In-process detection of weld defects using laser-based ultrasound

Stephen W. Kercel[*a], Roger A. Kisner[a], Marvin B. Klein[b], G. David Bacher[b], Bruno Pouet[b]

[a]Oak Ridge National Laboratory, P.O. Box 2008, Oak Ridge, TN 37831-6011

[b]Lasson Technologies, 6059 Bristol Parkway, Culver City, CA 90230

ABSTRACT

Laser-based ultrasonic (LBU) measurement shows great promise for on-line monitoring of weld quality in tailor-welded blanks. Tailor-welded blanks are steel blanks made from plates of differing thickness and/or properties butt-welded together; they are used in automobile manufacturing to produce body, frame, and closure panels. LBU uses a pulsed laser to generate the ultrasound and a continuous wave (CW) laser interferometer to detect the ultrasound at the point of interrogation to perform ultrasonic inspection. LBU enables in-process measurements since there is no sensor contact or near-contact with the workpiece.

The authors are using laser-generated plate (Lamb) waves to propagate from one plate into the weld nugget as a means of detecting defects. This paper reports the results of the investigation of a number of inspection architectures based on processing of signals from selected plate waves, which are either reflected from or transmitted through the weld zone. Bayesian parameter estimation and wavelet analysis (both continuous and discrete) have shown that the LBU time-series signal is readily separable into components that provide distinguishing features which describe weld quality. The authors anticipate that, in an on-line industrial application, these measurements can be implemented just downstream from the weld cell. Then the weld quality data can be fed back to control critical weld parameters or alert the operator of a problem requiring maintenance. Internal weld defects and deviations from the desired surface profile can then be corrected before defective parts are produced.

Keywords: laser-based ultrasonic, weld inspection, on-line inspection, Bayesian, wavelet

1. INTRODUCTION

The automotive industry is continually re-engineering its manufacturing processes to effect cost savings, enhance quality, reduce weight, and improve safety. One new process with all these attributes is associated with the fabrication of sheet metal panels for auto body manufacture. Different areas of the body have varying requirements for strength and corrosion resistance. Older manufacturing techniques either used single panels of sheet metal with stiffeners and protective coatings, or multiple panels with the proper characteristics which were attached separately to the body frame. In a new design-for-manufacturing approach, manufacturers are now producing large sheet metal panels or blanks made from smaller, individually engineered panels that are butt-welded together using a CO_2 or Nd:YAG laser welding process. With this approach, panels with differing thickness, metallurgy, or surface treatment can be joined to provide the desired attribute only in positions where it is required. Compared to conventional methods, the advantages of these tailor-welded sheet metal blanks are:

- less tooling and better integration of parts,
- forming with a single set of dies,
- reduction in manufacturing steps and in part count,
- superior dimensional control,
- reduction in overall weight,
- improved crash energy management,
- corrosion resistance, higher strength only where required,
- lower net cost, and
- better fit during assembly, resulting in less body noise.

[*] Correspondence: Email: kzo@ornl.gov; Telephone: 423-574-5278; Fax: 423-574-6663.

In the United States, the production of tailor-welded blanks for 1997 is estimated to be 8 million blanks worth approximately $25 million and made from 110,000 tons of sheet steel. The corresponding measure of welds produced is on the order of 10 miles per day. This production rate is expected to increase by a factor of 10 by the turn of the century.

In the blank welding process, the adjacent edges of the two panels to be welded are held at a fixed gap or in contact and the panels are then advanced under the stationary laser fixture. The most critical weld condition that must be controlled is the integrity and shape of the weld nugget formed between adjacent base metal surfaces. The desired transition from one surface to the other is smooth and flat, as shown in Figure 1. Either a concave or a convex surface is a sign of a lower strength weld. Typically, a concave surface can arise from poor fit-up or seam tracking, incomplete weld penetration or excessive weld penetration. Additionally, any crack or lack of fusion is a weld defect.

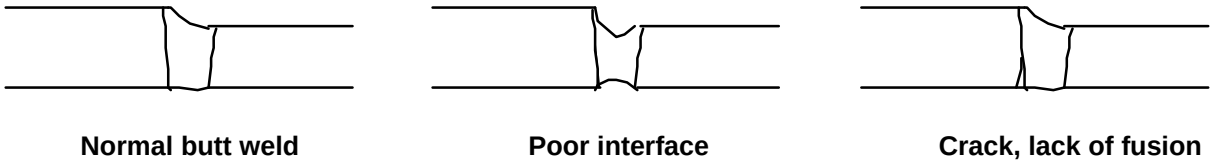

Figure 1. Schematic drawing of normal laser butt weld and poor weld with concave surfaces.

On-line monitoring of weld quality is a crucial unsolved problem in the manufacture of tailor-welded blanks. The goal is to be able to detect the full range of defects that may be present. At the current time, process monitors (e.g., weld speed, laser current) have been implemented, and control of these parameters to maintain the desired weld quality has been attempted. However, these open-loop controls are not sufficiently effective. As a further step in monitoring the weld process, a number of acoustic, optical, and ultrasonic sensors have been studied. Acoustic emission from the weld process (measured with an external microphone) was found to be noisy and not a full measure of weld quality. Optical sensors based on measurement of the spectral emission from the weld plasma or fusion zone have been investigated. Some correlation with weld quality has been observed, but the discrimination is not strong. Finally, ultrasonic inspection using electromechanical acoustic transducers (EMAT) has been studied, and has been shown to be unable to detect all the flaws.

It is clear the need remains to implement a sensor that could directly measure the nugget integrity and surface profile, and thus determine the strength of the weld. One approach is to measure the deflection angle of a laser beam reflected from the nugget surface. Such an approach has been attempted, but poor surface quality has prevented reliable measurement of the surface profile. In addition, both sides of the blank would have to be interrogated. In the absence of any reliable automated sensor, welded blanks in the factory are now inspected by visual means on a statistical basis. Destructive sectioning with micrographic examination is sometimes used to augment the visual inspection. If defective welds are found then the output from the most or the entire shift must be inspected and the welder must be recalibrated. Clearly, this approach leads to added cost in energy, materials, and labor. As an alternative to this *reactive* welder maintenance approach, it would be desirable to develop a sensor that would monitor the desired weld property directly and allow *proactive* intervention to correct the weld process in real-time or to alert the operator that maintenance is required.

