A Multiscale Transient Event Detector for Nonintrusive Load Monitoring

Steven B. Leeb James L. Kirtley, Jr.

Laboratory for Electromagnetic and Electronic Systems
Massachusetts Institute of Technology
Cambridge, MA 02139, USA

Abstract

The nonintrusive load monitor (NILM) determines the operating schedule of the major electrical loads in a building from measurements made at the utility service entry. This paper describes a multiscale transient event detection algorithm that extends the applicability of the NILM to challenging commercial and industrial sites. The performance of the algorithm is illustrated with results from a prototype event detector.

I. Background

Electric utilities and commercial facilities managers want to develop detailed electric power consumption profiles of their customers. Conventional metering of individual appliances is costly and inconvenient to the consumer. To deal with these concerns, utilities have sought a way of determining the operating history of an electrical load from measurements made solely at the utility service entry of a building. The residential nonintrusive load monitoring project undertaken in MIT's Laboratory for Electromagnetic and Electronic Systems has demonstrated the feasibility of a low cost, microprocessor-based recorder that competes favorably with conventional load monitoring schemes ([1], [2]). The advantages of nonintrusive load monitoring over conventional load monitoring include:

- ease of installation at the monitoring site, since the nonintrusive load monitor (NILM) requires a single set of electrical ties
- simplification of data collection, since the NILM automatically determines the electrical nature of the simple, "two-state" appliances in a target building without the need for a load survey or inspection
- facilitation of data analysis since, by definition, the NILM collects and potentially analyzes all data at a central location.

The NILM is more than a convenient and economical means of acquiring energy usage data. It is, for example, a potentially important platform for power quality monitoring. Many loads, such as computers and other office equipment, lighting fixtures, and adjustable speed drive systems, could draw distorted, nonsinusoidal input current waveforms ([3]). The NILM could track down power quality offenders by correlating the introduction of undesired harmonics with the operation of certain loads at a target site. The NILM could also serve as a platform for monitoring the performance of critical loads.

0-7803-0891-3/93$03.00 ©1993IEEE

The current implementation of the residential NILM tracks the operating schedules of the loads at a target site by looking for changes in steady state levels of real and reactive power. While informative in the residential setting, changes in steady state power levels are less revealing in commercial or industrial environments, where substantial efforts, e.g., power factor correction and load balancing, are made to homogenize the steady state behavior of different loads.

Fortunately, the transient behavior of many important load classes is sufficiently distinct to serve as a reliable indicator of load type. This paper describes a transient event detection algorithm, introduced in [4], that can be used to identify observed transient waveforms even when multiple transients overlap. This algorithm is suitable for incorporation into an advanced NILM which would be capable of monitoring demanding commercial and industrial sites. The performance of the algorithm is demonstrated with results taken from a prototype real-time event detector implemented with a digital signal processor.

II. Approach to Transient Recognition

The transient behavior of a typical load is intimately related to the physical task that the load performs. The load survey conducted in [4] indicates that nonlinearities in the constitutive relationships of the elements that comprise a load model, or in the state equations that describe a load, or both, tend to create interesting and repeatably observable turn-on transient profiles suitable for identifying specific load classes. The turn-on transients associated with a fluorescent lamp and an induction motor, for example, are distinct because the physical tasks of igniting an illuminating arc and accelerating a rotor are fundamentally different. Transient profiles tend not to be eliminated even in loads which employ active waveshaping or power factor correction.

However, direct examination of a current waveform or a closely related waveform like instantaneous power may fail to accentuate important features for pattern recognition. It is important to isolate key features from near constant frequency, "carrier wave" type signals like 120 Hertz instantaneous power so that slight errors in matching carrier frequency variation with a template do not dominate the results of a recognition system searching for a modulating envelope. An analog preprocessor in the prototype event detector eliminates carrier frequency artifacts from input data by averaging over at least one carrier frequency period to generate a short time estimate. The slow envelopes of the time average of instantaneous power, i.e., real power, and of reactive power, and of higher harmonic content, are found by mixing the observed current with appropriate sinusoids and then low-pass filtering. The event detector searches the slow envelopes or input data streams for known transient shapes as new data is acquired.

Searching for complete transients is an undesirable approach because it limits the tolerable rate of event generation. No two transient events could overlap significantly if each transient were to be identified correctly. Instead, the transient event detector searches for a time pattern of segments with significant variation, or *v-sections*, rather than searching for a transient shape in its entirety.

