Estimation of probability of detection curves based on theoretical simulation of the inspection process

Charles R A Schneider, Ruth M Sanderson, Capucine Carpentier, Lu Zhao and Channa Nageswaran

TWI
Cambridge CB1 6AL, UK
Telephone 01223 899000
Fax 01223 890952
E-mail charles.schneider@twi.co.uk

Abstract

Estimation of probability of detection (POD) curves by NDT typically relies on the manufacture of large numbers of realistic defect specimens, followed by practical trials of the inspection procedure. These are costly and time consuming activities. POD curves could be generated more rapidly and more cost-effectively if theoretical simulation of PODs were shown to be sufficiently representative of actual inspection performance.

This paper compares the predictions of two such simulation-based POD models with pre-existing evidence from EDF Energy's Capability Statement for manual ultrasonic testing (UT). One of the models is an in-house software tool called PODPEDGE, which calls the EDF Energy code PEDGE. The other POD tool is part of the CIVA code developed by CEA.

Overall the PODs predicted by both tools were consistent with pre-existing evidence. Where like-for-like comparisons were possible, there was also remarkably good agreement between them. In many respects, the CIVA software offers broader functionality than the PODPEDGE tool, eg different materials and component geometries. However, PODPEDGE has the advantage of faster run-times and can therefore produce more accurate point estimates of POD (within a specified run-time). PODPEDGE also includes an internal estimate of model accuracy, based on error flags available within the PEDGE code.

1. Introduction

Experimental determination of POD is challenging for industry because a large number of experimental trials are typically required in order to determine accurate estimates. Annis and Gandossi [1], for instance, recommend using at least 60 target flaws for hit/miss POD data. Such trials can be costly and time consuming. At the same time, the need for POD data is becoming greater since the use of probabilistic methods in safety justifications is becoming more widely accepted.

POD curves could, in principle, be generated more rapidly and more cost-effectively if theoretical simulation of PODs were to be made acceptably accurate, consistent and repeatable. Such models are called either model-based or model-assisted PODs [2], depending on whether the modelling results are supplemented by experimental data. The ultimate aim is to use simulation software to generate a distribution of flaw responses numerically from postulated distributions of so-called 'essential parameters', eg flaw locations and orientations [3]. The possibility of using theoretical considerations to estimate PODs in this way has been investigated by Walker et al [4,5]. Moreover, several organisations interested in the development of the topic of model-assisted POD (MAPOD) have set up a working group on this [6].

This paper considers two distinct simulation software tools in establishing the feasibility of generating theoretical POD curves. These two tools are described further in Section 3. Several case studies were explored to evaluate the capabilities of these simulation tools, a selection of which are presented in this paper. The first of these case studies was based on the Capability Statement produced by EDF Energy for manual UT of ferritic welds [7]. The overall aim of the work was to demonstrate the feasibility of estimating realistic POD curves for ultrasonic inspection.

2. The case studies

2.1 Case Study 1

2.1.1 Guiding principles

The initial Case Study was designed to be challenging for, but broadly compatible with, the conditions assumed in the Capability Statement produced by EDF Energy for manual UT of ferritic welds [7]. Our initial case studies also sought to use reasonably realistic distributions of flaw parameters, based on practical experience of centre-line solidification cracking in thick-section butt welds. Figure 1 shows the layout of Case Study 1.

