A Nonlinear Filter to Remove Impulse Noise from Meteorological Data

FRED V. BROCK*

*National Center for Atmospheric Research,** Boulder, CO 80307*

(Manuscript received 12 March 1985, in final form 17 June 1985)

ABSTRACT

Nonlinear median filters were modified to use threshold logic and used to remove impulse noise (spikes) from a set of meteorological data. The impulse noise in the dataset, which originated in the communications section of the Portable Automated Mesonet, could be characterized as random bit noise. Most of the pulses had a duration of one time interval, which in this case was one minute. The filters were effective irrespective of the frequency of occurrence and of the amplitude of the noise spikes. Pulses were removed even when the frequency of occurrence rose to every other data point as was observed in several short intervals. The amplitude of pulses removed ranged over three orders of magnitude.

1. Introduction

Data acquisition systems that transmit data in digital form through a communications channel are susceptible to errors due to impulse noise. The data are usually transmitted in bit serial form with amplitude, frequency, or phase modulation. Impulse noise as a primary source of errors is sporadic and may occur in bursts or discrete impulses. Impulse noise can occur in the communication channel, due to lightning for example, or may occur in the hardware at either end. Impulse noise randomly alters the state of bits in a message (from 0 to 1 or from 1 to 0). Any bit in the message can be affected with equal probability. If a datum is represented as a set of eight to sixteen bits with binary weighting, the effect on the data sequence will be a series of impulses ranging in amplitude from the weight of the least significant bit to that of the most. Since the bits can be changed either way, the impulses can be both positive and negative. Since the error probability is uniform with respect to the bit position in the datum, the resulting amplitude distribution of the impulse errors will be symmetrical with zero mean.

Bit errors in the communications system are quite different in nature from errors arising in other parts of a system, such as in the sensors. Typical errors originating in sensors are offset, calibration drift, hysteresis, and response to unwanted secondary inputs such as solar radiation on temperature sensors. These error sources affect only the specific sensor data, persist for some time, and show definite, although not necessarily obvious, patterns.

A special filter is required that will remove impulse noise while doing the least damage to the data. The median filter suggested by Beaton and Tukey (1974) is a special case of a class of nonlinear filters and is appropriate to this task.

This paper illustrates the impact of impulse noise on data in one particular case, describes a special modification of a median filter and shows how the noise was removed using this modified nonlinear filter.

2. The problem

The Portable Automated Mesonet II (PAM II) developed by the National Center for Atmospheric Research (Brock and Saum, 1983) was deployed in support of the PHOENIX and MAYPOLE projects conducted in May and June 1984. The stations transmitted via the Geostationary Operational Environmental Satellite (GOES) to a base station located in Boulder. A failure in the base station caused a highly significant increase in the bit error rate. The problem was eliminated but not before the dataset was contaminated with noise.

A sample plot of data from the project PHOENIX, contaminated with noise is shown in Fig. 1. This is a plot of one minute data for four hours taken from station 53 on 25 May 1984. The data plotted are pressure, dry- and wet-bulb temperature, accumulated rain, wind speed and maximum wind speed, and wind direction. Noise spikes that exceed the plot limits are plotted only to the plot border. Some of the noise amplitudes are much greater than indicated; one pressure noise impulse has an amplitude of 122 mb. Others are so small as to be barely noticeable in the plot, e.g., the pressure noise impulse just after minute 120.

The wind speed is the average wind speed over one minute, the maximum wind speed is the maximum

* Currently at The University of Oklahoma, Norman, OK 73019.

