

Hardening of Arches for Commercial Simulation of Industrial Flares: A Program Update

Marc Cremer, Minmin Zhou, Dave Wang
Reaction Engineering International
cremer@reaction-eng.com
801-875-4314

Jeremy Thornock
University of Utah
j.thornock@utah.edu

Abstract

Reaction Engineering International (REI) is working with the University of Utah to leverage the Uintah Computational Framework (UCF) for commercial simulation of industrial flares, with funding from the Department of Energy (DOE). The Arches component of the UCF provides a reacting large eddy simulation (LES) capability, which is a more fundamentally accurate description of turbulent mixing and combustion than is obtained in conventional Reynolds Averaged Navier Stokes (RANS) approaches. Since the application of Arches to the simulation of commercial flares presents many challenges to a potential user, including software compilation, case definition, case setup, simulation and post-processing in an HPC facility, there is a need for streamlining this process to make Arches and commercial HPC facilities more accessible to flare designers and end-users. This paper provides an update on the results of our DOE program focusing on three areas:

- Air Assist Flares Simulation of TCEQ tests
- Multipoint Ground Flare Simulation
- Uncertainty Quantification Analysis

The air-assist flare simulations are based on two specific tests from the Texas Commission on Environmental Quality (TCEQ) 2010 flare test program. The primary difference between these two test cases is the stoichiometric ratio: 1) Test A2.1, SR=15.9; 2) Test A2.4, SR=28.2. The increased excess air in test A2.4 led to a measured reduction in combustion efficiency from 96% to 89%. The LES-based Arches simulations of these two test conditions resulted in a statistically stationary combustion efficiency prediction of 90% and 85%, respectively. Both of these simulations included a refinement to the case setup associated with the operation of the continuous pilots and the description of the vent gas inlets, which significantly improved the prediction of combustion efficiency compared to the results for test A2.1, which we had previously reported in the 2018 conference.

To develop a procedure for simulation of multipoint ground flare, we have used a handoff procedure that involves detailed simulation of a single flare tip and a coarsened simulation of the flare field. We have demonstrated the procedure using geometry from a three steam-assist flare array.

In addition to the simulations of single flares, we have been carrying out uncertainty quantification (UC) analysis using a Bayes' law/probabilistic framework to quantify the uncertainty of experimental measurements and simulation predictions of the downstream combustion efficiency (CE) from John Zink SKEC steam-assisted flares at a high turndown condition. The current analysis includes variability in the crosswind speed and steam flow rate. Comparing the mean predictive results with the measurement indicates good agreement for conditions greater than or equal to a $NHV_{cz}=100$ BTU/scf. The predictions at NHV_{cz} of 50 BTU/scf are significantly lower than the observed measurements.

1 Introduction

Visual examination of an industrial flare clearly demonstrates the highly unsteady behavior of flare flames, and the large range of turbulent length scales that are at play within a flare, fundamentally impacting the overall flare performance and flare emissions. Due to the inherent unsteadiness and the large range of turbulent length scales involved with industrial flare operation, it is challenging for Reynolds Averaged Navier Stokes (RANS) models to provide robust predictions of flare performance over a large range of design, environmental, and operational parameters. Large-eddy simulation (LES) may provide a feasible approach to high fidelity computational evaluation of industrial flares.

Researchers at the University of Utah have previously demonstrated the application of Arches, an LES-based combustion code, to simulation of industrial flares using US Department of Energy (DOE) and University of Utah high performance computing (HPC) facilities [1] [2]. The computational requirements in terms of data storage and computational time to carry out these simulations can be daunting. In addition, the tasks associated with code compilation, case setup, simulation, and post processing on sophisticated HPC hardware requires expertise with the LES software as well with the HPC hardware.

The key objective of our DOE sponsored program is to harden the application of Arches to commercial simulation of industrial flares to make the software and commercial HPC facilities more accessible to designers and end users. One area of development, which is not discussed in this paper, is the integration of Arches within a web-based user-interface to ease the tasks associated with case definition, access to commercial HPC facilities, case monitoring, and post-processing. The focus of this paper is the analysis associated with single elevated flares and multipoint ground flares, which will serve as template cases within a case library of the interface.