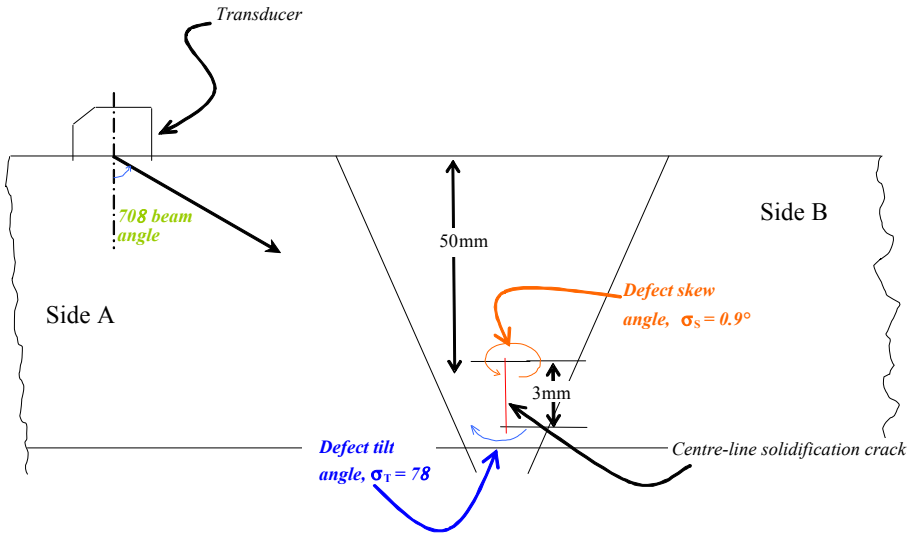

Figure 1. Inspection configuration for Case Study 1.

2.1.2 Flaw description

The assumed flaw description for the Case Study was as follows:

Location: Embedded, weld centre-line (ie well within the inspection volume).

Depth: 50mm to top edge (chosen to give limiting range of ~150mm with 70° probe).

Character: Smooth, planar, elliptical solidification crack.

Size: 3mm through-wall height by 15mm long.

Tilt: Beta distribution with range ±15° to the vertical, ie standard deviation $\sigma_T \approx 7°$. Beta distribution shape parameters used were $\alpha=2$ and $\beta=2$.

Skew: Normal distribution with mean of 0° and standard deviation $\sigma_S = 0.9°$.

This flaw description essentially satisfies the conditions stipulated by the EDF Energy Capability Statement, ie the flaw location, depth, character, size and tilt assumed in the Case Study are all within or equal to the values given in the EDF Energy Capability statement, as long as the weld is inspected from both sides (ie from both Sides A and B in Figure 1). In principle, the Case Study could generate defects with skews larger than the 3° assumed in the Capability Statement. In practice, however, the likelihood of this occurring for the specified distribution is ~0.1%. Thus, the aim of Case Study 1 was to confirm (or otherwise) and quantify the high value of the POD (close to 100%) that is anticipated under these conditions.

2.1.3 UT procedure

The assumed UT parameters were as follows:

Beam angle: 70° (other standard beam angles of 0°, 45° and 60° would make a negligible contribution to the detection of the above flaws).

Reporting threshold: 20dB more sensitive than a distance amplitude corrected (DAC) 3mm side drilled hole (SDH), ie 3mm SDH DAC −20dB.

Frequency: 5MHz.

Crystal size: 10mm diameter single crystal.

Shoe delay: Zero.

Pulse: A length of 3.5 cycles (between points 10% down from the peak amplitude) and a reasonably realistic pulse shape (based on Schneider and Chapman, 1997).

2.1.4 PODs generated

For this initial Case Study, a single POD value was computed, using PODPEDGE, for each of the following types of scan:

a) No raster scan – single probe position with beam axis lying on bottom edge of flaw;
b) A 1D line scan (transverse to the weld).

A line scan clearly requires more computation time than a single probe position, but may give a higher POD if the maximum UT response occurs at a different probe position from that in case a) above. The purpose of the line scan was to investigate the extent to which the POD tool needs to simulate the scanning of the UT probe to derive realistic POD estimates.

The POD was estimated as the proportion of simulations where the peak signal (from Side A and/or Side B as applicable) was above threshold. As for most of the PODPEDGE runs in this paper, the POD was calculated in three different ways:

- POD 'best estimate'. Here the peak signal is taken to be the largest of the signals listed in the tabular section of the PEDGE output file for which neither the shadow boundary flag 'B' nor the caustic flag 'C' is set (see Schneider and Chapman [8]).
- POD 'lower bound'. Here the peak signal is taken to be the largest of the signals listed in the tabular section of the PEDGE output file for which neither the shadow boundary flag 'B' nor the caustic flag 'C' nor the 'P' flag (which flags possible inaccuracy in the probe beam model) is set.
- POD 'upper bound'. This disregards all flags warning of possible model inaccuracy.

2.2 Case Study 2

2.2.1 Overview

Case Study 2 was very similar to Case Study 1, but used a normal distribution of defect skew, instead of a Beta distribution. This change was made once it was recognised that the Beta distribution is not currently an input option for the CIVA tool; the normal distribution was used instead in an attempt to facilitate comparison between PODPEDGE and CIVA. Figure 2a illustrates the layout of Case Study 2.

