

A MODEL FOR HIGH EXPLOSIVE COOKOFF *

Albert L. Nichols III, Andy Anderson, Rob Neely, and Brad Wallin

Lawrence Livermore National Laboratory
Livermore, CA 94550

We have continued to improve our ability to model the response of energetic materials to thermal stimuli and the processes involved in the energetic response. Several new algorithms have been developed to increase the accuracy and fidelity of the modeling process. These include a level set driven multi-material deflagration model, a multi-temperature mixed material treatment, self-consistent thermal-hydro coupling, full implicit quasi-static hydrodynamics, ale slide surfaces and ale slide deletion. These capabilities have allowed us to improve our ability to model the cookoff process from the initial application of heat to the final metal expansion.

INTRODUCTION

Understanding the violence of high explosive cookoff events is important for understanding explosive safety. The severity of these events can range from benign rupture of the confinement to a violent response nearing that of a detonation. By predicting the violence of the cookoff one will be able to determine the proper configurations and the level of controls required for safe handling of the explosive systems. In previous work¹ it was shown that the cookoff cannot be confidently modeled by separating the process into a thermal/chemical phase and a hydro/chemical phase. The physical change of the material during the initial heat up phase can be significant and therefore radically change the results. This has been demonstrated by recent experiments² where a slight change in the initial ullage in the explosive confinement can radically change the violence of the cookoff event.

The physical process that must be modeled in a cookoff starts with heat being applied to the container, which is then conducted to the explosive. The explosive then undergoes both physical and chemical change that must be accounted for by appropriate material motion. Finally, the explosive transitions to a thermal run away, where the energy release rate of the explosive is larger than can be removed by thermal conduction. This starts the burning or deflagration process, which transforms the explosive from a condensed phase to a gas in relatively short order. In our previous work, we demonstrated a level set scheme to model the deflagration of the high explosive, a revised thermal transport model that improves the thermal/hydro energy consistency, and a semi-implicit hydrodynamic model for mechanically moving the explosive while taking long time steps.

We have continued development of the models required to represent the cookoff process. We have added several thermal boundary conditions that more accurately describe the various physical heating mechanisms that are found in cookoff events. We have extended our improved thermal transport model to include elements consisting of multiple materials, also known as Eulerian mixed zones. This improves our energy conservation in mixed zones and also removes the artificial thermal transport between hot and cold materials that might co-exist in the same zone. The semi-implicit hydrodynamics scheme has been replaced by a fully implicit hydrodynamics scheme. By insuring that the hydrodynamic end state is fully converged, this scheme improves the stability of the overall model by not allowing the mechanical state to get too far out of equilibrium. It also allows us to take larger time steps than we were allowed to before. As the cookoff process proceeds, the size of time step required to accurately model the system must decrease. At some point, it is no longer necessary to use a fully implicit thermal transport model as we had previously used, since the heat is only able to flow a short distance during a time step. We have developed an explicit thermal treatment to take advantage of this fact. Finally, improvements have been made to the deflagration model to treat the fact that the interface between the product gases and the unreacted material can be much thinner than the size of any zone that would be feasible in a three dimensional model. To properly describe this process in the hydrodynamic framework, we have developed a model where the reacted and unreacted materials are considered distinct material regions. This allows the code to treat the deflagration front as a real discontinuity.

These new models will be fully described in the next section. Examples of their use on a spherical test object and on the Los Alamos Annular Cookoff experiment will be provided in the third section. We will conclude in the final section.

THE ALE3D CODE

The ALE3D³ code is a coupled thermal-hydro-chemical code that has been under development at LLNL for several years. The current version of ALE3D began as a 3D ALE hydrocode to which has been added several capabilities. These include implicit thermal transport, thermally driven reactions, models for both the thermal and mechanical properties of chemical mixtures, second order species advection, and implicit hydrodynamics. Many of these features were described in the previous paper¹. Here we describe changes that have been implemented to improve the cookoff models.

THERMAL BOUNDARY CONDITIONS

If we are to model the response of real systems to thermal insult, it is important that we provide a mechanism to describe the thermal loading. This is typically done by specifying a boundary condition that approximates the physical system to some




Figure 1. Example of a thermal profile on an aluminum sphere provided by a point source, line source, two line segment sources and a plane source, all controlled by a PID controller.

degree. In ALE3D, we have the standard fixed spatial coefficient convection, radiation, flux, and temperature boundary conditions. These are insufficient to model a real system, so we have added both an enclosure radiation mechanism from MONTE3D⁴ and an algorithm based on geometric optics to handle heating from a plane, line, line segment, and a point.

