

Progress on multifidelity simulations of a laser-ignited rocket combustor

By T. Zahtila, D. Passiatore AND G. Iaccarino

1. Motivation and objectives

The goal of the Stanford PSAAP-III Center (2020–2025) is to determine laser ignition reliability maps in subscale rocket combustors using multifidelity ensembles of numerical simulations (Wang *et al.* 2021). This report describes progress made on understanding the requirements for low- and high-fidelity simulations, which together compose the aforementioned multifidelity ensembles.

Uncertainties and variability are unavoidable in the operation of any engineering device. In the present context, despite the use of a subscale rocket combustor in a laboratory setting [Purdue Propulsion, Strelau *et al.* (2023)], laser-induced ignition is not a deterministic process, and reliability maps have been produced experimentally by repeating the same nominal experiment.

Similarly the target predictive simulations must account for the known source of uncertainty in the laboratory and propagate their effect to the output quantities of interest (QOIs). This process requires repeated simulations with carefully defined input perturbations to mimic the variability observed in the experiments. Therefore, the necessity of a multifidelity simulation strategy arises from (i) representing uncertainties that demand large numbers of simulations for determination of the QOI and (ii) the computational cost associated with the grid requirements for accurate numerical calculation of physical processes at play (Moin & Mahesh 1998). Specifically, the overarching aim of the bi-fidelity campaign is success/fail ignition maps for the combustor with robust confidence intervals.

In this work, low fidelity simply refers to aggressive coarsening of the combustor mesh and using simplified models to describe the physical phenomena (i.e., a reduced combustion mechanism). Therefore, it is expected that limited overall accuracy in low-fidelity simulations will be achieved. However, it is critical that these low-fidelity simulations preserve trends and correlations of the simulation results with respect to the imposed input variability. High-fidelity simulations, conversely, are required to be validated as physically accurate. To this end, the protocol for generation of a bi-fidelity set of results is to (i) run a large number of less expensive low-fidelity simulations spanning the expected range of uncertainty sources, which is typically accomplished via random sampling; (ii) construct a data matrix that collects the low-fidelity results in terms of columns (e.g., using time traces of the QOI in each low-fidelity simulation); (iii) perform a matrix interpolative decomposition to construct a reduced basis and an interpolation rule to reconstruct the full data matrix; (iv) run high-fidelity simulations corresponding to the inputs of the reduced basis; and (v) replace the low-fidelity reduced basis with the corresponding high-fidelity basis and generate a reconstructed data matrix that approximates a large number of high-fidelity simulations.

As mentioned earlier, a companion experimental campaign in an optically accessible rocket combustion chamber (Strelau *et al.* 2023) is used as validation for the present

simulations. Schlieren imaging and laser measurements are collected to facilitate the determination of the variance in the deposited energy kernel. Visualization of ignition events and flame growth are also obtained. These experiments also provide time series of pressure signals at probe locations near the wall of the combustion chamber. Validation in this sense strictly refers to a match between our high-fidelity simulation and experimental outputs; in other words, it is expected that low-fidelity simulations are unable to provide quantitative agreement with the measured data. The experiments also provide an indicative ignition reliability map; however, the low number of experimental trials that can be performed limits confidence in the resulting ignition probability estimate.

The model combustor experiments vary the nominal laser focal location in the stream-wise direction of the injector coflowing jet and the radial distance from the jet axis of symmetry. This is not an uncertainty, but rather a systematic study of the ignition likelihood in various regions in the combustor. It is easy to imagine that the physics of the deposited laser energy kernel will vary strongly between the jet potential core region, inner and outer regions and the shear layer. Two scenarios emerge. First, direct ignition may occur if the energy deposition site instantaneously spatially coincides with a favorable fuel-air equivalence ratio $\phi = (n_{fuel}/n_{oxidizer})/(n_{fuel}/n_{ox})_s$, where n denotes the number of moles of each chemical species and subscript s refers to the stoichiometric ratio. Second, even if the laser is fired at a location that is too lean, breakdown of the deposited energy kernel leads to an ejection jet that may lead to ignition via transport of hot gas (Wang 2020), which is termed indirect ignition.

