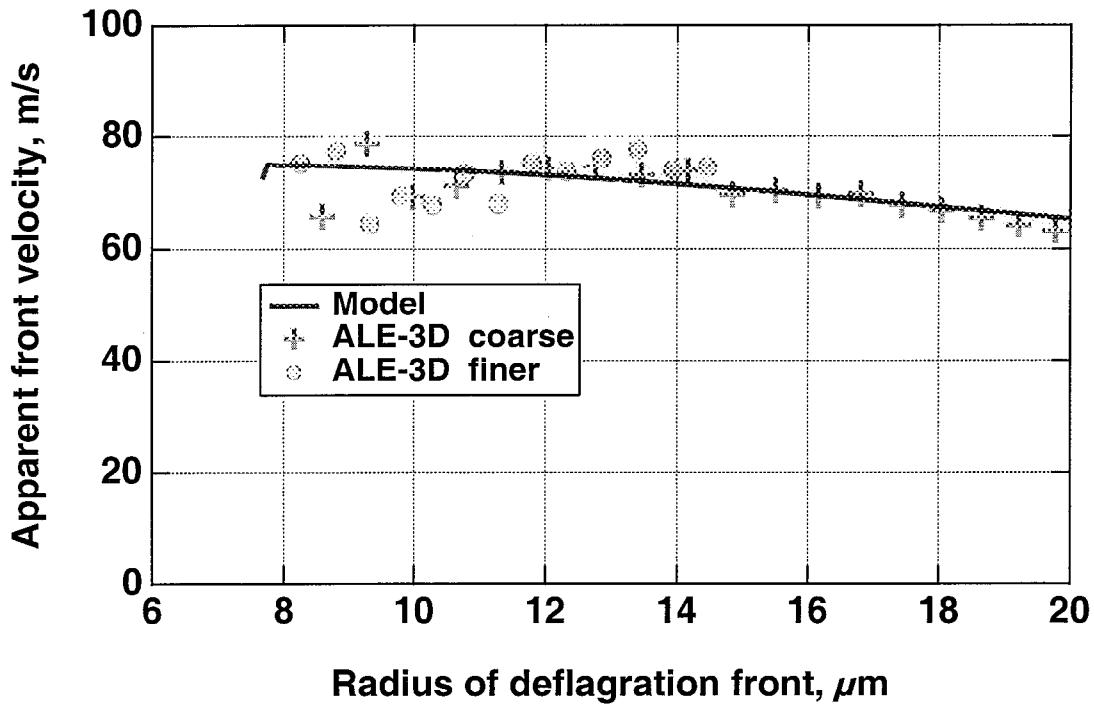


Figure 20. Model result for apparent velocity of deflagration front and front velocity from ALE-3D For both 0.5 and 0.67 μm zones, the agreement is good after the initial transient in ALE-3D. The input deflagration velocity is 60 m/s.

Growth of an assembly of hot spots

When the hot spots are isolated, for example a regular cubic array of hot spots, the initial behavior has precisely the same functional form as does the single void. The factor $3/r_0$ represents the initial surface to volume ratio of the hot spots.

When the hot spot centroids are located randomly, with no special care taken to forbid the overlap of neighbors, some hot spots will overlap from the beginning. As a result, the surface to volume ratio will increase over that of isolated spheres, so that the initial rate will be somewhat higher and the power law a will no longer be $2/3$.

For gun and rocket propellants, where the geometry is known, the form factor can be obtained analytically or numerically. For random arrays of hot spots, we rely on numerical simulations and fit the results to simple functional forms. For spheres burning from the outside in, the form factor is obtained analytically as

$$(1 - F)^{2/3}$$

The result of our simulations for the array of hot spots shown in Fig. 21 is that the form factor is given by

