Harpur Hill, Buxton, SK17 9JN
Telephone: 0114 289 2000
Facsimile: 0114 289 2050

HEALTH & SAFETY LABORATORY

On Defining a Safety Criterion for Flammable Clouds

HSL/2007/30

Project Leader: **Dr M. Ivings**
Author(s): **Dr D. Webber**
Science Group: **Fire and Explosion Group**

CONTENTS

EXECUTIVE SUMMARY

The question is investigated of whether Lower Flammable Limit or a half of that concentration should be used as the criterion in safety analyses. Relevant experimental data, model concepts, and theoretical ideas are reviewed, and the imprecision of the question is highlighted. The opinion formed in this paper is that requiring the instantaneous centre-line concentration to be less than half the lower flammable limit is reasonable, and that relaxing this criterion would involve some risk.

1 INTRODUCTION – THE QUESTION

The question as initially posed is

"Should one use as the criterion for the possibility of ignition of a cloud (a) that the concentration is below the Lower Flammable Limit (LFL) or (b) that it is below ½ LFL"

Here we shall examine the factors which affect this decision and make a recommendation. We note from the outset that the question is ill-defined. In particular, as we shall discuss in detail, it does not define what is meant by the "concentration". More fundamentally, it implicitly contains the idea that whether or not a cloud is ignitable, simply depends on (some measure of) concentration. We shall discuss this too.

Published research has addressed this kind of question from a number of angles, which we shall review here. It does not appear possible to point to one definitive result and say "here is the answer". Rather one must weigh the various pieces of evidence, make the question more precise where possible, and come to an opinion. This is what is offered here.

The meaning of the question is obviously dependent on the meaning of the term "Lower Flammable Limit" and of the term "the concentration" which we now go on to examine.

Lower Flammable Limit: This is measured by mixing a gas uniformly with air, to a given concentration at a given temperature, in a tube or other container, and then attempting to ignite it. There is a fairly sharp distinction between concentrations which will ignite and those which won't and this defines the Lower Flammable Limit. Typically, different methods may result in estimates of say 4.5% or 5.0% for the LFL of methane, which is not significant compared with what we shall discuss below. Concentration itself is not an issue when measuring LFL as, in the case of a uniform mixture, there is no ambiguity. There may be some dependence on temperature, but not much in the temperature range of interest. The way in which one defines "the concentration" (for gas clouds in the open) is likely to be far more problematic.

2 CONCENTRATION

In a uniform mixture in the laboratory, “concentration” is well defined, but what do we mean by “the concentration” of a gas cloud in the open? The answer is in fact very ill defined unless one qualifies the term quite strongly. This means that the question as posed above is itself very ill-defined, and much of the rest of this report will discuss how this can be resolved.

Concentrations in gas clouds fluctuate in time and space. They can be averaged over a certain time, which reduces both the fluctuations and the mean value more, the longer the averaging period. Equally they can be averaged over a certain volume. And of course there is a lot of variation of the mean concentration in general through time and space. Let us examine some of these features.

2.1 SPATIAL VARIATION OF CONCENTRATION THROUGH A CLOUD

For the moment, let us assume we have a trusted device for measuring concentration, and use the word “concentration” here to refer to what it measures.

A gas cloud in the open atmosphere is not a parcel of gas at uniform concentration. If we look at the concentration at different points in the cloud at the same time - $c(x,t)$ at different x and the same t - the results will be very different. An example is provided by the measurements of Liedtke and Schatzmann (1997) of a buoyant plume of helium and air rising in a calm atmosphere (actually in a laboratory). Figure 1 shows the (normalised) helium concentration as one moves laterally through the plume at various heights above the source. A similar picture with high centre line concentration decaying towards the edges emerges in the case of jet releases (plumes driven by a high initial velocity rather than by buoyancy). (Note that this concentration is clearly averaged over a time long enough to make it roughly independent of time for this steady source case.)

Common sense suggests that we should at a minimum demand that the centre-line concentration is below LFL in order for ignition not to take place. Even so, however, common sense may be frustrated if one does not look too closely at what concentration a model is predicting – see below.

