

Application of global sensitivity analysis to FDS simulations of large LNG fire plumes

Adrian Kelsey, Simon Gant and Kevin McNally, Health and Safety Laboratory, Harpur Hill, Buxton, SK17 9JN

Steven Betteridge, Shell Research Ltd, Shell Technology Centre Thornton, P.O. Box 1, Chester CH1 3SH

In 2009, a series of experiments were conducted by Sandia National Laboratories at a large-scale test complex in Albuquerque, New Mexico, in which Liquefied Natural Gas (LNG) was released onto the surface of a large pool of water and then ignited. These experiments involved the largest releases of LNG ever performed, with the spills covering a circular area up to 83 m in diameter. In the test with the largest LNG release rate, unexpected behaviour was observed. The flames did not cover the entire area of the LNG spill, but were instead limited to a smaller region, around 50 m in diameter. The height of the flames was also greater than expected from previous large-scale tests.

One possible explanation for this behaviour is that in this very large scale test, the speed at which air and fuel vapour was drawn into the fire exceeded the local burning velocity. The flames could therefore not propagate upwind, against the entrained flow, and ignite the whole surface of the LNG pool. The additional unburnt fuel vapour drawn into the fire may also account for the higher flame height.

To study these phenomena, this paper presents a simulation-based sensitivity analysis of entrainment rates and flame heights in large LNG fire plumes. Pool fire simulations are performed using the Fire Dynamics Simulator (FDS) software and five model inputs are varied: the fire diameter, burn rate, radiative fraction, computational grid cell size and turbulence model. To perform a global sensitivity analysis on these five inputs using standard Monte Carlo sampling would involve thousands of FDS simulations and a large amount of computing resource. The alternative approach taken here is to create a Gaussian process based emulator from just 127 FDS simulations. The emulator is checked for its reliability and then used to perform the sensitivity analysis.

The analysis identified fire diameter as having the greatest influence on flame height and entrainment velocity. Flame height was also affected by burn rate, while entrainment velocity was affected by burn rate, grid resolution and turbulence model. Interactions between multiple inputs being varied at the same time are also examined. These findings help to advance our understanding of the Sandia experiments, which may have important implications for hazard assessment of major LNG spills.

Keywords: LNG fire plumes, FDS, global sensitivity analysis, GEM, Sandia Phoenix experiments

Introduction

In recent years, the growing global demand for natural gas has led to an increase in the transport of LNG using ship tankers and the development of new technologies, such as Floating LNG (LNG) facilities. In parallel, concerns have been raised about the impact from potential incidents, which could result in large spills of LNG. A number of hazard studies have been performed to quantify the effects of large diameter LNG pool fires, but these have produced widely varying hazard consequence assessments (Blanchet *et al.*, 2011). A limiting factor has been the large degree of uncertainty in the prediction of flame height, fire diameter and, importantly, the degree of smoke shielding. These uncertainties have largely been due to the limited amount of experimental data available to develop the fire models, and the fact that all of the experiments were at a much smaller scale than the potential incidents. Up until 2009, the largest pool fire tests were the 35 m diameter LNG pool fire tests carried out at Montoir (Nedelka *et al.*, 1990).

To address this issue, Sandia National Laboratories was contracted by the US Department of Energy in 2008 to perform a series of large-scale LNG pool fire tests on water. These experiments, which became known as Phoenix tests (Blanchet *et al.*, 2011), involved the largest releases of LNG ever performed. In the second test, with the largest LNG release rates, unexpected behaviour was observed. The LNG spill covered a circular area up to 83 m in diameter, but the flames did not cover the entire area of the LNG spill, being limited to a smaller region, around 56 m in diameter. The flame height of 146 m was also much larger than expected.

One possible explanation for this behaviour is that in this very large scale test, the speed at which air and fuel vapour was drawn into the fire exceeded the local burning velocity. The flames could therefore not propagate upwind, against the entrained flow, and ignite the whole surface of the LNG pool. The additional unburnt fuel vapour drawn into the fire may also account for the higher flame height.

To examine the credibility of this hypothesis, two parallel projects have been undertaken at HSL over the last year. The first project, funded by Shell, has focused on understanding the physical behaviour in the Sandia tests using a relatively simple CFD model, where the fire has been represented as a volumetric heat source and vapour has been released from the surface of the pool (Betteridge *et al.*, 2014). The use of this model has been driven by the need to simulate the dispersion of fuel vapour from the pool that does not burn immediately when it mixes with air, and instead only burns once it is entrained into the fire. Standard CFD models for non-premixed combustion would instead assume that “mixed equals burned” and the entire surface of the pool is on fire. Another advantage of this model is that the rates of pool vaporisation and heat release rate from the fire can be controlled independently, to study their effects on entrainment rates.

Although this work has provided very useful results, one of its limitations is that the fire has been represented in an idealised way, as a uniform source of heat in a prescribed cylindrical region of space. To help assess the effect of this approximation on the predicted entrainment rates, a second project has been undertaken in parallel at HSL (which is presented in this paper) to study the behaviour of large fire plumes using the Fire Dynamics Software (FDS), (McGrattan *et al.*, 2013a).

The FDS fire model simulates the turbulent mixing processes between fuel vapour and air, and in a simple pool fire it should (in principle) provide more reliable predictions of the entrainment rate than a volumetric heat source model. However, one of the

limitations of using FDS to simulate the Sandia experiments is that it cannot simulate the dispersion of non-burning fuel from the pool surface, since the combustion model in FDS assumes that fuel combusts once it mixes with air. In the approach used here, the pool fire is modelled as a circular burning region, surrounded by ground. To account for the entrainment of additional fuel from the non-burning pool surface in the Sandia experiment, the fuel supply rate into the base of the fire has been varied in the model (using a range of values, some of which are higher than would be expected for an LNG pool fire).

