Automated Sizing and Classification of Defects in CANDU Pressure Tubes

Timothy Lardner[a], Graeme West[a,*], Gordon Dobie[a], Anthony Gachagan[a]

[a]*Department of Electronic and Electrical Engineering, University of Strathclyde, 204 George Street, Glasgow G1 1XW, United Kingdom*

Abstract

Pressure tubes within CANDU reactors are subject to frequent ultrasonic non-destructive examination to identify and characterize any defects that pose a risk of initiating Delayed Hydride Cracking (DHC), a well-known problem that occurs in zirconium components that are subject to high mechanical and thermal stresses. The analysis of the ultrasonic data gathered from the pressure tubes is often a long and repetitive process on the critical path to the restart of the reactor. Motivations for developing an automated system include saving time on the critical path, minimising human subjectivity from the process and increasing repeatability of measurements. An automated system providing decision support to analysts also reduces the risk of operator fatigue by minimising time spent processing routine defects. This paper describes a novel system for the automated analysis of CANDU pressure tube inspection data. The system automates the entire initial examination process including the complex decision-making procedures by implementing an expert system with an associated rule-base. Results from this system are detailed, illustrating the location and characterisation of defects and key features within the pressure tube with a high degree of accuracy while increasing repeatability through the removal of subjectivity in measurements.

Keywords: CANDU, Ultrasound, Automation, NDE
PACS: 29.85.Fj, 43.35.Zc, 43.40.Le

1. Introduction

A large proportion of nuclear power plants worldwide are entering the latter stages of their operational lifetimes. To ensure safe continued and extended operation, there is an increasing need to gather and assess data which provides evidence of the health of key components. Plant operators of nuclear plants need to monitor reactors for structural changes as part of ageing management programs (Aho-Mantila et al., 2012). As well as online data gathering, outages are scheduled to allow for critical components to be inspected. In recent years, non-destructive testing technology has evolved such that great volumes of data can be gathered in a relatively short period of time. While previously storage, bandwidth and computing power was a bottleneck for inspection systems, the speed of human operators limits the practical capabilities of modern systems.

This paper explores a specific case where automation can be utilised as part of the analysis procedure for a modern reactor inspection system, namely the initial assessment of detected flaws to determine whether they meet a set of acceptable criteria. The paper is organised into five sections. Section 1 introduces the problem area and provides an overview of the proposed solution. Section 2 outlines the current inspection procedure. Section 3 provides details of the automated analysis system and discusses the modular approach. Key results and their implications are discussed in Section 4, while Section 5 provides a conclusion to the paper.

*Corresponding author
Email address: graeme.west@strath.ac.uk (Graeme West)

1.1. CANDU Reactor Inspection

Understanding the health of critical components of a nuclear power plant is key to the ongoing operation and future lifetime extension of such plants. A cornerstone of this activity is thorough inspection during planned statutory outages. In CANDU (Canada Deuterium-Uranium) reactors a representative subset of the fuel channels, usually around 20 out of 400, are subject to detailed regular inspection (Kwak et al., 2005) with each reactor inspected at least once every three years. The selection of these channels is governed by a number of factors, such as the desire to revisit a channel that has been previously inspected, targeting channels exhibiting unusual behaviours as indicated by online monitoring, and the need to have a representative subset of the full core. The primary means of locating flaws within a CANDU fuel channel is by using a set of ultrasonic measurement devices and manually assessing the gathered data to identify and assess any defects with respect to the initiation of Delayed Hydride Cracking (DHC).

DHC is a major concern in reactor designs which utilize zirconium alloys and is caused by increased hydride concentration coupled with the increased mechanical stresses associated with these defects as well as the heating and cooling cycles of the reactor (Kim et al., 2003; Puls, 1997; Sahoo and Pandey, 2010; Radu and Roth, 2012). DHC can lead to leaks within the reactor core which can potentially damage the reactor (Puls et al., 1998). Aside from manufacturing flaws, the most common areas where hydride build-up can be identified are around pressure tube defects caused by in-service operation (Cheadle et al., 1987). Early identification of these defects is paramount to reducing the probability of unplanned outages and maximizing both safety and income.

To provide a high degree of confidence that all defects have been correctly identified and classified from the channel data gathered during an outage, two independent analysts manually assess the inspection data, with a third resolution analyst providing oversight and arbitration for any disagreements (Obrutsky et al., 2007). A case is made to restart the reactor based on the information gathered about the defects in the pressure tubes, ensuring that no flaws within the reactor will compromise the structural integrity over the number of expected thermal cycles.

1.2. Research Motivation

Classification and sizing of defects in pressure tubes is time-consuming and laborious, and lies on the critical-outage path as the reactor cannot be restarted until the inspection and subsequent analysis is complete. The desire for automation was highlighted in Obrutsky et al.'s (2007) work, however it is only comparatively recently that new inspection systems are able to record volumes of data that are impractical to manually analyse. An automated process saves time on the critical outage path while providing a repeatable decision support framework that is not affected by the subjectivity of a human analyst. Furthermore, such a system reduces the risk of operator fatigue by streamlining an analyst's workflow and ensuring they are spending the most time examining complex regions of the data.