** The National Center for Atmospheric Research is sponsored by the National Science Foundation.

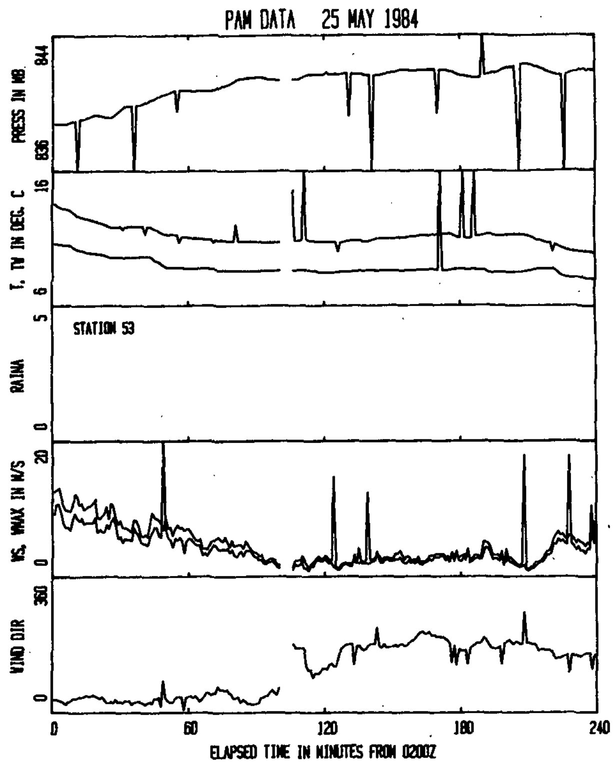

FIG. 1. Data from station 53 for a 4-hour period commencing at 0200 GMT 25 May 1984 in Project PHOENIX. Shows impulse noise present in 6 of the 7 data channels: pressure, dry- and wet-bulb temperature, max wind speed, wind speed and direction in degrees. Only the accumulative rain channel is noise-free in this set. There is missing data from minute 91 through minute 95 in the record.

10-second gust within the minute. The remote station transmitted the horizontal wind vector in orthogonal components and the speed and direction were computed at the base station. That is why the impulses in wind direction are relatively small; if the wind direction had been transmitted as such and subjected to communication noise, some of the wind direction impulses would have exceeded the 0 to 360 degree range.

3. Possible approaches

Most filters used in data processing are linear filters since, by virtue of the superposition principle, their performance can be defined in the frequency domain. However, the effects of impulse noise are spread throughout the frequency domain. Consider the Fourier transform of a rectangular pulse of height a and half-width b,

$$p(t) = \begin{cases} 0 & \text{for} \quad |t| > b \\ a & \text{for} \quad |t| \leq b \end{cases} \tag{1}$$

$$F[p(t)] = \int_{-\infty}^{\infty} p(t)\exp(-j\omega t)dt$$
$$= 2ab\sin(\omega b)/\omega b, \tag{2}$$

where $\omega = 2\pi f$. Most of the variance is at low frequencies but some ripple is apparent at higher frequencies. As the pulse amplitude grows and the width decreases, the pulse approximates the impulse function, for which the noise is named,

$$\delta(t) = \infty \quad \text{for} \quad t = 0, \quad 0 \text{ otherwise}$$
$$\int_{-\infty}^{\infty} \delta(t) = 1. \tag{3}$$

The Fourier transform is given by

$$F[\delta(t)] = \int_{-\infty}^{\infty} \delta(t)\exp(-j\omega t)dt = 1. \tag{4}$$

A single true impulse function affects the entire spectrum equally. The impact of a large, but finite, impulse on the true spectrum is similar to the ideal impulse function. Therefore the filter required cannot be defined in the frequency domain.

4. Median filters

The median filter was suggested by Beaton and Tukey (1974) for robust noise suppression. It is one of a class of nonlinear filters reported by Bednar and Watt (1984) and by Kundu et al. (1984). The properties of the median filter have been carefully defined by Gallagher and Wise (1981) and appear to be well-suited to this problem. The first-order median filter ($N = 1$, window length $L = 2N + 1 = 3$) is given by:

$$Y_I = \text{median}\ \{X_{I-1}, X_I, X_{I+1}\} \tag{5}$$

and the second-order median filter ($N = 2$, window length $L = 2N + 1 = 5$) by:

$$Y_I = \text{median}\ \{X_{I-2}, X_{I-1}, X_I, X_{I+1}, X_{I+2}\}. \tag{6}$$

The action of a first-order median filter ($N = 1$, window length $L = 3$) on a sequence of data is shown in Table 1. The window initially selects the first three points, finds the median value, 85, which is placed in the output sequence. In this example the first and last points have been copied directly to the output sequence. As the window moves across the data, the suc-

TABLE 1. Example of a first-order median filter. Line 0 is the filter input, lines 1 through 6 show the data in the window at each step as the filter is moved through the data, and line 7 is the filter output. Values that were changed as a result of filtering are underlined.