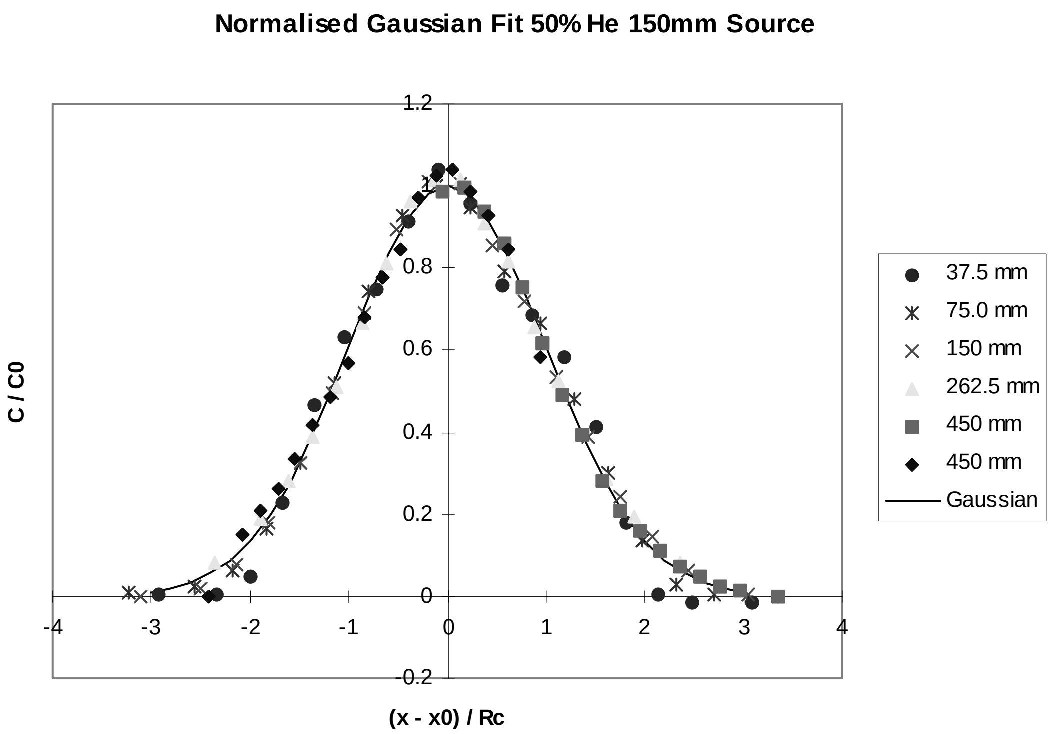

Figure 1: The lateral concentration profile in a buoyant helium plume as measured by Liedtke and Schatzmann 1997.

2.2 TIME-AVERAGED CONCENTRATIONS

Changing the subject only slightly: look at a plume from a chimney (outdoors) in a wind. It meanders. Imagine some contaminant in the plume, and picture a concentration detector at some fixed point in the sky. As the plume meanders, the detector will sometimes be in the centre of the plume and sometimes out of it, and sometimes at the edge. It will therefore be reading a high concentration or none at all, or anywhere in between – more or less at random.

In order to get a meaningful idea of “the concentration”, an average over a certain time is almost always quoted. For passive plumes in the atmosphere an averaging time of 10 minutes or half an hour may be considered appropriate for many purposes, and such averages are often presented in experimental measurements. However, one may validly ask what a concentration averaged over (say) 10 minutes can possibly have to do with whether a cloud is ignitable “instantaneously”.

This question can be examined in the context of heavy gas clouds near to a source. The Thorney Island Trials conducted in the mid 80’s released a cylinder of gas, 14m high and 13m in diameter. Typically, for most trials, the released gas was about twice the density of the ambient air. Concentrations were measured by detectors on fixed masts. Fluctuations in concentration were known to be a relevant feature of such clouds, and so the detectors were built to resolve concentrations on a time scale of 0.6s. The graphs of concentration against time with a point plotted every 0.6s oscillated wildly. So wildly as to defy any sensible direct interpretation of the expression “the concentration”. This is just the result of turbulent flow – which results in randomly varying concentration just as it results in randomly varying velocity (known as eddies). Figure 2 is the result obtained by Brighton (1985) when he averaged the data over 1 second, and then further averaged the concentration of all the sensors at a given height which were in the cloud. Even at the lowest measurements on the mast (0.4m above the ground)