1.3. Overview of Proposed Methodology

This paper presents a system for providing automated decision support to the analysts, by providing a robust, reliable and repeatable means of automatically identifying, sizing and classifying defects within the ultrasonic data.

Furthermore, this system automates the full analysis process from beginning to end requiring no input from an analyst. This process will hereafter be referred to as being 'end-to-end', taking the raw ultrasonic data as an input and outputting a file that can feed into the existing analysis workflow. The system adopts a knowledge-based approach to processing the data, incorporating algorithms and rule bases derived from the domain knowledge of experts, therefore allowing for the decisions at each stage of the process to be auditable and understandable by all involved in the process. This differs from other published work in this field (Sambath et al., 2010; D'Orazio et al., 2008) that either requires input from an analyst, or utilises machine learning where the factors involved in decisions are not always readily understood.

The result of the end-to-end process is a document of each feature and flaw found within a pressure tube. This is reported to the analyst as a 'channel file', a file compatible with the proprietary inspection software that can be used to review the ultrasonic data and the measurements made.

The novelty in this approach is the fact that the automation is not limited to a subset of the existing analysis process. Instead, the full analysis process from beginning to end is automated in a robust, repeatable and explicable manner that can be run in parallel with a human analyst. The results from this system are then presented in a way already familiar to analysts through a bridge that allows users to view results directly within the existing inspection software.

1.4. Key Related Work

Much of the work to date has focussed on the development of tools and techniques which support the analytical evaluation of flaws which have not met the initial acceptable criteria by examination. For example, Oh and Chang (2014) describes an integrated probabilistic approach to the assessment of these types of flaw, but focusses on those which have not met the initial acceptable criteria. Similarly Cho et al. (2016) provides an overview of recent work in the prediction of DHC from in-service inspection flaws, but also tackles those flaws already identified as not having met the acceptable criteria by examination. Tackling the initial task of identifying those flaws which have not met the initial acceptable criteria through the application of automated decision support techniques has received limited attention.Ciocan and Ciobanu (1998) demonstrated a prototype neural network that showed the potential to be able to detect a number of different sized flaws within pressure tube inspection data. The primary drawback of this technique is the black-box method in which the data is processed. For a nuclear application, explicability is paramount which is why an expert system was identified as the most suitable method of automating processes within the overarching system.

A review of expert systems and their applications within a CANDU system was undertaken by Uhrig (1989). This review focused on processing the output of condition monitoring systems rather than the systems themselves. Anderson (1990) had similar findings to Uhrig and also identified condition monitoring as a potential application for expert systems. However, no implementation similar to the proposed method was mentioned in either of these studies. More recently, Garcia et al. (2005) applied expert domain knowledge to failure mode and effects analysis, creating a set of logical inferences that inform a model designed to calculate the probability of failure rates. Yang et al. (2016) demonstrated a rule-based expert system for providing decision support to plant personnel to conduct risk monitoring. As well as presenting an effective risk management system, the authors also demonstrated that expert systems are a viable and relevant solution to automating and providing decision support within nuclear power plants. Coble et al. (2015) undertook a review of the current methods for prognostics and health management, citing Venkatasubramanian et al.'s (2003) series of papers on fault detection and diagnosis where the authors discuss applying expert systems to the field. It was stated that as such systems are widely used in an industrial setting, a key advantage of expert systems is the ease of development and transparent reasoning. Coble et al. also highlighted that aside from Meyer et al.'s (2014) work, there is little research into flaw detection and characterisation following non-destructive examination for passive components such as pressure tubes.

Murray et al. (2016) presented a method for processing visual inspection data from Advanced Gas-cooled Reactor cores utilising knowledge from domain experts to develop their technique, and have since incorporated feature detection algorithms into their software. This approach is similar to our proposed solution, in terms of applying signal and image processing techniques to the raw data and in parallel with existing human operators. Our solution interfaces with software that analysts currently use and are comfortable with, reducing barriers to adopting the new approach. West et al. (2012) implemented a knowledge-based expert system to process Fuel Grab Load Trace data in AGR cores. This approach is fundamental to many of the modules within the proposed end-to-end system and a the rule-based module described later in this paper uses the the same framework presented by West et al to implement logical rules based on the domain knowledge of experts.

2. Current Inspection and Analysis Procedure

2.1. Ultrasonic Data Acquisition

The inspection of CANDU pressure tubes are currently performed using a multi-purpose tool known as Channel Inspection and Gauging Apparatus for Reactors (CIGAR) though with the increased potential of transferring and storing large datasets, new avenues are being explored for exploiting this capability. Consequently, new systems have been developed that gather far more data than the existing, established inspection tools.

2.1.1. CIGAR

The CIGAR tool has six ultrasonic probes used for defect detection, sizing and classification (Trelinski, 2008). There are two normal incidence probes, one at 10MHz and another at 20MHz. The normal incidence probes are always operated in pulse-echo mode. There are four angled shear wave probes that operate at a centre frequency of 10MHz. These probes can be operated in either pulse-echo or pitch-catch mode and are orientated around the 20MHz normal incidence probe so that any point on the pressure tube will be inspected from above as well as from the front, back and each side via full-skip propagation.