This paper focuses on three areas of analysis:

- Two air-assist flare cases having different amounts of air-assist, from the Texas Commission on Environmental Quality 2010 flare study
- Demonstration of hand-off approach for simulation of multi-point ground flares
- Uncertainty quantification analysis

2 Methods, Assumptions, and Procedures

2.1 Air-Assisted Flare Simulations

The air assisted flare from the TCEQ 2010 study was Model LHTS-24/60 and was provided by John Zink. The technical drawing in the TCEQ final report provided the geometrical representation of the flare. The LHTS-24/60 Flare has a tip diameter of 24 inch with 3 pilots with maximum capacity of fuel flow rate of 144,000 lb/hr. This flare size and design configuration represents many flare models currently in the field. The flare exit tip is 33 feet above ground level and the flare bottom diameter is 12.3 feet. The flare tip detail is shown in Figure 1.

The operating conditions of tests A2.1 and A2.4 in the TCEQ 2010 study [3] are listed in Table 1. Both tests fire pure propylene at a flow rate of 354 lb/hr, ~0.25% of maximum capacity.

Table 1: Operating Conditions for tests A2.1 and A2.4

Parameter	A2.1 Test	A2.4 Test
Fuel		
Fuel C ₃ H ₆ Mole Fraction	1.00	1.00
Fuel Flow Rate (lb/hr)	354	354
Fuel Flow Velocity (ft/s)	0.50	0.50
Fuel Initial Temperature (°F)	94.	92.
Air Assist		
Air Flow Rate (klb/hr)	83.6	148.5
Air Flow Velocity (ft/s)	45.7	80.9
Ambient Temperature (°F)	81.0	80.9
Stoichiometric Ratio (SR)	15.9	28.2
Wind Velocity (mph)	13.0	10.1
Measured CE - ARI (%)	95.9	89.3
Measured CE - PFTIR (%)	93.1	82.1




Figure 1: Three-pilot configuration for air-assisted flare

This fuel gas flow rate is in the range of operation for typical flow rates (less than 0.5% of maximum capacity) used in industry. The key difference between the two operating conditions is the air assist flow rate, which is significantly higher on A2.4, leading to a stoichiometric ratio (SR) increase from 15.9 to 28.2.

The three pilots are operated continuously with a heat input of 69,607 Btu/hr each. The arrangement of the three pilots is as shown in Figure 1. Based on the 2010 TCEQ test report, the pilots are fueled by natural gas and operated with a stoichiometric quantity of combustion air. The representation of the continuously operating pilots in our recent analysis was a refinement compared to that which we performed in the Phase I program [4].

Figure 2 shows the full-scale simulation domain (60.6 ft x 39.0 ft x 18.0 ft). With reference to the full-scale simulation from the Phase I project [4], the overall domain was decreased from (66.0 ft x 48.0 ft x 36.0 ft). In accordance with lessons from the previously completed mesh sensitivity study the mesh size of this simulation is 0.91 in. (2.3 cm) [5].




Figure 2: Full-scale simulation domain for air-assisted flare

Combustion Efficiency (CE) is computed as the percentage of the total hydrocarbon stream entering the flare that burns completely to form only carbon dioxide and water. The TCEQ 2010 tests [3] employed two approaches to measure the CE, Aerodyne Research, Inc. (ARI) used gas chromatography (GC) analysis of the vent and stack gases, and Passive Fourier Transform Infrared (PFTIR) was used for remote measurements of the flare plume constituents. The strategy for flare plume extractive sampling from the TCEQ tests is shown in Figure 3. The design of flare plume sampling system is to continuously draw as large a sample of the plume as possible, homogenize the sample and then obtain a sample of this well mixed portion of the plume for analysis. The sample collector also is moved into position approximately in the center of plume and at a distance far enough downwind from the flare tip to ensure that combustion reactions had ceased. Numerically, if CO_2 in the air is assumed to be negligible, CE is calculated as

$$CE(\%) = \left(\frac{CO_2(plume)}{CO_2(plume) + CO(plume) + 3 * \sum C_3H_6(plume)} \right) \times 100, \quad \text{Eq. 1}$$



Figure 3: Air-assisted flare tip with sampling collector in TCEQ 2010 test

where, CE (%) = combustion efficiency (%); CO_2 (plume) = volume concentration of carbon dioxide in the plume (ppmv) after combustion has ceased; CO (plume) = volume concentration of carbon monoxide in the plume (ppmv) after combustion has ceased; \sum hydrocarbons (plume) = volume concentration of all the unburned hydrocarbons in the plume after combustion has ceased multiplied by the number of carbons in the hydrocarbon (ppmv).

