ELSEVIER

Applied Acoustics, Vol. 50, No. 4, pp. 279–296, 1997
Published by Elsevier Science Ltd
Printed in Great Britain
PII: S0003-682X(96)00040-0 0003-682X/97/$17.00+0.00

Quantifying Larval Infestation with an Acoustical Sensor Array and Cluster Analysis of Cross-Correlation Outputs

D. Shuman, D. K. Weaver and R. W. Mankin

United States Department of Agriculture, Agricultural Research Service, Center for Medical, Agricultural, and Veterinary Entomology, 1700 SW 23 Drive, Gainesville, FL 32608, USA

(Received 2 February 1996; revised version received 22 July 1996; accepted 16 August 1996)

ABSTRACT

An automated, computer-based system was designed to quantify infestation of internally feeding larvae in a grain sample by obtaining data correlated with the location of sound sources. Information related to the relative arrival times of insect feeding sounds to an array of acoustic sensors is obtained despite the low signal to noise ratios and the differential distortion induced by sound propagation through the non-uniform grain medium to the different sensors. This is achieved by employing parallel acquisition of all sensor outputs and cross-correlation analyses of all adjacent sensor pairs in the vicinity of the sensor with the largest signal. The peak location times of the resulting cross-correlograms cluster together for multiple sounds produced by the same insect but otherwise are more broadly distributed. A cluster analysis algorithm was developed to group sounds with similar 'fingerprints' (i.e. patterns of peak locations across several cross-correlograms). Each sufficiently large group of matching sounds indicates the presence of an insect. Published by Elsevier Science Ltd

Keywords: insect detection, stored-grain, sample inspection.

INTRODUCTION

The presence of insects in stored grain is a major factor in the determination of quality under current mandated industry standards. Currently, grain inspection involves manually counting the insects sieved from a defined sample, usually 1 kg. This procedure limits detection to externally feeding larvae and adults. However, larvae of some economically important species,

such as the rice weevil, *Sitophilus oryzae* (L.), and lesser grain borer, *Rhyzopertha dominica* (F.), feed inside kernels of grain and are not detected. If adults are not present, because they either have not yet emerged from infested kernels or have been removed by cleaning or other manufacturing processes, grain internally infested may be mistaken for uninfested grain. Current laboratory methods for the detection of internal feeders (e.g. X-ray, carbon dioxide production, resonance spectroscopy, and ELISA testing) are costly, time-consuming, and generally are not implemented (X-ray technology, which is unable to differentiate between live and dead insects, is sometimes employed by the milling industry). There is a need for a rapid, quantitative, and economical method for detecting both adults and larvae of major insect pests in grain.

Detection of insects in fruits and grain by amplifying their feeding and movement sounds was suggested by Brain,[1] but technical difficulties prevented the development of practical systems.[2–4] Recent technological advances (sensitive detectors, suitable band-pass filters, and inexpensive computers) have stimulated studies[5–8] directed at the development of practical acoustic detection systems for stored-product insects.[9] Although the latter studies have demonstrated a strong correlation between the number of insects in a sample and total acoustical activity, it is not sufficiently accurate for the rapid grading of an unreplicated grain sample. Insect size and distance to the transducer also strongly influence infestation estimates based on the measured acoustical activity.[10] A system that endeavored to minimize the influence of these factors by determining the number of loci within a sample from which sounds were originating was developed by Shuman *et al.*[11] This first-generation prototype was named the Acoustic Location Fixing Insect Detector (ALFID) system.

The ALFID principle of operation is that the transit time of a sound is directly proportional to the distance traversed. The prototype incorporated a linear array of 16 acoustic sensors mounted in one wall of a rectangular grain sample container. By employing amplitude threshold detection of the amplified sensor outputs, it attempted to identify the first and second adjacent sensors in the array to receive a particular sound and to determine the time-delay between these two detections. This would ideally localize the sound source to a (hyperbolic) surface in the grain container. The success of this method was dependent upon detecting corresponding single points on the sensors' output waveforms. However, the differences in the signal levels as a function of different source to sensor distances, the low signal to noise ratios and the differential distortion induced by propagation of the feeding sounds through the non-uniform medium to the different sensors in the array, made the time-delay data unusable due to its large variability. As a result, only the identities of the first and second adjacent sensors in the array to receive a

particular sound were used. Owing to the large acoustic attenuation of the grain, sometimes a sound was only threshold detected by one sensor which, at best, localized the sound's source within a volume twice as wide as the distance between sensors. This coarse resolution made differential detection of multiple insects or between an insect and grain settling sounds difficult. Even so, the performance of the system demonstrated conceptual feasibility and led to the development of a second-generation ALFID system whose design is described in this paper. Weaver *et al.*[12] compares the performance of the two ALFID prototypes.

SECOND-GENERATION ALFID SYSTEM DESIGN

Overview

The main intent of the redesign of the ALFID system was to obtain improved sound propagation time-delay data in order to increase its resolution of sound source locations. The technique of cross-correlation is well suited for extracting this time-delay information when the signals are contaminated with uncorrelated noise.[13] It involves shifting two sensor output waveforms relative to each other along the time axis until a best fit occurs. The magnitude of the shift is the sought time-delay and is based on entire signal waveforms and not just a single point on each waveform as is the case with amplitude threshold detection. This technique has been widely used for locating sound-producing organisms in noisy environments with inhomogeneous sound media.[14] This mathematically intensive approach necessitated the acquisition of the amplified sensor outputs and the subsequent computer processing of the sound data. Acquiring data from all 16 sensors, even when the output signal from only one sensor exceeds the (trigger) threshold level, reduces the problems previously associated with the large acoustic attenuation of the grain. This is because a subset of the sensors near the sound source will usually have low level but discernible information that can be extracted by the cross-correlation analysis. As will be seen, it was not necessary (or practical) to locate the insects but only to discriminate and match time-delay 'fingerprints' unique to individual sound source locations in order to determine the number of sound-producing insects. This matching of relative time-delay data eliminates errors due to absolute but consistent changes to acoustic signals introduced by the non-uniform medium (e.g. propagation velocity and multi-path distortion), as well as mismatches in the phase–frequency responses of the 16 sensor channels. Hence, the name of the ALFID system was changed from the first-generation prototype's 'Acoustic Location Fixing Insect Detector' to the second-generation's 'Acoustic Location Fingerprinting Insect Detector'.