$$F^a(1-F)^b$$

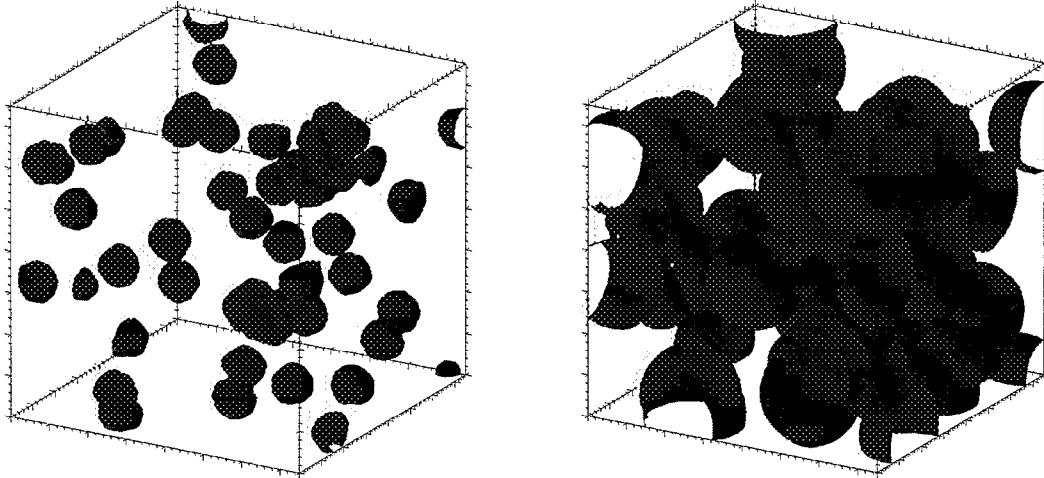


Figure 21. Growth of reaction from a number of $14 \mu\text{m}$ diameter hot spots in a $150\text{-}\mu\text{m}$ cube at early time (left) and at later time (right). Red spheres mark the flame fronts and enclose final product gas.

We performed a family of simulations of the deflagration of clusters of hot spots varying the initial hot spot volume fraction, diameter, and the deflagration velocity. The results of those simulations can be represented by

$$\dot{F} = Af_v^{1/3} \frac{3}{r_0} v F^a (1-F)^b, \quad (20)$$

where the prefactor A changes with the values of a and b . The parameter values for four such simulations are given in Table 1.

Table 1. Parameter values for the growth of reaction. Time is in μs .

Cube side, μm	Volume fraction	Hotspot radius, μm	A	a	b	Model symbol
300	0.0336	7	1.04	0.63	0.70	A
150	0.0190	7	0.78	0.54	0.45	
150	0.0356	5	0.73	0.48	0.64	B
150	0.0187	7	1.27	0.69	1.00	

The form factors from those simulations are shown in Fig. 22. Most of the scatter in these results, which were independently fit to the target functional form, are due to the scatter in the distance to neighbor hot spots in different random assemblies. This is noticed when the form factors are plotted (Fig. 22). It is observed that in the early stages of the reaction, all calculations give similar results. It is in the completion side of the form factor that the scatter is evident. As a result, we have chosen the values obtained from the largest cube (Model A in Table 1) to be our estimate of the form factor that results from our grain-scale simulations, and use the values of Model B for comparison. The form factors from our simulations show an initial transient at the beginning of the reaction (near $F=0$). This is caused by an idiosyncrasy of the deflagration model in ALE-3D. The hot spots are ignited numerically by the artifice of a heat source in the volume encompassed by the hot spots. When the temperature of the reactant reaches the preset trigger value, all material in all the zones so affected are ignited at once. This produces an artificially high rate at first. After the initial transient, the deflagration progresses smoothly.

