

Article

Nozzle Erosion Reconstruction Model for Data Analysis in Rocket Engines and Correlation with Chamber Pressure

Ryan J. Thibaudeau  and Stephen A. Whitmore 

Department of Mechanical and Aerospace Engineering, Utah State University, Logan, UT 84322, USA;
stephen.whitmore@usu.edu

* Correspondence: ryan.thibaudeau@usu.edu

Abstract

Graphite nozzles remain the dominant choice for small hybrid and solid rocket motors operating on laboratory and university budgets, owing to their low cost, ease of machining, and rapid turnaround during iterative design campaigns. These same programs, however, must contend with the fact that graphite erodes through coupled thermochemical and mechanical mechanisms when exposed to the oxidizing species generated by high-energy propellant combustion, and the resulting throat-area growth fundamentally alters the time histories of chamber pressure, thrust, and delivered specific impulse. This paper presents a nozzle-erosion reconstruction model that extracts the time-resolved throat area from coupled thrust and chamber-pressure measurements using the thrust coefficient relationship, scales the reconstructed area history against pre- and post-test throat measurements, identifies the onset and rate of erosion, and accounts for variable sensor lag between the thrust-stand and pressure-transducer signal chains. The model is exercised on two complementary sets of laboratory-scale GOX/ABS hybrid hot-fire data that together span roughly two orders of magnitude in total throat-area change and peak chamber pressures from 0.5 to 3.4 MPa: a controlled three-operating-point campaign conducted in support of the NASA Plume-Surface Interaction (PSI) program, and a set of higher-pressure firings from the laboratory development series in which the technique was matured. Reconstructed erosion-onset times, erosion rates, and total throat-diameter change are reported for each firing, the reconstruction accuracy is characterized as a function of erosion magnitude. A correlation of graphite erosion with chamber pressure is examined across the combined envelope. The results demonstrate the robustness of the reconstruction technique and provide a reusable framework for post-test reconstruction of transient nozzle geometry in rocket-engine ground testing.

Keywords: graphite nozzle erosion; hybrid rocket propulsion; thrust coefficient; data reconstruction; chamber pressure; ground test analysis; NASA MSFC



Academic Editors: Giuseppe Gallo,
Harunori Nagata and Landon Kamps

Received: 28 May 2026

Revised: 22 June 2026

Accepted: 23 June 2026

Published: 25 June 2026

Copyright: © 2026 by the authors.
Licensee MDPI, Basel, Switzerland.
This article is an open access article
distributed under the terms and
conditions of the [Creative Commons
Attribution \(CC BY\) license](https://creativecommons.org/licenses/by/4.0/).

1. Introduction

1.1. Graphite as a Nozzle Material for Small-Motor Development

Graphite nozzles have long been the default choice for small-scale rocket motor development, and they remain so for most university and laboratory programs. Graphite is inexpensive, available in suitable grades from industrial suppliers, and easy to machine to tight tolerances on conventional lathes and mills. Its high sublimation temperature and modest density give it respectable thermal survivability for short-duration firings, and it

tolerates the steep thermal gradients of transient hot-fire exposure without the through-thickness cracking that often plagues monolithic ceramic inserts [1,2]. For university laboratories, small companies, and research groups iterating on experimental propulsion hardware, those properties translate directly into a low barrier to entry: a new nozzle geometry can be drafted in the morning, machined and inspected in the afternoon, and fired the next day. That cadence is essential to the tight, hypothesis-driven test campaigns typical of early-stage propulsion research. More exotic nozzle materials such as refractory-metal alloys, carbon-carbon composites, and ceramic matrix composites cannot reasonably support that cadence because of cost, lead time, and machining complexity [1,3]. As a result, a substantial fraction of the small-motor data generated in the research community comes from graphite-nozzle motors, and the data-reduction challenges that go with graphite are effectively unavoidable.

Those challenges begin with the chemistry of graphite erosion itself, considered next.

1.2. Thermochemical Erosion Mechanisms

The central challenge is that graphite is not chemically inert at the temperatures and gas compositions produced by modern rocket propellants. When exposed to the combustion products of high-energy propellant systems, particularly those containing oxygen, fluorine, or hydrogen-rich species, the graphite surface is consumed by heterogeneous chemical reactions. The dominant attackers are water vapor, carbon dioxide, and hydroxyl radicals, with smaller contributions from atomic and molecular oxygen [2,4,5]. The classical reactions are the water–gas reaction, $C + H_2O \rightarrow CO + H_2$, and the Boudouard reaction, $C + CO_2 \rightarrow 2CO$, both of which are strongly endothermic and strongly activated [1,2]. The rate at which they consume the throat surface depends on the local boundary-layer chemistry, the wall temperature, and how quickly oxidizing species are transported from the freestream to the wall. In the low-pressure, mass-transfer-limited regime, erosion is controlled by diffusion through the thermal boundary layer; in the high-pressure, kinetically-limited regime, it is set by the Arrhenius behavior of the surface reactions. Both mechanisms accelerate with chamber pressure, and the transition between them is gradual rather than sharp. That gradual transition is one of several reasons that first-principles erosion-rate prediction for an arbitrary propellant-motor combination remains unreliable [3,5]. Under these conditions the throat geometry drifts during a single firing: throat diameter grows, expansion ratio drops, chamber pressure at fixed mass flow sags, and delivered thrust departs from the design intent in ways that depend on firing duration and operating point.

Limiting that geometry drift, rather than merely tolerating it, has been the aim of substantial materials and propellant research, reviewed next.

1.3. Mitigation Approaches and Their Limits

Considerable effort has gone into slowing or arresting this erosion so that graphite can be used at higher pressures and longer durations without unacceptable geometry change. Refractory coatings such as tungsten, rhenium, hafnium carbide, and zirconium carbide have been shown to reduce erosion rates substantially when applied via chemical vapor deposition, plasma spray, or pack cementation. The trade-off is that the coatings introduce new spallation and coating-substrate mismatch failure modes [1,5]. Pyrolytic graphite, with its highly anisotropic basal-plane structure, presents a much smaller through-thickness reaction cross-section and has been used in tactical motors for decades [4]. More recent work has investigated three-dimensional carbon-carbon composites, in which the reinforcing fiber architecture resists the surface recession that bulk graphite suffers; cost and fabrication lead time generally place such throats out of reach for laboratory testing [3]. On the propellant-chemistry side, work has examined the role of hydrogen-rich fuels, halogenated additives,

and low-flame-temperature oxidizers in lowering the concentration of the most aggressive oxidizing species at the throat wall [2]. All of these approaches deserve continued attention. None of them change the basic reality, though: for the overwhelming majority of small-scale hot-fire testing today, the motor is fired with a plain-graphite nozzle, the throat erodes during the firing, and any analysis of the measured performance data has to contend with a time-varying throat area.

Interpreting that data quantitatively therefore requires reconstructing the throat-area history, the motivation taken up next.

1.4. Motivation for Transient Throat Reconstruction

Given that reality, reconstructing the transient throat geometry from measured ground-test data is not a niche capability. It is a prerequisite for the quantitative interpretation of essentially every graphite-nozzle firing. Under the conventional fixed-throat-area assumption, the combustion efficiency, characteristic velocity, and specific impulse extracted from a test are all biased in systematic and non-trivial ways. A motor with a significantly eroded throat registers a lower measured chamber pressure than a non-eroding equivalent, and under the fixed-area assumption that drop is mistakenly attributed to reduced combustion efficiency, off-design oxidizer-to-fuel ratio, or a deficiency in theoretical performance modeling. With a correctly reconstructed throat-area history, the measured chamber pressure and thrust can be reinterpreted in terms of their true thermodynamic causes, and genuine performance anomalies can be separated from geometric artifacts of nozzle recession.

Beyond performance bookkeeping, the erosion history itself is a first-class observable. The reconstructed throat-diameter-versus-time curve encodes the instantaneous erosion rate, which is the quantity most directly relevant to design margins on burn duration, throat wall thickness, and nozzle-material selection. Correlating that reconstructed rate with the co-measured chamber pressure, mixture ratio, and combustion temperature gives an empirical basis for erosion-rate scaling laws in regimes where first-principles models are known to be unreliable.

How prior work has attempted to recover that history, and the limitations that motivate the approach taken here, are reviewed in the following subsection.

1.5. Prior Reconstruction Approaches

Earlier reconstruction efforts in the literature have taken one of two simpler approaches. The first assumes a constant erosion rate averaged over the firing, which masks the transient behavior. The second injects pre-test erosion-rate estimates into a forward performance model, which conflates model error with measurement [1,5]. More rigorous techniques instead invert the measured performance data directly for the throat-area history, and two such coupled-ballistics closures are in current use.

The first closure is built on the thrust coefficient, conventionally defined as

$$C_F = \frac{F}{P_c A_t}, \quad (1)$$

where F is the axial thrust, P_c is the stagnation chamber pressure, and A_t is the nozzle throat area [6]. In a non-eroding nozzle operating at steady state, Equation (1) is used as a forward relationship: C_F is computed from measured F and P_c and compared to the ideal thrust coefficient predicted from the propellant thermochemistry. When the throat is allowed to vary in time, the same relationship is rearranged into an inverse problem for the instantaneous throat area,

$$A_t(t) = \frac{F(t)}{P_c(t) C_F(t)}. \quad (2)$$

The conventional treatment of Equation (2) evaluates $C_F(t)$ from the ideal compressible-flow expression of Sutton and Biblarz,

$$C_{F,\text{id}} = \Gamma \sqrt{\frac{2\gamma}{\gamma-1} \left[1 - \left(\frac{P_e}{P_c} \right)^{\frac{\gamma-1}{\gamma}} \right]} + \left(\frac{P_e - P_\infty}{P_c} \right) \epsilon, \quad (3)$$

where γ is the specific-heat ratio at the throat, P_e is the static pressure at the nozzle exit, P_∞ is the ambient back pressure recorded during the firing, and $\epsilon = A_e/A_t$ is the instantaneous expansion ratio referenced to the fixed exit area A_e . The Vandekerckhove function is

$$\Gamma(\gamma) = \sqrt{\gamma} \left(\frac{2}{\gamma+1} \right)^{\frac{\gamma+1}{2(\gamma-1)}}. \quad (4)$$

The exit-to-stagnation pressure ratio follows from the isentropic exit Mach number M_e , which satisfies the area-Mach relation

$$\epsilon = \frac{1}{M_e} \left[\frac{2}{\gamma+1} \left(1 + \frac{\gamma-1}{2} M_e^2 \right) \right]^{\frac{\gamma+1}{2(\gamma-1)}}, \quad (5)$$

and

$$\frac{P_e}{P_c} = \left(1 + \frac{\gamma-1}{2} M_e^2 \right)^{-\frac{\gamma}{\gamma-1}}. \quad (6)$$

Equations (2)–(6) cannot be evaluated as a strict forward chain. The expansion ratio $\epsilon(t) = A_e/A_t(t)$ depends on the unknown throat area, so $C_{F,\text{id}}$ depends on A_t through the area-Mach relation. The specific-heat ratio γ at the throat is itself a function of the running chamber pressure and the instantaneous oxidizer-to-fuel ratio O/F . In a hybrid configuration, the in-test O/F depends on the fuel mass flow generated by grain regression, which depends on the oxidizer mass flux through the port and therefore on the throat area through its effect on chamber pressure. Closing the system requires inserting Equation (2) into a coupled solver that resolves, at each time sample, the throat area, the specific-heat ratio (typically from a NASA Chemical Equilibrium with Applications lookup [7,8]), the exit pressure ratio, the regression-rate-derived fuel mass flow, the resulting O/F , the characteristic-velocity efficiency η_{c^*} , and the thrust-coefficient efficiency η_{C_F} that lumps boundary-layer, divergence, and two-phase losses. Convergence of this coupled system is iterative even without erosion, with three to five passes typically required at each sample for a non-eroding nozzle and substantially more over intervals on which the throat is actively receding.

A second conventional closure, and the one developed most directly for hybrid nozzle erosion, is the nozzle-throat reconstruction technique (NTRT) of Kamps et al. [9,10]. Rather than routing the inverse through the thrust coefficient, the NTRT closes it through the characteristic-velocity identity,

$$c^* = \frac{P_c A_t}{\dot{m}}, \quad (7)$$

in which the total mass flow $\dot{m} = \dot{m}_{\text{ox}} + \dot{m}_f$ combines the directly metered oxidizer flow with the fuel flow implied by the instantaneous mixture ratio O/F . The theoretical characteristic velocity and specific-heat ratio are again drawn from a chemical-equilibrium lookup as functions of chamber pressure and mixture ratio, $c^* = f(P_c, O/F)$ and $\gamma = g(P_c, O/F)$, and the measured thrust is tied to one-dimensional nozzle-flow theory through a thrust-correction factor. The system is closed by requiring that the fuel mass obtained by integrating the reconstructed fuel flow over the firing match the fuel mass weighed before and after the test, with the final throat radius—or, when it cannot be measured, a wall-pressure

tap near the nozzle exit—fixing the absolute geometry. The NTRT thereby reconstructs the throat-radius history and the mixture-ratio history simultaneously. The technique has been scaled across 30–2000 N-thrust-class motors and used to assemble an empirical graphite erosion-rate model and a chemical-kinetic erosion-onset criterion from several hundred reconstructed data points [9,10], and its reconstructed data have since served to validate first-principles computational-fluid-dynamics predictions of throat ablation [11]; those numerical results reinforce a conclusion central to the present study, that nozzle erosion rate and throat wall temperature peak under oxidizer-rich rather than stoichiometric conditions and that the chamber-pressure dependence of the erosion rate is itself a function of mixture ratio. Its principal merit is precisely this physical grounding: every reconstructed quantity is anchored to a measured propellant mass or pressure, and the method recovers the full O/F history as a useful by-product of the erosion solve. The thrust-coefficient closure of Equations (1)–(6) and the characteristic-velocity closure of Equation (7) are both in current use, and each has genuine merit; they differ chiefly in which idealized performance relation is inverted and in whether the absolute geometry is fixed by an expansion-ratio model or by a fuel-mass-consumption constraint.

Although the two closures differ in this way, they share the same structural feature: each reconstructs the throat geometry in parallel with the in-test internal ballistics, co-solving the eroding throat area and the O/F history within a single coupled loop. The structural drawback of that coupling for erosion-focused analysis is that nozzle erosion is one of the largest dynamic perturbations in the very loop being used to infer it. Each coupled unknown carries its own uncertainty, and the fixed-point character of the solver propagates those uncertainties into the reconstructed throat area in ways that are difficult to bound analytically. The dominant contributors are familiar and common to both closures: the CEA-derived γ at the throat depends on an inferred O/F that is itself uncertain; the regression-rate model or fuel-mass-consumption constraint used to generate the fuel mass flow is a correlation rather than a direct measurement; and a lumped efficiency term—the thrust-coefficient efficiency η_{C_F} in the first closure and the characteristic-velocity efficiency η_{c^*} in the NTRT—is most often calibrated against pre-test bench-top conditions that do not capture the dynamic effects of an actively eroding throat. The reconstructed throat-area history is therefore biased by every term the coupled solver assumes or co-solves, and the resulting uncertainty band, when carried through to the erosion rate, is large enough to obscure the chamber-pressure correlation that is the central observable of the present study.

The present work shares this measurement-driven, inverse philosophy but departs from both closures by decoupling the throat-area inverse from the internal-ballistics loop, so that the uncertainty of an inferred O/F history is never propagated into the erosion measurement it is meant to support. It removes the inner thermochemical loop entirely, at the cost of accepting a single empirically calibrated thrust-coefficient scaling determined from direct throat-geometry measurements at two times during each firing. That decoupled reconstruction is developed in Section 2.

1.6. Scope and Contribution of the Present Work

This work develops such an approach and exercises it on two complementary hot-fire datasets to demonstrate its robustness across a wide range of operating conditions. The reconstruction model has been in development at the Utah State University Propulsion Research Laboratory (USU PRL) for several years, refined iteratively against an extensive series of laboratory-scale hybrid firings. The present paper draws on a small high-pressure subset of those PRL firings, together with a statistically meaningful sample of firings from the NASA Marshall Space Flight Center Component Development Area (CDA) campaign conducted in support of the Plume-Surface Interaction (PSI) program. The two datasets

bracket each other in chamber pressure and total throat-area change: the CDA campaign provides a controlled, three-operating-point test matrix at low to moderate chamber pressure, and the PRL firings extend the validation envelope to peak chamber pressures of 2.3–3.4 MPa and total throat-diameter changes exceeding 3 mm.

The reconstruction model itself uses the thrust-coefficient relationship of Equation (1), rearranged to solve for throat area as a function of the measured thrust and chamber-pressure histories. Pre- and post-test throat measurements provide two-point anchors that absorb systematic calibration biases in the load cell and pressure transducer. An erosion-onset criterion applied to the reconstructed throat-area history identifies the point at which the throat begins to recede from its pre-test geometry. A regression-based erosion-rate estimate is then extracted from the post-onset portion of the curve. A sensor-lag-aware masking interval handles the fact that the thrust-stand and pressure-transducer signal chains have different effective bandwidths, and that the lag between them is not constant across tests. The CDA dataset's three-operating-point structure isolates chamber-pressure effects on the measured erosion rate, and the PRL dataset extends those measurements above the CDA envelope. The results reported here therefore contribute both a reusable data-reduction methodology for graphite-nozzle ground testing and a quantitative characterization of the chamber-pressure correlation of graphite erosion across the GOX/ABS hybrid envelope exercised by the two datasets.

2. Materials and Methods

The conventional coupled-ballistics closures reviewed in Section 1.5 recover the eroding throat area only by co-solving the in-test internal ballistics, which propagates thermochemical and regression-rate uncertainty into the very quantity being measured. The reconstruction developed here deliberately avoids that coupling. This section presents it in a sequence of self-contained steps. Section 2.1 develops the decoupled throat-area reconstruction, which treats the thrust coefficient as an empirical calibration absorbed entirely by direct pre- and post-test throat measurements; Section 2.2 defines the erosion-rate extraction applied to the corrected throat-area history; and Section 2.3 describes the motor hardware, instrumentation, and test matrices of the NASA MSFC Plume-Surface Interaction (PSI) campaign and USU PRL series that supply the datasets exercised in Section 3. All data is analyzed using MATLAB.