To produce a clearer understanding of the model behaviour, various other model inputs have been varied to investigate their effect on the entrainment rate. These include the fire diameter, the proportion of heat that is released from the fire as radiation, the computational cell size and the turbulence model.

Both the Betteridge *et al.* (2014) volumetric heat source model and the FDS model used here have their strengths and weaknesses in modelling the Sandia Phoenix experiments. Although comparing two models against each other does not represent “validation”, it is useful to see whether they provide consistent results.

The most common way of investigating CFD model sensitivity, where there are multiple variable inputs, is to pick a baseline case and then to change each input variable one-at-a-time, whilst keeping all of the other inputs fixed. This approach is conceptually simple, but it has the disadvantages that the findings can be sensitive to the particular choice of the baseline case, and no information is provided on the interactions between multiple inputs varying at the same time. An alternative is to perform a *global* sensitivity analysis where all of the variables are varied simultaneously. One way of performing a global sensitivity analysis would be to use a Monte Carlo approach. However, in the case considered here, to run such an analysis using Monte Carlo sampling on five variable inputs would require thousands of FDS simulations. Since each FDS simulation typically takes 6 to 12 hours to run on a multi-processor desktop, this would involve very lengthy computing times (of the order of months).

To avoid these time constraints, the approach taken in the present work has been to construct a statistical emulator (essentially, a sophisticated curve fit) from a limited number of FDS simulations. The global sensitivity analysis has then been performed using this emulator, which is very fast to run, and the whole analysis including the 127 FDS simulations has taken just 2 weeks of computing time on a desktop computer. This use of a statistical emulator is relatively novel in the field of process safety. A similar approach was recently used by HSL to examine the sensitivity of the DNV Phast dispersion model (Gant *et al.*, 2013), but to our knowledge, this is the first time an emulator of the type used here (based on a Gaussian process) has been applied to study a CFD model in a practical process safety-related application. An important objective of this study has therefore been to understand the capabilities and limitations of Gaussian process emulators for global sensitivity analysis in a practical CFD engineering context. In fact, this was the main motive for undertaking the work.

Returning to the Sandia pool fire application, since there is relatively limited information in the literature on entrainment rates into large pool fires, the flame height has also been output from FDS in order to provide a second means of assessing the accuracy of the model predictions. This also enables the sensitivity analysis to consider two separate model outputs and examine important practical issues, such as “do the model inputs that have the greatest effect on the entrainment rate also have the greatest effect on the flame height”?

The paper proceeds by providing some background to the FDS model and the Gaussian process based emulator, before presenting the results of the global sensitivity analysis and concluding with the significant findings. The focus of this paper is principally upon application of the global sensitivity analysis. More detailed analysis of the Sandia experiments is provided in the companion paper by Betteridge *et al.* (2014).

Methodology

Fire plume simulations

The fire plume simulations were performed using FDS version 6 (McGrattan *et al.*, 2013a), which is a freely available CFD code that be downloaded from the National Institute of Standards and Technology (NIST) website¹. The software is widely used in the fire engineering community and is accompanied by comprehensive verification and validation manuals (McGrattan *et al.*, 2013b, 2013c). FDS has been used extensively to model pool fires. Examples can be found in the validation exercise for the US Nuclear Regulatory Commission (NRC, 2007) and the work of Trouvé (2008).

Turbulence modelling

The effects of turbulence are modelled in FDS using a Large-Eddy Simulation (LES) model. The simulations therefore resolve the time-dependent behaviour of large-scale turbulent motion, such as the characteristic puffing instabilities in buoyant plumes. The effects of small-scale (unresolved) turbulence are modelled using an enhanced (turbulent) viscosity. Various LES models are available in FDS version 6 and in the present work two models are tested: the Smagorinsky (1963) model with a constant coefficient of $C_s = 0.2$, which was the default turbulence model in previous versions of FDS, and the Deardorff (1972) model, which is the new default model in FDS version 6. Since the FDS simulations are inherently transient, they must be run for a sufficient duration to obtain stable statistics for the mean quantities of interest. This aspect of the analysis is discussed in more detail below.

Combustion and Radiation

Combustion is modelled using the default FDS single-step mixture-fraction based model, which is appropriate for well-ventilated fires, where it can be assumed that “mixed is burned”. In common with most large-scale fire simulations, the computational grid cells in the simulations presented here are too large to resolve the thin flame front. Since the heat released from the flame is distributed across a relatively large volume (the width of the cells), the maximum flame temperatures are likely to be under-predicted. To avoid this having an impact upon the thermal radiation predictions (due to the fourth-power dependence of the radiation intensity on

¹ <http://www.fire.nist.gov/fds>, accessed 10 December 2013

temperature) FDS uses a pragmatic approach, where the proportion of heat released as radiation is specified by the user, using a parameter called the “radiative fraction”. A high value of this parameter is suitable for sooty hydrocarbon fires whilst a lower value should be used for cleaner burning fuels. In the Sandia tests the fires appeared to burn with little cold soot production. To investigate the effect this may have had on the flow behaviour, FDS simulations are performed for a range of radiative fractions, from its default value of 0.35 (suitable for smoky hydrocarbon fires) to a value of 0.20.

