



Diffuse interface method for solid composite propellant ignition and regression

Maycon Meier^a, Emma Schmidt^a, Patrick Martinez^{a,b}, J. Matt Quinlan^a,
Brandon Runnels^{c,a,*}

^a Department of Mechanical and Aerospace Engineering, University of Colorado, Colorado Springs, CO, USA

^b Division of Engineering and Applied Science, California Institute of Technology, Pasadena, CA, USA

^c Department of Aerospace Engineering, Iowa State University, Ames, IA, USA

ARTICLE INFO

Article history:

Received 23 June 2023

Revised 4 October 2023

Accepted 9 October 2023

Keywords:

Solid rocket motors

Combustion

Diffuse interface methods

Computational physics

ABSTRACT

Solid Composite Propellants (SCPs) are extensively used in the field of propulsion for their chemical and mechanical stability in long-term storage, and for having simple production and operation processes. Computational modeling enables cost reduction, increased efficiency, and greater coverage of the configuration space in the SCP design process. However, accurate and efficient SCP modeling presents a number of numerical challenges. A primary obstacle in modeling these systems is capturing the complex evolving interface. Recently, it has been shown that the phase-field method has a strong ability to model the combustion behavior of SCPs, implicitly capturing the topological evolution at a relatively low computational cost. Initial phase-field methods show promise in their ability to do predictive regression modeling but require a number of approximations and heuristic modeling methods. This work presents a new formulation for the phase field regression model that combines a thermal solver with an Arrhenius rate law to model interface regression using a comprehensive and physics-based approach. This improves the capability of previous methods by increasing the number of kinematic forces that are accounted for, and allowing the study of thermal diffusivity in the system. It also enables the integration of a fully coupled solid-fluid interface. To demonstrate the efficacy of the model, it is applied and calibrated to a homogeneous monopropellant (ammonium perchlorate). The model is validated for a range of temperatures, with a reasonable quantitative match to experimental data. It was also demonstrated that the model recovered the relationship between ignition time and heat flux, with no fitting required. Overall, the method showed great capability of reproducing experimental data by matching temperature, burn rate, and thermal diffusivity profiles.

Statement of Significance: A new diffuse interface method is developed for calculating ignition and regression rates in SCPs. While this builds on previous work in diffuse interface modeling of SCPs, it is novel in its coupling to thermal transport, as well as its development of (and coupling to) an approximation for gas phase heat flux. The model is shown to reproduce experimental regression rate data with reasonable accuracy. Moreover, it quantitatively captures ignition behavior with no additional modifications to the model. The model is implemented in a high performance computational framework that is able to resolve large, three-dimensional mesostructures.

© 2023 The Combustion Institute. Published by Elsevier Inc. All rights reserved.

1. Introduction

Solid Composite Propellants (SCPs) have been used as rocket and missile propellants for over a century. The earliest SCPs were composed of a mixture of charcoal, potassium nitrate, and sul-

fur [1,2]. Over time, they have been further developed to enhance their propulsion properties. Despite increased efficiency and controllability of liquid propellants in the early 20th century, SCPs are still used because of their simplicity, reliability, longevity, and cost [3]. Modern SCPs mainly consist of a mixture of an oxidizer and a fuel, with additional chemical additives to improve stability, cross-linking, and thrust [3]. A common SCP is the combination of ammonium perchlorate (AP) for an oxidizer and hydroxyl-terminated polybutadiene (HTPB), which acts both as a binder and a fuel. Modern SCPs often have additive components such as aluminum

* Corresponding author at: Department of Mechanical and Aerospace Engineering, University of Colorado, Colorado Springs, CO, USA.
E-mail address: brunnels@iastate.edu (B. Runnels).




Fig. 1. Conceptual flame structure of a burning AP/HTPB SCP. The AP (orange) particles undergo phase change and closely coupled self-deflagration reactions (light pink), which produce oxygen as a product. Heat transfer to the HTPB (navy) causes polymer degradation and decomposition into gaseous, short-chain hydrocarbons. These two regions of gas diffuse together creating a triple-point flame (lean and rich premixed flames shown in light red and light blue and primary diffusion flame shown in red). The diffusion flame typically encapsulates the AP regions due to the overall rich mixtures. Owing to the high temperature of the diffusion flame and the low liftoff height near the triple-flame, the heat transfer to the solid is locally high, causing high local regression rates. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

particles to increase performance. Propellants such as AP/HTPB have a heterogeneous mesostructure, leading to a complex and time-varying burn behavior at the mesoscale (Fig. 1). Specifically, at the interfaces between the AP and HTPB a lifted diffusion flame sits close to the solid surface resulting in high heat transfer back to the solid. This high heat transfer locally promotes faster surface regression rates. The coupling between SCP composition and burn rate behavior is difficult to quantify, and has historically been done in an Edisonian approach, through costly and hazardous full-scale tests [4]. Computational simulations, on the other hand, are essential to efficiently and effectively design propellants with reduced environmental impact and lower cost.

Since the flame structure and chemical mechanisms are complex and the nonlinear interface topology must be accounted for, it is a challenge to computationally simulate the regression of SCPs. SCPs have a heterogeneous structure at the mesoscale, formed mostly by agglomerations of AP particles in a HTPB matrix. Since these materials decompose differently, the resulting flame structure depends on the concentration of the fuel and oxidizer, which in turn depends on the location of the materials' interface. This interface, where vigorous reaction and high heat transfer occurs, is geometrically complex and time-varying. Furthermore, the preferential regression near the interface, due to the heat transfer, tends to form an undulating surface geometry. This can lead to another difficulty in regression simulations, the solid surface topology can even change.

While AP is a solid crystalline compound that is stable at room temperature, at elevated temperatures it can act as an oxidizer or even react as a monopropellant. The specific decomposition mechanism of AP is complex due to its pressure and temperature dependence. At low pressure, the crystal structure of AP undergoes phase transformation before decomposing, but at pressures higher than 2 MPa the material directly sublimates. In this case, the AP can support self-deflagration. The flame temperature and surface temperature have been found to be approximately 1400 K and 800 K respectively [5,6]. HTPB, on the other hand, is a polymer compound formed by very long carbon chains. It is commonly used in SCPs, in part due to its mechanical properties and that it has a curing point lower than the decomposition temperature of AP [7]. The HTPB is

gasified in a pyrolysis process where the long carbon chains break into smaller hydrocarbons. Since this decomposition process is significantly slower than the burning of pure AP or the combustion of AP/HTPB at their interface, chunks of binder can be ejected from the SCP during the combustion [8,9]. Additionally, AP and HTPB are dense and have very low thermal conductivity, acting as insulators during the burning of SCPs. As a result, it has been observed that heat propagation during combustion only affects the temperature a short distance ahead of the surface while a large portion of the materials' temperature is unchanged [10].

Although the burn rate of SCP mixtures of AP/HTPB have been experimentally measured in the past using optical and break-wire methods [11,12], these methods can be costly and hazardous, and do not provide surface topology information. Therefore to understand the interfacial burning of a SCP more fundamentally, a sandwich structure is often used in which layers of fuel and oxidizer are burned.¹ Specifically, the test samples have a layer of HTPB sandwiched between layers of AP on either side. Compared with complex mixtures in a particle/matrix structure, these experiments are easier to perform and the results are more repeatable. In contrast to the flame structure for an AP/HTPB SCP that is three-dimensional and inherently unsteady, the sandwich propellant experiments' flame structure is nearly two-dimensional—which simplifies observations of the flame—and is nearly steady. As a result, the detailed surface geometry in relation to the interface can be quantified. In previous studies, the description of this combustion process has been divided into low-pressure premixed flame, where gasified binder and oxidizer are mixed at the molecular level, and high-pressure combustion where the flame behaviors split into oxidizing products + ammonia and into three different flames: (i) premixed binder rich flame, occurring at low pressure (ii) Oxidizing Products + Ammonia, and (iii) Binder products + AP oxidizing products, where (ii) and (iii) occur at high pressure [13,14].