2.1. Decoupled Reconstruction Algorithm

The reconstruction used in the present work begins from the unscaled inverse

$$\hat{A}_t(t) = \frac{F(t)}{P_c(t)}, \quad (8)$$

in which the thrust coefficient has been omitted entirely. Equation (8) differs from the true throat area by the factor $C_F(t)$ alone. Because C_F is a slowly-varying function of expansion ratio at small-motor operating conditions and varies by only a few percent over even a heavily-eroded firing, the time history $\hat{A}_t(t)$ has the same shape as $A_t(t)$ to within a near-constant multiplicative offset. Shape preservation is the key property the algorithm exploits: erosion-onset detection, lag-spike identification, and the qualitative features of the throat-area history are all unaffected by the absolute level and can be performed on $\hat{A}_t(t)$ directly. The absolute level is fixed at the end of the procedure by a two-anchor affine map onto the directly measured pre- and post-test throat areas. That map absorbs both the constant component of C_F and the slow drift of C_F with expansion ratio over the burn, without ever entering an internal-ballistics loop.

The reconstruction proceeds in five steps: array initialization with hard-coded boundary segments, removal of the ignition lag transient, detection of the erosion-onset time on the unscaled signal, two-anchor affine scaling to the directly measured throat areas, and erosion-rate extraction from the scaled signal. Each step is described in turn.

2.1.1. Array Initialization

Let t_{ro} denote the run-valve-open time and t_{eo} the engine-off time, both recorded directly by the data-acquisition system through dedicated discrete channels. The reconstructed throat-area history is initialized over the full data-acquisition window as a piecewise array,

$$\hat{A}_t(t) = \begin{cases} A_{t,pre} & t < t_{ro}, \\ F(t)/P_c(t) & t_{ro} \leq t \leq t_{eo}, \\ A_{t,post} & t > t_{eo}, \end{cases} \quad (9)$$

where $A_{t,pre}$ and $A_{t,post}$ are the throat areas computed from pin-gauge measurements of the throat before and after the firing. Hard-coding the pre- and post-firing segments to the directly measured throat geometry forces the reconstruction to be physically consistent at the boundaries by construction, and isolates the inverse calculation to the active firing window between t_{ro} and t_{eo} .

The thrust and chamber-pressure records that enter the F/P_c inverse are passed through a short zero-phase moving-average filter (window length on the order of 50 ms) before the ratio is taken, to suppress the residual broadband noise that the inverse would otherwise amplify when P_c is small. The unsmoothed records are retained for rise-time and tail-off-initiation detection, both of which depend on the original transient characteristics of the sensors.

2.1.2. Sensor-Lag Transient Removal

The thrust-stand load cell (LC) and the chamber-pressure transducer (PT) have different mechanical and electrical bandwidths. In the configuration used here, the load cell responds more rapidly to the ignition transient than the pressure transducer, so during the rising edge of the firing the measured thrust climbs to its quasi-steady value before the measured chamber pressure does. The unscaled inverse $F(t)/P_c(t)$ therefore exhibits a large positive spike across the ignition transient, with the numerator at or near its quasi-steady level while the denominator is still rising, before relaxing to its plateau value once both signals have stabilized. The spike is an instrumentation artifact and is masked from the rest of the analysis.

The masking interval is set by the slower of the two sensor rise times plus a brief settling buffer. Let τ_F denote the load-cell rise time and τ_{P_c} the pressure-transducer rise time, each defined as the elapsed time from t_{ro} to the first sample at which the corresponding signal first reaches a fraction f_{rise} (taken here as 0.95) of its plateau value, with the plateau value computed as the median of the signal over a short reference window centered on the nominal mid-burn time. The trustworthy-interval boundary is then

$$t_{rise} = t_{ro} + \max(\tau_F, \tau_{P_c}) + \Delta t_{settle}, \quad (10)$$

where Δt_{settle} is a settling buffer (taken here as 100 ms) appended after the rise event so that residual relaxation past the 95%-of-plateau threshold has decayed before the affine fit is applied. The reconstruction over $[t_{ro}, t_{rise})$ is hard-coded to the directly measured pre-test throat area $A_{t,pre}$. Hard-coding the masked samples, rather than filling them by linear interpolation in the unscaled domain, eliminates a residual artifact in which the affine

scaling subsequently applied in Section 2.1.4 (with $\alpha \neq 1$ and $\beta \neq 0$) would otherwise project the interpolated samples to non-physical values at the leading edge of the firing record. The settling buffer suppresses a related artifact in which \hat{A}_t has not yet relaxed to its plateau value at the rise-time threshold and is consequently anchored against an inflated value, producing a sharp upward jump in the scaled reconstruction at t_{rise} . Defining the masking interval from the slower of the two sensors, rather than from the apparent location of the spike itself, ensures that the reconstruction is not trusted until both signal chains have reached their respective steady states, regardless of which sensor was the rate-limiter on a given firing.

2.1.3. Erosion-Onset Detection by Smoothed Cumulative Differentiation

The erosion-onset time t^* is identified by piecewise-linear changepoint regression on the smoothed unscaled throat-area history $\tilde{A}_t(t)$. The pre-erosion plateau and the actively-eroding interval are modeled as separate segments connected at the changepoint, and t^* is the breakpoint that minimizes the joint residual sum of squares of the two-segment model. Two competing right-segment forms are evaluated in parallel: a linear segment (constant erosion rate) and a parabolic segment (linearly accelerating erosion rate). The form with the smaller residual is reported. The piecewise formulation captures erosion onset directly as a change in slope rather than indirectly through a deviation from a fixed plateau mean, and it is naturally robust against the slow C_F drift that contaminates fixed-threshold detection on $\tilde{A}_t(t)$ for low-erosion firings.

The masked, unscaled throat-area history is first passed through a zero-phase low-pass filter to produce a smoothed signal $\tilde{A}_t(t)$. The filter used here is a forward-backward fourth-order Butterworth implementation, with cutoff f_c chosen at one decade below the Nyquist frequency of the data-acquisition system. The two-segment model is

$$\tilde{A}_t(t) \approx \begin{cases} a + b_1(t - t^*) & t \leq t^* \\ a + b_2(t - t^*) & t > t^* \quad (\text{linear-right model}) \\ a + c_2(t - t^*)^2 & t > t^* \quad (\text{parabolic-right model}) \end{cases} \quad (11)$$

where the intercept a is shared between segments to enforce continuity at the changepoint, b_1 is the pre-erosion plateau slope (which absorbs any slow C_F drift), and the right-segment shape parameter is either a slope b_2 or a curvature coefficient c_2 depending on which model is selected. For each candidate t^* in the discretized search range $[t_{\text{rise}} + \Delta t_{\text{min}}, t_{\text{toi}} - \Delta t_{\text{min}}]$ (with $\Delta t_{\text{min}} = 500$ ms to avoid degenerate fits at the segment ends), Equation (11) is fit by ordinary least squares for both right-segment forms; the candidate t^* and right-segment form with the smallest joint residual sum of squares are taken as the working solution.

The changepoint is accepted as a real onset when the relative reduction in residual sum of squares over the null single-line fit exceeds a threshold,

$$\frac{\text{RSS}_{\text{null}} - \text{RSS}_{\text{cp}}}{\text{RSS}_{\text{null}}} > \rho_{\text{min}}, \quad (12)$$

where RSS_{null} is the residual sum of squares of a single linear fit to $\tilde{A}_t(t)$ over the trustworthy interval $[t_{\text{rise}}, t_{\text{toi}}]$ and RSS_{cp} is the residual sum of squares of the changepoint model. The relative-reduction form normalizes against the data variance and gives a unitless acceptance criterion that does not require an explicit noise-floor estimate. For the present analysis $\rho_{\text{min}} = 0.10$. Firings that fail Equation (12) are reported as “no onset detected,” and the changepoint search returns no t^* for those records; the affine scaling in Section 2.1.4 then proceeds with the early-burn fallback for the pre-anchor.

A consequence of the shape-preserving property of the unscaled inverse Equation (8) is that the changepoint location identified on $\hat{A}_t(t)$ is identical, to within the discretization of the candidate search, to the location that would be identified on the fully scaled $A_t(t)$, because the affine map preserves the shape of the underlying signal. Performing the detection on the unscaled signal therefore introduces no additional error, and the detection step precedes the scaling step in the algorithm so that t^* can be used to define the upper edge of the pre-erosion plateau over which the first scaling anchor is computed. The fit-coefficient values $\{a, b_1, b_2, c_2\}$ obtained from the unscaled regression are projected into the scaled domain by the affine map of Section 2.1.4 (with $a_{\text{scaled}} = \alpha a + \beta$, $b_{1,\text{scaled}} = \alpha b_1$, and similarly for b_2 and c_2) and are reported as the final analytical erosion model in Section 2.2.

Two constraints are imposed on the candidate-search and model-selection steps. First, the candidate range $[t_{\text{rise}} + \Delta t_{\text{min}}, t_{\text{toi}} - \Delta t_{\text{min}}]$ uses $\Delta t_{\text{min}} = 1.5$ s to keep the search from latching onto small noise features near either segment boundary; with the campaign-typical 6-s burn this leaves more than 3 s of valid candidates for t^* . Second, the linear right-segment model is preferred over the parabolic model unless the parabolic fit reduces the residual sum of squares by at least 10% relative to the linear fit. The structural form of the parabolic model imposes zero slope at t^* and a slope that grows linearly with time, which on a roughly-linear erosion ramp creates a bias toward placing t^* early so that the curvature has room to “ramp up” to the observed late-burn slope. Requiring a clear RSS-reduction margin before the parabolic model is selected eliminates this bias for genuinely-linear erosion histories while preserving the parabolic fit on firings where the curvature is real and dominant. The selection rule is therefore

$$\text{model} = \begin{cases} \text{parabolic} & \text{if } \text{RSS}_{\text{quad}} < 0.90 \text{RSS}_{\text{lin}} \\ \text{linear} & \text{otherwise.} \end{cases} \quad (13)$$

The detection is gated by a per-firing classification of whether erosion occurred. The default classification compares the directly measured throat-area change $\Delta A_{t,\text{meas}} = A_{t,\text{post}} - A_{t,\text{pre}}$ against a detection threshold $\Delta A_{\text{thr}} = 0.5$ mm² (about 0.04 mm in throat-diameter change at the nominal 8.45 mm throat). The threshold sits well above the pin-gauge precision of ± 0.01 mm and is set instead by the noise floor of the F/P_c reconstruction. When the campaign-record metadata distinguishes “Slightly” eroded firings from “No” eroded firings explicitly, as is the case for the PSI test campaign records used here, the metadata classification overrides the magnitude gate. This is necessary because $|\Delta A_{t,\text{meas}}|$ for “Slightly” eroded firings can fall in the same numerical range as for “No” eroded firings, and a magnitude-only gate cannot separate them. Firings classified as “No” eroded skip the changepoint regression entirely; any noise-induced “kink” the regression might find on those records would not correspond to a physical onset. Firings classified as “Yes” or “Slightly” eroded run a tiered detection sequence: the regression is first attempted at the standard $\rho_{\text{min}} = 0.10$ acceptance threshold; if that fails, a relaxed retry is performed at $\rho_{\text{min}} = 0.05$, on the rationale that the boundary measurements have already established that real erosion occurred and the algorithm should be more sensitive to weak signals; if the relaxed attempt also fails, a forced-onset fallback fixes $t^* = t_{\text{rise}}$ and fits a single post-onset segment over the entire trustworthy interval $[t_{\text{rise}}, t_{\text{toi}}]$ using the same linear-versus-parabolic model-selection rule of Equation (13).

The forced-onset fallback always produces a usable rate model on firings with measured erosion, but the resulting t^* should be treated as a placeholder rather than a confident detection, because the fallback overestimates the duration of erosion if the real onset was later than t_{rise} . The fit-info structure carries an explicit `onset_at_rise` flag that identifies these firings so that downstream rate-comparison analyses can treat their t^* values appro-

propriately. The tier of the accepted fit (standard, relaxed, or forced-onset fallback) is reported in the per-test diagnostic record alongside the fit equation.

2.1.4. Two-Anchor Affine Scaling

The unscaled reconstruction $\hat{A}_t(t)$ from Equation (8) differs from the true throat area $A_t(t)$ by the time-varying factor $C_F(t)$. Across a firing, C_F depends on the expansion ratio through Equations (3)–(6), and the expansion ratio decreases as the throat erodes. In the small-motor expansion-ratio range of interest here, this drift is on the order of a few percent of the plateau value of C_F between the start and end of an actively-eroding firing, and a single multiplicative correction applied uniformly across the burn is therefore insufficient to anchor both ends of the reconstruction to the directly measured pre- and post-test throat areas simultaneously. Empirically, a single-scale correction tuned to match the pre-test throat measurement leaves a systematic residual at the post-test boundary, with the magnitude of the residual scaling with the total amount of erosion. The reconstruction therefore uses an affine correction with two free parameters,

$$A_t(t) = \alpha \hat{A}_t(t) + \beta, \quad (14)$$

where α and β are determined by enforcing the corrected reconstruction to match the directly measured throat areas at two reference times within the firing.

The two reference windows are placed flush against the boundaries of the trustworthy interval, rather than insetted away from them, so that the affine map produces $A_t(t_{\text{rise}}) \approx A_{t,\text{pre}}$ and $A_t(t_{\text{toi}}) \approx A_{t,\text{post}}$ by construction. Insetted anchors were tried first and produced a systematic late-burn overshoot in which the unscaled signal continued to climb between the right edge of the inset post-anchor window and t_{toi} , leaving the affine output above $A_{t,\text{post}}$ over the last fraction of a second of the trustworthy interval and forcing a visible step at t_{toi} where the hard-coded post-tail-off segment took over. Flush anchors eliminate this overshoot.

The pre-anchor window is selected as follows. When erosion onset is detected and the resulting pre-erosion plateau is wide enough to support a stable median, the pre-anchor is the median of \hat{A}_t over the plateau interval, with the final five percent of the plateau dropped to avoid pre-onset curvature:

$$\bar{A}_{t,\text{pre}}^{\text{recon}} = \text{median}\{\hat{A}_t(t) : t \in [t_{\text{rise}}, t^* - \delta t_{\text{in}}]\}. \quad (15)$$

When the plateau is too narrow to support a stable median, or when no onset is detected (gradual erosion or no measurable change), the pre-anchor falls back to a fixed-duration early-burn window $[t_{\text{rise}}, t_{\text{rise}} + \Delta t_{\text{anchor}}]$, where Δt_{anchor} is the larger of one-tenth of the firing duration and 300 ms. The median is used in both cases in preference to the arithmetic mean to suppress sensitivity to short-duration excursions within the window.

The post-anchor window is the late-burn quasi-steady interval placed flush against the tail-off-initiation time. Let t_{toi} denote the tail-off-initiation time, defined as the time at which the cutoff transient first contaminates the F/P_c inverse. Operationally, t_{toi} is detected as the earlier of two independent threshold-based detections, one on the chamber-pressure derivative and one on the thrust derivative. For each signal $X \in \{P_c, F\}$, the detection time is the first sample after the burn midpoint at which $\dot{X}(t)$ drops below the threshold $\dot{X}_{\text{thresh}} = \dot{X}_{\text{plat}} - k_{\text{grad}} \sigma_{\dot{X},\text{plat}}$ and remains below threshold for a 30-ms sustain interval, where \dot{X}_{plat} and $\sigma_{\dot{X},\text{plat}}$ are the plateau median and MAD-based standard deviation of \dot{X} over the early-burn segment, and $k_{\text{grad}} = 6$ is a confidence multiplier. The dual-sensor formulation handles two distinct cutoff patterns observed in the campaign: in most firings the chamber pressure decays slightly before the load-cell signal, but in a minority the load

cell reaches its decline threshold first and produces a downward F/P_c excursion in the late-burn reconstruction if t_{toi} is set from \dot{P}_c alone. Taking $t_{\text{toi}} = \min(t_{\text{toi},P_c}, t_{\text{toi},F})$ excludes whichever sensor leads. The earlier formulation placed t_{toi} a fixed margin upstream of the most-negative \dot{P}_c event and produced nearly identical t_{toi} values across all firings because the campaign's cutoff sequence is geometrically similar across operating points; the threshold-based detector instead resolves per-firing variation in the actual onset of decline. A fallback path (most-negative \dot{P}_c minus a 50-ms margin) is retained for firings whose decline does not exceed threshold within the search range. The conservative placement of t_{toi} is necessary because the leading edge of the tail-off transient drives the load cell and the pressure transducer apart in time. The chamber-pressure record begins its decay slightly before the load-cell record does, and the resulting spread is amplified by the F/P_c inverse into a transient excursion above $A_{t,\text{post}}$ followed by a non-physical droop back below $A_{t,\text{post}}$ as the load-cell record catches up.

The post-anchor interval is $[t_{\text{toi}} - \Delta t_{\text{anchor}}, t_{\text{toi}}]$, and the anchor value is taken not as the simple median of \hat{A}_t over the window but as the value of a linear least-squares fit through the window evaluated at t_{toi} ,

$$\bar{A}_{t,\text{post}}^{\text{recon}} = [a_1 t + a_0]_{t=t_{\text{toi}}}, \quad \{a_1, a_0\} = \arg \min_{a_1, a_0} \sum_{t \in [t_{\text{toi}} - \Delta t_{\text{anchor}}, t_{\text{toi}}]} [\hat{A}_t(t) - (a_1 t + a_0)]^2. \quad (16)$$

The fit-at-edge form is used because the unscaled reconstruction is still rising through the late-burn window for actively-eroding firings, so a window median underestimates \hat{A}_t at the right edge of the trustworthy interval. An affine map calibrated against that median produces a residual overshoot of $A_t(t)$ above $A_{t,\text{post}}$ in the last fraction of a second of the firing. Linear extrapolation to t_{toi} recovers the value of the unscaled reconstruction at the boundary itself and eliminates the overshoot. The corresponding noise-floor estimate used in the SNR test of Section 2.1.6 is taken as the standard deviation of the residuals of the same linear fit, which separates true plateau noise from the slow drift induced by ongoing erosion within the window. The anchor value $\bar{A}_{t,\text{post}}^{\text{recon}}$ is matched to the direct measurement $A_{t,\text{post}}$. The affine coefficients are then

$$\alpha = \frac{A_{t,\text{post}} - A_{t,\text{pre}}}{\bar{A}_{t,\text{post}}^{\text{recon}} - \bar{A}_{t,\text{pre}}^{\text{recon}}} \quad (17)$$

$$\beta = A_{t,\text{pre}} - \alpha \bar{A}_{t,\text{pre}}^{\text{recon}}. \quad (18)$$

Equation (14) maps the unscaled reconstruction onto two independently measured anchor points and absorbs into α and β both the constant component of the thrust-coefficient bias and the slowly-varying drift of C_F as the expansion ratio decreases over the burn. The shape of the throat-area trajectory between the anchors is preserved by construction; only its absolute level and slope are adjusted. Because α and β are determined entirely from direct throat-geometry measurements and from the unscaled reconstruction, no element of the propellant thermochemistry or of the internal-ballistics solution enters the scaling step. That is the principal mechanism by which the present algorithm avoids the coupled-solver uncertainty discussed in Section 1.5.