Input	Line 0	87	85	84	21	86	11	86	85
	Line 1	87	85	84					
	Line 2		85	84	21				
	Line 3			84	21	86			
	Line 4				21	86	11		
	Line 5					86	11	86	
	Line 6						11	86	85
Output	Line 7	87	85	84	84	21	86	85	85

cessive windows, with replacement, are shown in lines 1 through 6 of the table. If the numbers 21 and 11 represent impulses, then the filter correctly replaced the value 11 with 86 but failed to remove the value 21. A second pass of the filter would remove the remaining impulse. Alternatively, the performance of the filter can be improved by modifying the filter to allow recursion, that is, the replacement of the data with the filter output as the filter window progresses. With recursion, the defining equations for the first- and second-order median filters become

$$Y_I = \text{median}\ \{Y_{I-1}, X_I, X_{I+1}\} \qquad (7)$$

$$Y_I = \text{median}\ \{Y_{I-2}, Y_{I-1}, X_I, X_{I+1}, X_{I+2}\} \qquad (8)$$

and the performance with the same input data sequence is shown in Table 2. Now both impulses, 21 and 11, have been replaced in one pass. If the number 86 following the 21 was a good datum the filter acted incorrectly in replacing it with 84. Evidently the filter can be useful but we must categorize and then modify its action.

Since the median filter is nonlinear, its performance cannot be defined in the frequency domain, however, it may be defined by the response of the filter to various patterns in the data. By analogy with the description of linear filters in the frequency domain, the data patterns that are passed unchanged by the median filter are considered to be in the "pass band" of the median filter. The patterns that are removed are in the "stop band" of the filter.

Figure 2 shows how a first-order median filter responds to a fabricated signal. The top panel shows the input data sequence, which comprises a monotonically increasing signal with impulse noise added, a constant region, a monotonically decreasing signal with noise added, several step functions with noise added, and a sinusoid with random, Gaussian signal fluctuations. The latter part of the input data sequence is considered to be noise free.

The second panel of Fig. 2 shows the filter output or pass band, that is, the input data sequence passed by the filter. The third panel shows the stop band, defined as the input data minus the filter output. The scale of the stop band has been magnified to enhance smaller differences. The monotonically changing signal and the constant regions are passed without alteration and all of the impulse noise spikes but one are stopped. The one impulse passed consists of two consecutive impulses with different amplitudes. The action of the filter is to truncate the larger impulse to the amplitude of the shorter impulse. This is characteristic of a median filter—it can not stop a pulse whose width is greater than N.

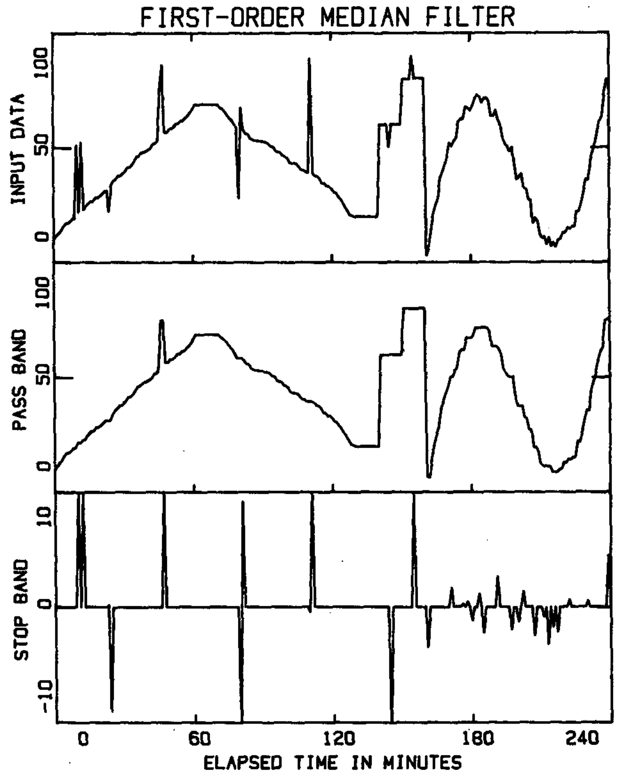

FIG. 2. Input data, passband and stopband of a first-order median filter operating on a special set of test data.

TABLE 2. Example of a first-order median filter operating *with* replacement. Line 0 is the filter input, lines 1 through 6 show the data in the window at each step as the filter is moved through the data, and line 7 is the filter output. Values that were changed as a result of filtering are underlined.

Input	Line 0	87	85	84	21	86	11	86	85
	Line 1	87	85	84					
	Line 2		85	84	21				
	Line 3			84	21	86			
	Line 4				<u>84</u>	86	11		
	Line 5					<u>84</u>	11	86	
	Line 6						<u>84</u>	86	85
Output	Line 7	87	85	84	<u>84</u>	<u>84</u>	<u>84</u>	<u>85</u>	85

Figure 3 shows the effect of a second-order median filter on the same input dataset. The double impulse has been stopped but the filter has also removed more valid signal.