During a training phase, either before installation or on-site, the event detector employs a change-of-mean detector ([5]) to segment a transient representative of a class of loads. This segmentation process delineates a set of v-sections that will represent a particular transient shape in each of the input data streams. The trace in Fig. 1, for example, shows the measured envelope of real power during the turn-on transient of an instant start fluorescent lamp bank. Figure 2 shows the measured envelope of real power on one phase during the turn-on transient of a three phase induction motor. The locations of the v-sections in the two waveforms, as computed by a change-of-mean detector implemented in MATLAB, are approximately indicated by the ellipse in Fig. 1 and the rectangles in Fig. 2. In practice, a more complicated set which included v-sections in other data streams such as reactive power would be used to represent the transient profile of a load.

A complete transient identification is made by searching for a precise time pattern of v-sections. As long as each of the v-section shapes overlaps with no more than a near-constant region, the event detector will be able to identify the patterns of v-sections and therefore the transients. For example, the overlap of the two transients from the induction motor and the instant start lamp bank shown in Fig. 3 is tractable because all of

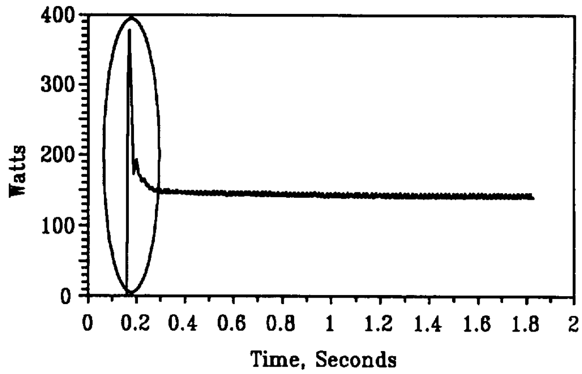

Figure 1: Measured Instant Start Lamp Bank Real Power Transient

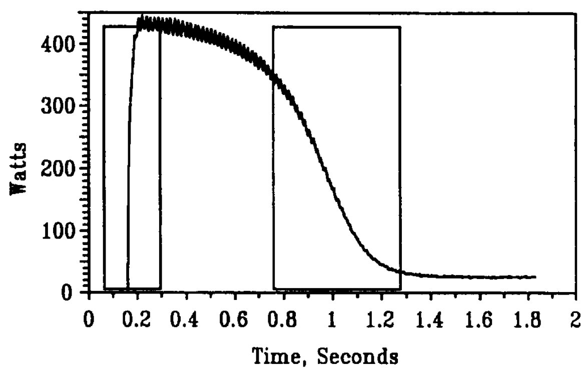

Figure 2: Measured Induction Motor Real Power Transient

the v-sections for both transients are distinguishable. The overlap condition in Fig. 4 would not generally be tractable, since the instant start v-section and the first induction motor v-section overlap severely. Since some degree of overlap is tolerable, the v-section set recognition technique will generally operate successfully in an environment with a higher rate of event generation than would a detector searching for whole, undisturbed transient shapes.

Classification of individual v-sections in the input data streams is determined by a set of pattern discriminator functions. These functions are used to compute a distance metric that locates a particular input vector in a region of a state space of known transient templates. Since a v-section may appear on top of a variably large static or quasi-static level created by the operation of other loads, the discrimination process focuses on only the "AC" or varying component of the v-section.

The prototype event detector employs a transversal or matched filter as a pattern discriminator, although other possibilities could be used and are discussed in [4]. The

