Fault Detection Using a Two-Model Test for Changes in the Parameters of an Autoregressive Time Series

P. Scholtz and P. Smyth
Communications Systems Research

This article describes an investigation of a statistical hypothesis testing method for detecting changes in the characteristics of an observed time series. The work is motivated by the need for practical automated methods for on-line monitoring of DSN equipment to detect failures and changes in behavior. In particular, on-line monitoring of the motor current in a DSN 34-m beam waveguide (BWG) antenna is used as an example. The algorithm is based on a measure of the information theoretic distance between two autoregressive models: one estimated with data from a dynamic reference window and one estimated with data from a sliding reference window. The Hinkley cumulative sum stopping rule is utilized to detect a change in the mean of this distance measure, corresponding to the detection of a change in the underlying process. The basic theory behind this two-model test is presented, and the problem of practical implementation is addressed, examining windowing methods, model estimation, and detection parameter assignment. Results from the five fault-transition simulations are presented to show the possible limitations of the detection method, and suggestions for future implementation are given.

I. Introduction

The motivation for this study is on-line performance monitoring of DSN electromechanical and hydraulic equipment, particularly the pointing system components of the DSN 70-m and 34-m antennas. Previous articles [1-3] have described in detail the motivation behind on-line monitoring: Essentially as the antennas get older and deep space missions are of longer duration, early detection of component failures becomes more critical.

Simple thresholding methods, whereby detection of change occurs when the magnitude of the observed sig-nal exceeds prespecified alarm limits are available in commercial off-the-shelf products and widely used for on-line monitoring. However, the simple thresholding approach is nonadaptive and may be susceptible to false alarms in the presence of noise. In this article more sophisticated methods, which not only account for changes in the signal amplitude, but can also detect changes in the underlying statistical characteristics of the signal in cases where no amplitude change is observable, are investigated.

Various methods which detect changes in the mean of a signal directly by utilizing statistical cumulative sum

(cusum) schemes have been thoroughly examined in [4–7]. These methods generally work well when there is sufficient prior knowledge about the magnitude of change. Mean-change detection algorithms have been well-defined for observations which consist of independent, identically distributed Gaussian random variables; however, when there is a significant time correlation in the observations (as is almost always the case in applications of interest), the usefulness of these methods is diminished, and other techniques must be utilized.

One promising technique for detecting parametric changes in the model of the process is well documented in [8,9]. This method computes a cumulative sum based on the information-theoretic distance between two estimated autoregressive (AR) models of the process. The focus of this article is to summarize the theory underlying this cusum algorithm and to examine the problems and necessary adaptations for practical implementation in a DSN application.

In choosing the two-model algorithm for implementation, a test was desired with the following properties:

(1) Few (not necessarily minimal) false alarms.

(2) Robustness with respect to variance in model estimates.

(3) Few a priori assignments of detection parameters and the use of on-line estimated parameters.

(4) Symmetry with respect to transitions from both low-to-high and high-to-low signal variance, and vice versa, and with respect to changes in the AR parameters.

The two-model algorithm described in this article utilizes two filters, a growing memory AR model and a local memory AR model, which are implemented by using two data windows: a reference window of growing length and a fixed-length window. A statistic based on conditional Kullback information measures the distance between the two models based on the innovations from both filters. The crossing of a threshold by the cumulative sum of this statistic detects the distribution change. An advantage of this technique is that it will detect a change in the actual AR parameters of the process or a change in the energy of the signal. Section II of this article presents the AR model change hypothesis. Section III focuses on the theoretical derivation of the cumulative sum, while Section IV examines the problems for implementing the algorithm for DSN applications. Results of simulated fault detection are discussed in Section V. Limitations of the method are examined in Section VI, and Section VII compares the method with an alternate approach using hidden Markov models.

II. Autoregressive Model Change Hypothesis

Suppose the observed signal (Y_n) can be modeled by an autoregressive process of order p, that is

$$Y_n = \theta_p \overline{Y}_{n-p}^{n-1} + \epsilon_n \tag{1}$$

where $\overline{Y}_{n-p}^{n-1}$ is the column vector of the past p values of the observation signal, θ_p is the AR(p) parameter vector, and the first term on the right is a vector dot product. It is well established in time series analysis that any stationary signal (or any piecewise stationary signal) can be modeled as an AR model of sufficient (finite) order, plus a deterministic term. To test for change, assume there exists a time r such that, for $n \leq r$

$$\theta_p = \theta_p^0 \quad \text{and} \quad \sigma_{\epsilon_n}^2 = \sigma_{\epsilon_n^0}^2 \tag{2}$$

and for $n > r$

$$\theta_p = \theta_p^1 \quad \text{and} \quad \sigma_{\epsilon_n}^2 = \sigma_{\epsilon_n^1}^2 \tag{3}$$

where $\sigma_{\epsilon_n}^2$ is the variance of the prediction error ϵ_n at time n. These changes in AR parameters reflect an underlying change in the probability laws governing the process. The problem is to test sequentially at each time n, the null hypothesis H_n^0 of no change in the probability law of the process against the alternate hypothesis H_n^1 that the above change in probability laws occurs at a time $r < n$.

III. Cumulative Sum Detection Statistic

The most obvious test for difference between two models is the likelihood ratio test. This tests the null hypothesis H_n^0 that all observations up to time n follow the *joint* probability law $p^0\left(Y_n, \overline{Y}^{n-1}\right)$ against the hypothesis that all observations up to time n follow $p^1\left(Y_n, \overline{Y}^{n-1}\right)$ [8]. However, the sequential test desired for application is the *conditional* hypothesis test, which tests the null hypothesis H_n^0 against the alternate hypothesis H_n^1 that the probability distribution changes from p^0 to p^1 at time $r < n$ (as stated in the introduction). Hence, letting y denote the value of the variable Y_k, the detection test is the sequential cumulative sum test:

undergoes a decrease from an initial mean of zero, reflecting the increasing difference in the two models after the change in distributions; hence a one-sided Hinkley test is used. Detection occurs at the first time when $D > h$ for a preset threshold h. For recursive computation on line, the value

$$U_n = U_{n-1} + T_n - \delta \qquad (11)$$

can be used [9]. Implementation of this scheme will be discussed in the next section.

IV. Implementation of the Two-Model Method

A. Model Order

Prior work modeling the elevation pointing system motor current from the DSS-13 34-m BWG antenna led to an autoregressive-exogenous input (ARX)(5,3,1) model as a first approximation to the proper model [1]; this model is defined by 3 poles, 5 zeros, and 1 delay in the transfer function. The order of this model was determined using Akaike's information criterion. Further work has shown that comparing the model parameters of the ARX(5,3,1) to AR(5) yields no significant difference in the estimated autoregressive coefficients. Hence, although the model underlying the process has yet to be properly identified, it appears that an AR(5) model may be sufficient for change detection purposes. However, it should be noted that since the detector relies directly on the goodness-of-fit of the model, the opportunity exists to potentially improve the model identification method in order to reduce false alarms due to incorrectly estimated prediction errors.

B. Window Implementation

There are two well-documented approaches to block window implementation of the two-model detection algorithm. The first relies on a fixed-length reference window which is allowed initially to stabilize in a controlled period of normal operation in order to estimate the AR(p) reference model. This model is then compared with the estimated model in a moving fixed-length window, with detection when the distance between these two models becomes sufficiently large. In the past, the difference between these two models has been measured by the mean quadratic difference between the two spectra [14]. Basseville and Benveniste [8] report that the disadvantages of this approach are a large variance in the metric and an asymmetrical test for increases or decreases in signal noise.