2. LASER ULTRASONICS

A technology that shows great promise for on-line measurement of weld quality in tailor-welded blanks is laser-based ultrasound (LBU). LBU is a technique for performing ultrasonic inspection using a pulsed laser to generate the ultrasound and a separate continuous wave (CW) laser interferometer to detect the ultrasound at the point of interrogation.[1] There are several significant advantages for the application of LBU. No sensor contact (or near-contact) with the workpiece is required, thereby allowing in-process measurements in the harsh welder environment (i.e., high temperatures, turbulent atmospheric environment, vibrating parts with rapid lateral motion). Both free-space and fiber-based delivery of optical energy can be implemented.

The authors are investigating the use of laser-generated plate (Lamb) waves propagating from one plate into the weld nugget as a means of determining the surface profile. The investigation is considering a number of inspection architectures based on processing signals from selected plate waves, which are either reflected from, or transmitted through the weld zone. One goal is to identify the simplest generation technique, wave mode, wave propagation geometry, feature extraction process, and

pattern recognition algorithm that will provide effective identification of defective welds. It is anticipated that these measurements can be implemented just downstream from the weld cell (see Figure 2), so that weld quality data can be fed back to control critical weld parameters (or alert the operator of a problem requiring maintenance). Any deviation or defects from the desired surface profile can then be corrected before defective parts are produced.

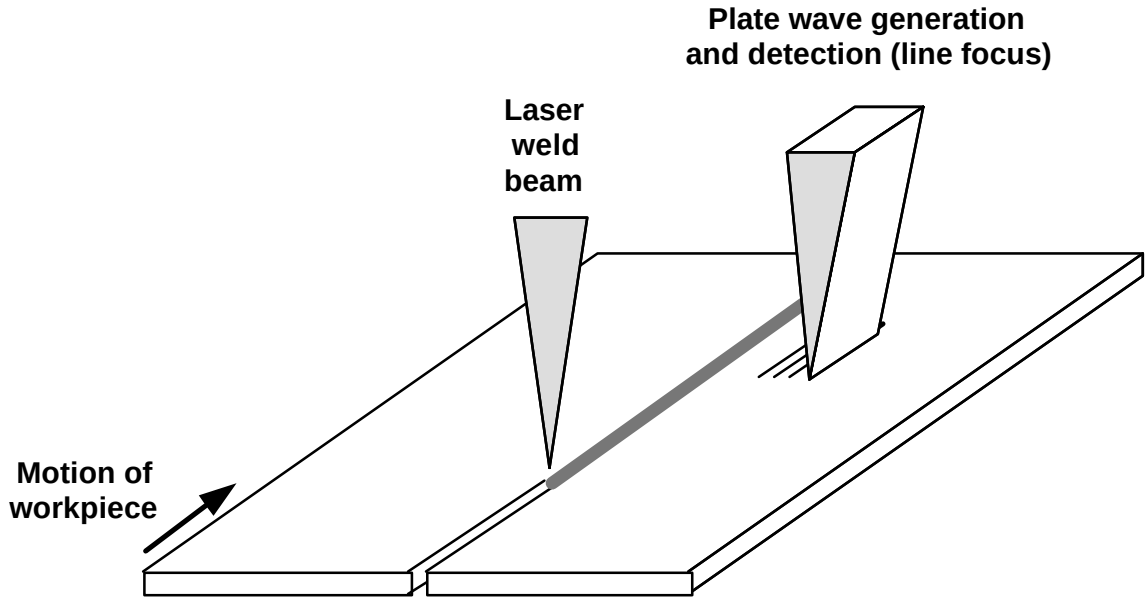

Figure 2. Schematic diagram of weld inspection geometry using reflected plate waves.
The gap between the sheet metal panels is exaggerated.

Recent advances in optical receiver technology can overcome some of the disadvantages of conventional laser-ultrasonic methods.[1] Adaptive interferometric receivers under development at Lasson allow the processing of speckled signal beams that result from the interrogation of machined surfaces. In conventional interferometric receivers, speckle is a source of phase error that can greatly reduce the receiver signal-to-noise ratio. In conventional LBU technology, multiple speckles can be avoided by focusing tightly to a single speckle, but this approach has the risk of signal "dropouts" due to the presence of "dark speckles." The other major advantage of the new adaptive receivers is that they are insensitive to noise at lower frequencies caused by workpiece vibration and by welder-induced turbulence in the path of the receiver probe laser.

One challenge in implementing LBU for process control is the relatively low signal-to-noise ratio. This problem could be overcome by using more sophisticated signal processing for LBU than has been used in the past.[2] The signals produced by the LBU receiver are nonstationary transients in the time-domain, and real-time wavelet analysis should lead to the most efficient method of feature extraction. Furthermore, embedded wavelet processes are implementible in real-time hardware and are effective for many pattern recognition problems for transient signals in severe noise.[3-5]

3. SIGNAL OVERLAP

One of the major problems in the analysis of signals from ultrasonic sensors is the fact that the signal of interest is often overlapped with many other signals.[6] The feature of interest (such as a localized weld defect) produces a signal that contains desired information (such as the distance of the weld defect from the point of ultrasonic excitation). However, other features (such as the edges of the workpiece) also produce signals that overlap the signal of interest, but contain no relevant information. In addition, the signals are inherently nonstationary, and equally inseparable in the time domain and the Fourier frequency domain. Finally, the signal of interest may be considerably weaker than the obscuring signals.

The key to the signal-processing problem is to separate the desired signal from the undesired signals while not destroying the desired information with the process. To appreciate the idea, consider a sensor signal that is the sum of four overlapping Gabor functions plus a low-level Gaussian-distributed random function. Also consider that the information of interest is contained in one of the Gabor functions. [Note: A Gabor function is a Gaussian-windowed sinusoid. Its waveform is

completely characterized by four parameters—the frequency and phase of the sinusoid and the peak location and width of the window.]