Inspection of pressure tubes using the CIGAR tool is a two stage process. The first stage acquires pulse-echo data from each probe and applies gating to the data so that the signals reflected from the pressure tube interface, outer diameter and inner diameter are isolated. These files are known as General Helical (GH) datasets and can be manipulated to allow an analyst to view a top-down, C-scan presentation of the data and to review the data for differences in amplitude response throughout the pressure tube which could indicate the presence of a defect.

The second stage of the acquisition process records the full signals from the 20MHz normal incidence probe, as well as both the axial and circumferential pitch-catch pairs. These datasets are known as B-scans and are used to measure the depth of a flaw once it has been located using the initial dataset.

2.1.2. BRANDE

A newly developed inspection system known as BRANDE uses the same ultrasonic probe configuration as CIGAR but records B-scan data over the entire pressure tube and at a higher resolution compared to CIGAR. The B-scan acquisition occurs at the same time as the C-scan which overcomes a disadvantage of CIGAR where misalignment can occur between the two scans, making it more difficult to evaluate a potential defect. The increased resolution of the BRANDE tool, and the fact that B-scan data is always recorded increases the size of BRANDE datasets to around 1000 times larger than those of CIGAR. The limitation of processing this additional data lies with both the existing analysis software which is not designed to cope with the larger datasets, and the analysts themselves who are not able to process the additional volume of data in the time-scales required.

2.2. Existing Analysis of Pressure Tube Data

The ultrasonic data recorded during an outage is initially processed in parallel by two independent analysts using a proprietary software package, FLAW (Ontario Power Generation, Canada), to visualize the data and to create and record measurements using inbuilt tools. Each analyst will review all of the available data in order identify any potential defects with the tubes. An analyst locates potential defects by finding regions where the amplitude response of any shear wave probe is greater than a predefined threshold. The analyst will measure the length, width and depth of any flaws as well as assigning a classification with the likely root cause of the defect, such as bearing pad fretting or debris.

Internal operator guidelines are followed for the sizing and classification of flaws, but a degree of subjectivity is introduced by analysts leading to potential differences between measurements. A third resolution analyst will therefore compare the measurements from each analyst and resolve any discrepancies between the two results, in consultation with the raw data and further discussions with the analysts. This provides the existing, robust method for assessing pressure tubes.

The final analysis result, known as the verified result, is then passed to a team of mechanical and structural engineers to determine if all defects fall within conservative tolerances and that the risk of hydride build-up is minimal over the number of expected thermal cycles. The engineers will first model each defect assuming a worst-case 'V-shaped' flaw using the measured dimensions. If the model indicates that there is a risk of DHC initiation within the defect region, a replica of the flaw is made using an impression tool. The replica measurements can be significantly more accurate than the ultrasonic measurement, allowing structural engineers to better analyse the flaw and make the restart case for the reactor.

The analysis process is time consuming and lies on the critical outage path, as the reactor cannot be brought back online until the inspection process is complete and all flaws have been verified to be within tolerance. Shortening the time taken for the analysis process would mean the restart case for the reactor could be completed sooner, thus maximising income.

While the FLAW software has a degree of automation to assist the analyst in locating key pressure tube features, it requires input from the operator for these automation procedures to initiate and is limited to processing a single dataset at a time. The proposed algorithms are designed to operate on an larger scale and to incorporate the knowledge of the expert such that reasoned decisions can be proposed with minimal input from a human operator.

3. End-to-End Processing

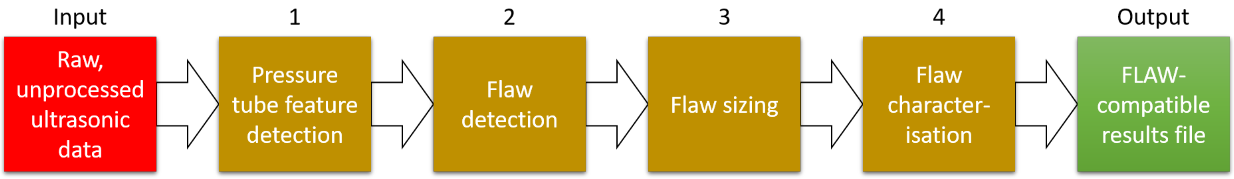

Figure 1: A high-level overview of automated process depicting each of the major modules within the system

The system presented in this paper implements the process depicted in Figure 1. This system will run alongside a human analyst and will provide decision support to the resolution analyst, who compares the results of each of the analysts and makes a final decision on the details of each of the identified flaws.

Our system begins with the same set of unprocessed ultrasonic data that a human operator receives, and will create reports from the data in a format compatible with the existing analysis software, FLAW. Each of these steps will now be examined in detail.

3.1. Pressure Tube Feature Detection

At each end of a pressure tube, there are three mechanical features that must be located prior to flaw detection, keeping to the existing analysis process as much as possible. This procedure is Step 1 in Figure 1's flowchart. These three features are the burnish mark, rolled joints and end of tube. The rolled joints and burnish mark are mechanical features of the pressure tube that exist as a result of the fitting of the pressure tube to the reactor core. Figure 2 shows a diagram of the mechanical fitting where the rolled joints and burnish mark can be observed.