A more robust, dynamical window system employs a growing-memory window for reference along with the fixed-length sliding window. First a growing reference window is allowed to attain a stable model. As soon as the window stabilizes, a shorter fixed-length sliding window begins to move along the time series with the reference window. At each time n, a model is estimated in each data window. When a transition occurs in the spectrum of the signal, the abrupt change is reflected in the local window, while the reference model remains relatively unchanged due to its long memory. The information metric between these two windows is measured by the statistic T of Eq. (9), which is integrated in the cumulative sum U of Eq. (8); then the Hinkley decision rule given by Eq. (10) is invoked. The use of this growing reference window instead of a static reference window greatly reduces the rate of false alarms by adapting to the dynamic nature of the system [8].

C. Model Estimation

The first attempt at implementing the algorithm employed a dynamic *block* window scheme for model estimation. Sequentially, at each time n an AR(5) model was estimated for both the growing window and the sliding window. Treating each window as a batch, each model was estimated using the "forward–backward" AR estimation algorithm and then the cusum was computed based on the model errors. As expected, choosing an apppropriate window size for the fixed-length window was critical in the implementation of the algorithm. An undersized window leads to unstable model estimation, creating largely varying prediction errors which lead to false alarms. On the other hand, oversized local windows lead to longer delays until detection and more computation. A window size of 200 to 400 data points (approximately 4 to 8 seconds at the sampling rate of 50 Hz) yielded a reasonable fit. Below a window size of 200, the AR model was highly unstable and produced unreliable results.

From a practical standpoint, this nonrecursive algorithm was computationally intensive and infeasible to implement on-line.

A second attempt utilized a recursive algorithm for model estimation. The Normalized Gradient approach of the recursive least mean squares (LMS) parameter estimation was used to fit the AR(5) model [15]. For the linear regression

$$\hat{Y}_n = \overline{Y}^{n-1} \hat{\theta} \qquad (12)$$

the parameter estimate $\hat{\theta}_n$ is

$$\hat{\theta}_n = \hat{\theta}_{n-1} + \gamma \overline{Y}^{n-1} \epsilon_n \tag{13}$$

where γ is the (constant) gain. For the normalized approach, γ is replaced by

$$\gamma' = \frac{\gamma}{|\overline{Y}^{n-1}|^2} \tag{14}$$

This is the same method employed by Eggers and Khuon [9] with their recursive LMS learning-parameter method, which they showed will converge to give the true AR coefficients. The dynamic windowing method described above is implemented by choosing the gains γ_1, γ_2 corresponding to the reference window and the fixed-size window, respectively, such that the model is weighted to reflect the information content of the most recent observations. Hence, the gains are chosen such that $0 < \gamma_1 < \gamma_2$, reflecting the abrupt change in spectral characteristics in the local model while leaving the reference model relatively unchanged [9]. The gain is usually chosen in the range 0.001 to 0.02 [15]. In the work reported here, γ_1 was chosen to be 0.001, while γ_2 was set to 0.02, for maximum distinction between the two models. The gains must be bounded by $1/tr(\overline{R}) = 1/pr_0$, where $\overline{R}$ is the autocorrelation matrix of the last $p - 1$ observations, p is the AR(p) model order, and r_0 is the autocorrelation sequence at lag 0. For example, one can obtain an estimate of this bound on-line by estimating the autocorrelation sequence r_0 after n samples by

$$\hat{r}_0^{YY} = \frac{1}{n} \sum_{k=1}^{n} |Y_k|^2 \tag{15}$$

as shown in Marple [16]. However, in the work presented here, γ_1 and γ_2 were fixed as described above for simplicity. Further investigation is required to determine a reasonable way to preset the gains for this approach.

Other recursive methods were considered, such as the recursive AR algorithm utilizing a Kalman filter scheme. However, these methods require more prior information, such as a knowledgeable guess of the covariance of the innovations. The gradient approach utilized here, as shown in Eqs. (12) through (14), is a logical choice for the case where there is little a priori information. A possible alternative estimation method is the fast recursive least-squares algorithm of Ljung, as described by Marple [16]. This algorithm employs a time-varying gain for the normalized

gradient approach resulting in lower MSE than the recursive LMS algorithm, with approximately the same computational efficiency.

D. Choice of the A Priori Detection Parameters

In on-line implementation, it was difficult to choose appropriately the a priori parameters h (threshold) and δ (drift) for Eqs. (8), (9), and (10). The most difficult detection parameter to assign is the drift δ of T_n. In cases of small changes in the signal energy or AR parameters, the cusum is particularly sensitive to the choice of δ. Hinkley [4] determined the value of δ to be

$$\delta = \frac{1}{2}(\mu_1^* - \hat{\mu}_0) \tag{16}$$

where $\hat{\mu}_0$ is the estimated initial mean, and μ_1^* is the minimum expected final mean of the process. In order to assign δ, there must be prior knowledge of the expected behavior of T_n, and therefore prior knowledge of the expected faults.

It should be noted that the cumulative sum can be run with an assigned drift bias equal to zero. Instead of using the Hinkley stopping rule of Eq. (10) to detect a large deviation from the maximum of the cusum, the decision to stop would occur when the cusum U given by Eq. (8) passes a set threshold value h'. However, this merely further complicates the problem of setting the decision threshold.

When the Hinkley stopping rule of Eq. (10) is utilized, the threshold h must be assigned a priori. In this experiment, a constant value for h was chosen, specific to the fault transition under examination. However, for practical implementation when multiple unknown changes can occur, a constant threshold does not appear feasible. Certainly a dynamic threshold h based on the variance of T seems appropriate. In fact, in another application Basseville and Benveniste [17] suggest a threshold

$$h = c\frac{\hat{\sigma}_n^2}{\delta} \tag{17}$$

where $c > 0$ is a constant, $\hat{\sigma}_n^2$ is the variance of the desired random variable (in this case T_n) estimated on-line, and δ is assigned the value of one-half the minimum expected magnitude of change in the mean. The usefulness of this threshold has not yet been examined in this context.

E. Signal Energy Change Considerations

As can be seen in the next section, low- to high-energy transitions in the signal are easily detected with the use of the detection statistics of Eqs. (8), (9), and (10). However, when the signal energy drops at the transition, detection is more difficult. By running a second test in parallel this obstacle is overcome. This alternate test requires switching the innovations and variances of model 0 in Eq. (9) with those of model 1 [9]. The drift in T_n is more delineated in this alternate test, allowing for better behavior of the cusum, and thus better detection.

V. Results

A. Simulated Fault Data

As mentioned previously, the motivation in examining various cumulative sum tests was to find a failure detection test which was feasible for on-line implementation. In a previous experiment [1], on-line readings from various sensors on the DSS-13 34-m BWG elevation pointing system were collected, both under normal conditions and under five hardware faults which were introduced in a controlled manner. The proposed change detection method (as described above) was tested on this data set. The introduced hardware faults were

(1) Noise in the tachometer feedback path.

(2) Tachometer failure.

(3) Torque share/bias loss to motor.

(4) Integrator short circuit.

(5) Rate loop compensation short circuit.

The simulation of these failures is described in greater detail in [1]. Of the 12 sensors on the control system, the most easily modeled with a time series was the motor current, and this was the sensor data selected for this study. In the past, a pattern recognition approach utilizing a hidden Markov model had been used to attempt detection of these five simulated fault transitions, with very accurate results [3]. However, this method requires training data for each fault a priori. In contrast, the purpose of the study reported here was to investigate alternative techniques which do not necessarily require training data for a set of faults which are known a priori. Nonetheless, as pointed out earlier, the two-model approach investigated here does in fact require some prior knowledge of the fault characteristics in terms of setting the drift parameters and choosing the order of the model.