One reason that simulations have been used to predict the surface regression of burning SCPs is to reduce the environmental impact of the combustion and improve their design at a relatively low cost. In addition, effectively computing this regression is a major step into understanding and simulating the combustion process of SCPs along with its acoustic waves and fluctuations [15–17]. Modeling this regression requires two major steps: (i) evaluate the regression rates of the propellants and (ii) evolve the surface morphology. Traditionally, the regression rates have been evaluated using a pressure power-law in the form of $r = \alpha(P/P_0)^n$ as proposed by Summerfield [18]. Although this method has performed well in previous studies, it is obtained by fitting the experimental data and fails to capture the physical unsteady flame behavior and its influence on the regression. Furthermore, predictions cannot be made on SCPs where the power-law behavior has not been experimentally measured. Another issue is that tracking the surface morphology can be particularly challenging for geometries such as packed AP particles in a HTPB matrix, where the interface between the two materials is complex. While some compositions permit a highly simplified representation of the surface [5], this is not sufficient in general. Previous approaches have fallen primarily into two categories. The first is to use a coordinate mapping to project the dynamic corrugated propellant surface from physical space onto a flat plane in computational space, subsequently employing one-sided differences to address the flux conditions occurring at the surface [4,19–21]. This method works well for surfaces that admit a smooth mapping to a flat surface, but is prone to er-

¹ It should be noted that the inherent two-dimensionality of the sandwich structure makes it an inadequate surrogate for three-dimensional behavior, and therefore does not truly qualify as a validation dataset in the general case. Rather, its purpose is to serve as an academic benchmark that is useful when developing a more generalized three-dimensional model in the condensed phase.

ror for interfaces with high curvature or that undergo topological transformation. Alternatively, the level-set method has been introduced to approach multi-valued cases [22–24], and is currently the most well-developed approach that uses the Eulerian view of the surface regression. This method approximates an isocountour mapping based on local regression rates, thus it effectively also uses a sharp interface model. The most severe limitations of the level-set methods are that it is difficult to capture topological interface changes.

In order to avoid these issues, the phase-field method for simulating regression was recently introduced by Kanagarajan et al. [8], in which a diffuse interface model is used to track the solid/gas interface. In the phase-field method, an order parameter field is used (usually called $\eta(\mathbf{x})$) to represent the domain occupied by the unburned material as the support of $\eta = 1$. If the range of η is $[0,1]$, then the support of $\eta = 0$ represents the domain occupied by the burned/gas phase materials. The boundary is consequently represented implicitly as the region in which $0 < \eta < 1$, with a (usually controllable) diffuse width of ϵ . The gradient of η is then only non-zero within the interfacial region, and is in the direction of the inward normal. The advantage of this approach, similar to the level-set method, is that the implicit boundary efficiently captures complex morphology and topological transitions. Unlike the level-set method's field function, however, the phase-field method's order parameter can be interpreted physically as a pointwise reaction coordinate that evolves according to a thermodynamic potential. Such a potential, generally called a free energy functional, may include contributions from multiple physical systems, such as chemical energy, stored mechanical energy, and thermal energy. The diffuse boundary formulation is also readily amenable to solving PDEs driven by complex boundaries, including solid mechanics [25], fluid mechanics [26], and heat transport (discussed below). Since the diffuse interface must be adequately resolved, the diffuse boundary methods do incur additional computational cost. This cost may be effectively limited, though, through the use of adaptive mesh refinement. Limitations for diffuse interface methods are normally associated with length scale parameters effects over the solution and computational complexity and cost. As it has been demonstrated in reference [8], the phase-field equations can be formulated to assure convergence of the simulation and that advanced meshing techniques can be used with the model to maintain computational efficiency.

This work presents a phase-field model for the regression of the burn interface during ignition and deflagration of an SCP with arbitrary morphology. This work builds on the previously introduced phase-field model in [8] is used, but the heuristic pressure power-law kinetics are replaced with a “full-feedback” coupling that accounts for the effect of mass and energy exchange across the boundaries between the fluid and solids phases. This reduces the number of fitting parameters, increases the physicality and predictiveness of the model. It also enables the ignition process to be simulated, which was previously not possible, and provides the complete thermal response throughout the domain. Though not explored here, it also makes it possible to couple to a reacting fluid solver so that the combustion effects on the regression can be directly calculated. While the implementation of a fluid solver is beyond the scope of this work, a fluid stand-in has been implemented as a placeholder to both simulate the effects of that solver and also to assure convergence of the solution.

The remainder of this paper is structured as follows: in Section 2 the phase-field model is constructed based on the chemical potential behavior of the solid propellant combustion model, and the treatment of multiple species is discussed. Furthermore, the derivation of the diffuse heat equation is shown and the construction of a stand-in fluid phase is shown. In Section 3 the adaptive mesh refinement and computational solution framework is

presented. In Section 4 a qualitative analysis of the model performance is discussed based on four application cases: (i) homogeneous substances regression, (ii) oxidizer and fuel binder heterogeneous sandwich, (iii) packed oxidizer spheres in a fuel binder matrix, and (iv) Miller packing of oxidizer in a fuel binder matrix. Finally, Section 5 synthesizes the contributions of this work.

2. Diffuse interface thermally driven regression model

In this section, a full-feedback phase model is developed and applied to the problem of burn interface regression. The model builds on the work originally proposed by Kanagarajan et al. [8], which is a general reference for this section. The theoretical underpinnings of the phase-field theory are first introduced, and then the diffuse boundary method for thermal evolution is presented, along with a surrogate model for estimating the heat flux at the burn interface.

2.1. Free energy functional and regression rate kinetics

Following prior work, a continuum region $\Omega \subset \mathbb{R}^n$, $n = 2, 3$, a simply connected measurable subset, is established containing SCP in both the solid and fluid phases. The order parameter, $\eta \in H^1(\Omega, [0, 1])$, is a field implicitly tracking the physical state of the material over Ω as a function of space and time. In the present work, material is considered to be either solid (when $\eta = 1$) or gaseous (when $\eta = 0$). No intermediate liquid melt phase is considered, though the present model may be extended to include it without adjustment to the method. The boundary between the solid phase and the gaseous phase is consequently tracked implicitly, and the representative sharp interface is taken as the set $A = \{x \in \Omega : \eta(x) = 0.5\}$. The diffuse boundary is defined as the open set $A^\epsilon = \{x \in \Omega, 0 < \eta(x) < 1\}$. The volume of the diffuse boundary is given by $|A^\epsilon| \approx |A| \times \epsilon$, where ϵ is a regularizing length scale parameter, and the approximation is exact as long as ϵ small compared to the smallest radius of curvature of A . The diffuse character of A^ϵ can be interpreted in two ways: the diffuse region may be regarded as a physical domain over which the solid-to-gas reaction occurs; i.e. interpreting η as a reaction coordinate along the lines of the thickened flame model [27,28]; alternatively, A^ϵ is considered strictly as a numerical regularization. The present work adopts this second interpretation, and thus the “error” of the solution, $\mathcal{O}(\epsilon)$, vanishes as $\epsilon \rightarrow 0$. This is referred to as the sharp interface limit.

The solid phase is a heterogeneous composite consisting of an AP oxidizer embedded in an HTPB matrix. The modeling framework is sufficiently robust to account for other species, such as metal fuels or burn rate enhancers. Here, the simplest SCP composition in order to streamline the formulation of the model is considered. As with the solid/gas interface (defined implicitly by η), it is advantageous to define the geometry of the SCP implicitly as well. A “diffuse species field” $\phi \in H^1(\Omega, [0, 1])$ is used, in which the region $\{x \in \Omega : \eta(x) = 1, \phi(x) = 1\}$ corresponds to that occupied by solid AP; $\{x \in \Omega : \eta(x) = 1, \phi(x) = 0\}$ corresponds to the region occupied by solid HTPB. (Note that the ϕ field is not assigned meaning in the gaseous region.) Like η , the diffuse species field contains diffuse boundaries that are regularized to have a diffuse width, designated ζ . The diffusiveness of ϕ is considered to be a purely numerical regularization since the physical boundary is nearly perfectly sharp. Unlike η , the diffuse species field is simply assigned rather than evolved, and so the diffusiveness constraint is constant and determined at initialization.