Acoustic sensors and amplifiers

The second-generation ALFID system uses a highly sensitive piezoelectric microphone element (Kobitone 25LM022, Mouser Electronics, Mansfield, TX). This sensor has a bandwidth of 3.5–5 kHz with a resonant peak at 4 kHz, which is a good match to the spectral content of the larval feeding sounds. A fine wire mesh over the face of the sensor helps keep grain dust out of the sound apertures in its aluminum casing. To minimize the lengths of its high impedance leads, and thus reduce electromagnetic pickup, each sensor is mounted on its own amplifier circuit board. The amplifier is an impedance matching, low noise design with 85 dB of gain and some bandpass filtering (1–10 kHz).

Grain sample container

The ALFID system still uses a 1-kg grain sample container with an array of 16 sensors,[11] but it is now designed to maximize sensitivity throughout the volume based on the sensor's measured spatial sensitivity. It achieves this by use of a PVC cylindrical tube with an inner vinyl sound barrier layer and two rows of eight sensors mounted directly across from each other (Fig. 1). This configuration also permits computation of relative time-delays in two orthogonal directions, from pairs of sensors across the tube as well as from pairs of sensors along its length. Theoretically, this would localize a sound source in the grain container to a line formed by the intersection of two hyperbolic surfaces. During sample testing, the container is oriented horizontally with the sensors on the sides to equalize the grain pressure on the sensors. For field use, the grain container is housed in a sound attenuation box to reduce the effects of ambient noise.[15] Four additional acoustic sensors are mounted on the outside of the container to provide a noise-masking function discussed later.

Sound data acquisition

The amplified sensors' output signals are sampled and acquired by a 16-channel, 1 MHz, D/A board (Flash-12 Model 1, Strawberry Tree, Sunnyvale, CA) installed into a PC computer (Fig. 2) and controlled by a custom-designed software driver. The insect sounds occur randomly and can last up to 20 ms (due to grain ringing). A larva (fourth instar) will typically produce less than five detectable sounds/min. For this reason, in order to not fill the PC's memory with incoming background noise, the D/A board is configured to transfer data being acquired into its own ring buffer memory to the PC's memory only when the board is triggered by a sufficiently loud sound. The

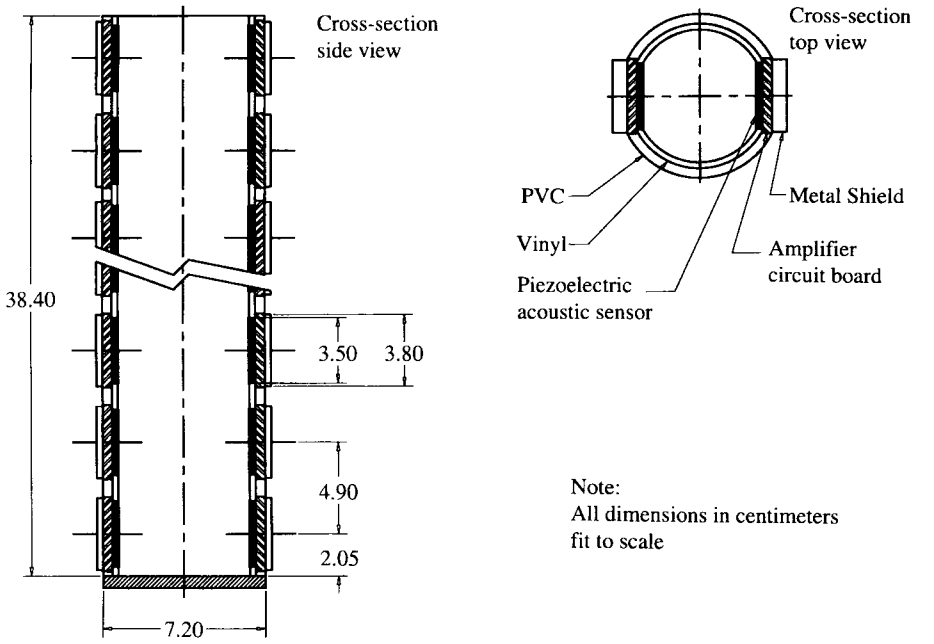

Fig. 1. Mechanical drawing of the ALFID grain container.

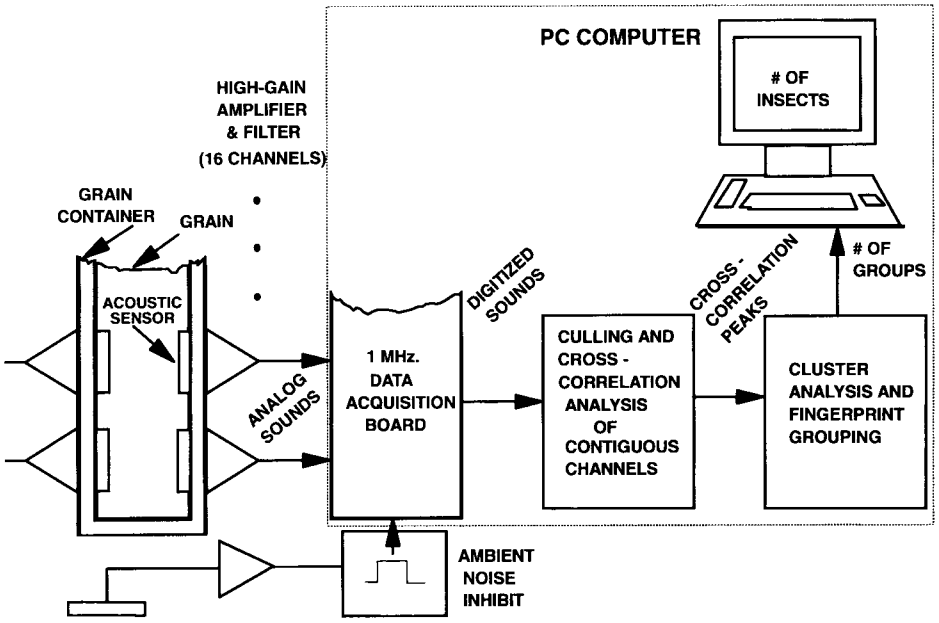

Fig. 2. ALFID system functional block diagram.