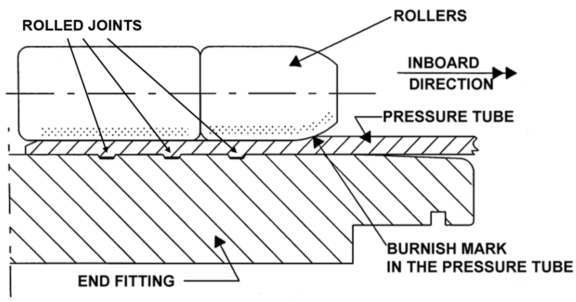

Figure 2: An annotated diagram of the end region of a pressure tube (Clendening, 2002). The three rolled joints are visible beneath the roller and join the pressure tube to the end fitting.

These features are most easily located using a 20MHz normal incidence ultrasonic probe. Analysts locate these features by identifying stripes of equal amplitude along the radial dimension of the pressure tube. This process is replicated in our feature detection algorithm by taking the mean value across each axial position.

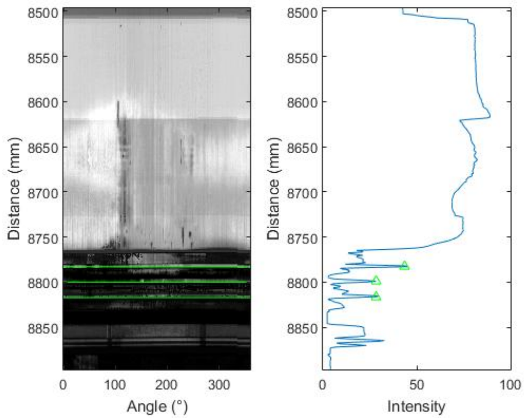

Figure 3: A top-down view of the pressure tube, side-by-side with the mean values for each axial position. The position of the rolled joint fittings are overlaid on each of the plots.

Figure 3 shows the planar (top-down) view of the pressure tube on the left-hand side, while the right shows the mean value of each axial position. The areas of high and low amplitude can be seen clearly on the right of Figure 3, with the group of three peaks representing the three rolled joints and the trough representing the burnish mark location. While the example depicted in Figure 3 shows clear peaks representing the rolled joints, the presence of noise can introduce uncertainty.

When multiple peaks are present in a signal, determined by a peak detection algorithm, the module will sort the peaks into groups of three using the knowledge that these rolls are found approximately 16mm apart and have a significantly higher amplitude response than the rest of the region. It is also known that there will be a maximum of one rolled joint set within a given dataset. A confidence can be assigned to the reported group of rolled joint peaks using the cumulative amplitude of each of the peaks, compared to the other candidates.

The burnish mark is then located, as the lowest amplitude region found approximately 16mm from the rolled joint feature closest to the centre of the pressure tube. The burnish mark is a impression left on the pressure tube as a consequence of the rolled joint machining and is used to determine the location of the critical area where the threshold for reporting flaws changes from -6dB to -14dB (Trelinski, 2008). The threshold is lowered due to DHC being more common in this region due to non-uniform stresses and therefore the need to identify and monitor flaws in this area is increased.

The final feature to be located is the end of tube. The end of tube refers to the physical end of each pressure tube and is located by determining the first position where the ultrasonic response from the inner pressure tube wall can be observed.

3.2. Flaw Detection

Once the pressure tube features are located, Step 2 is to process all of the available GH datasets to locate any potential flaws. Flaws are located by applying a threshold to the shear wave datasets and any region with a received signal above the given threshold is flagged for further investigation. This threshold is -6dB in the body of tube and -14dB within the critical area. The flagged regions are checked to determine if the signals are considered noise. This stage of the process is well-defined in inspection specification documentation.

Noise is also well-defined within the existing operating procedures, however there is scope for subjectivity when following these guidelines due to the interpenetration of the kind of signal that would constitute a response for a given probe. In order to determine whether or not a response is noise, the length, width and B-scan response of the flaw must also be taken into consideration. Flaws that measure more than one sample in length are automatically considered to be legitimate due to the fact that the probe's rotational nature means it has moved out of position and returned while observing a response.

Only a basic noise classification algorithm is applied initially, and subsequent modules in the system are designed such that candidate flaws can later be classified as noise should their characteristics indicate this.

3.3. Flaw Sizing

To find the dimensions of a flaw (Step 3 on the flowchart in Figure 1) the length and width are measured separately from the depth. To find the length and width, the response of the 10MHz normal incidence probe is considered. The sizing method known as the '6dB method' is well defined in the existing operating procedure. To find the 6dB size, the mean amplitude of the normal incidence amplitude response throughout the dataset must be found, as well as the minimum amplitude recorded within the flaw region.

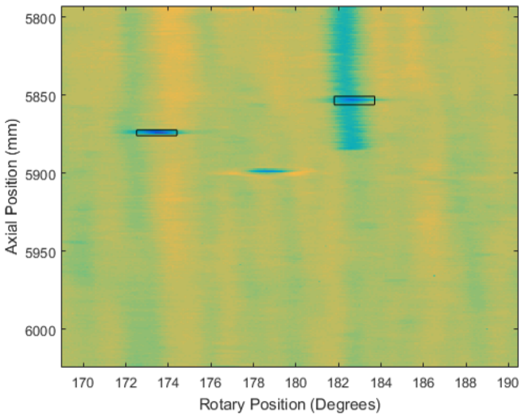

Figure 4: The size and location of two defects, represented by black rectangles, overlaid on the normal incidence ultrasonic dataset.

