An Image–Based Stochastic Model for Predicting Microstructure Effects on Shock-Induced Ignition and Burn in Pressed HMX*

Philip T. Melton † and K. A. Gonthier‡

The performance and safe handling of energetic solids is sensitive to variations in their microstructure. Predicting the shock response of these materials often involves the use of high-fidelity computational methods that attempt to fully resolve the coupled thermomechanical and reaction fields within the microstructure that influence shock-to-detonation transition (SDT). The high computational cost of these models, however, can hinder their usability for iterative design workflows. This paper describes a computationally efficient and energetically constrained stochastic dissipation model for predicting shock-induced ignition and burn of pressed granular HMX. The proposed approach statistically captures the leading-order microstructural influence on energy localization and burn wavelet propagation, enabling a direct and efficient linkage between microscale heterogeneity and macroscale burn behavior. By bypassing the computational expense of bridging resolved multiscale physics, this approach allows for the rapid virtual design and assessment of energetic microstructures for engineering applications.

I. Introduction

Energetic solids, such as propellants and explosives, can exhibit high energy release rates when burned under confinement or shocked. Shock-induced ignition and burn is key to understanding their safety and performance in applications. The fast and coupled multiphase and multiscale physics during shock loading pose significant challenges in theoretically, computationally, and experimentally characterizing both their micro and macroscale behavior, particularly for applications that involve comparatively large length and time scales [1]. This topic has been extensively studied in the past and represents an active area of ongoing research [2].

Particular emphasis has been placed on characterizing the influence of a material's microstructure on its shock-induced ignition-burn threshold in terms of shock strength, often referred to as a criticality threshold [3, 4]. Pressed materials often contain porosity that enhances shock sensitivity due to localized dissipation during pore collapse resulting in intense temperature fluctuations (i.e., hot-spots) that can trigger reaction. Various dissipative mechanisms, such as viscoplastic deformation and jetting, shear band formation, and gas compression, that are sensitive to pore shapes, orientations, interfacial topology, and distributions, are believed responsible for hot-spot formation. Real microstructures containing many randomly distributed and oriented defects may collectivey result in statistically rich interactions that occur when hot-spots overlap, merge, or compete which are difficult to describe using conventional approaches. Importantly, deterministic macroscale ignition and burn models do not explicitly capture this microstructural variability [5].

Recently, mesoscale modeling and simulation and molecular dynamics (MD) have been used to characterize shock-induced hot-spots due to pore collapse. Mesoscale simulations of image-derived two-dimensional (2D) microstructures have been used to explicitly resolve pore networks and track reactive hot-spot formation under shock loading [6]. While such simulations provide critical insight into defect–shock interactions, they remain computationally intensive and are typically restricted to a small number of images. Consequently, they are expensive to use on large ensembles of images required to obtain statistically converged predictions, nor can they be easily extended to three-dimensional (3D) microstructures where stochastic effects become even more pronounced. Molecular dynamics (MD) simulations remain largely constrained to isolated nanoscale defects due to prohibitive computational cost, limiting their ability to capture the collective effects of randomly distributed and oriented defects within real microstructures [7]. These constraints expose a critical capability gap: the absence of an efficient, scalable, and statistically rigorous methodology

*Submitted for presentation at the 2026 Region II Student Conference at the University of South Carolina in March 2026

†M.S. Graduate Assistant, Department of Mechanical and Industrial Engineering, Louisiana State University, Baton Rouge, LA 70803. AIAA Student Member 1900186. Email: *Philip.Melton@lsu.edu*

‡Professor, Department of Mechanical and Industrial Engineering, Louisiana State University, Baton Rouge, LA 70803. AIAA Senior Member 99599. Email: *kgonth1@lsu.edu*

for predicting shock-induced dissipation, ignition, and burn across realistic material ensembles. Bridging this gap is important for enabling the next generation of predictive, uncertainty-aware energetic material design.

Our research directly addresses this need by developing a new stochastic reduced-order framework that links microstructure to shock-induced ignition and burn behavior in energetic materials. The approach incorporates several key innovations. First, a dissipation localization element (DLE) methodology is formulated that embeds mesoscale physics through statistically constrained dissipation fields rather than explicit high-fidelity resolution, enabling dramatic reductions in computational cost while remaining anchored in macroscale shock energetics. Second, the inherently probabilistic structure of the framework provides a natural mechanism for uncertainty quantification by representing dissipation, ignition, and reaction variability through ensemble statistics. Third, the formulation is extensible, supporting iterative refinement as new insights emerge from MD simulations, mesoscale modeling, and advanced experimental diagnostics.

This paper presents a preliminary formulation and partial validation of a stochastic reduced-order modeling framework applied to idealized pore microstructures. The objective is to predict hot-spot temperature fields, ignition onset, and burn evolution under prescribed shock pressures and mass-specific dissipation consistent with macroscale compaction models. Section II describes the SEM-based image processing, constrained stochastic dissipation element formulation, and ignition–burn modeling approach. Section III presents results for representative simple pore configurations, and Section IV summarizes key findings and future directions.

II. Methodology

SEM micrographs, such as the HMX-FEM (fluid energy milled) microstructure shown in Fig. 2a [†] , are treated as statistically representative elementary volumes (REVs) for analyzing shock-induced ignition and burn. The methodology comprises three steps: (1) segmentation of the REV using image processing; (2) stochastic placement and superposition of dissipation localization elements (DLEs) to generate hot-spot temperature fields consistent with the effective macroscale compaction shock dissipation; and (3) extraction of ignition and burn time-of-arrival fields to construct microstructure-dependent burn relations as functions of shock strength.