The ignition probability for the direct ignition case is straightforward because in low- and high-fidelity simulations, ignition has occurred in every trial. Constructing a probability map for indirect ignition is more challenging as ignition success is typically $Pr(ign) \approx 0.5 - 0.7$. Thus, the focus of this brief is oriented toward requirements for the systematic determination of ignition probability for indirect ignition.

In generating a multifidelity data ensemble, of primary concern is degree of correlation between the low- and high-fidelity simulations. In the present application, there is a heightened sensitivity to the spatial structure of the turbulent state at laser deposition time. The intermittency of coherent structures and turbulent fluctuations translates to complex energy kernel-turbulence interactions. Accordingly, a part of this brief is dedicated to the challenges in matching turbulent states between fidelities. Accounting for the influence of instantaneous turbulence realization constitutes a significant extension of the bi-fidelity approximations to data sets collected until now (Doostan *et al.* 2016; Fairbanks *et al.* 2020).

2. Formulation and computational setup

2.1. Combustor computational model

The combustor computational schematic is shown in Figure 1. The essential features are an O_2 - CH_4 gas-gas coflowing injector jet, a developing recirculation region and a shear layer. Initially the combustor is quiescent and at atmospheric temperature and pressure when the injector commences. Thereafter, the features of an axisymmetric jet develop and are quasi-steady, but the recirculation region in the combustion chamber continues to develop for the duration of the pre-ignition phase, which is of time interval $\Delta t = 1000 \mu s$ in the present campaign. At this point, at a random snapshot of turbulence between $t = 1000 \mu s$ and $t = 1284 \mu s$, the laser is deposited at one of two sites, which are marked on the schematic. The simulation progresses, and probes in the computational domain




Figure 1. Streamwise cross section of the computational domain showing the combustion chamber, and coflowing injector. The red circles indicate the present laser deposition sites investigated. Dashed lines demarcate boundary between shear layer and outer region.

record the pressure time series in the combustor. In subsequent campaigns, the phases of the air-fuel mixture in the coflowing jet will be varied, as will the initial quiescent state inside the combustor.

2.2. Sources of uncertainty

Input uncertainties are a feature of any real-world system, and in the present campaign, we focus on a restricted set of input uncertainties. The uncertainty vector $\xi = (\xi_0, \xi_1, \dots, \xi_m)$ is currently composed of laser deposition uncertainties that are summarized in Table 1 with their respective probability distributions. The laser beam and computational discretization were described previously (Wang *et al.* 2021). However, for the benefit of present context, it is worth mentioning that ξ_0, ξ_1 simply refer to the difference in the nominal laser deposition site and the actual realized site.

The laser shape is modeled as a kernel with two lobes with radius R_1 , which is nearer to the shear jet axis of symmetry, and R_2 . The ratio of these lobe radii is captured by ξ_2 , where $R_2/R_1 = 1$ would be a symmetric kernel. The total energy deposited uncertainty, ξ_3 , is a function of this laser shape, as the varying the lobe radii changes the volume of the kernel, but also a laser intensity parameter that linearly scales the deposited energy. Finally, ξ_4 refers to a lag between the target laser deposition time and actual deployment.

The probability distributions are inferred from the experimental campaign measurements, for example, by comparing the measured and nominal spark locations. These data are first visually classified as either distributed according to a uniform or normal distribution, and subsequently a least squares fit determines the distribution. This procedure was followed for parameters ξ_0, ξ_1 and ξ_3 , for which experimental measurements are available. The experimental data are not without limitations, however, as (i) there are apparent biases in the measurements that vary with nominal focal location, and (ii) only a small number of repetitions are carried out at each site. Some aspects of the determined probability distributions in Table 1 remain difficult to interpret, such as why the bias (μ) in the streamwise (ξ_0) and radial (ξ_1) directions are similar in magnitude but the variance (σ^2) is much smaller in the streamwise direction.