(a) Case Study 2 **(b) Case Study 3**

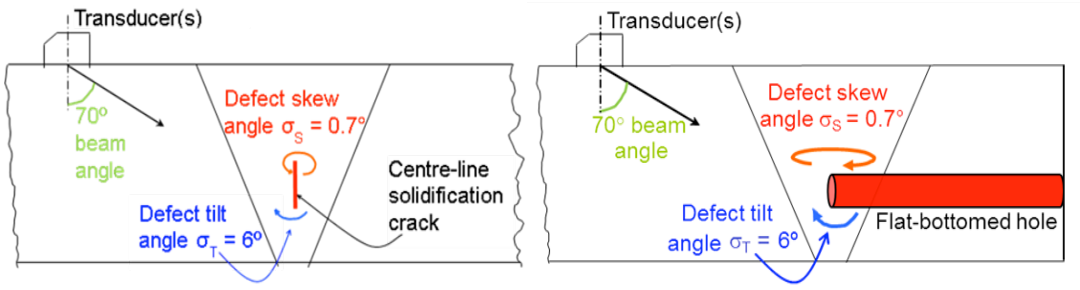

Figure 2. Inspection configuration for Case Studies 2 and 3.

The original intention was to simulate the same Case Study using both PODPEDGE and CIVA, but it later became apparent that the current release of CIVA was not able to simulate elliptical flaws. Because of this, semi-elliptical flaws were used instead in CIVA for just one of the PODs in this Case Study (corresponding to item 1 under Section 2.2.3 below). Case Study 3 was subsequently introduced to allow direct comparison between the PODPEDGE and CIVA tools for a simpler shape of reflector (flat-bottomed holes).

The component and basic UT parameters were identical to those in Section 2.1 (for Case Study 1). The flaw description was also identical (initially), with the following exceptions:

Tilt: Normal distribution with zero mean and standard deviation $\sigma_T = 6°$

Skew: Normal distribution with zero mean and standard deviation $\sigma_S = 0.7°$

2.2.3 PODs generated

1. A single value POD for flaws 3mm through-wall by 15mm long at 50mm depth, with a reporting threshold 20dB more sensitive than a 3mm SDH, as used in standard EDF Energy procedures, for each of the following types of scan:
   a) No raster scan – single probe position with beam axis lying on bottom edge of flaw;
   b) 1D line scan (transverse to the weld);
   c) 2D raster scan.

   In each case, the PODs were calculated for the case of inspection from both sides of the weld (A and B). As in Case Study 1, the purpose of the 1D and 2D scans was to investigate the extent to which the tool needs to simulate the scanning of the UT probe to derive realistic POD estimates.
2. POD curve based on successively less sensitive reporting thresholds (no raster scanning). For this case onwards, inspection was from just one side of the weld (to ensure compatibility with later CIVA runs).
3. POD curve varying through-wall size, but retaining a fixed aspect ratio of 5:1. The top edge of the flaw was fixed at 50mm depth, and the reporting threshold was 10dB more sensitive than a 3mm SDH (designed to maximise the observed variations in POD). 2D raster scanning was undertaken.

In each case, the POD was estimated, using PODPEDGE, as the proportion of simulations where the peak signal was above threshold. As in Case Study 1, the PODs were calculated in three different ways, ie 'best estimate', 'lower bound' and 'upper bound'.

2.3 Case Study 3

Case Study 3 was essentially the same as Case Study 2 (see Figure 2b), but used flat-bottomed hole (FBH) reflectors instead of elliptical (or semi-elliptical) flaws, thus allowing direct comparison between PODPEDGE and CIVA. The top edge of the reflectors was at a fixed depth of 50mm and a POD curve vs. through-wall size was generated, as for the final POD curve produced for Case Study 2 (ie item 3 above). For this Case Study, however, we reverted to the original reporting threshold of 20dB more sensitive than a 3mm SDH (as used in the standard EDF Energy procedures). This was designed to make the variations in POD more evident, on the basis that the FBHs generally gave smaller signals than the elliptical flaws. 1D scans (transverse to the weld) were used for this Case Study, primarily to make the CIVA run times more manageable.

3. The models

3.1 PODPEDGE

The PODPEDGE model is designed to predict the POD of smooth planar flaws of a specified size, based on conventional pulse-echo UT. Flaw size is usually treated as the independent variable for POD curves (called the 'characteristic' parameter by Benoist and Calmon, 2011). The material under test is assumed to be ferritic steel with negligible attenuation (ie fine-grained). The scanning surface is assumed to be smooth and flat. The flaw is also assumed to be smooth and flat with sharp edges and no liquid or oxide inclusion.

