IMAGE PROCESSING AND ARTIFICIAL INTELLIGENCE FOR DETECTION AND INTERPRETATION OF ULTRASONIC TEST SIGNALS*

Keith S. Pickens, John C. Lusth, Pamela K. Fink,
Karol K. Palmer, Ernest A. Franke

Southwest Research Institute
San Antonio, Texas

INTRODUCTION

Detection of flaws is an important industrial concern. For example, aircraft and nuclear-power reactor owners and regulatory authorities need effective means of detecting flaws that could pose a threat to public safety. Operators of costly equipment require information on service-induced flaws to be able to make run-or-retire decisions. As the cost of parts and concerns for public safety increase, the importance of flaw detection and size estimation has likewise escalated.

Ultrasonic nondestructive evaluation (NDE) is one of the primary tools in the inspection for flaws (such as voids, nonmetallic inclusions, and cracks). Because ultrasonic testing (UT) uses acoustic waves for detection, it can inspect the interior of a thick material and reach inaccessible surfaces. Detection of a flaw, however, is only one step in the process. Flaw sizing is equally important and becoming more so with the increasing concern about lifetime prediction.

Current UT technology does not provide adequate sizing information in all cases. At Southwest Research Institute (SwRI), this problem was addressed by the development of the patented SLIC (shear and longitudinal waves to inspect for cracks) multibeam technology, which greatly extends ultrasonic inspection accuracy [1-4].

The importance with which the development of the SLIC technology must be viewed is evidenced by the recently announced result of an international round-robin test. The round robin was conducted to measure the efficiency of organizations and NDE methods for correctly detecting and characterizing flaws contained in test blocks representing nuclear-power reactor pressure-vessel components. These test blocks were examined by more than thirty of the world's foremost organizations who individually applied their own best technology and efforts. The results speak for themselves. When practiced by an SwRI expert, the SLIC technology outperformed every other technique and international organization.

*Funded by the Southwest Research Institute Internal Research Program

The limiting factor in the wide application of SLIC technology to industrial problems is, ironically, the large amount of information it generates. An analyst requires a high level of expertise to understand and interpret a SLIC module signal. The lack of skilled analysts has limited the application of SLIC technology to the many problems for which it can provide a solution.

OVERVIEW

The technical approach to the recognition and interpretation of the SLIC multibeam-multipulse signals involves three major subject elements--UT data acquisition, image analysis, and analysis of image features by an expert system. These elements separately represent known and established technology within their limited domain. The strength of this system lies in the synthesis of these three different fields to form an interdisciplinary solution to a difficult problem.

One of the most productive areas of work in artificial intelligence (AI) has involved the production of programs with expert-level performance. These systems operate in their limited domains by incorporating the knowledge of a human expert, usually in the form of decision rules. Rules embody the relationships that a human expert uses in a formal statement of his knowledge. Thus, an expert system provides access to this specific expertise in a consistent and reproducible manner and on a much larger scale.

The expert system is provided with a description of the key features extracted from ultrasonic signals received during an examination. To facilitate extraction of key features, the examination data are organized as an image. The UT signals do not form an image in the classic sense, but it is possible to treat them as a two-dimensional image by using time as one axis. Forming an "image" allows the powerful and well-developed tools of image processing to be applied to extract a description of the signals for analysis by the expert system.

The acquisition of UT signals to form an "image" (see Figure 1) is the basis upon which the image processing and expert system rest. Computer-based data-acquisition systems allow the necessary large-scale signal acquisition. The data acquisition system technology used with the expert system is described in another paper in these proceedings [5].

Image Enhancement

Image processing [6-7], used to extract a description of the signals, reduces the amount of data to be analyzed by the expert system by two orders of magnitude. This reduction greatly increases overall system performance. The image processing itself involves three major steps. First, noise from the image formed by the SLIC signals is filtered out which enhances clarity of the image. Second, the characteristics that define the features of interest are accentuated. Third, these features are analyzed to determine the characteristics that describe the SLIC signals. The characteristics thus determined are then passed to the expert system for analysis.