The edge amplitude is a threshold value, found using Equation 1 where A_E is the edge amplitude, A_f is the minimum signal amplitude of the flaw and A_{AV} is the average back wall signal amplitude in the non-defective region of the pressure tube. The region where the amplitude of the normal incidence amplitude response is lower the edge amplitude is defined as the area of the flaw. The length and width of a flaw are the dimensions of this area in the axial and radial plane respectively.

$$A_E = A_f + \frac{A_{AV} - A_f}{2} \tag{1}$$

Figure 4 shows a subset of a pressure tube containing two defects. These defects were located using the thresholding operation and subsequently sized using the 6dB method. It can be observed that the defect areas highlighted within the figure correspond to areas with high amplitude response in the ultrasonic dataset.

Initially, to determine the depth of a flaw, metadata is loaded from the B-scan datasets to establish which datasets have coverage of the current flaw. Once the correct B-scan dataset is located, the full dataset is loaded into memory and the subset of data covering the defective region extracted.

The depth finding algorithm was developed by combining the procedure outlined in the existing operating procedure document with the domain knowledge of experienced analysts. The algorithm first locates potential signal features indicating a defect by analysing data from a 20MHz normal incidence ultrasonic probe, as well as the signals recorded on two pairs of transducers in a pitch-catch configuration. A confidence measurement for each signal feature is assigned via the iteration through a series of rules. This derives the most likely measurement of depth for each A-scan and the process is repeated for every A-scan within the defective region. Figure 5 shows the depth map for the 20MHz probe for a debris defect within a pressure tube. Depth maps are also constructed for each of the pitch-catch pairs and finally, a second rule base determines which measurements should be reported.

This depth map allows an analyst to observe, not only the maximum depth of a flaw, but how the depth varies with position in relation to the pressure tube. Flaws with a steep depth gradient will have an

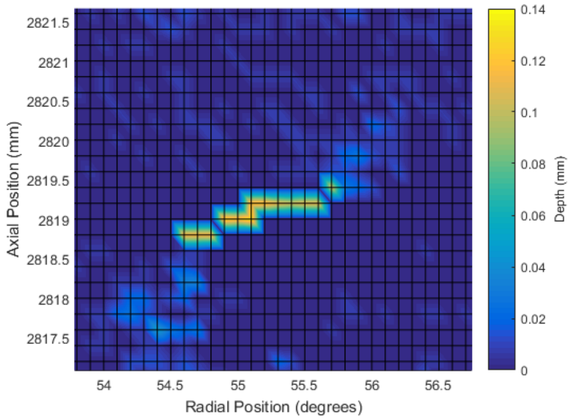

Figure 5: A map generated from the application of an expert system methodology to obtain a depth measurement from ultrasonic data.

increased mechanical stress at the tip leading to hydride build-up and a risk of DHC. Conversely, flaws that are shallow or flat undergo far less stress and therefore pose less of a risk of complication.

3.4. Flaw Classification

With the knowledge of defect location, length, width and depth, classification (Step 5) now takes place using a set of rules elicited from analysts which are used to assign a likely root cause of a flaw based on its location in the channel, its dimensions and shape. Currently, the classification system is in a rudimentary stage and differentiates between fuel bundle bearing pad frets, debris and axial scrapes, although it is anticipated that the classifier will be further developed in order to identify a greater range of potential defects. At this stage it is possible to conclude with a high level of confidence whether or not a flaw is noise, based on the available data.

The suspected root cause of the defect will determine whether or not the flaw is reportable (i.e. of potential concern and should be monitored during future inspections). For example, scratches in the critical area of the pressure tube are not reportable unless they have an amplitude response of greater than -6dB. This differs from debris and bearing pad frets which are reportable if their amplitude response is greater than -14dB. The sizing and classification process is repeated for each flaw found within the body of tube, then the critical area is processed using the same method but applying a -14dB threshold for flaw detection.

4. Results

In order to test the end-to-end system, and the individual sub-modules, a total of 126 datasets from individual pressure tubes were input to the system and the results compared to the results from the existing manual process. Table 1 is concerned with correctly identifying the location of the Rolled Joint, End of Tube and Burnish Mark features (Step 1 of the end-to-end process). The feature was judged to be successfully identified if the location was within 3mm of the position identified by the analyst. Each channel has two sets of these 5 features, one at each end of the tube, so 252 sets of features were classified overall. Overall performance was very good, with two features being identified in the correct location 100% of the time.

Table 2 is concerned with assessing the performance of locating (axial and rotary start position) and sizing (length, width and depth) of all potential defects (Step 3 of the end-to-end process). There are no true negatives included in the results; these would represent the defect free regions of the pressure tube and describing these regions in terms of position, area and (lack of) depth does not make sense. From approximately 1457 defects, there were 1288 true positive identifications, 155 false positive identifications and 169 false negative identifications. Given the lack of true negative information, the balanced F1-score is chosen as an accuracy measure, and with a value of 0.89 it represents a good overall accuracy, though there are areas where improvements could be made. These are discussed later.