As described below, the methodology is energetically constrained by the effective compaction shock pressure and dissipated work. While the spatial distribution and intensity of hot-spots depend strongly on microstructural details, it is assumed that the effective dissipated work depends primarily on the initial porosity and shock strength. This motivates treating bulk dissipation as a macroscopic constraint, while microstructure geometrically controls the spatial localization and statistics of dissipation that are important for shock induced ignition and burn.

Effective shock quantities may be obtained from macroscale continuum modeling, mesoscale simulations, and/or experimental measurements. In the present study, shock loading is determined using a macroscale continuum formulation that resolves the internal compaction-zone structure of a steady shock for a prescribed initial material state and shock speed [8, 9]. As an example, Fig. 1 illustrates the predicted spatial structure of a right-propagating steady shock by a macroscale compaction model for a shock speed $D = 3100$ m/s, particle speed $u_p = 800$ m/s, and pressure $P_s = 3.98$ GPa in pressed HMX with 15% initial porosity. This loading is sufficient to fully eliminate the initial porosity in the compacted state. Because the shock propagates supersonically relative to the ambient sound speed of HMX, it exhibits a leading discontinuous solid-phase shock, as reflected in the pressure and dissipation-rate $\overline{\dot{w}}_d$ profiles. The corresponding mass-specific dissipated work is $\overline{w}_d = 211.5$ kJ/kg, and the shock width, estimated from the dissipation-rate profile, is approximately 0.15 mm.

Image Processing: SEM micrographs are processed using a structured segmentation pipeline for reliable feature extraction. A Gaussian smoothing filter $(\sigma_x, \sigma_y) = (25, 25)$ is first applied to homogenize grain interiors while preserving pore morphology, and background normalization is performed by dividing the grayscale image by its smoothed counterpart to correct for uneven illumination. Adaptive Sauvola thresholding is then used to accommodate local contrast variations, followed by morphological cleaning with a 3×3 elliptical structuring element to remove isolated high-frequency artifacts while preserving macroscale features. Because smaller pores display higher ignition thresholds than larger voids [10], conventional threshold-based segmentation provides adequate resolution for the present stochastic analysis. More advanced approaches, including machine-learning-based methods, may be adopted if higher-fidelity pore characterization is required. The procedure is problem-specific and may require recalibration depending on material system and imaging conditions.

[†] Image provided by Dr. Chris Molek, AFRL/RWTE, Eglin Air Force Base, Florida, 2023.

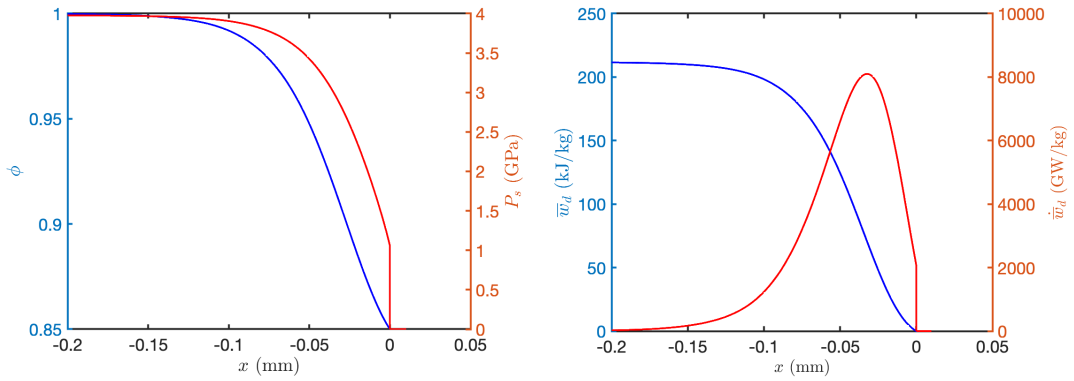

Fig. 1 Steady compaction shock structure for a shock particle speed of $u_p = 800$ m/s in pressed HMX with 15% initial porosity.

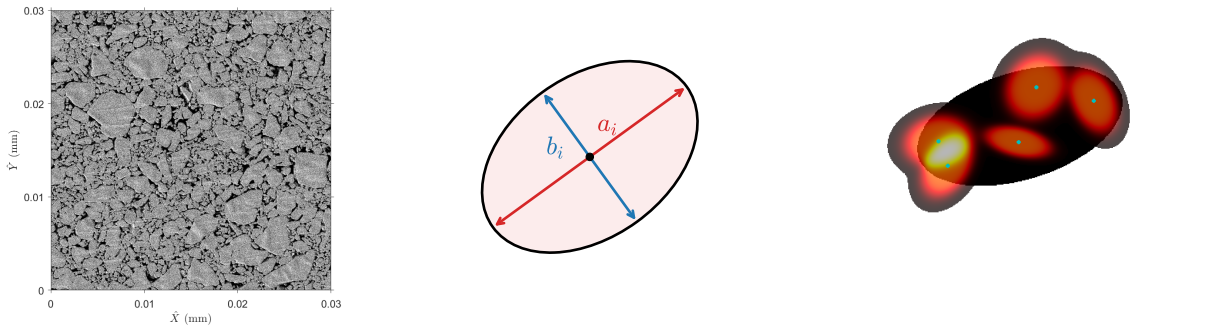

(a) HMX-FEM microstructure. (b) Dissipation localization element geometry. (c) Overlaid stochastic dissipation elements.

Fig. 2 Dissipation localization element geometry with stochastic distribution across a void region.