For ξ_2 , which is uncertainty in the depositing lobe radii, the experiments are not able to measure this quantity directly. The corresponding uncertainty model is informed by a previous investigation into laser energy kernels (Wang *et al.* 2020) that elucidated the

Uncertainty ID	Description	Probability distribution
$\xi_0: \Delta x_l$	Imprecision in focal location, streamwise direction	$N \sim [0.29 \text{ mm}, 0.04 \text{ mm}^2]$
$\xi_1: \Delta x_r$	Imprecision in focal location, radial direction	$N \sim [-0.54 \text{ mm}, 0.20 \text{ mm}^2]$
$\xi_2: R1/R2$	Lobe radii ratio	$U \sim [1.1, 2.0]$
$\xi_3: E$	Total energy deposited	$U \sim [18 \text{ mJ}, 30 \text{ mJ}]$
$\xi_4: \tau_L$	Lag time	$U \sim [0 \text{ } \mu\text{s}, 284 \text{ } \mu\text{s}]$

Table 1. Description of the uncertainty parameters of the laser pulse for the indirect ignition case. Probability distributions are denoted as either $U \sim [a,b]$, uniform probability between a and b, or $N \sim [\mu, \sigma^2]$, normal distribution with mean μ and variance σ^2 .

relationship between ejection character with kernel shape. Further, ξ_4 , is not measured experimentally and although nominally associated with laser lag manifests as a shift in the instantaneous turbulent state in which the laser energy kernel is injected.

To ensure efficient and purposeful use of computational resources, it is imperative to determine whether input uncertainties propagate through to variation in the output QOI. As a preliminary work, 2^m low-fidelity simulations were run at the extrema of the ranges of the uncertainty parameter. It was found that variation in all uncertainty parameters had a non-negligible effect on pressure traces, primarily having the effect of a shift in the ignition time delay. Consequently, all uncertainty parameters were retained.

2.3. Low- and high-fidelity computational models

The Hypersonics Task-based Research solver described in Di Renzo *et al.* (2020) for compressible chemically reacting flows is used to carry out simulations. The computational details concerning the solver parallelism, simulation domain, boundary conditions and numerics can be found in previous work as part of PSAAP-III (Wang *et al.* 2021).

Two mesh resolutions are employed in this study. The low-fidelity mesh is composed of 10M grid points and the high-fidelity counterpart is composed of 220M grid points. The same subgrid-scale models are employed in both cases (Smagorinsky 1963). A fully resolved direct numerical simulation would require approximately 100 billion grid points by comparison, but presently the PSAAP project has achieved satisfactory agreement between measured quantities in experiments and output statistics from the present resolution used in high-fidelity simulations. As is typical, an added cost associated with an increased spatial resolution is a greater restriction placed on time-stepping. The chemistry in the low-fidelity simulations is simplified to a four-species, one-step mechanism (CERFACS 2023) as opposed to a twelve-species mechanism (Tao *et al.* 2018) employed in high-fidelity simulations. Consequently, the high-fidelity simulation is carried out on 96 GPU cards, with approximately one week of wall-clock time, compared to the low-fidelity simulations that take two days of wall-clock time on two GPUs (a cost ratio of approximately 150).

It is important to note that simulations are composed of pre- and post-ignition phases, and the pre-ignition, though expensive, only needs to be run once and stored. The laser energy is then deposited at various restart times from the stored data. In the pre-ignition phase, the coflowing jet commences injection of the O_2 and CH_4 coflow into an initially quiescent combustor domain. This proceeds for an interval of $\Delta t = 1000 \text{ } \mu\text{s}$, and the recirculation region can be seen developing continuously. As an additional exploratory case, one simulation is run with the low-fidelity mesh and the twelve-species mechanism.

Set	Symbol	Mesh fidelity	Chemistry (no. of species)	Nominal laser radial deposition site (mm)	No. of runs
LF-DI	●	Low	4	0.0	25
LF-IDI	●	Low	4/12	9.0	50
HF-DI	■	High	12	0.0	1
HF-IDI	■	High	12	9.0	4

Table 2. Description of the different simulations conducted. Nominal laser streamwise distance for all simulations is 19.05 mm. LF denotes low fidelity, HF denotes high fidelity, DI denotes direct injection and IDI denotes indirect injection.