Table 2 presents the results for successfully identifying the location and dimensions of the true positive defects. A tolerance value is provided which describes how close to the ground truth result the measured value needed to be in order to be classed successful. These tolerance values were determined in conjunction with the analysts as representative of the variation that would be typically expected between different analysts, based on experience. From the results it can be seen that module which determines the length of the defect had the poorest performance. Discussions with analysts revealed that often the defect would have a tail, for example if caused by a piece of debris being trapped under the bearing pad and being scraped along the bottom of the channel when the fuel bundle was moved. Determining the length of this tail was very subjective relying on the individual analyst's interpretation of the C-scan image to determine the end point, and is an area where further investigation and refinement of the algorithm could be undertaken. Flaws within the critical area are not processed due to known limitations with the 6dB sizing technique for these flaws. Because flaws within critical area have an amplitude approaching that of the surrounding pressure tube, setting the 6dB threshold using the edge amplitude calculation results in a greatly oversized defect and interferes with other measurements made.

Feature	Success Rate %
Rolled Joint 1	94.4
Rolled Joint 2	100
Rolled Joint 3	91.7
Burnish Mark	100
End of Tube 1	97.2

Table 1: Measurements from the body of tube of six pressure tubes, compared to the verified result within a threshold.

Feature	Tolerance	Success Rate (%)
Axial Start Position	3mm	85.8
Rotary Start Position	3 degrees	94.9
Length	3mm	75.3
Width	3 degrees	92.4
Depth	0.1mm	82.3

Table 2: Measurements from the body of tube of six pressure tubes, compared to the verified result within a threshold.

The key result is that an unprocessed channel of CANDU data can be input to this system and all available datasets can be rapidly processed before saving the results in a format compatible with the existing analysis software. Figure 6 shows the FLAW software with an ultrasonic dataset loaded into it. The output file from our system has been input to the software and the defect overlay can be seen, showing the location and size of the defect.

4.1. Analysis of 'Error'

In order to provide an insight into the discrepancies in measurements between the algorithm and human operators, the results from the pressure tube feature detection module are further explored.

Due to the transformation of the helical motion of the probe to a conventional 2D image, a shift appears in the data. This shift can occur at any radial position in a given dataset and causes a measurement error of one sample. A visualisation of this is shown in Figure 7, which shows that a one-sample offset can be introduced depending on which side of the pressure tube a measurement is made from. To put this error in to perspective, a single pixel equates to a physical distance of 0.8mm.

Subjectivity between analysts is a larger source of discrepancy in the feature detection technique when comparing the results of the algorithm to that of the expert. In techniques not tightly defined within the

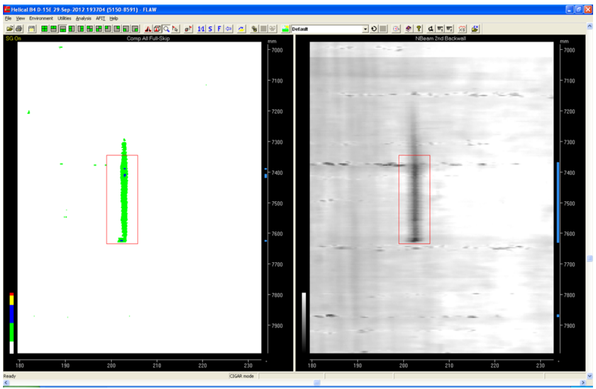

Figure 6: A screenshot of the FLAW software showing a region of the pressure tube containing a defect. The left hand pane shows the response from the shear wave probe, while the right shows the normal incidence beam response. The red bounding box indicates the size as determined by the end-to-end system.

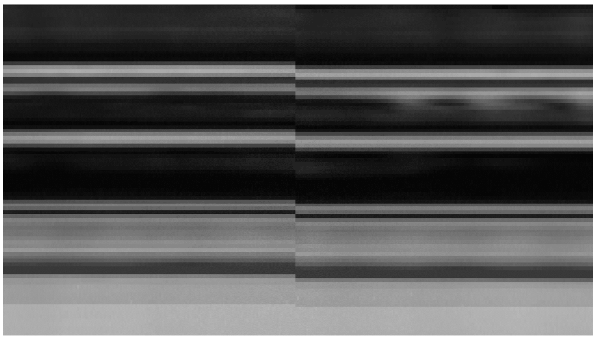

Figure 7: A screen-shot of the data visualised within the FLAW analysis software, with a visible shift in the data around the rolled joint region.

inspection specification documentation, analysts will use their expert judgement to make measurements. For features such as a rolled joint or burnish mark, analysts may choose to measure from any point within the feature. As a feature can potentially occur over 3 axial positions (4 samples, if accounting for the offset error), discrepancies in the measurement can be introduced. An analyst is consistent about the choice of reference point against which the measurements are made (for example always choosing the bottom of the feature) however, the choice of reference may vary between analysts resulting in discrepancies between recorded measurements. Furthermore, analysts have access to full historical records for every given channel and a measurement is made not only using current information, but using records from past outages. It is therefore important to note that these discrepancies are not necessarily errors, as analysts make a subjective judgement with access to more information than is available to the algorithms.