PODPEDGE is based on the EDF Energy PEDGE model which uses the GTD scattering theory. It is well known that (at least in its simplest form) GTD fails at two types of mathematical singularity, called 'reflection boundaries' and 'caustics' respectively. It is not yet clear how important 'caustics', in particular, are to the prediction of PODs in general. Earlier unpublished work by TWI showed that caustic crossings had a dramatic effect on the predicted signal amplitudes for two particular inspection configurations (which were closely related to Case Studies 1 and 2 of this paper). Section 0 further discusses the effect of reflection boundaries and caustics on the PODs predicted for the selected case studies.

PODPEDGE is able to use Monte Carlo simulation to calculate the response from flaws with a set of distributions for the following randomised parameters: flaw skew, beam angle and flaw tilt. In this paper, however, only the flaw skew and flaw tilt were randomised, ie defined by probability distributions. The beam angle was treated as fixed; in practice, random variations in beam angle have a very similar effect on POD as the equivalent variations in flaw tilt. In PODPEDGE, a given flaw is considered to be detected if and only if at least one of the signals received from it lies above a defined reporting threshold. Each POD predicted in this paper by PODPEDGE was based on 2000 simulations (or 'plays').

3.2 CIVA

CIVA (version 10.0 onwards) also offers the possibility of using its deterministic modelling capability to generate POD curves using the Monte Carlo concept. There are two approaches available in the software to generate POD: the 'hit/miss' and the 'signal response' methods. In this paper only the 'hit/miss' method will be considered. The 'signal response' method effectively assumes that signal amplitudes follow a linear trend vs. the independent variable (eg flaw size), which is not considered appropriate for the case studies in this paper.

In the 'hit/miss' approach (as utilised in this study), the detection probabilities are first estimated as the ratio between the number of hits (where the signal amplitude was above the threshold) and the total number of simulated attempts (as in PODPEDGE). CIVA also fits a log-odds function through the estimated detection probabilities. This effectively serves as a smoothing function, and is of the form [9]:

$$POD(a) = \left[1 + \exp\left(-\frac{\pi}{\sqrt{3}}\left(\frac{\ln(a) - \mu}{\sigma}\right)\right)\right]^{-1} \quad (1)$$

where a is the flaw size, μ is the natural logarithm of the crack size for which there is 50% detectability, and σ is the scale parameter that determines the flatness of the POD function. The parameters μ and σ are computed as maximum likelihood estimates using the estimated detection probabilities. In this paper, we focus on the raw estimates of the detection probabilities (before any smoothing), since this provides a like-for-like comparison with PODPEDGE. In practice, however, the smoothing function would help to mitigate against random sampling errors in the detection probabilities, which can be

important because CIVA is unable to use such large sample sizes for the Monte-Carlo simulation as PODPEDGE can (mainly because it requires longer computation times).

In CIVA, the POD is calculated for a range of values for the characteristic parameter, for example the height of the flaw; unlike PODPEDGE, CIVA does not allow a point estimate of POD based on a single value of the characteristic parameter. The CIVA runs in this paper were performed using 31 characteristic values. In order to be comparable with the initial cases run with PODPEDGE, the characteristic value was initially defined between 1 and 8mm, with a distribution that was concentrated around 3mm (which was the through-wall size of interest used for the initial PODPEDGE runs). The POD was then determined from the number of simulations or 'plays' implemented for each value of the characteristic parameter. 200 plays were simulated for each flaw height in Case Study 2, whereas 150 plays were simulated for each flaw height in Case Study 3.

The randomised parameters (called 'uncertain' parameters in CIVA) were then specified which, in our case studies, were the flaw tilt and the flaw skew (with distributions as specified in Section 0). The uncertain parameters are distributed randomly over the grid of characteristic parameter values and the defined plays.

All the CIVA results in this paper were generated using the Kirchhoff module only; it was not possible to compute any POD data using the GTD module in CIVA for the case studies in this paper.