Computational grid resolution

The computational grid in FDS is Cartesian and cubic cells have been used for all of the simulations presented here. To characterise the grid resolution in fire plumes, it is common to relate the cell size Δx to the characteristic fire length scale $D^* = (Q^*)^{2/5}D$, where D is the fire diameter (m) and Q^* is the dimensionless heat release rate, given by:

$$Q^* = \frac{\dot{Q}}{\rho_{\infty} c_p T_{\infty} \sqrt{g} D^{5/2}} \quad (1)$$

where \dot{Q} is the heat release rate (J s^{-1}), ρ_{∞} the ambient density (kg m^{-3}), c_p is the specific heat ($\text{J kg}^{-1} \text{K}^{-1}$), T_{∞} the ambient temperature (K), and g is the acceleration due to gravity (m s^{-2}). In the fire plume simulations performed for the US Nuclear Regulatory Commission validation study, a range of grid sizes with $D^*/\Delta x$ from 4 to 16 was used (NRC, 2007), whilst Chung and Devaud (2008) later suggested $D^*/\Delta x = 40$ as a good compromise between accuracy and computing time. In the present work, the grid resolution has been varied from $D^*/\Delta x = 16$ to 40, to assess its impact on the results.

Pool fire

The pool fire is prescribed in FDS as a circular opening, through which a specified mass flow rate of methane is released, treating the release as pure methane is a simplified representation of LNG. Since a Cartesian grid is used, the circular opening has a jagged, stair-stepped outline, but a sufficient number of cells has been used that provides a reasonably accurate representation of a circle (see Figure 1).

To investigate the effect of the fire size on the entrainment rates into the plume, simulations are performed using a range of circular pool diameters, from 10 to 100 m. This range encompasses the conditions present in both the first and second Phoenix experiments, which had LNG spill diameters of 21 m and 83 m, respectively (Blanchat *et al.*, 2011).

The specified flow rate of methane through the pool surface area (henceforth called “burn rate”) is varied from $0.05 \text{ kg m}^{-2} \text{ s}^{-1}$ to $0.5 \text{ kg m}^{-2} \text{ s}^{-1}$. This range includes the average burn rate of $0.126 \text{ kg m}^{-2} \text{ s}^{-1}$ (with a range from 0.084 to $0.336 \text{ kg m}^{-2} \text{ s}^{-1}$) recommended in the review of LNG fire plumes by Luketa *et al.* (2008); the average burn rate from the first Phoenix test of $0.147 \text{ kg m}^{-2} \text{ s}^{-1}$ with a reported uncertainty of 20.4% (Blanchat *et al.*, 2011); and the range from 0.0945 to $0.49 \text{ kg m}^{-2} \text{ s}^{-1}$ suggested in a white paper from the Mary Kay O'Connor Process Safety Center (2008). In the second Phoenix test, the equivalent burn rate across the 56 m diameter fire from burning the vapour released across the whole of the 83 m diameter spill surface area was estimated by Betteridge *et al.* (2014) to be $0.30 \text{ kg m}^{-2} \text{ s}^{-1}$, at the time when the flame height was at its maximum.

Computational Domain Size

The FDS simulations use a rectangular computational domain with dimensions of (3 x 3 x 8) pool diameters in the (length x width x height) directions. The height of 8 fire diameters was needed to ensure that all of the heat from the fire was released within the domain. The total number of grid cells varied from 42,000 to 4.2 million.

Duration

Medium to large scale plumes are characterised by the repetitive shedding of coherent vortical structures at a well-defined frequency, a phenomenon known as “puffing” (DesJardin *et al.*, 2004). For axisymmetric fire plumes, Cetegen and Ahmed (1993) found the following relationship between the duration of a puffing cycle in seconds, t , and the diameter of the burner or source in metres, D :

$$t = \frac{\sqrt{D}}{1.5} \quad (2)$$

This equation is used here to specify the time duration of the FDS simulations, which are set to 30 puffing cycles. The first 5 cycles are used to allow the flow to become established and the next 25 cycles are used to calculate the averaged statistics. Checks are performed to ensure that this provides a sufficient averaging time (see below).

FDS Model Outputs

Two quantities are output from FDS for the sensitivity analysis: the flame height and the radial entrainment velocity. The flame height is defined as the height over which 99% of the heat from the fire is released. This is the same approach and value used by McGrattan *et al.* (2013c), but it should be recognised that the choice of 99% is somewhat arbitrary. Flame heights based on 95% would give a flame heights in the range of 15-25% smaller. In the Phoenix tests, Blanchat *et al.* (2011) measured the flame lengths using a different approach that involved image processing of the videos of the fire plumes. The comparison of the flame heights from these experiments and from the FDS results is therefore subject to some uncertainty. To provide an additional means of assessing the results, the flame height correlation from Heskestad (2002) is also shown in the results presented below.

The radial entrainment velocity of interest is the maximum radial velocity at the edge of the fire. This was calculated from FDS by extracting velocity data along four vertical sampling lines at 90° intervals around the edge of the pool (shown in Figure 1). The

maximum radial velocity on each of the vertical lines was found and an average was calculated of these four values over time. The entrainment velocity was not measured in the Phoenix tests.

To assess whether the FDS simulations are run for a sufficiently long period of time to generate reliable mean values, the 95% confidence intervals on the mean values of the flame height and entrainment rate are calculated using a moving block bootstrap method, the approach used to analyse the previous LES simulations of Gant (2010). This use of confidence intervals to judge statistical convergence of LES results was originally proposed by Celik *et al.* (2006).

FDS Model Inputs

A summary of the model inputs used in the global sensitivity analysis is provided in Table 1.