Table 2 details the generated data sets featured in this brief. With reduced chemistry and a lower mesh fidelity, an order of magnitude more simulations have already been obtained for the low-fidelity models, which does not represent full-capacity production runs at this stage. As was noted in the cost ratio across fidelities, the capability currently is to generate two orders of magnitude higher number of simulations, but the present aim is to still further improve this cost ratio.

2.4. Bi-fidelity approximation via high-fidelity basis and low-fidelity interpolation

The QOI in the present study is chosen as the pressure trace obtained by averaging six locations in the computational domain that match probe sites in the experimental setup (we assume that m time samples are considered). The bi-fidelity procedure commences with construction of an ensemble of pressure traces corresponding to n different realizations of the input uncertainties $\boldsymbol{\xi}$, i.e., a random realization of the input uncertainties is generated and the corresponding simulation is carried out. The pressure trace obtained is then collected into one column of a data matrix

$$\mathbf{U}^L = \{u_1, u_2, \dots, u_N\}, \quad (2.1)$$

where \mathbf{U}^L is an $m \times n$ matrix. At this point, an interpolative decomposition algorithm factorizes the data matrix into basis vectors and a correlation matrix

$$\mathbf{U}^L \approx \mathbf{U}_C^L \mathbf{C}^L, \quad (2.2)$$

where \mathbf{U}_C^L is an $m \times r$ matrix referred to as the column skeleton of \mathbf{U}^L , and \mathbf{C}^L is an $r \times n$ matrix that contains the interpolation coefficients to approximate \mathbf{U}^L . r is the estimated rank of \mathbf{U}^L . The construction of the matrix interpolative decomposition requires the specification of a reconstruction accuracy ϵ for a pivoted rank-revealing QR factorization (Doostan *et al.* 2016). This is the controllable parameter that the user can use to determine their preferred trade-off between accuracy and speed. The central idea is to swap the low-fidelity column skeleton \mathbf{U}_C^L with high-fidelity counterparts

$$\begin{array}{ccc} \text{Output} & \text{Input } (r \text{ HF simulations}) & \\ \downarrow & \downarrow & \\ \mathbf{U}^{BF} & = \mathbf{U}_C^H \mathbf{C}^L. & \end{array} \quad (2.3)$$

The important details are that the rank r of the matrix interpolative decomposition is typically an order of magnitude lower than the number of low-fidelity simulations n and