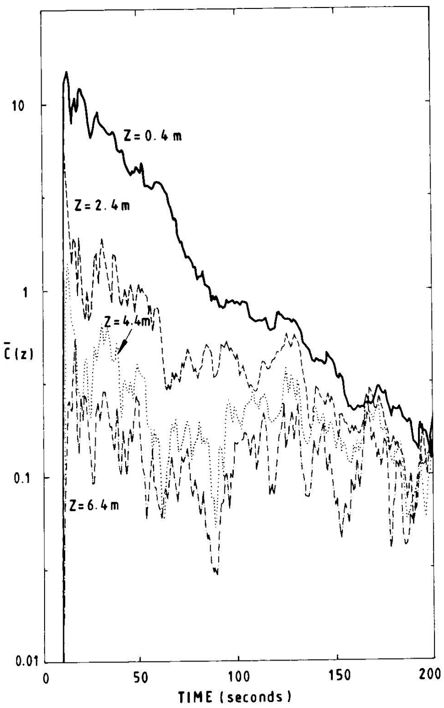

Figure 2. Area-averaged, one-second time averaged concentrations, at different heights z, in an instantaneously released heavy gas cloud.

one sees a "non systematic" component of the concentration variation amounting to approximately a factor of 2. At higher levels where there is more intermittency, the fluctuation is rather more prominent. Individual sensors' one second averages will fluctuate much more than this.

2.3 MORE ON SPATIAL VARIATION OF CONCENTRATION

The cloud with such rapidly fluctuating concentration at the fixed position of a detector, is of course flowing past the detector (with a randomly fluctuating turbulent velocity). The variation in time is therefore accompanied by a random variation in space of the "instantaneous" concentration at any fixed time (superimposed of course on a general trend where the concentration is usually larger in the middle and small on the outside – sometimes with sharp edges).

This situation leads to the possibility of some pockets of cloud being above LFL, some below, and with perhaps a connected path through the cloud (and in particular back to the source) above LFL but perhaps not. This leads in turn to the possibility that the whole cloud will all burn in a flash fire, or perhaps an attempt at ignition will cause a small region to burn and then go out quickly. The spatial average concentration over a large region may not be directly relevant for ignition.

2.4 MORE ON TIME AVERAGES

Attempts have been made to relate different time average concentrations. If one samples at high enough frequency, obtaining what one supposes to be "instantaneous" concentration readings, then one can average them with whatever longer averaging time one chooses. An empirical relationship which is commonly quoted is

$$\frac{c_1}{c_2} = \left(\frac{t_1}{t_2}\right)^{-0.2} \qquad (1)$$

where c_1, is the measured concentration averaged over an interval t_1, and c_2 is the concentration at the same time and place averaged over an interval t_2. This relationship obviously depends on the structure of the turbulence and even if true in some circumstances may not be in all. We shall discuss it below.

Accepting it for the moment, let us note by way of illustration, if t_1=18.75s and t_2=10min (=600s) then c_1/c_2=2.0. Thus, if it were thought that the concentration measured by a detector which implicitly averaged its input over 18.75 seconds should be below LFL, then the 10min average concentration would have to be below ½ LFL.

Caveat: this simple formula can only have limited validity – otherwise very short time averages would produce arbitrarily high concentrations – which is clearly nonsense. One has the concept of an "instantaneous" concentration which should be measurable with a detector of high enough time resolution – albeit with a spatial averaging determined by the size of the volume over which its nozzle (or whatever) does the sampling! The formula is discussed at some length by Wilson (1995) with lots of caveats, but there is a certain amount of support for it for times in the range 10s to 10min.

This example is taken from life and provides a very direct example of how the meaning of concentration can very directly affect the answer to our original question. We'll come back to it. But for the moment let us simply note that if person A says the concentration must be below LFL and person B says the concentration must be below ½ LFL, then they might both be right (or both wrong of course) – if A is talking about an 18.75 sec average and B is talking about a 10 minute average.

2.5 THE BOOK BY WILSON (1995)

Time averages are the subject of the extensive monograph of D Wilson (1995) published by the American Institute of Chemical Engineers' Center for Chemical process Safety. There is much interesting material here, but it is difficult always to be quite sure what to make of it. For example the very first sentence of Chapter 1 "Background and Objectives" is "Most plume dispersion models used in hazard assessments are based on environment protection models developed by regulatory agencies for predicting the long-term impact of pollution sources." It is very hard to reconcile this statement with the intensive research into heavy gas cloud dispersion and assorted source terms which took place over the two decades between the Flixborough accident and the publication of Wilson's work. This research includes major programmes of the HSE and European Commission and during this period dozens of heavy gas dispersion models were published focusing significantly on immediate fire and explosion hazards. Numerous, field and wind-tunnel experiments on heavy gases during this period most certainly did not just measure the 1-hour average concentrations Wilson discusses in his introduction; to do so would have been meaningless. It is important therefore not to take such statements too literally, especially as Wilson actually goes on later in the book to discuss some of the research. He includes short chapters on dense and buoyant clouds and on jets, pointing out that the fluctuations may not be much smaller than in passive plumes in the atmosphere, though he does give the impression that atmospheric fluctuations are his prime concern (which is possibly because there may be a greater quantity of experimental data in that area).