In Figs. 2 and 3 the filter has removed some of the random, Gaussian signal fluctuations. Clearly, the filter doesn't view the data as one would wish. The problem is to define the difference between valid signal, which is to be passed, and noise, which is to be stopped. The stop band display shows that some of the random signal looks like impulse noise.

This example was fabricated such that the noise impulses are large and easily distinguished. In the real dataset it is not so easy to distinguish noise from data. Careful examination of the data revealed cases where very small amplitude impulse noise was present. If

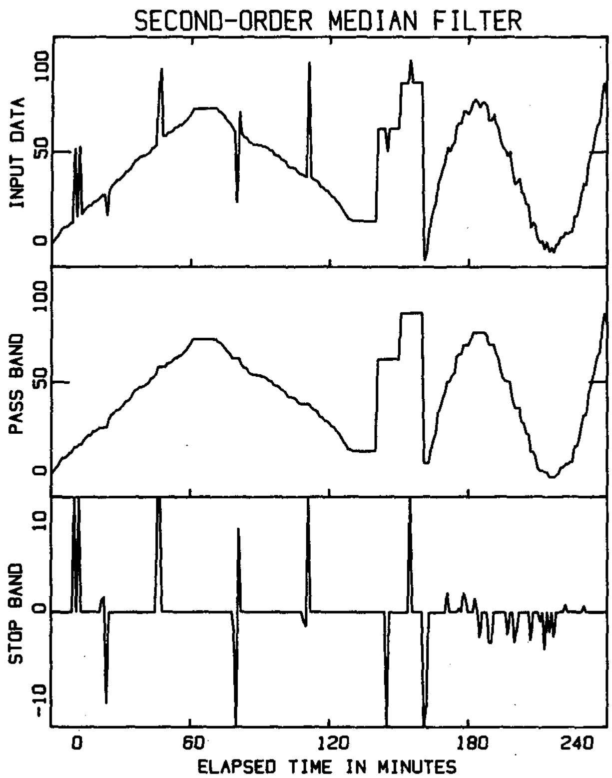

FIG. 3. Passband and stopband of a second-order median filter operating on a special set of test data.

flipping of the least significant bits is as probable as flipping of most significant bits, then there must be noise present with amplitudes down to the weight of the least significant bit and it will not be possible to remove these.

Properties of median filters:

1) A median filter of order N has a window length $2N + 1$. The filtering operation consists of sliding this window across the data and the filter output is the median of the values in the window and is associated with the time sample at the center of the window.

2) The pass band and stop band of a median filter can be defined by the signal patterns in the time domain, which are passed unchanged or removed, respectively.

3) A median filter will pass a signal that consists only of edges and constant neighborhoods consisting of at least $N + 1$ consecutive points (Gallagher and Wise, 1981). By extension, a signal will be invariant to median filtering if it consists of monotonically increasing and monotonically decreasing regions separated by constant regions of length at least $N + 1$, as shown in Figs. 2 and 3.

4) The stop band of a median filter will include an impulse whose width is less than or equal to N.

5) As the frequency of occurrence of impulse noise increases, the performance of median filters remains unchanged.

6) The performance of median filters is not affected by the amplitude of impulse noise.

7) The efficiency of median filtering is improved by recursion, i.e., replacement of data as the filter progresses. This was illustrated in Tables 1 and 2.

5. A modified median filter with threshold logic

A median filter can be further modified by using threshold logic (Pasian and Crise, 1984) to decide whether an impulse is noise or simply part of the normal signal fluctuations. The threshold prohibits replacement of most of the random signal fluctuations at the cost of passing small noise spikes.

To determine the threshold, the filter is first passed over the data without replacement. A histogram is constructed from the difference values between the data and the filter outputs. If the data were free of impulse noise, the filter output would have an approximately normal distribution. With impulse noise added, there would be positive and/or negative lobes in the distribution. Figure 4 shows a typical histogram of data with impulse noise. Zero differences are contained in bin 12. A square root function is used to compress the frequency scale so that noise impulses of large amplitude, but with low frequency of occurrence, can be seen. The central distribution contained in bins 6 through 14 is assumed to be the desired signal while the impulse noise is contained in bins 0, 19 and 24. The desired threshold is a value just large enough to encompass the central lobe of the distribution and exclude the spikes.