In this example, suppose each of the Gabor functions is produced by a different physical feature and that we can tell something about the feature (for example, its location in space) by examining the parameters of the Gabor function associated with that specific feature. The problem is that we must disentangle the underlying Gabor functions given the overlapped signal. As shown in Figure 3, the signal is the sum of a low level of Gaussian noise and four Gabor functions with normalized frequencies of 0.025, 0.05, 0.025, and 0.05, and window peaks at time delays of 600, 1400, 1400, and 600 respectively. The components are not conveniently separable in either the time or frequency domains.

Figure 3. Signal overlapped in time and frequency.

Suppose we have prior knowledge from our understanding of the underlying physical process that the signal should contain one or more Gabor functions, but we have no prior knowledge of the parameters such as frequency or the time of the window peak. By the methods of Bayesian parameter estimation, we can use the Gabor function as a model, and guess a frequency and a time.[7] We can project the signal shown in Figure 3 onto the model, and compute the log likelihood that the Gabor function with the guessed parameters fits the data. If we repeatedly guess sets of parameter values, and plot the resulting log likelihoods against the guessed parameter values, we obtain the plot shown in Figure 4.

Another way of saying this is that Figure 4 is the projection into log likelihood space of the signal shown in Figure 3. When we examine this composite signal in log likelihood space, we see four well-separated components, and expect that we should be able to recover each component, one by one. The two parameters are ω, the oscillation frequency (0-0.06) and τ, the time of occurrence (0-2000) of the event. The vertical dimension is the log-likelihood that the observed data contains a Gabor function with the given pair of the parameter values.

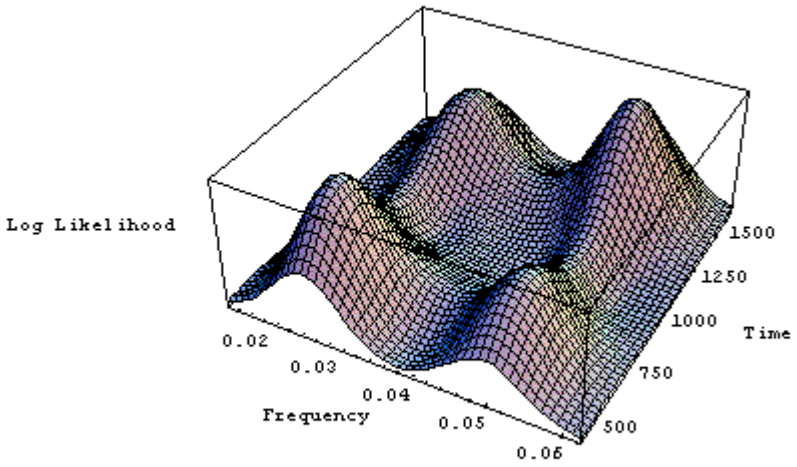

Figure 4. Likelihood of fit of signal to Gabor function.

Given the signal shown in Figure 3, the most likely Gabor function that fits the signal, is the one whose parameters lead to the greatest log likelihood. Using a global optimizing algorithm, with the objective function being the log likelihood as a function of the time and frequency parameters, we can readily find the optimal combination of parameters. [Note that local optimization algorithms such as gradient descent are vulnerable to being trapped in local optima.] As indicated by the peak in Figure 4, the optimum value of the objective function occurs at a time of 605.4 and a frequency of 0.0248.

The most likely Gabor function in the signal shown in Figure 3 is plotted in the left-hand plot in Figure 5. This is the Gabor function whose parameters are found at the global optimum in Figure 4. When this estimated signal is subtracted from the signal in Figure 3, the residual in the right-hand plot of Figure 5 is obtained. Note that the most likely Gabor function has been separated from the signal without disturbing the other information in the signal.

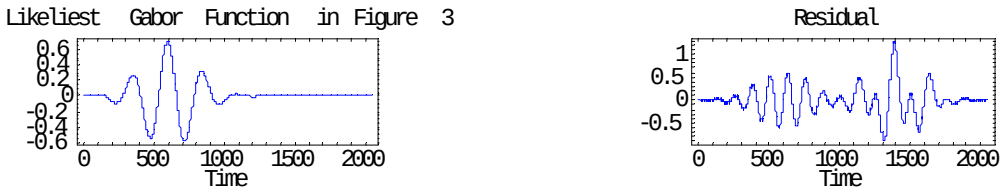

Figure 5. Most likely signal and residual.

Since we have reason to believe that the signal should contain several Gabor functions plus noise, we can apply precisely the same methods to the residual and obtain the next most likely signal and its residual with the next likeliest Gabor function removed. In Figure 6 we see the results of repeating this process until the residual is reduced to noise. The second through fourth likeliest Gabor functions have times of 1397.9, 612.2, and 1400.9 and frequencies of 0.0251, 0.050, and 0.050 respectively. The residual after removing the four likeliest Gabor functions is noise and cumulative rounding error.

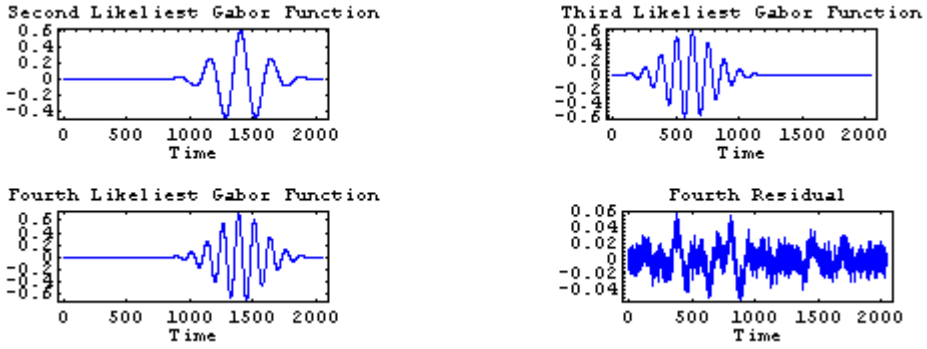

Figure 6. Other components of overlapped signal.