As with most such models, this phase field model is built upon the definition of a free energy functional $f : H^1(\Omega, [0, 1]) \rightarrow \mathbb{R}$ that drives the evolution of the order parameter. Here (following [8]),




Fig. 2. Problem schematic illustrating the role of the two field variables η and ϕ . $\eta = 1$ represents the solid/unburned region and is the initial condition setting for the problem, while $\eta = 0$ implicitly defines the gaseous region. The diffuse interface A^ϵ located in the region $0 < \eta < 1$. c). $\phi = 1$ indicates the oxidizer AP, while $\phi = 0$ indicates the binder HTPB. The transition region between oxidizer and binder, indicated by $0 < \phi < 1$, is also diffuse in order to smoothly transition material properties with no grid dependence [8].

the following free energy functional is used:

$$f[\eta] = \int_{\Omega} \left[\lambda w(\eta) + \frac{1}{2} \epsilon^2 \kappa |\nabla \eta|^2 \right] dx, \quad (1)$$

where w is the chemical potential, λ can be interpreted as a Lagrange multiplier or non-dimensionalization factor, κ is the energy at the solid/gas interface. The parameter ϵ is included to allow control of the diffusive length scale without altering the physical behavior of the system. (See Section Appendix A for a detailed discussion.) The governing equation for η follows by the adoption of a kinetic relation. In some cases, a nonlinear kinetic law may be required [29,30]. In most cases including this, however, linear kinetics are sufficient [31]. Therefore, the order parameter evolution is set proportionally to the gradient of the free energy, leading to the following kinetic evolution equation

$$\frac{d\eta}{dt} = -\frac{L}{\epsilon} \frac{\delta f}{\delta \eta} = -L(\mathbf{x}) \left[\frac{\lambda}{\epsilon} \frac{dw(\eta(\mathbf{x}))}{d\eta} - \epsilon \kappa \Delta \eta(\mathbf{x}) \right], \quad (2)$$

where L is a rate constant referred to as “mobility.” The mobility directly controls the interface velocity, and it is essential that it be computed correctly in order to produce an accurate regression rate. Historically, the regression rate has been connected to pressure via a power-law relation, which correctly captures the pressure dependence of regression in a number of different SCP configurations. Previous work [8] showed that a species-dependent pressure power-law can be used to estimate regression in a phase-field model by assigning different power-law constants to the fuel binder, the oxidizer, and the interface region. While this heuristic approximation produces reliable results, it ultimately fails to capture the full effect of the coupling between solid and gas phases; this renders the model unreliable for scenarios in which more complex solid/gas interactions are at play. In this updated model, the pressure power law is replaced with an Arrhenius law,

$$L = A(\phi(\mathbf{x})) \exp \left(\frac{-E_a(\phi(\mathbf{x}))}{R_u T(\mathbf{x})} \right), \quad (3)$$

where A is the pre-exponential factor, E_a is the activation energy, R_u is the universal gas constant, and T is the temperature. This form for the mobility is more physically sound, and relies on physical constants that have physical interpretability. Of particular note is the absence of an explicit pressure term from the mobility calculation. Here, explicit pressure dependence is replaced by implicit dependence through the temperature evolution, which is in better keeping with the physics of the problem.

As it has been demonstrated in [8], the model accuracy has low dependence on the choice of chemical potential and can effectively be fitted as a fourth order polynomial function, with prescribed values and zero derivative at $\eta = 0$, $1/2$, and 1 . To track multiple species, the properties at the interface between materials are computed using a mixture rule on the diffuse species parameter ϕ (Fig. 2). In addition, it has been demonstrated in [8] that in the limit these equations retrieve sharp interface formulation, and ref. [26] has shown that the computational effects of the interface length-scale are dependent on the total amount of cells in the interface region, with a minimum of 8 cells for proper evaluation of the interface effects.

2.2. Thermal evolution driven by diffuse boundary conditions

As interface regression relies on temperature (Eq. (3)), it is essential to accurately evolve temperature within the solid phase as driven by the heat flux from the gas phase (denoted here and subsequently as q_0) across the diffuse boundary. Special care must be taken to impose a flux boundary condition using a diffuse boundary. One way is to parameterize sharp interface as the isosurface $A = \{\mathbf{x} \in \Omega : \eta = 1/2\}$ and then drive the thermal evolution accordingly; this is cumbersome, however, and ultimately counteracts the advantage of the diffuse boundary formulation. Instead, boundary conditions on diffuse boundaries can be imposed through source terms obeying the constraint that the sharp interface solution is recovered exactly in the limit as $\epsilon \rightarrow 0$. This was proposed in [32], and developed more recently for the Euler equations [26] and the equations of linear elasticity [25], and the formulation for the heat equation is presented here. Fourier's law and the thermal diffusion equations for a discrete boundary are given by

$$k \nabla T - \mathbf{q} = \mathbf{0} \quad \mathbf{x} \in \Omega_{\text{unburned}} \quad (4)$$

$$\rho c \frac{\partial T}{\partial t} - \nabla \cdot \mathbf{q} = 0 \quad \mathbf{x} \in \Omega_{\text{unburned}} \quad (5)$$

$$T - T_0 = 0 \quad \mathbf{x} \in \partial_D \Omega_{\text{unburned}} \quad (6)$$

$$\mathbf{n} \cdot \mathbf{q} + q_0 = 0 \quad \mathbf{x} \in \partial_N \Omega_{\text{unburned}} \quad (7)$$

where Ω_{unburned} is the limit of the support of η as $\epsilon \rightarrow 0$, $\partial_D \Omega_{\text{unburned}}$ and $\partial_N \Omega_{\text{unburned}}$ are the boundaries over which Dirichlet and Neumann conditions are applied respectively, q_0 is the heat

flux determined by solving the gas phase equations or (as discussed in the following section) through approximation, and $\mathbf{n}$ is the outward-facing normal. T_0 is the prescribed temperature if the boundary condition is Dirichlet. In the diffuse boundary context, following [26], the explicit boundary conditions are replaced with source terms. If the boundary condition is Neumann, then the diffuse boundary formulation is:

$$k\eta\nabla T - \eta\mathbf{q} = \mathbf{0} \quad (8a)$$

$$\eta\rho c \frac{\partial T}{\partial t} - \nabla \cdot (\eta\mathbf{q}) = |\nabla\eta| q_0. \quad (8b)$$

In this work, only Neumann conditions are imposed over the diffuse boundary. However, if a Dirichlet boundary condition is imposed, then the diffuse boundary formulation is

$$k\nabla(\eta T) - \eta\mathbf{q} = \nabla\eta k T_0 \quad (9a)$$

$$\eta\rho c \frac{\partial T}{\partial t} - \eta\nabla \cdot \mathbf{q} = 0. \quad (9b)$$

In the diffuse boundary formulation, the boundary-wise prescribed values for temperature T_0 and heat flux, are modified to be functions over the diffuse boundary, having constant value along the normal direction equal to the prescribed values in the discrete boundary case. It can be shown that both converge to the discrete boundary solution in the limit as the diffuse boundary goes to zero. Applying the product rule, and observing that $\nabla\eta = -\mathbf{n}|\nabla\eta|$, yields

$$\underbrace{\eta(k\nabla T - \mathbf{q})}_{f_1} = |\nabla\eta| \underbrace{\mathbf{n}k(T - T_0)}_{g_1} \quad (10)$$

$$\underbrace{\eta\left(\rho c \frac{\partial T}{\partial t} - \nabla \cdot \mathbf{q}\right)}_{f_2} = |\nabla\eta| \underbrace{(\mathbf{n} \cdot \mathbf{q} + q_0)}_{g_2}. \quad (11)$$

It was demonstrated in [26] that, assuming the quantities indicated as f_1 , f_2 and g_1 , g_2 are bounded, and that ∇g_1 and ∇g_2 are bounded in the normal direction, then f_1 , $f_2 = 0$ for all $\mathbf{x} \in \Omega_{\text{unburned}}$ and g_1 , $g_2 = 0$ for all $\mathbf{x} \in \partial\Omega_{\text{unburned}}$ in the limit as $\epsilon \rightarrow 0$. Since f_1 , f_2 and g_1 , g_2 in Eqs. (10) and (11) correspond to the governing equation and boundary conditions, respectively, for Fourier's law and the heat equation, Eqs. (9a) and (9b) are concluded to be equivalent to Eqs. (4) to (7) in the sharp interface limit.

Special care must be taken to ensure boundary condition consistency in the diffuse interface formulation as on a discrete boundary domain. When Neumann (heat flux) boundary conditions are desired, then the governing equations to be solved are Eqs. (4) and (9b). It is worth remarking that one should not attempt to solve Eqs. (9a) and (9b) simultaneously, because this is the same as imposing Neumann and Dirichlet conditions simultaneously, which is unphysical and would lead to numerical error. In the event that one desires to enforce a Dirichlet boundary condition across the diffuse interface, one should solve Eqs. (6) and (9a). For other cases such as mixed or Robin boundary conditions, there are numerous strategies for consistently applying diffuse boundary conditions. In the present work, however, only heat flux conditions are of interest, and so the final form of the heat equation to be used here is

$$\eta \frac{\partial T}{\partial t} = \frac{1}{\rho c} \nabla \cdot (\eta k \nabla T) + |\nabla\eta| \frac{q_0}{\rho c} \quad (12)$$

Equation (12) effectively evolves the temperature in the solid phase of the system, but thermal evolution is undefined where $\eta = 0$ and can easily become unstable as only the total tempera-

ture field across both phases is well-defined. For the remainder of the section, the temperature T will be referred to as T_{solid} , and T_{total} will denote the temperature field across the entire domain. In order to construct T_{total} , an artificial “stand-in” fluid state to evolve the temperature outside the support of η is introduced. It must be emphasized that, here, the stand-in does not affect the solution in any way other than to prevent arithmetic overflow due to division by $\eta = 0$. The stand-in is implemented here in the same manner as that proposed in [26]; the total temperature evolved by introducing an auxiliary equation:

$$T_{\text{total}} = \eta T_{\text{solid}} + (1 - \eta) T_{\text{fluid}}, \quad (13)$$

where T_{fluid} is an assigned value for the fluid stand-in and T_{solid} is the temperature in the solid phase. The quantity (ηT_{solid}) is evolved according to a forward finite difference equation

$$\left. \frac{d(\eta T_{\text{solid}})}{dt} \right|_t = \eta|_{t-1} \left. \frac{\partial T_{\text{solid}}}{\partial t} \right|_t + T_{\text{solid}}|_{t-1} \left. \frac{\partial \eta}{\partial t} \right|_t, \quad (14)$$

which ensures the stability of T_{total} without affecting the solution inside the unburned domain. It is worth noting that while T_{fluid} is an assigned value in this work, it can easily be replaced by a full fluid-phase solver.