trigger threshold level, adjusted to be above the noise level, corresponds to a sound pressure level of about 23 dB (±6 dB due to the range of the sensor/amplifiers' gain within the band of predominant insect sound spectral energy). Since the identity of the sensor channel with a signal large enough to cause triggering is unknown, a trigger generation circuit was designed to logically 'OR' the output of all the amplified sensors' outputs. A generated trigger signal that rises above the board's trigger amplitude threshold level initiates the transfer of all 16 sensors' acquired output signals, including channels with signals that may remain below this threshold level (Fig. 3), to the PC's memory. Once triggered, the board's trigger input is disabled so that subsequent peaks of a sound's received waveforms (as manifest in the trigger signal) do not initiate further data transfers to the PC's memory until the system is ready for a new sound.

The grain sample test duration is an input variable for the custom software driver that controls the data acquisition operation of the ALFID system. A longer test duration increases the probability that an insect will be detected but slows down the throughput of grain samples. The durations for tests with rice weevil larvae typically range from 10 to 30 m.

A noise-masking feature (Fig. 2) was designed to reduce the possibility of loud ambient sounds not sufficiently blocked by the sound attenuation box from being interpreted as an insect sound. If the outputs of any of four

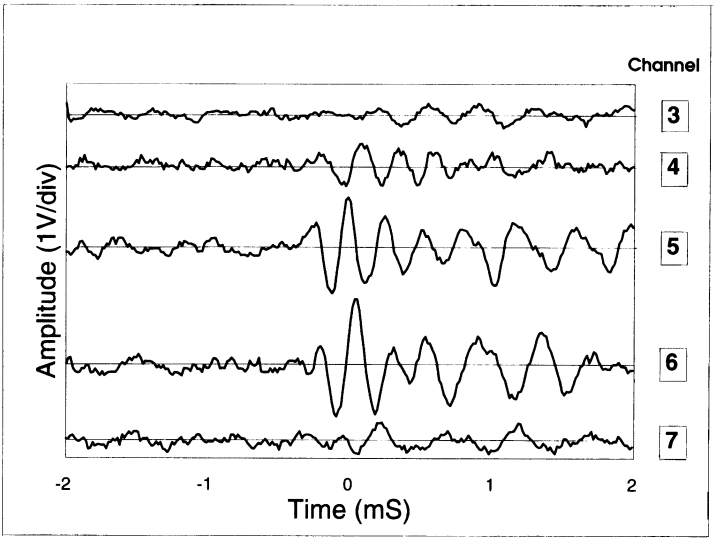

Fig. 3. Typical signal waveforms from a single insect sound as seen on the amplified outputs of different sensors. The stored data begins 2 ms before, and ends 2 ms after, the triggering of acquisition occurred.

acoustic sensors mounted on the outside of the grain container exceed an amplitude threshold, an inhibit signal is generated which blocks the transfer of the acquired ambient sound data (in the D/A board's ring buffer) to the PC's memory.

The first few cycles of sound waveforms are the most important segments of the sensors' output signals for determining the propagation time-delay because they are less likely to be contaminated by multipath and reflected waves or the lower frequency resonant 'ringing' of the grain. To insure acquisition of these first few cycles, the D/A board is configured to provide pretrigger data. It is continuously acquiring sampled sensors' output signal data into a ring buffer memory so that new data is overwriting recently (as a function of sampling rate and memory size) acquired data. When a trigger occurs, an additional 2 ms of sampled signal data is written into the ring buffer memory. The board then transfers 4 ms of this sampled data to the PC, beginning with sampled data that was written into the ring buffer 2 ms before the time that the trigger occurred (time = 0 in Fig. 3). The sampling rate of the D/A board is set to 62.5 kHz/channel to provide a 16 μs time shift resolution for the cross-correlation analysis. Therefore an acquired 4 ms block of data fills 4 k samples of buffer memory. To insure no loss of data due to the relatively slow transfer rate of data from the A/D board's buffer memory to the PC's memory in the event of a burst of insect sounds, the A/D board hardware was custom-modified to segment its 64 k sample ring buffer into sixteen 4 k sample ring buffers, each independently addressable and able to store a full 4 ms block of data. The A/D board hardware was also custom-modified to re-enable its trigger input only after 7 ms have elapsed since the last time the trigger signal exceeded its threshold level. This feature prevents a single long sound whose generated trigger signal stays below the trigger amplitude threshold level for short durations (less than 7 ms) from being interpreted as multiple sounds.

In the present ALFID system implementation, the sound analysis phase does not begin until after data acquisition is completed. The analysis can continue for several minutes depending upon the number of sounds collected and the power of the PC computer used. Since the sound data acquisition software is driven by interrupts generated by the trigger signal, future versions of the system could potentially begin analysis of the data during the PC's relatively idle periods between insect sounds.