4. Results and discussion

4.1 Case Study 1: 3 x 15mm flaw, 3mm SDH –20dB threshold

Firstly, the Case Study was computed using PODPEDGE with inspection from both sides of the flaw (ie Sides A and B in Figure 1), with no raster scan, for initial comparison with the Capability Statement [7]. The result was a POD of 98.7% (with lower/upper bounds of 98.6% and 99.1% respectively). This POD is regarded as sufficiently high to be in reasonable agreement with the Capability Statement. The same POD values were predicted when a 1D scan was performed.

Note that the Capability Statement takes credit for 2D scanning of the probe. If a 2D scan had been performed in this Case Study, it may be that the predicted PODs would have been even closer to 100%.

4.2 Case Study 2: 3 x 15mm flaw, 3mm SDH –20dB threshold

4.2.1 PODPEDGE

This case was very similar to Case Study 1, but with slight changes to the distributions of defect tilt and skew (as described in Section 2.2.2). Table 1 summarises the PODPEDGE results for Runs 1a to 1c of Section 2.2.3, for inspection from both sides of the weld (A and B).

These predictions are again broadly consistent with the EDF Energy Capability Statement [7]. For this Case Study, it is clearly not necessary to simulate the scanning of the UT probe (in 1D or 2D) to derive reasonably realistic POD estimates.

Table 1. PODPEDGE results for Runs 1a to 1c of Section 2.2.3 for inspection from both sides of the weld (A and B).

	Predicted PODs		
	Best estimate	Lower bound	Upper bound
No raster scan	94.7%	94.7%	100%
1D line scan	95.7%	95.4%	100%
2D raster scan	95.5%	94.9%	100%

4.2.2 CIVA

The CIVA tool predicts PODs based on inspection from just one side of the weld (Side A or Side B). However, unlike PODPEDGE, CIVA does not currently combine the results from inspecting both sides of the weld. Three separate POD analyses were therefore carried out, each with a different range of tilts derived from the same normal distribution of zero mean and standard deviation of 6°. The different distributions used are summarised in Table 2. The first two scenarios use just half of the parent normal distribution. The distribution for Scenario 1 is that half of the distribution with negative (favourable) tilt, whereas that for Scenario 2 is that half of the distribution with positive (unfavourable) tilt; the sign of the tilt is such that negative tilt values are those closer to normal incidence, whereas positive values are for flaw orientations further away from normal incidence to the beam. POD Scenario 1 is comparable with the PODPEDGE runs from both sides of the weld if it is assumed that the maximum signal is always from the side of the weld with the lower (more favourable) angle of incidence. For each POD analysis the value and distribution of the skew is the same.

Table 2. Distribution of the uncertain parameters (tilt and skew) for the three scenarios.

	POD Scenario 1		POD Scenario 2		POD Scenario 3	
	Tilt	Skew	Tilt	Skew	Tilt	Skew
Maximum value	–9.995°	–2.096°	0	–2.096°	–9.993°	–2.096°
Minimum value	–0.002°	2.538°	9.987°	2.538°	9.994°	2.538°
Mean of parent normal	0°	0°	0°	0°	0°	0°
Standard deviation of parent normal	6°	0.7°	6°	0.7°	6°	0.7°

Using the Kirchhoff model with a threshold of 20dB from 3mm SDH at the same range as the flaw, the POD of a 3mm high flaw is 100% for Scenario 1 (when the value of the tilt was negative), 96% for Scenario 2 (when the value of the tilt was positive) and 99% when both tilt orientations are taken into account. These results were without any raster scanning. The POD predicted for Scenario 1 is consistent with the EDF Energy Capability Statement and the range of PODPEDGE predictions (95%-100%) in Table 2, even though the assumed flaw shape is here semi-elliptical rather than a full ellipse.

4.3 Case Study 2: POD curve vs. inspection sensitivity for a 3 x 15mm flaw

Figure 3 illustrates the variation in POD with inspection sensitivity predicted by PODPEDGE. In this case study, the predicted ‘best estimate’ POD was equal to the

lower bound in every case. Figure 3 shows the expected trend of increasing POD with increasing sensitivity. The difference between the lower and upper bound POD predictions can be attributed to flaws where PEDGE indicates inaccuracy due to the proximity of a mathematical singularity called a 'caustic'. For inspection sensitivities above DAC −9dB, the potential inaccuracy only amounts to ~5% in the resulting POD, but the uncertainty in the prediction becomes increasingly important at lower sensitivities.