2.6 ENSEMBLE AVERAGES

Time averages are not the only form of averaged concentration. One can measure the "instantaneous" concentration at the same time after release and the same place in each of a series of "identical" releases. One gets a different result each time, but one can average over 10, 100, or 1000 releases or however many (in principle). The concept of Lower Flammable Limit is again seen to be more complicated. If some of the experiments produce a concentration above it, and some below, then what does it matter whether the average is above or below? We are not attempting to ignite the average, but rather to one particular realisation.

2.7 CONCENTRATIONS PREDICTED BY MODELS

Life is further complicated by the fact that the COMAH safety case assessor is not dealing with measured concentrations, but with concentrations predicted by models. It may therefore be asked "what precise form of concentration is the model predicting"? Indeed, from what is discussed above, one sees immediately that this must be asked.

One thing is clear: no model currently available makes direct predictions of "instantaneous" concentrations. This is true of integral models and of CFD calculations. The former implicitly assume some averaging and the latter, the ones used in hazard analysis, explicitly make some form of Reynolds decomposition of fluctuating fields, which in itself assumes that a meaningful time average can be found with an averaging time longer than the important turbulent "fluctuations", but shorter than the time over which the "mean" concentration changes in an evolving flow.

So what form of concentration does any given model predict? Do we even know? There are various possibilities with different consequences for the assessor. In the list below:

- the "modeller" is the person who derived the mathematical model, wrote the computer code, and validated it against experimental data;
- the "analyst" is the person who used the model to make predictions for a safety case, and wrote the safety case;
- the "assessor" is the person, from the regulatory body, attempting to assess whether the safety case makes sense.

All or any of these may be teams of people, rather than individuals, but the distinction is enough for current purposes.

So what form of concentration does any given model predict? Various possibilities (some dire) are:

- The model has never been compared with data. In this case no-one can know.
- The modeller doesn't know.
- The modeller does know but never wrote it down and so the analyst, who has read the model document, doesn't know.
- The modeller does know and wrote it down in the documentation, but the analyst is inexperienced and doesn't understand the significance and ignores it. The modeller may even have offered a choice (say) of time averages, ignored by the analyst who always chooses the default value.
- The modeller does know, the analyst knows, and makes well defined choices, but doesn't write down what the choices are in his safety case.
- Finally it is possible that the modeller does know, the analyst knows, and makes well defined choices, and documents those choices in the safety case.

It will be readily appreciated that the only satisfactory situation for the assessor is the last one. It will also be readily appreciated that we haven't yet answered the original question.

Coming back to models. Even if the modeller and analyst are less diligent than those in the last case above, we can say something about the predicted concentrations, in all but case 1. When a model is compared with experimental data, various free parameters are tuned to optimise the fit, and then left at their optimal values for making future predictions. Entrainment coefficients are one case in point (but there are others). In fitting say 10min average experimentally measured concentrations the modeller is (possibly implicitly) making the choice that this is the form of concentration he wants to predict. Had he fitted 6sec average concentrations, then he'd have obtained different values for his entrainment coefficients, and the model would be tuned to predict 6sec average concentrations. So the answer to what the model attempts to predict lies entirely in its validation programme. Thus, for example, the decision of whether one should use LFL or ½ LFL as a safety criterion with respect to the model's predicted concentrations, must depend on what concentration the model has been tuned to predict. And different models are tuned differently. Modellers can offer a choice of what is predicted, and one way is to include an option by building in a formula like equation 1. There are others, which we shall discuss below.

2.8 SUMMARY OF CONCENTRATION CONSIDERATIONS

Quoted concentrations can be any of a whole host of different things, whether they are measured or predicted concentrations. Real, instantaneous, point concentrations fluctuate a lot, and usually some sort of average is measured or predicted.

In order to estimate flammability we have to

- know exactly what concentration is measured/predicted;
- be able to relate that concentration by way of the LFL to an estimate of whether or not the cloud can ignite.

The first of these points is the responsibility of the analyst, who must use a model where it has been explicitly considered by the modeller.

The second has been the subject of research which we shall now review.