When a changepoint has been detected, the affine map is applied only over the post-onset segment $[t^*, t_{\text{toi}}]$, and the pre-onset segment is hard-coded to the directly measured pre-test throat area: $A_t(t) = A_{t,\text{pre}}$ for all $t \leq t^*$, including the lag-spike interval $[t_{\text{ro}}, t_{\text{rise}}]$ and the pre-erosion plateau $[t_{\text{rise}}, t^*]$. The physical interpretation of t^* as the onset of erosion implies that the throat geometry has not changed between the pre-test measurement and t^* ; hard-coding the pre-onset segment therefore replaces the residual noise on $\tilde{A}_t(t)$ in this interval with the boundary measurement and removes the apparent waviness of the pre-

erosion plateau without distorting the post-onset reconstruction. The post-firing segment ($t_{\text{toi}}, t_{\text{eo}}$) is similarly hard-coded to $A_{t,\text{post}}$. The cutoff transient carries no information about erosion (the throat geometry is fixed once the run valve closes), so excluding it from the F/P_c inverse keeps spurious oscillation out of the reported $A_t(t)$ trajectory.

When no changepoint is detected (low-SNR or no-erosion firings), the affine map is applied over the full trustworthy interval $[t_{\text{rise}}, t_{\text{toi}}]$ as in the fallback formulation, and only the lag-spike interval and the post-tail-off interval are hard-coded.

2.1.5. Final Smoothing and Boundary Continuity

The scaled reconstruction $A_t(t)$ over the trustworthy interval is finally passed through a zero-phase Gaussian smoother with a 500-ms kernel, applied only over $[t_{\text{rise}}, t_{\text{toi}}]$ so that the hard-coded boundary segments are not affected. The kernel width is a compromise between suppressing sub-second oscillation associated with unmodelled internal-ballistics drift (which dominates the low-erosion firings, where the reconstructed signal would otherwise show waviness comparable in magnitude to the real erosion change) and preserving the gross erosion-ramp shape, which evolves over multiple seconds and is therefore unaffected by a 500-ms low-pass operation. The smoothing step is structurally cosmetic, since the onset-detection step in Section 2.1.3 and the affine fit in Section 2.1.4 both run on the un-smoothed reconstruction. The smoothed output is, however, the form most directly suitable as an input to downstream motor-ballistics analysis. The end-points $A_t(t_{\text{rise}})$ and $A_t(t_{\text{toi}})$ are then reset to $A_{t,\text{pre}}$ and $A_{t,\text{post}}$ respectively to enforce exact continuity with the hard-coded segments at both boundaries.

2.1.6. Robustness and Fallback Cases

Four special cases are handled explicitly. The first is a firing in which erosion onset is not detected by Equation (12), either because erosion is gradual rather than abruptly initiated, or because the total erosion is at the noise level of the unscaled reconstruction. In this case the two-anchor affine fit is still applied, with the pre-anchor taken over the early-burn window described above and the post-anchor unchanged. The reported onset time is set to “not detected,” but the boundary values $A_{t,\text{pre}}$ and $A_{t,\text{post}}$ are still matched, and the reconstructed in-burn $A_t(t)$ history reflects whatever gradual erosion is encoded in $\hat{A}_t(t)$.

The second case is a firing in which the pre-test and post-test throat measurements are identical to within their measurement uncertainty, indicating no measurable erosion.

The third is a low-signal-to-noise firing in which the direct measurements indicate a small but nonzero erosion while the unscaled reconstruction shows no inter-anchor difference above its own noise floor. Without protection, this case drives $\bar{A}_{t,\text{post}}^{\text{recon}} - \bar{A}_{t,\text{pre}}^{\text{recon}}$ toward zero in Equation (17) and produces an unreasonably large α that amplifies broadband noise on $\hat{A}_t(t)$ into oscillations of tens of mm^2 in the scaled output. The algorithm therefore tests the inter-anchor difference against the standard deviation of $\hat{A}_t(t)$ over the pre-anchor window, requiring

$$\left| \bar{A}_{t,\text{post}}^{\text{recon}} - \bar{A}_{t,\text{pre}}^{\text{recon}} \right| > \text{SNR}_{\text{min}} \sigma_{\hat{A}_t}^{\text{plateau}}, \quad (19)$$

and additionally bounds the affine slope to physically reasonable values $\alpha_{\text{min}} \leq \alpha \leq \alpha_{\text{max}}$. The bounds are tied to the typical small-motor thrust-coefficient range $1.2 \lesssim C_F \lesssim 1.7$, which implies $\alpha = 1/C_F$ in the absence of C_F drift, so $\alpha_{\text{min}} = 0.3$ and $\alpha_{\text{max}} = 3$ are conservative limits. When either Equation (19) or the slope bounds are violated, the algorithm falls back to a single-anchor scaling against the pre-test measurement, with $\alpha = A_{t,\text{pre}} / \bar{A}_{t,\text{pre}}^{\text{recon}}$ and $\beta = 0$. The truly-non-eroded second case falls into this branch automatically, since Equation (19) is violated by construction. The single-anchor reconstruction preserves the shape of $\hat{A}_t(t)$ in the trustworthy interval without imposing any specific erosion-rate model;

the trace will show whatever shape is encoded in the unscaled reconstruction, including any unmodeled internal-ballistics drift not separable from real erosion at the inter-anchor signal-to-noise ratio of the firing. The post-test boundary $A_{t,\text{post}}$ is then reached only at the hard step at t_{toi} , the magnitude of which is the unresolved portion of the total erosion. For the present analysis $\text{SNR}_{\text{min}} = 4$.

The fourth case is a firing in which the post-test throat geometry was destroyed or compromised by post-fire handling. In that case α is held at the campaign-mean value from sibling tests at the same chamber-pressure operating point, and β is set to enforce only the pre-test anchor.

2.2. Erosion-Rate Extraction

The erosion rate is reported in two complementary forms. The first is the analytical instantaneous rate derived from the changepoint regression of Section 2.1.3, evaluated in the scaled domain by projecting the unscaled-domain fit coefficients through the affine map of Section 2.1.4. For the linear-right model, the throat-area erosion rate is the constant

$$\dot{A}_t(t) = b_{2,\text{scaled}} \quad \text{for } t \in [t^*, t_{\text{toi}}], \quad (20)$$

and for the parabolic-right model the rate grows linearly with time after onset,

$$\dot{A}_t(t) = 2c_{2,\text{scaled}}(t - t^*) \quad \text{for } t \in [t^*, t_{\text{toi}}]. \quad (21)$$

In both cases, the throat-diameter rate is recovered from the area rate by

$$\dot{D}_t(t) = \frac{\dot{A}_t(t)}{\sqrt{\pi A_t(t)}}, \quad (22)$$

in which $A_t(t)$ is itself given by the changepoint fit, so the analytical $\dot{D}_t(t)$ trajectory is internally consistent with the reconstructed $A_t(t)$ trajectory rather than introducing additional numerical-differentiation noise. The diameter rate is the form most directly comparable to the pin-gauge boundary measurements and to the erosion-rate correlations in the literature [1,2,5]. The choice between the linear and parabolic right-segment models is the same model selected by the changepoint regression: whichever form gave the smaller residual sum of squares is the form whose analytical rate is reported.

The total throat-diameter change over the firing,

$$\Delta D_t = 2\sqrt{A_{t,\text{post}}/\pi} - 2\sqrt{A_{t,\text{pre}}/\pi} = D_{t,\text{post}} - D_{t,\text{pre}}, \quad (23)$$

is reported as the second, independent erosion metric. By construction of the affine scaling in Section 2.1.4, ΔD_t is anchored directly to the pin-gauge measurements of the throat and is therefore not subject to the absolute-level uncertainty of the reconstructed history. Comparing the rate-from-fit and total-diameter-change metrics across tests at a common operating point provides one of the principal diagnostics for the chamber-pressure sensitivity reported in Section 4: the rate-from-fit captures both the magnitude and the time evolution of erosion within each firing, while ΔD_t provides a model-independent measurement-anchored consistency check.

2.3. Hot-Fire Validation Datasets

The reconstruction technique developed in this work has been in development for several years at the Utah State University Propulsion Research Laboratory (USU PRL), where it has been exercised on a wide range of laboratory-scale hybrid firings as the algorithm was iteratively refined. The present paper draws on two complementary subsets

of those firings to demonstrate the algorithm's behavior across the operating envelope of practical interest. The first subset is a small group of high-pressure ambient tests from the original USU PRL development series, in which 75-mm hybrid hardware was exercised at chamber pressures from approximately 1.4 to 3.4 MPa peak (200–500 psia). The second subset is a statistically meaningful sample of firings from the NASA MSFC Plume-Surface Interaction (PSI) vacuum-test campaign, which targeted three distinct chamber-pressure operating points and provided a controlled basis for between-firing comparison. The combination spans roughly an order of magnitude in chamber pressure and total throat-area change, which together let the reconstruction be evaluated both in its high-signal regime (the PRL tests) and in its low-erosion regime (the CDA low- and medium-pressure points). A full performance analysis of the PSI campaign hardware is reported separately by Whitmore and Thibaudeau [12] and is not duplicated here; the hardware, instrumentation, and test matrices used to exercise the reconstruction are summarized in the subsections below.

2.3.1. Motor Hardware

Both validation datasets use 75-mm motor casings burning gaseous oxygen (GOX) against an acrylonitrile-butadiene-styrene (ABS) fuel grain, fitted with a single-use graphite nozzle. The nozzle divergent section has a 3-degree conical half-angle and a nominal expansion ratio of 2.75 across the firings reported here. Pre-test throat geometry is measured with an expandable pin gauge whose reading is verified with calibrated calipers; multiple measurements are taken at orthogonal angular positions to confirm circularity. Post-test throat geometry is measured by the same procedure on the recovered nozzle. The CDA firings use a straight-port ABS grain of nominally 19.05 mm initial port diameter, with the grain length adjusted between the high-, medium-, and low-pressure operating points to keep the target oxidizer-to-fuel ratio near 2.0. The PRL firings considered here use the same casing and propellant combination but a helical-port grain geometry (2-turn or 4-turn), which produces a higher fuel regression rate and therefore a different chamber-pressure trajectory at the same oxidizer feed condition; the helical grain does not affect the throat-wall thermochemistry that drives erosion. Nominal hardware and operation parameters are summarized in Table 1; the per-test pre- and post-throat measurements are reported with the per-test results.

Table 1. PSI hybrid-motor nominal hardware configuration and representative CDA operating conditions.

Parameter	Value
Oxidizer	GOX
Fuel	ABS
Motor casing diameter	75 mm
Nozzle material	Bulk graphite
Pre-test throat diameter, $d_{t,pre}$	8.45 mm (0.333 in)
Exit-cone diameter, d_e	14.35 mm
Nominal expansion ratio, ϵ_0	2.75
Nozzle conical half-angle	3°
Fuel grain initial port diameter	19.05 mm (0.75 in)
Nominal burn duration	6 s
Oxidizer mass-flow rate, \dot{m}_{ox} ^a	≈35 g/s
Oxidizer mass flux ^a	≈123 kg m ⁻² s ⁻¹
Mixture ratio, O/F ^a	≈1.8–2.0
Adiabatic flame temperature ^{a,b}	≈3300 K

^a Average CDA-campaign values; the parameters above are nominal hardware values common to both PRL and CDA datasets other than what is defined in Table 2. ^b NASA CEA equilibrium adiabatic value at the nominal mixture ratio and campaign chamber pressure, accounting for combustion efficiency.

2.3.2. Instrumentation

The PRL test cart uses a calibrated load cell for axial thrust, a head-end-mounted piezoresistive chamber-pressure transducer, and a calibrated venturi for oxidizer mass flow, all logged to a LabVIEW data-acquisition system at $f_s = 100$ Hz on a common time base. The CDA test cell uses a six-degree-of-freedom thrust stand that records three orthogonal force components (L_x, L_y, L_z) and three orthogonal moment components (M_x, M_y, M_z) about the stand origin; only the axial component L_x is used for the reconstruction reported here. The CDA thrust stand is calibrated against a NIST-traceable load cell at the start of the test series and is not recalibrated between individual firings; pre-firing tare-load comparisons confirmed that no calibration drift was observed across the campaign. The CDA cell also instruments a head-end piezoresistive chamber-pressure transducer, a calibrated venturi with multiple piezoresistive pressure transducers at the upstream-static, throat, and downstream-static stations, and a barometric reference for ambient back pressure. The CDA data-acquisition system samples at $f_s = 2$ kHz. The reconstruction results reported here are insensitive to a $\pm 10\%$ perturbation of f_s for either dataset, which indicates that the reconstruction bandwidth is set by the sensor chain rather than by the data-acquisition system. The algorithm itself is sample-rate-agnostic by construction, with all of its internal smoothing and detection windows specified in time units.

2.3.3. USU PRL Ambient Test Series

The four PRL firings selected for the present analysis are listed in Table 2. All four were ambient-pressure tests at the USU PRL test cart, used a GOX/ABS propellant combination with a 0.125-inch axial injector, and ran on helical fuel-grain geometry. The peak chamber pressure spans approximately 2.3–3.4 MPa (327–498 psia), all above the chamber-pressure operating points sampled by the CDA campaign, and the total throat-diameter change spans 1.64 to 3.53 mm. These firings sit firmly in the algorithm’s high-signal regime and are included primarily to demonstrate reconstruction quality on firings where the erosion ramp is the dominant feature of the F/P_c trace.

Table 2. USU PRL ambient test matrix used in the present analysis. Peak P_c is taken from the chamber-pressure transducer record during the active firing interval; the achieved chamber pressure droops continuously over the burn as the throat erodes, so a single “achieved” value does not represent the trajectory well.

Test	Avg O/F	Peak P_c (MPa)	Burn Duration (s)
PRL Test 10	3.1	3.40	6.54
PRL Test 15	2.2	2.26	6.27
PRL Test 17	2.0	3.34	6.66
PRL Test 18	2.4	3.44	8.05

The four PRL firings used graphite nozzles of slightly different geometry; pre- and post-test throat diameters for each firing are reported with the per-test results in Section 3.

2.3.4. NASA MSFC CDA Vacuum Test Campaign

The CDA firings exercised in this work were performed as part of the NASA MSFC Plume-Surface Interaction (PSI) test campaign, conducted under NASA’s Human Landing System (HLS) program in support of the Artemis lunar-return effort [12]. The PSI campaign matrix targets a common performance specification at three chamber-pressure operating points: a high-pressure point at $P_c \approx 1.72$ MPa (250 psi), a medium-pressure point at $P_c \approx 0.86$ MPa (125 psi), and a low-pressure point at $P_c \approx 0.55$ MPa (80 psi). The motor hardware and instrumentation are common across the campaign, but the fuel grain length

and the oxidizer feed pressure were adjusted between operating points so that the achieved oxidizer-to-fuel ratio remained near 2.0 in every group [12]. Within each operating point, multiple firings were performed to provide a statistical handle on test-to-test variability. The eighteen firings selected for the present analysis are listed in Table 3, along with the per-operating-point design parameters and the achieved chamber pressure on each firing.

Table 3. NASA MSFC CDA hot-fire test matrix used in the present analysis. Target P_c , fuel grain length, and feed pressure are operating-point design parameters reported in the companion campaign analysis [12]. Achieved P_c is the time-median of the chamber-pressure record over the trustworthy reconstruction interval $[t_{\text{rise}}, t_{\text{toi}}]$ from the present reconstruction, and is consistent with the ensemble means reported in [12].

Operating Point	Test	Target P_c (MPa)	Grain Length (mm)	Feed Pressure (MPa)	Achieved P_c (MPa)
High pressure	CDA Test 7	1.72	229	2.83	1.89
	CDA Test 9				1.80
	CDA Test 10				1.86
	CDA Test 11				1.86
	CDA Test 12				1.83
	CDA Test 14				1.83
	CDA Test 15				1.85
	CDA Test 16				1.86
Medium pressure	CDA Test 24	0.86	163	1.31	0.87
	CDA Test 25				0.88
	CDA Test 26				0.85
	CDA Test 27				0.87
	CDA Test 28				0.87
Low pressure	CDA Test 17	0.55	129	0.90	0.52
	CDA Test 19				0.56
	CDA Test 20				0.55
	CDA Test 22				0.54
	CDA Test 23				0.54

3. Results

The reconstruction algorithm of Section 2 was exercised on the two complementary validation datasets introduced in Section 2.3. Section 3.1 summarizes the directly measured firing records and the boundary throat metrology from the CDA campaign, which supply the inputs and the two scaling anchors of the reconstruction. Section 3.2 presents the per-test CDA reconstruction in the six-panel diagnostic format produced by the algorithm and tabulates the resulting fit outputs, scaling coefficients, and onset-detection diagnostics. Section 3.3 aggregates the per-test CDA outputs into a campaign-level comparison across the three operating points and reports the erosion-rate metrics that feed the chamber-pressure correlation analysis of Section 4. Section 3.4 reports the corresponding outputs for the four PRL firings, which extend the validation envelope to higher chamber pressure and higher total throat-area change. Throughout the results and discussion, physical throat erosion is reported in diameter terms—throat diameter in millimeters and diameter-erosion rate in millimeters per second—with total throat change additionally expressed as a fractional throat-area change in percent; area-based quantities in SI units (m^2 , m^2/s) are retained only within the reconstruction algorithm of Section 2, where the inverse is formulated on throat area.