Previously, we described the implementation of a material heat generation option based on a Proportional-Integral-Derivative (PID) thermal controller. The PID controller option requires the nodal location of a virtual thermocouple. The difference between the actual temperature at the thermocouple location and the target temperature defines an error θ . The rate of energy delivery to the PID controlled elements is:

$$\dot{E}(t) = a\theta(t) + b\dot{\theta}(t) + c\int_0^t \theta(t')dt' \quad (1)$$

where a , b , c are the PID constants. The flux has both upper and lower limits. Thus, when the system becomes exothermic, the PID heat generation option simply stops adding energy to the heater elements. We have extended this technique so that it can apply to any thermal boundary condition.




Figure 2. Comparison between set point and actual temperature for the test object in Figure 1. Temperature is in Kelvin, time is in microseconds.

The use of these boundary conditions on a simple spherical test object is shown in Figure 1. The sphere is subjected to fluxes from a planar source, two line-segment sources, a infinite line source and a point source, as well as a uniform radiation boundary condition to 300 K. The flux sources were controlled by a PID controller. Comparison between the actual temperature and target temperature is shown in Figure 2.

THERMAL/MECHANICAL INTERACTIONS

ALE3D accomplishes thermal/mechanical coupling by a sequence of alternating mechanical and thermal steps. The mechanical steps move the nodes while holding the entropy, S , constant. The thermal step moves heat between nodes holding the nodal locations fixed. The mechanical energy is modified by the change induced by thermal transport. The ϕ treatment described previously to communicate the hydrodynamic contribution to the temperature change has been expanded to include second order contributions:

$$\begin{aligned}\phi &= \frac{\Delta T}{T} \\ &= -\gamma \frac{\Delta V}{V} + \left[\gamma^2 - \frac{\partial^2 P}{\partial V \partial e} + \frac{\partial^2 P}{\partial e^2} \right] \left(\frac{\Delta V}{V} \right)^2 \\ &\quad + \Delta \underline{\underline{\varepsilon}} : \left(\frac{\partial \underline{\underline{\xi}}}{\partial e} \right)_{\varepsilon, V}\end{aligned}\quad (2)$$

where γ is the Grünesen gamma function, $\underline{\underline{\xi}}$ is the second order derivative contribution, $\underline{\underline{\varepsilon}}$ is the deviatoric stress and $\Delta \underline{\underline{\varepsilon}}$ is the deviatoric strain increment. Inclusion of the second order terms has significantly improved the temperature consistency between the thermal and hydrodynamic packages, which maintain nodal and zonal temperatures, respectively.

IMPLICIT-EXPLICIT TREATMENTS

The terms implicit and explicit are often used to describe a method of solution for the hydrodynamic or thermal transport equations. Unfortunately, these terms are often overused, leading to potential confusion. In this work, we describe two types of

implicit vs. explicit calculations: spatial and temporal. The choice of which type of scheme should be used is generally dependent on the size of the time step, with smaller time steps favoring explicit methods and larger steps favoring implicit methods. Part of the reason for the confusion between these two types of schemes is that most codes' implementations are either both temporal and spatial implicit or both explicit. However, this connection is often not necessary and can lead to unnecessary computational overhead.

The first type we consider is spatial implicit versus explicit. The spatial explicit solution of a problem makes the assumption that a perturbation in one element of the problem can only effect those elements that are directly connected to it. The implicit solution assumes that such a perturbation can affect all elements in the system within a single time step. Because every element (or node) is connected to every other element (or node), one generally uses a matrix to solve these problems.

The second type we consider is temporal implicit versus explicit. The temporal explicit solution of a problem makes the assumption that the properties of the system do not change rapidly during a time step. The implicit solution assumes that the properties of the system can vary significantly during a time step. Because the properties of the system can change during the time step, it is necessary to evaluate the properties at some time during the time step. If taken at the half step, these systems are second order but conditionally stable, while if taken at the end of the time step they are first order accurate but unconditionally stable. Another way this process is described is the difference between a linear and non-linear solution scheme.