B. Data Preprocessing

The raw data from the motor current sensor had a significant amount of sensor noise and outliers due to sensor faults. First, linear detrending was performed over the entire sample to remove low-level linear trends. Then the raw data was bandpass filtered with a tenth-order Butterworth filter to remove any further outliers. The passband of the filter was from 0.5 to 10 Hz for effective smoothing.

C. Fault Transition Detection

In the current detection experiment, transitions from normal to fault conditions were examined for all five faults. Data were compiled with normal conditions for the first 1000 samples, and with the desired fault conditions for the remaining 1000 data samples. At a 50-Hz sampling rate, this represents only approximately 2/3 of one minute of real data since the data records used were shortened for computational reasons. Transitions to faults 1 and 2 increased the signal energy, transitions to faults 3 and 4 preserved the signal energy, and transitions to fault 5 decreased the signal energy. Also, fault 1 conditions were scaled to equal energy with normal conditions to ascertain whether the test could detect a change in AR parameters without a signal energy change.

A summary of the important detection parameters is presented in Table 1. For this table, the AR parameters of each process were estimated using a block window AR model algorithm with window size 200. Figure 1 shows a typical fault transition statistic T and cusum U for a low to high energy transition. Note the decrease in mean of T at the time of jump, sample number 1001, and the drop in cusum U from its maximum near transition. Figures 2 through 7 present the fault transitions examined and the behavior of the cusum U prior to detection. The filtered signal shown in the top view of these figures is proportional to motor current. The cusum U (dimensionless) shown is equal to zero for the first 200 points, corresponding to the initial model stabilization period, and then varies with time until detection, when it is reinitialized to a value of zero.

1. Faults 1 and 2. Faults 1 and 2 are the simplest to detect; they have both large increases in signal energy and distinct changes in the AR parameters θ. Figure 2 shows the behavior of the filtered signal and cusum U for a transition to fault 1. Detection occurs at sample 1093, a delay of 93 sample points (or about 2 seconds), which should be acceptable for DSN operational requirements. Similarly, the transition to fault 2 conditions is detected with a delay of only 11 points (0.3 seconds) (Fig. 3).

2. Faults 3 and 4. Faults 3 and 4, the bias loss and integrator short circuit, show little deviation from nominal conditions in the AR parameters or signal energy, as seen in Table 1. Although components were physically removed or altered, no significant change in sensor data or pointing performance was observed due to the redundancy in the system. Figures 4 and 5 show the signal and appropriate cusum for these transition detectors. Recall that the fault transition occurs at sample point 1001; clearly there is no visible change in the signal. As might be expected, the transition was undetected for both faults.

3. Fault 5. The compensation short circuit, fault 5, is a transition from high to low signal energy, with changing AR parameters. For this type of transition, the alternate T_n statistic described in Section IV.E is used, which increases the model distance measure for better detection. Figure 6 shows the signal under consideration and the alternate cusum U'. In this case, detection is possible and occurs almost instantaneously at transition, with a delay of only 8 sample points. However small this delay may be, it should be noted that in this case detection is highly sensitive to the choice of the threshold h; increasing the threshold by a small amount may lead to a missed detection, while decreasing the threshold by a small amount may lead to multiple false alarms. Thus, it is clear that the algorithm does detect faults that cause decreases in signal energy, but with some limitations.

4. Scaled Fault 1 Case. A simulation was run using a transition to fault 1 conditions which were scaled by 0.25 to match the signal energy of the nominal data. This test was performed to determine whether a change in AR parameters *without* a change in signal noise could be detected. Figure 7 shows the signal transitions and cusum behavior for this case; detection occurred at sample 1151, a reasonable delay of 150 points (3 seconds). Again there is some doubt as to exactly how small a transition can be detected, but at least the test works with the appropriate choice of parameters.

VI. Limitations of the Approach

A. Prior Knowledge Requirements

Although not conclusive, the results of the five fault detection simulations point out some limitations of the two-model method. First, some knowledge of the manner in which faults will affect the AR parameters describing the data is required. Since this detector relies on the prediction errors of the AR model, prior knowledge of the

model order for both normal and fault conditions is required. Moreover, this method is not designed to detect changes in the model order, which may occur for fault transitions. Such order changes could be detected by a more sophisticated algorithm which dynamically tries to fit multiple-order models to the data—however, this would be both computationally intensive and potentially difficult to stabilize.

B. Parameter Assignment

A more detailed examination of the preassigned parameters h (threshold) and δ (drift bias), which are fault-specific parameters, would have to be conducted to determine a systematic method of assignment. A global time-varying threshold for all the parallel tests based on the variance of T_n, as given by Eq. (17), is the logical candidate for improving the parameter h. The drift bias δ, corresponding to a change in the mean of the statistic T, can be difficult to assign. Furthermore, the detection test is highly sensitive to the choice of δ, so proper assignment is critical. However, with a larger record of available data, a small number of trials for each expected fault type is expected to yield feasible values of δ for on-line implementation.

C. False-Alarm Rate

The false-alarm rate has not been derived analytically for the recursive AR estimation approach for the two-model approach. For the block window implementation as presented in Section IV.C, the rate of false alarms was derived by Basseville in [8] and is equal to the inverse of the expected value of the detection time D_h

$$E(D_h) = \frac{1}{\delta}\left[\frac{2}{\delta N}(e^{2Nh} - 1) - h\right] \qquad (18)$$

where N is the size of the fixed-length block window and δ is chosen as a function of the distance of the two probability laws describing the process before and after the change. However, the exact relationship between the block window model estimation and the recursive normalized-gradient model estimation is not yet clear, so an analytic expression for the false-alarm rate has not yet been derived. In any case, derivation of false-alarm rates for these types of models requires a complete model of the fault to be detected and, thus, is of questionable utility for practical purposes.

VII. Comparison to the Hidden Markov model method

In [2,3] a hidden Markov model (HMM) method for fault detection was reported. The HMM method assumes that the set of faults is known in advance and training data are required for each fault. However, once trained the model is quite robust and does not require any additional parameters to be set or calibrated. On the other hand, the two-model method described here does not specifically require training data in advance; rather, some prior knowledge about the possible faults is required. It would appear that if training data are available the HMM method is more robust and accurate as a detector than the two-model hypothesis testing approach. However, since training data for specific faults are unlikely to be available in many applications of interest (particularly at the 70-m antennas where experimentation is difficult due to operational commitments), methods such as the two-model approach may be a more practical alternative in the long run. Hybrid models which combine the better features of each approach are worth further investigation.

VIII. Conclusions

The use of a two-model cumulative sum detection scheme has been investigated for on-line detection of faults in a dynamic system. This method involves detecting the change in the mean of a function defined on the prediction errors of two recursive AR models estimated sequentially on-line. The algorithm detects changes in the AR model parameters or the signal energy. Experimental results were presented for this two-model method for five controlled hardware faults at the DSS 13 34-m BWG antenna control assembly. The following conclusions can be made in summary:

(1) This method is feasible if prior knowledge of faults is available.

(2) This method can be sensitive to parameter choices. Thus, making more robust detectors which require fewer parameter choices and prior assumptions would be useful.

References

[1] J. Mellstrom and P. Smyth, "Pattern Recognition Techniques Applied to Performance Monitoring of the DSS-13 34-m Antenna Control Assembly," *TDA Progress Report 42-106*, vol. April–June 1991, Jet Propulsion Laboratory, Pasadena, California, pp. 30–51, August 15, 1991.

[2] J. Mellstrom, C. Pierson and P. Smyth, "Real-Time Antenna Fault Diagnosis Experiments at DSS-13," *TDA Progress Report 42-108*, vol. October–December 1991, Jet Propulsion Laboratory, Pasadena, California, pp. 96–108, February 15, 1992.