To be consistent with the experimental “plume-average”, our analysis assumed that the flare plume occupied all spatial regions where the mixture fraction exceeded 1 ppmv. We integrated the mass flow rate of all species at outlet boundaries of the simulation domain to provide the species information for the CE computation. For example, the CO_2 distribution selected by the mixture fraction criteria at the outlet is shown in Figure 4.




Figure 4: CO_2 distribution selected by 1 ppmv mixture fraction criteria at the outlet of simulation domain

2.2 Scaling Up Single-Flare Simulations to Multi-Flare Systems

A scaling procedure has been implemented into Arches for scaling up near flare-tip simulations to larger, more coarse domains. This procedure allows one to capture near flare-tip dynamics with a localized domain and communicate that information to a second simulation that simulates the larger scale mixing and the final downstream combustion efficiency. Since Arches relies on a static, Cartesian mesh, this method offers a way to perform a one-way coupled mesh refinement calculation without an expensive development investment in automatic mesh refinement or implementation of a stretched meshes.

The procedure for using the new feature is as follows:

- Perform a fine scale LES simulation of the near flare tip
- Time average the important data, and define some cutting plane that will serve as the handoff location to the coarse level simulation
- Export the data from the fine simulation at the specified hand-off cutting plane
- Conservatively (preferably) coarsen the fine data to match the resolution of the coarse mesh
- Use the coarsened data to serve as a boundary from which fluxes of mass, momentum and energy are injected into the coarse simulation

2.3 Uncertainty Quantification Analysis (Bayesian Approach)

Bayes probability law [6] offers a method for determining posterior probability distributions of input parameters given data from a model and data from measurement. Bayes law is

$$p(X|Y) \propto p(X)p(Y|X), \quad \text{Eq. 2}$$

where $p(\cdot)$ is a probability distribution, X is an input and Y is an output. Note that (X, Y) can also be defined as arrays of inputs and outputs. The term on the left ($p(X|Y)$) is the posterior. The terms on the right are the prior ($p(X)$) multiplied with the likelihood function ($p(Y|X)$). The function is a proportionality because the normalizing distribution ($p(Y)$) has been dropped for brevity.

In this uncertainty quantification analysis, we consider a single quantity of interest (QOI): the downstream flare combustion efficiency. The combustion efficiency (η) is defined as,

$$\eta = \frac{\phi_{CO_2}}{\phi_{CO_2} + \phi_{CO} + \phi_{HC}}, \quad \text{Eq. 3}$$

where ϕ_i is the integrated concentration of each species ($ppm \cdot m$) along a line of site through the sight downstream post-combustion plume. The input parameters (X) that dictate the many parameters. To reduce the overall dimensionality of the problem, we selected a subset of parameters for analysis. These parameters are:

- The mean crosswind speed (\overline{u}_{wind})
- The mean steam feed rate (\overline{m}_{steam})
- The mean crosswind direction ($\overline{\theta}_{wind}$)
- The PFTIR target characterized by the height above the flare stack and the distance downstream (h, d) at the center cutting plane orthogonal to the wind direction.

Note that we assume that the vent gas flow rate is held constant without variability. The first three parameters are self-explanatory. How these parameters contribute to the modeling are discussed next.

2.3.1 LES Simulation data

The raw modeled data are obtained from the Arches LES simulations. Various QOIs from Arches simulations are easily extracted, time averaged and reported for visualization or data reduction.

Geometric details of the SKEC flare head were scarce [7]. The simulated version was constructed using photos, online promotional information, and flow area information reported with the PFTIR data. Wind data including speed and direction were collected during the test but it was uncertain where the information was collected relative to the flare stacks. The average windspeed and direction were used for the simulated boundary conditions. These were held constant through the simulation. Other important boundary conditions, such as the vent gas conditions, flow rates and steam flow rates were provided in the report. For the simulation, we assumed the vent gas flow rate was known exactly and held constant. The steam flow rate, on the other hand, was variable.