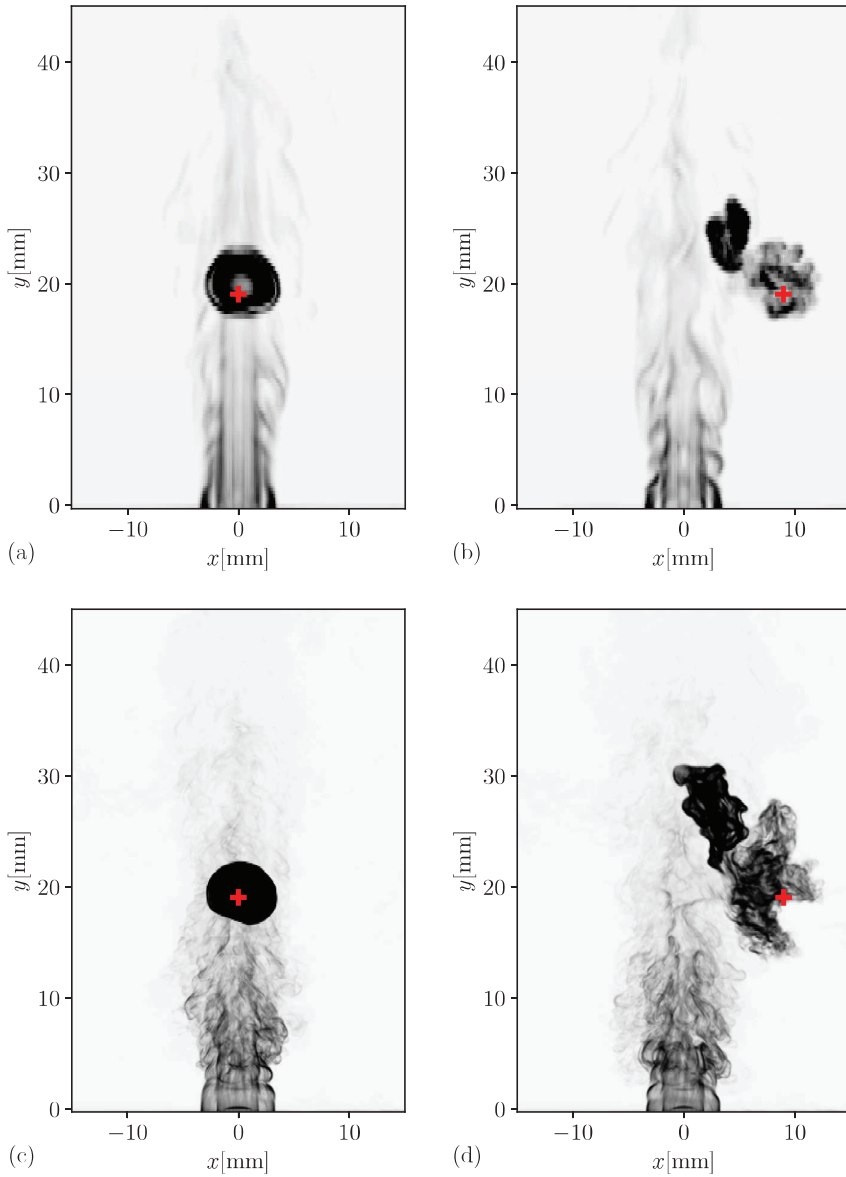


Figure 2. Computational Schlieren of two ignition modes for two-fidelities: (a) low-fidelity direct ignition, (b) low-fidelity indirect ignition, (c) high-fidelity direct ignition and (d) high-fidelity indirect ignition. Red markers indicate the nominal laser energy deposition site.

so the computational cost to construct the bi-fidelity data matrix \mathbf{U}^{BF} is substantially reduced compared to running many high-fidelity simulations. The bi-fidelity data matrix \mathbf{U}^{BF} represents an approximation to the higher-fidelity data matrix \mathbf{U}^H , where basis vectors from a low number of high-fidelity runs are guided by the low-fidelity correlation matrix.




Figure 3. Analysis of low-fidelity data. (a) Normalized singular values of pressure traces for direct and indirect ignition scenarios. Truncated singular value decomposition reconstruction of the pressure traces obtained from low-fidelity simulation output for both direct and indirect ignition. Number of leading singular value decomposition modes used: (b) $N = 1$, (c) $N = 3$ and (d) $N = 4$. LF denotes low fidelity, HF denotes high fidelity, DI denotes direct injection and IDI denotes indirect injection.

3. Results

3.1. Computational Schlieren

Computational Schlieren has been developed as part of the PSAAP program in order to facilitate close comparison with experimental flow visualization. Here we focus on qualitative comparisons between ignition modes under investigation. In Figure 2, instantaneous snapshots of direct and indirect igniting cases are shown. The direct ignition case featured energy deposition site coordinates of $xy = (0.0, 19.05)$, and the indirect case corresponds to energy at $xy = (9.0, 19.05)$. In panels (a,b) of Figure 2, the snapshots are taken from low-fidelity simulations, which is why the images appear clearly coarse-grained. The darkened and saturated grayscale pixels in the images are associated with combustion and flame regions. It is observed that for the direct ignition case, combustion occurs near to the laser deposition site. However, for the indirect case, at the deposition site, the remnants of the energy kernel ejection can be seen and the combustion region is nearer to the jet axis of symmetry. In contrast to the early stage combustion kernel, the direct ignition energy kernel appears spherical as opposed to the irregular shape of the indirect case. Ignition occurring near to the laser deposition site and the more spherical combustion kernel shape are hints that the direct ignition case will have considerably less variability associated with its output statistics. Additionally, high fidelity output of igniting cases is shown in panels (c,d) of Figure 2 and there appears to be good agreement between low- and high-fidelity ignition kernel characteristics.