The point of the foregoing discussion is to demonstrate that overlapping signals can be separated without disturbing each other. Bayesian parameter estimation provides the optimal estimates of the models underlying the signal, however, it is computationally costly. For the purposes of on-line, real-time monitoring, the goal is to find a method that is almost as good as Bayesian, but much faster.

In laser ultrasonics, the signals of interest are oscillating bursts. This suggests that wavelet analysis might produce acceptable performance for an on-line instrument. The wavelet basis function is an oscillating burst and the discrete wavelet transform is implemented as a bank of computationally inexpensive finite impulse response (FIR) digital filters.

The idea behind wavelet analysis is that the signal can be considered as the weighted sum of overlapping wavelet functions.[8] In fact, any signal of finite bandwidth and finite duration can be completely characterized as a weighted sum of a finite number of scaled and shifted versions of the underlying wavelet. The concept is similar to Fourier analysis, in which the time series signal can be considered a weighted sum of sinusoids at various frequencies, with the transform coefficients being the weights. The practical meaning of the wavelet transform of a signal is that each coefficient of the transform is the weight, or relative amount of information the wavelet at that particular value of scale and shift contributes to the overall signal.

For the results reported in the next section, the wavelet analysis was performed with the Daubechies 10-coefficient least asymmetric discrete wavelet.[9] Discrete wavelets are not expressible in closed form. Plots of the wavelet and its corresponding scaling function were computed with Daubechies cascade algorithm, and are shown in Figure 7.

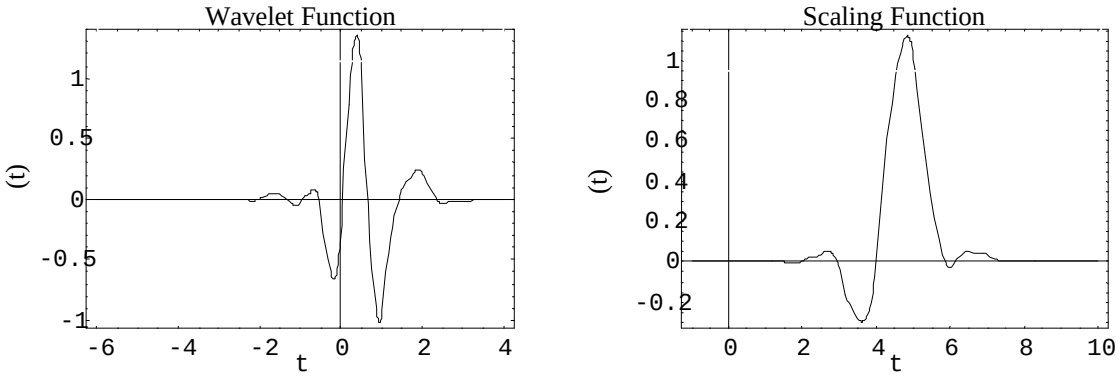

Figure 7. Wavelet and scaling functions for 10-coefficient least asymmetric wavelet.

Suppose that a time-domain input consists of a list of 960 evenly-spaced samples of a band-limited signal. The discrete wavelet analysis results in a list of 960 wavelet-domain coefficients output in response to each 960-element time series input signal. As sinusoids at different frequencies are orthogonal to each other, so also are scaled and translated versions of these wavelet functions orthogonal to each other. This means that Parseval's theorem holds for discrete wavelet transform; the amount of energy in the signal in the wavelet domain is exactly the same as the amount of energy in the signal in time domain. Energy is regarded as proportional to the information in the signal.

For a transient signal, such as an oscillating burst, it is expected that most of the energy in the signal will be concentrated in relatively few of the wavelet coefficients, with all the others having values very close to zero. The energy contributed by each coefficient is the square of the coefficient. The fraction of the total signal energy contributed by each coefficient is the energy of the coefficient divided by the total energy of the signal (the sum of the squares of all the samples of either the time series or all the wavelet coefficients). Sorting the wavelet coefficients from greatest to smallest, the cumulative energy of each coefficient is the fraction of total signal energy contributed by that coefficient, plus the fractions of signal energy contributed by all the larger coefficients. For a typical data set collected in this research, the largest coefficient (out of 960) contains over 30% of the total energy of the signal, and that the ten largest contain over 98% of the total signal energy.

In many cases, once the signature of interest is separated from everything else in the sensor signal, the detection of flaws in the workpiece is straightforward. Typically, the workpiece is scanned across a physical range by the LBU apparatus. Continuity of the mathematical properties of the signature across scans suggests an unflawed workpiece. The appearance of an abrupt discontinuity, such as a dramatic localized change in peak of window location (often suggesting a localized change in time of flight of an ultrasonic echo) suggests the presence of a flaw.

Consistent interpretation of weld signatures must be model-based, where the model is in some reasonable sense a description of the underlying physical reality. The various bursts that appear in the sensor output have physical causes, and in ultrasonics, the physics of the causes is typically well understood. Signal processing by the methods described in this section exploit this knowledge to wring the maximum of new and relevant information from the sensor output data, and provide an indication of confidence in the results. This is in contrast with the widely popular paradigm of signal processing in which an empirical model is surmised without respect to the underlying process. The empirical approach works for some data sets, but give no indication of when it might break down for the next unknown trial.

4. RESULTS

A typical result of LBU inspection of a tailor-welded blank is shown in Figure 8. Each plot is the output voltage of a LBU photo-detector as a function of time, as the LBU system makes 30 successive scans across a workpiece with a flawed weld. It would be very difficult to detect the flaw simply from a visual inspection of these raw data.

From the data for each of the 30 scans, the most likely model (or dominant component) was computed using Bayesian parameter estimation. For each scan, the waveform of the most probable model is a chirped Gaussian-windowed sinusoid. The most probable is taken to be the dominant component of the signal. These are plotted in Figure 9. This dominant component appears to be a biasing effect characterizing the experimental setup. As shown in Table 1, this component contains approximately 95% of the signal energy for each scan.