We varied \bar{u}_{wind} and \bar{m}_{steam} over a range that encompassed the measured data, resulting in a set of 20 LES simulations. The sample points are shown in Figure 5. Note that Cases 1-6 were run with conditions that matched the mean of the experimentally measured wind speeds and steam flow rates for Replicates 1 as reported in the report. The remaining samples were chosen using a Latin Hypercube sampling technique to fill out the space. Each simulation was run for about 2-3 days on anywhere on 800 - 1500 cores on a local compute cluster housed at the University of Utah. The simulation time was chosen so that statistically steady values of concentration were obtained.




Figure 5: LES samples across the wind speed and steam flow rate variables.

2.3.2 PFTIR Instrumental Modeling

As mentioned above, concentrations of the various species were obtained from line-of-sight PFTIR measurements aimed at the downstream combustion plume. Details of the measurement and the procedures are outlined in the Clean Air Engineering report on the performance of the steam and pressure-assisted ground flares. To compare simulated LES results to the measurement data, LES data must be processed in a way that replicates the PFTIR device. To do this, we developed an instrumental model representing a simplified PFTIR that processes LES data to arrive at a combustion efficiency prediction.

The PFTIR instrument model is a function of three variables, $\eta = f(\bar{\theta}_{wind}, \bar{h}, \bar{d})$. Given an a priori specification of the position of the PFTIR at ground level and these three variables, a line-of-sight is run from the PFTIR through the time-averaged LES data. Along the line, concentration values of CO_2 , CO , and HC , are obtained that are then used to compute the combustion efficiency value for that line. The integration along the line of sight uses the fuel mixture fraction to weight the contribution of each local species:

$$\phi_i = \frac{\sum_j f_j \phi_j}{\sum_j f_j}, \quad \text{Eq. 4}$$

where i is the species index and j is the line-of-sight position index. The variable f is the total fuel mixture fraction (m_{fuel}/m_{total}). The PFTIR instrument model follows the same heuristics developed for the actual experiment, in that the optimal line-of-sight should be aimed within one flame length of the downstream hot combustion plume and not through any combustion zone.

The PFTIR model sampling is accomplished by varying the relative position of the PFTIR to simulate changing wind conditions and varying the target location at the center cutting plane (orthogonal to the wind direction). The wind direction is modeled by rotating the PFTIR ground position at a fixed radius about the flare stack. For a given single set of operating parameters, this procedure allows us to sample a single LES calculation from several PFTIR views, as if the wind had switched direction. Additionally, the position of the target (h, d) is varied. These parameters are then sampled over a set of predetermined bounds to produce several PFTIR predictions of η . The sampling is constructed using a Latin Hypercube design with around 300 samples. Samples that violate the sampling heuristics described above are discarded. Figure 6 and Figure 7 illustrate the sampling for the PFTIR model.




Figure 6: Example of the PFTIR instrument model sampling showing the variation in the targeting position and the wind direction effect. The flare stack is indicated with the green bar. The intersecting targets are shown with red dots. The starting and stopping locations of the PFTIR ray are indicated with black dots while the ray path is shown with the dashed blue line.




Figure 7: Instantaneous (left) and time-averaged (right) LES temperature plots with the PFTIR instrument model target samples shown with pink markers.

The PFTIR model operates on the LES data as a post-processing procedure after the LES simulation has reached statistically steady state. One sample of the PFTIR model is obtained in approximately one second. Because of the low expense, the PFTIR model is sampled several times ($O(100)$) over the parameter space.

A final note on the PFTIR model; while the wind speed, wind direction, and steam flow rate are crucial environmental and operational parameters, the PFTIR targeting may seem a bit of a nuisance. Indeed, the targeting parameters are recognized formally as statistical nuisance

parameters. That is, while the targeting is important for modeling the PFTIR so that comparisons to the measured data can be drawn, the resulting values of the target parameters are not particularly interesting. Had the targeting been reported, the range of these parameters may have been minimized or the parameters eliminated altogether.