The threshold can be found by searching in both directions from the center of the histogram for that value for which the count is minimum. Then the filter can be passed over the data again and the threshold

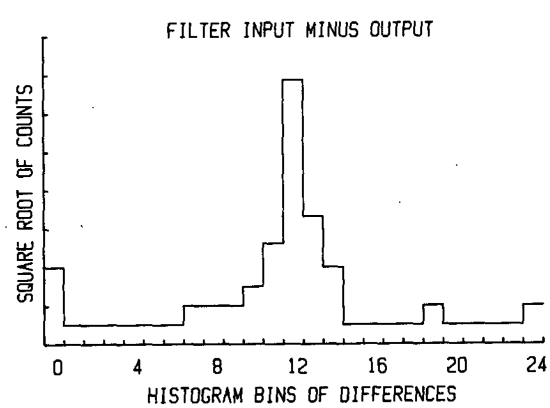

FIG. 4. Typical histogram of the difference between the filter output and the input data. Includes data and noise.

used to determine whether to replace data with the filter output.

The operation of the modified filter is as follows:

(i) Given an input sequence $\{X_I\}$, pass the filter, first- or second-order, over the data without replacing any data.

(ii) Compute the histogram of the sequence $\{D_I\}$ which is $D_I = X_I - Y_I$ where Y_I is the filter output.

(iii) Determine the threshold by searching the histogram for the first minima from the center. Use this value as the threshold, DT.

(iv) Use the differences from step (i) above or pass the filter over the data again, finding the difference at each step and applying threshold logic:

$$X_I = Y_I, \quad \text{if} \quad |D_I| > DT$$

$$X_I = X_I, \quad \text{otherwise}$$

and the sequence $\{X_I\}$ becomes the filter output.

While the original median filter had only one parameter, the filter order, the modified filter has another parameter—the size of the histogram bins, since the number of bins was held constant. If the bin size is too small, artificial gaps can form because the data have finite resolution determined by the weight of the least significant bit when the data were digitized. Then the filter threshold will be too small and too much acceptable data will be removed. If the bin size is too large, too much of the noise will be retained.

A procedure for selecting the proper bin size is to initially set it to twice the resolution, i.e. twice the weight of the least significant bit when the data were digitized. The number of bins can be fairly small, say about 25. It should be an odd number because the mean of the distribution of the differences, which is expected to be zero, should appear in the middle bin. After the distribution has been generated, search for the first bin from the center, in either direction, that contains zero counts. If no bin is found with zero counts, double the bin size and recompute the histogram. In this way, a histogram with the smallest practical bin size will be generated and the appropriate threshold found. The threshold increases approximately as the variance of the data in the absence of noise. As was noted above, many of the impulses are so large that this histogram procedure gives better results than one based on direct variance calculation.

6. Results

The filter implementation selected for the PAM data used both a first-order and a second-order median filter, both modified with threshold logic. Before filtering, the missing data, which were represented by the value −99., were replaced to avoid confusion with impulses. Each missing data value was replaced with the median of the three adjacent points, excluding other missing data in the process. The median was used to avoid inadvertent propagation of impulses that might be adjacent to the missing data. In addition, the end points of the record were replaced with the median of the three points at either end of the record. This was done to handle the possibility of impulses at the beginning or end of the record.

When PAM data are formatted for a user, the wind vector is represented as wind speed and direction. The wind vector data were in orthogonal component form when subjected to impulse noise. The later conversion to polar coordinates altered the characteristics of the noise, making the filters less effective. Therefore, prior to filtering, the wind vector was converted back to orthogonal components. In a post-filtering pass, the filtered orthogonal components were again converted to speed and direction.

The effects of the filtering can be seen in Figs. 5–9. Figure 5 shows the pressure data from station 53. The top panel shows the unfiltered data. The next panel shows the pass band, that is, the filter output, plotted to the same scale; the bottom panel shows the stop band, that is, the original data minus filtered data plotted to an expanded scale of −1 to 1 mb. The filter