2.3. Reduced heat flux approximation

The temperature, which determines the evolution of the burn surface (Eq. (12)), is driven in the condensed phase by the heat flux of the reaction, transmitted to the surface from the gas. The most accurate way to determine the heat flux is to run a concurrent gas phase simulation that resolves the flow field and chemistry, coupling to the condensed phase either by reporting a heat flux or requiring continuity of temperature (see [4,22,33] and others). In this work, we propose a simplification that facilitates the approximation of the heat flux from the gas phase, using pre-calculated data, reducing the usual computational expense. This does not preclude, but rather complements, the use of a full gas-phase coupling: rapid screening can be accomplished using the simplified model, and then full-fidelity simulations (which shall be explored in future work) can be used to achieve greater accuracy for specific cases of interest.

In this section, we propose a fit to previously reported computational data that provides an estimate of heat flux based on composition ϕ , pressure P , and mass flux $\dot{m}$. By assuming a constant fluid pressure over the course of the regression, the surface heat flux can be approximated as a function of pressure and concentration of species. Here, heat flux data is taken from Gross and Beckstead [34], who reported the surface heat flux during AP/HTPB regression. By modeling a sandwich structure, it is possible to estimate the heat flux for pure HTPB (Fig. 3 a, dashed lines—left), pure AP (Fig. 3 a, dashed lines—right), and the spike in heat near the AP/HTPB interface due to the close-proximity diffusion flame (Fig. 3 a, dashed lines—center). Heat fluxes for pure AP, pure HTPB, and the peak heat flux at the interface are fitted using a least squares fit to a linear form (for pure AP and pure HTPB) and quadratic form (for the interface):

$$\hat{q}_{\text{AP}}(P) = C_1^{\text{AP}} P + C_0^{\text{AP}}, \quad (15a)$$

$$\hat{q}_{\text{HTPB}}(P) = C_1^{\text{HTPB}} P + C_0^{\text{HTPB}}, \quad (15b)$$

$$\hat{q}_{\text{D}}(P) = C_2^{\text{D}} P^2 + C_1^{\text{D}} P + C_0^{\text{D}}, \quad (15c)$$

where C are fitting parameters. These are assembled into a general approximation function for the heat flux as a function of species




Fig. 3. Heat flux and computational domain.

composition ($\phi \in [0, 1]$), pressure P , and mass flux $\dot{m}$:

$$q_0(\phi, P, \dot{m}) = \left[\hat{q}_{AP} \phi + \hat{q}_{HTPB} (1 - \phi) + \left(\frac{\zeta_0}{\zeta} \right) \left(\hat{q}_D \left(\frac{\zeta_0}{\zeta} \right)^2 \phi^{(1-\phi)} - 1 \right) \right] \times \tanh \left[\frac{\dot{m}(\phi, \eta)}{\dot{m}_0(\phi)} \right], \quad (16)$$

where $\dot{m}_0 = \phi \dot{m}_0^{AP} + (1 - \phi) \dot{m}_0^{HTPB}$ is the species-dependent maximum (or reference) mass flux, based on fit parameters $\dot{m}_0^{AP}$, $\dot{m}_0^{HTPB}$. The parameters ζ_0 and ζ correspond to the approximate width of the interface heat flux region (indicated in Fig 3a), and the diffuseness of the species-species interface, respectively. In this work, $\zeta = \zeta_0 \approx 100 \mu\text{m}$. However, the ratio is included so that the total heat flux, when integrated from $\phi = 0$ to $\phi = 1$ over an interval of width ζ , scales with ζ . This ensures that the total heat flux at the AP/HTPB interface does not depend on the diffuseness of the species-species boundary. The mass flux term approximates the heat flux in the unknown domain between the no-burn case (where heat flux is zero) and the steady case (as reported). The form provides an interpretable yet close fit to the computational data (Fig 3a) when $\dot{m} = \dot{m}_0$, implying that the burning is occurring evenly and at maximum rate at all regions. The mass flux is determined by,

$$\dot{m}(\phi, \eta) = \rho(\phi) \frac{\partial \eta}{\partial t}, \quad (17)$$

where $\rho(\phi)$ returns the density of AP or HTPB based on a linear mixture rule, and the derivative of η is determined by Eq. (2). This closes the feedback loop: the heat flux is determined by the regression rate, determined by the Arrhenius rate law, determined by thermal diffusion, determined by heat flux.

3. Computational methods

The accuracy of phase field models depends on sufficient grid resolution across the diffuse interface. It has been previously shown that 8 to 16 cells can be used to provide optimal results for SCPs regression modeling [8,26]. Implementing this level of refinement meshing onto large domains with relatively small interface length scales can become computationally expensive. Adaptive mesh refinement (AMR) is a general class of discretization methods that allows for the construction of a meshing grid with multiple levels of refinement that can be dynamically adjusted based on refinement and stability criteria. Block-structured AMR (BSAMR)

is a particular type of AMR that enhances performance by splitting the refinement levels into independent overlapping grids, using averaging and interpolation methods to establish communication between different levels of meshing. This streamlines computation and promotes excellent scalability on distributed memory platforms. BSAMR also enables the use of temporal sub-stepping, which in turn allows for the number of refinement levels to be changed arbitrarily without violating the stability condition.

The work presented here is based on methods implemented in the authors' in-house software (Alamo [35]), which is based on the open-source library AMReX [36]. The use of BSAMR to the problem of SCP regression is particularly effective due to the insulating properties of AP and HTPB (as described in [10]) that result in localization of the heat transport as well as the order parameter evolution. It should be noted that BSAMR is used to resolve only the burn front, not necessarily the entire geometry of the SCP simultaneously, which would lead to prohibitive cost and minimal improvement in accuracy. To effect refinement, every 10 timesteps, cells are tagged or de-tagged if the refinement criterion is met (or not). Here, the refinement criterion is:

$$|\nabla \eta| |\Delta x| \geq 0.001 \quad \text{or} \quad |\nabla T| |\nabla x| \geq 5.0 \quad \text{and} \quad \eta > 0, \quad (18)$$

where $|\Delta x|$ is the local magnitude of the grid spacing. This triggers refinement at the burned/unburned interface (where the gradient of η is nonzero) and wherever there are large thermal gradients within the solid phase. Thus, high resolution is maintained at the interfaces and the effects of thermal diffusivity are captured at sufficient resolution to ensure accuracy (Fig 3b). The $\eta > 0$ condition guarantees that spurious gradients in the gas phase do not trigger refinement unnecessarily; that is, it restricts refinement to the condensed phase (where $\eta > 0$). If a full gas phase is coupled to this solver, additional and separate criteria must be added.

The formulation is fully explicit, and so the solver uses an explicit forward Euler time integrator to evolve both the kinetic equation (Eq. (2)) and the heat equation (Eq. (12)). Spatial derivatives of first and second order are computed using central finite difference methods, and one layer of ghost cells is used to enforce domain boundary conditions and to communicate between course/fine levels. Explicit forward-Euler time integration is first-order accurate, and is only stable as long as the appropriate stability conditions are met [37]. Phase field models such as this are diffusion-limited, and so the stability condition² may be expressed

² Phase field literature sometimes refers to this as the CFL condition.

as

$$\frac{L\epsilon\kappa\Delta t}{\Delta x^2} < 1. \quad (19)$$

A similar condition must also hold for the heat equation, but since $L\epsilon\kappa \ll \alpha$ in general, it is (almost) always satisfied automatically. The computational domain for all cases $4\text{ mm} \times 1\text{ mm} (\times 1\text{ mm})$ for 2D (3D) simulations.

VisIt [39] is used for postprocessing and visualization. All regression rates are calculated by determining the farthest progression of the $\eta = 0.5$ isosurface. This mimics the experimental “trip wire” method, that also bases regression rates on the farthest progression of the interface (rather than, for instance, the average location of the interface).