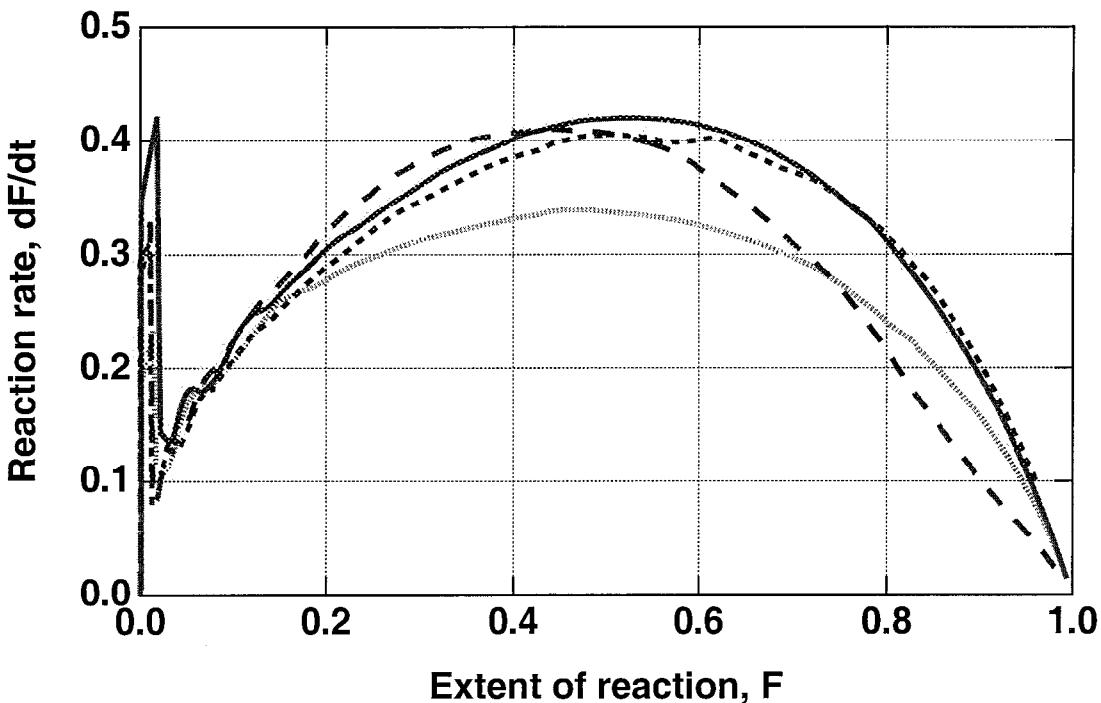


Figure 22. Calculated form factors for arrays of hot spots with differing radii and volume factor. Curves have been scaled by $3vf_v^{1/3}/r_0$

Application of initial continuum model

Summary

We have used a reactive flow continuum model to describe the behavior of an HMX explosive in small diameter, unconfined charges. The baseline model parameters are unpublished values from Craig Tarver that result in reasonable fits to the measured change in detonation speed with charge diameter. We also use values for the parameters taken from our grain-scale simulations to compare with the baseline results.

Baseline results

We used approximately 7-diameter long explosive cylinders as our computational acceptor charges, which varied between 1 and 4 mm diameter. The donor charges were all 10 mm long, and had the same diameter as the acceptor charge, but were detonated as an ideal explosive. As a result, the acceptor charges are overdriven. We observed computationally that for those diameters in which a detonation propagated, the propagation speed decreased down the length of the charge and reached a steady value within 4 or 5 diameters of propagation. In all simulations, the radial and axial zone size was the same, 0.02 mm. Tarver and we confirmed that the results did not change when 0.01 mm zones were used. The increase in detonation speed with increasing charge diameter (Fig. 23) is in substantial agreement with the experiments performed at LANL by Campbell and Engelke.

Results with parameters developed from grain-scale simulations

We use both the ignition model (Eq. 1) and the form factor model (Eq. 20) to obtain the reactive flow parameters for PBX-9404. We used a power-law fit to the diamond anvil cell experiments to obtain the deflagration velocity as a function of pressure.

$$v = 1.05p^{1.726} \text{ m/s}$$

where the pressure, p , is in GPa. The volume fraction of hot spots is taken to be the void fraction of PBX-9404, $f_v=0.01$, which we obtain from the ratio of the theoretical maximum density (TMD) of the explosive to the measured density. Whereas Tarver's model can be summarized by

$$\begin{aligned} \dot{F} = & H\mu^d(1-F)^e, F < 2f_v \\ & + GF^a(1-F)^b p^c \end{aligned} \tag{21}$$

the model that results from our grain scale simulations has the same form, except that we identify the parameters H and G by

$$H = I \frac{f_v}{r_0}, \quad (22)$$

$$G = A f_v^{1/3} \frac{3}{r_0} v_0$$

and where v_0 is the deflagration velocity at unit pressure.