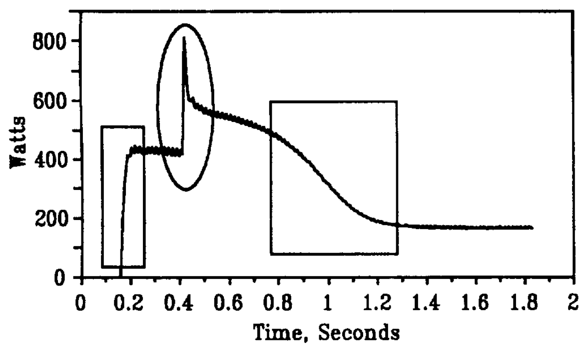

Figure 3: Acceptable Overlap

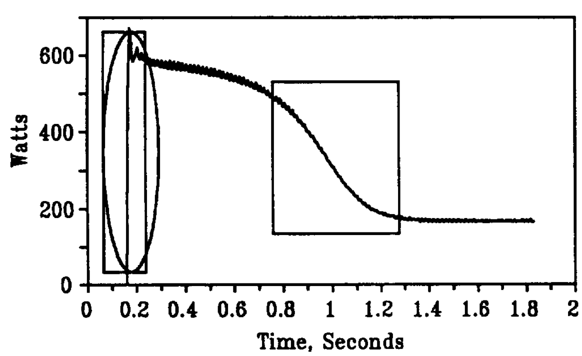

Figure 4: Intractable Overlap

impulse response of a transversal filter is proportional to the time-reversed signal for which the filter is designed to search. The transversal filter is an attractive signal processing construct for performing pattern discrimination because there are many efficient hardware solutions available for implementing a transversal filter. Each v-section is positively identified by checking the outputs of two different transversal filters.

The first transversal filter scans an input data stream for a particular shape. Let t denote a vector of N samples of a v-section of interest, such as one of the v-sections marked in Figs. 1 and 2. The vector t consists of elements

$$t[i], \quad i = 0 \ldots N-1.$$

An "AC coupled" and amplitude normalized version of the v-section, designated t_{ac}, can be computed as

$$t_{ac} = \frac{\left[t - \frac{1}{N}\sum_{i=0}^{N-1} t[i]\right]}{\| \left[t - \frac{1}{N}\sum_{i=0}^{N-1} t[i]\right] \|} \qquad (1)$$

where the denominator in Eq. 1 is the 2-norm of the numerator. Thus, t_{ac} is a unit length vector with zero mean. This shape transversal filter operates by sliding an N-point window across the input data stream. At any time, the window contains an AC-coupled and amplitude normalized vector of points x_{ac} of an N-point section or vector x of the input data. The vector x_{ac} is computed from the vector x using an equation structurally identical to Eq. 1. The output of the shape transversal filter is the inner product of the template vector t_{ac} and the data vector x_{ac}.

This output corresponds to the cosine of the angle between the two unit length vectors t_{ac} and x_{ac}. An output of unity indicates a perfect match between the template vector and the input data. Naturally, noise and slight variation in the repeatability of the v-sections will make a perfect match unlikely. In a practical system some degree of imperfection will be tolerated and any inner product within a certain tolerance of unity will constitute a match.

When a segment in the input data stream is found that matches the shape of a template v-section, a second transversal filter is employed to check the amplitude of the segment. This amplitude transversal filter computes the inner product of the template vector t_{ac} and the input data vector x, rather than the inner product of t_{ac} and x_{ac}. Checking the amplitude is essential in conjuction with checking shape to ensure that a small wiggle or noise pattern that is fortuitously close in shape to a v-section template is not mistaken for an actual v-section.

III. Searching for Patterns Over Many Time Scales

Loads in a particular class which spans a wide power range often exhibit transient profiles that are identical in shape but scaled in amplitude and duration. The transversal filter is suitable for identifying transient shapes over a narrowly defined time scale. The prototype event detector employs a tree-structured decomposition to search efficiently over many time scales with the transversal filters. The use of the tree-structured decomposition is inspired by recent applications of subband coding ([6]) and the discrete-time wavelet transform ([7], [8]).

A tree-structured signal decomposition is shown schematically in Fig. 5. An atomic section, or coder, of the tree shown would consist of the upsampler with upsampling rate l, operations $\mathbf{R}$ and $\mathbf{D}$, and the trailing downsampler with downsampling rate m. Each upper path, or *resolving path*, in any particular coder in the tree alters the resolution of the input signal with operation $\mathbf{R}$, and also alters the scale of the input signal with the up/downsampling operations at rates l and m. The lower *discriminating* path implements operation $\mathbf{D}$ which identifies patterns at a particular time scale in the input data. In the prototype event detector, a transversal filter computes the discriminating coefficients, U_s, for example, that indicate the possible presence of a v-section.

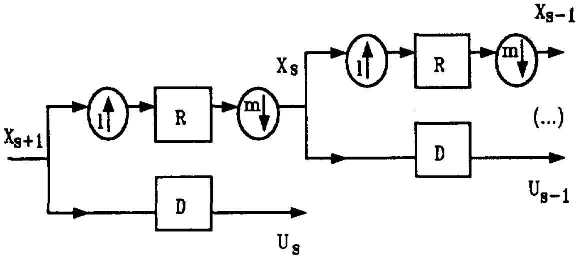

Figure 5: Tree-Structured Signal Decomposition

Experience gathered from the load survey in [4] and from experiments with the event detector indicate that operation **R** may be implemented with a linear filter for the v-sections associated with many commercial and industrial loads. Constraints for designing filters that minimize the impact of distortions in a pattern recognition setting, and conditions when nonlinear filters might be desirable, are presented in [4]. An adaptive algorithm for minimizing the effects of amplitude distortion and shift-variance is also presented in [4].