Systems where material properties do not change can be solved temporally explicitly, even if the time step is large and an implicit spatial treatment is required. On the other hand, complex processes, exothermic chemical reactions for example, may require a full temporal implicit treatment even though the time step is sufficiently small and the communication between elements is so slow that an explicit spatial treatment is sufficient.

Systems with energetic materials should always be modeled with an implicit temporal solver. Their

capability of rapid energy release and phase change would drive explicit solvers into time scales that are completely impractical.

From the previous discussion, it is clear that depending on the scale of the time step, implicit and explicit spatial solution schemes need to be available to solve both the mechanical motion and the thermal transport. Previously, the ALE3D code supported both types of schemes for the hydrodynamic problem, but only an implicit scheme for the thermal problem. An explicit solution scheme has been added for the thermal problem by making a lumped mass matrix approximation. This diagonalizes the matrix, allowing us to solve the thermal problem with a trivial inverse solution.

MODELING LONG TIME SCALES

Previously we described an implicit spatial hydrodynamics scheme that would allow us to model the long time scales associated with the cook-off response. It was shown that this method captured significant physical changes in the system that could not be modeled with the more traditional method that modeled the thermal response of energetic materials with a thermal/chemical code until the chemical reaction went into thermal run-away, and then transitioned to some form of a deflagration model.

Our previous implicit hydrodynamic solution scheme can be described as implicit in space but explicit in time. We have improved the implicit quasi-static hydrodynamics method that we described previously by adding an outer iteration loop, transforming it to an implicit in time and space solution scheme. As stated before, the implicit hydrodynamics method solves the mechanics problem by solving for the change in the location of the nodes over a time step. In our previous scheme, these changes were dependent on both the forces and their derivatives at the beginning of the time step. In the new scheme, we use the values at the end of the time step. Since these values are not known at the beginning of the process, it is necessary to iterate until a consistent set of values is obtained. We can solve these with either a direct or iterative matrix solver routine. The matrices produced by the implicit hydrodynamics can be ill conditioned and difficult to solve for certain classes of problems, such as when there is a slide surface. A method has been developed that significantly improves the

condition number for slide surface problems by incorporating a Schur reduction that reduces the complexity of the set of linear equations.

THERMAL/MECHANICAL CONSISTENCY

In the original formulation, a single ALE3D time step would first calculate a hydrodynamic step, followed by a thermal transport step, and finish off with an advection step. It was found that this process could lead to serious stability problems, especially when the thermal process was either changing the material properties of the explosive material or generating a substantial quantity of energy. In essence, it was possible for the thermal-chemical process to generate sufficient change that it would be impossible for the implicit hydrodynamic formulation to find a solution. A similar problem exists where it is possible for the hydrodynamic system to provide sufficient change so that the thermal solver could not converge. This problem had been handled by allowing the thermal process to subcycle the time step until it reached convergence. Such a process would not work with the hydro step because the thermal energy had to be added to the problem at the end of the previous step, therefore, no amount of subcycling would achieve convergence. Thus, we found that it became necessary to control the time step so that it would be impossible for the thermal-chemical process to over-drive the hydrodynamic solver.

The process of successive hydro-thermal steps also leads to a degradation of either second order accuracy or thermal consistency. Simply stated, a choice must be made as to how the thermal energy is deposited in the hydrodynamic process. If it is added to the hydrodynamic energy at the end of the time step then there will be a pressure discontinuity between the end of one time step and the beginning of the next. This results in a loss of second order accuracy in the energy calculation. On the other hand, if the energy is added throughout the subsequent time step, then there will not be a pressure discontinuity, but there will be a thermodynamic inconsistency between the thermal and hydrodynamic phases, which will also lead to loss of second order accuracy.

To fix these problems, we have developed an improved thermal-hydro coupling scheme. In this process we calculate a single hydrodynamic non-linear iteration followed by a single thermal non-

linear iteration. The hydrodynamic load from the current hydrodynamic step is applied to the thermal calculation, and the thermal load from the previous thermal step is applied to the hydrodynamic step. After both steps, the convergence criteria for both iterative schemes are checked, and if either has not been satisfied, the process is repeated. During the subsequent hydro steps, the thermal load that has been calculated for the current step is used, in a time centered manner, thereby recovering second order accuracy and thermodynamic consistency. Also, since it is still possible for either of the packages to generate changes that the other cannot handle, we allow for the fully coupled system to automatically reduce the time step, enabling the code to maintain consistency and stability without continuous user intervention. In essence, we have brought the hydrodynamic calculation into the thermal non-linear iteration.