[3] P. Smyth, "Hidden Markov Models for Fault Detection in Dynamic Systems," submitted to *Pattern Recognition*, May 1992.

[4] D. Hinkley, "Inference About the Change Point From Cumulative Sum Tests," *Biometrika*, vol. 58, no. 3, pp. 509–523, 1971.

[5] M. Basseville, "On-Line Detection of Jumps in Mean," *Detection of Abrupt Changes in Signals and Dynamical Systems*, edited by M. Basseville and A. Benveniste, Berlin: Springer-Verlag, 1986.

[6] M. Basseville, "Detecting Changes in Signals and Systems—A Survey," *Automatica*, vol. 24, no. 3, pp. 309–326, 1988.

[7] M. Basseville, "Edge Detection Using Sequential Methods for Change in Level— Part II: Sequential Detection of Change in Mean," *IEEE Trans. Acoust. Speech Sig. Process.*, vol. ASSP-29, no. 1, pp. 32–50, February 1981.

[8] M. Basseville and A. Benveniste, "Sequential Detection of Abrupt Changes in Spectral Characteristics of Digital Signals," *IEEE Trans. Inf. Theory*, vol. IT-29, no. 5, p. 709–723, September 1983.

[9] M. Eggers and T. Khuon, *Adaptive Preprocessing of Nonstationary Signals*, Technical Report 849, MIT Lincoln Laboratory, Cambridge, Massachusetts, 1989.

[10] S. Kullback, *Information Theory and Statistics*, New York: Wiley and Sons, 1959.

[11] J. Segen and A. Sanderson, "Detecting Changes in a Time-Series," *IEEE Trans. Inf. Theory*, vol. IT-26, no. 2, pp. 249–254, March 1980.

[12] M. Basseville, "The Two Models Approach for the On-Line Detection of Changes in AR Processes," *Detection of Abrupt Changes in Signals and Dynamical Systems*, edited by M. Basseville and A. Benveniste, Berlin: Springer-Verlag, 1986.

[13] E. Page, "Continuous Inspection Schemes," *Biometrika*, vol. 41, pp. 100–115, 1954.

[14] G. Bodenstein and H. Praetorious, "Feature Extraction from the Encephalogram by Adaptive Segmentation," *Proc. IEEE*, vol. 65, pp. 642–652, 1977.

[15] L. Ljung, *System Identification*, Englewood Cliffs, New Jersey: Prentice-Hall, 1987.

[16] S. Marple, *Digital Spectral Analysis With Applications*, Englewood Cliffs, New Jersey: Prentice-Hall, 1987.

[17] M. Basseville and A. Benveniste, "Design and Comparative Study of Some Sequential Jump Detection Algorithms for Digital Signals," *IEEE Trans. Acoust. Speech. Sig. Process.*, vol. ASSP-31, no. 3, pp. 521–535, June 1983.

Table 1. Detection parameters.

Conditions	Threshold h	Drift δ	Detection time D_h units of τ, $\tau = 20$ msec	AR parameters $\theta_1, \theta_2, \theta_3, \theta_4, \theta_5$	Signal energy[a]
Nominal	—	—	—	−0.41, −0.21, −0.12, −0.11, −0.05	0.9897
Fault 1	500	−0.04	1093	−1.61, +0.97, −0.27, −0.13, −0.10	3.9677
Fault 2	500	−8.0	1011	−0.85, −0.37, −0.10, +0.07, +0.29	3.4507
Fault 3	50	−0.01	N/A	−0.40, −0.17, −0.13, −0.14, −0.05	1.1861
Fault 4	50	−0.01	N/A	−0.37, −0.18, −0.14, −0.14, −0.06	0.9333
Fault 5	35	−0.10	1009	−0.02, +0.03, −0.02, 0.0004, −0.06	0.9891