Stochastic Dissipation Model: The stochastic dissipation framework addresses the intrinsic complexity of pore-scale energy localization in heterogeneous microstructures, where shock transmission and reflection generate highly localized inelastic strain and heating fields that are expensive to directly resolve (e.g., MD simulations of a single 50 nm pore may exceed 10^6 CPU-hours [11]). To mitigate this cost, dissipation localization elements (DLEs) are introduced as computational surrogates for pore-scale hot-spots observed in mesoscale and MD studies. DLEs are probabilistically seeded near localization-prone features (e.g., pores and high-curvature interfaces), with geometry sampled from bounded, physically informed distributions. Each element is assigned a smooth dissipation kernel, and the ensemble is normalized to enforce consistency with prescribed bulk compaction work. Superposition of overlapping elements produces spatially intermittent dissipation fields, while repeated sampling enables efficient extraction of ignition-burn statistics. The result is an energetically consistent, microstructure-aware predictive framework at dramatically reduced computational cost.

In this approach, DLEs are randomly distributed within pores using a weighted sampling strategy in which the pore selection probability Ω_k increases with pore perimeter ℓ_k, reflecting the enhanced likelihood of plastic dissipation along larger collapse interfaces. Each DLE is represented by an elliptical level-set geometry whose parameters are sampled as functions of local pore morphology and imposed shock strength. For the i^{th} DLE ($i = 1, \ldots, N$), a dimensionless

distance function is defined as

$$\Gamma_i(\mathbf{x}) = \left[\frac{-(x-x_i)\sin\theta_i + (y-y_i)\cos\theta_i}{a_i}\right]^2 + \left[\frac{(x-x_i)\cos\theta_i + (y-y_i)\sin\theta_i}{b_i}\right]^2, \tag{1}$$

where $\mathbf{x} = (x, y)$ is position within the 2-D REV, (x_i, y_i) denotes the element center, and θ_i is a random orientation angle. The element boundary corresponds to $\Gamma_i = 1$, with $\Gamma_i < 1$ defining the interior. The ellipse axes are stochastically determined subject to microstructural and energetic constraints. The microstructural axis is defined as

$$a_i = r_{h,i}\,\Upsilon(0,1), \tag{2}$$

where $r_{h,i}$ is the hydraulic radius of the selected pore and $\Upsilon(0,1)$ is a uniform random variable on $[0,1]$. The dissipation axis is estimated from a randomized partition of the effective compaction work,

$$b_i = \frac{\rho_s\,\overline{w}_d}{\sigma_y}\left[\frac{\Upsilon(0,1)}{\ell_T/A_{\mathrm{REV}}}\right], \tag{3}$$

where ρ_s is the solid density, σ_y is the yield stress, ℓ_T/A_{REV} is the area-specific total pore perimeter within the REV, and $\overline{w}_d$ is the prescribed effective dissipated work. To ensure physically realistic localization and numerical robustness, the geometry is constrained such that $a_i b_i \le A_k$ and $1/5 \le b_i/a_i \le 5$, where A_k is the area of the selected pore.

Each DLE is assigned a smooth spatial dissipation kernel supported on the ellipse,

$$\phi_i(\mathbf{x}) = \frac{1}{2}\left[1 - \tanh\left(\eta\left(\sqrt{\Gamma_i(\mathbf{x})} - 1\right)\right)\right], \tag{4}$$

where η controls the sharpness of the dissipation transition near the element boundary ($\eta = 4$ in this study). This tanh-level-set kernel captures the sharp yet continuous dissipation gradients commonly observed in mesoscale simulations of shocked porous solids. The stochastic dissipation map is constructed by superposition,

$$\Phi(\mathbf{x}) = \sum_{i=1}^{N} \phi_i(\mathbf{x}), \tag{5}$$

and subsequently scaled to enforce global energetic consistency with the prescribed compaction work:

$$\mathcal{D}(\mathbf{x}) = \alpha\,\Phi(\mathbf{x}), \qquad \alpha = \overline{w}_d\left(\frac{1}{A_{\mathrm{REV}}}\int_{A_{\mathrm{REV}}} \Phi(\mathbf{x})\,dA\right)^{-1}. \tag{6}$$

This normalization guarantees that the spatially integrated dissipation matches the macroscale mass-specific work input. To determine the required number of elements, N, an iterative convergence analysis is performed in which N is increased by a factor of 1.5 with each iteration. The final value of N is selected once both the maximum and standard deviation of the dissipation field, $\mathcal{D}(\mathbf{x})$, satisfy a 5% convergence criterion. The functional form of the kernel, together with controlled overlap among DLEs, reproduces the spatial intermittency and clustering behavior observed in mesoscale collapse simulations (as illustrated in Fig. 2c).

Finally, the hot-spot temperature field is estimated directly from the dissipation field as

$$T(\mathbf{x}) = T_0 + \frac{\mathcal{D}(\mathbf{x})}{c_v}, \tag{7}$$

where c_v is the specific heat capacity and T_0 is the initial material temperature.

Ignition and Burn Model: The DLE-generated temperature field $T(\mathbf{x})$ establishes the ignition and burn dynamics by modeling how hot-spots ignite, interact, and form propagating burn wavelets. An advantage of this framework is that it does not assume simultaneous ignition of all hot-spots behind the shock, unlike many simplified ignition and growth models. Instead, each stochastic DLE realization produces a spatially varying ignition-time field based on hot-spot temperatures and established ignition criteria, capturing the time-staggered nature of hot-spot activation arising from spatially heterogeneous dissipation. For now, it is assumed that all material contained in the REV is instantaneously shocked resulting in a stochastic realization of the dissipation field that subsequently burns. This fast shock assumption can be later relaxed to account for finite shock propagation times through the REV microstructure.