3 Results and Discussion

3.1 Air-assisted Flare Simulations

Figure 8 and Figure 9 show an instantaneous flow field of the simulations of TCEQ tests A2.1 and A2.4, respectively, after each simulation reached a statistically stationary state. To clearly present the flow field, three surfaces were selected as $Z=0$ m (center plane, $X=9.5$ m and $X=11$ m, where the velocity profile are plotted in Figure 8 and Figure 9. $X=9.5$ m plane shows the typical turbulent flow field around the cylinder, in which the vortices are shed in the wake of the cylinder and further turbulent eddies are dissipated downstream.

Overall, the flare flow field has been derived by the interaction between the vertical air-inlet flow with the horizontal wind flow. The flame structure has been bent towards the horizontal direction due to the horizontal crosswind flow. In contrast, buoyancy brought by the combustion reaction lifts the flame from its horizontal position. Compared to A2.4 test, A2.1 test bends with a larger degree. The reason contributes to that is the momentum ratio of air assisted flow to the crosswind flow. Comparing the air assist flow rate of A2.4 test with vent gas flow rate and their very comparable molecular weights, the air assist exit velocity at A2.4 is the dominant driver in characterizing the momentum of the vent gas/air assist mixture exiting the tip of flare.




Figure 8: Velocity Magnitude distribution for A2.1 flare simulation: a) center plane, b) $X=9.5$ m plane, c) $X=11$ m plane




Figure 9: Velocity Magnitude distribution for A2.4 flare simulation: a) center plane, b) X=9.5m plane, c) X=11m plane

Figure 10 shows the temporal evolution of unburned fuel mass fraction at the central plane for test A2.1, and Figure 11 shows the evolution of CO₂ mass fraction for test A2.4. These serial plots show the history of the dynamic behavior of flare development. The snapshot at 0.4 sec for the earlier stage shows the less fuel concentration locates near the flare tip. A primary reason is the “hotspot” initialization which is terminated after 0.6 s. After approximately 1.8 sec for the A2.1 test, the flare approaches to the domain exit which indicates the large eddies turn over time is about 1.8 sec.

Noticeable unburned fuel at 2.6 sec for both cases is observed leaving the domain downstream outlet. Aerodynamic interactions associated with flare tip with the crosswind flow, along with the interactions with the plume of hot combustion gas cause turbulent eddies to strip unburnt fuel away from the combustion zone [1]. Subsequent dilution with air and cooling effect of the stripped fuel results in combustion inefficiency. The comparison of A2.1 and A2.4 indicates how CE is affected by crosswind flow, which is the consistent with the previous flow field analysis.




Figure 10: Temporal evolution of unburned fuel mass fraction field for full-scale simulation of A2.1




Figure 11: Temporal evolution of unsteady CO₂ mass fraction field for full-scale simulation of A2.4

Figure 12 the local stripping factor for the A2.1 simulation and A2.4 simulation. The local stripping factor is defined as the ratio of unreacted fuel to the total amount of fuel available locally if no reaction occurred. This is a useful parameter for identification of the location where combustion inefficiency is exhibited. Significant combustion inefficiency is occurring downstream of the flare plume for the simulation as the combustion ceases. Comparison of A2.1 and A2.4 shows the stripping factor for case A2.4 is approximately 0.2 at the domain outlet, whereas it is approximately 0.1 for case A2.1, which is consistent with the corresponding predicted CE for the two cases.




Figure 12: Volume rendered image of stripping factor assessment for A2.1 and A2.4 simulation

Spatial evolution of time-averaged CE along flare developing direction is shown in Figure 13. The results show that the domain outlet, the CE is the stationary stage after 7 sec. Both plots in Figure 13 show that the calculation of CE asymptotes to a constant value as the downwind calculation plane moves downstream to the model exit, which indicates that the simulation domain setup is sufficient for accurate calculation of CE.




Figure 13: Spatial evolution of time-averaged CE along flare developing direction

Figure 14 shows the corresponding time dependent plume-averages of CE for A2.1 and A2.4 simulations. The results show higher CE for case A2.1, consistent with the lower SR for that case. This trend is also consistent with the TCEQ results showing higher combustion efficiency for A2.1 than for A2.4. This indicates that correcting the continuous pilot operation and by correcting the fuel inlet specification for the simulation of A2.1 lead to a significantly improved prediction of combustion efficiency, compared to the reported prediction of 65% during the Phase I project.