The net effects of this change are that it provides an increased accuracy for the results, and it allows us to run problems with larger time steps since the code is now able to adapt to changing physical processes without human intervention.

IMPLICIT-EXPLICIT TRANSITION

With both implicit and explicit solution schemes available for both the thermal and hydrodynamic problems, it is necessary to define criteria for the transition between them. The criteria are different for the thermal and hydrodynamic systems because of the approximations used. As described previously, the transition for the hydrodynamic solver is based on the computational time difference between the implicit or explicit solve, so that we transition when it is more efficient to calculate with an explicit scheme instead of an implicit one. This transition is always in this one direction. The implicit hydrodynamics scheme is a quasi-static solution to the hydrodynamic problem. When we transition, we are actually changing the problem that is being solved. For the thermal process, the implicit and explicit schemes are actually solving the same problem, just in different time regimes. Thus, the transition between the two solution schemes is allowed to be much smoother and also bi-directional. The implicit solution scheme is maintained until the time step drops to the point where the explicit calculation would be as accurate, at which point the solution scheme is changed.

BURN PROPAGATION

Once the high explosive system has ignited, we use a front propagating capability based on level set models⁵. To propagate the burn front with these models, one creates a field ψ that is zero at the surface of ignition, and monotonically increasing away from it. One then solves:

$$\frac{\partial \psi}{\partial t} = -c(P, T, \{N_i\}) |\nabla \psi| \quad (3)$$

which is a reformulation of the wave equation. Here c is the experimentally determined burn speed of the high explosive under the appropriate conditions.

Typically, level set techniques are used on Eulerian systems. That we need to use it on a Lagrange mesh adds certain complications. The first complication is that the mesh will move, carrying along the level set field with it. This can cause severe distortion in the field such that the magnitude of the gradient is no longer approximately one. To correct for this, we have added a renormalization step to our level set process that adjusts the level set values away from the level zero surface to achieve the unit magnitude of the gradient. This has a side benefit of allowing us to only initialize the level set in a small portion of the mesh around the ignition point. The renormalization process will then propagate the level set field to the rest of the explosive.

Whereas the first complication occurs mainly on the unreacted side of the level zero surface, the second complication occurs on the product side. On the product side, the gases move away from the zero level set surface, flattening out the level set as they move. As the deflagration process occurs within a zone, this tends to drag the level zero surface with the gas. To prevent this, we calculate the average solid velocity and adjust the level set based on the gas flow relative to the solid velocity. This prevents the level zero surface from moving where it does not belong.

A third complication occurs due to species advection. In our original formulation, the deflagration process would change the composition of the explosive material from reactant species to product species. However, during the advection phase of the calculation, the second-order monotonic

advection would cause reactant species to be advected through the level zero surface into the product gases. The reason for this is simply that the advection scheme assumed that all species concentrations are continuous. However, this is not the case through a deflagration. This led us to the scheme that is described in the next section.

MULTI-MATERIAL DEFLAGRATION MODEL

In the original formulation, a single high explosive material would be assigned to the explosive region of the problem. As the explosive burned, the composition of the material would change, but it would still be considered a single material with a single temperature and pressure within any zone. As noted previously, this would lead to concentration errors during advection. To correct this problem, we implemented a scheme where the level set divides the explosive into two material regions, burned and unburned. This separation allowed us to define different physical states for the two regions, even within the same mesh element. Thus, the reactant region would have a cold temperature while the product region would be hot. By having distinct temperatures, one could also allow the reactant and product regions to continue to have thermally based reactions even while undergoing the burn process. Previously, it was necessary to turn off the reaction in the burning zones because the temperature was not appropriate to either phase. Another benefit of dividing the explosive into two regions was that the advection scheme could now treat them as such, thereby not advecting material from one region into the other.

This method also brought along certain complications as well. Since the pressures in the two regions within an element need not be equal, the pressure equilibration phase could change the volume fractions of the explosive regions. In order to keep the zero level set and the material boundaries synchronized, it was necessary to realign the level set in the mixed elements. This was done by computing the volume fraction defined by the level set and adding an element based correction to the level set until the level set volume fraction and material volume fraction agreed. This adjustment was then averaged at the nodes to obtain the final adjustment.