3.1. CDA Firing-Record and Boundary-Metrology Overview

The achieved chamber pressure of each CDA firing, taken as the time-median of P_c over the trustworthy reconstruction interval $[t_{\text{rise}}, t_{\text{toi}}]$, is the right-most column of Table 3 and confirms that the three operating points were reached within $\pm 5\%$ of their nominal targets across the campaign. The eight high-pressure firings cluster between 1.80 and 1.89 MPa with a group median of 1.86 MPa; the five medium-pressure firings span 0.85 to 0.88 MPa with a group median of 0.87 MPa; and the five low-pressure firings span 0.52 to 0.56 MPa with a group median of 0.54 MPa. The tightness of these windows is consistent with the matched motor hardware and matched run-valve schedules used across the campaign, and is the basis on which the operating-point label is treated in subsequent analyses as a clean independent variable.

The recorded firing duration $t_{\text{eo}} - t_{\text{ro}}$ averaged 6.68 ± 0.16 s at the high-pressure operating point, 6.24 ± 0.14 s at the medium-pressure point, and 5.73 ± 0.08 s at the low-pressure point, with the run-valve sequence held nominally identical across all firings. The systematic shortening of the recorded duration with decreasing chamber pressure traces directly to the rise- and tail-off-time detectors of Sections 2.1.2 and 2.1.4: as the plateau value of P_c decreases, the load-cell and pressure-transducer rise events both occur later relative to t_{ro} , and the trustworthy interval $[t_{\text{rise}}, t_{\text{toi}}]$ correspondingly contracts from a high-pressure mean of 5.34 s to a low-pressure mean of 4.88 s. The slower sensor response at the low-pressure operating point shows up in the rise-time statistics: the load-cell rise time τ_F averaged 0.82 s at high pressure and 1.10 s at low pressure, with the pressure-transducer rise time τ_{P_c} tracking it to within 30 ms in both cases, consistent with the lag-aware masking interval defined by Equation (10).

The pin-gauge measurements of the throat before and after each firing supply the two scaling anchors of the algorithm and provide a model-independent metric of total erosion. Table 4 reports the pre-test throat diameter $d_{t,\text{pre}}$, the post-test throat diameter $d_{t,\text{post}}$, the diameter change $\Delta d_t = d_{t,\text{post}} - d_{t,\text{pre}}$, and the test-record erosion classification used to gate the changepoint regression. The high-pressure group exhibited a mean diameter change of 0.374 mm with a standard deviation of 0.174 mm, the medium-pressure group a mean of 0.052 mm with a standard deviation of 0.089 mm, and the low-pressure group a mean of 0.034 mm with a standard deviation of 0.024 mm. The order-of-magnitude separation between the high-pressure group and the two lower-pressure groups is the principal feature of the boundary metrology. It is also the largest single signal that constrains the chamber-pressure correlation analysis.

The scatter of Δd_t within each operating point is shown in Figure 1. The high-pressure group is well-separated from the bulk of the medium- and low-pressure groups: with CDA Test 26 ($\Delta d_t = 0.21$ mm) excluded as a medium-pressure outlier, all remaining medium- and low-pressure firings lie below $\Delta d_t = 0.08$ mm, which is the lowest single value in the high-pressure group (CDA Test 15). CDA Test 26 is the only firing at sub-high-pressure conditions whose Δd_t falls within the high-pressure range. The high-pressure group itself shows a within-group spread of more than a factor of eight, from $\Delta d_t = 0.08$ mm to $\Delta d_t = 0.69$ mm, despite the achieved chamber pressures sitting within a 0.09 MPa band. This intra-group spread at fixed nominal operating point is a recurring feature of the campaign and is examined in Section 3.2 in terms of the within-firing time evolution of the throat geometry.

Table 4. Boundary throat metrology and erosion classification for the CDA validation campaign.

Test	Group	$d_{t,pre}$ (mm)	$d_{t,post}$ (mm)	Δd_t (mm)	Eroded
CDA Test 7	High	8.44	8.85	0.41	Yes
CDA Test 9	High	8.45	8.92	0.47	Yes
CDA Test 10	High	8.53	8.82	0.29	Yes
CDA Test 11	High	8.43	8.81	0.38	Yes
CDA Test 12	High	8.45	8.83	0.38	Yes
CDA Test 14	High	8.47	9.16	0.69	Yes
CDA Test 15	High	8.43	8.51	0.08	Yes
CDA Test 16	High	8.50	8.79	0.29	Yes
CDA Test 24	Medium	8.42	8.42	0.00	No
CDA Test 25	Medium	8.43	8.45	0.02	Yes
CDA Test 26	Medium	8.44	8.65	0.21	Yes
CDA Test 27	Medium	8.41	8.44	0.03	Yes
CDA Test 28	Medium	8.47	8.47	0.00	No
CDA Test 17	Low	8.45	8.49	0.04	Yes
CDA Test 19	Low	8.43	8.49	0.06	Yes
CDA Test 20	Low	8.41	8.46	0.05	Yes
CDA Test 22	Low	8.42	8.44	0.02	Yes
CDA Test 23	Low	8.42	8.42	0.00	No

Group denotes the nominal chamber-pressure operating point of the CDA test matrix (Table 3): High \approx 1.72 MPa, Medium \approx 0.86 MPa, and Low \approx 0.55 MPa target P_c . **Eroded** is the test-record erosion classification used to gate the changepoint regression (Section 2.1.3): “Yes” = measurable throat recession, “No” = no measurable recession.

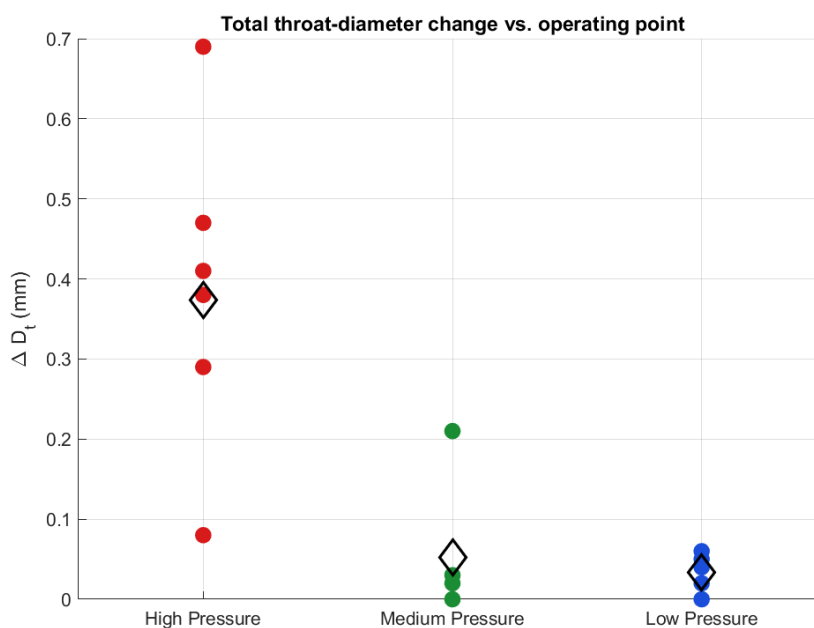


Figure 1. Total throat-diameter change by operating point, CDA campaign. The black diamond is the average value for each set.

3.2. Per-Test CDA Throat-Area Reconstruction

The reconstruction algorithm produces a six-panel diagnostic for each firing that records the directly measured signals, the unscaled and scaled reconstruction stages, the changepoint regression, and the analytical erosion-rate model in a single layout. The top-left panel overplots the axial-thrust and chamber-pressure records on a common time axis, with vertical markers placed at t_{ro} , t_{rise} , the detected onset t^* , and the tail-off-initiation time t_{toi} . The top-center panel shows the unscaled reconstruction $\hat{A}_t(t) = F(t)/P_c(t)$ over the full data-acquisition window, including the lag-spike interval that is masked from downstream analysis. The top-right panel shows the result of the two-anchor affine map of Equations (14)–(18) applied without the pre-onset hard-coding of Section 2.1.4, and

overlays the changepoint regression model fit. The bottom-left panel shows the same scaled reconstruction with the pre-onset segment hard-coded to the directly measured pre-test throat area, which is the form used as the input to all downstream analyses. The bottom-center panel converts the area history into a throat-diameter history $D_t(t) = 2\sqrt{A_t(t)}/\pi$ for direct comparison with the pin-gauge measurements. The bottom-right panel reports the analytical erosion rate evaluated from the fit coefficients projected through the affine map, using either Equation (20) or Equation (21) depending on the model selected by the changepoint regression.

Four representative CDA firings are presented to illustrate the range of behaviors encountered across the campaign. Figure 2 shows CDA Test 10, a high-pressure firing on which the changepoint regression selected the linear right-segment model and the two-anchor affine scaling was applied without fallback. Figure 3 shows CDA Test 14, the high-pressure firing with the largest measured Δd_t in the campaign, on which the changepoint regression selected the parabolic right-segment model with a relative residual reduction of 38% over the linear model. Figure 4 shows CDA Test 17, the only low-pressure firing on which a confident two-anchor scaling was supported by the unscaled-signal SNR test of Equation (19). Figure 5 shows CDA Test 24, a medium-pressure firing on which the boundary metrology recorded zero erosion and the algorithm correctly fell through to the no-changepoint, single-anchor fallback path.

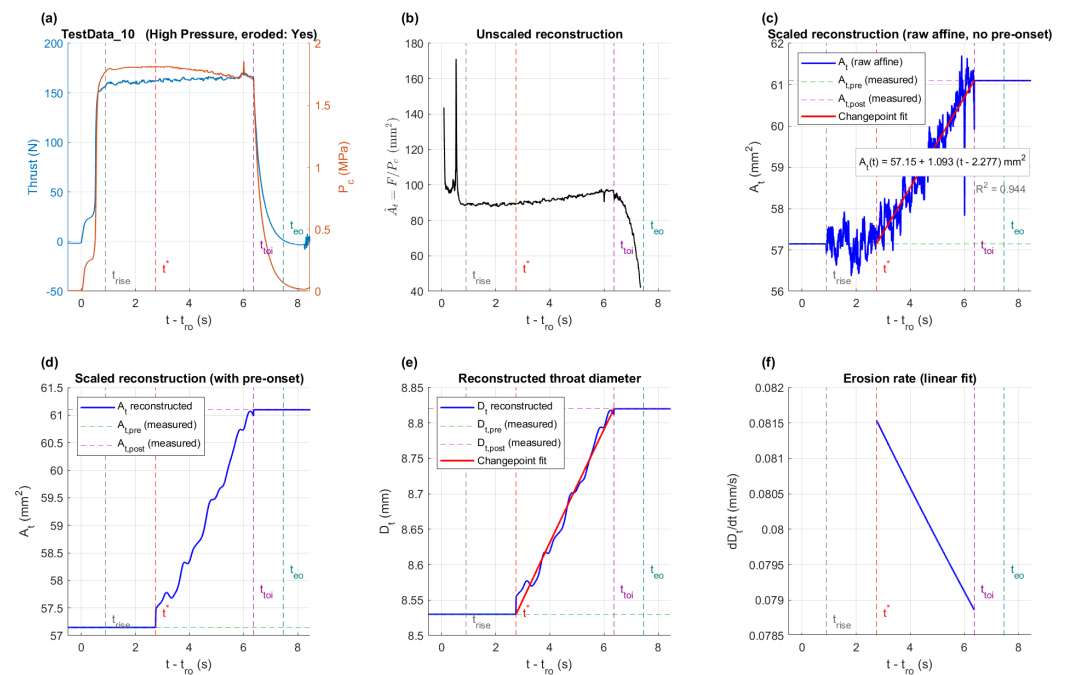


Figure 2. Reconstruction diagnostic for CDA Test 10 (high pressure, linear fit). Panels: (a) measured thrust and chamber pressure; (b) unscaled inverse $\hat{A}_t = F/P_c$; (c) scaled reconstruction (raw two-anchor affine) with measured anchors and changepoint fit; (d) scaled reconstruction with pre-onset hard-coding; (e) reconstructed throat diameter D_t ; (f) post-onset erosion-rate fit.

The fit coefficients, scaling coefficients, and onset-detection diagnostics for all eighteen CDA firings are reported in Tables 5–7, partitioned by operating point so that the within-group statistics can be summarized at the foot of each table. Three of the eighteen firings are classified by the campaign metadata as not eroded; on these, the algorithm applies the single-anchor scaling and reports no fit model. Five firings (CDA Tests 25, 27, 19, 20, and 22) are classified as eroded by the metadata but fail the inter-anchor SNR test of Equation (19). The algorithm applies the single-anchor fallback while still attempting a changepoint regression on the unscaled signal; the resulting fit is reported, but the post-

onset diameter-erosion rate on these firings is below 0.01 mm/s and is not separable from the residual unmodeled drift on the unscaled signal at the given SNR. The remaining ten firings (eight high-pressure, one medium-pressure at CDA Test 26, and one low-pressure at CDA Test 17) cleared both the metadata gate and the SNR test; these are the firings that supply quantitative erosion-rate measurements to the campaign-level analysis.

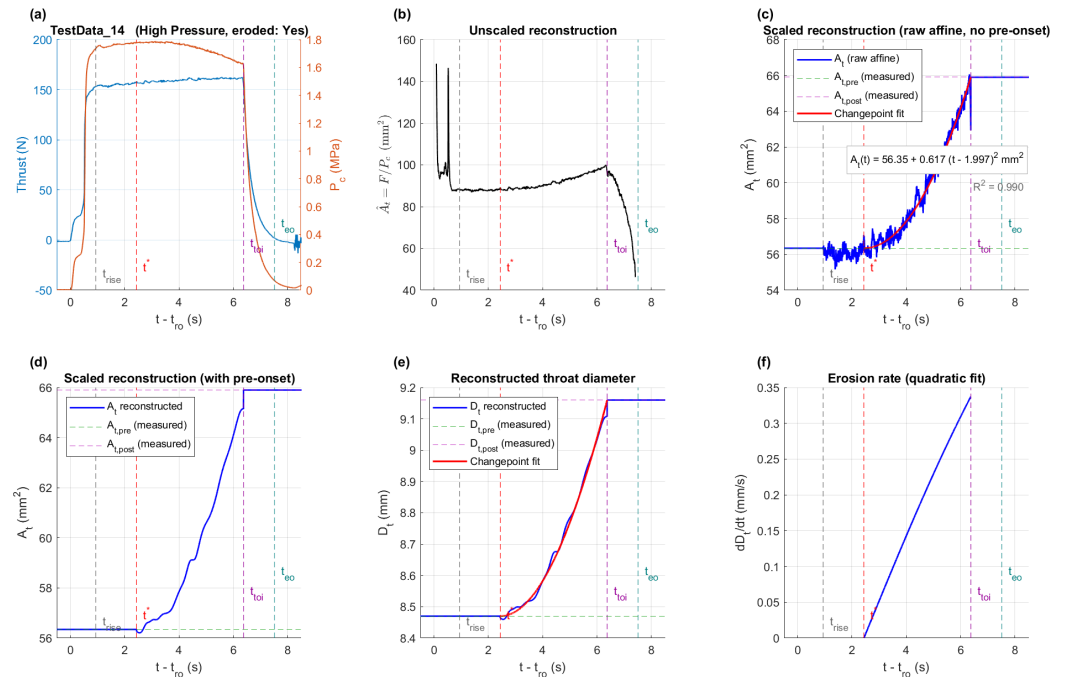


Figure 3. Reconstruction diagnostic for CDA Test 14 (high pressure, parabolic fit). Panels: (a) measured thrust and chamber pressure; (b) unscaled inverse $\hat{A}_t = F/P_c$; (c) scaled reconstruction (raw two-anchor affine) with measured anchors and changepoint fit; (d) scaled reconstruction with pre-onset hard-coding; (e) reconstructed throat diameter D_t ; (f) post-onset erosion-rate fit.

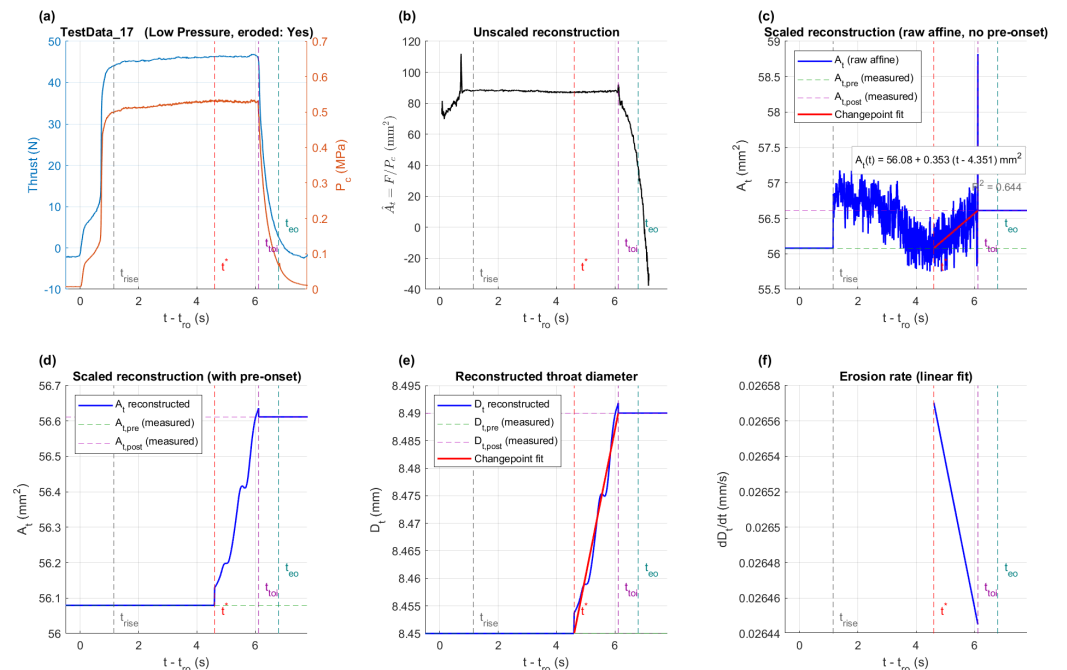


Figure 4. Reconstruction diagnostic for CDA Test 17 (low pressure). Panels: (a) measured thrust and chamber pressure; (b) unscaled inverse $\hat{A}_t = F/P_c$; (c) scaled reconstruction (raw two-anchor affine) with measured anchors and changepoint fit; (d) scaled reconstruction with pre-onset hard-coding; (e) reconstructed throat diameter D_t ; (f) post-onset erosion-rate fit.