Figure 14: Temporal evolution of time-averaged combustion efficiency at flare plume

Comparisons of the predicted and measured combustion efficiencies for tests A2.1 and A2.4 are shown in Figure 15. As shown, both the predictions and the measurements show a significant decrease in combustion efficiency due to the increase in flare SR. Although both the ARI GCM measurement and the PFTIR measurement for test A2.1 yield combustion efficiency higher than the predicted value of 90%, the agreement is much improved due to the corrections to the updated

A2.1 simulation which included the addition of the continuous pilots along with the correction to the fuel inlet specification. For test A2.4, the ARI data show a 6.6% reduction in combustion efficiency (i.e. 89.3% vs. 95.9%) compared to test A2.1. This decrease is reasonably consistent with the predicted decrease of 5% (i.e. 85% vs. 90%). The PFTIR data, however, show a larger decrease of 10% between the two cases (i.e. 82.1% vs. 93.1%).

In addition to likely inaccuracies in the simulations, there are also multiple sources of uncertainty in the measurements of combustion efficiency, which can contribute to the differences between the predicted and measured CEs. For the measurements, one area of uncertainty is the method of sampling. For the ARI measurements, sampling is dependent on the proper location of the sampling probe in the flare plume (i.e. Figure 3). Similarly, for the PFTIR measurements, variability in the line of sight leads to variability in the measured CE.




Figure 15: Comparison of measured and predicted combustion efficiency

3.2 Multipoint-ground flare simulation

A demonstration of the hand-off procedure was carried out for the SKEC flare tip. The fine-scale resolution of the flare tip used $dx = 3$ mm. Using a cutting plane directly above the vent gas exit, data was time averaged and extracted to disk. These data were then coarsened in a conservative manner to preserve total flux through the plane. The coarse mesh resolution was $dx = 1$ cm. Figure 16 shows two contours plots of the momentum flux through the plane. The filled contour plot uses the fine data while the dashed line contour plot is the coarsened version of the data.

Using the handoff data, a three-SKEC array was constructed and simulated, modeling the experimental conditions of the Marathon tests. Examples of the vorticity, temperature, and mixture fraction output are shown in the volume rendered images in **Figure 17**. With the handoff capability demonstrated, future work includes providing a quantitative analysis of the effect of the handoff on computed downstream combustion efficiency.




Figure 16: Fine and coarse representations of the same total flux through a plane. The fine data were obtained from a simulation of the SKEC flare tip using 3mm resolution. It was coarsened to a 1cm resolution for the far field simulation.




Figure 17: Volume rendered images of the coarse level simulations using the fine level information. Here we show vorticity (top-left), temperature (top-right) and total mixture fraction (bottom).

3.3 Bayesian analysis of flare Simulations

3.3.1 PFTIR Instrumental model results

For the PFTIR measurement method in section 2.3.2 , we chose normal distributions with identical parameters as our best guess prior distributions for the target parameters (h , d). The target locations are centered and normalized given the mean h and d location from the sampled points as,

$$\bar{h} = \frac{h - \mu_h}{\sigma_h} \quad \text{and} \quad \bar{d} = \frac{d - \mu_d}{\sigma_d} \quad \text{Eq. 5}$$

where the mean and standard deviation of the sample points are indicated with μ and σ respectively. Given the normal parameterization ($\mathcal{N}(0, .5)$) of the target parameters along with the centering and scaling, we arrive at the prior distribution shape shown in Figure 18.

To apply Bayes' law, we construct a normally-distributed likelihood function. Here, we assume a normal likelihood form,

$$p(Y|X) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left[\frac{Y - Y_{sim}}{\sigma}\right]^2} \quad \text{Eq. 6}$$

where Y is the experimentally measured data and Y_{sim} are the data from the PFTIR instrument model ($Y_{sim} = f(\theta_{wind}, h, d)$). The variance (σ) in Eq.9 is decomposed using errors in variables approach that results in two contributions, the variance of the wind direction $\sigma_{\theta_{wind}}$ and the uncertainty of the combustion efficiency, σ_{η} . Since these parameters are unknown, they also require prior distributions. Here, we use lognormal distributions for each parameter to ensure positive values. These distributions are shown in Figure 18. As described above, the intent of the application of Bayes law is to inform a priori distributions using the observed data (including the input measurements [e.g., θ_{wind}]). The result is the posterior distribution conditioned on the data, $p(X|Y)$. Posterior distributions that replicate the prior distributions indicate that the priors were not informed by the observations. Posterior distribution that are different than the priors that learning has occurred.