THERMAL TRANSPORT THROUGH MIXED ELEMENTS

Along a vein similar to the pressure, we needed to extend our thermal diffusion treatment to handle the thermal diffusion within a mesh element that has multiple temperatures. Previously, we found that the addition of a zone centered excess temperature to the standard eight-gauss-point finite element thermal treatment improved the agreement between the temperature calculated in the thermal solver and that calculated from the energy of the element. This excess temperature is associated with a basis function that has compact support within the element, but is zero outside. Energy from hydrodynamic motion or chemical reactions would be added to the excess temperature and would then be conducted to the nodes.

In a mixed zone, it became apparent that if there were only one excess temperature, then the temperatures of the two materials in that zone would have to be the same. In order to fix this problem, we added an excess temperature per material in the element. These excess temperatures would conduct to the nodes, not to each other. This prevents the instantaneous heating of the cold solid by the hot gases.

ALE-SLIDE SURFACES AND SLIDE DELETION

One of the most important new features in ALE3D is ALE slide surfaces. A slide surface is a contact surface between two potentially discontinuous sections of the finite element mesh. An ale slide surface connects two discontinuous regions that are required to have nodal locations that are co-incident across the slide surface. During the lagrange phase of the time step, the two regions are able to move tangentially with respect to each other – they slide. During the advection phase of the step the nodes on the slave surface are returned to the location of their corresponding master node. The special relationship between the two slide surfaces allows the code to correct for volume flux errors that normally occur during advection of non-flat slide surfaces.

Since the nodes of the elements across the slide surface are guaranteed to be co-located at the end of the time step, it is also possible to remove the slide surface and thereby allow advection in the direction normal to the former slide surface. This is especially

useful when the slide surface has begun to fold, or when the elements on one side of the slide surface have become thin due to expansion on the other side. An example of where these features are required is the cookoff of a cased explosive. During the initial heat up phase, it is important to have a slide surface between the explosive and the case, because the thermal expansion of the explosive will cause lateral motion between them. Later, when the explosive has started to transition to a gas, the case elements will start to compress due to the expansion of the gas volume. At this point, removing the slide surface will allow the explosive products to expand into the elements that were formerly occupied by the case material.

EXAMPLE PROBLEMS

We show two models of the energetic material response using ALE3D.

SAMPLE PROBLEM 1: DEFLAGRATION IN A SPHERICAL SHELL

For our first example, we return to spherical pressure vessel geometry we used previously. The initial configuration is shown in Figure 3. The system is a pressure vessel consisting of a 1 cm spherical shell of Aluminum surrounding a 2 cm spherical shell of HE surrounding a 1 cm spherical shell of Aluminum. The HE is treated as a four-component chemical mixture material. The material properties for the pristine HE and its first decomposition product in this system have the properties of unreacted LX-04, an HMX-based explosive. The second and third decomposition products are treated with an LEOS⁶ equation of state based on CHEETAH^{7,8,9} calculations for the early product gases of HMX and the fully equilibrated product gases. The initial composition is the fully unreacted material. The energetics and reaction rates are those derived from the work of Tarver et. al.¹⁰ The deflagration model is ignited at the point when the temperature of either an element or a node has exceeded 1000 K. The deflagration progresses at a rate that is ten times that derived from experiments of Maienschein et al¹¹ on pristine LX-04. This increased rate is used to take into account the increased surface area of heated explosive. The outside of the outer shell has both thermal radiative and convective boundary conditions, and is heated with a uni-axial flux along the negative z-axis. The




Figure 3. Initial configuration of the spherical shells deflagration example.

heat flux is controlled by a PID-controller located at the bottom of the outer shell so that the temperature increases at a rate of ten degrees C per minute. A slide surface has been inserted between the explosive and the outer aluminum shell.