Let $\tau(\mathbf{x})$ be a local temperature-dependent ignition-time field computed from $T(\mathbf{x})$ using an expression for adiabatic reaction times of HMX. In this work, adiabatic ignition times were described using the expression $\tau = \frac{1}{B}\exp(E_a/R_uT)$ provided by Henson and Smilowitz [12], where $B = 5.9 \times 10\ \mathrm{s}^{-1}$ and $E_a = 148.9$ kJ/mol. This expression was chosen due to its simplicity compared to other models and its ability to accurately predict shock criticality envelopes in mesoscale simulations for pressed HMX [13]. When a point $\mathbf{y}$ within the REV ignites at time $\tau(\mathbf{y})$, it initiates an infinitely thin circular burn front propagating at a constant speed c that can be related to the effective shock pressure P_s. The pressure-dependent burn speed is approximated by $c(\mathrm{mm/s}) = 10 \times P_s(\mathrm{GPa})^{1.81}$ based on HMX deflagration data provided in Refs. [14, 15]. The arrival time of a burn front originating from $\mathbf{y}$ at an arbitrary point $\mathbf{x}$ is given by

$$\tau(\mathbf{y}) + \frac{1}{c}\,|\mathbf{x}-\mathbf{y}|.$$

Because each point within the REV may either self-ignite or be consumed by a burn front from an earlier ignition site, an effective burn-arrival time field $\mathcal{T}(\mathbf{x})$ can be defined by

$$\mathcal{T}(\mathbf{x}) := \inf_{\mathbf{y}\in\Omega}\left[\,\tau(\mathbf{y}) + \frac{1}{c}\,|\mathbf{x}-\mathbf{y}|\,\right], \tag{8}$$

where $\inf[\cdot]$ is the greatest lower bound of the enclosed quantity which selects the earliest possible arrival among all paths through $\mathbf{y}$.

The region burned by time t is then the sublevel set of $\mathcal{T}$, given by $A_{\mathrm{burn}}(t) = \{\mathbf{x} \in A_{\mathrm{REV}} : \mathcal{T}(\mathbf{x}) \le t\}$, and the corresponding burned-area fraction is

$$F(t) = \frac{1}{A_{\mathrm{REV}}}\int_{A_{\mathrm{REV}}} H(t-\mathcal{T}(\mathbf{x}))\ dA, \tag{9}$$

where $H(\cdot)$ denotes the Heaviside function. Differentiating Eq. (9) gives the instantaneous burn rate

$$\dot{F}(t) = \frac{dF}{dt}(t) = \frac{1}{A_{\mathrm{REV}}}\int_{A_{\mathrm{REV}}} \delta(t-\mathcal{T}(\mathbf{x}))\ dA. \tag{10}$$

This formulation naturally captures wavelet coalescence, shielding, and sequential ignition, and it can be adapted to account for preferential propagation along microstructural interfaces as part of our ongoing work.

Though not addressed in this paper, time can be eliminated from Eqs. (9) and (10) to obtain a reduced-order ignition–burn relation of the form

$$\dot{F}(F) = R_{\mathrm{ign}}(F) + R_{\mathrm{burn}}(F), \tag{11}$$

where the ignition term $R_{\mathrm{ign}}(F)$ reflects spatially distributed, time-staggered hot-spot activation and is directly correlated with the effective shock dissipated work that drives local heating. The burn-wave contribution $R_{\mathrm{burn}}(F)$ follows a KJMA-type nucleation-and-growth formulation. Figure 3 gives a simple illustration of the propagation of burn wavelets originating at two ignition sites. This approach is similar to that of Hill [16, 17] but accounts for time-staggered ignition sites based on specific distributions of pores/defects within a material's microstructure. Ensemble analyses across many DLE realizations can yield statistically converged ignition–burn curves, parameter distributions, and uncertainty bounds, enabling a microstructure-dependent family of burn-rate laws suitable for uncertainty-aware predictive modeling.

III. Results and Discussion

This section gives representative results demonstrating the predictive capability and partial validation of the DLE methodology for idealized pore configurations. The stochastic temperature and ignition–burn fields are shown to reproduce key features observed in mesoscale simulations, including spatial intermittency and variability in ignition delay. The implications of these results for constructing a physically grounded, microstructure-dependent reduced-order ignition–burn relation, $\dot{F}(F)$, are briefly examined.

Single Pore Configuration: As a preliminary benchmark, we examine the ignition–burn response of a single 50 nm diameter pore corresponding to 8.73% porosity within the REV. This configuration serves as a canonical test case in many mesoscale and molecular dynamics investigations of hot-spot formation [18–20]. The computed effective shock

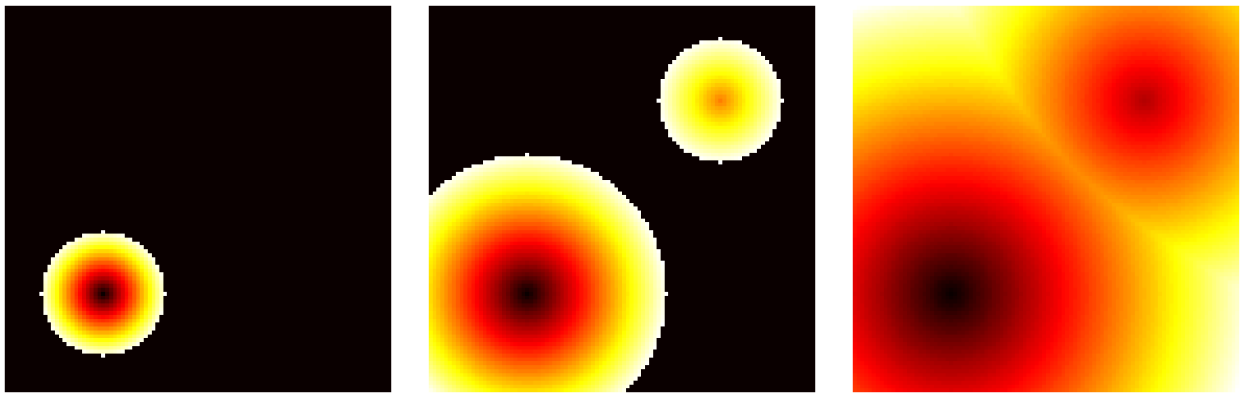

(a) Early Stage ($t = t_1$) (b) Intermediate Stage ($t = t_2$) (c) Merging Stage ($t = t_3$)