3 EXPERIMENTS

Let us here briefly examine some relevant experiments

3.1 EXPERIMENTS AT WARREN SPRING LABORATORY AND AT TNO

Hall et al (1991) at Warren Spring conducted experiments making many “identical” instantaneous releases of an advecting heavy gas cloud and measured the (effectively instantaneous) concentration at given points for each instance. From a detailed statistical analysis of this, and other data taken at TNO, Davies (1992,1994) found an approximately log-normal probability distribution in concentrations at different given points. Jones et al (1994) went on to embody this empirical result into the UKAEA computer code DRIFT as described in the section on of modelling below.

3.2 JET EXPERIMENTS BY BRITISH GAS

British Gas’s research division conducted a number of experiments on natural gas jets which addressed the question of ignition.

Birch et al (1981) made very detailed, small scale measurements of mean and fluctuating components of concentration in a single phase jet and found that their ability to ignite it was not solely dependent on the mean, or root mean square, concentration.

Figure 3 summarises their main result. It maps a jet with its axis shown vertical and the radial direction horizontal. The points show the boundary of the flammable region; the lines are concentration contours.

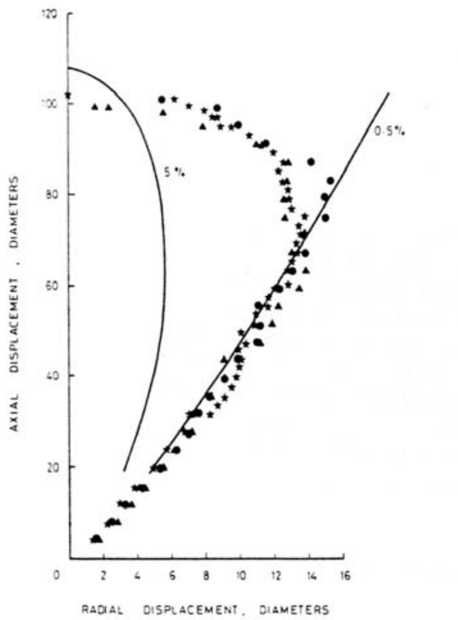

FIG. 2. The flammable boundary for complete light-up of the jet ● Re = 20900, U_0 = 32.2 m/s; * Re = 16700, U_0 = 25.7 m/s; ▲ Re = 12500, U_0 = 19.2 m/s: Solid lines give contours of mean volume fraction concentration.

Figure 3. The flammable boundary for complete light-up of a jet reproduced from Birch *et al.* (1981)

Two features are clear: (a) along the centre line of the jet the ability to ignite it fails close to the point where it dilutes to the LFL (about 5%) but (b) just a little way short of this it is ignitable at some distance from the centre line down to mean concentrations of 1/10 of LFL.

Clearly there are indeed other factors than just mean concentration which determine whether the jet will ignite, and so it would be imprudent a priori just to accept the centre line conclusion and look no further.

Smith et al (1986) did more experiments on gas jets (with a variety of gases) this time focussing on centre line concentration only. At very high jet velocities they were not able to light the jet (on the centre line) to form a stable flame. At more moderate jet velocities the jet was ignitable on the axis roughly down to ½ LFL and could be lit (on the axis) so as to burn back to the source anywhere above LFL.

Birch et al (1989) pursued a similar approach using natural gas jets emerging vertically in a cross-flow. Again they emphasise that ignitability is not simply correlated with mean concentration, but found regions of significantly less than ½ LFL which could be lit would light back to the source.

Taken together these experiments indicate that, whilst it may be difficult to light a very fast jet in a calm environment on the centre line, various factors in a more realistic scenario – a slower jet, ignition off centre, and a jet in a cross-flow – all contribute to making it easier to light at a mean concentration well below LFL.

3.3 EXPERIMENTS BY SHELL RESEARCH

Evans and Puttock (1986) report experiments which tested the ignitability of a heavy advecting propane cloud. Their concentration measurements consisted of recorded analogue signals digitised at 10Hz, and then smoothed with a 1s moving average. This they describe as a “peak” concentration and assume a “mean” concentration 1.4 times smaller, appealing to measurements of similar clouds in earlier experiments.

They found, in terms of this mean concentration, that that sustained ignition could occur at 0.9 LFL and that smaller flames could be produced at 0.6 LFL.