The results of our project have reduced the number of free parameters to the single one, r_0 , which represents the characteristic dimension of the voids and subsequent hot spots in the explosive. Although it would be more pleasing to have established a value for the void/hotspot dimension by examining the as-pressed explosive, at present we treat that value as a free parameter. We show the results for models A and B (Table 1.) using values for the defect radius, r_0 , of $1.7\mu\text{m}$, and also two variants of the model A result for smaller ($1.1\mu\text{m}$, A1) and larger ($2.2\mu\text{m}$, A2) values of the void dimension. We list the parameter values for Tarver's original set, and for our models in Table 2. The ignition parameters of Tarver and our new ignition parameters have about the same effect for this geometry. Tarver reported that the high power (value of d) was needed to capture the initiation behavior for thin flyer plates.

Table 2. Parameter values for reactive flow models for PBX-9404. Units are g, cm, μs , Mbar.

Model	H	d	e	G	a	b	c	r_0
Tarver	7.43E11	20	0.667	1940	0.667	0.667	2	
A	65	2	0	1215	0.629	0.7	1.726	$1.7\mu\text{m}$
A1	7.43E11	20	0.667	1820	0.629	0.7	1.726	1.1
A0	7.43E11	20	0.667	1215	0.629	0.7	1.726	1.7
A2	7.43E11	20	0.667	908	0.629	0.7	1.726	2.2
B	7.43E11	20	0.667	914	0.476	0.637	1.726	1.7

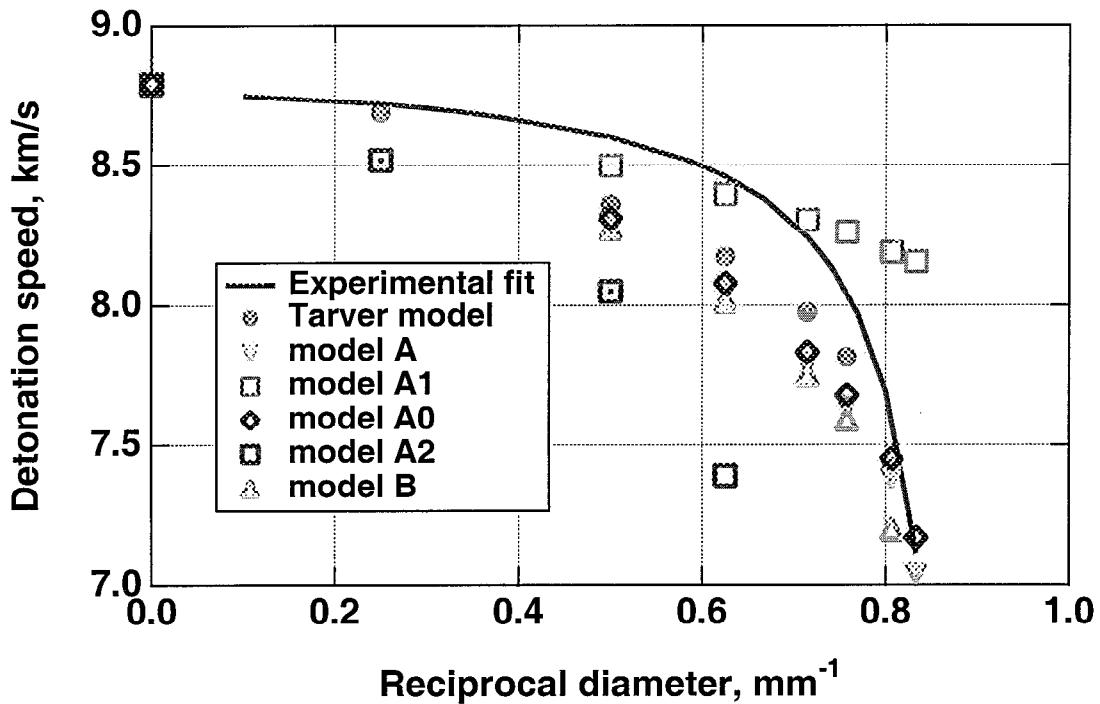


Figure 23. Detonation speed as a function of reciprocal diameter for Tarver's parameters (solid dot) and the new parameters (other symbols) compared with the experimenters' fit to their data (solid).