Culling and cross-correlation analysis

The cross-correlation of two sensors' output signals in discrete (digitized) form, $A(I)$ and $B(I)$, where I is the sample number, is described by the equation

$$R_{A\times B}(j) = \sum_{k=1}^{N} A(k)B(k+j)$$

where j is the number of shifts of A relative to B, and N is the total number of samples of each signal. The range of j depends on the extent of time-delay that needs to be considered. For the sensor spacing in the grain container and the speed of sound propagation in the grain, the j values required are only from -32 to 32 shifts corresponding to -512 to 512 μs. The value of N is 250 samples for the 4 ms interval of acquired sound. Digitized signal values for sample numbers outside the range of 1–250 were zero-padded. The direct method of calculating the cross-correlation with this equation (as opposed to the 'short' FFT method) is more efficient for the small values of N and j. The plot of this equation is referred to as the cross-correlogram (Fig. 4) and the time location of its peak indicates the time-delay for a best fit of the sound waveforms. With 16 sensors arranged in two lines of eight sensors, there are 36 possible contiguous pairs of sensors, including adjacent (14), across (8), and diagonal (14) pairs, for which cross-correlations can be calculated (Fig. 5). These cross-correlations provide information related to a sound's source location in only two dimensions, but the redundancy of the diagonal pairs can help with fingerprinting in this low signal/noise context.

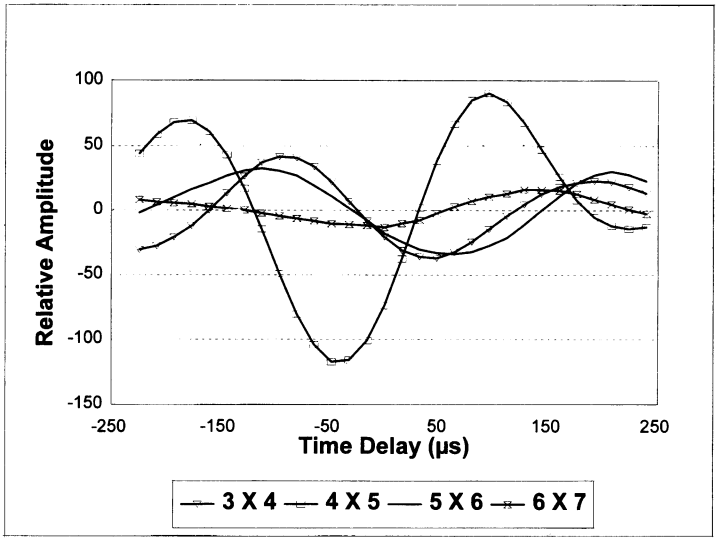

Fig. 4. Cross-correlograms calculated using the signal waveforms from a single insect sound shown in Fig. 3. The full time-delay window used in the cross-correlation analysis (-512 μs to 512 μs prior to imposing the time-delay boundaries) is not shown.

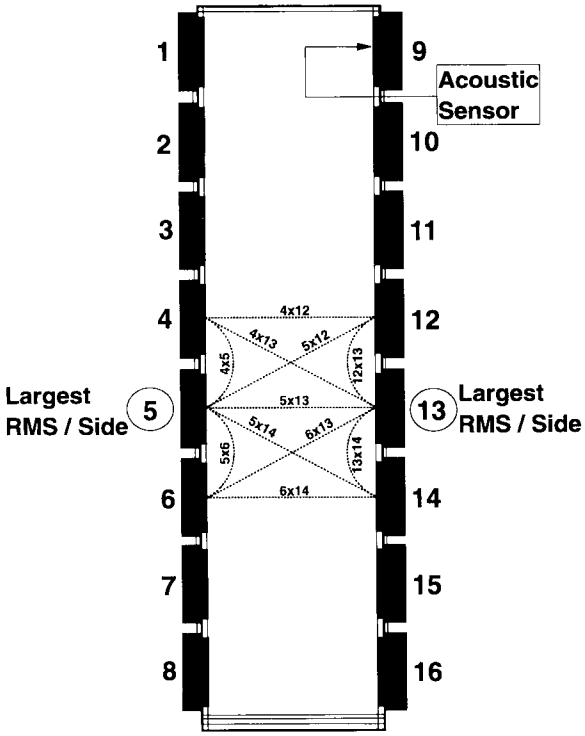

Fig. 5. Depiction of different contiguous sensors pairs for which cross-correlations are calculated, including adjacent, across, and diagonal pairs. To cull potentially error-producing data, only the sensors with the largest RMS signal outputs on each side (for each acquired sound), and those adjacent to them, determine the subset (e.g. the 11 pairs shown of the 36 possible pairs) of cross-correlations to be considered.

Ideally, samples of all 16 sensors' output signals would be obtained simultaneously to prevent time-skew errors in the cross-correlograms. However, the A/D board chosen uses sequential sampling with a 1-μs sampling period. For the 36 possible contiguous pairs of sensors used, the maximum cross-correlation time-skew error is 9 μs (e.g. for 2×11). With a cross-correlogram resolution of 16 μs, the resulting maximum absolute error in peak time locations is one data point or 16 μs. However, the error would be consistent for multiple sounds from the same source location, and so it would not effect the matching of these sounds.

The cross-correlograms usually have two to four peaks (Fig. 4) within their $-512\,\mu$s to $512\,\mu$s time-delay window due to the periodic nature and spectral content of insect sound waveforms. The true peak, having a time location corresponding to a sound source's true location, does not necessarily have the largest amplitude because of differences in waveform shapes (from a single sound emanation) on the outputs of different sensors (Fig. 3). For this

reason the two largest peaks are considered, to increase the probability that at least one of them will match in time-delay with a peak on subsequent sounds produced by the same insect. It is because of these multiple peaks that the true source location is indeterminate. That is, if all the unculled (see below) peaks in all the cross-correlograms resulting from a sound were used in an attempt to establish its source location (by a mathematical analysis similar to triangulation), the result would be a multiplicity of virtual locations. However, the consistency of cross-correlogram peak time locations across sounds produced by one insect (Fig. 6), especially when observed for several different sensor pairs, supports the methodology of using a fingerprint for each acquired sound consisting of two peak time-delay locations for each sensor pair cross-correlation. It is for these reasons that the original goal of 'fixing' a sound source's location to be compared with the locations of other sounds was aborted in favor of establishing a 'fingerprint' related to a sound source's location to be compared to the fingerprints of other sounds, all for the purpose of matching sounds that are from the same source. An overview of the ALFID system algorithm (subsequently described) for reducing the raw acquired data to sound fingerprints that are then compared by a clustering analysis is presented in Fig. 7.