4 MODELLING APPROACHES TAKEN IN THE LITERATURE

4.1 A PROBABILISTIC STUDY

At the outset it is worth noting that some approaches try to avoid the problems discussed above. Spencer and Rew (1997) and Spencer, Daycock and Rew (1998) provide an interesting example.

Their approach is based on the concept of a correlation of the probability that a cloud has ignited with its area. The domain of interest is divided into a number of cells with different distributions of ignition sources of different types. They define graphs of ignition probability against area for different kinds of domain. It is not emphasised, but it appears that “area” in this case means the area of ground covered by that part of the cloud which is above LFL. There is considerable discussion of the nature and statistical distribution of ignition sources, but essentially none of whether LFL is the appropriate criterion, nor of which concentration is being discussed.

One point which does come out in these papers is that the estimates of ignition probabilities for a given cloud area, when all the factors which are considered are folded in, can vary by many orders of magnitude between say a rural night-time scenario and an urban day-time one. Viewed from this perspective, it may appear that the unaddressed question of LFL or 1/2LFL is not too significant. But we must be careful. If the calculations are, for example, being done for purposes of assessing the safety of a new urban development, then the concentration of the cloud in the area of the new development is important and at critical distances from the source, the portions of cloud above LFL or ½ LFL may define very different cloud areas, and thus in the framework of this model yield very different ignition probabilities (assuming all other, very uncertain, things are equal).

Spencer et al (1996,7) thus, whilst considering ignition sources in some detail, do not give any help on the effect of cloud concentrations on ignition.

4.2 THE APPROACH OF HEGADAS

The nature of the concentration is considered in the heavy gas code HEGADAS (also embodied in the suite “FRED”) of Shell Research. A user of HEGADAS may specify what concentration he wishes to see in the output with different time averages being related in the code by equation 1, discussed above. As already noted, a 10min average gives half the 18.75sec average concentration. These times do indeed fall within the area of applicability (10s-10min) of the formula discussed by Wilson (1995) and it may be possible to regard 18.75 sec as an “instantaneous” concentration which must be below LFL in order that the cloud is not ignitable. In that case the 10min average (common in many models and analyses of passive plumes) would have to be below 1/2LFL. Even so it is incumbent on the analyst to note (and the assessor to appreciate) which concentration he has used.

4.3 THE APPROACH OF DRIFT

The code DRIFT of AEA Technology approaches concentration fluctuations differently from HEGADAS. It uses the analysis of Davies of the results of Hall et al in order to estimate how far out its mean predictions of concentration might be – based purely on the knowledge that different “identical” releases will give rise to different measured concentrations.

DRIFT's mean concentration predictions are (by dint of the data against which it has been tuned) a cloud area average of 6sec average concentrations in the near field (heavy gas region) moving smoothly to 10min averages in the far field (passive plume region). The ad hoc justification is that short time averages will be more relevant for flammability in the near field and possibly longer ones for questions of toxicity in the far field.

In this way DRIFT can estimate not only the "expected" concentration c, but also percentiles c+ and c- (based on this prediction as a central value) such that the concentration would be expected to fall between c- and c+ (at a given position and time) in x% of nominally identical releases. Investigation indicated that 95% of cases should result in observation of concentrations between 1.8 and 0.46 times the mean predicted value.

If LFL is taken to be a physical criterion, then it may therefore be wise to demand that the upper 95% percentile is below LFL rather than the predicted mean. But then one asks why 95% instead of 90% or 99%, and one is thrown back to the inescapable fact that this is a probabilistic activity, and one with a considerable degree of uncertainty. Among this uncertainty, Jones at al (1994) observed the fluctuations in the open atmosphere are unlikely to be less significant than in a wind tunnel, and may be considerably more so.

Adopting this procedure would imply that the mean 6sec cloud area average concentration (NB not the 6sec average at a given point) should be below about 1/2 LFL.

4.4 MODEL ACCURACY

It is possibly productive to look at the question from another viewpoint. Models predict some form of averaged concentration. Measured concentrations from sensors in experiments fluctuate. Concentrations in different "identical" releases differ. The code DRIFT, discussed above, attempts to cope with this by comparison with cloud-area-averaged concentrations in the Thorney Island Trials, and by extracting an idea of variability from the repeated experiments of Hall et al.

But suppose we just take model predictions and compare them with the measurements of individual sensors in a selection of experiments. What should we expect? Well what we should not expect is to fit the data: we are predicting (some kind of) average and comparing with actual realisations. Hanna et al (1996) show the results of such a comparison which we reproduce here in Figure 4.