Significance of r_0

The value of r_0 that we infer for the hot spots is much smaller than the imperfections seen by Hoffman in HMX crystals (Fig. 7). It corresponds to the smallest of the grains in our bimodal distribution (Fig. 6). It is possible that the pressing process introduces this size imperfection. However, the molding powder beads from which the parts are pressed are much larger than this size, of order mm. Although voids on the order of μm might be introduced by pressing, they would be separated by dimensions on the order of mm, and so be relatively rare. This size defect must be present within the individual molding powder beads.

The imperfections seen by Hoffman present their own difficulty of interpretation. Defects of that size, about $20\ \mu\text{m}$, can only be contained in crystals significantly larger than that. As a result, there are significant regions that are defect-free (See Fig. 6) This requires, in turn, that hot spots must connect up by deflagration in not only the crystalline explosive, but also through regions of defect-free but binder coated crystals. We presume that this would result in a lower deflagration

velocity, and so we would infer a still-smaller characteristic dimension for PBX-9404. For LX-04, our nominal PBX, the Viton binder has higher volume content, and is more resistant to decomposition than the PBX-9404 binder. We expect the influence of binder to more significant for LX-04 than for PBX-9404. In support of that point, we note that the reaction zone for LX-04 is approximately 0.6 mm. For our (and Tarver's) simulations in PBX-9404, the reaction zone as measured by edge lag is approximately 0.2 mm. We need to measure the deflagration in HMX and/or other explosive crystals in the presence of binder. We also need to visualize the defects in as-pressed explosive to determine their morphology, size, separation, and to determine whether they are gas-filled or liquid-filled.

Summary and Conclusions

We derived a mixture rule based on numerical simulation of the growth of reaction by flame spread from ignition sites. The rule is to treat computational zones containing both unreacted explosive and product gases as a mixture at pressure equilibrium, but not at thermal equilibrium. As part of that work we demonstrated that for a JWL product equation of state, the choice of mixture rule does not affect the result. We performed simulations of the growth of reaction from hot spots at the grain scale. From these simulations we inferred the form factor to be used in our initial continuum model. We used the measured deflagration rate in our initial continuum model because our direct numerical simulations of the deflagration rate at high pressure under-predict the measured values. We determined the additional physics modules and revised kinetic steps needed in ALE-3D to resolve the pressure dependence of deflagration. We applied the initial continuum model (reactive flow model) to the calculation of detonation in small diameter, unconfined explosive cylinders, for which the reaction zone is a significant fraction of the diameter, so that the non-ideality is an important feature of the results.

In addition to meeting the milepost objectives, we have achieved a number of technical firsts and additional accomplishments. We developed methods by which the transport properties of non-ideal gas mixtures can be determined that are consistent with the equation of state at high temperature and density. We developed methods to produce realistic particle size distributions for assemblies of plastic bonded explosives. In so doing, we are near the limit imposed by present computer software and hardware to using the full richness of these assemblies of particles in our simulations. We have begun a computational study of the ignition process around intra-crystalline defects, both air-filled and solvent filled, and the pressure dependence of the ignition process. These features have been incorporated into the continuum model in a preliminary way. We have begun an experimental project funded under the joint DoD/DOE Munitions Technology Development Program. The experimental results from that work, in conjunction with the computational and theoretical work under direct ASCI sponsorship, have been crucial in identifying weaknesses in our numerical simulations, and those additional physics models needed to advance our work.

Along the way, the results of this project have been beneficial to other parts of the ASCI program, and to other parts of this laboratory. The ASCI materials project benefits from our work in developing methods to describe assemblies of crystals for meso-scale modeling. The ASCI ALE-3D development project benefits from our continued exploration of the non-linear coupling of

temperature, heat transfer, and chemical reaction in significantly deforming materials. A summary of the ALE-3D improvements resulting from this project is included in the Appendix. In addition to uncovering errors and programming inefficiencies, the models have become much more robust as a result of our exercising them alone and in combinations with other models.