Figure 4. Comparison of the (a) low-fidelity and (b) high-fidelity pre-ignition density field taken at $t = 1025 \mu\text{s}$. The length of the co-flowing jet inviscid core for the low-fidelity case is marked on both images. This facilitates visual assessment of the artificially lengthened inviscid core in the low-fidelity mesh.

3.2. Singular value decomposition

Once the $m \times n$ low-fidelity data matrix \mathbf{U}^L has been assembled, there is no *a-priori* optimal rank (or error tolerance) to proceed with for the matrix interpolative decomposition. As was noted in a previous bi-fidelity approximation campaign (Fairbanks *et al.* 2020), the singular values of \mathbf{U}^L bound the error of reconstruction of \mathbf{U}^L via the matrix interpolative decomposition. Accordingly, the singular values of \mathbf{U}^L are plotted in Figure 3(a). For both the direct and indirect ignition scenarios, it is clear that the majority of the energy lies in the zeroth mode. Therefore, in panels (b,c,d), we plot the columns of \mathbf{U}^L (the raw pressure traces) and plot their reconstruction based off the leading $\sigma = 1, 3$ and 4 modes. Remarkably on the direct ignition reconstruction, the pressure traces are well captured by reconstruction based off a single mode. Actually, with the constraint of one mode-reconstruction, the profile may only be linearly modulated with an offset. By contrast, the indirect ignition scenario contains bifurcating non-igniting and igniting branches of solutions; consequently, the zeroth mode reconstruction offers a family of curves that are not observed in the raw pressure traces. In Figure 3(c), reconstruction based off the first three modes successfully splits the bifurcating ignition fates for the indirect scenario. Reconstruction based off the leading $\sigma = 4$ modes captures well the ignition fate, the ignition delay time (onset of rapid pressure rise) and thereafter the plateau of pressure.

3.3. Mesh fidelity, discretization of the laser energy kernel and ejection velocity

Figure 4 shows a comparison of instantaneous snapshots of the density field for the low- and high-fidelity simulations. There are key differences between the visualized flow fields. It is apparent that many scales are missing in the low-fidelity case, and the length inviscid core of the jet is dramatically overpredicted by the low-fidelity mesh.

The established picture of laser-induced energy kernels is rapidly ejected hot gas that can travel significant distances (Wang *et al.* 2020), but this understanding has been developed in the context of quiescent background flow. In the present scenario, there are shear layer effects arising from the proximity of the deposition site to the co-flowing jet, and there are also effects associated with the evolving recirculation zones in the




Figure 5. Variation in simulation mesh fidelity and kernel asymmetry. Left panels are taken after laser energy deposition has completed and right panels after $\Delta t = 100 \mu\text{s}$. Symmetric kernel on (a) high fidelity and (b) low fidelity mesh. Asymmetric kernel on (c) high fidelity and (d) low fidelity mesh.

combustor. In Figure 5, visualizations of both mesh resolutions first show the energy kernel soon after laser injection has been completed, and secondly after $\Delta t = 100 \mu\text{s}$ has passed and the motion of the jet ejecta has completed. Note, there are slight differences in the output time snapshots in low- and high-fidelity runs resulting in the slight change in the shock front location that is observed in the right panels of Figure 5. For reference, the laser energy deposition interval occurs on the fast time scale of order $\Delta t_{\text{dep}} = \mathcal{O}(100 \text{ ns})$. Figure 5 presents the result of two lobe radii ratios, $R_2/R_1 = 1.15$ and 1.98 , which constitute a near-symmetric and asymmetric kernels, respectively.