4. Results

This section presents results for the application of the model to AP monopropellant and AP/HTPB solid composite propellants. It should be noted that the form of the model is general and not unique to these particular SCP species. In this work, however, the examples are limited to AP/HTPB. All parameters are kept the same across all tests; the number of calibrated parameters are minimal and limited mainly to the exponential pre-exponential factor and activation energy.

Physical properties are taken from known values. To determine the values of the fit parameters for AP, a one-dimensional calibration is performed. The pressure range for which the model is considered to be valid is 2 MPa to 6 MPa, as this is the range in which AP can sustain self-deflagration for those pressure levels (and for which experimental data is available). Calibration was accomplished through a grid search in parameter space and the judicious use of the bisection method to identify correct parameter values. In total, approximately 700 simulations were performed in order to evaluate best-fit values for the z_0 and $\dot{m}_0$ over the valid pressure range. General model parameters are presented in Table 1; parameters specific to particular species are presented in Table 2.

4.1. Thermal transport and deflagration of pure AP

A one-dimensional test of heat transfer in AP is considered to validate the diffuse boundary method for heat-flux driven temperature evolution. The domain is a homogeneous AP sample, and the interface is diffused but not permitted to regress (i.e., the pre-exponential factor is set to zero). The boundary conditions at all edges of the domain (that is, not the diffuse solids/gas boundary) are homogeneous Neumann (zero gradient/flux). At the diffuse

Table 1
AP and HTPB parameters for the solid phase model [8].


Parameter name	Symbol	Value
Base pressure [MPa]	p_p	1.0
Burned energy density [MPa]	w_0	0.0
Unburned energy density [MPa]	w_1	1.0
Activation energy density [MPa]	$w_{1/2}$	2.0
State interface width [mm]	ε	0.0005
Material interface width [mm]	ζ	0.0001
Chemical potential factor [–]	λ	0.001
Interface energy [MPa]	κ	1.0
Shape constant [–]	γ	0.02726

boundary, a Neumann heat flux condition is imposed via source term with a value of $1 \times 10^6 \text{ W/m}^2$. The systems is then evolved by Eq. (12) over 0.5 s with AMR active. The thermal evolution is measured over time, and compared to the sharp-interface solution for transient conduction described in equation 12 of [40] with $\alpha = 1.65 \times 10^{-7} \text{ m}^2/\text{s}$ and same length/time units for both numerical and analytic solution (Fig 4a). Very close agreement is observed between the diffuse boundary solution and the sharp interface solution. The main difference is a slight delay in the diffuse boundary solution, causing the temperature to be slightly lower. This can be attributed to a time shift that results from the diffusiveness of the interface: the initial heat flux is slightly limited and this causes the evolution to slightly lag the exact profile. The effect of the diffusive length scale parameter, and the convergence of the temperature profile to the exact solution, is presented in more detail in Section Appendix B.

Next, the evolution of temperature driven by heat flux with a moving boundary is considered in order to determine the steady-state behavior of the model for the mass flux $\rightarrow$ heat flux $\rightarrow$ thermal evolution $\rightarrow$ regression rate $\rightarrow$ mass flux cycle. Unlike previous phase field models for regression [8], the present model is stable in the unburned configuration, i.e., because there is no mass flux, there is no heat flux. In the current simulations, it is necessary to initiate the regression of the solid through an imposed heat flux, mimicking the physical process of ignition. Starting from a equilibrium, room-temperature solid, the sample is subjected to a heat flux of $1 \times 10^6 \text{ W/m}^2$, which initiates a rise in temperature and subsequent regression of the interface (Fig 4b) Thermal behavior of the solid propellant is as expected, with the temperature at the interface similar to that described in the literature, staying within the range of 500 K to 900 K [3]. The thermal profile in the solid region (Fig 4b—solid lines) is consistent with the thermal conductivity of AP, and reaches a steady state after about 10 ms. The thermal profile is included for the stand-in gaseous phase as

Table 2
Thermal properties of AP and HTPB and combined diffuse interface region. Values with references indicate literature from which material properties were taken. Values without reference were measured or calibrated as part of this work.

Parameter name	Symbol	Combined	AP	HTPB
Thermal Conductivity $[\frac{\text{J}}{\text{m s K}}]$	k		0.4186 [38]	0.1463 [38]
Density $[\frac{\text{kg}}{\text{m}^3}]$	ρ		1957 [22]	920 [22]
Specific Heat $[\frac{\text{J}}{\text{kg K}}]$	c_p		1297.9 [38]	2418.29 [38]
Thermal Diffusivity $[\frac{\text{m}^2}{\text{s}}]$	α		1.6×10^{-7}	6.5×10^{-8}
Pre-exponential Factor [–]	A		1.45×10^5 [22]	1.04×10^3 [22]
Activation Energy [K]	E_A		11,000 [4]	7500 [4]
Base Mass Flux $[\frac{\text{kg}}{\text{m}^2}]$	$\dot{m}_0$		1000	5000
$\hat{q}$ constant [–]	C_0	$C_0^D = -0.3715 + \frac{\zeta_0}{\zeta}$	$C_0^{\text{AP}} = 0.42$	$C_0^{\text{HTPB}} = 0.323$
$\hat{q}$ pressure factor $[\text{MPa}^{-1}]$	C_1	$C_1^D = 2.01$	$C_1^{\text{AP}} = 0.46$	$C_1^{\text{HTPB}} = 1.114$
$\hat{q}$ pressure square factor $[\text{MPa}^{-2}]$	C_2	$C_2^D = -0.09906$		
Ref. interface width $[\mu\text{m}]$	ζ_0	3.0	3.0	3.0



(a) Comparison between Alamo results (markers) and the exact solution (solid lines) for the temperature evolution for a 1D pure AP under constant heat flux. (b) Evolving temperature profile over 50 ms for an external pressure of 4 MPa and stand-in fluid temperature of 600 K. The dashed line indicates the temperature profile in the fictitious gas-phase stand-in.

Fig. 4. Verification and demonstration of the diffuse boundary heat transport model.




Fig. 5. Alamo regression rate prediction. (a) Comparison of experimental data for pure AP with Alamo's calculated average burn rate for the pressure range of 1 to 6 MPa. (b) Comparison of Alamo results and experimental data for a sandwich set up for pressure ranges of 0.8 MPa through 4.0 MPa.

well (Fig. 4b—dashed lines) to illustrate the effect of the stand-in: without the stand-in, the thermal evolution in the burned phase generally becomes unstable. If this model is coupled to a fully reacting fluid solver, the fluid stand-in would be replaced with the simulated fluid temperature in the $\eta = 0$ domain.

Regression rates were measured for pure AP in the 1–6 MPa pressure range and compared to experimental data from Price et al. [41] (Fig. 5a). The correlation between the model and experimental data is good. For instance, small changes in the Arrhenius rate law tended to have little effect on the regression, which appears to be due to the fact that changes in the regression rate in turn changes the temperature as well. Interestingly, the regression rate was relatively insensitive to the fitted model parameters. A substantial improvement in the present model is its ability to approximate the pressure cutoff for AP deflagration, although the model predicts a small burn rate at 1 MPa, it is close to the observed behavior. This can be attributed to the physically realistic feedback between the mass flux and heat flux in the model.

4.2. Burning of AP/HTPB sandwich composite

The “sandwich” configuration is useful for experimentally or numerically investigating deflagration behavior in a simple yet nontrivial configuration. In the present work, the sandwich structure provides a useful test case on which to further calibrate the model. Here, a laminate of HTPB with thickness equal to 1×10^{-4} meters is placed in an AP matrix. All of the parameters for AP in the previous section are left unchanged; only the HTPB parameters are varied. As with the pure AP case, domain boundary conditions are homogeneous Neumann. Therefore, the only variables calibrated using this case are the length scale of the diffuse

interface between AP and HTPB ζ_0 , and base mass flux for mixed substance $\dot{m}_0$. The size of the interface plays an important role in assuring convergence of the results, since it controls the model's ability to capture the effects of higher heat fluxes at the AP/HTPB interface. A minimum of eight interface cells is required for proper evaluation of those effects.

Regression rates were calculated for the AP/HTPB sandwich in the 0.8–4 MPa pressure range, and compared to corresponding experimental data from [42] (Fig. 5b). Since physical properties at the AP/HTPB interface are computed using the linear mixture rule, the burning rate at the material interface varies based on the burn rate of pure HTPB. It is observed that the regression rates for the sandwich case show a good match between the experiments and simulations.