IV. Event Detection Algorithm

A flow chart of the complete event detection algorithm is shown in Fig. 6. This algorithm is implemented in the *NILMscope* software that is executed by the prototype event detector. The *NILMscope* software is written in two components. One component performs the event detection algorithm and is executed on a TMS320C30 DSP system. The other is a user interface responsible for initializing the DSP system and for report generation. This second component is executed on an 80486-based personal computer. Each numbered item in the following list is a description that details a numbered step of the algorithm in the flow chart.

1. Data Acquisition: During this data acquisition step, the DSP system collects a window of samples which will be searched for known transient patterns.
2. Tree-Structured Decomposition: Once a full window or vector of samples has been acquired, the DSP system performs a tree-structured decomposition on the data as described in the previous section. In the current implementation, the input data for each coder step in the tree is computed before any pattern discrimination occurs at any scale step. A tree structure with a total of three 2 to 1 coder or scale steps proved sufficient for identifying all of the transients associated with the loads in the test stand described in the next section.
3. Set Scale Steps: Next, the DSP system searches at each scale for all of the transient types that could appear. There are three scale steps, and a scale step variable M is initialized to the value of 3 to count down the search. First the finest sampled scale step is inspected, followed by the middle and coarse scales.

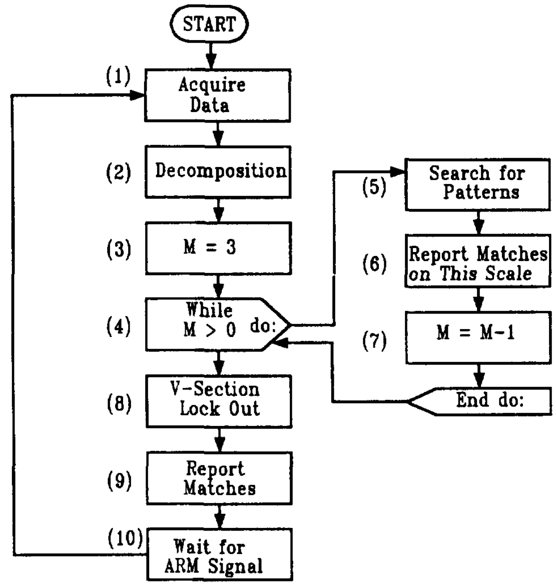

Figure 6: The *NILMscope* Event Detection Algorithm

4. Initiate Pattern Search: A loop in the program flow at this point in the *NILMscope* software controls the search for patterns over all three scale steps.
5. Hierarchical Pattern Search with V-Section Lock Out on Scale M: During each pass through the pattern search loop, the DSP system searches for the v-sections associated with the known transient events on a single scale. The pattern search is hierarchical, in that the DSP system searches first for the patterns with the most v-sections. When all of the v-sections for a pattern are found with both the shape and amplitude transversal filtering operations, the complete transient pattern is presumed to be present in the input data, and an event is recorded.

   A v-section lock out is performed at each scale. If a complex pattern is found in the input data, the location of the v-sections of the pattern are recorded. The identification of any subsequent, less complex patterns will not be permitted based on the detection of v-sections at the previously recorded, "locked out" locations.
6. Report Generation for Scale M: If all of the v-sections are found for a particular pattern, the transient pattern is presumed to be present in the data at the current scale M, and an event type and time is communicated to the interface component of the *NILMscope* software running on the PC.
7. Decrement Scale Step: The scale counter M is decremented, and the pattern detection loop is repeated until all remaining coarser scales have been searched.
8. V-Section Lock Out Over All Scales: The PC component of the *NILMscope* software performs a final check to ensure that v-sections from a complex but coarse scale pattern were not used to match a less complicated, finer scale pattern.
9. Final Report Generation: The PC component of the *NILMscope* software collates a final report of the type and time of occurrence of all positive event detections.
10. Standby: After reporting any contacts, the PC waits for the user to issue an arming command.