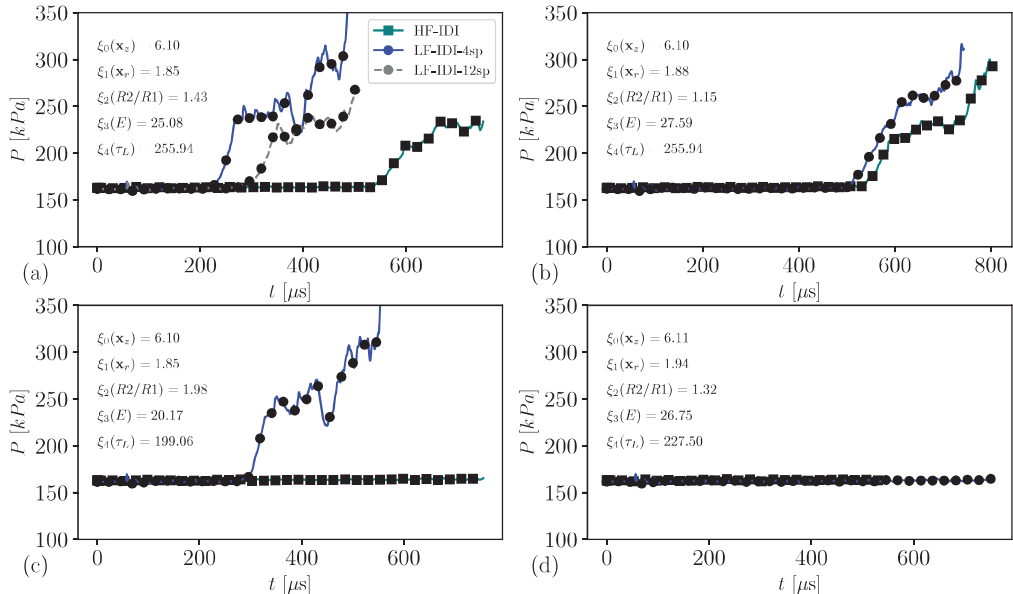


Figure 6. Four cases from the indirect ignition data matrix that are identified by the matrix interpolative decomposition procedure as skeleton cases of the low-rank approximation. The corresponding high-fidelity output is also plotted. The dashed grey line in (a) corresponds to the low-fidelity simulation with 12-species chemistry. LF denotes low fidelity, HF denotes high fidelity and IDI denotes indirect injection.

The symmetry or asymmetry of the kernel is well captured by the low-fidelity mesh at early time intervals after deposition, but there are visible artefacts due to under-resolution. Additionally, the laser-generated shock wave front propagates consistently across fidelities, as can be seen from the reflected shock waves in the later time instances. At later times, the characteristics of the hot gas kernel significantly diverge between mesh fidelities. For the high-fidelity results, the initially symmetric kernel with near-equal lobe radii results in a horizontal dumbbell shape, and there appears to be a limited ejection velocity. Conversely, the initially asymmetric kernel shows the ejected hot gas profile consistent with the quiescent core case, but it is clearly interrupted by velocity fluctuations from the shear layer.

When comparing the highly-resolved energy kernel with the low-fidelity result, the discrepancy in inviscid core length and scale information makes it impossible to distinguish between resolution effects on the kernel ejection mechanism and the fluid–kernel interaction, and this will be the basis of future targeted study. Nevertheless, in the low-fidelity simulations, there does not appear to be strong differences in the kernel shape at the latter inspected time, so the effect of kernel asymmetry does not seem to be properly captured.

3.4. Comparison of low- and high-fidelity QOI output

Figure 6 shows a comparison of the pressure traces obtained for the low-fidelity simulations for the cases identified by the matrix interpolative decomposition procedure; the corresponding high-fidelity simulation data are also plotted. The ignition outcomes can

be classified as either standard ignition (a,c), delayed ignition (b) or ignition failure (d). One-to-one agreement between low- and high-fidelity simulation output is not required for the present workflow, but establishing correlation across model fidelities is a prerequisite for bi-fidelity data ensembles (Geraci *et al.* 2017). To this point, Figure 6(c) shows a mismatch in ignition fate between low- and high-fidelity output. For that case, $\xi_2 = 1.98$ is close to the maximum asymmetry in the laser lobe radii. However, this case is also a different turbulent realization in which the laser energy was deposited compared to (a,b), where there was agreement across fidelities. The challenge in predicting ignition through a QOI such as the pressure trace is that this time series is dependent on the binary outcome of ignition success. Future work will examine whether there are more informative quantities of interest, e.g., joint probability distributions between hot regions of the flow and favorable equivalence ratio regions.