The second likeliest component for the data of each of the 30 scans is plotted in Figure 10. For each scan, the first residual is computed by subtracting the estimated data for the likeliest component from the original data for the scan. The most probable model of the residual computed by Bayesian parameter estimation is taken as the second likeliest component of the original signal. For each scan, the waveform of the second likeliest component is a non-chirped exponentially rising sinusoid. This appears to be a reflection of an exponentially decaying sinusoidal component off the back of the workpiece. As shown in Table 1, the second likeliest component contains approximately 2–3% of the signal energy for each scan.

The third likeliest component for the data of each of the 30 scans is plotted in Figure 11. For each scan, the second residual is computed by subtracting the estimated data for the second likeliest component from the first residual for the scan. The most probable model of the second residual computed by Bayesian parameter estimation is taken as the third likeliest component of the original signal. For each scan, the waveform of the third likeliest component is the sum of several Gaussian pulses. This appears to be a reflection off the weld. As shown in Table 1, the third likeliest component typically contains less than 1% of the signal energy for each scan.

Figure 11 reveals some interesting information about the workpiece. For the first 21 scans, the third likeliest component is located consistently (except for scans 14 and 15) in the neighborhood of time = 500. This is consistent with a visually detectable pinhole flaw in the weld in the vicinity of scan 14. The inconsistent location of the third likeliest component in scans 22 through 30 suggests other flaws in the workpiece that are not revealed by visual inspection.

The Bayesian parameter estimation algorithm has comparatively high computational complexity. For an on-line, real-time system, it is desirable to find an algorithm that gives performance that is almost as good as Bayesian, but with substantially lower computational cost. For example, discrete wavelet processing is almost as good as Bayesian for detecting short bursts and is fast enough to be implemented on a presently available digital signal-processing chip. For a given setup, the real-time algorithm might be wavelet-based and the Bayesian analysis might be performed only occasionally, to verify calibration or other guidance on how to exploit the wavelet output.

For these data, the tendency of the signal energy to concentrate in relatively few wavelet basis functions is shown in Table 2. Discrete wavelet transforms were computed of each of the 30 scans, using the 10-coefficient Daubechies least-asymmetric wavelet. For the first scan, the 8 largest wavelet coefficients contain 97% of the signal energy. The reason for asking which wavelet coefficients contain 97% of the signal energy is that the Bayesian analysis suggests that 97% of the signal energy is concentrated in the two likeliest components of the signal, and that these two components do not contain information about the weld. If we can find a cheap method of identifying this part of the signal, we can discard it with reasonable expectation that we are not discarding very much of the useful information in the signal. For the same scan, the 52 largest wavelet coefficients contain 99.99% of the signal energy. As shown in Table 1, other scans show a similar, but not identical energy distribution.

This energy distribution is consistent with what was gleaned from the Bayesian analysis and provides a convenient way to parse the signal to eliminate irrelevant information. The largest 7 or 8 wavelet coefficients contain the 97% of the signal energy that does not include information about the weld. The largest 60 (or thereabouts) wavelet coefficients contain in excess of 99.99% of the total signal energy. This suggests that we can zero out the other 900 smallest and the 7 or 8 largest wavelet coefficients in each signal with practically no loss of information.[10] The information we seek resides in the part of the signal that remains.

Figure 12a was constructed as follows. The wavelet transform was computed for each signal. In wavelet space, the largest coefficients containing the first 99.99% of the signal energies were retained and the others zeroed out. Then each zeroed-out set was inverse wavelet transformed to recover the approximate time series. Then each time series was subtracted from the corresponding original time series. The resulting residuals are plotted as a density plot. The vertical axis corresponds to the 30 spatial locations on the workpiece. The horizontal axis corresponds to time. The gray level is strength. The residuals are mostly (but not entirely) noise.