When the output of a sensor that is mostly or entirely noise is used in an cross-correlation calculation, the result is a cross-correlogram with peaks in random locations. With a large collection of sounds, some of these random peak locations could erroneously match different sounds' fingerprints together,

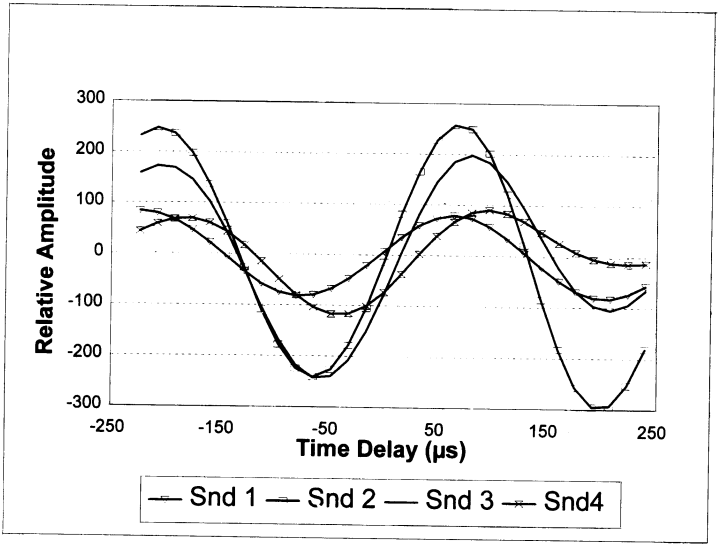

Fig. 6. Cross-correlograms calculated from the outputs of a single pair of sensors (5×6) for multiple sounds produced by the same insect. The peaks tend to be aligned with each other.

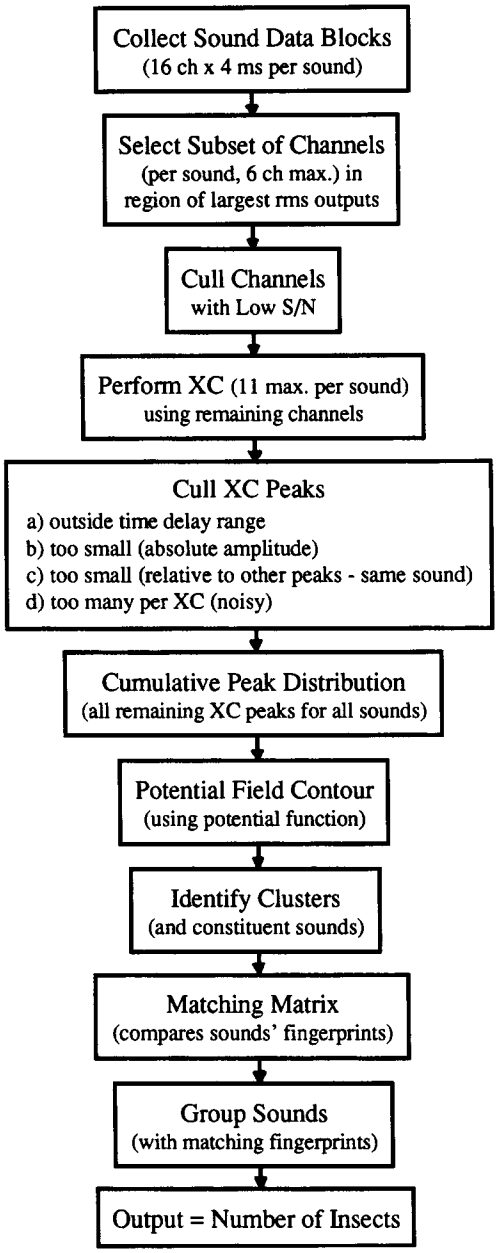

Fig. 7. ALFID system algorithm flowchart.

giving an incorrect assessment of insect infestation. To reduce this occurrence, various methods of culling potentially problematic data are incorporated.

Two methods of culling a subset of sensor channels are employed prior to the cross-correlation analysis. Insect sounds generally reach only a few of the

nearby sensors and so, for the eight sensors on each side of the grain container, only the sensor with the largest RMS amplitude signal and its two adjacent sensors are considered for a maximum of six out of 16 sensor channels or 11 out of 36 contiguous pairs of sensors (Fig. 5). Also, for any one of these six sensor channels to be used, its output signal must first pass a check for an acceptable signal to noise ratio. This is accomplished by specifying a minimum acceptable value (e.g. 1.4) for the ratio of the RMS amplitude of a later segment of the waveform, where the sound would be expected to be found, to the RMS amplitude of an initial pretrigger segment of the waveform where only the channel's noise output would be present (Fig. 3). The location of these segments and the minimum RMS ratio value can be adjusted to empirically optimize the system's performance for a given hardware implementation.[12]

After the cross-correlation analysis is performed, cross-correlogram peaks are culled based on their amplitudes and time-delay values. Peaks with amplitudes below a specified absolute threshold value, indicating a questionable best fit of shifted waveforms, are culled. Also, for each acquired sound, cross-correlogram peaks with amplitudes below a specified percentage of the largest cross-correlogram peak amplitude obtained with that sound are culled. This effectively is adjusting the culling threshold for the loudness of a sound, which is beneficial because the cross-correlation of a large sound signal on one channel with a purely noisy output on another channel can still result in substantial peak amplitudes. These absolute and relative peak amplitude culling thresholds can also be adjusted to empirically optimize the system's performance for a given hardware implementation.[12] Cross-correlogram peaks are also not used if their time-delay values are greater than what is physically possible, given the distances between sensors and the speed of sound propagation in the grain. By substituting the next smaller peak for one known to be false, this time-delay boundary (which is empirically determined independently for adjacent, across, and diagonal sensor pairs) increases the probability that one of the two peaks used is the true peak and, therefore, provides a more reliable fingerprint of the sound. It has also been observed that if the number of peaks in a cross-correlogram is greater than 5, it is a good indication that one of the employed sensor channel outputs is predominantly noise and this cross-correlogram is therefore culled.