Noise Reduction

The first step in processing the image is noise reduction. This is accomplished in the system by convolution with two digital filters. The digital filters take advantage of the fact that design of the SLIC module and the accompanying data acquisition process result in signals that form

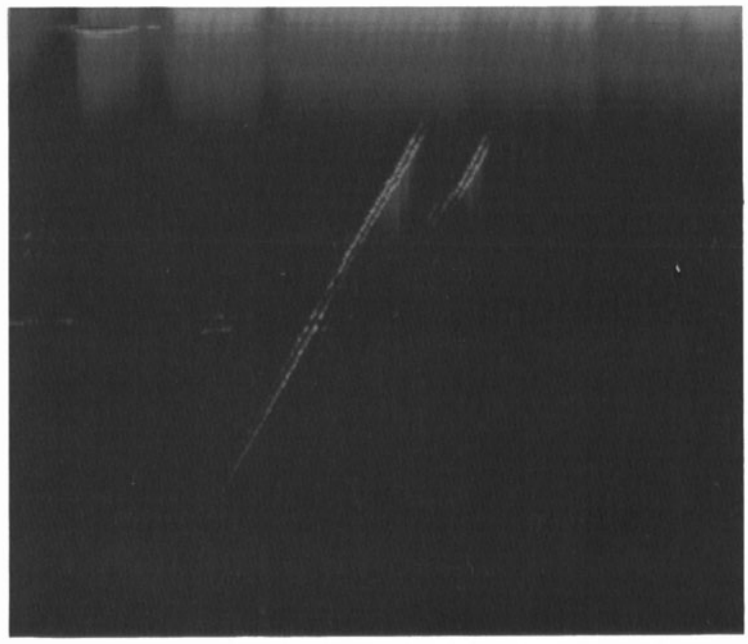

Fig. 1. Image formed by a SLIC-50 transducer before processing

lines of known slope. One filter emphasizes lines with positive slopes; the other filter emphasizes lines with negative slopes; and both filters reduce the background noise. Since different filters are used for the positive and negative slope lines, the image-processing module tries both and selects the one that produces the "brightest" image.

Conversion to a Binary Image

To identify the lines, each pixel in the image is placed in one of two groups. One group contains pixels within the line; the other contains pixels not in a line. This binary image can then be used to build a set of parameters for the lines. To convert the enhanced image to binary, a threshold must be selected. A histogram of pixel intensities shows a large number of background pixels with low intensity and a smaller number of brighter pixels that form the enhanced lines (see Figure 2).

A threshold value must be selected that will suppress the background while leaving the line pixels. If the threshold selected is too low, the binary image will contain extraneous (noise) pixels that can distort the features, introduce extra features, and increase processing time. If too high a threshold is used, portions of lines or some entire lines may be suppressed. The program determines the threshold based on an empirically determined number of bright pixels. By basing the threshold on the number of pixels rather than a given level, good results can be obtained on a wide range of images.