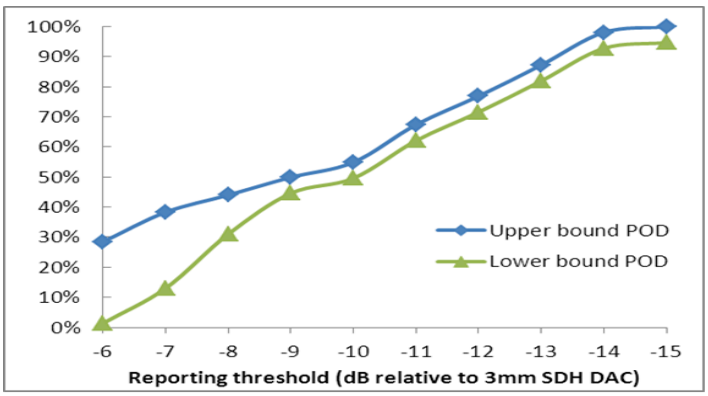

Figure 3. PODs predicted by PODPEDGE vs. inspection sensitivity for 3 x 15mm elliptical flaw.

Figure 3 indicates that 50% POD is reached at a threshold of about DAC −10dB. This threshold was therefore used for the ensuing study of POD vs. flaw size to allow maximum discrimination of either any increases or any decreases in POD relative to this starting level of ~50%.

4.4 Case Study 2: POD curves vs. flaw height for elliptical flaws

Figure 4 illustrates the variation in POD with flaw height predicted by PODPEDGE for elliptical flaws having a fixed aspect ratio (ie length:height) of 5:1. Best estimate PODs, together with upper and lower bounds, are given for three different scanning increments in the secondary scanning direction (ie along the length of the weld). This illustrates that, at least in some cases, the assumed scan pattern has an important effect on the POD. The best estimate POD in Figure 4 at a flaw height of 3mm, for instance, varies from 36% to 73% according to the secondary scanning increment. A 1D scan would result in a smaller POD than any of these 2D scans so, by implication, the use of a 1D scan (or no scanning) could also have a substantial impact on the predicted POD.

The POD has a local maximum (corresponding to ~70% POD) at a flaw height of ~3mm and then gradually drops to ~20% POD (at a flaw height of ~5mm), after which it steadily increases again, following a more intuitive trend for larger flaw heights. The presence of a local maximum in the POD is counter-intuitive and contrary to the generally held belief that POD increases monotonically with increasing flaw height. It is, nevertheless, believed to be a genuine effect (at least for the idealised elliptical flaws considered here). It arises essentially because the bottom edge response is enhanced when the probe lies near the caustic, ie the line that lies at right angles to the plane of the flaw and that passes through its centre of curvature [10].

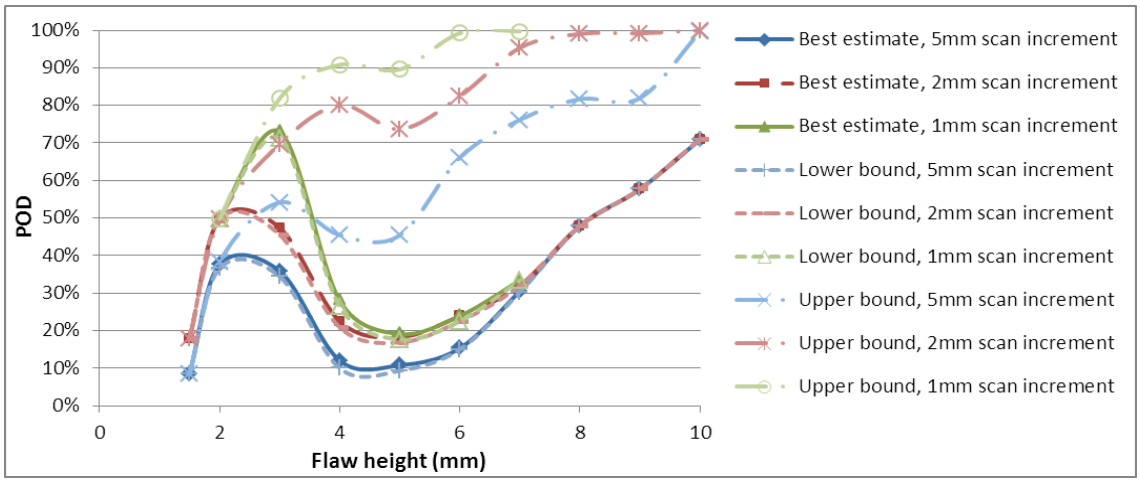

Figure 4. PODs predicted by PODPEDGE vs. flaw height for elliptical flaws having a fixed aspect ratio (reporting threshold DAC –10dB).