4.3. Ignition in pure AP

The model's capability is further demonstrated with two sets of example simulations: (i) The model's ability to sustain deflagration, and (ii) the effects of laser power on the ignition process. In the first one, the model is applied in a 3-dimensional domain with the laser turned on for the initial 0.01 s of integration over a smaller region of the body's left face (0.1 by 0.1 mm square face). As with the previous cases, homogeneous Neumann boundary conditions are used on all domain boundaries. The burn front propagation is simulated over a 0.1-second interval (Fig. 6). It is observed that the burning starts locally while the laser is on but the system can self-sustain after it is shut off at $t = 0.025$ s. It can also be observed that areas outside of the initial laser spot eventually receive lateral heat through the mass/heat flux feedback cycle and conduction,




Fig. 6. The ignition energy (laser-induced heat flux) is applied to a middle square in a pure AP domain for the first 0.025 s and then shut off (a,b, and c). Once the ignition temperature is achieved the burn can self-propagate through the surface and regress throughout the domain (d).




Fig. 7. (a) The laser is turned on at time $t = 0$ s and turned off at time $t = 0.025$ s. If the laser power is high, the burn rate increases. If the laser energy is too low, no ignition occurs. If ignition occurs, the burn proceeds at a uniform rate after laser is deactivated even if over driven by a high input power. (b) Comparison between the simulations and the values from literature [43] indicates realistic ignition times and scaling.

causing it the regression interface to propagate along the body's left face.

An analysis of the effect of laser power on the system is also performed. The laser is set as a pulse that starts adding energy to the system at 0.0 s and stops at 0.025 s, with a total integration time of 0.05 s. It is observed that the AP would sustain self-deflagration if the total energy added to the system reaches the activation energy. For the pure AP case, the total energy added is enough to raise the surface temperature to 400 K, and the minimum laser energy to reach that temperature under 0.025 s was found to be $0.9 \times 10^6 \text{ W/m}^2$. If the laser power is below that value, the temperature peaks at 0.025 s and then cools. This result is particularly interesting since similar have previously been reported [6], and no further calibration from the ones performed in Sections 4.1 and 4.2 were made to match this observed ignition temperature. Moreover, it is observed that the burn rate directly depends on the amount of energy added to the interface as it regresses (Fig. 7a). It is also apparent that after the laser is shut off, the burn rate of all simulations that reached the ignition energy approach the self-sustained burn rate. Oscillations in the effective burn rate can be observed as a consequence of the slow propagation of energy through conduction and quick local temperature increase as the evolving interface regresses.

Evaluating the solid propellant response to different levels of ignition energy is crucial for allowing the proper design of igniters. Baer and Ryan have developed correlations between ignition time and ignition energy based on variables such as surface heat flux, pressure, and material composition [43], which relates the ignition time and the ignition energy flux to ignition time by assuming a homogeneous, semi-infinite body with uniform initial temperature and constant heat flux (see Equation 4 in [43]). Since the proposed assumptions are natural conditions for a computational simulation, simulations are performed for pure AP cases in the pressure range of 2 MPa through 6 MPa for six different laser power levels, for a total of 24 samples. Lower laser power was tested but did not

initiate regression. Baer's equation is not directly dependent on pressure but its effects are captured by the surface reaction frequency factor and the transmitted energy per unit reaction, which is equivalent to the pre-exponential factor in the Arrhenius law. Simulation results are compared to the solution of Baer's equation relating ignition time to heat flux, and a reasonable agreement is observed (Fig. 7b). Again, it should be noted that the model used the calibration coefficients found in the previous section to produce these results.

4.4. Burning of packed AP spheres in HTPB matrix

A key advantage of the diffuse interface approach is its ability to account for various propellant compositions with arbitrary morphology. To demonstrate this, a number of 3-dimensional simulations were run to determine regression rates in an SCP consisting of AP spheres packed into a HTPB matrix. The packed sphere configurations were generated using a Miller packing scheme, as described in [44,46], which consists of inserting spheres in the domain with an initial low volume fraction and then resizing the spheres to achieve the target mass fraction. To ensure accurate comparison to experimental data, 4 different 3-dimensional packings have been created with mass fractions between 60% and 80% and total AP spheres of 500 and 1000, for a total of 40 different packings (Fig. 8). To generate 2-dimensional packings, the 10 3-dimensional samples are sliced into 10 cross-sections. Slices with a 2-dimensional packing fraction less than 40% were removed, leaving a total of 80 2-dimensional samples. As with the previous cases, homogeneous Neumann boundary conditions are used on all domain boundaries. Ignition was effected with a laser power of $1.0 \times 10^6 \text{ W/m}^2$ applied to the boundary during the initial 0.025 s. The time step on the coarsest level of the simulation is set to 5×10^{-6} s to ensure that the stability criterion is met. The samples were simulated for 4 different pressures, resulting in a total of 320 2-dimensional simulations and 160 3-dimensional simula-

Diffuse interface method for solid composite propellant ignition and regression




Fig. 8. 3D samples of Miller Packed AP spheres [44] and particle distribution. (a) 60% AP mass fraction. (b) 83% AP mass fraction. (c) AP particle size histogram for 60% AP mass fraction. (d) AP particle size histogram for 83% AP mass fraction.




Fig. 9. Burn positions (left) and instantaneous interface velocities (right), calculated using central difference. The position is evaluated as the furthest position where $\eta = 0.5$.

tions. The burn rate represents the average value of all instantaneous burn rates between every time step and the error bar shows its standard deviation.

The burn surface location, for purposes of regression rate calculation, is determined by evaluating the extreme value (in the burn direction) for which $\eta \leq 0.5$. A rolling averaged central difference stencil is then used to calculate the instantaneous regression rate (Fig. 9). This method is used because it is a numerical analog to the experimental break-wire method, and because it is conservative (i.e., cannot underestimate the regression rate). It also eliminates the potential of regression rate pollution by unburned chunks of material left behind, which would affect the interface position if calculated based on an average.

The regression rate results, along with error bars indicating standard deviation in instantaneous burn rate, are compared with experimental data from [45] (Fig. 10). Other work (e.g., [47]) may be used to investigate more subtle trends connecting particle size to regression rate; this will constitute future work. A significant

variation in the average burn rate across different samples with the same pressure and composition is observed. This is expected since the burn rate is highly influenced by particle distribution and particle size. Of particular interest is the variation (or lack thereof) in regression rate with temperature for low loadings for 2D vs 3D simulations: whereas 3D exhibit a substantial burn rate at high pressure for all cases, 2D exhibits almost no burn rate at high pressure for $\leq 75\%$ AP. This is because of the much higher interface-to-volume ratio in 3D vs 2D, which results in better burn efficiency and greater heat transfer. The average temperature was determined to be in the range of 700 K through 1200 K, which is within the expected range, (Fig. 11).

The burn rate dependency on physical characteristics of SCPs, such as mass fraction and particle size, are well reported, and the simulations in this work have been able to mimic those behaviors. For instance, as the AP mass fraction decreases below 55% the HTPB volume fraction is high enough that deflagration ceases and only a very slow regression is observed. This is often observed in




Fig. 10. Comparison of Alamo burn rate with experimental data from Kohga [45] for a packed spheres set up. (a) Average and standard deviation for 320 samples of 2D simulations with mass fractions in the range of 55% through 85%; (b) Average Burn rate for 40 Samples of 3D simulations with AP mass fractions of 60% through 80%; Spheres represent pickings with 1000 AP particles and squares are packings of 500 AP particles.




Fig. 11. Evolution of a 65% AP mass fraction samples over the course of 0.2 s. Color corresponds to surface temperature.

low-pressure simulations where the thermal and chemical forces are not strong enough to initiate the regression if no AP particles are present in the front interface. If this pure HTPB interface is located in the middle of the domain, the result is a high increase in the local temperature with minimal or no interface movement. In such cases, the local temperature at the interface propagates very quickly, which can cause the results to diverge if the time step is too large. Robustness requires careful setup of integration time or an adaptive time step based on local properties. Particle size impact on burning rate has also been demonstrated in these simulations, decreasing the average particle size reflected in an increase in burn rate. This phenomenon is widely reported in the literature [9,48,49] as the increase in the surface area between AP and HTPB accelerates the regression rate due to the close proximity of the premixed flame. This is observed to be true for most cases, but large AP particles in low-pressure configurations increase the burn rate since the self-deflagration of AP has a faster burn in those cases.

Overall, the results show reasonable accuracy when compared to experimental data, especially considering the minimal number of fitted parameters. For the higher AP loadings, the interface is more planar, and therefore more geometrically similar to the sandwich configuration from which the heat flux data is measured. For lower loadings, on the other hand, the interface complexity is substantially increased, which alters the amount of heat transferred into the solid.