We additionally tested both four- and twelve-species chemistry mechanisms on the low-fidelity mesh for a single case as shown in Figure 6(a). It was found that the reduced chemistry overpredicted the maximum temperature associated with the energy kernel by a factor of two. However, in that test, the ignition fate did not change.

4. Conclusion

In this study, we report the progress in developing a multifidelity framework for predictions of a laser-ignited rocket combustor. The approach integrates low- and high-fidelity models to analyze complex interactions between laser energy, turbulent flows and ignition processes. The computational Schlieren visualization and the singular value decomposition of the low-fidelity simulation matrix both show that indirect rather than direct ignition requires further attention for accurate determination of ignition probability maps.

Experimental data from a model rocket combustor have informed the simulation uncertainty parameters. The necessary prerequisites have been investigated for the bi-fidelity approximation method, which allows for exploration of a broader parameter space. However, discrepancies between low- and high-fidelity simulations in ignition predictions indicate areas where progress is required for full deployment of the bi-fidelity workflow.

Acknowledgments

This investigation was funded by the Advanced Simulation and Computing program of the U.S. Department of Energy’s National Nuclear Security Administration via the PSAAP-III Center at Stanford, Grant No. DE-NA0003968.

REFERENCES

- CERFACS 2023 *CERFACS Chemical database*, accessed October 2023, <https://chemistry.cerfacs.fr/en/chemical-database/mechanisms-list/1-step-methane-air-mechanism/>.
- DI RENZO, M., FU, L. & URZAY, J. 2020 HTR solver: An open-source exascale-oriented task-based multi-GPU high-order code for hypersonic aerothermodynamics. *Comput. Phys. Commun.* **255**, 107262.
- DOOSTAN, A., GERACI, G. & IACCARINO, G. 2016 A bi-fidelity approach for uncertainty quantification of heat transfer in a rectangular ribbed channel. In *Turbo Expo: Power for Land, Sea, and Air*, , Vol. 49712. American Society of Mechanical Engineers.

- FAIRBANKS, H. R., JOFRE, L., GERACI, G., IACCARINO, G. & DOOSTAN, A. 2020 Bi-fidelity approximation for uncertainty quantification and sensitivity analysis of irradiated particle-laden turbulence. *J. Comput. Phys.* **402**, 108996.
- GERACI, G., ELDRED, M. S. & IACCARINO, G. 2017 A multifidelity multilevel Monte Carlo method for uncertainty propagation in aerospace applications. *AIAA Paper*, 2017-1951.
- MOIN, P. & MAHESH, K. 1998 Direct numerical simulation: A tool in turbulence research. *Annu. Rev. Fluid Mech.* **30**, 539–578.
- SMAGORINSKY, J. 1963 General circulation experiments with the primitive equations: I. the basic experiment. *Mon. Weather Rev.* **91**, 99–164.
- STRELAU, R., FREDERICK, M., SENIOR, W. C., GEJJI, R. & SLABAUGH, C. D. 2023 Modes of laser spark ignition of a model rocket combustor. In *AIAA Paper*, 2023-2377.
- TAO, Y., SMITH, G. P. & WANG, H. 2018 Critical kinetic uncertainties in modeling hydrogen/carbon monoxide, methane, methanol, formaldehyde, and ethylene combustion. *Combust. Flame* **195**, 18–29.
- WANG, J., DI RENZO, M., WILLIAMS, C., URZAY, J. & IACCARINO, G. 2021 Progress on laser ignition simulations of a CH_4/O_2 subscale rocket combustor using a multi-GPU task-based solver. *Annual Research Briefs*, Center for Turbulence Research, Stanford University, pp. 129–142.
- WANG, J. M. 2020 Ignition and hydrodynamic ejections by laser-induced breakdowns. PhD thesis, University of Illinois at Urbana-Champaign.
- WANG, J. M., BUCHTA, D. A. & FREUND, J. B. 2020 Hydrodynamic ejection caused by laser-induced optical breakdown. *J. Fluid Mech.* **888**, A16.