Cluster analysis and fingerprint grouping algorithm

The cross-correlation analysis provides a fingerprint for each acquired sound that consists of up to two peak time-delay values for a subset of the 36 possible cross-correlations. Each cross-correlation is treated as an independent

dimension. For the fingerprints of any two individual sounds to be considered a match (meaning they emanated from the same source location), their peak time-delay values should match (within some amount of variability) within several different cross-correlation dimensions. A number of sounds from the same source location should form clusters within these cross-correlation dimensions. A cluster analysis and fingerprint grouping algorithm was developed on the premise that the matching of individual sounds can be inferred from their peak locations being members of the same clusters.

The first step in the algorithm is the plotting of the cumulative peak distribution (CPD) for each cross-correlation dimension (Fig. 8). To determine whether a resulting scattering of peaks is due to sounds emanating from different source locations (and should not be clustered together) or due to the inherent experimental variability from a single source location (and should be clustered together), an approach based on a potential field density method for delineating clusters (Massart and Kaufman[16]) is employed. To perform this discrimination, the method incorporates the peaks' inherent experimental variability in 'smoothing' the peak data by replacing each peak with a potential function (based on this experimental variability).

To generate the potential function, a distribution of cross-correlogram true peak locations was empirically derived from 2200 sounds acquired with

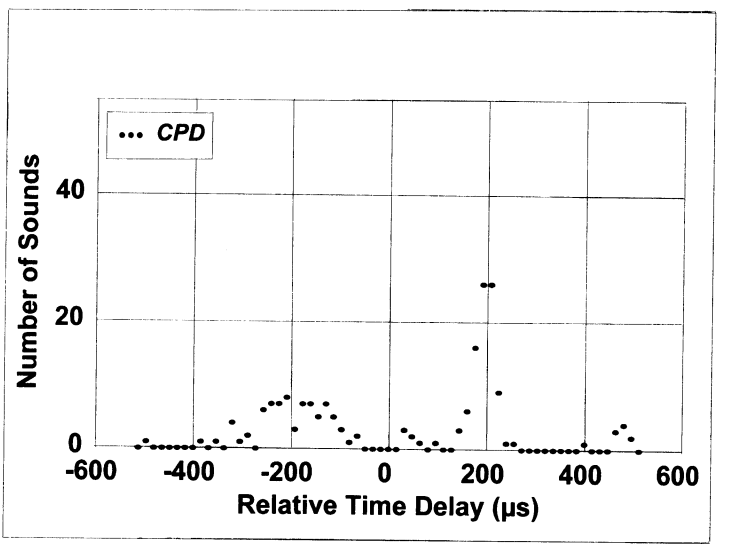

Fig. 8. A typical cumulative distribution of cross-correlations peaks (CPD) in one cross-correlation dimension, obtained with a grain sample containing one infested kernel. Each CPD element indicates the number of sounds that have a cross-correlation peak with that relative time-delay value.

single insects [fourth instars of the rice weevil, *Sitophilus oryzae* (L.)] at various locations in a grain sample container filled with wheat. Individually obtained discrete distributions were aligned (to compensate for different insect locations) along centroids that most evenly split their areas, and then summed. The shape of the resulting composite discrete distribution can be observed as the (scaled) data points in Fig. 9. This composite discrete distribution (representing the expected variability) can be used as a discrete potential function except that the discrete points are not necessarily at the relative time-delay values needed (i.e. the peak of the potential function normalized to 1 at 0 μs, and additional potential function points at $\pm 16\,\mu$s, $\pm 32\,\mu$s, etc.) for proper operation. To derive this discrete potential function, the empirically derived composite discrete distribution was used with TableCurve 2D automated curve fitting software (Jandel Scientific, San Rafael, CA) under the constraints of symmetry and 0 asymptote. The best fitting curve was a Gaussian Lorentzian Cross Product distribution that was subsequently normalized (shown with its algebraic expression in Fig. 9). The

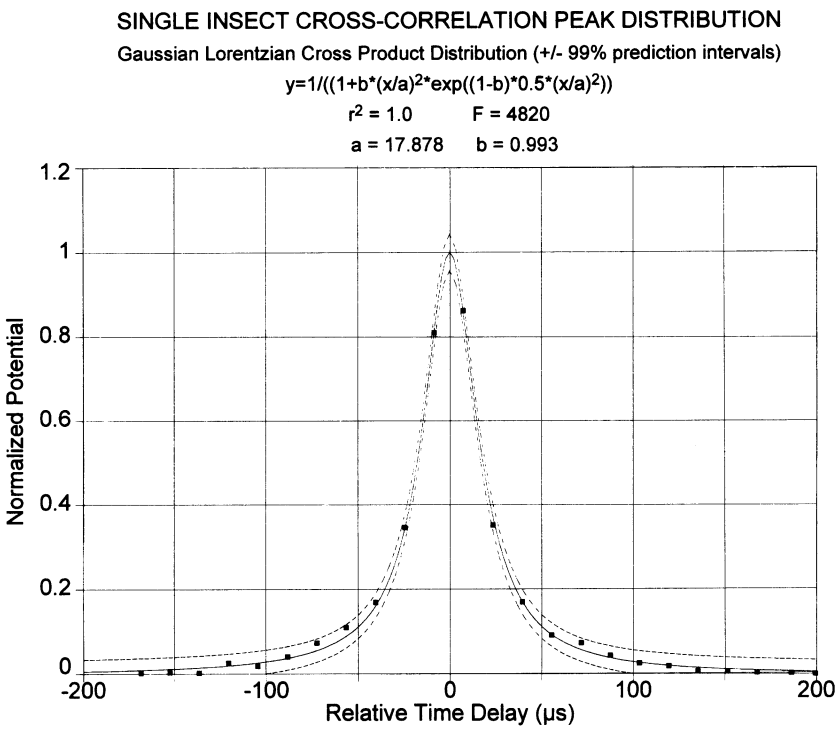

Fig. 9. Empirically derived normalized distribution of cross-correlation true peaks consolidated from tests with single insects. Values taken from the curve of best fit (shown with its 99% prediction intervals) taken at 16-μs intervals (nine values centered around $t = 0$) are used as the discrete potential function.

needed discrete potential function is then this curve evaluated at 16-μs intervals centered around its peak. To reduce computations, the potential function is truncated to encompass only nine discrete sample points since more than 90% of the correlogram true peak locations from single insects are within this range (a window that is eight sample periods or 128 μs wide).