Although the path leading from grain-scale simulations to continuum models has proven to be a more difficult one than any of us had envisioned, and involved more theoretical, experimental, computational and numerical expertise than we expected, it has been a fruitful path that has provided us with many insights into the orderly way in which a plastic bonded explosive proceeds from mechanical stimulus to explosion.

Future Directions

As should be expected in any research project we have not only experienced success, but also exposed areas of ignorance that must be addressed to make continued progress to our long term goals, both in the physics and chemistry of the processes that lead to detonation, and in the numerical methods used to simulate those processes. They are 1) the morphology of defects in PBXs, and the interaction of those defects with low-amplitude stress as well as high amplitude shock waves; 2) resolution of the differences between our simulations and the measured deflagration speed, which will involve revised chemical kinetic rates as well as species diffusion in ALE-3D; and 3) the effect of binder on the propagation of deflagrations.

Our initial calculations of ignition in the wake of a shock wave collapsing an air-filled hole are flawed by the numerical interaction of the advection of hot product into zones that contain cooler reactant, and the chemistry module that is based on the assumption that all species are at a single temperature. The resolution of these numerical issues required to advance this project will also improve the general treatment of mixed zones in ALE-3D.

We anticipate that the Munitions Technology Development Program will continue to play an important role in providing the touchstone experimental results that are crucial in evaluating a materials-based project. Our out-year plans for that work include continuing work on measuring burn velocity at high pressure, starting work on methods to visualize the defects in addition to the optical methods developed by Mark Hoffman for visualizing inter-crystalline defects, and beginning work on visualizing the deflagration that starts from a shock-loaded defect.

Resolution of the differences between our measured and calculated deflagration rates will require us to incorporate species diffusion to calibrate our kinetic rates from low pressure experiments, and appropriately inhibiting that diffusion as the pressure increases. We also anticipate that we will need to refine our kinetic rates to more closely match the insights from Manaa's QMD simulations. According to Manaa's results, the appearance of elemental carbon is delayed. That delay, in turn, changes the equilibrium mixture of species, and so changes the equation of state of the products at early times. Since that also changes the temperature of the product gases in a constant pressure flame significantly, the effect on calculated flame speed is likely to be large as well.

Acknowledgements

This project has enjoyed the enthusiastic support of Randy Simpson, the leader for explosive research in both ASCI and the Munitions Technology Development Program, during this project's formation, and throughout its life. His support has been crucial in maintaining focus on a long-term project with ambitious objectives.

Stew Keeton, Estella McGuire and Tom Reitter have all spent enormous amounts of time and energy in performing the numerical simulations for this project, using state-of-the-art computer hardware (that was sometimes being beta-tested by them) to do large scale computing, and using versions of ALE-3D that were not always up to the severe tasks we were asking of them.

The ALE-3D team, under the leadership of Richard Sharp, has supported this work by modifications, improvements, and the addition of new features throughout. The technical requirements of our simulations have used features in ALE-3D in ways that were not envisioned by the module designers, and have used combinations of features in ways that sometimes resulted in overwhelmingly non-linear response. As a result, the ALE-3D simulation program is much more robust than it would otherwise be. Of that team, we especially want to recognize several individuals. Gary Friedman created dramatic improvements in the efficiency of generating shapes, programmed the algorithms for several of the early methods of developing shapes, and served as the point of contact and interface between this project and the team. In addition, he met the challenge of our bimodal distribution with his customary initiative, enthusiasm and skill. Al Nichols was responsible for the chemistry, deflagration, and heat-transfer modules in ALE-3D and repaired them as frequently as we broke them. Bob Cooper, Andy Anderson, Brad Wallin, Rob Neely, and Scott Futral were also instrumental in making our progress possible.