The calculation begins with implicit thermal and hydrodynamic calculations. When the time step dropped to the millisecond range the thermal solver transitioned to the explicit formulation. The code transitioned to explicit hydrodynamics when the time step dropped into the 10-microsecond range. The explosive ignited after approximately 1390 seconds. In Figure 4 we show the log of the temperature, the pressure, the velocity field, and material boundaries between the explosive reactant phase, product phase, and internal aluminum shell. There are several interesting features in this figure. As should be expected, the temperature of the product gases is significantly higher than the reactant phase, while the pressure is essential continuous. Although the explosive ignited at the very bottom of the system, it is interesting to note that it is spreading much faster along the external shell boundary than it is propagating toward the center. When examining the velocity field, the reason for this becomes obvious. The pressure from the product gases is causing the external aluminum shell to cantilever away from the center of the burn, creating a place for the explosive gases to migrate. This migration spreads the burn in a direction tangential to the nominal burn direction. This feature is significantly damped if there are no slide surface between the explosive and the outer shell.




Figure 4. A slice through the spherical shells deflagration example shortly after initiation of deflagration. The log of the temperature between 350 K and 2000K is displayed on the left and the pressure between +/- .3 Gpa is displayed on the right. The internal circle is boundary between the inner aluminum shell and the unreacted explosive. The other line is the boundary between the reactant and product phases. The relative velocity field is also displayed.

SAMPLE PROBLEM 2: A MODEL FOR THE LOS ALAMOS ANNULAR CONFINEMENT TEST

Los Alamos recently conducted an experiment where they carefully heated a right circular annulus of PBX9501 in a brass confinement fixture within a vacuum chamber¹². Heating tape was wrapped around the external surface, and the interior was instrumented with several thermocouples. After approximately 10 hours the explosive ignited, resulting in a relatively uniform collapse of the internal cylindrical confinement wall and the formation of a jet that eventually destroyed their X-ray diagnostics. The initial configuration is shown in Figure 5.

We have been able to model this system by making several assumptions. Burn rate measurements on PBX9501 have shown deconsolidative burning, resulting in apparent burn rates several orders of magnitude larger than that found for LX04. The lower binder concentration of

PBX9501 causes it to have a higher porosity than LX04. This porosity allows hot gases to permeate through the bed of unreacted explosive. In our model, we use the level set model to propagate the location of the burn front through the explosive. However, unlike the previous example, we will only ignite 0.1% of the explosive as the burn front passes through. We found that it was necessary for the burn front to progress at a velocity of 1mm/ μ s to maintain the required symmetry.

The rest of the material is burned with the following geometric rate law:

$$\frac{dx_g}{dt} = Ax_s^{.667} (1 - x_s)^{.323} x_g^{.01} R(P) \quad (4)$$

where x_s is the solid mass fraction, x_g is the gas mass fraction, $R(P)$ is the burn rate as a function of pressure as determined from the stand burner experiments. The $Ax_s^{.667}$ term is set to model the burning of 100-200 micron HMX particles found in PBX9501. The $(1-x)^{.323}$ term is to model the process of the initial flame expanding to cover those particles, and the $x^{.01}$ term is to prevent the explosive from burning before it has been ignited.

We use the same reactant and product equation of state as used in the previous example. To simplify the modeling process, we did not model the entire heat up of the system, but instead ignited a small region 3 mm in radius at the experimentally




Figure 5. The initial material configuration of the Los Alamos Annular Cookoff Test. The pink, green and light blue regions are the brass, and the dark blue region is the unreacted explosive. The system is symmetric about the left side.




Figure 6. Model of the Los Alamos Annular Cookoff Test. We show five stages in the collapse of the inner cylindrical liner. The top figure in each set is the material layout cut through the axis showing the position of the interior cylinder and explosive products. The lower figure shows the pressure on the left side and the material configuration on the right. Note that although it takes $\sim 100 \mu\text{s}$ to start, the collapse process takes roughly $50 \mu\text{s}$, and results in an axial jet in essential agreement with experiment.

indicated ignition site. During early attempts to model this experiment, it was found that the presence of significant strength in the explosive would break the symmetry of the core collapse, and so strength was removed from the explosive material model.

The results of our calculation are shown in Figure 6. Immediately after ignition, the system slowly starts to pressurize as the burn front propagates through out the explosive media. After $\sim 100 \mu\text{s}$ the inner cylindrical liner begins to collapse, completing that process in $\sim 50 \mu\text{s}$, in

agreement with the experimental results. It should be noted that this model does not represent a detonation of the high explosive. In fact, the highest extent of reaction is only about $2/3$ when the jet is formed. This result implies that it is not necessary for the explosive to detonate in order to have both a symmetrical collapse and jet formation.