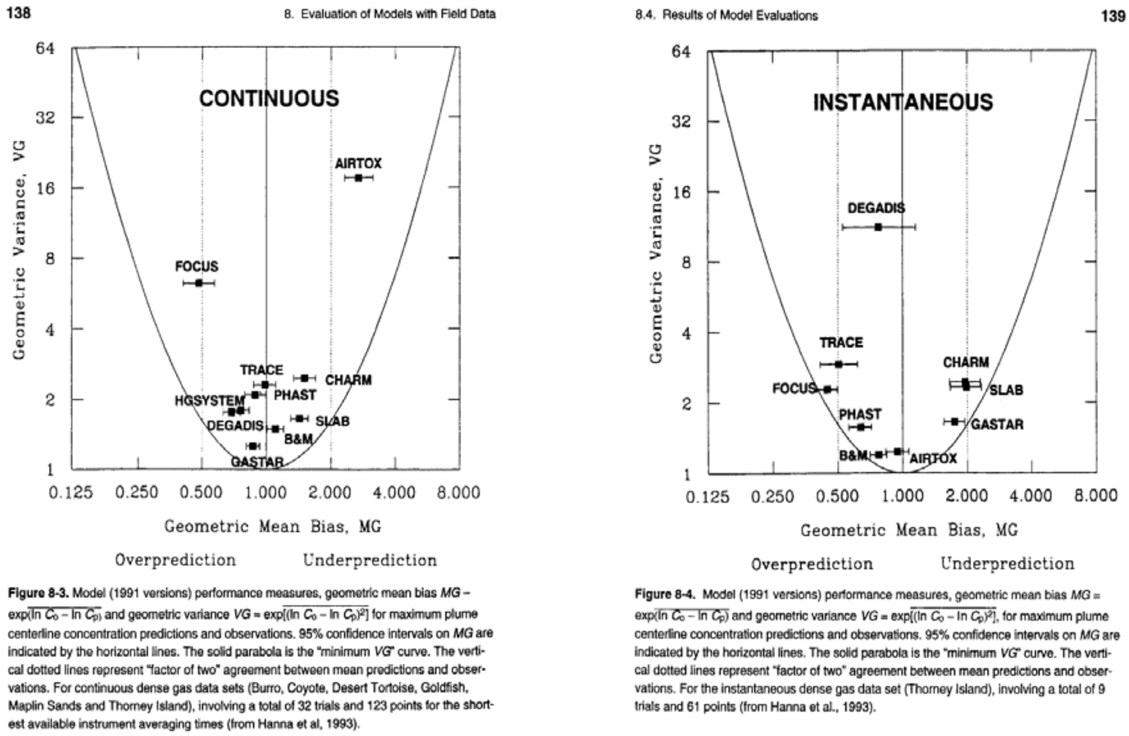

Figure 4. Comparison of model predictions against experiments reproduced from Hanna et al. (1996)

Models predicted mean concentrations typically within a factor of two (and better for continuous than for instantaneous releases) but showed a scatter from the data “geometric variance” also typically of order a factor of two. Interpreting this is difficult and one has to be cautious. It could be that this reflects fluctuations in the data, and this may be the biggest source of difference, especially if one considers that these are idealised experiments in the open, and that some of the models may have had their free parameters fixed by comparison with some of the same data. Allowing for a more complicated release scenario, which models will find more “difficult”, it is probably fair to conclude that a good model should be expected to do no better than fit measured concentrations to within a factor of two, and that a mediocre one may do a lot worse. Thus, if one believed LFL to be the appropriate criterion for the actual concentration (which we have seen may not always be the case) then a conservative interpretation of the results from one of the better models, would indicate that ½ LFL might be more appropriate for the predicted concentration.

5 CONCLUSIONS

There are degrees of safety. In deciding what is “safe enough” the assessors of COMAH safety cases have to draw an artificial line. Currently, in COMAH assessments involving fire safety, the cloud concentration is considered low enough when it is below ½ the lower flammable limit. Should one relax that and allow the lower flammable limit itself to become the decision point?

From the diverse evidence reviewed here, some things are clear.

1. For flammability questions one should be looking at the instantaneous concentration at a point. Much lower 10 minute (say) averages are not relevant. Cloud averages may have something to say about whether a flame will propagate through the whole cloud, but for purposes of ignition one should look at a single point.

2. Good models may predict the concentration to within a factor of two or so (either way), for releases which approximate well to the idealisations on which the models are based.