Table 1 Model inputs and their ranges used in the global sensitivity analysis

<i>Model Input Parameter</i>	<i>Range</i>
Pool fire diameter	10 – 100 m
Burn rate	0.05 – 0.5 kg m ⁻² s ⁻¹
Radiative fraction	0.20 – 0.35
Grid resolution ($D^*/\Delta x$)	16 – 40
Turbulence model	Deardoff or Smagorinsky



Figure 1 Photograph of Phoenix test 2 (Blanchat *et al.*, 2011) showing 83 m diameter pool and 56 m diameter fire (left), and sample results from an FDS simulation showing radial velocity sampling positions (right)

Global sensitivity analysis

The purpose of the global sensitivity analysis presented here is twofold:

- 1) To identify which model input parameters have a significant effect on the model outputs
- 2) To identify how the outputs vary as a function of the inputs

To address Item 1, a variance-based analysis has been performed. The variance is the average of the squared distance between the individual model predictions and the average over all of the simulations. The contribution to the variance from each input, when their value is varied over their specified range, indicates their importance. Important parameters explain a large proportion of the output variance.

Two quantities are output from the variance-based analysis to identify the importance of each input parameter: the main effect and the total effect. The main effect of a parameter describes its influence on the output due to changing it alone. The total effect of a parameter includes any additional influence due to its interactions with other varying parameters.

This type of variance-based analysis of main and total effects is popular in the literature, principally because these quantities are relatively easy to calculate and understand. Also the importance of the interactions usually decreases rapidly with increasing order (i.e. interactions between three parameters varying simultaneously are often weaker than between two varying parameters), so the difference between main and total effects often gives a good indication of the important joint effects.

The analysis presented here uses an emulator-based approach. A small number of FDS runs are used to create an emulator, which is essentially a sophisticated curve fit linking the inputs to the outputs. It is much faster to run than FDS, taking less than a second instead of many hours. The sensitivity analysis is then performed using this emulator, rather than FDS. With clever design and analysis, some of the sensitivity measures needed to calculate main and total effects can be evaluated analytically from the emulator, further reducing the cost of the sensitivity analysis.

To address Item 2, the emulator is used to show the mean behaviour of the output as a function of each input (where the mean is obtained by averaging over all of the other varying inputs). For instance, it can show how the flame height varies as a function of the pool fire diameter. Equally, the results can show how the mean output varies as a function of two inputs varying at the same time. In each case, thanks to the type of emulator used here, a measure of the emulator fit to the underlying FDS model is provided, in terms of confidence intervals. This quantification of the uncertainty in the emulator fit to the underlying FDS model is another important benefit of a Gaussian process based emulator over alternative response surface models.

GEM

The global sensitivity analysis is performed using the Gaussian Emulation Machine (GEM) software (Kennedy, 2005), which is freely-available for non-commercial use. The software features an easy-to-use Graphical User Interface (GUI) and good documentation. Some background to the use of Gaussian emulators is provided by Oakley and O'Hagan (2004) and O'Hagan (2006).

A fundamental assumption of the emulator is that the output is a smooth, continuous function of the input parameters (Figure 2). Formally, the method is based on a multivariate normal distribution, with specified mean and covariance functions that depend on the relative positions of the training data points. The method used to derive the emulator is Bayesian in that initially, a simple function is used to relate model inputs to outputs (a prior distribution) and this is updated using a relatively limited number of training data points to derive the emulator (the posterior distribution).

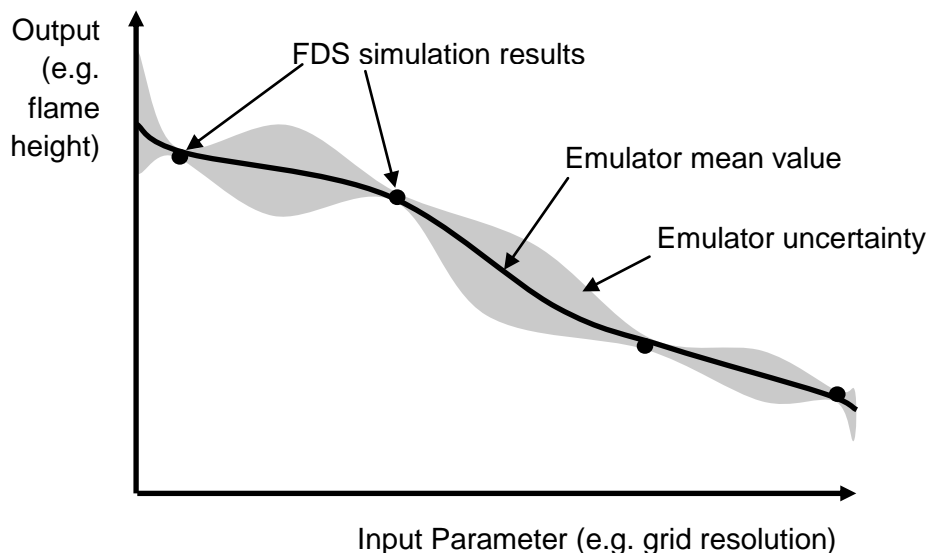


Figure 2 Illustration of Gaussian emulator fit to training data points

In designs for computer experiments, the aim is to use the smallest number of points to sample the parameter space to minimise the number of simulations needed. Space-filling designs are therefore chosen that cover the input ranges by varying all of the parameters simultaneously without repeating any of them. Often, some of the input parameters have little effect on the output and therefore a design where no value is repeated is efficient, because a parameter that has no effect still provides useful information to assess the effect of other varying inputs.

In the LNG pool fire application considered here, four of the model inputs are continuously varying parameters and the fifth (the turbulence model) is a switch that is either Deardorff or Smagorinsky (see Table 1). The design used was suggested by Challenor

(2011) and uses a Sobol' sequence with $n = 2^m - 1$ points in d dimensions, where m is a user-defined integer that controls the sample size. The first $2^{m-1} - 1$ points are used for one option of the switch (here the Deardorff model) and the remainder for the other option (Smagorinsky). The full set of n points is a Latin hypercube design, with good coverage, and both subsets also have these properties. For the case considered here, the experimental design used $m = 7$ to give a design of $n = 127$ points in $d = 4$ dimensions. This number of points was chosen to be in excess of the rough rule-of-thumb of Loeppky *et al.* (2009), which states that at least 10 times the number of input parameters is needed to identify key factors.