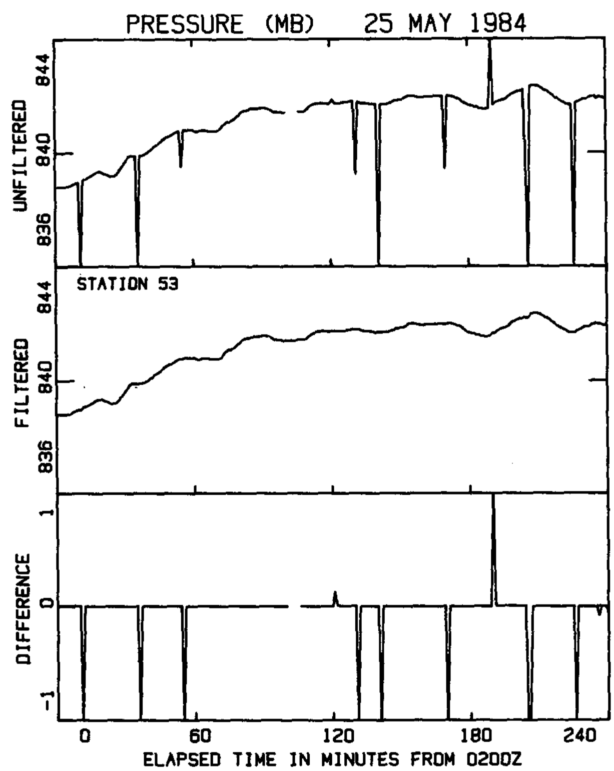

FIG. 5. Effects of the filtering operation on pressure data from station 53. Unfiltered data is the filter input. Filtered data is the filter output, difference shows the impulse noise removed by the filter on an expanded scale of −1.0 to 1.0 mb.

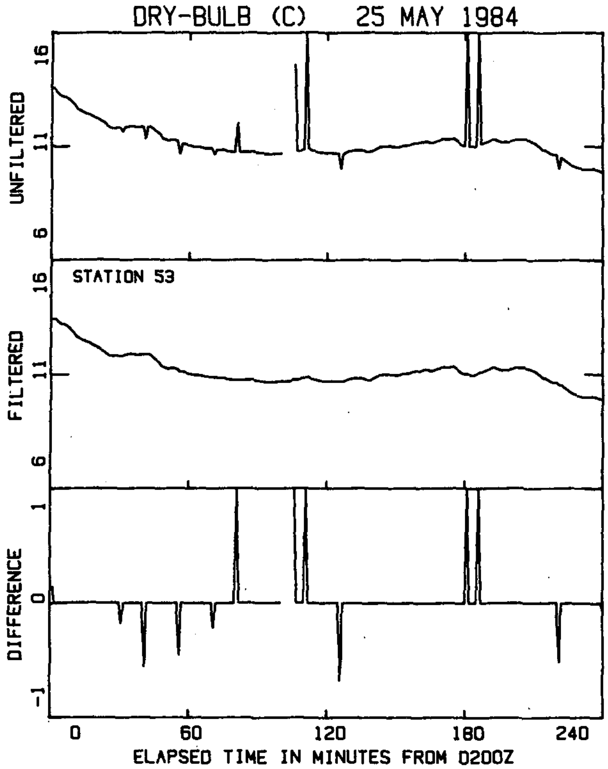

FIG. 6. Effect of the filters on dry-bulb temperature. The difference plot scale is −1° to 1°C.

removed all of the major impulse noise spikes and two smaller ones at minute 121 (amplitude = 0.12 mb) and minute 236 (−0.08 mb). The largest impulse removed was at minute 11 and had an amplitude of −122.81 mb. Unlike linear filters, the median filters remove impulses irrespective of size, above the threshold, without any side effects.

Figure 6 shows the dry-bulb data from station 53. Here the difference scale is −1° to 1°C. Eleven impulse spikes were removed ranging in amplitude from −0.18° to +49.66°C. In addition, the end-point treatment described above modified the first and last points in the dataset.

Analysis of the filter performance is more complicated for wind speed, (Fig. 7), wind direction and maximum wind speed. As noted before, the wind speed and direction data were converted to orthogonal components before filtering and then restored after filtering. In each case the clearly identifiable spikes were removed including some impulses that may have been valid data. The adaptive nature of the filter algorithm is evident here in that the filter threshold has changed in response to the data. The algorithm set the threshold much higher for the wind data with its noticeably greater variance than for the pressure or temperature data. In Fig. 7 there are four spikes that are clearly noise and many more that are probably noise. For example, the five spikes removed between minutes 120 and 180 appear to be acceptable filter action. On the other hand, most observers would likely only remove the single large spike in the first 60 minutes of the record. The filter removed 9 spikes.

Figure 8 shows the dry-bulb data from station 24. It had 45 noise impulses removed by the filter. As was noted, the performance of the filter is unaffected by the noise amplitude, as long as it is above the threshold. This figure shows that the filter performance is not affected by the noise frequency. It is possible that a second pass of the filters would remove some of the smaller noise impulses left in the data since the threshold would be recomputed using a new histogram.