Fig. 3 Illustration of time-staggered hot-spot ignition. (a) An early hot-spot ignites at $t_{ign} = t_1$ and propagates a flamelet while the secondary site remains in a thermal induction phase. (b) The second hot-spot ignites at $t_{ign} = t_2$ and propagates a flamelet. (c) The thin flamelets coalesce and establish the time-of-arrival field.

dissipated energy for this case is $\overline{w}_d = 238$ kJ/kg corresponding to a shock pressure of $P_s = 7.16$ GPa and particle speed of $u_p = 1.0$ km/s. The dissipation-induced temperature field predicted by the stochastic DLE method is shown in Fig. 4a. The resulting field exhibits a mean temperature of 458.3 K and a peak temperature exceeding 1400 K. Although the detailed spatial distribution of temperature differs from that reported in Ref. [18], the average temperature, peak temperature, and characteristic hot-spot dimensions are comparable, yielding a similar integrated thermal mass. From an engineering perspective, these results indicate that the stochastic DLE framework produces a physically reasonable hot-spot temperature field for this benchmark configuration. Notably, the full stochastic evaluation of a temperature field required approximately 0.09 s of CPU time on a standard laptop workstation, highlighting the computational efficiency of the reduced-order approach.

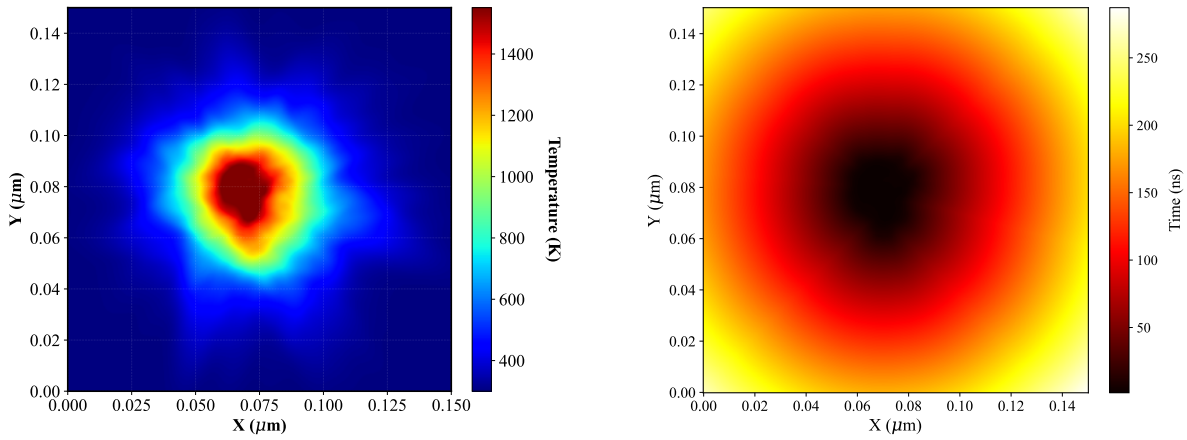

(a) Temperature Contours from Stochastic Dissipation Placement on a Single Pore Microstructure.

(b) Time-of-Arrival Burn Field.

Fig. 4 Stochastically generated temperature field and subsequent time-of-arrival field for a 50 nm circular void (8.73% porosity) with a shocked particle velocity $\mathbf{u_p} = 1.0$ km/s.

Figure 5 presents the ensemble ignition–burn response predicted by Eqs. (9) and (10) for stochastic variations in the hot-spot temperature and burn time-of-arrival fields shown in Fig. 4. Results from 20 independent realizations are summarized, with the ensemble mean and corresponding 95% confidence interval indicated. The predicted ignition

delays, ranging from approximately 50–100 ns, are comparable to the critical ignition times reported for similar pore sizes and temperatures in HMX [10].

Figure 5a shows pronounced high-amplitude, high-frequency burn-rate fluctuations during the ignition phase, reflecting the intermittent and spatially heterogeneous hot-spot temperature fields. As the reaction transitions to sustained burn, these fluctuations diminish and the burn rate evolves more smoothly, indicating a shift from localized, temperature-sensitive ignition dynamics to a more spatially integrated reaction regime. The ignition and burn-front phases are also evident in the $\dot{F}(F)$ curves in Fig. 5b. The ignition process rapidly consumes approximately 8% of the REV mass, corresponding to the initial porosity (a common assumption in ignition–burn modeling) before transitioning to slower sustained reaction. These trends are consistent with mesoscale results reported in Ref. [21] and resemble reaction-rate behavior associated with populations of Arrhenius-type reactors [2].

The total CPU time required to compute the time-of-arrival maps for all realizations was approximately 7.9 minutes (23.77 seconds per realization). Though not currently addressed, the ignition and burn timescales extracted from these ensemble simulations can be systematically isolated and used to construct a microstructure-dependent ignition–burn functional decomposition of the form given in Eq. (11).