Figure 18: Prior distributions for the instrument model parameters, θ_{wind} , \bar{h} , and \bar{d} .

Relevant experimental data reported for this analysis includes the wind direction, steam flow rate, and PFTIR combustion efficiency measurement. For each replicate, the data were reported for all variables on a minute-by-minute basis over roughly a 20-minute period. Three replicates were performed for each of the NHVcz conditions. Even with a fixed set point, process variability in the steam flow rate was observed for each replicate resulting in variability in the NHVcz. Recorded environmental variability in the wind speed and direction occurred naturally. We constructed confidence intervals of 95% using all nonzero data for each replicate across the six conditions. Results are shown in Figure 19. We note that the variability of the data increases as the NHVcz decreases.

We used Bayes' law to get posterior distributions under each of the six conditions using mean NHVcz and crosswind values. As mentioned above, only the PFTIR instrument model parameters were varied to examine the effect of the PFTIR instrument model. Examples of two posterior results along with the variable cross interactions are shown in Figure 20. These conditions represent NHVcz values of around 300 BTU/scf (Case 1) and 50 BTU/scf (Case 6), which are the extreme ends on the tested operational steam flow rate. The uncertainty in σ_η increases from Case 1 (mode of $\sigma_\eta = .02$) to Case 6 (mode of $\sigma_\eta = .26$), reflecting the increasing uncertainty in the η value as the NHVcz drops. This trend is also observed in the predictive posterior values as shown in Figure 21 and Figure 22. These figures also included the confidence interval predictions from each case for reference.




Figure 19: Plot of confidence intervals using all nonzero reported experiment data for each condition and for each replicate within that condition.




Figure 20: Posterior ($p(X|Y)$) results for Case 1 (top) and Case 6 (bottom).




Figure 21: Predictive posterior for Case 1 (left) along with 95% confidence intervals (right) constructed from the measured data.




Figure 22: Predictive posterior for Case 6 (left) along with 95% confidence intervals (right) constructed from the measured data.

3.3.2 Full Parametric Analysis

We next extend the analysis to include variability in the crosswind speed and steam flow rate. This process increases the dimensionality of the problem, while retaining the same Bayesian workflow: (1) proposing prior distributions for the additional two parameters, (2) obtaining the posterior distributions via Bayes law, and (3) feeding these distributions forwards through the model to produce the predictive posterior values. The analysis, however, now requires the use of an ensemble of LES calculations due to the new parameters. The LES ensemble (see Figure 5) is used to create a surrogate response surface to enable rapid calculation of the model output given the inputs over the crosswind speed and steam flow rate parameters.

An additional modification to the analysis is that the errors-in-variables step is extended to include the steam flow rate and wind speed measurements. With these additions, the total variance in the likelihood function is a function of σ_η , $\sigma_{\theta_{wind}}$, $\sigma_{u_{wind}}$, and $\sigma_{\dot{m}_{steam}}$. We characterized the prior distributions for the two new parameters as log-normal to retain positive-only values. Figure 23 shows the posterior distributions of the parameters.

Figure 24 shows the final predictive posterior values with the measured data. In this plot each 1-min observation for each replicate across all conditions is plotted with a green X. The mean of the predictive posterior distribution is shown with the red circle. Note that for the experimental data, all nonzero values were used. Zero values were not included in this analysis. We emphasize that the results of the predictive posterior contain other statistical information because they are distributions. The full distributions, however, are omitted here for brevity.

Comparing the mean predictive posteriors with the measurement indicates good agreement for conditions greater than or equal to a $NHV_{cz}=100$ BTU/scf. The predictions at $NHV_{cz} \approx 50$ BTU/scf are significantly lower than the observed measurements. The results are also in stark contrast to Figure 21 and Figure 22, wherein just the PFTIR instrument model parameters were evaluated. At this point, more work is required to determine the source of the deviation, but we suspect that the response surface is not a good representation of the simulation data in this regime.