In order to determine if a technique is accurate a threshold must be defined, outside of which measurements are deemed to be inaccurate pending manual review. A 3mm threshold was chosen for tube feature detection which equates to an error of fewer than four samples. Each of the 10 features within a pressure tube were located including the 6 rolled joints (labelled RJ1 through RJ6), the two burnish marks (labelled BM1 and BM2) and the two ends of tube (labelled EoT1 and EoT2). In order to obtain more statistically significant results, tests to measure the discrepancy between measurements took place over 87 channels of data. Figure 8 shows the mean and standard deviation of errors for feature detection within these pressure tube channels, while Figure 9 shows the percentage of times that the measurement was made within the threshold. It is expected that the rolled joints and burnish marks are located with fewer discrepancies compared to the ends of tube. Each rolled joint within the fitting is located 16mm from the next and the burnish mark occurs 16mm from the last rolled joint. With this in mind, the measurements made by analysts are

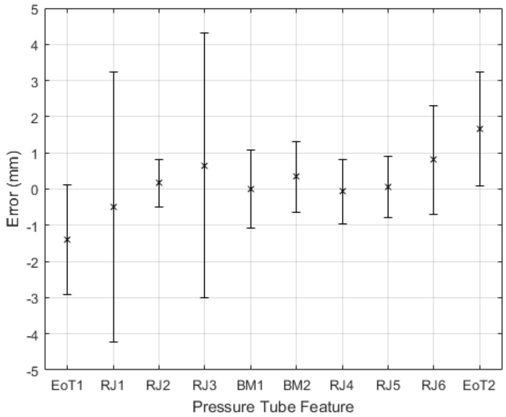

Figure 8: Error bars for each of the features located within the pressure tube. Calculated by comparing the algorithm output to the analysts' verified result over 87 channels of data.

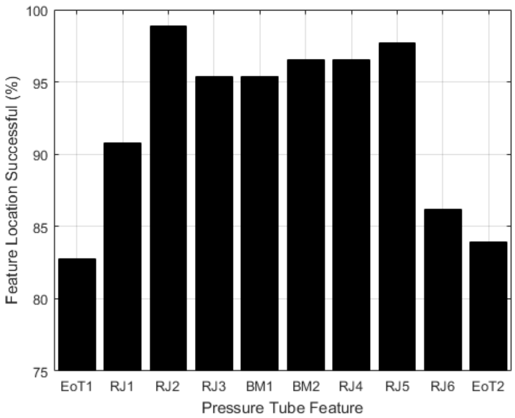

Figure 9: The percentage success rate for locating pressure tube features within a tolerance of 3mm, over 87 channels of data.

less subjective. The end of tube is defined as the point where ultrasonic data starts to be returned from the probe and this measurement can be subject to if an analyst determines that a returned signal is noise.

5. Conclusion

A novel system has been presented for the automation of analysis of CANDU pressure tubes. All of the key features of the tube were accurately located and all expected flaws found. Length, width and depth sizing for these flaws were successful, though it was noted that the accuracy of some of these measurements could be improved through future research. The end-to-end system was run successfully and without the input of a human operator. Furthermore, the results were output in a format allowing for the end-to-end system to provide resolution analysts with decision support.

The approach taken to develop the modules for this end-to-end process was based on the expert system. Using knowledge of the analysis process, rules were developed that follow the decision making procedures performed by the analysts themselves. This process is logged at each step, providing an result which can be audited and the cause of a given output traced. There is also potential to explore a data-driven approach to the processing of the data and to use machine learning to statistically analyse CIGAR datasets. The drawback of this approach is the need for large volumes of labelled data which does not currently exist.

Future work will involve a wider test of this system. The proposed end-to-end system has been tested on six channels of CIGAR data, however historical data is also available for analysis. An automated testing procedure has been planned that will process all of the available datasets and their associated verified results to compare the verified result to the result of the new system. This testing procedure will build confidence in the new system, allowing an easier case to be made for deploying this system during planned outages. Deployment of the algorithms developed within the overarching system will to be made available for the analysis process via direct integration with the current analysis software.

Reflecting on Bertovic's (2016) paper in which it was stated that automation poses a potential risk of error within non-destructive testing, the error is largely mitigated with the current operational procedure of two analysis streams and the fact that while the procedure that an analyst follows is now automated, the procedure itself is remains auditable. Furthermore, the industry application of this research is decision support. An analyst will use the developed modules as tools that are able to provide reasoned case for measurements, providing a robust solution for defect sizing and classification within pressure tubes while also able to review the flaws most at risk of delayed hydride cracking.

Acknowledgments

The authors wish to thank Bruce Power for funding this work. The opinions expressed herein are those of the authors alone.

References

Aho-Mantila, I., Cronvall, O., Ehrnstén, U., Keinänen, H., Rintamaa, R., Saarenheimo, A., Simola, K., Vesikari, E., 2012. Lifetime prediction techniques for nuclear power plant systems. In: Feron, D. (Ed.), Nuclear Corrosion Science and Engineering. Woodhead Publishing Series in Energy. Woodhead Publishing, pp. 449–470.

Anderson, J. W. D., 1990. The use of expert systems for operations support in the CANDU nuclear power supply system. In: Use of expert systems in nuclear safety. Vienna, pp. 237–244.

Bertovic, M., Feb. 2016. A human factors perspective on the use of automated aids in the evaluation of NDT data. In: AIP Conference Proceedings. Vol. 1706. AIP Publishing, Minneapolis.