Larry Fried provided early versions of CHEETAH and sorely needed advice on its effective use in developing equation of state tables. Sorin Bastea provided equations of state with CHEQ, and also developed methods to calculate consistent transport properties for the explosive products, which we have been using in our simulations. His contributions to improving the accuracy of our material properties have led us to much more accurate simulations, and a real understanding of the physics and chemistry issues we need to address in the future. Riad Manaa's Quantum Molecular Dynamics simulations of the decomposition of HMX at high temperature and pressure have given us a benchmark for testing our simplified chemical kinetics reaction scheme at high pressure and temperature.

Craig Tarver provided his unpublished parameters for reactive flow modeling of PBX-9404 in small diameter charges, and freely shared his expertise in applying reactive flow models to calculations of macro-scale phenomena.

Joe Zaug and his colleagues have provided expertise for the measurements of the propagation of deflagrations in the diamond anvil cell, which remain to date as the closest experimental touchstone to our complex computer simulations.

Mark Hoffman has extended his techniques for visualizing flaws in explosive crystals and provided the photomicrographs of HMX and RDX for us.

The developers of QHULL at the University of Minnesota geometry center have made our efforts to describe the structure of PBXs much easier. Similarly, we have used the ideas of Edward Garboczi at NIST to speed our development of the description of realistic particle-size distributions.

Project Documentation

U. S. Department of Energy Defense Programs, "Accelerated Strategic Computing Initiative (ASCI) Program Plan," DOE/DP - 99-000010592, January 2000

Lawrence Livermore National Laboratory, "Joint DoD/DOE Munitions Technology Development Program FY 2000 progress report, UCRL-ID-103482-00, May 31, 2001

Daniel L. Farber, Anthony P. Esposito, Joseph M. Zaug, John E. Reaugh, and Chantel M. Aracne, "First results of reaction propagation rates in HMX at high pressure," in Shock Waves in Condensed Matter, Atlanta, GA, July 2001

John E. Reaugh, "Delivery of initial macro-scale reactive flow model for high-explosive detonation derived from grain-scale dynamics," UCRL-ID-146734, December 27, 2001

Lawrence Livermore National Laboratory, "Joint DoD/DOE Munitions Technology Development Program FY 2001 progress report, UCRL-ID-103482-01, June 28, 2002

S. Bastea, "Transport properties of fluid mixtures at high pressures and temperatures. Application to the detonation products of HMX," 12th International Detonation Symposium, San Diego, CA, August, 2002

J. E. Reaugh and E. L. Lee, "Isochoric burn, an internally consistent method for the reactant to product transformation in reactive flow," 12th International Detonation Symposium, San Diego, CA, August 2002

References

Dale P. Bentz, Edward J. Garboczi, and Kenneth A. Snyder, "A Hard Core / Soft Shell Microstructural Model for Studying Percolation and Transport in Three-Dimensional Composite Media," National Institute of Standards and Technology report NISTR 6265, Gaithersburg, MD, January 1999.

M. R. Manaa, L. E. Fried, J. E. Reaugh, and M. Elstner, "The decomposition of HMX at extreme conditions." Proceedings of the international workshop on New Models and Hydrocodes for Shock Wave Processes in Condensed Matter, Edinburgh, Scotland, May 2002

M. R. Manaa, L. E. Fried, C. F. Melius, M. Elstner, and Th. Frauenheim, "The decomposition of HMX at extreme conditions: A Molecular Dynamics Simulation," *Journal of Physical Chemistry A*, **2002**, 106, p9024

R. R. McGuire and C. M. Tarver, "Chemical Decomposition models for the Thermal explosion of confined HMX, TATB, RDX, and TNT explosives," 7th Symposium (International) on Detonation, Annapolis MD, June, 1981

Ya. B. Zel'dovich and G. Barenblat, *Combustion and Flame*, **3**, p61 (1959)

A. W. Campbell and Ray Engelke, "The diameter effect in high-density heterogeneous explosives," 6th Symposium (International) on Detonation, White Oak, MD, August 1976

F. P. Bowden and Y. D. Yoffe, *Initiation and growth of explosion in liquids and solids*, Cambridge University Press, Cambridge, 1952

Appendix: Improvements to ALE-3D that have resulted from this project

Information in this appendix was assembled by the ALE-3D team from their documentation of the history of ALE-3D development.