A direct consequence of the single-anchor fallback is that the reconstructed in-burn $A_t(t)$ trace does not in general line up with the post-test nozzle measurement at t_{toi} . CDA Tests 19, 23, and 28 illustrate the effect most clearly. On CDA Test 19 (low pressure, $\Delta d_t = 0.06$ mm classified as eroded but failing the SNR test), the reconstructed diameter sits at approximately 8.43 mm through the full firing window and a hard step takes the trace to the post-test value of 8.49 mm at t_{toi} . On CDA Tests 23 and 28 (low and medium pressure, both classified as not eroded), the reconstruction wanders by a few percent of A_t within its noise envelope and likewise meets the post-test boundary at t_{toi} via the hard-coded post-tail-off segment rather than by smooth convergence. The visible step is not a deficiency of the reconstruction; it is the direct expression of the single-anchor fallback, in which only the pre-test boundary is matched by the affine map ($\beta = 0, \alpha = A_{t,pre} / \bar{A}_{t,pre}^{recon}$) and the post-test boundary is reached only by the hard-coded segment. The gap between the in-burn trace and the post-test value at t_{toi} is the portion of the actual erosion that the unscaled signal could not resolve at the inter-anchor signal-to-noise ratio of the firing, and it is the unresolved residual that the algorithm preserves rather than smoothing over.

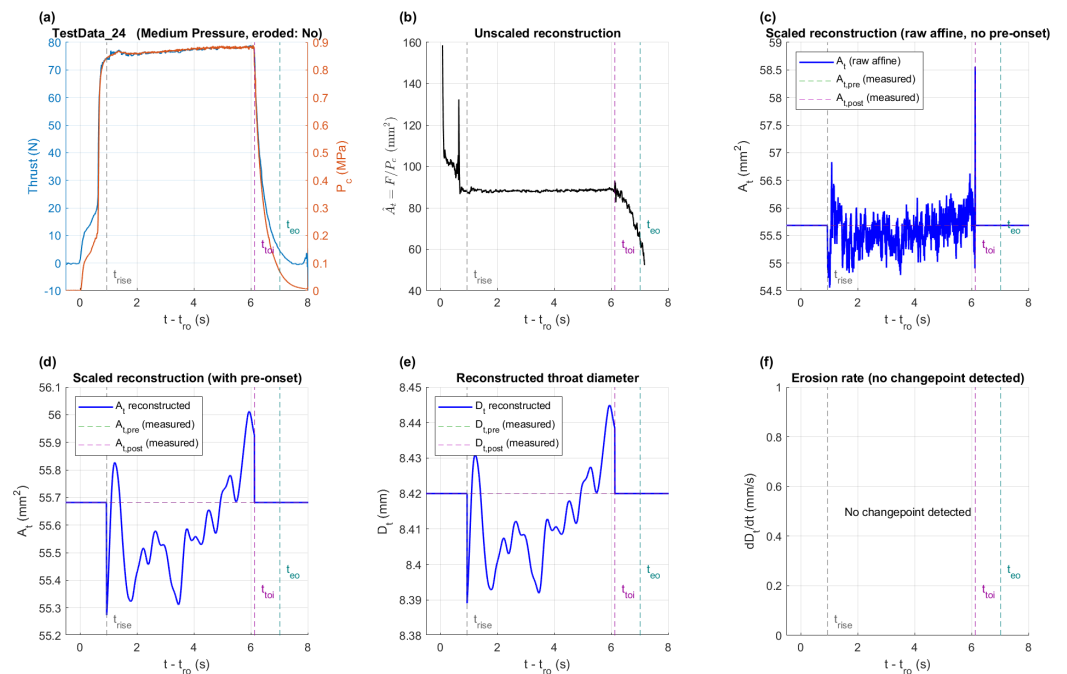


Figure 5. Reconstruction diagnostic for CDA Test 24 (medium pressure, no erosion). Panels: (a) measured thrust and chamber pressure; (b) unscaled inverse $\hat{A}_t = F/P_c$; (c) scaled reconstruction (raw two-anchor affine) with measured anchors and changepoint fit; (d) scaled reconstruction with pre-onset hard-coding; (e) reconstructed throat diameter D_t ; (f) post-onset erosion-rate fit.

Three of the eight high-pressure firings (CDA Test 12, CDA Test 14, CDA Test 16) selected the parabolic right-segment model. The selection is correlated with the magnitude of Δd_t : CDA Test 14 produced the largest Δd_t in the campaign, and CDA Test 12 and CDA Test 16 are above the high-pressure group median. The five high-pressure firings on which the linear model was retained are clustered at smaller total erosion (Δd_t between 0.08 and 0.47 mm). The pattern is consistent with the parabolic right-segment shape capturing a real late-burn acceleration in the highest-erosion firings, where the kinetically-limited surface-reaction regime accelerates with rising wall temperature and rising local oxidizer concentration as the throat opens. Firings that did not reach the geometric and thermal conditions required to trigger that acceleration retained a linear post-onset rate. The implication for the chamber-pressure correlation analysis is that the time-averaged erosion rate is not a single-valued function of chamber pressure within the high-pressure group

itself, and that the changepoint regression captures a genuine variation in the time-evolution of erosion that the boundary metrology alone could not resolve.

A side-effect of the parabolic-versus-linear model selection is visible in the detected onset times themselves. The three high-pressure firings on which the parabolic right-segment was retained (CDA Tests 12, 14, and 16) report t^* in a tight band of 2.00 to 2.07 s, while the five linear-fit firings span 2.28 to 3.67 s with a group mean closer to 2.84 s. The parabolic model places t^* approximately 0.8 s earlier than a linear fit to the same data would, and the offset is structural rather than physical. As discussed in Section 2.1.3, the parabolic right-segment imposes zero slope at t^* and a slope that grows linearly with elapsed time. On a roughly-linear erosion ramp, the regression therefore has to push t^* to the left of the true onset to give the curvature room to “ramp up” to the observed late-burn slope. The 10%-RSS-reduction margin in Equation (13) suppresses this bias on records where the underlying erosion is genuinely linear, but cannot remove it on records where the parabolic form is correctly preferred. Boundary-anchored metrics such as Δd_t and the integrated total throat-area change are insensitive to this placement bias because they depend only on the pre- and post-test geometry, but the absolute value of t^* on parabolic-fit firings should be interpreted as a fit parameter rather than as a confident timestamp of the physical onset of erosion.

The pre-onset interval is short relative to the total firing duration on the high-pressure firings, with the detected t^* falling between 2.0 and 3.7 s after t_{rise} and a group-mean pre-onset duration of 2.5 s. The medium- and low-pressure firings that supplied a confident detection ($t^* = 2.01$ s for CDA Test 26 and $t^* = 4.35$ s for CDA Test 17) bracket this range, with the medium-pressure record showing a pre-onset duration comparable to the high-pressure group and the low-pressure record showing a markedly delayed onset. The latter is consistent with the lower wall-heat-flux and lower oxidizer concentration at the low-pressure operating point delaying the transition from inert to actively-eroding throat geometry; the former indicates that even at the medium-pressure point, where most firings recorded sub-resolution erosion, the one firing that did erode (CDA Test 26, $\Delta d_t = 0.21$ mm) followed the same onset timing as a high-pressure firing.

Table 5. CDA high-pressure reconstruction outputs.

Test	t^* (s)	Model	$b_{2,\text{scaled}}$ (mm^2/s)	$c_{2,\text{scaled}}$ (mm^2/s^2)	α	β (mm^2)	R^2	Status
CDA Test 7	2.75	linear	1.789	—	0.720	−8.08	0.939	two-anchor
CDA Test 9	2.68	linear	2.021	—	0.703	−6.83	0.988	two-anchor
CDA Test 10	2.28	linear	1.093	—	0.515	11.29	0.944	two-anchor
CDA Test 11	2.81	linear	1.677	—	0.999	−32.81	0.988	two-anchor
CDA Test 12	2.07	parabolic	—	0.354	0.876	−22.51	0.994	two-anchor
CDA Test 14	2.00	parabolic	—	0.617	0.831	−16.98	0.990	two-anchor
CDA Test 15	3.67	linear	0.486	—	0.394	21.87	0.942	two-anchor
CDA Test 16	2.00	parabolic	—	0.267	0.816	−16.64	0.987	two-anchor
Mean ($n = 8$)	2.53	—	—	—	0.732	—	0.971	—
Std. dev. ($n = 8$)	0.57	—	—	—	0.197	—	0.025	—

Table 6. CDA medium-pressure reconstruction outputs.

Test	t^* (s)	Model	$b_{2,scaled}$ (mm^2/s)	$c_{2,scaled}$ (mm^2/s^2)	α	β (mm^2)	R^2	Status
CDA Test 24	—	—	—	—	0.629	0	—	no-erosion, single-anchor
CDA Test 25	3.48	linear	0.100	—	0.645	0	0.203	low-SNR, single-anchor
CDA Test 26	2.01	linear	0.730	—	0.629	3.20	0.916	two-anchor
CDA Test 27	4.34	linear	0.356	—	0.655	0	0.412	low-SNR, single-anchor
CDA Test 28	—	—	—	—	0.656	0	—	no-erosion, single-anchor
Mean ($n = 3$ fits)	3.28	—	—	—	0.643	—	0.511	—
Std. dev. ($n = 3$)	1.18	—	—	—	0.013	—	0.366	—

Table 7. CDA low-pressure reconstruction outputs.

Test	t^* (s)	Model	$b_{2,scaled}$ (mm^2/s)	$c_{2,scaled}$ (mm^2/s^2)	α	β (mm^2)	R^2	Status
CDA Test 17	4.35	linear	0.353	—	0.579	5.67	0.644	two-anchor
CDA Test 19	—	linear	0.010	—	0.655	0	0.008	low-SNR, forced-onset
CDA Test 20	—	linear	0.076	—	0.665	0	0.384	low-SNR, forced-onset
CDA Test 22	3.87	linear	−0.108	—	0.643	0	0.097	low-SNR, single-anchor
CDA Test 23	—	—	—	—	0.672	0	—	no-erosion, single-anchor
Mean ($n = 2$ detected)	4.11	—	—	—	0.611	—	0.370	—
Std. dev. ($n = 2$)	0.34	—	—	—	0.045	—	0.387	—

3.3. CDA Campaign-Level Erosion Summary

The reconstructed throat-diameter histories from all firings are overplotted in three group panels in Figure 6, with the time axis shifted to a common reference at t_{rise} on each firing so that the pre-onset plateaus of all records align at the origin. The three operating points produce visibly distinct erosion-ramp magnitudes. The high-pressure panel shows $D_t(t)$ growing by 0.1 to 0.7 mm over the active firing interval, with the largest excursion belonging to CDA Test 14. The medium-pressure panel shows $D_t(t)$ growing by less than 0.05 mm on all but one firing, and the one exception (CDA Test 26) is the firing whose boundary metrology recorded $\Delta d_t = 0.21$ mm. The low-pressure panel shows $D_t(t)$ growing by less than 0.05 mm on all firings, with the trajectories sitting close to the noise floor of the reconstruction.

The medium- and low-pressure panels of Figure 6, and the corresponding panels of the per-test reconstruction diagnostics for the same firings, present a deceptively noisy visual appearance that should be read in the context of the y-axis autoscaling rather than the absolute uncertainty of the reconstruction. The y-axis on these panels is set to the span of the reconstructed trace itself, which on the lowest-erosion firings is on the order of 0.1 mm in diameter or 1 mm^2 in area, so the broadband fluctuation that is invisible against the 0.7-mm growth of CDA Test 14 fills the entire vertical extent of the medium- and low-pressure plots. In absolute terms, the peak-to-peak excursion of the in-burn trace from the pre-test boundary value remains under approximately 0.05 mm of diameter on every firing in these two groups, which is roughly 0.6% of the nominal 8.45-mm throat. That excursion is several times the pin-gauge precision of ± 0.01 mm, so the boundary anchors are tight and the apparent “noise” on these panels reflects the in-burn reconstruction itself rather than uncertainty in the metrology. The apparent waviness is a plot-scaling artifact, not a statement about the absolute fidelity of the algorithm.

The analytical post-onset erosion-rate models from Equations (20) and (21) are overplotted in Figure 7, with the time axis on each firing shifted to a common reference at the

detected onset t^* so that the post-onset interval begins at the origin. The high-pressure panel includes all eight firings, with the parabolic-fit firings rising linearly from zero at t^* and the linear-fit firings sitting at constant rates between 0.04 and 0.18 mm/s; the parabolic firings cross the linear-fit envelope between $t - t^* \approx 1$ and 2 s and continue rising. The medium-pressure panel includes only the three firings that produced an accepted change-point fit, of which two (CDA Tests 25 and 27) sit at rates below 0.03 mm/s. The low-pressure panel includes the four firings that produced a fit, of which only CDA Test 17 sits clearly above the noise floor.




Figure 6. CDA reconstructed throat-diameter histories by operating point.

The mid-burn instantaneous throat-diameter erosion rate, defined as the analytical model evaluated at $t = (t^* + t_{toi})/2$, is the single-number summary used in the chamber-pressure correlation analysis of Section 4. Restricted to the ten firings that cleared the SNR test and supplied a quantitative two-anchor reconstruction, the mid-burn rate averages 0.111 ± 0.046 mm/s at the high-pressure operating point ($n = 8$, range 0.037 to 0.179 mm/s), 0.055 mm/s on the single qualifying medium-pressure firing (CDA Test 26), and 0.027 mm/s on the single qualifying low-pressure firing (CDA Test 17). The factor of approximately four between the high-pressure mean and the low-pressure single qualifying value, and the factor of approximately two between the high-pressure mean and the medium-pressure single qualifying value, are the principal quantitative results that the discussion section interprets in terms of chamber-pressure scaling laws and the diffusion-limited-versus-kinetically-limited regime structure of graphite erosion.

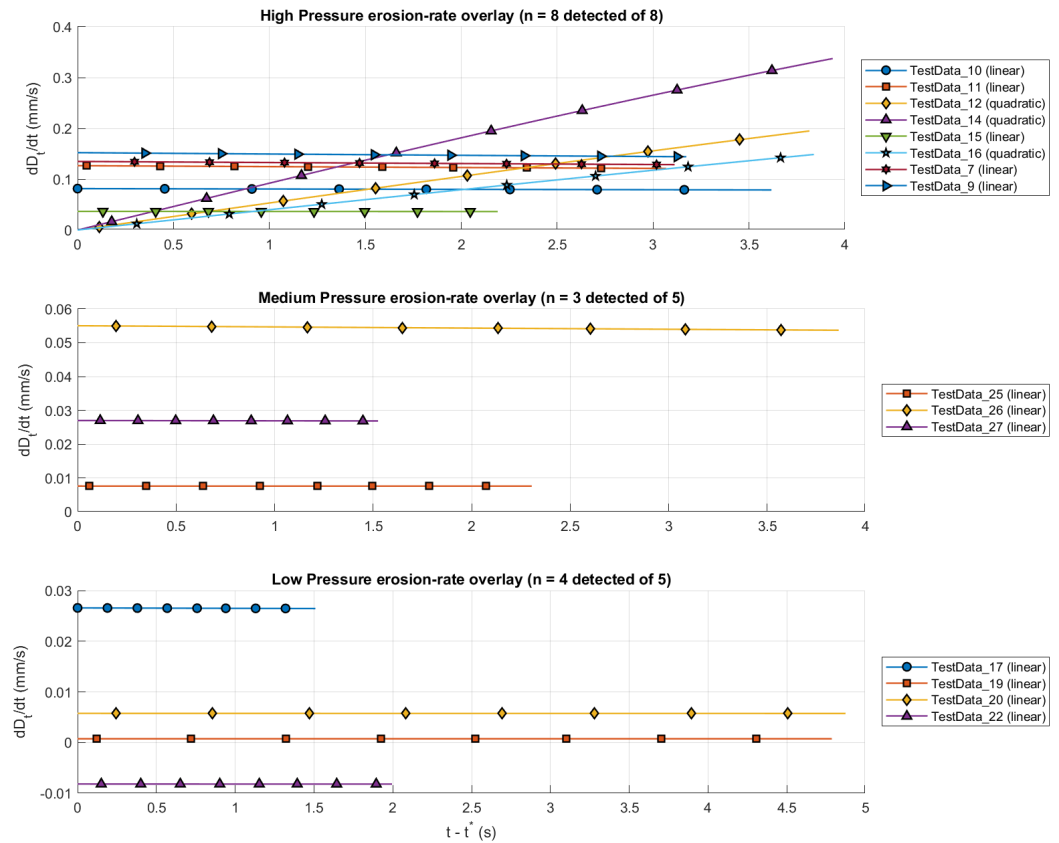


Figure 7. CDA post-onset diameter-erosion rate by operating point.

3.4. USU PRL High-Pressure Reconstruction Results

The four PRL firings listed in Table 2 were processed through the same reconstruction pipeline used for the CDA dataset. The per-firing pre- and post-test throat diameters, fit coefficients, and onset diagnostics are reported in Table 8. Total throat-diameter change spans 1.64 to 3.53 mm across the four firings, corresponding to fractional throat-area changes of approximately 31% to 103%—roughly an order of magnitude above the largest CDA firing (CDA Test 14 at 17%). The changepoint regression selected the linear right-segment model on all four firings, with $R^2 \geq 0.996$ in every case. Mid-burn instantaneous throat-diameter erosion rates range from 0.43 mm/s (PRL Test 15, peak $P_c \approx 2.3$ MPa) to 0.81 mm/s (PRL Test 17, peak $P_c \approx 3.3$ MPa), a factor of 4 to 7 above the high-pressure CDA mean of 0.11 mm/s.

Table 8. PRL high-pressure reconstruction outputs. PRL Test 10 used the forced-onset path of Section 2.1.3 ($t^* = t_{rise}$) because the chamber-pressure droop began within the ignition transient, faster than the standard 1.5-s changepoint search margin can resolve.