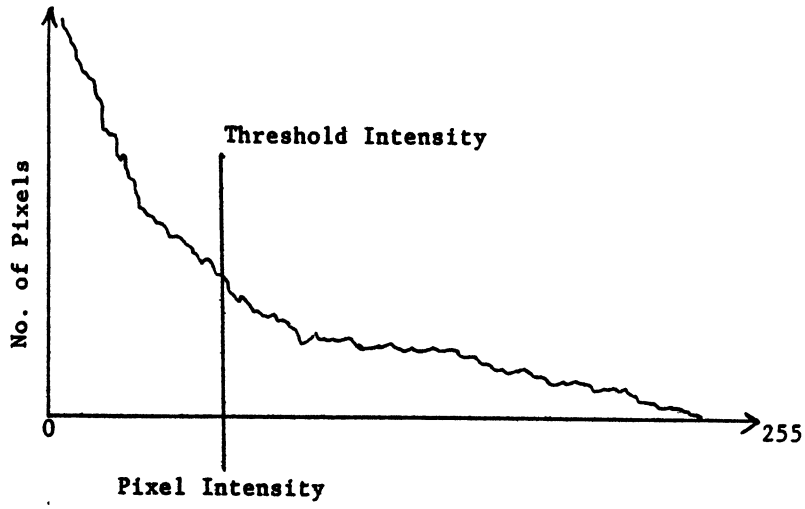

Fig. 2. Typical histogram of pixel intensities after digital filtering

Blob Labeling

After a binary image is obtained, the next step is to group connected pixels into discrete groups, referred to as "blobs." This is done by scanning the binary image pixel by pixel and constructing a new image in which each pixel is associated with a unique blob.

When a binary image point is found, the adjacent points that have already been scanned are examined to determine if any of them belong to a blob. If so, the pixel under consideration is labeled with that same blob identifier. The scan pattern is from top to bottom and left to right for images with positive slope lines on the monitor, and from right to left for negative slope images. This scan pattern was selected to reduce the possibility of fragmenting lines into multiple blobs.

During blob labeling, the number of pixels in each blob is counted. When labeling is completed, small blobs with fewer than four pixels are discarded.

Feature Extraction

The last section of the image processing module scans the labeled blob image and accumulates data to compute parameters identifying lines.

After these data are accumulated, final line parameters are computed by least squares regression. A file is then written to disk containing (for each line):

- length of line
- slope of line
- x, y coordinates of the line start point and its intensity
- x, y coordinates of the line end point and its intensity
- x, y coordinates of the maximum intensity.

Figure 3 shows the computed lines overlaid on a SLIC image.

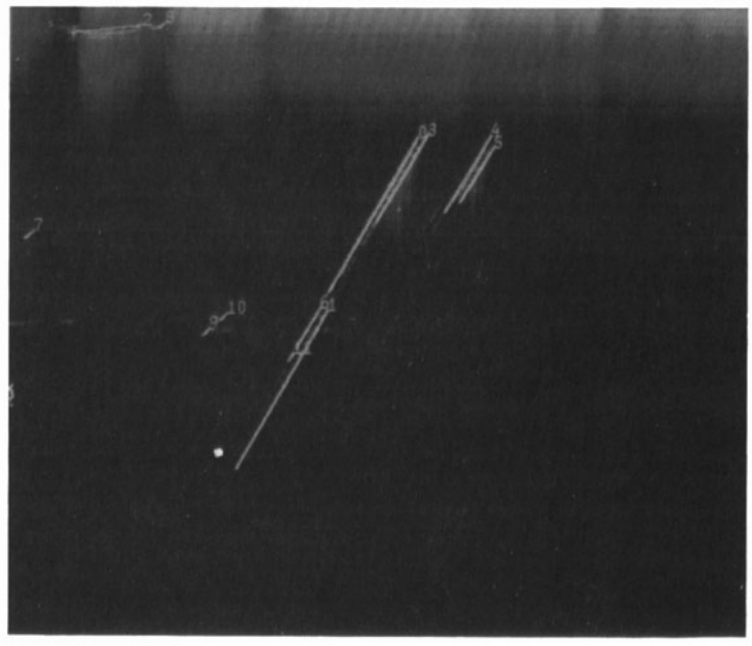

Fig. 3. Lines identified by the image-processing module. Fig. 1 shows the same image before processing.

Artificial Intelligence Module

The artificial intelligence [8-10] module interprets the image parameters produced by the image-processing module. From the image parameters received, this module must distinguish between data which represents actual artifacts and data which is the result of interference or noise. The analysis is implemented in three steps: line merging, artifact identification, and crack depth computation.