[a] Data units are proportional to one Joule; the value of the proportionality constant is unknown.
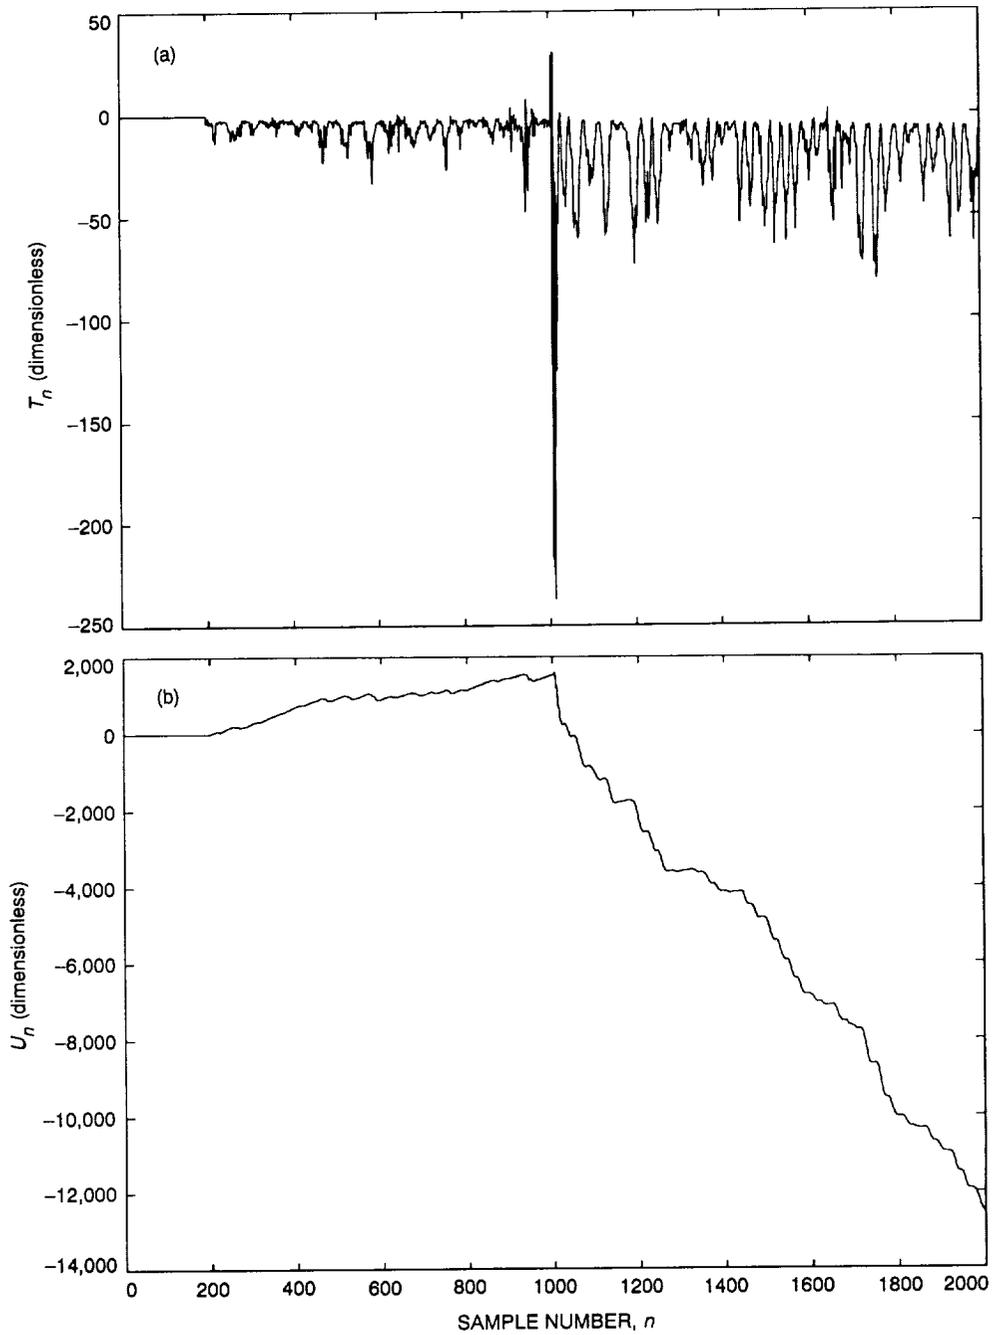
Fig. 1. Low-to-high signal energy transition: (a) typical T statistic and (b) cumulative sum U.
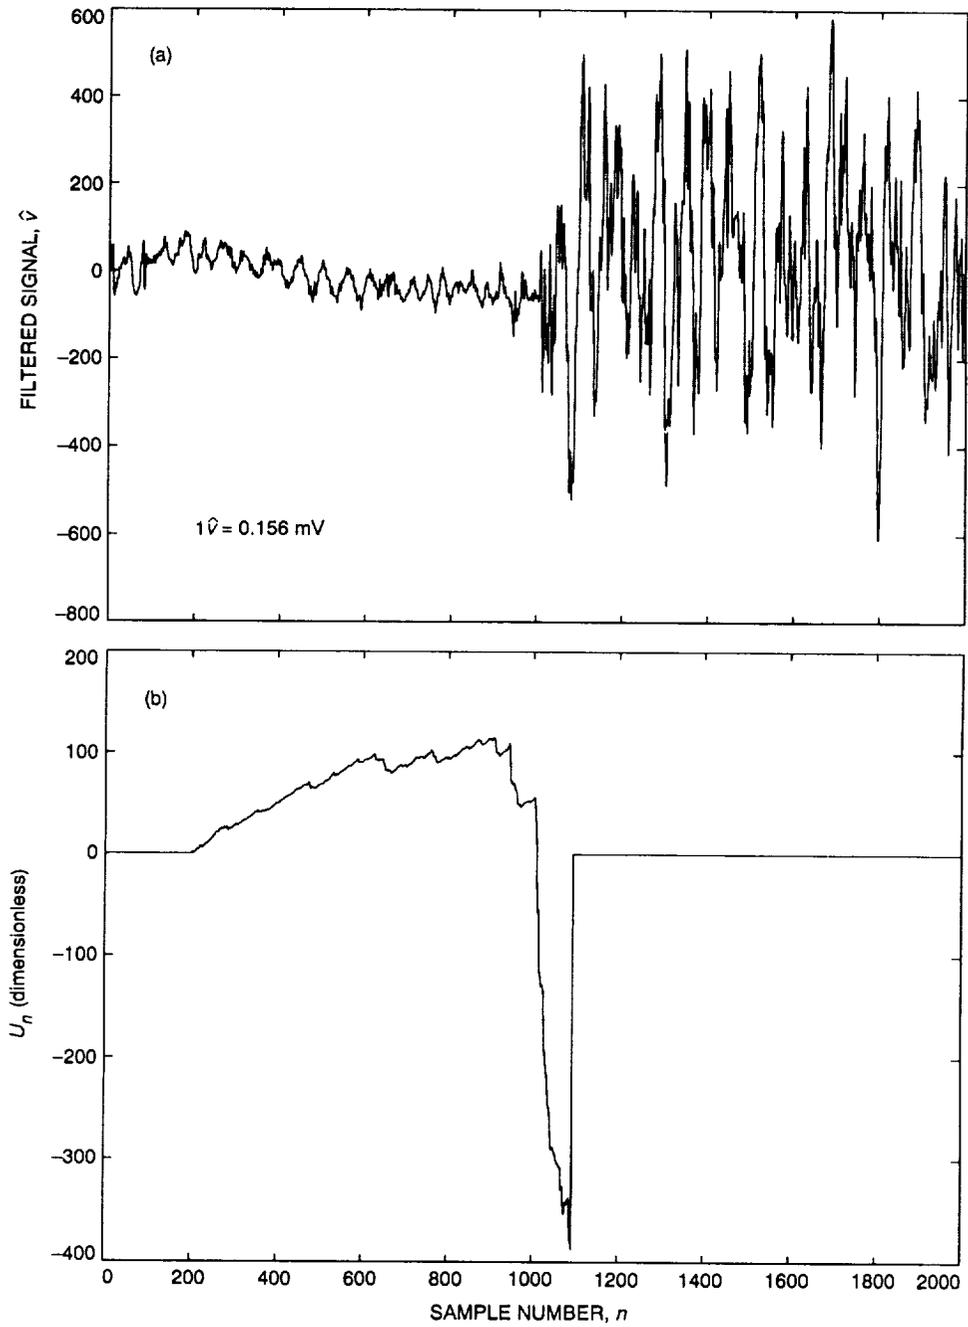
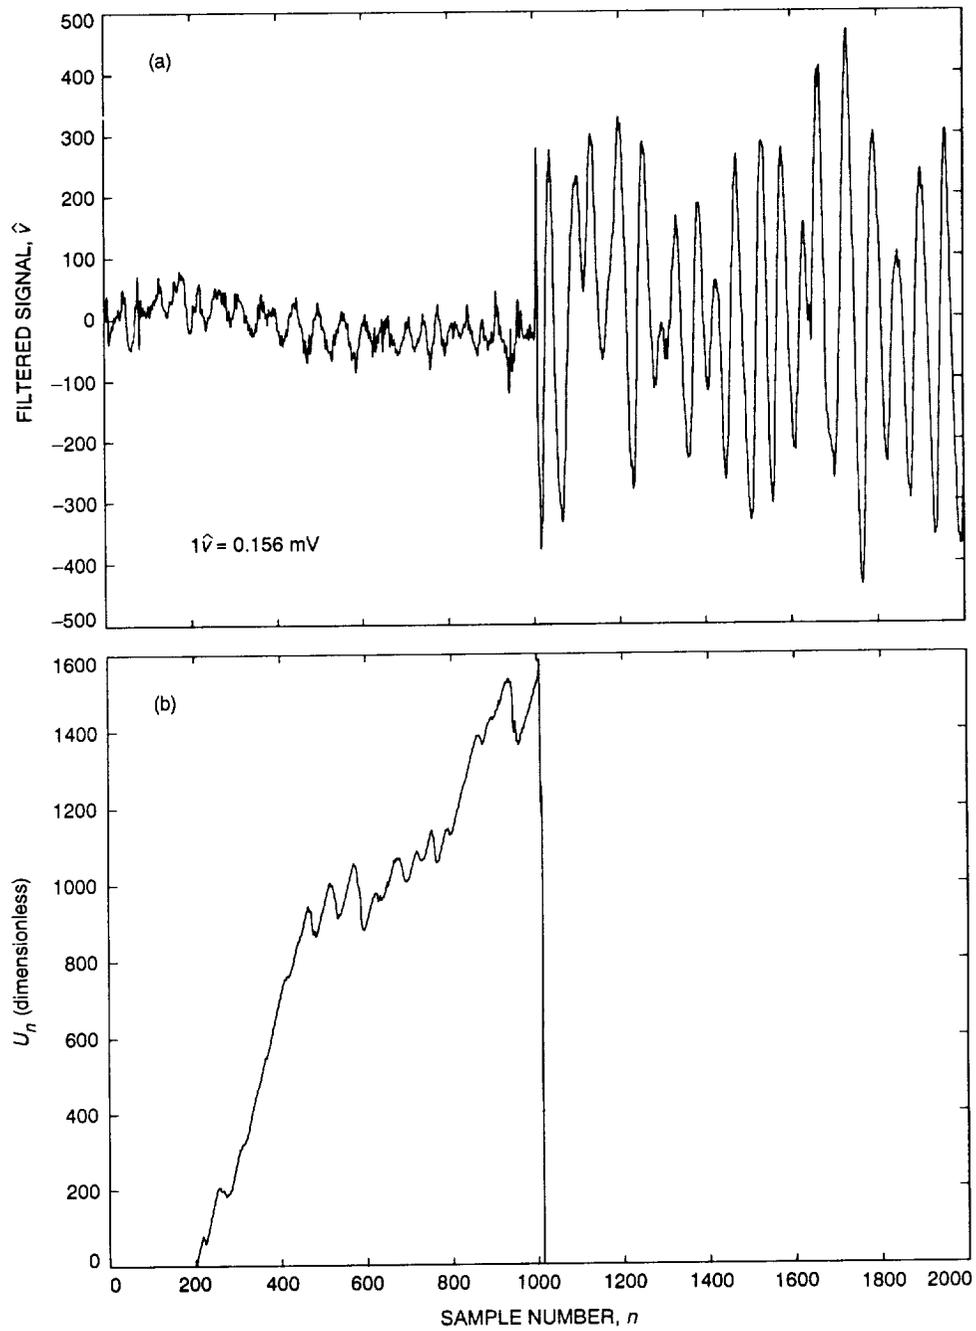Fig. 2. Normal to fault 1 transition: (a) a filtered signal and (b) cumulative sum U.