Test	$d_{t,pre}$ (mm)	$d_{t,post}$ (mm)	Δd_t (mm)	$b_{2,scaled}$ (mm ² /s)	α	R^2	Mid-Burn Rate (mm/s)	Status
PRL Test 10	8.28	11.81	3.53	9.31	0.704	0.996	0.72	two-anchor, forced-onset
PRL Test 15	11.16	12.80	1.64	7.60	0.976	0.996	0.43	two-anchor
PRL Test 17	8.66	11.27	2.61	11.04	0.785	0.999	0.81	two-anchor
PRL Test 18	8.66	11.31	2.65	10.21	0.731	0.999	0.75	two-anchor

Figures 8–11 show the six-panel per-test diagnostic for each PRL firing. The visual signature of erosion is strikingly different from the CDA panels: the chamber-pressure trace decays continuously over the burn rather than holding a plateau, the unscaled F/P_c inverse rises steeply and monotonically over the trustworthy interval, and the reconstructed throat-area history grows by tens of mm^2 rather than by the 1–10 mm^2 scale seen on the CDA firings. The reconstruction quality is correspondingly cleaner: the affine-mapped trace tracks the changepoint fit to within the line thickness over the entire post-onset segment, and the residual between the fit and the underlying data is comparable to the broadband noise floor of the unscaled signal alone.

The PRL reconstructed throat-diameter histories are overlaid on a single panel in Figure 12 (time referenced to t_{ro}) and the post-onset analytical erosion rates in Figure 13 (time referenced to t^*). The four firings cluster within a factor of two of each other in both diameter trajectory and rate. PRL Tests 17 and 18, which reached the highest peak chamber pressures, sit at the top of the rate envelope; PRL Test 15, with the lowest peak chamber pressure of the four, sits at the bottom. The ordering by peak P_c is consistent with the high-pressure end of the chamber-pressure scaling implied by the CDA data, which is examined further in Section 4.

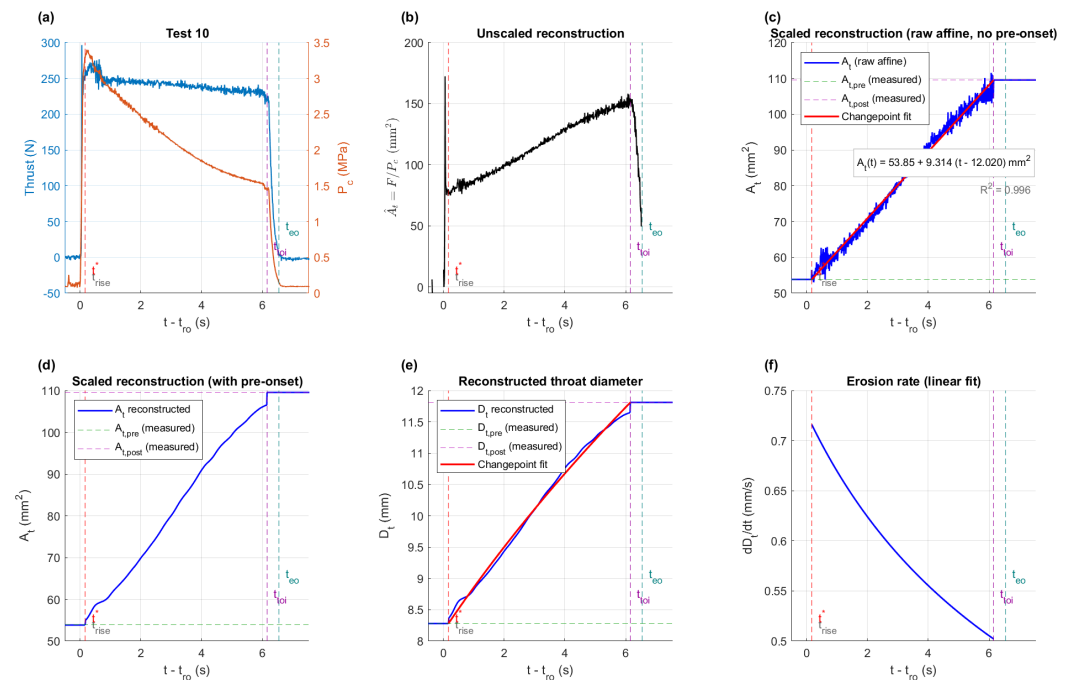


Figure 8. Reconstruction diagnostic for PRL Test 10 (peak $P_c \approx 3.4$ MPa, helical 4-turn grain). Forced-onset path used because erosion begins within the ignition transient. Panels: (a) measured thrust and chamber pressure; (b) unscaled inverse $\hat{A}_t = F/P_c$; (c) scaled reconstruction (raw two-anchor affine) with measured anchors and changepoint fit; (d) scaled reconstruction with pre-onset hard-coding; (e) reconstructed throat diameter D_t ; (f) post-onset erosion-rate fit.

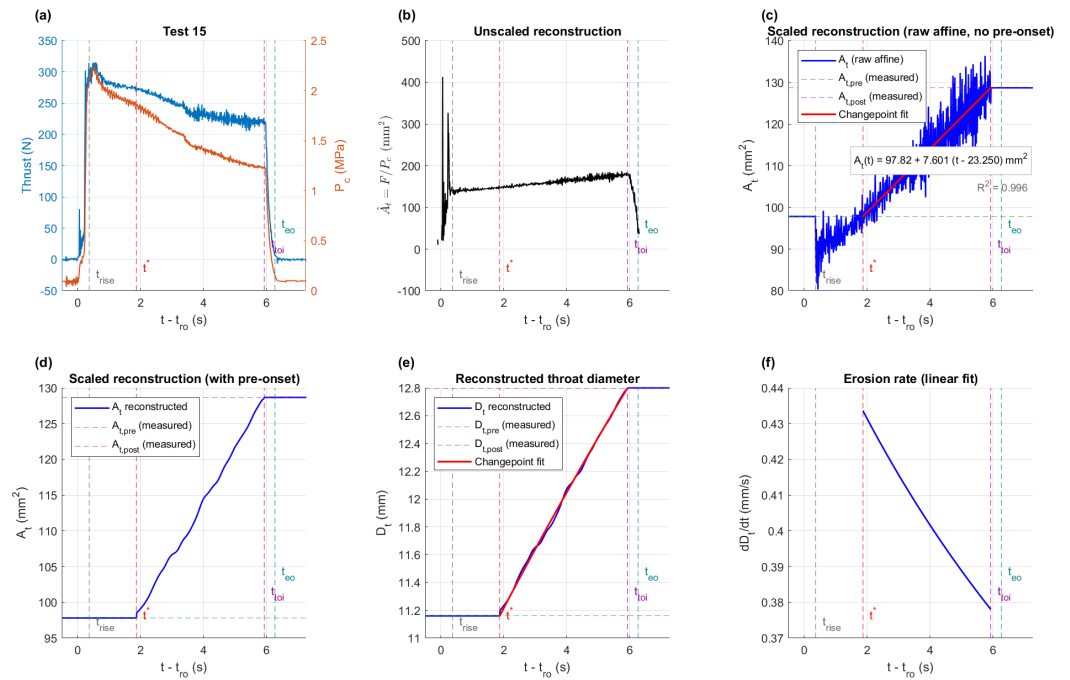


Figure 9. Reconstruction diagnostic for PRL Test 15 (peak $P_c \approx 2.3$ MPa, helical 2-turn grain). Panels: (a) measured thrust and chamber pressure; (b) unscaled inverse $\hat{A}_t = F/P_c$; (c) scaled reconstruction (raw two-anchor affine) with measured anchors and changepoint fit; (d) scaled reconstruction with pre-onset hard-coding; (e) reconstructed throat diameter D_t ; (f) post-onset erosion-rate fit.

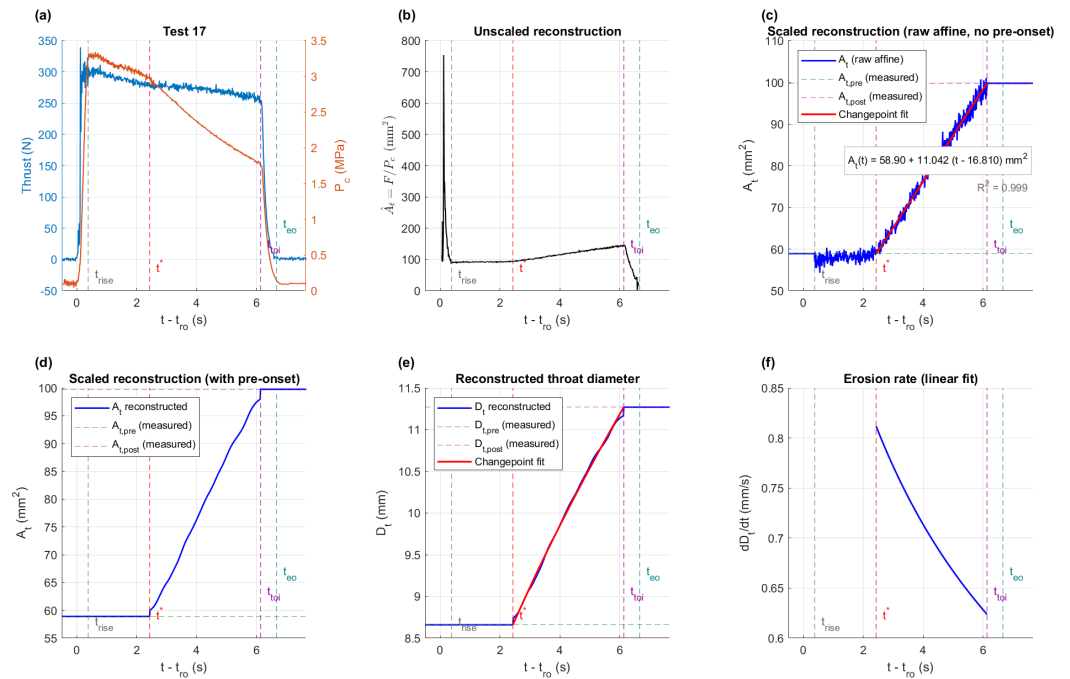


Figure 10. Reconstruction diagnostic for PRL Test 17 (peak $P_c \approx 3.3$ MPa, helical 2-turn grain). Panels: (a) measured thrust and chamber pressure; (b) unscaled inverse $\hat{A}_t = F/P_c$; (c) scaled reconstruction (raw two-anchor affine) with measured anchors and changepoint fit; (d) scaled reconstruction with pre-onset hard-coding; (e) reconstructed throat diameter D_t ; (f) post-onset erosion-rate fit.

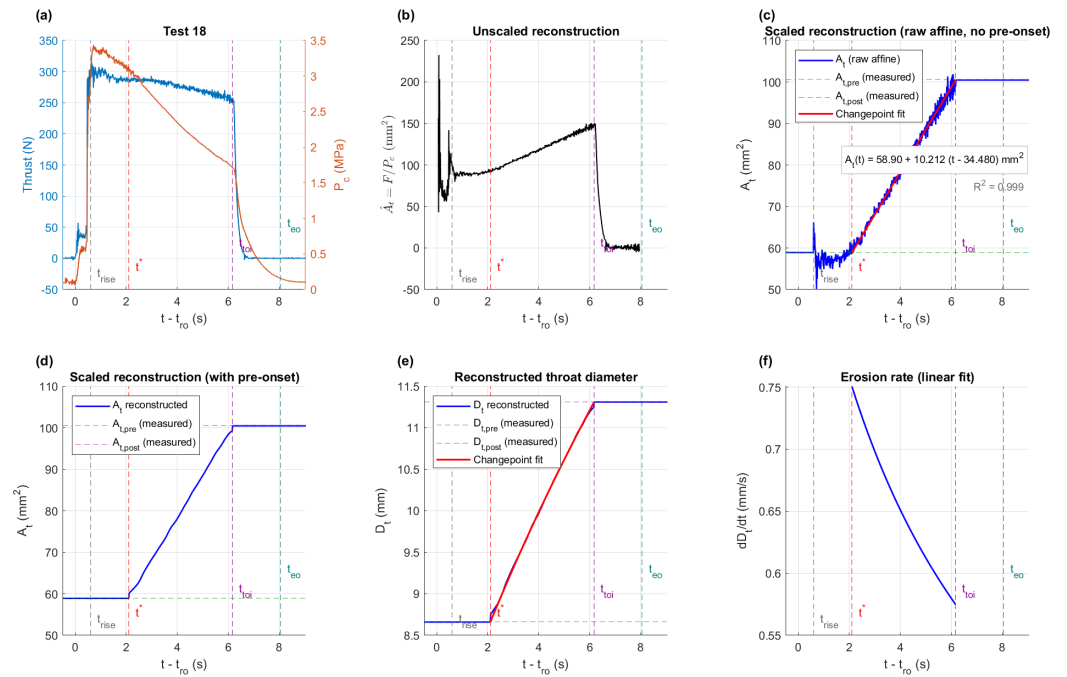


Figure 11. Reconstruction diagnostic for PRL Test 18 (peak $P_c \approx 3.4$ MPa, helical 2-turn grain). Panels: (a) measured thrust and chamber pressure; (b) unscaled inverse $\hat{A}_t = F/P_c$; (c) scaled reconstruction (raw two-anchor affine) with measured anchors and changepoint fit; (d) scaled reconstruction with pre-onset hard-coding; (e) reconstructed throat diameter D_t ; (f) post-onset erosion-rate fit.

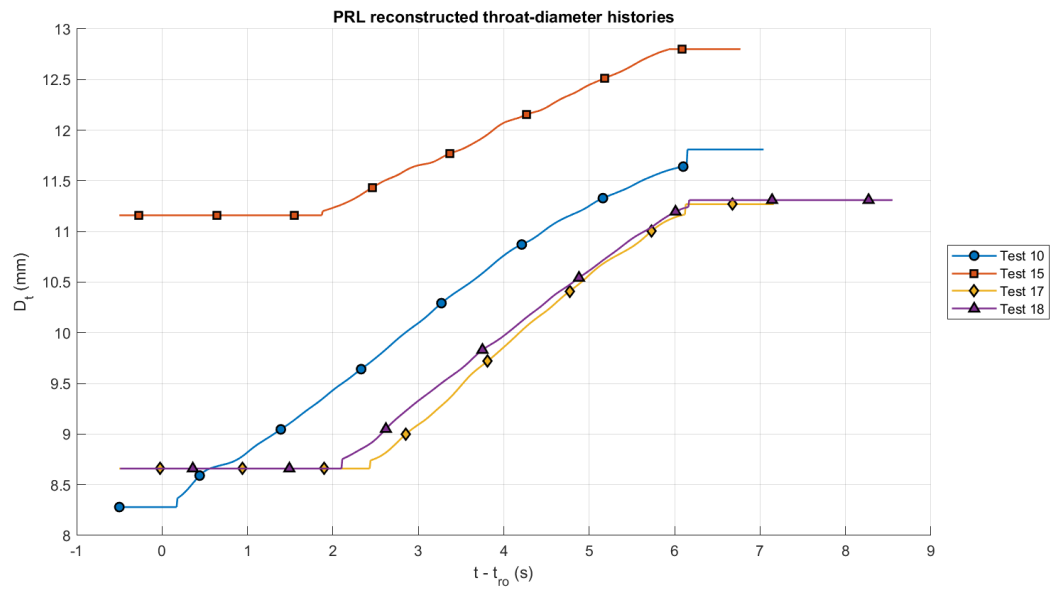


Figure 12. PRL reconstructed throat-diameter histories.

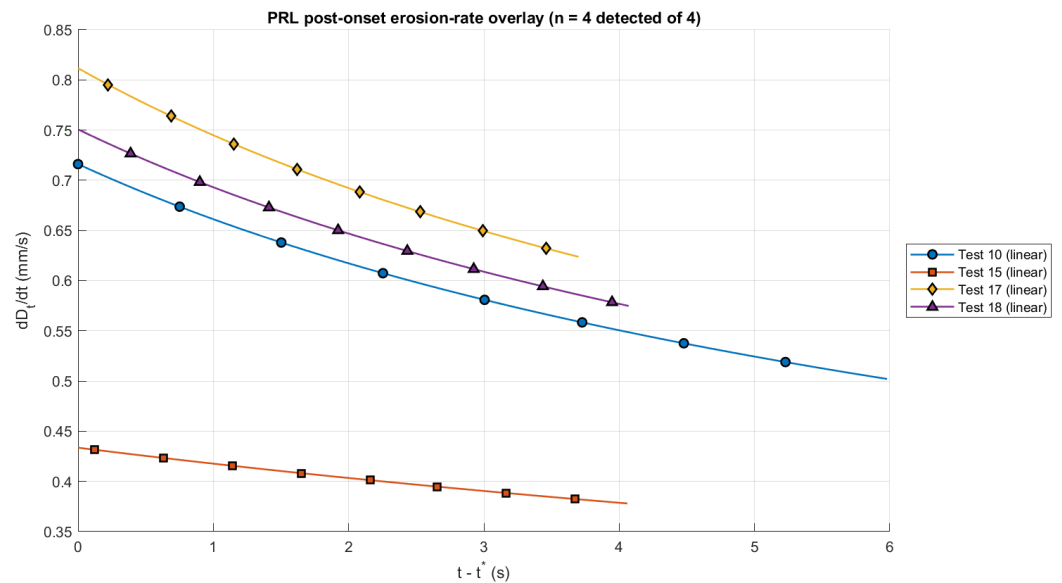


Figure 13. PRL post-onset diameter-erosion rates.

4. Discussion

The reconstruction outputs reported in Section 3 are interpreted in this section in four steps. Section 4.1 characterizes the algorithm's performance as a function of the total amount of erosion observed in a firing, drawing on the combined CDA and PRL datasets which together span roughly two orders of magnitude in fractional throat-area change. Section 4.2 treats chamber pressure as the principal independent variable and characterizes the pressure-erosion relationship using the CDA's three controlled operating points and the PRL's higher-pressure points to bracket the knee-point behavior. Section 4.3 pivots to the secondary drivers that chamber pressure alone cannot capture, using the within-group spread of the high-pressure CDA firings as the primary evidence and the PRL geometric variations as a consistency check. Section 4.4 closes the discussion with the modelling assumptions, the dominant uncertainty contributors, and the applicability envelope of the reconstruction.