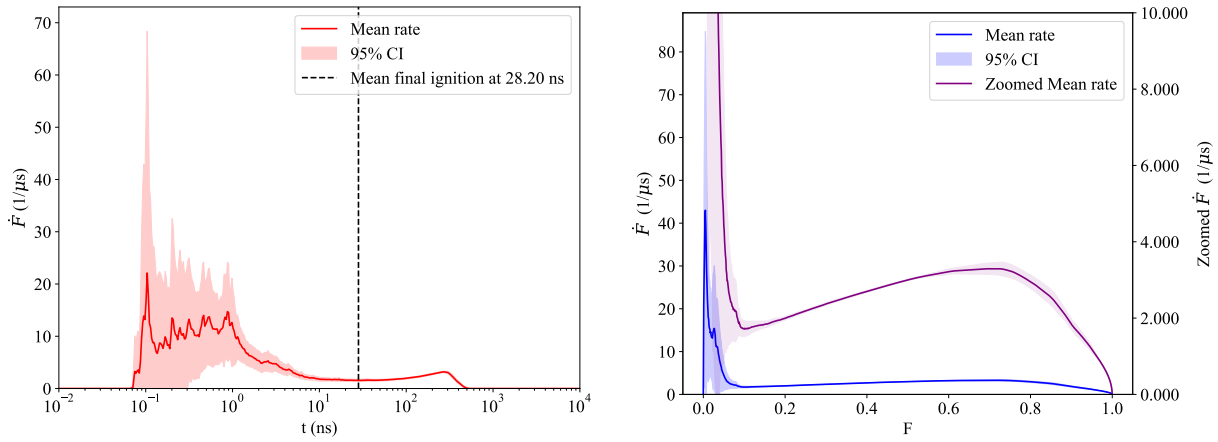

(a) Rate of Area Fraction Burned $\dot{F}(t)$ vs. Time for Single Pore Microstructure(t)

(b) Rate of Area Fraction Burned $\dot{F}(t)$ vs. Area Fraction Burned $F(t)$ for Single Pore Microstructure

Fig. 5 Rate of ignition-burn reaction rate F with respect to time t and fraction of domain area burned F for the single pore microstructure (u_p = 1.0 km/s).

Multi-Pore Configuration: We next examine the ignition–burn response of a random distribution of 15 pores, each with a diameter of 11.28 μm, corresponding to an overall porosity of 15% within the REV. The effective shock dissipated energy for this case is $\overline{w}_d$ = 83 kJ/kg corresponding to a shock pressure of P_s = 1.84 GPa and particle speed of u_p = 0.5 km/s. This multi-pore arrangement was selected to facilitate direct comparison with the inert mesoscale collapse simulations reported in Ref. [22], which examined interaction effects during pore collapse under similar loading conditions. Figure 6a gives a representative stochastic DLE realization of the resulting temperature field computed in 0.02 s on a standard laptop workstation. Consistent with the mesoscale simulations of Ref. [22], thermally activated interaction between adjacent hot-spots is observed when pores are in close spatial proximity. The stochastic DLE framework naturally produces variability in peak temperature and hot-spot intensity amongst and within individual pores, reflecting local geometric effects and stochastic dissipation partitioning. This variability is qualitatively consistent with trends reported in the mesoscale literature for comparable pore configurations.

Figure 7 presents the ignition–burn response predicted by Eqs. (9) and (10) for 20 stochastic realizations of the multi-pore temperature field shown in Fig. 6a. The ensemble mean and corresponding 95% confidence interval are indicated to quantify stochastic variability. The total computational time for all realizations was approximately 8.75 minutes (26.24 s per realization). The overall response exhibits features similar to the single-pore configuration; however, the presence of multiple interacting hot-spots broadens the distribution of ignition delays. This increased spread manifests as a substantially wider confidence interval during the ignition phase, reflecting the amplified variability

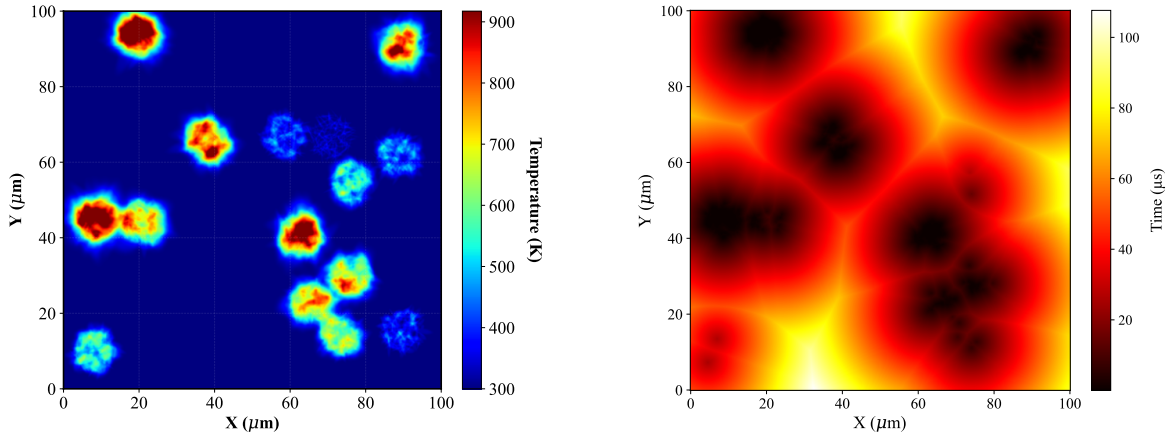

(a) Temperature Contours from Stochastic Dissipation Placement on a Multi-Pore Microstructure.

(b) Time-of-Arrival Burn Field.

Fig. 6 Stochastically generated temperature field and subsequent time-of-arrival field for a set of 11.28 μm pore (15% porosity) with a shocked particle velocity u_p = 0.5 km/s

introduced by staggered hot-spot activation.