5. Conclusion

In this work, a computational model is presented that predicts regression rate and interface morphology using a diffuse interface model. This work builds on the previously developed phase field model for regression, but replaces the heuristic pressure-dependence with a physically-motivated Arrhenius law informed by local temperature. The heat equation is formulated, using the diffuse boundary method, to impose heat flux boundary conditions

using an order parameter-based source-term. This is coupled with a stand-in model of the fluid phase in order to regularize the solution and prevent numeric overflow. The local temperature determines the local mobility of the regression surface, and is now directly linked to physical properties and can account for both thermal and chemical kinematic effects. The thermal behavior of the burning of solid composite propellants is successfully modeled and the temperature profile has great agreement with experimental testing of those materials. Regression rates were validated using pure AP and AP/HTPB sandwich cases. Based on those calibrations 2- and 3-dimensional simulations of packed spheres cases were performed and the regression behavior has shown reasonable qualitative results that followed experimental data for different pressure and material compositions. It was also demonstrated that the model accurately captures ignition behavior of AP when subjected to thermal loading.

Though not the central focus of this work, the runtime performance of the present model was compared to that of the previous work [8]. It should be emphasized that a direct comparison is difficult to obtain, as the mesh changes dynamically. An overall run time increase of 20% to 30% was observed for the present model compared to [8] for equivalent problems. This is not surprising, as the present model requires solving the heat transport equations in addition to phase field, and as a result, produces a greater amount of useful output. Part of the increase in runtime is because of additional mesh refinement triggered by the thermal evolution equations, causing the mesh to be larger. This was verified by calculating the processing time per cell, which (surprisingly) decreased by between 10% and 20%. We attribute this increase in performance to improvements that were made in other unrelated parts of Alamo and AMReX.

The primary limitation of the present model is its reliance upon a reduced-order estimate of heat flux from computational sandwich tests. This simplification is responsible for the error in the predicted regression rate, which is consistent with the expected correlation between regression surface complexity and error. To obtain heat flux information that is accurate for complex regression surfaces, the gas phase must be resolved simultaneously. Such an implementation is outside the bounds of the present work; however (as previously stated) the two phases can be readily coupled across a diffuse boundary [25,26]. Such a coupled solver constitutes future work; however, the present results justify the use of the model for analysis of SCP regression and morphological dependence.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRediT authorship contribution statement

Maycon Meier: Formal analysis, Writing – original draft, Software, Validation. **Emma Schmidt:** Formal analysis. **Patrick Martinez:** Formal analysis. **J. Matt Quinlan:** Supervision, Writing – review & editing. **Brandon Runnels:** Supervision, Methodology, Funding acquisition, Writing – review & editing.

Acknowledgments

Authors MM, ES, PM, and BR acknowledge support from the Office of Naval Research, grant # N00014-21-1-2113. PM acknowledges support from the Caltech Summer Undergraduate Fellowship (SURF) Program. This work used the INCLINE cluster at the Uni-

versity of Colorado Colorado Springs. INCLINE is supported by the National Science Foundation, grant #2017917.

Appendix A. Governing equation invariance with respect to diffusive length scale

In this section we briefly discuss the role of the length scale parameter ϵ in the governing equation, which is derived from Eq. (1):

$$f[\eta] = \int_{\Omega} \left[\lambda w(\eta) + \frac{1}{2} \epsilon^2 \kappa |\nabla \eta|^2 \right] dx,$$

and depends on ϵ . The length scale parameter is included in this way in order to tune the diffusiveness of the interface without altering the physics. Therefore, proper inclusion of ϵ is necessary to guarantee that the energetics and the kinetics of the system are invariant with respect to it. (It should be noted that some authors write the free energy functional alternately as

$$g[\eta] = \int_{\Omega} \left[\frac{1}{\epsilon} \lambda w(\eta) + \frac{1}{2} \epsilon \kappa |\nabla \eta|^2 \right] dx, \quad (\text{A.1})$$

where the kinetics are expressed equivalently as

$$\frac{\partial \eta}{\partial t} = -L \frac{\delta g}{\delta \eta}, \quad (\text{A.2})$$

which may arguably be more aesthetically pleasing. The difference between f and g is that f is a *bulk* free energy (i.e., energy per volume) whereas g is a *interfacial* free energy (i.e., energy per area). Neither is more correct than the other; however, f remains bounded in the limit as $\epsilon \rightarrow 0$, whereas g does not. Therefore, we prefer to use f .) Next we consider the invariance of the system with respect to ϵ . For this, it is useful to assume compact support of $|\nabla \eta|$. (This is not strictly true in the general case, but it is a very reasonable approximation.) This allows us to split the domain into:

$$\Omega = \Omega_{int} + \Omega_{ext} + \partial_{\epsilon} \Omega \quad (\text{A.3})$$

where Ω_{int} is the support of $\eta = 1$, Ω_{ext} the support of $\eta = 0$, and $\partial_{\epsilon} \Omega = A \times [-\epsilon/2, \epsilon/2]$ the support of the gradient. It is immediately apparent that the driving force inside Ω_{int} and Ω_{ext} is zero (recall that $w(\eta)$ is locally minimized at both values), meaning that only the behavior inside $\partial_{\epsilon} \Omega$ must be considered. As long as ϵ is sufficiently smaller than the smallest radius of curvature of the interface, the boundary portion of the integral can be approximated to arbitrary precision as

$$f_{\partial_{\epsilon} \Omega}[\eta] = \int_A \int_{-\epsilon/2}^{\epsilon/2} \left[\lambda w(\eta) + \frac{1}{2} \epsilon^2 \kappa \left(\frac{d\hat{\eta}}{ds} \right)^2 \right] ds dA \quad (\text{A.4})$$

where s is the distance to the midpoint of the diffuse interface and $\hat{\eta}$ is the profile of the order parameter normal to the boundary. Performing a change of coordinates:

$$\hat{s} = \frac{s}{\epsilon} \quad d\hat{s} = \frac{ds}{\epsilon} \quad \hat{s}(\epsilon/2) = 1/2 \quad \hat{s}(-\epsilon/2) = -1/2 \quad (\text{A.5})$$

yields

$$f_{\partial_{\epsilon} \Omega}[\eta] = \epsilon \int_A \int_{-1/2}^{1/2} \left[\lambda w(\eta) + \frac{1}{2} \kappa \left(\frac{d\hat{\eta}}{d\hat{s}} \right)^2 \right] d\hat{s} dA. \quad (\text{A.6})$$

Substituting into the kinetic equation (Eq. (2)) shows that,

$$\frac{\partial \eta}{\partial t} = L \frac{\delta}{\delta \eta} \left[\int_A \int_{-1/2}^{1/2} \left[\lambda w(\eta) + \frac{1}{2} \kappa \left(\frac{d\hat{\eta}}{d\hat{s}} \right)^2 \right] d\hat{s} dA \right], \quad (\text{A.7})$$

which is independent of length scale. Therefore, the local velocity of the interface is independent of ϵ , as long as (recall) ϵ is sufficiently smaller than the smallest radius of curvature. For a more




Fig. B.1. The effects of the interface size on the thermal solver were evaluated for different levels of ϵ .

technical discussion of the invariance of speed with respect to ϵ , we refer the reader to [8]. The reader may also consider [26] for a discussion of scaling behavior of diffuse interface methods with respect to boundary diffuse width.

Appendix B. Impact of diffusive boundary width on heat transfer

Convergence of the diffuse interface method to the exact solution with respect to diffuse interface width is shown by comparing the solution for decreasing values of ϵ (Fig. B.1).