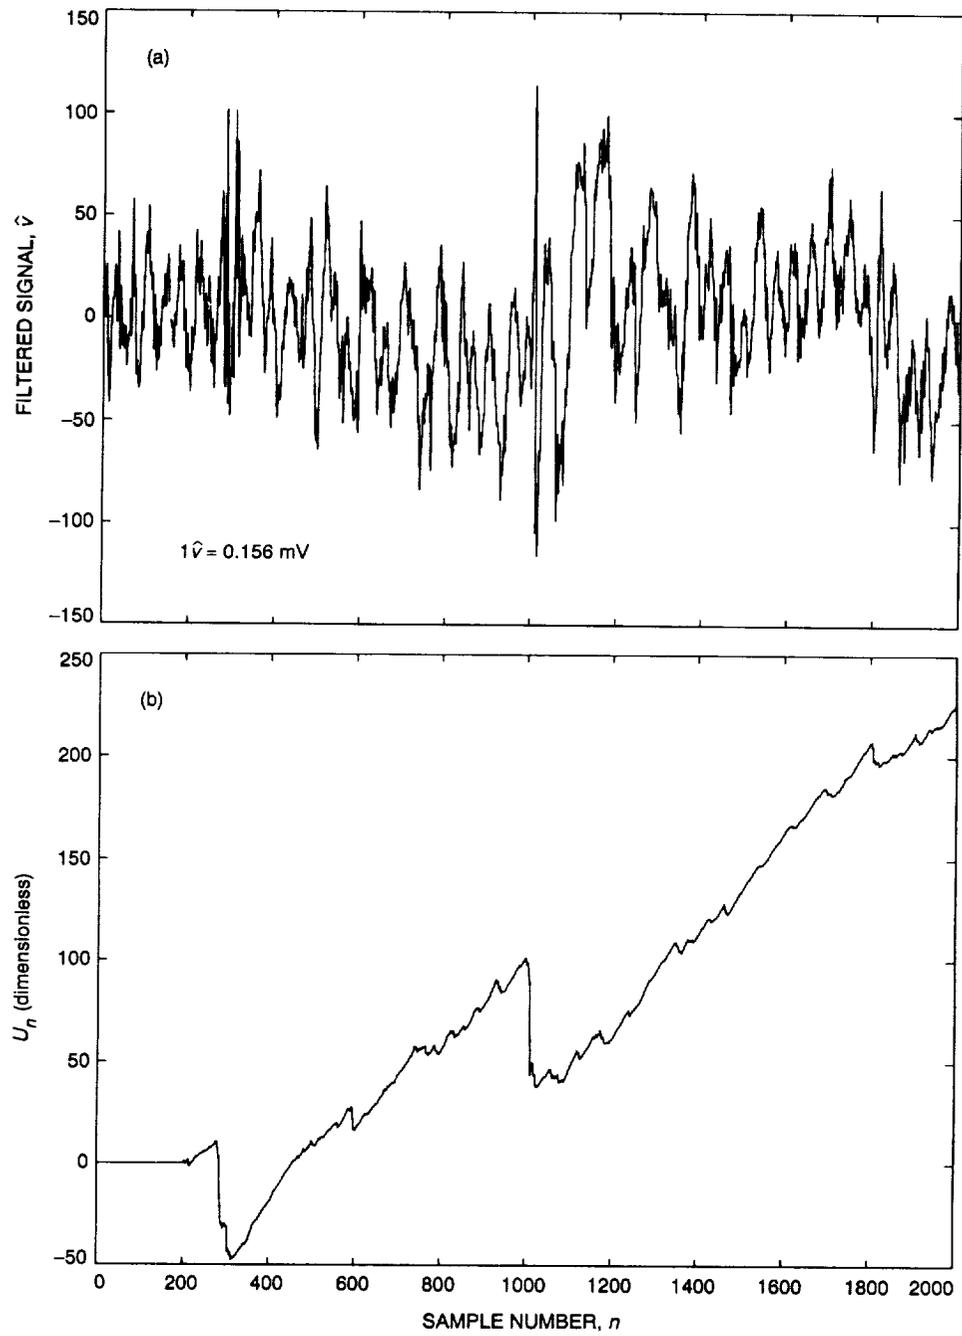Fig. 3. Normal to fault 2 transition: (a) filtered signal and (b) cumulative sum U.