The dry-bulb data from station 12, Fig. 9, was noteworthy in that there were about 40 impulses of very small amplitude, about −0.12°C. The filter removed about half of these. The threshold level must have been very close to the impulse amplitude. To the filter, the impulse amplitude is the length of the shortest leg. If the temperature is increasing or decreasing, an impulse of a given amplitude will be perceived by the filter as being shorter than a impulse of the same amplitude when the temperature is constant. Notice that fewer impulses are removed when the temperature trace is changing relatively rapidly.

There was a triple impulse in the wet-bulb temperature for station 42, that was not removed; it was passed

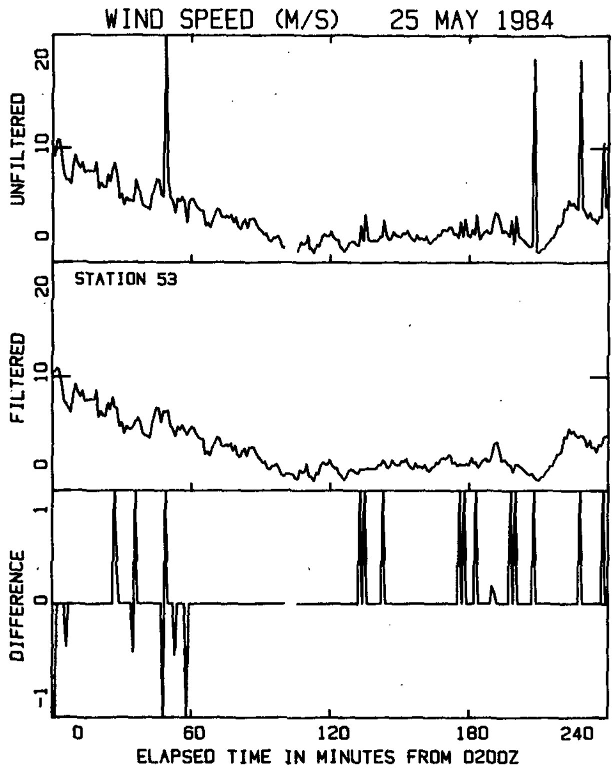

FIG. 7. The wind speed signal contains many valid fluctuations that resemble impulse noise. Some of the pulses removed may have real signal and not noise.

by both first- and second-order filters. A third-order median filter could have been used to eliminate this impulse but would have had undesirable side effects especially on the wind data. It is likely that somewhere in the data there are four sequential pulses, maybe even five. This kind of filter can not remove all bad data.

7. Conclusions

The specific problem that occurred in the PAM base station was unusual and is unlikely to recur. The nature of the noise introduced, however, was fairly typical of random bit noise, which can be generated in a variety of ways. In the broad sense, every data acquisition system includes a communication system, since even tape storage is a communication system. While systems are designed to have acceptably low bit error rates, failures can occur and when they do, random bit noise may be the result. The techniques developed here are applicable to the class of random bit noise problems and are, therefore, quite general.

The combination of first- and second-order median filters with threshold logic performed satisfactorily. The filters were effective in removing impulses whose amplitude exceeded the threshold and their performance was not affected by large impulse amplitudes or by a high frequency of occurrence. The threshold logic effectively adapted to the underlying data variance so that pulses apparent to the observer were effectively removed whether in the low variance temperature data or in the higher variance wind speed data. The first- and second-order filters failed to remove three or more consecutive impulses. Higher order filters could be used but they remove more of the valid signal fluctuation.

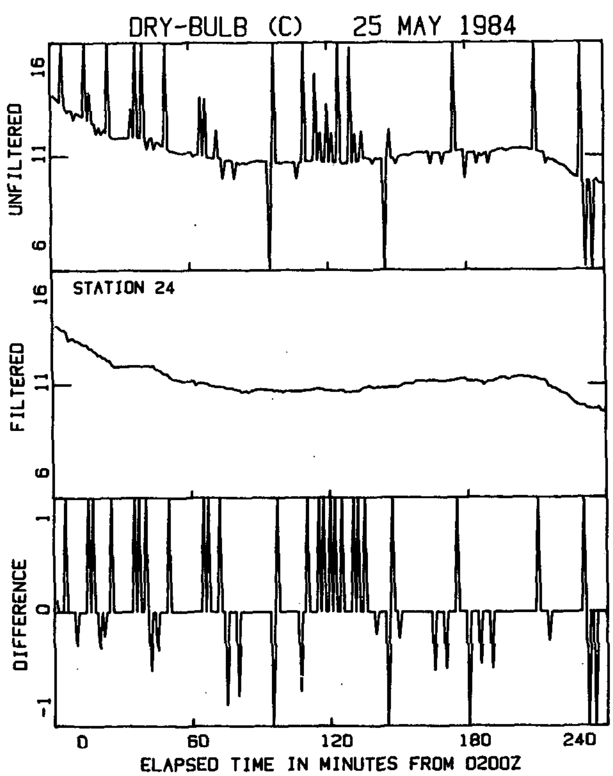

FIG. 8. Filter operation on dry-bulb temperature from station 24. The filter removed 45 pulses from the data but left some detectable fluctuations that were probably pulses.