Figure 23: The posterior distributions with windspeed and steam flow rate included.




Figure 24: Means of the predictive posterior (red circles) compared against the measured data (green cross).

4 Conclusions

The key objective of this research is to harden the UCF, particularly the Arches LES model component, to the commercial simulation of industrial flares. Our efforts have focused on three

specific areas: 1) expanding the elevated flare simulation library, 2) development of a new functionality to simulate multipoint ground flares, and 3) Uncertainty Quantification analysis of flare simulation results and experimental data.

Regarding item 1, the Arches case setup and simulation of test A2.1 from the TCEQ 2010 flare study has been refined to include representation of the continuous pilots and to improve upon the representation of the vent gas inlet composition. The predicted combustion efficiency from this simulation is in significantly improved agreement with the measured CE compared to that from the phase I program. Test A2.4 was also simulated. The primary difference between these two test cases is the stoichiometric ratio: 1) A2.1, SR=15.9; 2) A2.4, SR=28.2. The increased excess air in test A2.4 led to a measured reduction in combustion efficiency from 96% to 89% in the TCEQ tests, based on the PFTIR measurements. The Arches simulations show a reduction in predicted combustion efficiency from 90% to 85%, in reasonably good quantitative agreement.

Regarding item 2, we have developed a handoff strategy to enable the simulation of a multipoint ground flare. Our initial demonstration is a three-SKEC array. Future work will include a quantitative analysis of the effect of the handoff on computed downstream combustion efficiency.

Regarding item 3, a Bayes' law/probabilistic framework has been used to quantify the uncertainty of experimental measurements and simulation predictions of the downstream combustion efficiency (CE) from John Zink SKEC steam-assisted flares at high turndown scenario. This analysis includes variability in the crosswind speed and steam flow rate. Comparing the mean predictive results with the measurement indicates good agreement for conditions greater than or equal to a $NHV_{cz}=100$ BTU/scf. The predictions at NHV_{cz} of 50 BTU/scf are significantly lower than the observed measurements.

References

- [1] P. Smith, A. Jatale, J. Thornock and S. Smith, "A validation of flare combustion efficiency simulations," in *AFRC Industrial Combustion Symposium*, Salt Lake City, 2012.
- [2] J. Thornock, P. Smith and A. Chambers, "LES Simulations of Sour Gas Flares in Western Canada," in *AFRC Industrial Combustion Symposium*, Houston, 2009.
- [3] D. T. Allen and V. M. Torres, "TCEQ 2010 Flare Study Final Report," Texas Commission on Environmental Quality. PGA No. 582-8-862-45-FY09-04, August 1, 2011.
- [4] M. Cremer, D. Wang, J. Thornock and M. McGurn, "Leveraging the Uintah Computational Framework for Commercial Simulation of Industrial Flares," in *AFRC Industrial Combustion Symposium*, Salt Lake City, 2018.
- [5] M. Cremer, M. Zhou, M. McGurn, A. Cowley, D. Wang, J. Thornock, "Hardening of Arches for Commercial Simulation of Industrial Flares," in *AFRC Industrial Combustion Symposium*, Hawaii, 2019.
- [6] J. H. D. Gelman A. Carli, *Bayesian Data Analysis- Third Edition*, Boca Raton, FL: Chapman & Hall/CCRC Press, 2014.

- [7] Clean Air Engineering, "Performance Test of Steam-Assisted and Pressure Assisted Ground Flare Burners with Passive FTIF-Garyville," Garyville, 2013.

5 Acknowledgement and Disclaimer

This material is based upon work supported by the U.S. Department of Energy, Office of Science, under Award Number DE-SC0017039.

This paper was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

The authors would like to acknowledge the Argonne Leadership Computing Facility for their allocation of core hours from their discretionary allotment.

The authors would also like to acknowledge the contributions of Dr. Ralph Price of CP Chem, Mr. Markus Martin of Dynamite Digits, and Mr. Scott Evans and Mr. Daniel Pearson of Clean Air Engineering.