3. Repeated trials in a wind tunnel result in a distribution of measured concentrations which has a width of order a factor of two about the mean.

4. Even cloud area averages of measured 1s time averages exhibit fluctuations in concentration of up to a factor of 2.

5. Flames have been seen to propagate through the whole cloud when measurements near the ignition point reveal a concentration below the lower flammable limit. Flames have been seen to billow out and then extinguish when measurements near the ignition point reveal a concentration as low as half the lower flammable limit.

From this and the general discussion above it might just about be argued that a flame will probably not propagate through the whole cloud if it is ignited at a time when the nearest centre line concentration (at ground level in the case of non buoyant clouds) is below the lower flammable limit. But for this author the statistics of the spread around the (ensemble) average in an ensemble of nominally identical releases would make such an argument rather risky.

Add to that (if you will) the uncertainties in modelling (for idealised and other releases) and it becomes very difficult to argue with any conviction at all that the HSE’s procedure of regarding ½ LFL as the deciding line should be relaxed at all. The recommendation is that the HSE’s criterion should not be relaxed, and that the ½ LFL criterion should apply to an instantaneous (or very short time average) concentration. (And this implies an even lower criterion if long time averages – eg 10min – are predicted.)

6 REFERENCES

A D Birch, D R Brown, and M G Dodson (1981) “Ignition probabilities in Turbulent Mixing flows” Eighteenth Symposium (International) on Combustion, The Combustion Institute.

A D Birch, D R Brown, M Fairweather, and G K Hargreave (1989) “An experimental study of a turbulent natural gas jet in a cross-flow” Combust. Sci. and Tech. Vol 66 pp217-232

P W M Brighton (1985) “Area Averaged concentrations Height Scales and Mass Balances” J Hazardous Materials 11 (1985) pp189-208

J K W Davies (1992) “An analysis of some replicated wind tunnel experiments on instantaneously released heavy gas clouds dispersing over solid and crenellated fences” HSE Report IR/L/HA/92/8

J K W Davies (1994) “An analysis of data from some TNO-IMET continuous ground level release wind tunnel experiments with applications in dense gas dispersion modelling” HSE Report IR/L/RAM/93/03

A Evans and J S Puttock (1986) "Experiments on the Ignition of dense flammable gas clouds" 5th Intl Symposium on Loss Prevention and Safety Promotion in the Process Industries, Cannes, Sept 1986.

D J Hall, R A Waters, G A Marsland, and S L Upton (1991) “Repeat variability in instantaneously released gas clouds – some wind tunnel model experiments” Warren Spring laboratory Report LR 804 (PA).

D J Hall, V Kukadia, S L Upton, and G A Marsland (1991) “Repeat variability in instantaneously released gas clouds dispersing over fences – some wind tunnel model experiments” Warren Spring laboratory Report LR 805 (PA).

S R Hanna, P J Drivas, and J J Chang (1995) “Guidelines for use of vapor cloud dispersion models” Published by the centre for Chemical Process Safety of the American Institute of Chemical Engineers. (Interestingly the names of the authors of this book do not appear on the cover, nor on the inside title page(s) but only on p xii under “acknowledgements” where their authorship is acknowledged, presumably by someone other than the authors of the book.)

S J Jones, G A Tickle, and D M Webber (1994) “Incorporation of concentration fluctuations within dense gas dispersion models” - Vol 2, Part VII of UKAEA Report AEA/CS/FLADIS/1994 “Complex Features in Dense Gas Dispersion Modelling” (1994) by D M Webber, S J Jones, D Martin, G A Tickle, and T Wren.

J Liedtke and M Schatzmann (1997): Dispersion from Strongly Buoyant Sources. Final report, EU-project EV5V-CT-93-0262, University of Hamburg, Meteorological Institute (125 pages).

H Spencer and P J Rew (1997) “Ignition probability of flammable gases” WS Atkins Contract research Report 146 to HSE.

H Spencer, J Daycock and P J Rew (1998) “A model for the ignition probability of flammable gases” WS Atkins Contract research Report 203 to HSE.

M T E Smith, A D Birch, D R Brown, and M Fairweather (1986) “Studies of ignition and flame propagation in turbulent jets of natural gas, propane and a gas with a high hydrogen content” Twenty-first Symposium (International) on Combustion, The Combustion Institute.

D Wilson (1995) “Concentration Fluctuations and Averaging Time in Vapor Clouds” Published by the centre for Chemical Process Safety of the American Institute of Chemical Engineers.