4.1. Algorithm Performance and Detection Limits

The reconstruction algorithm's fit quality depends strongly on how much erosion actually occurred during a firing. Plotting the changepoint-regression R^2 against the fractional throat-area change $\Delta A_t / A_{t,pre}$ across both validation datasets makes the relationship clear. The eight high-pressure CDA firings, with $\Delta A_t / A_{t,pre}$ ranging from 1.9% on CDA Test 15 to 17.0% on CDA Test 14, all produced $R^2 \geq 0.94$ and a clean two-anchor reconstruction. CDA Test 26 at medium pressure ($\Delta A_t / A_{t,pre} \approx 5\%$) reported $R^2 = 0.92$. The four PRL firings, all sitting well above the high-pressure end of the CDA range with $\Delta A_t / A_{t,pre}$ from 31% on PRL Test 15 to 103% on PRL Test 10, all produced $R^2 \geq 0.996$. The marginal cases sit at roughly 1% area change: CDA Test 17 (0.95% area change, $R^2 = 0.644$), CDA Test 20 (1.2%, $R^2 = 0.384$, forced-onset fallback), and CDA Test 19 (1.4%, $R^2 = 0.008$, forced-onset fallback). The smallest measured changes—CDA Tests 22 and 25 at roughly 0.5% area, and CDA Tests 23, 24, and 28 at essentially zero—either failed the SNR test outright or produced R^2 values below 0.3.

Empirically, the transition between unreliable and reliable reconstruction sits at approximately $\Delta A_t / A_{t,pre} \approx 2\%$, corresponding to a throat-diameter change of about 0.08 mm at the nominal 8.45 mm throat. Above this threshold the changepoint-regression R^2 consistently exceeds 0.9 and the in-burn shape of $A_t(t)$ is well resolved. Below the threshold the algorithm still produces an output, but the in-burn trace is dominated by F/P_c reconstruc-

tion noise rather than by a real erosion signal, and the single-anchor fallback of Section 2.1.6 typically takes over.

The drop in fit quality at low erosion sounds worse than it is in practice. The throat-area history feeds subsequent internal-ballistics analysis through the predicted thrust, the predicted total mass flow, the inferred fuel mass flow, the running O/F ratio, and the CEA-derived γ , c^* , and C_F that close the iterative coupled solver. A fractional error $\delta A_t/A_t$ propagates as a fractional error of comparable magnitude into F and \dot{m} . Propagation into \dot{m}_f and O/F amplifies by roughly a factor of $(1 + O/F)$, which for the present GOX/ABS hardware is about 3 (running O/F settles at about 2, as reported in the companion campaign-performance analysis [12]). Propagation onward into c^* and C_F through CEA is mild, since both quantities vary slowly with mixture ratio near optimum O/F . A 1% error in A_t therefore produces roughly 1% error in F and \dot{m} , roughly 5% error in \dot{m}_f and O/F , and 1% or less in c^* and C_F .

These propagation factors define a second threshold: the minimum fractional area change above which an erosion correction makes a meaningful difference to the internal-ballistics analysis. The relevant noise floors in the present measurement chain are roughly $\pm 0.25\%$ on the thrust signal, $\pm 0.5\%$ on the chamber-pressure signal, and on the order of 1–2% on the venturi-derived oxidizer mass flow. Below approximately $\Delta A_t/A_{t,\text{pre}} \approx 1\text{--}2\%$, the throat-area contribution to the propagated quantities is comparable to or smaller than these noise floors, so ignoring erosion in the internal-ballistics reduction introduces a bias indistinguishable from the measurement noise.

Both thresholds—the reconstruction-reliability floor and the internal-ballistics-impact floor—sit near $\Delta A_t/A_{t,\text{pre}} \approx 2\%$. That coincidence is the saving grace of the algorithm. On firings where the algorithm cannot fully resolve the in-burn erosion shape, the erosion is also small enough that ignoring it in the downstream reduction introduces an error below the noise floor of the directly measured quantities. The algorithm is therefore most reliable on exactly the firings where the result matters most, and is least reliable on the firings where the result has the smallest impact on the broader analysis. The PRL dataset, with $\Delta A_t/A_{t,\text{pre}}$ between 31% and 103% and $R^2 \geq 0.996$ on every firing, sits firmly above this threshold and demonstrates that the reconstruction quality continues to improve monotonically as the erosion signal grows—there is no upper bound to the reliable-detection range that the present validation envelope has revealed.

4.2. Correlation of Chamber Pressure with Erosion

The relationship characterized in this section is a correlation, not a causal law: chamber pressure is treated as a convenient summary correlate of the underlying thermochemical drivers identified in Section 4.3, and the trends below are described as pressure-correlated erosion behavior. The CDA and PRL datasets are reduced and reported separately; they are brought together here only to characterize this correlation across the combined chamber-pressure envelope. The four PRL firings extend the pressure range above the CDA envelope and corroborate the CDA trend—they are not pooled with the CDA data into a single regression.

Across the CDA dataset, the data show a clear separation between the high-pressure firings and the medium- and low-pressure firings. The high-pressure CDA group reports a mean boundary-metrology Δd_t of 0.374 mm and a mean mid-burn rate-from-fit of 0.111 ± 0.046 mm/s ($n = 8$). The medium- and low-pressure CDA groups sit at 0.052 mm and 0.034 mm respectively, with single qualifying rate-from-fit values of 0.055 mm/s on CDA Test 26 (medium) and 0.027 mm/s on CDA Test 17 (low). The two metrics, derived independently from the same firing record, agree on a factor-of-four separation between high pressure and the lower-pressure groups. The medium- and low-pressure populations

are nearly indistinguishable from each other. CDA Test 26 is the only firing in the medium-pressure group that produced measurable erosion above the algorithm's reliable-detection floor, and without that one outlier the medium-pressure group mean would sit within 0.02 mm of the low-pressure group.

The PRL dataset extends the chamber-pressure coverage to higher peak pressures and provides four additional independently-reconstructed rate measurements at chamber pressures above the high end of the CDA range. The mid-burn instantaneous diameter erosion rates from the PRL reconstructions are 0.43 mm/s (PRL Test 15, peak $P_c \approx 2.3$ MPa), 0.72 mm/s (PRL Test 10, peak $P_c \approx 3.4$ MPa), 0.75 mm/s (PRL Test 18, peak $P_c \approx 3.4$ MPa), and 0.81 mm/s (PRL Test 17, peak $P_c \approx 3.3$ MPa). These values are 4 to 7 times the high-pressure CDA mean of 0.111 mm/s, against a peak-chamber-pressure ratio of approximately 1.5 to 2. The variation of erosion rate with chamber pressure above the CDA high-pressure point is therefore strongly nonlinear, with erosion rate climbing much faster than chamber pressure as P_c rises beyond the CDA's nominal 1.72 MPa (250 psi) target.

Combining the two datasets, the pattern across the full envelope looks more like a knee-point or threshold function than a continuous power-law dependence. Below approximately 0.9 MPa (the CDA medium- and low-pressure points), erosion sits at or below the metrology resolution on essentially every firing. At approximately 1.7 MPa (the CDA high-pressure point), erosion is measurable and reaches roughly 0.4 mm of diameter change over a 6-s burn, with a mid-burn rate near 0.11 mm/s. Above 2 MPa (the PRL firings), erosion rate climbs rapidly to 0.4–0.8 mm/s and total diameter change reaches 1.6–3.5 mm over the same 6-s burn duration. The transition between the low-erosion regime below 0.9 MPa and the rapid-erosion regime above 2 MPa is bracketed by the CDA high-pressure point at 1.72 MPa, but the present datasets do not finely resolve the transition shape between 0.9 and 1.7 MPa. A follow-on test series with a finer chamber-pressure grid across that band would be required to characterize whether the apparent threshold is sharp or gradual.

Statistical testing on the CDA dataset alone is consistent with the threshold interpretation. A Kruskal–Wallis rank-sum test on Δd_t across the three CDA groups returns $H = 13.7$ on two degrees of freedom ($p < 0.002$), confirming that the three populations are not drawn from a single underlying distribution. Pairwise Mann–Whitney comparisons separate the high-pressure group from both lower-pressure groups at $p < 0.005$. The medium- and low-pressure groups, however, are not rank-separable from each other at the available sample size, which fits the picture of both populations sitting below the chamber-pressure threshold for measurable erosion.

The practical implication of the pressure–erosion correlation is large. At the CDA high-pressure operating point, the graphite nozzle is approaching its geometric tolerance limit at six seconds: the mean fractional throat-area change at high pressure is approximately 9% over the 6-s burn, and the deepest-erosion firing (CDA Test 14) recorded a 17% area change. At the CDA medium- and low-pressure points, the throat-recession contribution to performance bookkeeping sits at or below the resolution of routine post-test metrology, and the firings retain at least an order of magnitude of remaining burn-duration margin. At the higher PRL chamber pressures, the throat-area change reaches 31% to 103% over the same burn duration, well past any reasonable geometric tolerance for sustained operation. An analyst applying a single average erosion rate to a multi-operating-point campaign would systematically misrepresent both ends of the envelope.

The chamber-pressure correlation captured here does not, on its own, explain the variability seen in the data. The high-pressure CDA group exhibits a within-group spread of more than a factor of eight in Δd_t despite the achieved pressures sitting within a 0.09 MPa band, and CDA Test 26 sits well above the rest of the medium-pressure group at essentially

the same nominal chamber pressure. These observations point to secondary thermochemical drivers that chamber pressure cannot fully resolve, which Section 4.3 addresses.

4.3. Secondary Drivers and Within-Group Variability

Chamber pressure is not the only thing driving erosion in these datasets. The CDA high-pressure group exhibits a within-group spread of more than a factor of eight in Δd_t , from 0.08 mm on CDA Test 15 to 0.69 mm on CDA Test 14, despite all eight firings sitting within a 0.09 MPa achieved-pressure band. That spread cannot be attributed to chamber pressure. The coefficient of variation on Δd_t at the high-pressure CDA operating point works out to approximately 47%, large enough that a single-firing assessment can produce an erosion measurement that differs by a factor of two from the population mean. The PRL firings show a comparable spread: PRL Test 15 reaches 1.64 mm diameter change at peak $P_c \approx 2.3$ MPa, while PRL Test 17 reaches 2.61 mm at peak $P_c \approx 3.3$ MPa, and PRL Test 10 reaches 3.53 mm at peak $P_c \approx 3.4$ MPa. The PRL ordering with peak chamber pressure is monotonic, but the rate differences are larger than the chamber-pressure differences alone would predict.

The thermochemistry already discussed in Section 1.2 explains why chamber pressure alone cannot close the picture. Graphite erosion is fundamentally a thermochemical surface reaction, and the actual drivers are the wall temperature and the partial pressure of the attacking species (H_2O , CO_2 , and OH) at the throat wall. Chamber pressure correlates with both. Higher P_c raises wall heat flux and therefore wall temperature, and it raises species partial pressure at the wall in proportion to the bulk density. But P_c is not itself a driver. It is a convenient single variable that happens to move both of the actual drivers in the same direction, which is why it serves as a useful summary correlate of erosion across operating points.

Several other variables influence wall temperature and species concentration at the throat, and these were not held constant firing-to-firing within either dataset. The companion CDA performance analysis [12] provides concrete numbers that constrain which of those variables are the dominant contributors. The instantaneous oxidizer-to-fuel ratio O/F sets the combustion temperature and the equilibrium mole fractions of the attacking species. In a hybrid configuration, O/F at any instant is a function of the running fuel mass flow, which is itself set by the regression rate of the ABS grain and the local heat flux into the port wall. The companion analysis reports a burn-to-burn O/F dispersion of approximately $\pm 16\%$ within the high-pressure CDA group at fixed nominal chamber pressure, the largest dispersion of any directly measured campaign parameter and substantially larger than the $\sim 5\%$ dispersion on thrust, chamber pressure, and oxidizer mass flow [12]. Combustion efficiency η^* , by contrast, was reported at approximately 98–100% across the related USU PRL ambient characterization tests of the same hardware [12], which effectively rules combustion efficiency out as a significant within-group driver.

The motor also exhibits a characteristic lean-to-rich O/F shift over the burn, with O/F starting near 2.5 and settling to approximately 1.8–2.0 by mid-burn [12], so the time-history of O/F at the throat differs from its mean even within a single firing. The CDA campaign was designed to vary P_c as the independent variable; O/F was targeted to be nominally constant within each operating-point group through fuel-grain length selection [12], but the measured burn-to-burn O/F dispersion shows that this control was incomplete. The within-group Δd_t spread on the high-pressure CDA firings is consistent with O/F as the dominant secondary driver and combustion efficiency as a negligible contributor.

The PRL dataset adds a complementary observation. The four PRL firings used helical-port fuel-grain geometry (4-turn on PRL Test 10, 2-turn on PRL Tests 15, 17, and 18), distinct from the straight-port grain used across the CDA campaign. The helical geometry

changes the regression-rate response of the grain to oxidizer mass flux, which in turn affects O/F and the chamber-pressure trajectory, but it does not change the throat-wall thermochemistry that drives erosion. The reconstruction algorithm produced consistent, high-confidence linear fits ($R^2 \geq 0.996$) on all four PRL firings despite the geometric variation. Read together with the high-confidence fits across the CDA grain configurations, this is direct evidence that the reconstruction is insensitive to fuel-grain geometry and that the algorithm's performance scales with the erosion signal magnitude rather than with the hardware configuration. It also reinforces that the secondary-driver picture in the CDA data is a property of the operating point, not of the reconstruction.

The practical consequence is that design margins for graphite-nozzle hardware operating at the upper end of the laboratory envelope should not be set against a single-firing erosion measurement. Multiple firings per operating point are required to populate the spread, and a campaign that depended on a single firing at high P_c could under- or over-predict the population-mean erosion by roughly a factor of two. The reconstruction algorithm recovers the time-resolved erosion trajectory on each firing and flags which firings sit at the upper end of their group, but the firing-to-firing scatter itself is a property of the test campaign rather than of the algorithm. Closing the gap between the present empirical fit and a deeper mechanistic understanding would require a follow-on campaign that varies O/F independently of P_c , ideally with instrumented wall-temperature or chamber-temperature measurements.

4.4. Limitations and Sources of Uncertainty

The reconstruction model rests on the thrust-coefficient relationship of Equation (1), which carries the standard assumptions of one-dimensional inviscid expansion through the nozzle, a stagnation pressure recovered at the head-end transducer location, and a quasi-steady momentum balance in which the unsteady thrust contribution from the eroding-throat mass injection is neglected. Each of these assumptions deserves a brief assessment of its quantitative contribution to the reconstructed throat-area history. The one-dimensional assumption is the same one underlying the published C_F -based reduction of essentially every small-motor ground test. Its dominant departure is non-uniform throat heating and the resulting non-axisymmetric erosion, which is not separable from the pin-gauge-anchored reduction in the present formulation. The pin gauge measures the throat at a single angular position and therefore averages over any azimuthal variation by construction. Post-test inspection of the recovered throats from both datasets is consistent with axisymmetric erosion at the resolution of the pin gauge, but a campaign that sustained strong azimuthal asymmetry in throat temperature would require an explicit two-dimensional reconstruction and is outside the validation envelope of the present model. The stagnation-pressure assumption is supported by the head-end mounting location of the chamber-pressure transducer and by the low subsonic Mach number of the chamber gas at the transducer port; the dynamic-pressure correction at the port is below 1% of P_c for the entire campaign envelope and is absorbed into the affine scaling along with the thrust-coefficient bias. The quasi-steady assumption is the most consequential of the three. The eroding-throat surface acts as a distributed mass source that perturbs the local momentum and energy balance at the nozzle entrance. The mass-injection rate at the highest-erosion firing of either dataset (PRL Test 10, mid-burn diameter rate approximately 0.72 mm/s) corresponds to a graphite mass-loss rate of order 5×10^{-4} kg/s, still approximately two orders of magnitude below the propellant mass flow rate at the same instant. That ratio is below the threshold at which the mass injection could perturb the thrust-coefficient relationship at the resolution of the present measurement chain.

The boundary throat metrology supplies the two scaling anchors of the algorithm and therefore sets the absolute-level uncertainty of the reconstruction. The expandable pin gauge used in both datasets, verified against calibrated calipers, resolves throat diameters to approximately ± 0.01 mm, which corresponds to roughly ± 0.13 mm² at the nominal 8.45 mm throat. That is tight enough that the boundary anchors are not the limiting factor for the low-erosion CDA firings: the measured Δd_t values on those tests (0.02 to 0.06 mm) sit several times above the instrument precision, even though they fall below the detection threshold $\Delta A_{\text{thr}} = 0.5$ mm² used in the changepoint-gate logic of Section 2.1.3. The threshold itself is set by the noise floor of the F/P_c reconstruction, not by the metrology. The pin gauge does have one important limitation. It captures the throat inner diameter at a single angular position and cannot resolve azimuthal asymmetry. If the throat erodes oblong rather than uniformly, the recorded diameter is an effective axisymmetric value rather than a true geometric mean of the inner-diameter trace. Post-test inspection of the recovered throats is consistent with axisymmetric erosion at the resolution of the pin gauge, but a campaign that sustained a strong azimuthal asymmetry—for example from a swirl-induced hot spot or from non-axisymmetric grain-port geometry—would require a higher-dimensional boundary measurement such as optical profilometry, a coordinate-measuring-machine probe, or a three-dimensional (structured-light) scan of the recovered nozzle. We recommend such higher-resolution post-test characterization as standard practice for future campaigns, both to quantify any non-axisymmetric recession or ovalization that the single-axis pin gauge cannot resolve and to support extension of the reconstruction to a two-dimensional throat model. The algorithm is structured to accept any such improved boundary measurement as a drop-in replacement of the pin-gauge value without modification of the reconstruction pipeline.