A potential field contour (PFC) is constructed by replacing each peak in a CPD with this discrete potential function and summing the results at each discrete time-delay value (Fig. 10). Each valley of a PFC is checked for validity as a cluster boundary based on a set of rules involving the valley's ordinate value relative to the ordinate values of adjacent PFC peaks. If a resulting cluster domain (delineated by a pair of cluster boundaries) contains outliers and is therefore excessively wide, it could bridge with other clusters and erroneously link sounds emanating from different source locations. Such a wide cluster domain could result from noise or from a very large number of sounds produced by one insect. To prevent this, a maximum allowed cluster width is imposed to cull outliers. A window of this maximum width is placed on any wide cluster domains and moved along the Relative Time-delay axis of the PFC to encompass the maximum number of (peak) data points (Fig. 11). The width of this window can be adjusted to empirically optimize the system's performance, a trade-off between false positives and false negatives in detecting insects.[12] Finally, only clusters that contain some minimum

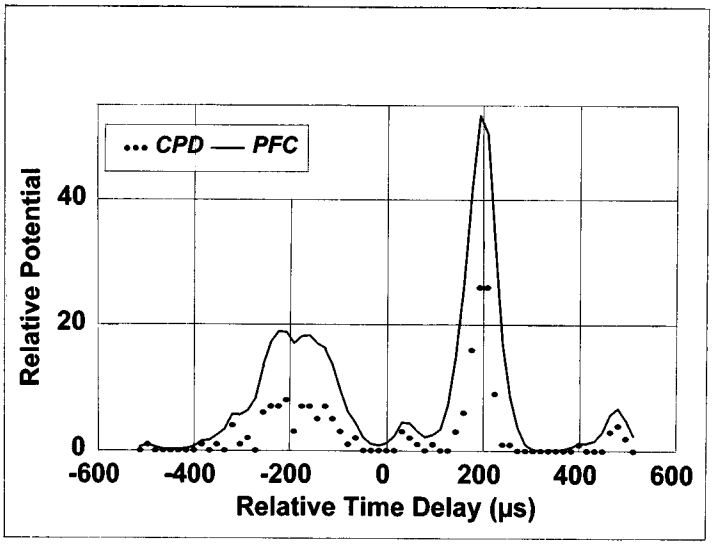

Fig. 10. A typical potential field contour (PFC), derived from the CPD shown in Fig. 8 and repeated here for clarity. The relative potential value at a given time-delay value is the sum of the weighted (by the discrete potential function) ordinate values of the nine CPD elements within ±four sample periods of that time-delay value.

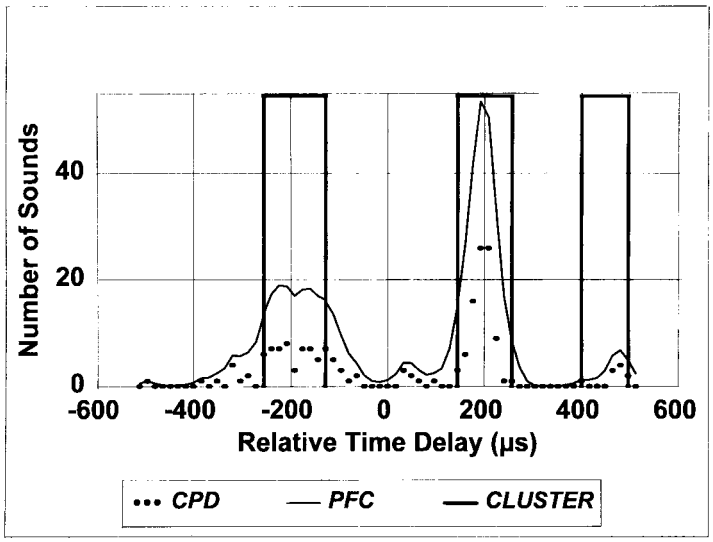

Fig. 11. Delineation of cluster domains from the PFC shown in Fig. 10 and repeated here with the CPD for clarity. The ordinate axis values relate to the CPD and not to the cluster domains that only show which CPD elements are grouped into clusters. The boundaries of the cluster centered about $-192\,\mu s$ were set using a maximum cluster-width window of $128\,\mu s$.

number of sounds (that are expected from each insect) are considered valid and therefore this number is proportional to the grain sample test duration.

Remembering that each acquired sound's fingerprint is a subset of an array of two peak time-delay locations in each cross-correlation dimension, a matching matrix is created to summarize quantified matches of fingerprints for every pairing of sounds (Fig. 12). After all the valid clusters in all the cross-correlation dimensions and their constituent sounds have been identified, each pair of sounds is assigned a fingerprint matching value equal to the number of clusters (a maximum of one per cross-correlation dimension) within which both are members. A pair of sounds that has a fingerprint matching value equal or greater than the (adjustable) matching value threshold[12] are considered to have emanated from the same source location and are therefore grouped together. This grouping process is continued until it reduces the matching matrix to a set of groups where each sound is present in only one group and a group can have any number of sounds. The final output of the ALFID algorithm is the number of individual larvae present in the sample as determined by the number of groups containing some minimum number of sounds. This minimum group size is adjustable[12] and its optimal value varies with test duration since an insect generally will produce more sounds over time.