Fig. 4. Normal to fault 3 transition: (a) filtered signal and (b) cumulative sum U.
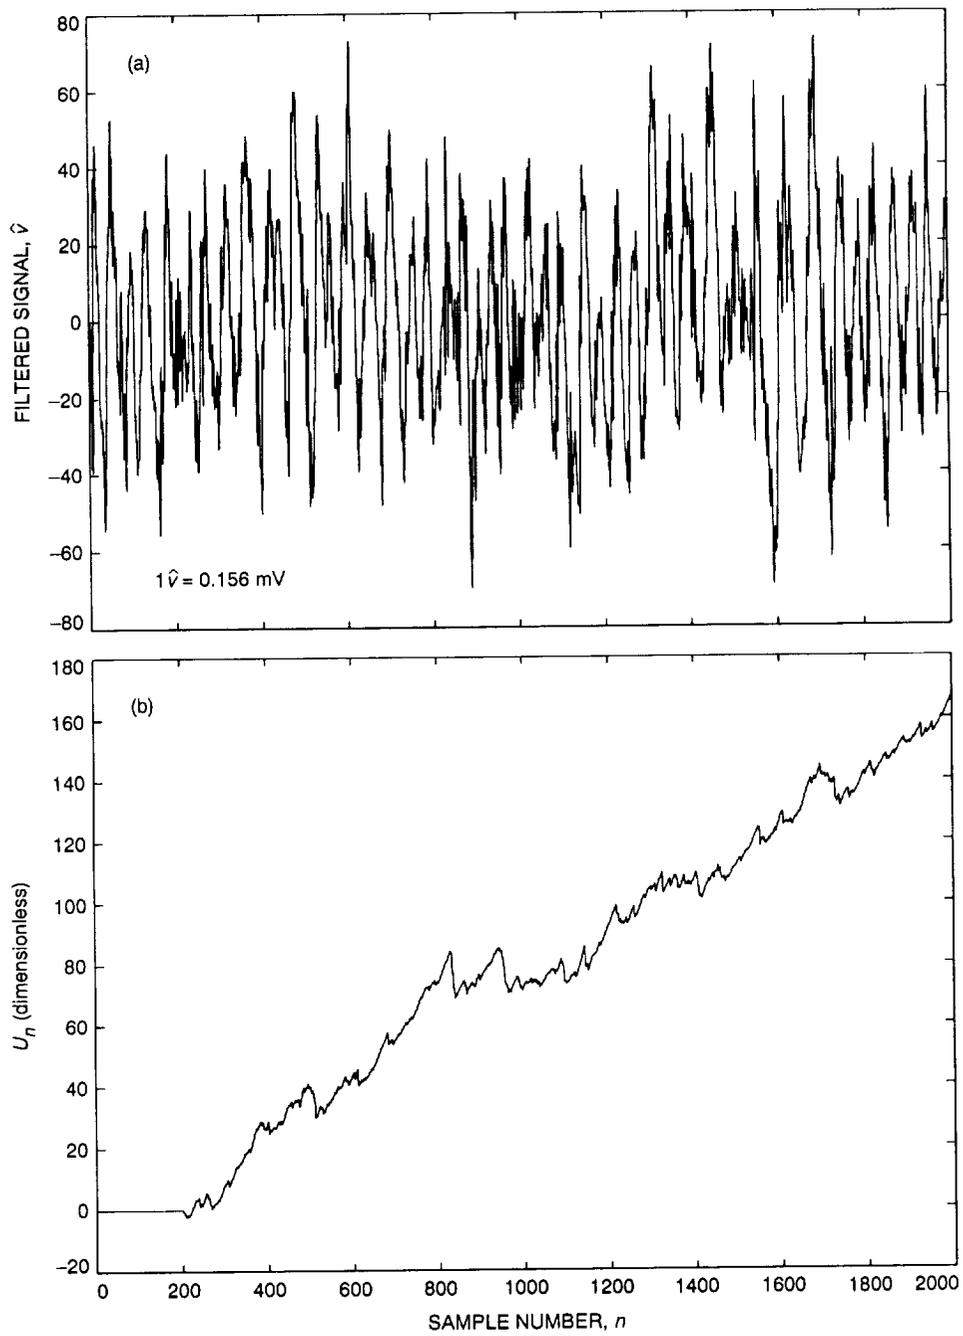
Fig. 5. Normal to fault 4 transition: (a) filtered signal and (b) cumulative sum U.
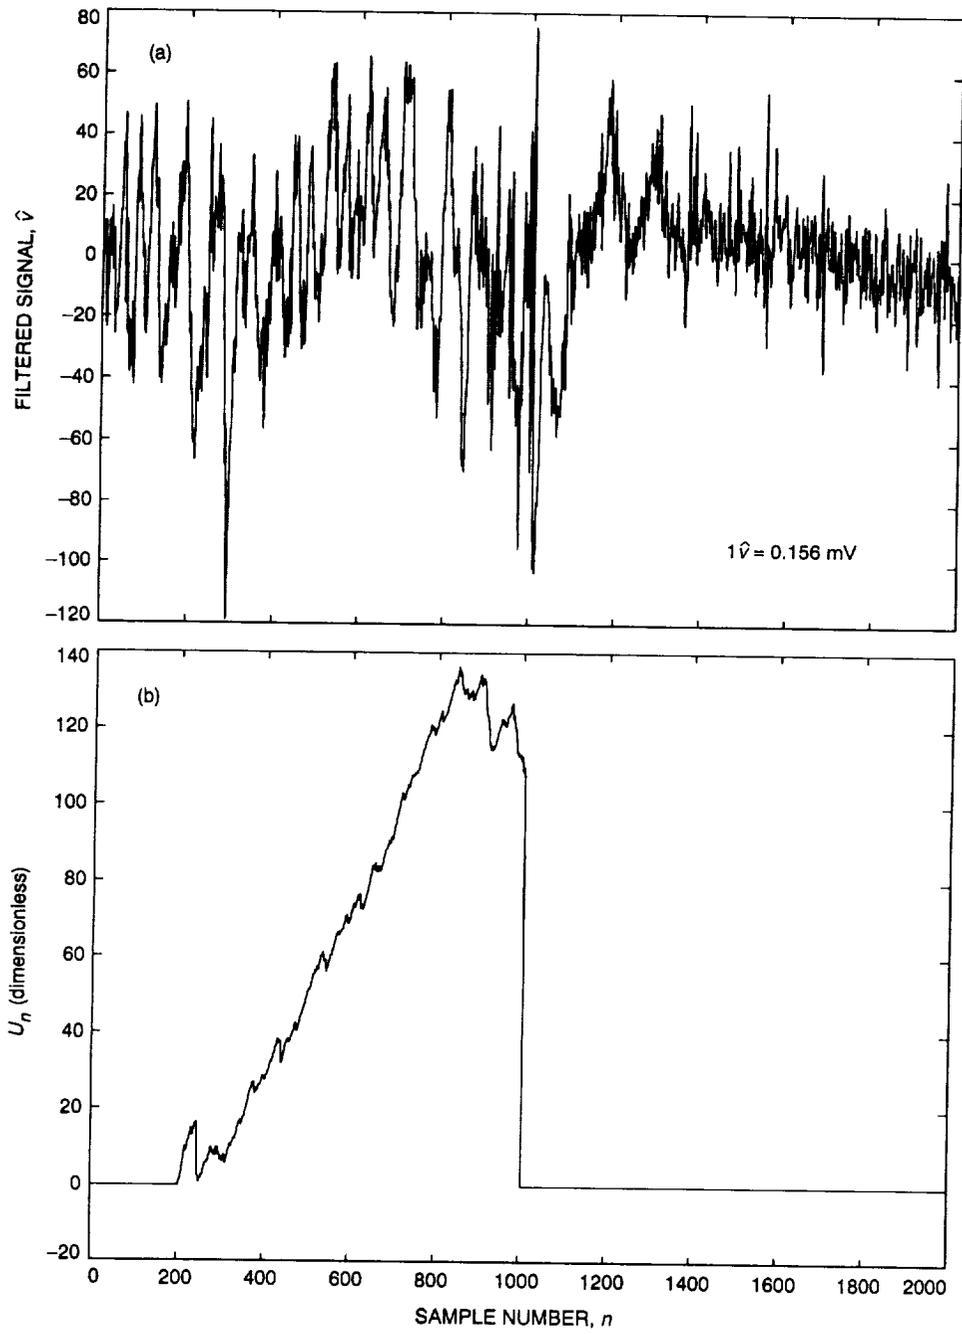
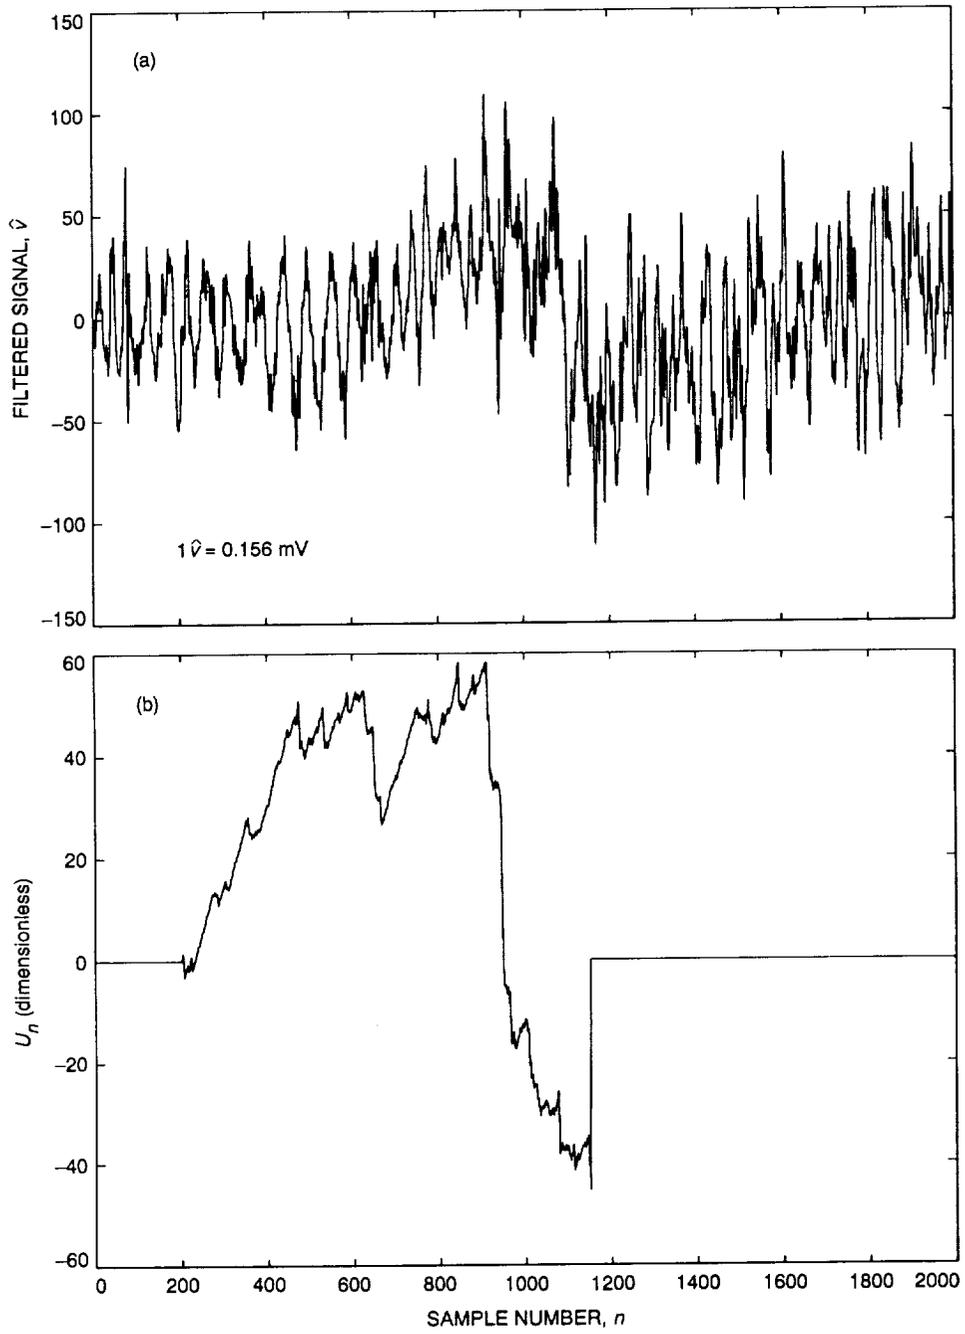Fig. 6. Normal to fault 5 transition: (a) filtered signal and (b) cumulative sum U.