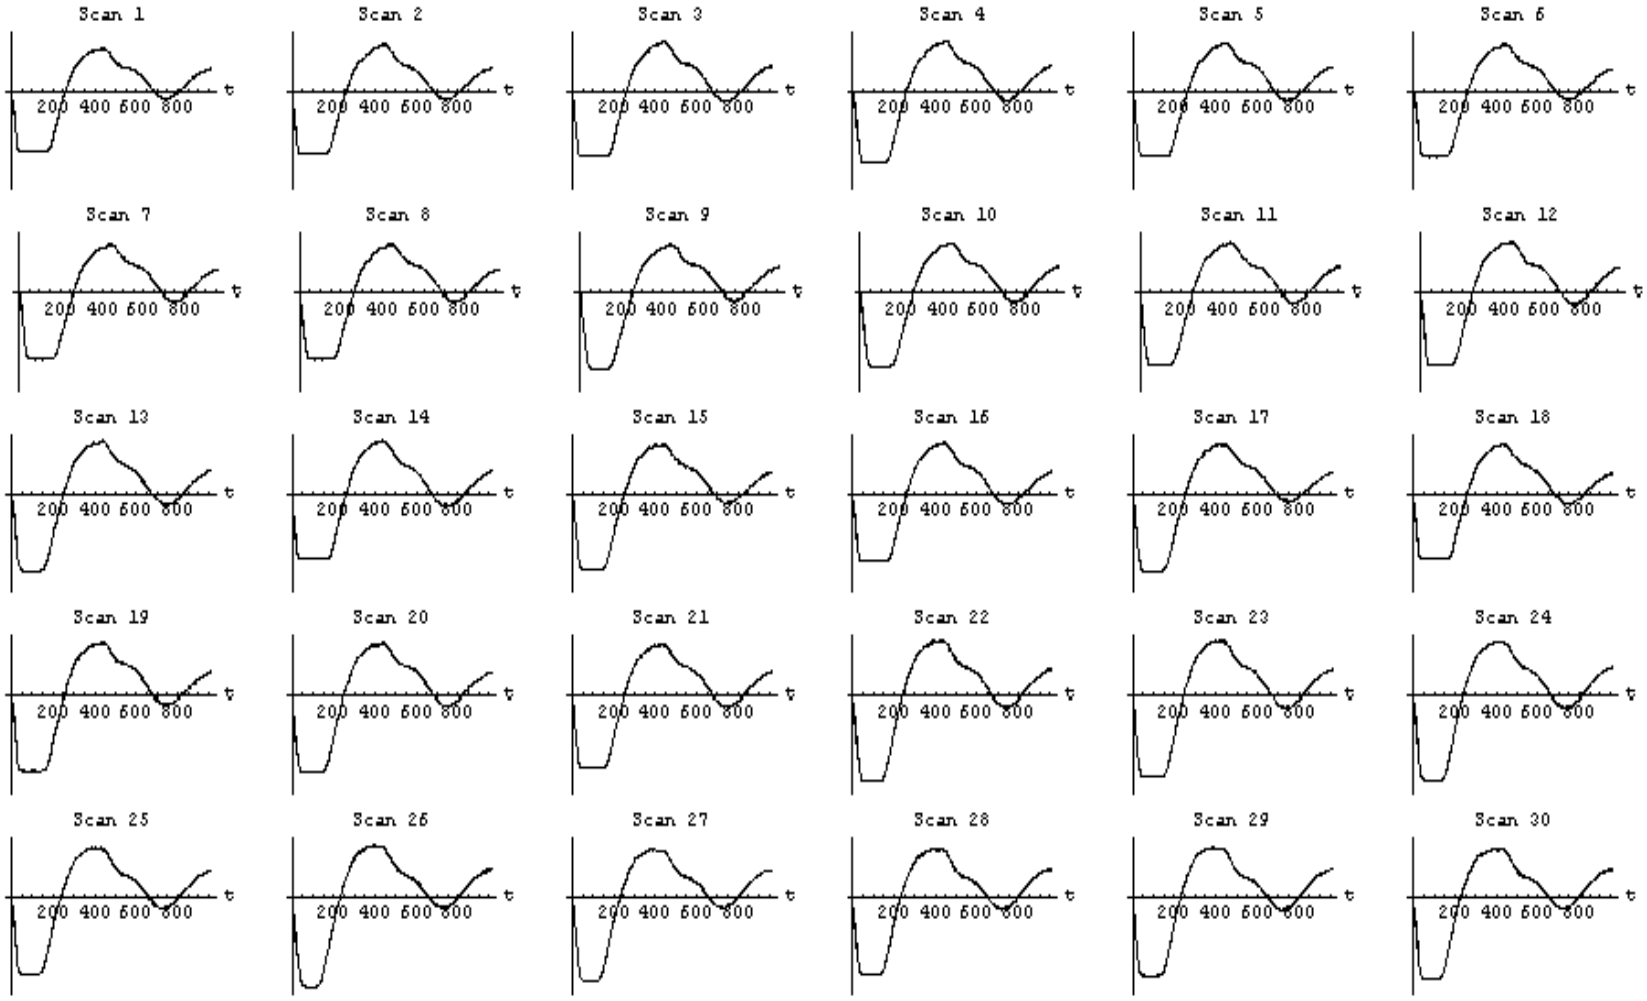

Figure 8. Successive scans across a flawed weld.

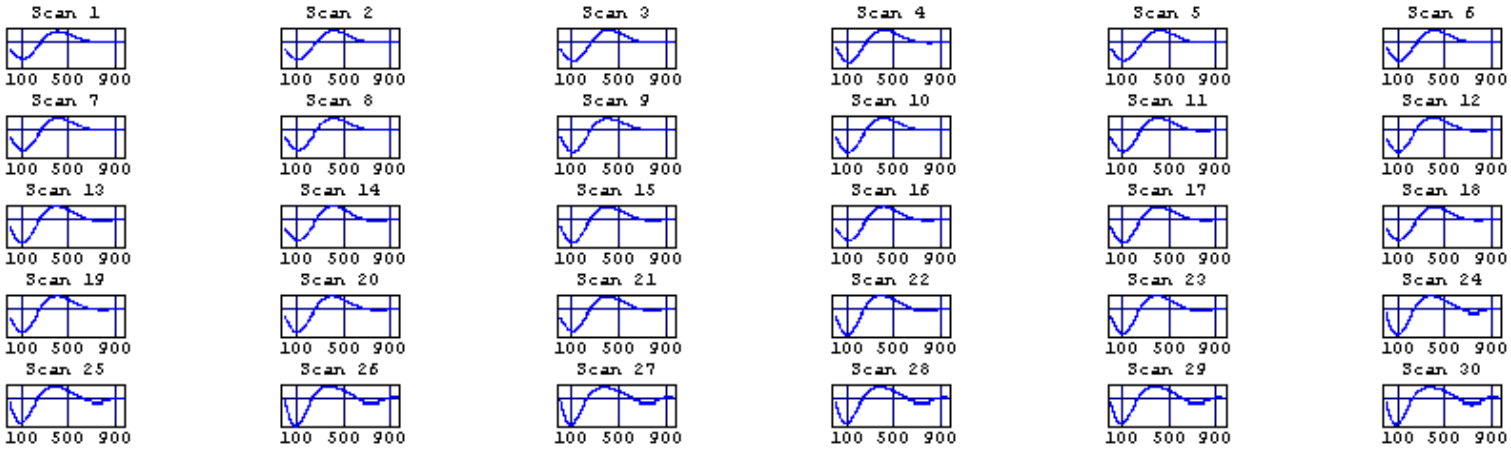

Figure 9. Likeliest component of each scan.