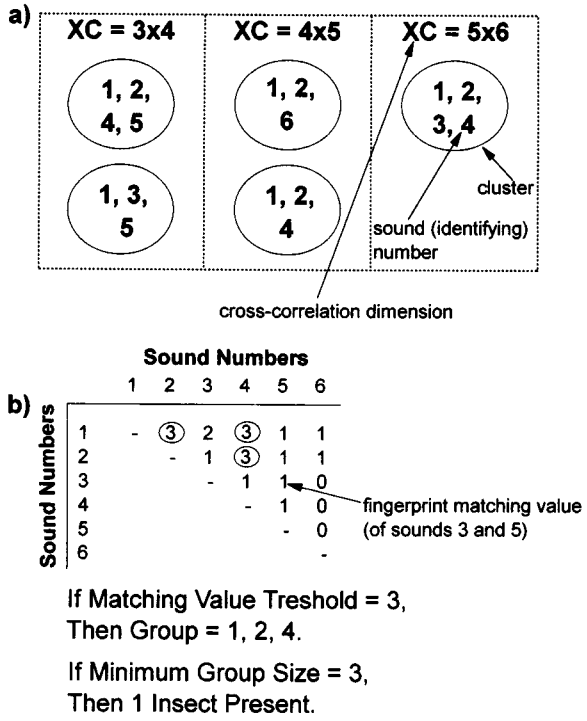

Fig. 12. Example of the fingerprint matching algorithm with six acquired sounds. In (a), clusters with their constituent sounds (sounds' identifying numbers encircled) are displayed for a sample of three (out of 36 considered) cross-correlation dimensions. In (b), the derived matching matrix shows three pairs of sounds with sufficiently large matching values for grouping. They form a single group of sounds that meets the minimum group size criterion, thus indicating the presence of an insect in the sample.

STATUS

The ALFID system described here has a number of parameters whose values were initially selected to provide an operational baseline and then were empirically adjusted to optimize performance.[12] Future research will focus on further tuning of the system, establishing performance abilities in the laboratory, and field testing it at a commercial grain elevator.

ACKNOWLEDGEMENTS

We thank the Grain Inspection Service, Packers, and Stockyard Administration, Kansas City, MO for their interest and support of this project. This research was supported in part by a specific cooperative agreement between the United States Department of Agriculture and the Agricultural and Biological Engineering Department at the University of Florida.

REFERENCES

1. Brain, C. K., Preliminary note on the adaptation of certain radio principles to insect investigation work. *Ann. Univ. Stellenbosch*, 1924, **2**(A), 45–47.
2. Adams, R. E., Wolfe, J. E., Milner, M. and Shellenberger, J. A., Aural detection of grain infested internally with insects. *Science*, 1953, **118**, 163–164.
3. Bailey, S. W. and McCabe, J. B., The detection of immature stages of insects within grains of wheat. *J. Stored Prod. Res.*, 1965, **1**, 201–202.
4. Street, W. M., Jr, A method for aural monitoring of in-kernel insect activity. *J. Ga. Entomol. Soc.*, 1971, **6**, 72–75.
5. Hagstrum, D. W., Webb, J. C. and Vick, K. W., Acoustical detection and estimation of *Rhyzopertha dominica* (F.) larval populations in stored wheat. *Fla. Entomol.*, 1988, **71**, 441–447.
6. Hagstrum, D. W., Vick, K. W. and Webb, J. C., Acoustical monitoring of *Rhizopertha dominica* (Coleoptera: Bostrichidae) populations in stored wheat. *J. Econ. Entomol.*, 1990, **83**, 625–628.
7. Hagstrum, D. W., Vick, K. W. and Flinn, P. W., Automated acoustical monitoring of *Tribolium castaneum* (Coleoptera: Tenebrionidae) populations in stored wheat. *J. Econ. Entomol.*, 1991, **84**, 1604–1608.
8. Vick, K. W., Webb, J. C., Weaver, B. A. and Litzkow, C., Sound detection of stored-product insects that feed inside kernels of grain. *J. Econ. Entomol.*, 1988, **81**, 1489–1493.
9. Hagstrum, D. W., Flinn, P. W. and Shuman, D., Automated monitoring using acoustical sensors for insects in farm-stored wheat. *J. Econ. Entomol.*, 1996, **89**, 211–217.
10. Hagstrum, D. W. and Flinn, P. W., Comparison of acoustical detection of several species of stored-grain beetles (Coleoptera: Curculionidae, Tenebrionidae, Bostrichidae, Cucujidae) over a range of temperatures. *J. Econ. Entomol.*, 1993, **86**, 1271–1278.
11. Shuman, D., Coffelt, J. A., Vick, K. W. and Mankin, R. W., Quantitative acoustical detection of larvae feeding inside kernels of grain. *J. Econ. Entomol.*, 1993, **86**, 933–938.
12. Weaver, D. K., Shuman, D. and Mankin, R. W., Optimizing and assessing the performance of an algorithm that cross-correlates acquired acoustic emissions from internally feeding larvae to count infested wheat kernels in grain samples. *Appl. Acoust.*, 1997, **50**, 297–308.
13. Helstrom, C. W., *Statistical Theory of Signal Detection.* Pergamon, New York, 1975.
14. Spiesberger, J. L. and Fristrup, K. M., Passive localization of calling animals and sensing of their acoustic environment using acoustic tomography. *Am. Naturalist*, 1990, **135**, 107–153.
15. Mankin, R. W., Sun, J. S., Shuman, D. and Weaver, D. K., Shielding against noise interfering with quantitation of insect infestations by acoustic detection systems in grain elevators. *Appl. Acoust.*, 1997, **50**, 309–323.
16. Massart, D. L. and Kaufman, L., *The Interpretation of Analytical Chemical Data by the Use of Cluster Analysis.* R. E. Krieger, Florida, 1989.