The lower bound POD is almost identical to the best estimate POD in Figure 4. However, the upper bound POD has a much smoother appearance and, for a 1mm secondary scanning increment, is almost monotonic. This confirms that the trough in Figure 10 at ~5mm flaw height is largely due to the proximity of the caustic. PEDGE is known to be potentially inaccurate in these regions, so the true peak amplitude is unknown. However, it is considered likely that the true peak amplitudes in these regions would be above threshold due to the favourable focussing of the diffracted rays, ie it is possible that the upper bound predictions in Figure 4 are realistic even though the PEDGE amplitudes on which they are based are likely to have been overestimated.

In many of the cases studied, there is very little difference between the 'best estimate' POD and the upper and lower bounds. However, Figure 4 shows that, in some cases, the difference between the 'best estimate' POD and the 'upper bound' POD can be as large as 70% (due to the proximity of the caustic). This highlights the importance of calculating both the 'best estimate' and the 'upper bound' PODs, in practice, as an internal check on possible model inaccuracy (at least for GTD-based models).

4.5 Case Study 3: POD curves vs. flaw height for FBHs: 3mm SDH –20 dB threshold

Figure 5 compares the PODs predicted by PODPEDGE and CIVA vs. FBH diameter. Unlike the comparison in Section 4.2, this was a direct like-for-like comparison of the two models. A threshold level of 3mm SDH – 20dB was assumed and the probe followed a 1D line scan in each case.

The overall agreement between the two models is remarkably good, especially given that they are based on different semi-analytical approximations (GTD and Kirchhoff theory). The PODPEDGE POD curve has a smoother appearance than the point estimates predicted by CIVA because each PODPEDGE result was based on 2000 simulated flaws, as compared to just 150 simulations for each CIVA prediction. The random sampling errors in the predictions will follow a binomial distribution. This means that the standard deviation of the random sampling errors in the CIVA predictions is typically ~4% in POD as compared to ~1% for PODPEDGE. In practice,

application of the log-odds smoothing function provided by the CIVA software is expected to substantially reduce the effect of such sampling errors (see Section 3.2).

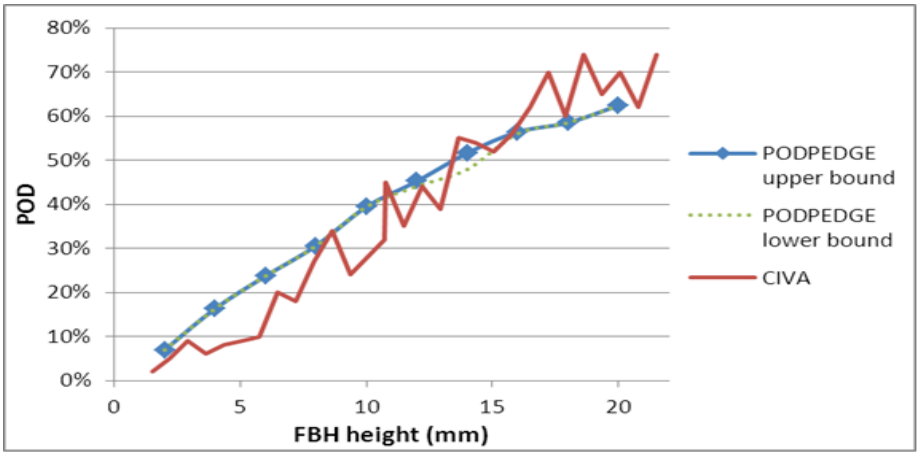

Figure 5. Comparison of PODPEDGE and CIVA predictions for FBHs.