Prior to implementation of these modules, the merging process was mainly considered a trivial task requiring no expertise. The expertise, however, that was believed to lie in the merging process (which seems simple, if not intuitive, to a human) was found to be more complex than expected. Findings such as this are not unusual when attempting to automate an intuitive process.

The line merging program detects and corrects two types of problems existing in the data produced by the image-processing module. First, the method used by the image-processing module tends to break line segments into two or even more segments. These segments must be reconnected in order to achieve a proper representation of the data resulting from the SLIC test. Second, the input data sometimes consisted of under-filtered video signals; a single signal appears to be multiple, closely spaced signals. These signals must be blended into the meaningful line segments. The merging program uses two sets of rules to determine if either of these situations exists.

Line segments that actually belong to the same original line tend to have certain attributes in common. For example, the endpoint of one line segment will tend to be close to the endpoint of the adjacent segment, the orientation of the two-line segments will tend to be the same (i.e., the line segments will be close to parallel), and intensity of the points along the two-line segments should produce a smooth change from low intensity to high intensity and back to low intensity.

Thus, three rules are used to determine whether line segments have been broken apart. First, the endpoints of two-line segments must lie within a certain distance (d1) of each other. Second, both endpoints of the shorter line segment must lie within a certain distance (d2) of an extension of the longer line segment. Both distances d1 and d2 are currently constant values which were empirically determined. In future implementations, these distances may be adaptively determined.

The third rule used to determine whether line segments have been broken apart involves the intensities of the various points on the line segments. When graphed, intensities along a typical line segment would form a curve with a single peak. Therefore, the intensities along the line segment that would result from the merge must also form a curve with a single peak.

Under-filtered video signals appear in the image as line segments which are parallel to and lie very close to another line segment. The merging routine uses two rules for identification of echoes. First, the two-line segments must overlap. Second, both endpoints of the shorter line segment must lie within a certain distance of an extension of the longer line segment.

The merge program tests each line segment against every shorter line segment for broken segments and echoes. Segments are tested beginning with the longest and working in descending order by length. This process is reiterated until no more broken segments or echoes can be identified.

For both broken segments and echoes, the line segments are merged by extending the longer line segment to the projection of the far endpoint of the shorter segment. The shorter segment is then deleted. The points of greatest intensity of the two-line segments are compared, and the one with the larger value is preserved.

From the line segments resulting from the merging program, a third program finds pairs of line segments that meet the criteria for a flaw. Currently, the criteria for cracks detected by the SLIC-50 method with their origin at the cladding-to-base metal interface are implemented. The criteria for other SLIC-50 cracks and for SLIC-40 cracks will be implemented in the future.

There are two rules currently implemented for determining whether a pair of line segments have the characteristics of a crack detected by the SLIC-50 method. The first rule dictates that the slopes of the two segments be nearly equal. The second rule involves a check against the known geometry of cracks within the clad sample. The depth of a crack can be measured in two ways. The first method is simply to measure the distance between the longer line segment and the maximum intensity point on the shorter line segment. The second method relates the location of the cladding-to-base metal interface to the location of the shorter line segment in the image. This assumes the origin of the crack is at the cladding-to-base metal interface. These two rules are applied to the line segments in descending order starting with the longest segment.

In an early implementation, the first pair of segments found that had the characteristics of a crack was determined to be the actual crack. This method is adequate for images in which there is very little noise and the crack is the most significant feature present. However, in more realistic scenarios, varying amounts of noise as well as spurious artifacts, such as a weld prep, may appear in the image. Therefore, all pairs having the characteristics of a crack are found. The "best" crack is then chosen from among these based on a set of rules that distinguish a true crack from other artifacts. Currently, two rules have been implemented for this set.

The first of these rules is based on the observation that for SLIC-50 images the "top" line segment of the pair of line segments representing the crack is consistently longer and brighter than any line segment representing noise or interference. Therefore, the longest and brightest line segment is initially assumed to be the "top" line segment of a crack. If no matching "bottom" line segment exits for this line, it is disqualified as the "top" line segment of a crack.