If an FDS simulation was repeated with the same input conditions, the mean flame height and entrainment rate predicted by the model might vary by a small amount due to a combination of numerical errors, the sensitivity of the turbulence model to initial conditions and the process used to obtain averaged statistics. A consequence of this small error is that the response surface produced from the FDS training data points would not be smooth if it was forced to zigzag through all the points. To enable a smoother fit, a small uncertainty bound is therefore introduced around each data point during the emulator fitting process, using what is known in the statistical literature as a "nugget". The size of this error is estimated by GEM from the variance in the training data. Justification for use of a nugget has been provided by Gramacy and Lee (2012).

It is important to check that the emulator is constructed using a sufficient number of training data points. To achieve this, "cross-validation" tests are performed in which the emulator is fitted with one of the training data points left out. The emulator is then used to predict this missing training data point, and the emulator predictions compared to the FDS predictions at that point. This process is repeated over all of the training data points, leaving each one out in turn, to obtain an overall picture of the emulator accuracy across the whole sample space.

Results

FDS simulations

To demonstrate that the FDS simulations were run for a sufficiently long period to generate reliable statistics, 95% confidence intervals were calculated for the flame height and entrainment velocity. Across all 127 simulations, the 95% confidence interval was found to be around 5% of the median value. In other words, if the median flame height was 100 m, the 95% confidence intervals were at 97.5 m and 102.5 m. It was therefore considered that averaging over 25 plume puffing cycles was sufficient to provide reliable mean values of flame height and entrainment rate.

In Figure 3 predicted mean flame heights from FDS can be compared to measurements from the Phoenix tests and the Heskestad (2002) correlation:

$$\frac{H}{D} = 3.7(Q^*)^{0.4} - 1.02 \quad (3)$$

where Q^* is the non-dimensional heat release rate, H is the flame height and D is the fire diameter. Multiple FDS results have the same value of Q^* , since the fire diameter and burn rate have been varied independently, and Q^* is proportional to $\dot{Q}/D^{1/2}$. The flame heights from the Phoenix tests are shown with error bars representing the measurement uncertainty.

Table 2 Measurement uncertainties on Phoenix test vapourisation rates and flame heights.

Test	Pool diameter	Fire diameter	Mean pool vapourisation rate	Fire vapourisation rate	Flame height
	m	m	kgm ⁻² s ⁻¹ mean (range)	kgm ⁻² s ⁻¹ mean (range)	m mean +/- 2sd
1	21	21	0.147 (0.117, 0.177)	0.147 (0.117, 0.177)	70 +/- 5
2	83	56	0.147 (0.117, 0.177)	0.330 (0.270, 0.400)	146 +/- 16

Figure 3 shows that the FDS results are in reasonably good agreement with the first Phoenix tests. The Deardorff and Smagorinsky models appear to show similar behaviour (this is examined further in the sensitivity analysis presented below). Overall, there is a general trend for FDS to predict slightly higher flame heights than were measured, the FDS simulations also show higher flame heights than predicted by the Heskestad correlation, but with a similar trend. These differences in flame height should be viewed within the context of the difficulty in measuring flame heights in both experiments and in CFD simulations. A more detailed discussion of flame heights in the Phoenix experiments is given in the companion paper by Betteridge *et al.* (2014).




Figure 3 Comparison of mean flame height (H), expressed in terms of fire diameters (D), as a function of the dimensionless fire size dimension Q^* . Results from FDS using the turbulence models of Smagorinsky (+ symbol) and Deardorff (\times symbol), measured flame heights in the Phoenix tests (\oplus symbol) and the empirical correlation from Heskestad (— line)

Emulator validation

The 127 FDS results were used to construct four emulators: one each for the Smagorinsky and Deardorff models, and one each for the flame height and entrainment velocity outputs. Cross-validation checks on the emulator fit to the underlying FDS output were performed, and some sample results are shown in Figure 4. Each graph shows the emulator's prediction of the flame height or entrainment velocity against that predicted by FDS, with error bars showing two standard deviations (95% confidence intervals) either side of the emulator predictions. While most of the comparisons are reasonable there are some outliers. These mostly represent simulations where the inputs were close to the minimum or maximum values given in Table 1. Towards the ends of the range of any parameter, the training points become sparse and emulator uncertainty increases. In effect, at the edge of the parameter space, the cross-validation emulator is being used to extrapolate, rather than to interpolate its value from neighbouring points. Care should therefore be taken with emulator predictions towards the end of the ranges of any parameter.

Variance-based sensitivity analysis

The variance based analysis shows that the fire diameter has the greatest effect on both the flame height and the entrainment rate. Varying the fire diameter across its full range (from 10 m to 100 m) explains over 75% of the total output variance.

For the flame height, the burn rate explains a further 25% of the variance, with some interaction between fire diameter and burn rate, but not a large amount. Varying the mesh resolution and radiative fraction has almost no effect on the predicted flame height. The turbulence model also does not significantly affect the relative contributions from the other input parameters.

The entrainment velocity is dominated by the fire diameter input, but the burn rate, grid resolution and turbulence model choice all have an effect. The grid resolution has a greater effect when the Deardorff turbulence model is used than when the Smagorinsky model is used. The levels of interaction between the input parameters are smaller than they were for the flame height.