Cheadle, B., Coleman, C., Ambler, J., Jan. 1987. Prevention of Delayed Hydride Cracking in Zirconium Alloys. In: Adamson, R., Van Swam, L. (Eds.), Zirconium in the Nuclear Industry. ASTM International, Philadelphia, pp. 224–241.

Cho, D., Xu, S. X., Scarth, D. A., Shek, G., 2016. Development of evaluation procedures for crack initiation from in-service flaws in candu zr-2.5nb pressure tubes with hydrogen. In: Proceedings of the ASME 2016 Pressure Vessels and Piping Conference PVP2016.

Ciocan, R., Ciobanu, D., 1998. Neural Network Classification of Flaws Detected During Ultrasonic Inspection of Candu Pressure Tubes and Brazed Joints in Candu Fuel. In: ECNDT '98. Copenhagen.

Clendening, W. R., Dec. 2002. Pressure Tube to End Fotting Rolled Joints.

Coble, J., Ramuhalli, P., Bond, L., Hines, J. W., Ipadhyaya, B., Jan. 2015. A Review of Prognostics and Health Management Applications in Nuclear Power Plants. International Journal of Prognostics and Health Management 6, 16–31.

D'Orazio, T., Leo, M., Distante, A., Guaragnella, C., Pianese, V., Cavaccini, G., Mar. 2008. Automatic ultrasonic inspection for internal defect detection in composite materials. NDT & E International 41 (2), 145–154.

Garcia, P. A. A., Schirru, R., Frutuoso E^Melo, P. F., Jan. 2005. A fuzzy data envelopment analysis approach for FMEA. Progress in Nuclear Energy 46 (3), 359–373.

Kim, Y.-J., Kwak, S.-L., Lee, J.-S., Park, Y.-W., Jul. 2003. Integrity evaluation system of CANDU reactor pressure tube. KSME International Journal 17 (7), 947–957.

Kwak, S.-L., Lee, J.-S., Kim, Y.-J., Park, Y.-W., Aug. 2005. Sensitivity analysis of input parameters for pressure tube integrity evaluation. Nuclear Engineering and Design 235 (17–19), 1909–1918.

Meyer, R. M., Bond, L. J., Ramuhalli, P., 2014. Online Condition Monitoring to Enable Extended Operation of Nuclear Power Plants. In: Yoshikawa, H., Zhang, Z. (Eds.), Progress of Nuclear Safety for Symbiosis and Sustainability. Springer Japan, pp. 161–176.

Murray, P., West, G., Marshall, S., Mcarthur, S., Apr. 2016. Automated in-core image generation from video to aid visual inspection of nuclear power plant cores. Nuclear Engineering and Design 300, 57–66.

Obrutsky, L., Renaud, J., Lakhan, R., Jan. 2007. Steam generator inspections: Faster, cheaper and better, are we there yet? In: Proc. IV Conferencia Panamericana de Ensayos no Destructivos. Buenos Aires, pp. 1–17.

Oh, Y.-J., Chang, Y.-S., 2014. Integrated probabilistic assessment for dhc initiation, growth and leak-before-break of phwr pressure tubes. Nuclear Engineering and Design 275, 30–43.

Puls, M. P., Jul. 1997. Assessment of aging of Zr2.5nb pressure tubes in CANDU reactors. Nuclear Engineering and Design 171, 137–148.

Puls, M. P., Wilkins, B. J. S., Rigby, G. L., Mistry, J. K., Sedran, P. J., Oct. 1998. A probabilistic method for leak-before-break analysis of CANDU reactor pressure tubes. Nuclear Engineering and Design 185, 241–248.

Radu, V., Roth, M., Dec. 2012. Probabilistic fracture mechanics applied for DHC assessment in the cool-down transients for CANDU pressure tubes. Nuclear Engineering and Design 253, 211–218.

Sahoo, A. K., Pandey, M. D., Mar. 2010. A risk-informed approach to assessment of DHC initiation in pressure tubes of CANDU reactors. Nuclear Engineering and Design 240, 630–638.

Sambath, S., Nagaraj, P., Selvakumar, N., Dec. 2010. Automatic Defect Classification in Ultrasonic NDT Using Artificial Intelligence. Journal of Nondestructive Evaluation 30, 20–28.

Trelinski, M., 2008. Inspection of CANDU Reactor Pressure Tubes Using Ultrasonics. In: 17th World Conference on Nondestructive Testing. Shanghai, China.

Uhrig, R. E., 1989. Use of expert systems in nuclear power plants. Tech. rep., Oak Ridge National Lab.

Venkatasubramanian, V., Rengaswamy, R., Kavuri, S. N., Yin, K., Mar. 2003. A review of process fault detection and diagnosis: Part III: Process history based methods. Computers & Chemical Engineering 27, 327–346.

West, G., McArthur, S., Towle, D., Jun. 2012. Industrial implementation of intelligent system techniques for nuclear power plant condition monitoring. Expert Systems with Applications 39, 7432–7440.

Yang, J., Yang, M., Wang, W., Li, F., Aug. 2016. Online application of a risk management system for risk assessment and monitoring at NPPs. Nuclear Engineering and Design 305, 200–212.