V. Prototype Testing and Performance

The prototype event detector consists of three components: an analog preprocessor, a DSP card, and a personal computer. The event detector monitors the voltage and current waveforms on a three phase electrical service that powers a collection of loads representative of important load classes in typical, medium to large size commercial and industrial buildings. The prototype event detector is used to identify the turn–on time and type of the various loads. The event detector has, of course, no *a priori* knowledge of the operating schedule for the loads.

Four loads were selected for inclusion in the test facility: a bank of instant start fluorescent lamps, a bank of rapid start fluorescent lamps, a $\frac{1}{4}$ horsepower induction motor, and a $\frac{1}{3}$ horsepower induction motor. The electrical hookup to the loads is routed through an electronically switched circuit breaker panel that activates loads with flexibility in relative timing. The pattern templates for the loads were captured during a one-time "walk through" of the test stand. However, *no data at all was collected from the large motor.* Because the large and small induction motors are members of the same load class, a single transient template, appropriately scaled in amplitude and duration, was expected to prove satsifactory for identifying both motors.

Figures 7 through 11 show screen prints from the PC running the *NILMscope* user interface software during five of the experiments conducted with the test stand. The graph windows in each figure show estimates of the envelopes of real and reactive power on one phase of the three phase service. For example, the windows in Fig. 7 display the data collected during the turn–on of the instant start lamp bank in the test stand.

In the lower left-hand corner of each screen, any transient events that the event detector has been able to identify appear under the heading `Contacts`. Events or contacts are reported by identity, time of occurrence, and scale. There are three scales listed in the contact window: fine, mid, and coarse. Any event that was identified by a known v-section set at the initial, highest sampling rate, i.e., at the first coder stage in the tree-structured decomposition, will be listed directly under the heading `Fine Scale` in the contact window. By design, it is anticipated that events associated with the small motor and both lamp banks will be listed as fine scale events when they appear, assuming that the event detector functions properly. Events found at the next two coder stages in the tree-structured decomposition will be listed under the headings `Mid Scale` or `Coarse Scale`. For example, any events recognized by the properly working detector which are caused by the turn–on of the large induction motor should appear under the coarse heading.

The three tests shown in Figs. 7 through 9 record the performance of the prototype when challenged individually with the turn–on events of the instant and rapid start lamp banks and the small induction motor, respectively. In each case, the observed event has been properly identified in the contact window. Figure 10 shows an example where both lamp banks and the small motor turn on so that all three transient events overlap. No key v-sections overlap with each other. All three events are correctly recorded at the finest time scale in the contact window.

In the final experiment shown in Fig. 11, the small induction motor turns on and completes its transient, followed by the turn–on transient of the large induction motor and the instant start lamp bank. No key v-sections overlap with each other. All of the events are correctly identified at the appropriate time scales. Recall that the template for the turn–on transient of event type `Motor` was generated from a single example of the small motor *only*. Nevertheless, the event detector correctly classified both the small and large motors.

VI. Conclusions

The examples reviewed in the previous section are representative of several hundred experiments conducted with the test stand. Provided the assumptions made in the development of the event detection algorithm are satisfied, the prototype detector performs remarkably well. This performance is perhaps more impressive in light of the fact that fairly little effort was made to "tune" the detector for the loads in the test stand.

In practice, the robustness of the detector could be enhanced by working with more complicated v-section pattern sets for each load. Employing more complicated sets of interesting v-sections will enhance the reliability of the event detection algorithm but will also require more computational capability. Fortunately, the search for different v-sections could be conducted in parallel. This suggests that a commercial NILM based on the event detection algorithm could be constructed around a parallel processing architecture composed of inexpensive microprocessors or microcontrollers.

Acknowledgements

This research was funded by the Electric Power Research Institute under EPRI project number RP2568-15. The authors gratefully acknowledge the valuable advice, support, and encouragement of Laurence Carmichael, Dr. Richard Tabors, Professors George Verghese, Jacob White, and Les Norford, and Kurt Levens.

References

[1] G.W. Hart, "Nonintrusive Appliance Load Monitoring," *Proceedings of the IEEE*, Vol. 80, No. 12, December 1992, pp. 1870–1891.

[2] R. Tabors, et. al., "Residential Nonintrusive Appliance Load Monitor," EPRI final report, to appear.