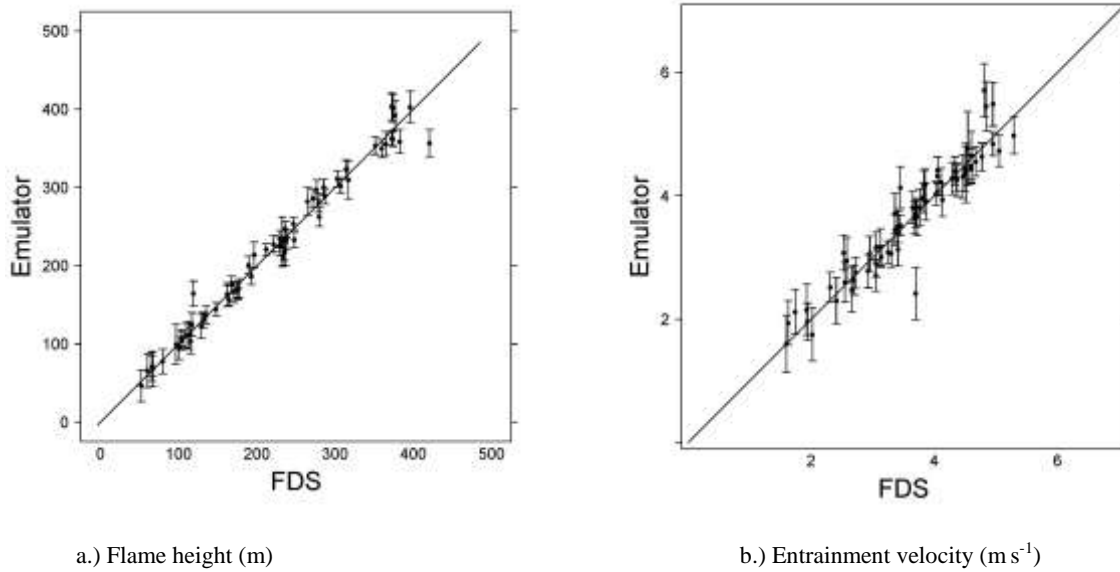


Figure 4 Results from cross-validation for the Smagorinsky model on flame height and entrainment velocity

Flame height sensitivity

The response of flame height to changes in the fire diameter and burn rate is shown in Figure 5 using a lattice plot. Each individual plot shows the predicted flame height on the vertical axis against the fire diameter on the horizontal axis. In addition, the three plots from left to right show the influence of burn rate varying from 0.1 kg m⁻² s⁻¹, to 0.275 kg m⁻² s⁻¹ and to 0.45 kg m⁻² s⁻¹. The analysis shows that the predicted flame height increases with both fire diameter and burn rate. The rate of increase with fire diameter is close to linear but may be decreasing towards the top of the range of fire diameters.

The results from both Deardorff and Smagorinsky turbulence models in Figure 5 are similar, confirming the finding from the variance-based analysis. The uncertainty in the emulator predictions, expressed in terms of 95% confidence intervals, is also plotted for each curve as a grey region, which shows that the emulator uncertainty is small, across the ranges presented here.

The results shown in Figure 5 are all based on grid resolution of $D^*/\Delta x = 28$, and the effect of the radiative fraction has not been shown. The variance-based analysis showed that these two inputs had little effect on the flame height.

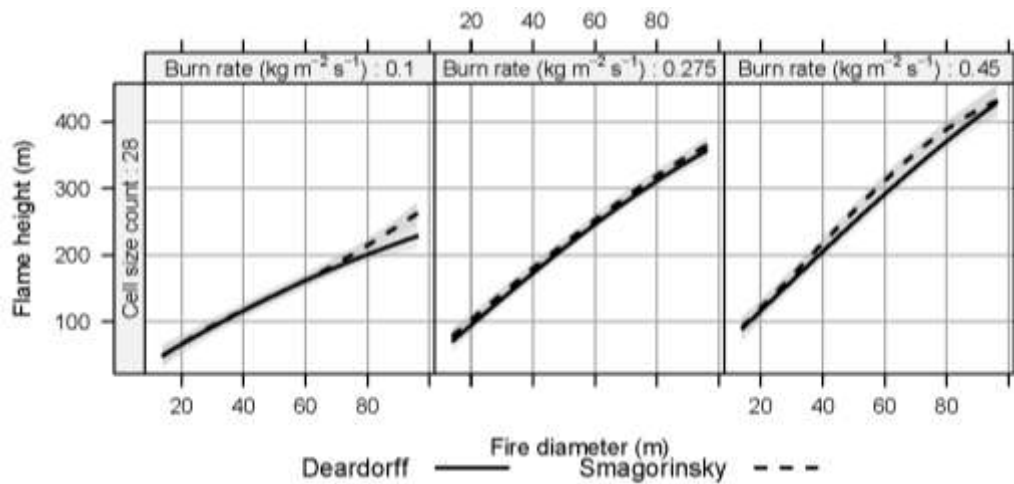


Figure 5 Emulator predictions of flame height against fire diameter with emulator uncertainty, showing influence of burn rate

Entrainment velocity sensitivity

Figure 6 shows the response of entrainment velocity to the fire diameter, burn rate and mesh resolution. Each individual plot shows the entrainment velocity on the vertical axis against the fire diameter on the horizontal axis. From left to right, the plots show the influence of the burn rate. The effect of the mesh resolution is examined by comparing the top row (fine mesh) with the bottom row (coarse mesh). The radiative fraction is fixed at the middle of its range, since the variance-based analysis showed it had little effect on the entrainment velocity.

The results show that fire diameter has the dominant effect on the entrainment rate. The burn rate has a secondary effect, but its influence is less than was observed in the flame height predictions (Figure 5). At the fine mesh resolution, the results from the Deardorff and Smagorinsky turbulence models are almost identical, but for the coarser mesh, the predictions from the two models

differ. Both models predict the entrainment velocity to reduce with increasing mesh resolution, which it would be useful to examine further with respect to the wall treatment in FDS.