In panel (a): $1\,\hat{v} = 0.156$ mV

Fig. 7. Normal to scaled fault 1 transition: (a) filtered signal and (b) cumulative sum U.

Appendix

Derivation of Detection Statistic T

The derivation of the two-model distance metric T, as presented by Basseville [12], is replicated for the reader's convenience here.

Recall Eq. (5) which gives the distance metric T_k at time k, based on the Kullback conditional information-theoretic distance between probability laws $p^0\left(Y_k|\overline{Y}^{k-1}\right)$ and $p^1\left(Y_k|\overline{Y}^{k-1}\right)$

$$T_k = \int p^0\left(y|\overline{Y}^{k-1}\right) \log \left(\frac{p^1\left(y|\overline{Y}^{k-1}\right)}{p^0\left(y|\overline{Y}^{k-1}\right)}\right) dy$$

$$- \log \left(\frac{p^1\left(Y_k|\overline{Y}^{k-1}\right)}{p^0\left(Y_k|\overline{Y}^{k-1}\right)}\right) \qquad \text{(A-1)}$$

When the observations Y_k are independent identically distributed Gaussian random variables, the conditional probabilities are given by Eq. (7):

$$p^0\left(Y_k|\overline{Y}^{k-1}\right) = \frac{1}{\sqrt{2\pi}\sigma_{\epsilon^0}} e^{-\left[\left(Y_k - \theta^0\overline{Y}^{k-1}\right)^2/2\sigma_{\epsilon^0}^2\right]}$$

$$\text{(A-2)}$$

$$p^1\left(Y_k|\overline{Y}^{k-1}\right) = \frac{1}{\sqrt{2\pi}\sigma_{\epsilon^1}} e^{-\left[\left(Y_k - \theta^1\overline{Y}^{k-1}\right)^2/2\sigma_{\epsilon^1}^2\right]}$$

With direct substitution, T_k becomes

$$T_k = -\frac{1}{2}\log\frac{\sigma_{\epsilon^0}^2}{\sigma^2_{\epsilon^1}} + \frac{\left(y - \theta^0\overline{Y}^{k-1}\right)^2}{2\sigma_{\epsilon^0}^2}$$

$$- \frac{\left(y - \theta^1\overline{Y}^{k-1}\right)^2}{2\sigma_{\epsilon^1}^2} + I_{k-1} \qquad \text{(A-3)}$$

where I_{k-1} is given by

$$I_{k-1} = \int \frac{1}{\sigma_{\epsilon^0}\sqrt{2\pi}}\exp\left(\frac{-\left(y - \theta^0\overline{Y}^{k-1}\right)^2}{2\sigma_{\epsilon^0}^2}\right)$$

$$\times \left[\frac{1}{2}\log\frac{\sigma_{\epsilon^0}^2}{\sigma_{\epsilon^1}^2} + \frac{\left(y - \theta^0\overline{Y}^{k-1}\right)^2}{2\sigma_{\epsilon^0}^2}\right.$$

$$\left. - \frac{\left(y - \theta^1\overline{Y}^{k-1}\right)^2}{2\sigma_{\epsilon^1}^2}\right] dy \qquad \text{(A-4)}$$

Integrating,

$$I_{k-1} = \frac{1}{2}\log\frac{\sigma_{\epsilon^0}^2}{\sigma_{\epsilon^1}^2} + \frac{1}{2} - \frac{1}{2\sigma_{\epsilon^1}^2}I\left(\theta^0\overline{Y}^{k-1}, \theta^1\overline{Y}^{k-1}\right) \quad \text{(A-5)}$$

where

$$I(A, B) = \int \frac{1}{\sigma_{\epsilon^0}\sqrt{2\pi}}e^{[-(y-A)^2]/2\sigma_{\epsilon^0}^2}$$

$$\times (y - B)^2 \, dy \qquad \text{(A-6)}$$

$$= \sigma_{\epsilon^0}^2 + (B - A)^2 \qquad \text{(A-7)}$$

Moreover, by observation,

$$\theta^1\overline{Y}^{k-1} - \theta^0\overline{Y}^{k-1} = \epsilon^0 - \epsilon^1 \qquad \text{(A-8)}$$

Equations (A-5) through (A-8) show that Eq. (9) holds

$$T_k = \frac{1}{2}\left[1 - \frac{\sigma_{\epsilon^0}^2}{\sigma_{\epsilon^1}^2} + \frac{\left(\epsilon^1\right)^2}{\sigma_{\epsilon^1}^2} - \frac{\left(\epsilon^0\right)^2}{\sigma_{\epsilon^0}^2} - \frac{\left(\epsilon^1 - \epsilon^0\right)^2}{\sigma_{\epsilon^1}^2}\right] \quad \text{(A-9)}$$