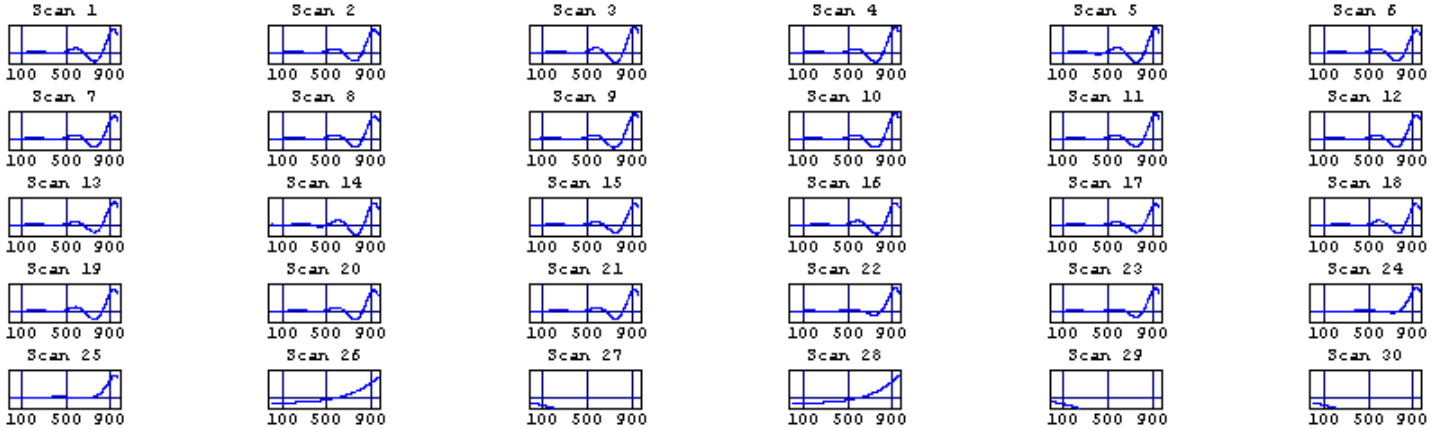

Figure 10. Second likeliest component of each scan.

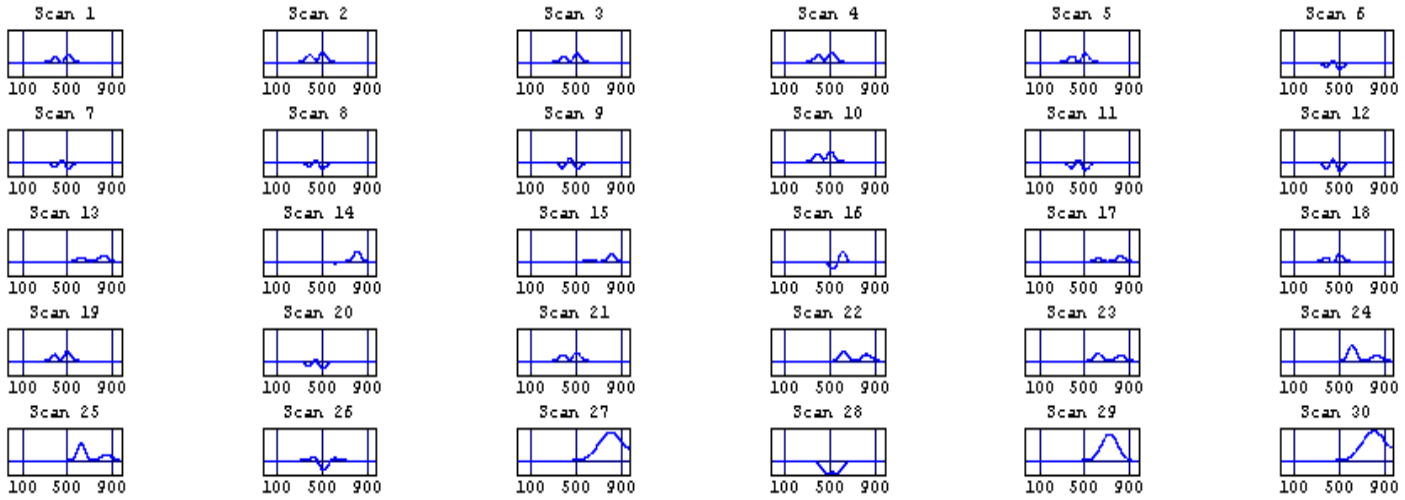

Figure 11. Third likeliest component of each scan.

Table 1. Percentage of energy in three likeliest components of each scan.

Scan	1st	2nd	3rd	Scan	1st	2nd	3rd	Scan	1st	2nd	3rd
1	94.6	3.4	0.2	11	95.1	3.	0.2	21	96.	2.3	0.2
2	94.8	3.2	0.2	12	95.4	2.7	0.2	22	96.3	1.9	0.2
3	94.5	3.5	0.2	13	96.2	2.1	0.1	23	96.1	2.	0.1
4	94.9	3.2	0.2	14	95.6	2.5	0.1	24	95.8	2.	0.6
5	94.7	3.4	0.2	15	96.4	2.1	0.1	25	95.6	2.	0.9
6	95.3	2.8	0.1	16	95.7	2.5	0.2	26	96.4	2.	0.2
7	95.3	2.8	0.1	17	96.6	2.	0.1	27	95.8	2.3	1.6
8	95.3	2.8	0.1	18	96.	2.3	0.1	28	94.9	2.7	0.3
9	95.2	3.1	0.2	19	96.3	2.1	0.2	29	95.4	2.3	1.7
10	95.1	3.1	0.2	20	96.3	2.1	0.2	30	95.7	2.3	1.6

Table 2. Percentage of energy in largest wavelet basis functions of each scan.

Scan	97%	99.99%	Scan	97%	99.99%	Scan	97%	99.99%
1	8	52	11	7	60	21	7	61
2	8	58	12	7	68	22	7	57
3	8	53	13	7	62	23	7	66
4	7	56	14	8	54	24	8	51
5	8	62	15	7	63	25	8	62
6	7	49	16	7	53	26	8	69
7	7	49	17	7	54	27	8	60
8	7	49	18	7	51	28	8	61
9	7	56	19	7	62	29	8	57
10	7	53	20	7	58	30	8	65

In wavelet space, the 8 largest coefficients (containing 95–96% of the signal energies) were retained, and the others zeroed out. Then each zeroed-out set was inverse wavelet transformed to recover the approximate time series. Then each time series was subtracted from the corresponding original time series. The resulting 4–5% residuals are contour-plotted below in Figure 12b.