Figure 6 Emulator predictions of entrainment velocity against fire diameter with emulator uncertainty, showing influence of burn rate and mesh resolution

Comparison to Betteridge *et al.* (2014) entrainment rates

One of the reasons for undertaking this study was to compare the entrainment rate predictions of FDS against the model predictions of Betteridge *et al.* (2014). For a 50 m fire diameter with a burn rate of 0.38 kg m⁻² s⁻¹ the entrainment velocities predicted using emulators are 3.4 m s⁻¹, with a 95% confidence interval of +/- 0.2 m s⁻¹. The equivalent result from Betteridge *et al.* (2014) predicts an entrainment velocity of 3.7 m s⁻¹. The fact that the two models produce similar results provides some additional level of confidence in the results.

As noted by Betteridge *et al.* (2014), these predicted entrainment velocities are of the order of 8 to 10 times larger than the laminar burning velocity, and larger than flame spread velocities of 2 m s⁻¹ for liquid pool fires reported by Gottuk *et al.* (2002). The results therefore tend to support the hypothesis that the fire did not spread across the whole surface of the LNG spill in the Phoenix second test because the entrainment velocity was in excess of the local burning velocity. However, further modelling and experimental work is needed to quantify the local burning velocities in LNG pool fires to be sure that this explains the observed behaviour.

Discussion

The sensitivity analysis has shown that the fire diameter has a dominant effect on both flame height and entrainment velocity. However, the influence of the other input parameters differed, depending upon which model output was analysed. The flame height was almost completely dependent on fire diameter and burn rate, whereas the entrainment velocity was sensitive to the mesh resolution and turbulence model as well. This result is important from a CFD fire modelling perspective. Choosing a grid resolution with FDS that is sufficiently fine to have no effect on the flame height may not guarantee that the entrainment rate will be grid-independent.

The use of a mean based sensitivity analysis allowed this behaviour to be examined in some detail. It showed that with finer meshes the predictions of the entrainment velocity from the Smagorinsky and Deardorff turbulence models became almost identical. This behaviour is consistent with the fact that a greater proportion of the turbulence energy should be resolved with a finer grid (and less of it modelled), so the overall influence of the turbulence model should diminish with a finer mesh.

It should be noted that the results from a variance-based sensitivity analysis are sensitive to the choice of the parameter ranges. Whilst the fire diameter and burn rate have been varied by an order of magnitude in this study, the radiative fraction has only been varied from 0.20 to 0.35 (Table 1). Sensitivity to parameters will vary but the ranges used will have an effect, with larger ranges potentially showing more influence. If required, the input ranges can easily be reduced within the parameter ranges and the sensitivity analysis repeated using the emulator, without necessarily having to undertake further FDS simulations. If the particular area of the input space that is being re-examined has poor coverage, this should be readily apparent in the uncertainty confidence intervals output from the emulator. This is another reason why the type of Gaussian process based emulator used here has some advantages over alternative sensitivity analysis tools.

The Sobol’ sequence design used here to select the combinations of inputs used in each of the 127 FDS simulations does not repeat any points. Since the flame height predictions have been shown to be insensitive to the choice of turbulence model, it should be

possible to use all of the simulations (for both turbulence models) to construct a single emulator. This would give improved coverage of the other varying inputs and help to reduce the uncertainty in the emulator predictions. This analysis, together with further investigation on the treatment of switches in emulators, will be undertaken in future work.

Conclusions

A global sensitivity analysis of large scale LNG fire plumes has been undertaken to examine the response of the flame height and entrainment velocity predicted by a fire model to five varying model inputs: the fire diameter, burn rate, radiative fraction, computational grid resolution and turbulence model. The fire modelling was performed using FDS and the sensitivity analysis using GEM.

The results from the analysis showed that the fire diameter had the greatest influence on both the flame height and entrainment velocity. The effect of the other inputs was different, depending upon which output was considered (flame height or entrainment velocity). For the flame height, the burn rate had a secondary effect, but the other three input parameters had little or no influence. For the entrainment velocity, there were secondary effects from the burn rate, grid resolution and turbulence model, but the radiative fraction had little effect.

Results were presented to show how the flame height and entrainment rate varied in response to changing multiple inputs at the same time. Examination of the model sensitivity showed that predictions from the two LES turbulence models converged as the grid resolution was increased.

The results showed that care should be taken in assessing the degree of grid dependence in CFD models. Choosing a grid resolution that was sufficiently fine to have no effect on the flame height did not give grid independent results for the entrainment rate.

The type of global sensitivity analysis presented here can be used to improve our understanding of physical behaviour and model predictions. Using software tools like GEM and making use of continued advances in computing power, it is now straight-forward to apply these techniques to practical process safety-related problems. There is the potential for this to significantly improve the use of consequence modelling tools in quantitative risk assessment and incident investigation. HSL intends to lead a joint-industry project to develop the use of sensitivity and uncertainty analysis in process safety in the coming years.

Acknowledgements

The authors would like to thank Dr. Marc Kennedy (Food and Environment Research Agency) for use of the Gaussian Emulator Machine (GEM) and the developers of FDS at NIST for providing their software freely and supporting its use. The photograph of Phoenix test 2, Figure 1, is taken from Blanchat et al (2011), the authors would like to thank T. Blanchat, SANDIA National Laboratories for permission for its reproduction. The contributions made to this paper by Adrian Kelsey, Simon Gant and Kevin McNally were funded by the UK Health and Safety Executive (HSE). The contents, including any opinions and/or conclusions expressed, are those of the authors alone and do not necessarily reflect HSE policy.