The thrust-stand and pressure-transducer calibrations contribute the second-order uncertainty terms. The CDA thrust stand was calibrated against a NIST-traceable load cell at the start of the test series and was not recalibrated between firings; pre-firing tare-load comparisons established that the calibration drift across the campaign was bounded by the manufacturer's combined linearity, hysteresis, and repeatability specification of approximately $\pm 0.25\%$ of full scale. The piezoresistive chamber-pressure transducer carries a comparable manufacturer-stated combined accuracy of $\pm 0.5\%$ of full scale. The PRL test cart uses comparable load-cell and pressure-transducer hardware with similar manufacturer-stated accuracies. Both contributions enter the unscaled inverse F/P_c multiplicatively. In the absence of erosion they would produce a constant offset that the affine map absorbs entirely into α , and in the presence of erosion they produce a slowly-varying drift that the affine map likewise absorbs along with the C_F drift. The transducer-bandwidth contribution to the lag estimate is the more consequential coupling: the rise-time-based masking interval of Equation (10) accommodates the dominant lag mode, but small per-firing variations in the lag cannot be calibrated from the rise event alone and contribute a residual jitter in the placement of t_{rise} . The 100-ms settling buffer is conservative against this jitter at the campaign-typical 6-s burn duration, but it would consume a non-negligible fraction of the trustworthy interval in shorter-duration tests. Firings of duration less than approximately 2 s are outside the validation envelope of the present masking strategy and would require either a faster pressure transducer or an explicit cross-correlation lag estimator.

The applicability envelope of the reconstruction algorithm itself is broad. The decoupled formulation of Section 2.1 explicitly avoids the propellant-thermochemistry assumptions that would otherwise tie the algorithm to a specific chemistry or operating regime, and the two-anchor affine scaling absorbs the thrust-coefficient bias regardless of its absolute value. The combined CDA and PRL datasets exercise the algorithm across roughly two orders of magnitude in fractional throat-area change (0% to 103%), peak chamber pressures

from 0.5 to 3.4 MPa, straight and helical fuel-grain configurations, and ambient and vacuum test-cell environments. The reconstruction succeeded on every firing in both datasets, with the only special handling required being the forced-onset path of Section 2.1.3 for PRL Test 10, on which erosion began within the ignition transient and the standard 1.5-s changepoint search margin was too long. The algorithm is therefore expected to work for any rocket motor whose performance is captured by the thrust-coefficient relationship $F = C_F P_c A_t$, regardless of propellant combination, hardware scale, burn duration, or operating pressure. Two real limitations remain. The first is the minimum erosion magnitude required for the reconstruction to be quantitatively reliable, which sits near $\Delta A_t / A_{t,pre} \approx 2\%$ as established in Section 4.1. Firings with smaller erosion still pass through the algorithm cleanly, but their in-burn reconstruction is dominated by F/P_c noise rather than by a real erosion signal. As also discussed in Section 4.1, the corresponding contribution to downstream internal-ballistics analysis is small, so the algorithm's reduced fidelity at small erosion magnitudes is unlikely to be a binding limitation in practice. The second limitation is geometric: the algorithm assumes an axisymmetric throat and operates on scalar throat area in the inverse calculation. It begins to fail when the throat cross-section departs significantly from a circle, such as in nozzles that develop strongly asymmetric heating during the firing or in throats whose post-test geometry is non-circular. Outside of those two limits, the algorithm is robust across propellant combinations, motor scales, burn durations, and chamber-pressure ranges.

Table 9 places the present method in the context of the other approaches in current use for reducing graphite-nozzle hot-fire data, summarizing the inputs each requires, whether it couples the erosion reconstruction to the in-test internal ballistics, and the resulting advantages and limitations.

Table 9. Comparison of throat-erosion data-reduction approaches for graphite-nozzle hot-fire tests.

Approach	Required Inputs	Couples Erosion with Internal Ballistics?	Advantages/Limitations
Fixed-throat-area reduction	F, P_c (A_t held constant)	N/A—no reconstruction	Simplest; biases c^* , η , and I_{sp} by attributing erosion-driven pressure and thrust loss to efficiency or O/F errors; unsuitable when erosion is significant.
Prescribed/averaged erosion	Pre- and post-test throat, or an assumed rate model	None—erosion imposed, not inferred	Gives total or mean recession; masks transient onset and rate; the forward-model variant conflates model error with measurement.
Conventional coupled-ballistics (C_F)	F, P_c , CEA, regression-rate model	Full—co-solves A_t , γ , O/F , η	Physically complete; iterative and model-dependent; erosion uncertainty entangled with the quantity being inferred.
Characteristic-velocity NTRT [9,10]	\dot{m}_{ox} , P_c , F , fuel mass, CEA, final radius/exit P	Full—co-solves throat radius and O/F	Strong physical grounding; recovers O/F history; validated across scales; still couples erosion to internal ballistics and CEA/ η uncertainty.
Present decoupled reconstruction	F, P_c , pre- and post-test throat anchors	None—throat-area inverse decoupled	No thermochemistry or regression inputs; shape-preserving affine scaling isolates the chamber-pressure correlation; requires two direct throat measurements and yields geometry and rate, not O/F .

5. Conclusions

This work developed and exercised a measurement-driven reconstruction model for the transient throat geometry of an eroding graphite nozzle. The algorithm computes the unscaled inverse $\hat{A}_t(t) = F(t)/P_c(t)$, masks the ignition-lag transient using a rise-time-based settling interval defined by the slower of the load-cell and pressure-transducer signal chains, identifies the erosion-onset time by piecewise-linear changepoint regression with a parabolic right-segment alternative selected when the residual reduction warrants it, and applies a two-anchor affine map that scales the unscaled signal onto the directly measured pre- and post-test throat areas. The structural property motivating the formulation is that the unscaled inverse differs from the true throat area by the slowly-varying thrust coefficient

alone, so onset detection, lag masking, and tail-off-initiation detection can be performed on the unscaled signal and are immune to thermochemistry and internal-ballistics uncertainty by construction.

The reconstruction was exercised on two complementary hot-fire datasets that together span roughly two orders of magnitude in fractional throat-area change and peak chamber pressures from 0.5 to 3.4 MPa. The first dataset is eighteen firings from the NASA MSFC CDA validation campaign, in which 75-mm GOX/ABS hybrid hardware was fired at three chamber-pressure operating points of 1.86, 0.87, and 0.54 MPa with operating-point-specific fuel grain lengths chosen to target $O/F \approx 2.0$ in every group. The second dataset is four high-pressure firings from the earlier USU PRL ambient test series, in which the same 75-mm hardware was exercised on helical-port fuel grains at peak chamber pressures from 2.3 to 3.4 MPa. The CDA boundary-metrology total throat-diameter change averaged 0.374, 0.052, and 0.034 mm at the three operating points; on the ten CDA firings that cleared the inter-anchor signal-to-noise test, the mid-burn diameter erosion rate averaged 0.111 ± 0.046 mm/s at the high-pressure point and reported single qualifying values of 0.055 mm/s and 0.027 mm/s at the medium- and low-pressure points respectively. The PRL reconstructions returned mid-burn erosion rates of 0.43 to 0.81 mm/s, a factor of four to seven above the CDA high-pressure mean. The boundary-anchored and rate-from-fit metrics, derived from independent reductions of the same firing record, agreed on the chamber-pressure ordering and on the magnitude of the chamber-pressure effect in every case. The combined data are more consistent with a knee-point or threshold behavior in chamber pressure than with a continuous power-law scaling: the CDA medium- and low-pressure groups sit at or below the metrology resolution, the CDA high-pressure group reaches an order of magnitude above them, and the PRL points climb rapidly with peak chamber pressure once above the CDA envelope. A follow-on test series with a finer chamber-pressure grid between 0.9 and 1.7 MPa would be required to resolve the transition shape.

The practical implication is that the throat-recession contribution to the apparent chamber-pressure history, C^* efficiency, and specific impulse on graphite-nozzle firings at the upper end of the laboratory envelope is large enough to dominate the post-test reduction. The fixed-throat-area reduction still customary in many small-motor data pipelines therefore systematically attributes throat erosion to combustion-efficiency and mixture-ratio departures that did not occur. The reconstruction algorithm developed here removes that bias at the cost of two boundary throat measurements per firing and one rise-time-aware analysis pass. The within-group coefficient of variation of approximately 47% on Δd_t at fixed nominal chamber pressure in the CDA dataset is large enough that adequate margin assessment requires multiple firings per operating point rather than single-firing assessments.

The ability to resolve the erosion onset time, curve fit the erosion rate, and two-anchor affine scaling with high reliability and accuracy is dependent on the amount of erosion observed. This algorithm begins to struggle to resolve changes below approximately 2% increase in throat area. Above this threshold, onset time, erosion-rate fitting, and scaling resolve with high consistency, including on the PRL firings where $\Delta A_t / A_{t,pre}$ exceeded 100% and the changepoint-regression R^2 exceeded 0.996 on every firing. A 2% change in nozzle area not properly reconstructed propagates to an uncertainty in internal-ballistics calculations below approximately 5%, comparable to the noise floor of the directly measured thrust and chamber-pressure signals. The reconstruction-reliability floor and the internal-ballistics-impact floor therefore coincide, which is what makes the technique useful in practice: the algorithm is most reliable on the firings where the

result matters most, and least reliable on the firings where the result has the smallest downstream impact.

The PSI campaign is continuing at the NASA Langley Research Center 60-foot vacuum sphere, where the same hybrid motor hardware is being exercised against lunar-regolith-analogue targets at flight-relevant ambient pressure. The reconstruction algorithm developed here will be applied to the LaRC firing records as they become available, both as a direct extension of the chamber-pressure dependence reported above and as a calibrated input to the plume-surface-interaction analyses that are the program's primary deliverable.

Author Contributions: Conceptualization, R.J.T.; methodology, R.J.T.; software, R.J.T.; validation, R.J.T.; formal analysis, R.J.T.; investigation, R.J.T.; resources, S.A.W.; data curation, R.J.T. and S.A.W.; writing—original draft preparation, R.J.T.; writing—review and editing, R.J.T. and S.A.W.; visualization, R.J.T.; supervision, S.A.W.; project administration, S.A.W.; funding acquisition, S.A.W. All authors have read and agreed to the published version of the manuscript.

Funding: This development was supported in part by the NASA Plume-Surface Interaction (PSI) project at Marshall Space Flight Center, conducted under the Human Landing System (HLS) program in support of the Artemis lunar campaign, NASA Contract Number 80NSSC20M0083.

Data Availability Statement: The data presented in this study are available on request from the corresponding author. The data are not publicly available due to ongoing project use restrictions.

Acknowledgments: The authors gratefully acknowledge the staff of the NASA Marshall Space Flight Center Component Development Area for test support during the PSI hot-fire campaign. During the preparation of this manuscript, the authors used AI LLM Claude Opus 4.7 to improve the language, flow, and structural organization of the text. Claude was also used to simplify and improve the structure of the original code to be polished for this publication. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Nomenclature

The following symbols, subscripts, accents, and abbreviations are used in this manuscript.

Roman Symbols

A_e	Nozzle exit area	m^2
A_t	Nozzle throat area (corrected reconstruction)	m^2
\hat{A}_t	Unscaled reconstructed throat area, $\hat{A}_t = F/P_c$	m^2
\tilde{A}_t	Low-pass-filtered unscaled throat-area history	m^2
$A_{t,\text{pre}}$	Pre-test throat area from pin-gauge measurement	m^2
$A_{t,\text{post}}$	Post-test throat area from pin-gauge measurement	m^2
$\bar{A}_{t,\text{pre}}^{\text{recon}}$	Median of \hat{A}_t on the pre-erosion plateau anchor window	m^2
$\bar{A}_{t,\text{post}}^{\text{recon}}$	Median of \hat{A}_t on the late-burn anchor window	m^2
\dot{A}_t	Throat-area erosion rate	m^2/s
C_F	Thrust coefficient, $C_F = F/(P_c A_t)$	–
$C_{F,\text{id}}$	Ideal thrust coefficient (Sutton–Biblarz)	–
c^*	Characteristic exhaust velocity	m/s
d_e	Nozzle exit diameter	m
$d_{t,\text{pre}}, d_{t,\text{post}}$	Pre-test and post-test throat diameters	m
F	Axial thrust ($=L_x$ from the 6-DOF thrust stand)	N
L_x, L_y, L_z	Force components measured by the 6-DOF thrust stand	N
M_x, M_y, M_z	Moment components measured by the 6-DOF thrust stand	$\text{N}\cdot\text{m}$
f_c	Low-pass-filter cutoff frequency	Hz

f_s	Data-acquisition sample rate	Hz
f_{rise}	Rise-time fractional threshold (taken as 0.95)	–
f_{toi}	Tail-off-initiation fractional threshold	–
I_{sp}	Specific impulse	s
k_σ	Onset-detection confidence multiplier	–
k_{grad}	Tail-off-detection confidence multiplier on \dot{P}_c and \dot{F}	–
$\dot{P}_{c,plat}, \dot{F}_{plat}$	Plateau medians of \dot{P}_c and \dot{F}	Pa/s; N/s
$\sigma_{\dot{P}_{c,plat}}, \sigma_{\dot{F}_{plat}}$	Plateau MAD-based standard deviations of \dot{P}_c and \dot{F}	Pa/s; N/s
SNR_{min}	Minimum signal-to-noise ratio for two-anchor affine acceptance	–
M_e	Isentropic exit Mach number	–
\dot{m}_{ox}, \dot{m}_f	Oxidizer and fuel mass-flow rates	kg/s
\dot{m}	Total propellant mass-flow rate, $\dot{m} = \dot{m}_{ox} + \dot{m}_f$	kg/s
N	Number of samples in the data-acquisition record	–
O/F	Oxidizer-to-fuel mass ratio	–
P_c	Stagnation chamber pressure	Pa
P_e	Static pressure at the nozzle exit	Pa
P_∞	Ambient back pressure	Pa
D_t	Throat diameter	m
\dot{D}_t	Throat-diameter erosion rate	m/s
a	Intercept (continuity value) of the changepoint fit at t^*	m ²
b_1	Pre-erosion plateau slope of the changepoint fit	m ² /s
b_2	Post-onset linear slope (linear-right changepoint model)	m ² /s
c_2	Post-onset curvature coefficient (parabolic-right changepoint model)	m ² /s ²
t	Time	s
t^*	Erosion-onset time	s
t_{eo}	Engine-off time	s
<i>Greek Symbols</i>		
α	Multiplicative coefficient of two-anchor affine map	–
β	Additive coefficient of two-anchor affine map	m ²
$\Gamma(\gamma)$	Vandenkerckhove function	–
γ	Specific-heat ratio of combustion products at the throat	–
$\delta(t)$	Deviation of $\tilde{A}_t(t)$ from the plateau mean	m ²
ΔD_t	Total throat-diameter change over the firing	m
Δt_{anchor}	Anchor-window width	s
<i>Subscripts and Accents</i>		
pre	Pre-test value	
post	Post-test value	
plat	Plateau-region quantity	
cp	Changepoint-fit quantity	
$\dot{}$	Time derivative	
$\hat{}$	Reconstructed quantity	
$\tilde{}$	Filtered quantity	
<i>Abbreviations</i>		
MAD	Median absolute deviation	
SNR	Signal-to-noise ratio	
DOF	Degrees of freedom	

References

1. Bianchi, D.; Nasuti, F. Numerical Analysis of Nozzle Material Thermochemical Erosion in Hybrid Rocket Engines. *J. Propuls. Power* **2013**, *29*, 547–558. [\[CrossRef\]](#)
2. Thakre, P.; Yang, V. Chemical Erosion of Carbon-Carbon/Graphite Nozzles in Solid-Propellant Rocket Motors. *J. Propuls. Power* **2008**, *24*, 822–833. [\[CrossRef\]](#)
3. Bianchi, D.; Leccese, G.; Nasuti, F.; Onofri, M.; Carmicino, C. Modeling of High Density Polyethylene Regression Rate in the Simulation of Hybrid Rocket Flowfields. *Aerospace* **2019**, *6*, 88. [\[CrossRef\]](#)

4. Klager, K. The Interaction of the Efflux of Solid Propellants with Nozzle Materials. *Propellants Explos. Pyrotech.* **1977**, *2*, 55–63. [[CrossRef](#)]
5. Acharya, R.; Kuo, K.K. Effect of Pressure and Propellant Composition on Graphite Rocket Nozzle Erosion Rate. *J. Propuls. Power* **2007**, *23*, 1242–1254. [[CrossRef](#)]
6. Sutton, G.P.; Biblarz, O. *Rocket Propulsion Elements*, 9th ed.; John Wiley & Sons: Hoboken, NJ, USA, 2016.
7. Gordon, S.; McBride, B.J. *Computer Program for Calculation of Complex Chemical Equilibrium Compositions and Applications I: Analysis*; Technical Report NASA RP-1311; NASA Lewis Research Center: Cleveland, OH, USA, 1994.
8. McBride, B.J.; Gordon, S. *Computer Program for Calculation of Complex Chemical Equilibrium Compositions and Applications II: Users Manual and Program Description*; Technical Report NASA RP-1311-P2; NASA Lewis Research Center: Cleveland, OH, USA, 1996.
9. Kamps, L.; Saito, Y.; Kawabata, R.; Wakita, M.; Totani, T.; Takahashi, Y.; Nagata, H. Method for Determining Nozzle-Throat-Erosion History in Hybrid Rockets. *J. Propuls. Power* **2017**, *33*, 1369–1377. [[CrossRef](#)]
10. Kamps, L.; Hirai, S.; Sakurai, K.; Viscor, T.; Saito, Y.; Guan, R.; Isochi, H.; Adachi, N.; Itoh, M.; Nagata, H. Investigation of Graphite Nozzle Erosion in Hybrid Rockets Using Oxygen/High-Density Polyethylene. *J. Propuls. Power* **2020**, *36*, 423–434. [[CrossRef](#)]
11. Bianchi, D.; Migliorino, M.T.; Rotondi, M.; Kamps, L.; Nagata, H. Numerical Analysis of Nozzle Erosion in Hybrid Rockets and Comparison with Experiments. *J. Propuls. Power* **2022**, *38*, 389–409. [[CrossRef](#)]
12. Whitmore, S.A.; Thibaudeau, R.J. Development of a GOX/ABS Hybrid-Gas Generator System for the NASA Plume Surface Interaction (PSI) Experimental Campaign. In *Proceedings of the AIAA SciTech 2026 Forum*; AIAA: Orlando, FL, USA, 2026. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.