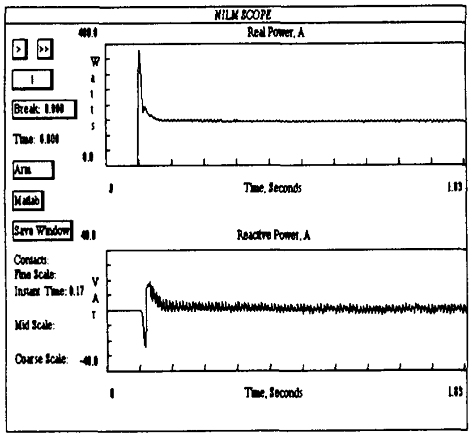

Figure 7: *NILMscope* Contact Report

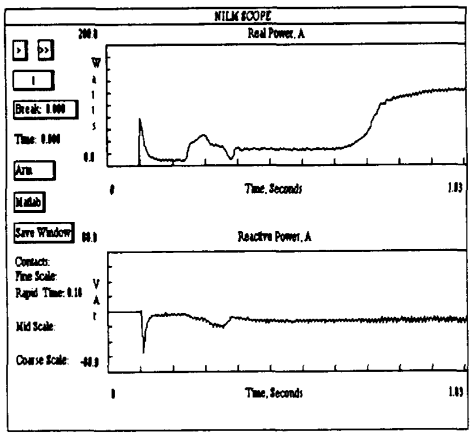

Figure 8: *NILMscope* Contact Report

[3] B. Wilkenson, "Power Factor Correction and IEC 555-2," *Powertechnics Magazine*, February 1991, pp. 20–24.

[4] S.B. Leeb, "A Conjoint Pattern Recognition Approach to Nonintrusive Load Monitoring," M.I.T. Ph.D. Dissertation, Department of Electrical Engineering and Computer Science, February 1993.

[5] M. Basseville and A. Benveniste, *Detection of Abrupt Changes in Signals and Dynamical Systems*, Springer–Verlag, 1980.

[6] P.P. Vaidyanathan, "Multirate Digital Filters, Filter Banks, Polyphase Networks, and Applications: A Tutorial," *Proceedings of the IEEE*, Vol. 78, No. 1, January 1990, pp. 56–93.

[7] S.G. Mallat, "A Theory for Multiresolution Signal Decomposition: The Wavelet Representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 11, No. 7, July 1989.

[8] Y. Meyer, "Orthonormal Wavelets," *Wavelets: Time–Frequency Methods and Phase Space*, Proceedings of the International Conference, Marseille, France, December 14–18, 1987, Springer–Verlag, pp. 21–37.

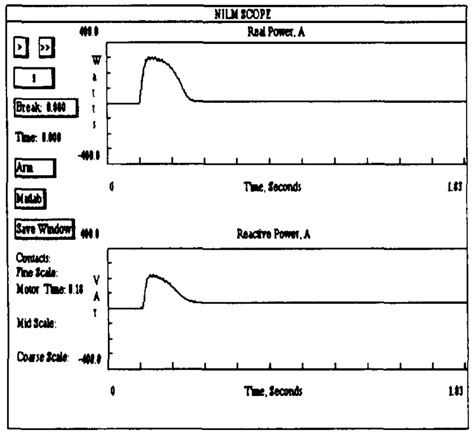

Figure 9: *NILMscope* Contact Report

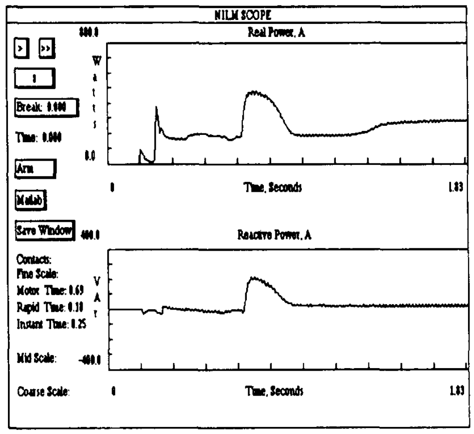

Figure 10: *NILMscope* Contact Report

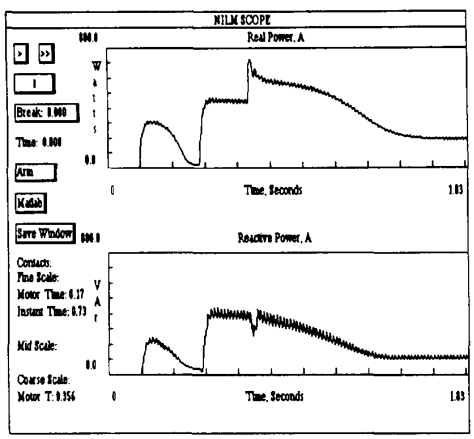

Figure 11: *NILMscope* Contact Report