5. Model assisted PODs

In the context of this paper, a POD is regarded as 'realistic' if it is representative of the performance that could be achieved under carefully-controlled laboratory conditions, ie the inherent performance of a given UT procedure. Thus both of the POD tools described in this paper predict the inherent likelihood that a given UT procedure will detect flaws having parameters drawn from specified distributions. Neither tool makes any allowance for experimental variability (eg due to temperature variations or variable coupling) or for human error. Inclusion of experimentally observed factors such as these is what distinguishes the 'model-assisted' probability of detection (MAPOD) approach from the 'model-based' approach in this paper. In practice, these extra factors (in particular, the human factors) mean that it is never possible to achieve 100% POD. In principle, the POD tools described in this paper could be adapted to implement the MAPOD approach. This could be done by experimentally measuring the UT signals from a set of known flaws and fitting a probability distribution $f(\delta)$ to the differences δ between the measured signals and the predictions of a deterministic model (eg PEDGE). The POD tool (eg PODPEDGE) would then be modified to adjust each predicted signal by an amount randomly selected from the distribution $f(\delta)$ to simulate the extra effect due to experimental factors.

6. Conclusions

- Both the PODPEDGE and CIVA tools predicted reasonably realistic PODs for the selected case studies. Where like-for-like comparisons were possible, there was also remarkably good agreement between the two tools and with pre-existing evidence in the literature.
- The PODPEDGE tool provides an internal estimate of model accuracy. This feature is particularly valuable in quantifying the uncertainties in the predicted POD that can arise due to mathematical singularities called 'caustics'.

- The assumed scan pattern can have a major impact on the predicted PODs for elliptical flaws. A realistic scan pattern should therefore be simulated wherever practicable.
- In general, the CIVA software offers greater functionality than PODPEDGE, eg different materials and component geometries. However, PODPEDGE has the benefit of faster run-times and an internal estimate of model accuracy. PODPEDGE is also able to simulate different defect shapes from CIVA (eg elliptical defects). Where possible, the accuracy of the CIVA tool should be benchmarked against PODPEDGE for an equivalent problem that lies within the scope of the PODPEDGE tool.

Acknowledgements

The authors are grateful to Dr A E Walker (Rolls-Royce), Dr R K Chapman (EDF Energy) and Prof. J F Knott for their advice and assistance. The work was funded by Industrial Members of TWI, as part of TWI's Core Research Programme.

References

1. Annis C and Gandossi L, 2012: 'Influence of sample size and other factors on hit/miss probability of detection curves'. European Network for Inspection and Qualification (ENIQ) report no 47, EUR 25200 EN – 2012, January, Publications Office of the European Union, Luxembourg. ISBN 978-92-79-23018-9.
2. Benoist P and Calmon P, 2011: 'Reliability of nondestructive techniques based on a statistical simulation approach'. Proc. IIW Meeting, Chennai, document V-1509-11.
3. ENIQ, 2005: 'ENIQ Recommended Practice 1: Influential/essential parameters'. European Network for Inspection & Qualification (ENIQ) report 24, Issue 2, EUR 21751 EN, Office for Official Publications of European Communities, Luxembourg, June. http://capture.jrc.ec.europa.eu/docs/eniq/rec_prac/RP1%20(EUR21751).pdf [accessed 26 April 2012].
4. Walker A E, Stevens A and Cool T, 2009: 'Use of Monte-Carlo methods to derive quantitative probability of defect detection data', Nuclear Future, Vol 5 No 1, pp 31-21, January, ISSN 17452058.
5. Walker A E and Daniels W, 2009: 'Use of Monte-Carlo methods to derive quantitative probability of defect detection data', 4th European-American Workshop on Reliability of NDE, Berlin, 24-26 June.
6. CNDE, 2012: Iowa State University Centre for Nondestructive Evaluation (CNDE) MAPOD Working Group website www.cnde.iastate.edu/MAPOD [accessed 26 April 2012].
7. Chapman R K and Bowker K J, 2001: 'The production of capability statements for standard NDT procedures', Insight, Vol 43 No 1, pp36-38, January.
8. Schneider C R A and Chapman R K, 1997: 'NDT Group computer programs for theoretical modelling of ultrasonic inspection of smooth planar defects – user guide', Nuclear Electric report EPD/GEN/REP/0066/96, Issue 3.
9. Jenson F, Iakovleva E and Reboud C, 2009: 'Assessment of inspection performance using simulation supported POD curves'. *PVP2009*, July 26-30, Prague.
10. Chapman R K and Coffey J M, 1982: 'Ultrasonic scattering from smooth flat cracks: edge wave signals from finite cracks with curved edges'. CEGB report, NW/SSD/82/0034/R.