The second rule of this set is the last rule to be executed. It selects, from the set of all line-segment pairs that pass the criteria of all the previous rules, the pair which represents the largest crack. This rule's success lies in the fact that the amount of noise present in an image is empirically found to decreases with depth. Therefore, line segments found deeper in the image are more likely to represent a crack than to be the result of noise or interference.

Finally, if a crack is present, a third program determines the depth of the crack. The depth is calculated by measuring the distance between the first line segment and the maximum intensity point of the second segment. This distance is then converted from pixels to millimeters by a constant scaling factor determined through a calibration.

To date, this system has been tested against two sets of images. The first set has of five scenarios. Each scenario consists of a flaw and the corresponding image or images containing relatively clear images of both line segments constituting the crack. These scenarios also have some noise. The system successfully identifies and measures the crack in all five of these scenarios.

The second set also consists of five scenarios. These have much fainter images of the line segments constituting the cracks. The "bottom" lines of the pairs are especially faint, in some cases less bright than the surrounding noise. They also contain more noise than the first set. The system successfully identifies and measures the cracks in three of these five scenarios.

For the other scenarios where the system does not correctly identify the crack, it chooses a "bottom" line segment that is actually noise. In both cases the reported crack depth was greater than that of the actual crack. The first of these scenarios represents the case where a small crack is present. The correct "bottom" line segment was never eliminated from the list of segments having the characteristics of a crack. The noise is chosen over the correct segment only because it lies deeper in the image. In this case, it is believed that a rule that addresses the intensities of the second segments along with their position in the image may resolve this problem. In the second of these two scenarios, the "bottom" segment of the line pair lies among a large cluster of line segments caused by the signals received from the edge of the sample. It is suspected that in this case it would be very difficult for even the human expert to identify the crack. Further consultation with the expert is planned for this scenario.

CONCLUSION

The combination of image processing techniques with an expert system has thus far proven to be extremely successful. "Real world" signals with both noise and artifacts have been handled. At this writing, the expert system uses only a few rules, yet it can deal with the entire class of SLIC-50 signals. Extension to other members of the SLIC family and to conventional UT can be accomplished by the addition of new rules to the expert system.

REFERENCES

1. "Ultrasonic Satellite-Pulse (Observation) Technique for Characterizing Defects of Arbitrary Shape," U.S. Patent No. 4299128, November 1981.
2. "Ultrasonic Multibeam Inspection Technique for Detecting Cracks in Bimetallic or Coarse-Grained Materials," U.S. Patent No. 4435984, March 1984.
3. "Multibeam Satellite-Pulse Observation Technique for Characterizing Cracks in Bimetallic or Coarse-Grained Components," U.S. Patent Application Serial No. 588898.
4. "Detection of Surface Cracks in Bimetallic Structures - A Reliability Evaluation of Six Ultrasonic Techniques," published in the Conference Proceedings of DARPA/AF Review of Progress in Quantitative NDE, University of Colorado, Boulder, Colorado, August 2-7, 1981.
5. H. L. Grothues, D. R. Hamlin, R. H. Peterson, and Keith S. Pickens, "A Highly Interactive System for Processing Large Volumes of Ultrasonic Testing Data," to be published.
6. Ramakant Nevatia, Machine Perception, New Jersey, Prentice-Hall, Inc., 1982.
7. Dana H. Ballard and Christopher M. Brown, Computer Vision, New Jersey, Prentice-Hall, Inc., 1982.
8. Avron Barr and Edward Feigenbaum (eds), The Handbook of Artificial Intelligence, Vols. 1-3, Los Altos, California, William Kaufmann, Inc., 1981.
9. Nils Nilsson, Principles of Artificial Intelligence, Palo Alto, California, Tioga Publishing Co., 1980.
10. Patrick Henry Loinston, Artificial Intelligence (2nd ed.), Reading, Massachusetts, Addision-Wesley, 1984.