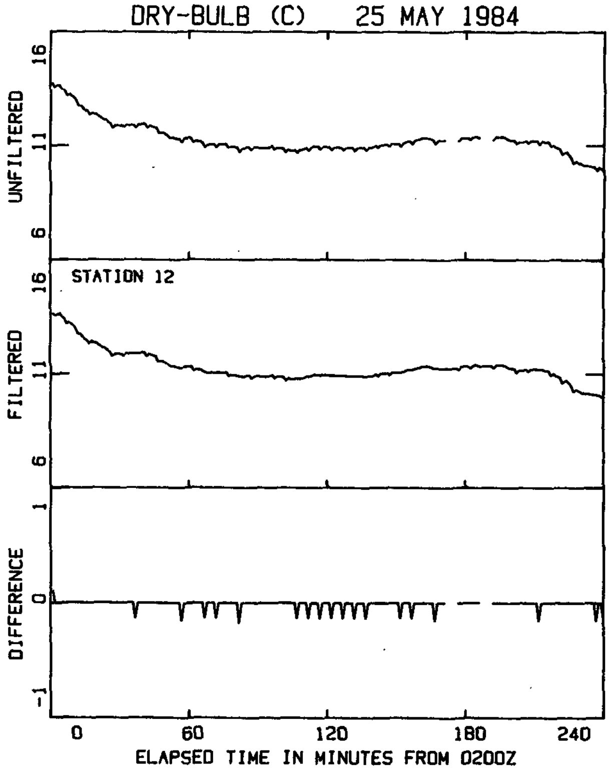

FIG. 9. In this case, the temperature signal had about 40 small impulses and the filter removed only 19. The rest were below the threshold, some because of an uneven base. The smallest pulse removed had an amplitude of −0.12°C.

Missing data were filled in but it would be easy to convert the original missing data points back to missing data. The data results plotted in the figures contain the replacement values for the impulse noise as selected by the filters, which may not be optimum for some purposes. As the median filter is passed over the data, it tends to propagate the last valid datum when replacing a series of impulses. Therefore it might be desirable to filter the data, in a final pass, with a low-pass filter designed to attenuate the highest frequency components to ameliorate the effects of the median filtering. This was not done in order to show clearly the results of the median filters. As an alternative, it would be possible to keep track of all data values replaced during filtering and convert them to missing data values as the final step.

Acknowledgments. I wish to acknowledge the assistance rendered by John Militzer of the NCAR Field

Observing Facility in providing very timely access to the data and Richard Carbone and Robert Serafin also of NCAR, for providing the encouragement, ways and means to proceed with this project.

This work was accomplished while visiting the Department of Meteorology of The University of Hawaii at Manoa, Honolulu, Hawaii. I would like to thank Profs. Wan-Cheng Chiu and Anders Daniels for their assistance in reviewing this article.

REFERENCES

Beaton, A. E., and J. W. Tukey, 1974: The fitting of power series, meaning polynomials, illustrated on band-spectroscopic data, *Technometrics,* **16,** 147–192.

Bednar, J. B., and T. L. Watt, 1984: Alpha-trimmed means and their relationship to median filters. *IEEE Trans. Acoust. Speech Signal Proc.,* **ASSP-32,** 145–153.

Brock, F. V., and G. H. Saum, 1983: Portable automated mesonet II. *Fifth Symp. Meteorological Observations and Instrumentation,* Toronto, Amer. Meteor. Soc., 314–320.

Gallagher, N. C., Jr., and G. L. Wise, 1981: A theoretical analysis of the properties of median filters, *IEEE Trans. Acoust. Speech Signal Proc.,* **ASSP-29,** 1136–1141.

Kundu, A., S. K. Mitra and P. P. Vaidyanathan, 1984: Application of two-dimensional generalized mean filtering for removal of impulse noise from images. *IEEE Trans. Acoust. Speech Signal Proc.,* **ASSP-32,** 600–609.

Pasian, F., and A. Crise, 1984: Restoration of signals degraded by impulse noise by means of a low-distortion nonlinear filter. *Signal Processing,* **6,** 67–76.