CONCLUSIONS

This paper has described the some of the improvements that have been added to the ALE3D code to refine our ability to model the cookoff process. As we have refined our capability, these models have elucidated the need for further improvements. Such is the history of the multiple-region deflagration model, for example. These improvements have allowed us to examine the cookoff process and its features with greater fidelity and thus increase our understanding of the underlying physics.

We have demonstrated these capabilities with two classes of models. The first attempts to model the entire process from the beginning to the end. This is done by automatically changing the way the physics of the problem is modeled as the calculation progresses, making the appropriate approximations at each stage. This procedure allows us to reduce the number of assumptions as to the intermediate configurations. As has been noted previously when we modeled the VCCT experiment¹, there can be significant material motion before the explosive begins thermal runaway. We have seen the same effect even in the closed packed system used for the first example.

The second way we have used these capabilities is to scope the problem. This was shown with the second example, as it would have been impractical to model the initial heat up of the system as we were developing the response model for the final deflagration and explosive consumption. That example showed that it is possible for the explosive system to have a very violent response without actually detonating.

It is clear that there must be more work done on the material models associated with the chemically reacting mixture material. The strength of the explosive in the partially decomposed state can have a profound effect on the ensuing reaction. The higher the strength of the explosive, the more it will

resist the expansion of the decomposition products. This will increase the rate at which the explosive burns, turning what could have been a benign event into a catastrophic one. On a different direction, the current material model assumes that all of the species are of one uniform temperature and pressure. This is valid during the initial phase of the cookoff, when time scales are long compared to the time it takes to exchange energy between the phases, but is less so as the time scale gets small.

ACKNOWLEDGMENTS

The authors would like to acknowledge the work of the entire ALE3D team in bringing this work to fruition.

*This work performed under the auspices of the U.S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

REFERENCES

1. Nichols, A. L. III, Couch, R., McCallen, R. C., Otero, I., and Sharp, R., "Modeling thermally driven energetic response of high explosives", Eleventh International Detonation Symposium, Office of Naval Research ONR 33300-5, Snowmass, CO, 1998, p. 862.
2. Wardell, J. and Maienschein, J., "The scaled thermal explosion experiment", Submitted to Twelfth International Detonation Symposium.
3. Dube, E., Neely, R., Nichols, A., Sharp, R., Couch, R., & The ALE3D Team, "Users Manual for ALE3D An Arbitrary Lagrange/Eulerian 3D Code System, Version 3.2.0", Internal Publication, Lawrence Livermore National Laboratory, Livermore, CA. (2002).
4. Maltby, J. D., and Burns, P. J., "MONT2D and MONT3D user's manual," Internal Publication, Department of Mechanical Engineering, Colorado State University, 1988.
5. Sethian, J.A., "Level Set Methods and Fast Marching Methods Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Material Science", Cambridge University Press, Cambridge, United Kingdom, 1999.
6. Hill, E., "LEOS Users Manual", Internal Publication, Lawrence Livermore National Laboratory, Livermore, CA, (2002)
7. Fried, L., and Howard, W. M., "An accurate equation of state for the exponential-6 fluid applied to dense supercritical nitrogen", J. Chem. Phys., 109, 7338 (1998)
8. Fried, L., and Howard, W. M., "The equation of state of supercritical HF, HCl, and reactive supercritical mixtures containing the elements H, C, F, Cl", J. Chem. Phys., 110, 12023 (1999)
9. Fried, L., and Howard, W. M., "Explicit Gibbs free energy equation of state applied to the carbon phase diagram", Phys. Rev. B, 61, 8734 (2000)
10. Tarver, C. M., Chidester, S. K., and Nichols, A. L. III, J. Phys. Chem. 100, 5794 (1996).
11. Maienschein, J., and Chandler, J., "Burn rates of pristine and degraded explosives at elevated pressures and temperatures", Eleventh International Detonation Symposium, Office of Naval Research ONR 33300-5, Snowmass, CO, 1998, p. 872.
12. Howe, P., Dickson, P., Henson, B., Asay, B., Smilowitz, L., and Greenfield, M., "Thermal response of explosives: slow heat", LA-UR-002859, Los Alamos National Laboratory, Los Alamos, NM (2000)