Level set Issues

Added a minimum distance that a zone must be away from the front before it can ignite. This prevents the thermal conduction from the burn front from igniting zones before the level set can get to them. Generally occurs when the ignition criteria are set too low.

Improved the algorithm to insure that the level set would move with the material as the material went through pressure relaxation.

Addressed and fixed several issues with small volume fractions causing ignition.

Moved deflagration out of the thermal package so that deflagration can run in a hydro-only calculation. This was important, since the time scale of many of the calculations were such that thermal conduction was insignificant, and thus it would be a waste of computational resources to run thermal just to support the deflagration.

The remap operation was using a scaled variable to determine the convergence of the parametric co-ordinate calculation. For systems with small length scales, the scaled value would allow significant error in the final results. Changed code to test on relative error.

Improved level set calculations in constant volume situations where slope limiters were being incorrectly applied.

Level set methods for simulating mixed material elements, i.e. HE reactants, products, and inert spectator materials, were developed for another calculation, but were used extensively in HE grain-scale calculations.

Significantly improved the algorithm for advancing the level set in the presence of multiple ignition sites.

General Material Issues:

Developed a material viewer that allows the user and developer to examine the material models used in the code. Used this to uncover improperly defined species materials that resulted in poor chemical material convergence.

Added Livermore Equation of State (LEOS) thermal models directly for HE grain-scale calculations.

Added ability to read LEOS table data from public and private databases.

Added internal extrapolation off of the LEOS tables rather than use the LEOS library's extrapolation routines. This made using LEOS materials as chemical species stable.

Chemical Material/Reaction Issues

Changed initial energy guess used in the species pressure/temperature equilibration routine from one based on the reference energy to one based on the species energy from previous state. Also, added coding to advect this energy.

Added advection controls on the rate of reaction and strength of the chemical discontinuity. This was to prevent a situation where components of the chemical system would be inconsistently advected due to the monotonic advection algorithm.

Modified ALE-3D so that the compression reactions could be run without other reactions.

Found and corrected a long-standing error in the ALE-3D implementation of the reactive flow model that gave negative pressures even with extensive gas fractions.

Thermal Solver Issues

Found and corrected an error in the time step calculation, eliminating time step crashes.

Corrected warning and error messages to more easily analyze crashes caused by rapid temperature change.

Improved recovery from failed thermal iterations.

IO issues

Corrected end-cases in writing ALE-3D restart files for multi-domain calculations where some domains did not contain all the materials.

General stuff

Fixed various memory leaks.

Fixed various end-cases that caused the code to crash, e.g. trying to do chemistry with no chemical reactions defined.

Problem Generation

About 200,000 ALE-3D shape commands are used to place the HE, binder and voids into the initial mesh for HE grain-scale calculations. The ALE-3D parallel generator was not designed for such use of the shape commands and much work has been performed to speed up this section of the generation process.

Generation time has been decreased from many tens of hours to a less than an hour for a three million element mesh. Recently, a 219 million element was generated in less than four hours.

A computer program was written based on University of Minnesota's Q-Hull that created ALE-3D shape commands for HE grains and binder.

The ALE-3D generator was modified to process "external shape files" that contained the shape commands for HE grain-scale calculations. The concept of a "file bounding box" was added to decrease generation time. The concept of "pre-ordered shape commands" was added to decrease memory usage and generation time. Various changes were made to the volume of intersection ("npic") algorithm in the ALE-3D generator to improve its performance and reduce memory usage.

The generator was modified to correctly process small scale meshes.

A Perl script was written to divide up the shape commands into many smaller input files in an attempt to exploit the use of "bounding box" information to reduce generation time. A similar Perl script was written to perform the same operation and to preserve the order of shape commands.

A program was written using a parallel algorithm that sorted the shape commands into smaller input files so that the "bounding box" of each input file would be minimized.

An option was implemented in the ALE-3D generator that caused shape commands to be sorted in parallel to the processors whose domains overlapped the shape command.