References

- Betteridge, S., Hoyes, J. R., Gant, S. E. and Ivings, M. J., 2014, Consequence modelling of large LNG pool fires on water, Submitted to IChemE Hazards 24, Edinburgh, UK, May 2014.
- Blanchat, T., Helmick, P., Jensen, R., Luketa, A., Deola, R., Suo-Anttila, S., Mercier, J., Miller, T., Ricks, A., Simpson, R., Demosthenous, B., Tieszen, S. and Hightower, M., 2011, The Phoenix series large scale LNG pool fire experiments, SAND2010-8676, Sandia National Laboratories, Albuquerque, NM.
- Celik, I.B., Klein, M., Freitag, M., and Janicka, J., 2006, Assessment measures for URANS/DES/LES: an overview with applications, *Journal of Turbulence*, 7(48): 1-27.
- Cetegen, B. M. and Ahmed, T. A., 1993, Experiments on the periodic instability of buoyant plumes and pool fires, *Combustion and Flame*, 93: 157-184.
- Challenor, P., 2011, Designing a computer experiment that involves switches, *Journal of Statistical Theory and Practice*, 5: 47-57.
- Chung, W. and Devaud, C. B., 2008, Buoyancy-corrected k-ε models and large eddy simulation applied to a large axisymmetric helium plume, *International Journal of Numerical Methods in Fluids*, 58: 57-89.
- Deardorff, J. W., 1972, Numerical investigation of neutral and unstable planetary boundary layers, *Journal of Atmospheric Sciences*, 29: 91-115.
- DesJardin, P. E., O'Hern, T. J. and Tieszen R., 2004, Large eddy simulation and experimental measurements of the near field of a large turbulent helium plume, *Physics of Fluids*, 16: 1866-1883.
- Gant, S. E., 2010, Reliability issues of LES-related approaches in an industrial context, *Flow, Turbulence and Combustion*, 84: 325-335.
- Gant, S. E., Kelsey, A., McNally, K., Witlox, H. W. M. and Bilio, M., 2013, Methodology for global sensitivity analysis of consequence models, *Journal of Loss Prevention in the Process Industries*, 26: 792-802.

- Gottuk, D.T. and White, D.A., 2002, Liquid Fuel Fires, Chapter 15, Section 2, SFPE Handbook of Fire Protection Engineering, Third Edition. National Fire Protection Association, Inc., Quincy, Massachusetts, USA.
- Gramacy, R. B. and Lee, H. K. H., 2012, Cases for the nugget in modelling computer experiments, *Statistics and Computing*, 22: 713-722.
- Heskestad, G., 2002, Fire Plumes, Flame Height, and Air Entrainment, Chapter 1, Section 2, SFPE Handbook of Fire Protection Engineering, Third Edition. National Fire Protection Association, Inc., Quincy Massachusetts, USA.
- Kennedy, M. C., 2005, GEM-SA, version 1.1 software: Gaussian Emulation Machine for Sensitivity Analysis, Available from: <http://www.tonyohagan.co.uk/academic/GEM/index.html>. Accessed 3 December 2013.
- Loepky, J. L., Sacks, J. and Welch, W. J., 2009, Choosing the sample size of a computer experiment: a practical guide, *Technometrics*, 51: 366-376.
- Luketa, A., Hightower, M. and Attaway, S., 2008, Breach and safety analysis of spills over water from large liquefied natural gas carriers, SAND2008-3153, Sandia National Laboratories, Albuquerque, NM.
- Mary Kay O'Connor Process Safety Center, 2008, LNG pool fire modeling, White paper, September 2008.
- McGrattan, K., Hostikka, S., McDermott, R., Floyd, J., Weinschenk, C. and Overholt, K., 2013a, Fire Dynamics Simulator Technical Reference Guide Volume 1: Mathematical Model, NIST Special Publication 1018, Sixth Edition.
- McGrattan, K., Hostikka, S., McDermott, R., Floyd, J., Weinschenk, C. and Overholt, K., 2013b, Fire Dynamics Simulator Technical Reference Guide Volume 2: Verification, NIST Special Publication 1018, Sixth Edition.
- McGrattan, K., Hostikka, S., McDermott, R., Floyd, J., Weinschenk, C. and Overholt, K., 2013c, Fire Dynamics Simulator Technical Reference Guide Volume 3: Validation, NIST Special Publication 1018, Sixth Edition.
- Nedelka, D., Moorhouse, J. and Tucker, R. F., 1989, The Montoir 35 m diameter LNG pool fire experiments, *Proceedings 9th International Conference on LNG*, Nice, 17-20 October, 1989, Published by the Institute of Gas Technology, Chicago, Volume 2, III-3, 1-23 1990.
- NRC, 2007, Verification and Validation of Selected Fire Models for Nuclear Power Plant Applications, NUREG-1824, U.S. Nuclear Regulatory Commission, Washington D.C.
- Oakley, J. E. and O'Hagan, A., 2004, Probabilistic sensitivity analysis of complex models, *Journal of the Royal Statistical Society, B* 66: 751-769.
- O'Hagan A., 2006, Bayesian analysis of computer code outputs: a tutorial, *Reliability Engineering and System Safety*, 91: 1290-1300.
- Trouvé, A., 2008, Challenges in CFD modeling of large-scale pool fires, *Proceedings of Advanced Research Workshop on Fire Computer Modeling*, Santander, October 2007.
- Sobol', L., 1967, On the distribution of points in a cube and the approximate evaluation of on the distribution of points in a cube and the approximate evaluation of integrals, *USSR Computational Mathematics and Mathematical Physics*, 7: 86-112.
- Smagorinsky, J., 1963, General circulation experiments with the primitive equations. I. The basic experiment, *Monthly Weather Review*, 91: 99-164.