In the same wavelet space, the 9 largest coefficients (containing 96–97% of the signal energies) were retained and the others zeroed out. Then each zeroed-out set was inverse wavelet transformed to recover the approximate time series. Then each time series was subtracted from the corresponding original time series. The resulting 3-4% residuals are contour-plotted in Figure 12c. Note there is a little difference between Figures 12b and 12c.

The reasonable place to search for features of weld defects is in the region of the signal between the bias (biggest eight coefficients of each signal, or thereabouts) and the noise (smallest 900 coefficients of each signal). In Figure 12d, 97.6% of the signal energy is assumed to be attributed to biasing effects. This bias is subtracted from the signal approximation constructed from the wavelet coefficients constituting of 97.8% of the signal energy. This difference constitutes 1.2% of the original signal energy and has a fairly dramatic global minimum whose contour is plotted in Figure 12d. This shows up in scan 14 and time 750. This corresponds to the change in the third likeliest component of the signal as revealed by Bayesian analysis and to the pinhole defect in the weld in the workpiece.

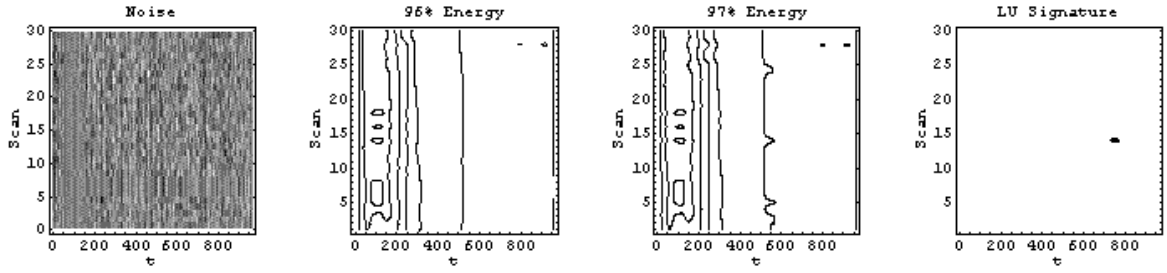

12a. 12b. 12c. 12d.

Figure 12. Separation of signal energies by wavelets.

5. CONCLUSIONS AND FURTHER RESEARCH

In conclusion, we have demonstrated the capability to detect localized weld defects using a computationally efficient processing approach. This paper reports the initial results of our experiments and analysis of data for LBU inspection of tailor-welded blanks. In the ongoing research we will identify the connection between the Bayesian-derived models and the underlying physical processes. We will also investigate the comparative reliability and computational cost of wavelet and Bayesian feature extraction methods.

In future work we will seek to classify individual defects, including those encountered in the welding process. We will also devise a practical implementation of the signal-processing algorithm in real-time. The ultimate goal of this research is to construct a prototype of an on-line LBU weld inspection device for tailor-welded blanks.

This work could lead to spin-offs for other on-line inspection in other processes. Additional follow-on research might include the examination of other types of continuous seam joints for distinguishing features. In addition, the Lamb-wave modeling and feature extraction would be directly applicable to inspection of other products fabricated from thin metal sheet.

ACKNOWLEDGMENTS

This research is a joint effort between Lasson Technologies, TWB, Inc., and Oak Ridge National Laboratory, managed and operated by Lockheed Martin Energy Research Corporation for the U.S. Department of Energy under Contract No. DE-AC05-96OR22464, and is sponsored under ER-LTR CRADA ORNL 98-0524.

REFERENCES

1. M. B. Klein, G. J. Dunning, P. V. Mitchell, T. R. O'Meara, and Y. Owechko, "Remote Laser-Based Ultrasonic Inspection of Weld Joints for High Volume Industrial Applications," *Review of Progress in Quantitative Nondestructive Evaluation*, D. O. Thompson and D. E. Chimenti, Editors, Volume 15, pp. 2257–2264, 1996.
2. J. Bussiere, S. McQueen, R. Shannon, J. Spicer, and R. Russo, *Report of the Workshop on Industrial Applications of Laser Ultrasonics*, produced by U.S. Department of Energy Office of Industrial Technologies, pp. 30–32, Draft, January 1998.
3. S. W. Kercel, L. E. Labaj, and V. M. Baylor, "Comparison of Enclosed Space Detection System with Conventional Methods," American Defense Preparedness Association, *Proceedings of 13th Annual Security Technology Symposium*, Virginia Beach, June 9–12, 1997.
4. S. W. Kercel, W. B. Dress, A. D. Hibbs, and G. A. Barrall, "Investigation of Wavelet-Based Enhancements to Nuclear Quadrupole Resonance Explosives Detectors," in *Wavelet Applications IV*, Harold H. Szu, Editor, Proc. SPIE 3391, pp. 424–434, 1998.
5. W. B. Dress and S. W. Kercel, "Wavelet-Based Acoustic Recognition of Aircraft," in *Wavelet Applications*, Harold H. Szu, Editor, Proc. SPIE 2242, pp. 778–791, 1994.
6. W. A. Simpson and D. J. McGuire, *Phase and Group Velocities for Lamb Waves in DOP-26 Iridium Alloy Sheet*, ORNL/TM-12749, Oak Ridge National Laboratory, Oak Ridge, July 1994.
7. L. Bretthorst, *Bayesian Spectrum Analysis and Parameter Estimation*, pp. 1–5, Springer-Verlag, Berlin, 1988.
8. S. W. Kercel, *A Near-Real-Time Instrument for Wideband Magnetic Field Monitoring with Simultaneous Time and Frequency Localization by Multiresolution Filter Bank*, pp. 109–113, UMI, Ann Arbor, 1996.
9. I. Daubechies, *Ten Lectures on Wavelets*, pp. 199–209, Society for Industrial and Applied Mathematics, Philadelphia, 1992.
10. D. Donoho, "De-noising by Soft-thresholding," *Technical Report*, Department of Statistics, Stanford University, 1992.