References

- [1] A. Davenas, Development of modern solid propellants, *J. Propul. Power* 19 (2003) 1108–1128.
- [2] B. Mason, C. Roland, Solid propellants, *Rubber Chem. Technol.* 92 (2019) 1–24.
- [3] S. Chaturvedi, P.N. Dave, Solid propellants: AP/HTPB composite propellants, *Arab. J. Chem.* 12 (2019) 2061–2068.
- [4] L. Massa, T. Jackson, M. Short, Numerical solution of three-dimensional heterogeneous solid propellants, *Combust. Theor. Model.* 7 (2003) 579.
- [5] W. Cai, P. Thakre, V. Yang, A model of AP/HTPB composite propellant combustion in rocket-motor environments, *Combust. Sci. Technol.* 180 (2008) 2143–2169.
- [6] P.W.M. Jacobs, H. Whitehead, Decomposition and combustion of ammonium perchlorate, *Chem. Rev.* 69 (1969) 551–590.
- [7] J. Chen, T. Brill, Chemistry and kinetics of hydroxyl-terminated polybutadiene (HTPB) and diisocyanate-HTPB polymers during slow decomposition and combustion-like conditions, *Combust. Flame* 87 (1991) 217–232.
- [8] B. Kanagarajan, J.M. Quinlan, B. Runnels, A diffuse interface method for solid-phase modeling of regression behavior in solid composite propellants, *Combust. Flame* (2022).
- [9] C. Dennis, B. Bojko, On the combustion of heterogeneous AP/HTPB composite propellants: a review, *Fuel* 254 (2019) 115646.
- [10] Y. Chen, D.R. Gueldenbecher, K.N. Hoffmeister, M.A. Cooper, H.L. Stauffacher, M.S. Oliver, E.B. Washburn, Study of aluminum particle combustion in solid propellant plumes using digital in-line holography and imaging pyrometry, *Combust. Flame* 182 (2017) 225–237.
- [11] B. Crawford, C. Huggett, F. Daniels, R. Wilfong, Direct determination of burning rates of propellant powders, *Anal. Chem.* 19 (1947) 630–633.
- [12] N. Eisenreich, H. Kugler, F. Sinn, An optical system for measuring the burning rate of solid propellant strands, *Propell. Explos. Pyrot.* 12 (1987) 78–80.
- [13] M.W. Beckstead, R. Derr, C. Price, A model of composite solid-propellant combustion based on multiple flames, *AIAA J.* 8 (1970) 2200–2207.
- [14] N.S. Cohen, L.D. Strand, An improved model for the combustion of ap composite propellants, *AIAA J.* 20 (1982) 1739–1746.
- [15] L. Massa, T. Jackson, J. Buckmaster, F. Najjar, Fluctuations above a burning heterogeneous propellant, *J. Fluid Mech.* 581 (2007) 1–32.
- [16] S. Gallier, M. Plaud, A model for solid propellant burning fluctuations using mesoscale simulations, *Acta Astronaut.* 158 (2019) 296–303.
- [17] A. Di Mascio, E. Martelli, M. Bernardini, F. Stella, A. Neri, Boundary conditions for inhomogeneous combustion simulation in solid rocket motors, *AIAA J.* 61 (5) (2023) 2314–2318.
- [18] M. Summerfield, Burning mechanism of ammonium perchlorate propellants, *Solid Propellant Rocket Research* (1960), pp. 141–182.
- [19] A. Hegab, T. Jackson, J. Buckmaster, D. Stewart, Nonsteady burning of periodic sandwich propellants with complete coupling between the solid and gas phases, *Combust. Flame* 125 (2001) 1055–1070.
- [20] T. Kumar, C.H. Rycroft, T.L. Jackson, Eulerian thermo-mechanical simulations of heterogeneous solid propellants using an approximate projection method, *Combust. Flame* 219 (2020) 198–211.
- [21] C. Vijay, P. Ramakrishna, Estimation of burning characteristics of AP/HTPB composite solid propellant using a sandwich model, *Combust. Flame* 217 (2020) 321–330.
- [22] X. Wang, T. Jackson, L. Massa, Numerical simulation of heterogeneous propellant combustion by a level set method, *Combust. Theor. Model.* 8 (2004) 227.

- [23] J. Choi, M. Cakmak, S. Menon, Simulations of composite solid propellant combustion using adaptive mesh refinement, 49th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition (2011), p. 417.
- [24] X. Xue, Y. Yu, Z. Ye, Heat and mass transfer mechanism of micro-combustion system with dual-fuel at high environmental load, *Appl. Therm. Eng.* 200 (2022) 117698.
- [25] V. Agrawal, B. Runnels, Robust, strong form mechanics on an adaptive structured grid: efficiently solving variable-geometry near-singular problems with diffuse interfaces, *Comp. Mech.* (2023).
- [26] E.M. Schmidt, J.M. Quinlan, B. Runnels, Self-similar diffuse boundary method for phase boundary driven flow, *Phys. Fluids* (2022).
- [27] F. Proch, A.M. Kempf, Numerical analysis of the cambridge stratified flame series using artificial thickened flame LES with tabulated premixed flame chemistry, *Combust. Flame* 161 (2014) 2627–2646.
- [28] O. Colin, F. Ducros, D. Veynante, T. Poinso, A thickened flame model for large eddy simulations of turbulent premixed combustion, *Phys. Fluids* 12 (2000) 1843–1863.
- [29] M. Gokuli, B. Runnels, Multiphase field modeling of grain boundary migration mediated by emergent disconnections, *Acta Mater.* 217 (2021) 117149.
- [30] L. Guin, D.M. Kochmann, A phase-field model for ferroelectrics with general kinetics, Part I: model formulation, *J. Mech. Phys. Solids* 176 (2023) 105301.
- [31] V.L. Ginzburg, L.D. Landau, On the theory of superconductivity, *On Superconductivity and Superfluidity*, Springer (2009), pp. 113–137.
- [32] H.-C. Yu, H.-Y. Chen, K. Thornton, Extended smoothed boundary method for solving partial differential equations with general boundary conditions on complex boundaries, *Model. Simul. Mater. Sci.* 20 (2012) 075008.
- [33] T.L. Jackson, Modeling of heterogeneous propellant combustion: a survey, *AIAA J.* 50 (5) (2012) 993–1006.
- [34] M.L. Gross, M.W. Beckstead, Steady-state combustion mechanisms of ammonium perchlorate composite propellants, *J. Propul. Power* 27 (2011) 1064–1078.
- [35] B. Runnels, V. Agrawal, W. Zhang, A. Almgren, Massively parallel finite difference elasticity using block-structured adaptive mesh refinement with a geometric multigrid solver, *J. Comput. Phys.* 427 (2021) 110065.
- [36] W. Zhang, A. Almgren, V. Beckner, J. Bell, J. Blaschke, C. Chan, M. Day, B. Friesen, K. Gott, D. Graves, et al., AMReX: a framework for block-structured adaptive mesh refinement, *J. Open Source Softw.* 4 (2019). 1370–1370.
- [37] S. DeWitt, S. Rudraraju, D. Montiel, W.B. Andrews, K. Thornton, PRISMS-PF: a general framework for phase-field modeling with a matrix-free finite element method, *npj Comput. Mater.* 6 (1) (2020) 29.
- [38] L. De Luca, M. Summerfield, Nonsteady burning and combustion stability of solid propellants, vol. 143, AIAA, 1992.
- [39] H. Childs, E. Brugger, B. Whitlock, J. Meredith, S. Ahern, D. Pugmire, K. Biagas, M. Miller, C. Harrison, G.H. Weber, H. Krishnan, T. Fogal, A. Sanderson, C. Garth, E.W. Bethel, D. Camp, O. Rübel, M. Durant, J.M. Favre, P. Navrátil, Visit: an end-user tool for visualizing and analyzing very large data, *High Performance Visualization—Enabling Extreme-Scale Scientific Insight* (2012), pp. 357–372, doi:10.1201/b12985.
- [40] S.M. Pineda, G. Diaz, C.F. Coimbra, Approximation of transient 1d conduction in a finite domain using parametric fractional derivatives, *J. Heat Transf.* 133 (2011).
- [41] C. Price, T. Boggs, R. Derr, The steady-state combustion behavior of ammonium perchlorate and cyclotetramethylenetetranitramine, 17th Aerospace Sciences Meeting, 1979.
- [42] G. Knott, M. Brewster, Modeling the combustion of propellant sandwiches, *Combust. Sci. Technol.* 174 (2002) 61–90.
- [43] A. Baer, N. Ryan, Ignition of composite propellants by low radiant fluxes, *AIAA J.* 3 (1965) 884–889.
- [44] A.L. Dye, J.E. McClure, C.T. Miller, W.G. Gray, Description of non-darcy flows in porous medium systems, *Phys. Rev. E* 87 (2013) 033012.
- [45] M. Kohga, Burning characteristics and thermochemical behavior of AP/HTPB composite propellant using coarse and fine ap particles, *Propell. Explos. Pyrot.* 36 (2011) 57–64.
- [46] S. Williams, A. Philipse, Random packings of spheres and spherocylinders simulated by mechanical contraction, *Phys. Rev. E* 67 (2003) 051301.
- [47] J.C. Thomas, G.R. Morrow, C.A. Dillier, E.L. Petersen, Comprehensive study of ammonium perchlorate particle size/concentration effects on propellant combustion, *J. Propul. Power* 36 (1) (2020) 95–100.
- [48] C. Fong, R. Smith, The relationship between plateau burning behavior and ammonium perchlorate particle size in HTPB-AP composite propellants, *Combust. Flame* 67 (1987) 235–247.
- [49] K. Fujimura, A. Miyake, Effect of the particle size and specific surface area of ferric oxide catalyst on the burning rate of AP/HTPB solid propellant, *Sci. Technol. Energ. Mater.* 71 (2010) 65–68.