The staggered ignition of multiple hot-spots strongly influences the ignition–burn response shown in Fig. 7. Following the high-amplitude ignition region, the bimodal behavior observed in Fig. 7b arises from continued ignition of lower-temperature, yet still critical, hot-spots after the first ignition event but prior to their complete consumption by burn waves emanating from earlier reaction sites. This regime therefore reflects a superposition of ignition- and burn-controlled contributions, i.e., a mixture of $R_{\text{ign}}(F)$ and $R_{\text{burn}}(F)$. Beyond the first bimodal peak, the burn rate exhibits continued growth toward full consumption, corresponding to a regime dominated primarily by $R_{\text{burn}}(F)$ as reaction fronts propagate and merge throughout the REV.

Owing to the lower shock strength in this case, the mean final ignition time in Fig. 7a is approximately three orders of magnitude greater than that observed for the single-pore configuration in Fig. 5a subjected to substantially higher shock pressure. Significant stochastic variability in ignition delay is again observed during the ignition phase, with residual variability extending into the early burn regime. This behavior reflects the sensitivity of multi-pore thermal interactions to local geometric configuration and dissipation partitioning under weaker shock loading.

IV. Conclusion

This work presents a stochastic reduced-order framework for predicting shock-induced hot-spot formation, ignition onset, and burn evolution in heterogeneous energetic microstructures derived directly from SEM images. The methodology integrates image-based representative elementary volumes (REVs), probabilistic dissipation localization elements (DLEs), and energetically constrained normalization to ensure consistency with macroscale compaction shock dissipation. In this manner, the framework establishes a controlled multiscale linkage between continuum shock energetics and microstructure-resolved ignition and burn behavior.

Controlled stochastic seeding of DLEs near geometrically relevant microstructural features produces spatially intermittent and clustered hot-spot temperature fields that qualitatively reproduce mesoscale collapse characteristics at a fraction of the computational cost. Ensemble simulations demonstrate statistical convergence of ignition and burn metrics using a modest number of realizations, enabling quantification of variability through confidence intervals and extraction of reproducible ignition–burn relations.

Benchmark studies of single- and multi-pore configurations illustrate several key findings. For a canonical 50 nm pore subjected to strong compaction loading, predicted ignition delays and peak temperatures are comparable to values reported in mesoscale and MD studies. Multi-pore configurations exhibit broadened ignition spectra due to staggered hot-spot activation and thermal interaction, leading to bimodal behavior in $\dot{F}(F)$ during the transition from

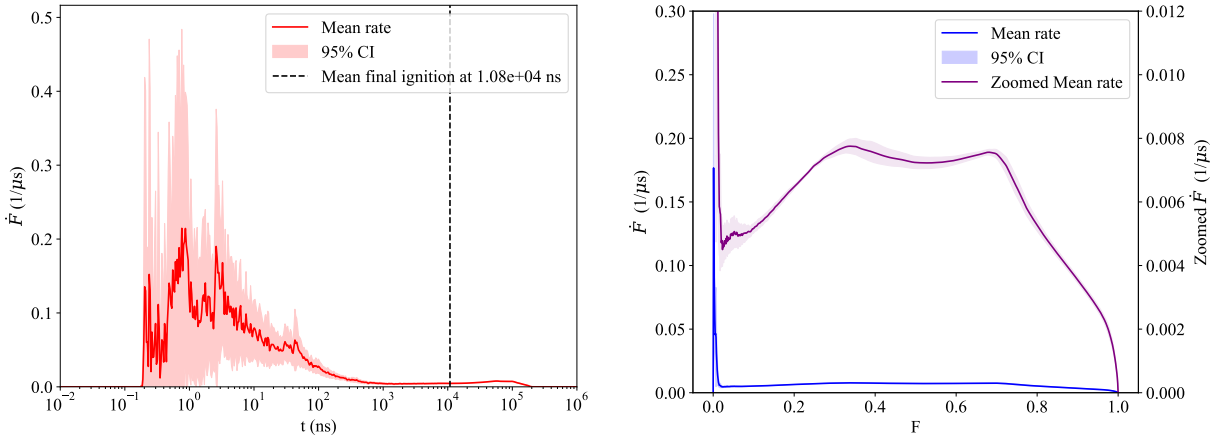

(a) Rate of Area Fraction Burned $\dot{F}(t)$ vs. Time (t) for multi-pore microstructure.

(b) Rate of Area Fraction Burned $\dot{F}(t)$ vs. Area Fraction Burned $F(t)$ for multi-pore microstructure.

Fig. 7 Rate of ignition-burn reaction rate F with respect to time t and fraction of domain area burned F for the multi-pore microstructure (u_p = 0.5 km/s).

ignition-controlled to burn-controlled regimes.

Future work will focus on extending the methodology to fully image-derived heterogeneous microstructures, incorporating additional physics such as strength evolution and thermal transport where appropriate, and systematically calibrating localization statistics against high-fidelity simulations and experiments. The present results demonstrate that controlled stochastic dissipation modeling offers a physically grounded and computationally efficient pathway for bridging macroscale shock loading and microstructure-dependent ignition behavior in pressed energetic materials.

References

[1] Field, J. E., ''Hot spot ignition mechanisms for explosives,'' *Accounts of chemical Research*, Vol. 25, No. 11, 1992, pp. 489–496.

[2] Handley, C., Lambourn, B., Whitworth, N., James, H., and Belfield, W., ''Understanding the shock and detonation response of high explosives at the continuum and meso scales,'' *Applied Physics Reviews*, Vol. 5, No. 1, 2018.

[3] Menikoff, R., ''Pore collapse and hot spots in HMX,'' *AIP Conference Proceedings*, Vol. 706, American Institute of Physics, 2004, pp. 393–396.

[4] Rai, N., and Udaykumar, H., ''Void collapse generated meso-scale energy localization in shocked energetic materials: Non-dimensional parameters, regimes, and criticality of hotspots,'' *Physics of Fluids*, Vol. 31, No. 1, 2019.

[5] Kittell, D. E., Yarrington, C. D., Lechman, J. B., Damm, D. L., and Baer, M. R., ''Modelling the fluctuations of reactive shock waves in heterogeneous solid explosives as stochastic processes,'' *Propellants, Explosives, Pyrotechnics*, Vol. 45, No. 2, 2020, pp. 295–315.

[6] Rai, N. K., and Udaykumar, H. S., ''Mesoscale simulation of reactive pressed energetic materials under shock loading,'' *Journal of Applied Physics*, Vol. 118, No. 24, 2015, p. 245905. https://doi.org/10.1063/1.4938581, URL https://doi.org/10.1063/1.4938581.

[7] Li, C., Verduzco, J. C., Lee, B. H., Appleton, R. J., and Strachan, A., ''Mapping microstructure to shock-induced temperature fields using deep learning,'' *npj Computational Materials*, Vol. 9, No. 1, 2023, p. 178.

[8] Gonthier, K. A., ''Modeling and analysis of reactive compaction for granular energetic solids,'' *Combustion Science and Technology*, Vol. 175, No. 9, 2003, pp. 1679–1709.

[9] Powers, J. M., ''Two-phase viscous modeling of compaction of granular materials,'' *Physics of Fluids*, Vol. 16, No. 8, 2004, pp. 2975–2990.

[10] Hamilton, B. W., and Germann, T. C., ''Thermal Gradient Effects on Local Hotspot Ignition in 1,3,5,7-Tetranitro-1,3,5,7-tetrazocane (HMX),'' *Propellants, Explosives, Pyrotechnics*, Vol. 50, No. 9, 2025, p. e70011. https://doi.org/https://doi.org/10.1002/prep.70011, URL https://onlinelibrary.wiley.com/doi/abs/10.1002/prep.70011.

[11] Herrin, J., Okafor, C., Picu, C., Sewell, T., Brennan, J., Larentzos, J., and Udaykumar, H., ''Length scale and grid resolution effects in the simulation of shear and energy localization during pore collapse in shocked energetic crystals,'' *Journal of Applied Physics*, Vol. 137, 2025. https://doi.org/10.1063/5.0257287.

[12] Henson, B. F., and Smilowitz, L. B., *The Chemical Kinetics of Solid Thermal Explosions*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2010, pp. 45–128. https://doi.org/10.1007/978-3-540-87953-4_3, URL https://doi.org/10.1007/978-3-540-87953-4_3.

[13] Parepalli, P., Nguyen, Y., Sen, O., Hardin, D., Molek, C., Welle, E., and Udaykumar, H., ''Multi-scale modeling of shock initiation of a pressed energetic material III: Effect of Arrhenius chemical kinetic rates on macro-scale shock sensitivity,'' *Journal of Applied Physics*, Vol. 135, 2024. https://doi.org/10.1063/5.0187735.

[14] Zaug, J. M., Young, C. E., Long, G. T., Maienschein, J. L., Glascoe, E. A., Hansen, D. W., Wardell, J. F., Black, C. K., and Sykora, G. B., ''Deflagration rates of secondary explosives under static MPa—GPa pressure,'' *AIP Conference Proceedings*, Vol. 1195, American Institute of Physics, 2009, pp. 420–423.

[15] Saunier, J., Chinnayya, A., Kaeshammer, E., Reynaud, M., and Genetier, M., ''Mesoscale modeling of the Shock-to-Detonation Transition of pressed-HMX based on a surface regression model,'' *Propellants, Explosives, Pyrotechnics*, Vol. 49, No. 9, 2024, p. e202400125.

[16] Hill, L. G., Zimmermann, B., and Nichols III, A. L., ''ON THE BURN TOPOLOGY OF HOT-SPOT-INITIATED REACTIONS,'' *AIP Conference Proceedings*, Vol. 1195, American Institute of Physics, 2009, pp. 432–435.

[17] Hill, L. G., ''The shock-triggered statistical hot spot model,'' *AIP Conference Proceedings*, Vol. 1426, American Institute of Physics, 2012, pp. 307–310.

[18] Nguyen, Y. T., Okafor, C., Zhao, P., Sen, O., Picu, C. R., Sewell, T., and Udaykumar, H., ''Atomistics-consistent continuum models and pore-collapse-generated hotspot temperatures in energetic crystals I: beta-1, 3, 5, 7-tetranitro-1, 3, 5, 7-tetrazocane (beta-HMX),'' 2024.

[19] Perera, D., ''Molecular dynamics simulations of shock-induced pore collapse in single crystal [beta]-HMX,'' Ph.D. thesis, University of Missouri–Columbia, 2023.

[20] Okafor, C. O., ''High Fidelity Numerical Simulations and Atomistics-Informed Continuum Mesoscale Modeling of Heterogenous Energetic Materials,'' Ph.D. thesis, 2024. URL https://www.proquest.com/dissertations-theses/high-fidelity-numerical-simulations-atomistics/docview/3162737753/se-2, copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2025-11-06.

[21] Udaykumar, H., Nguyen, Y. T., and Kumar Seshadri, P., ''Physically evocative meso-informed burn model: Topology of evolving hotspot fields,'' *Journal of propulsion and power*, Vol. 38, No. 6, 2022, pp. 920–934.

[22] Kapahi, A., and Udaykumar, H., ''Dynamics of void collapse in shocked energetic materials: physics of void–void interactions,'' *Shock Waves*, Vol. 23, No. 6, 2013, pp. 537–558.