

AD-A031 475

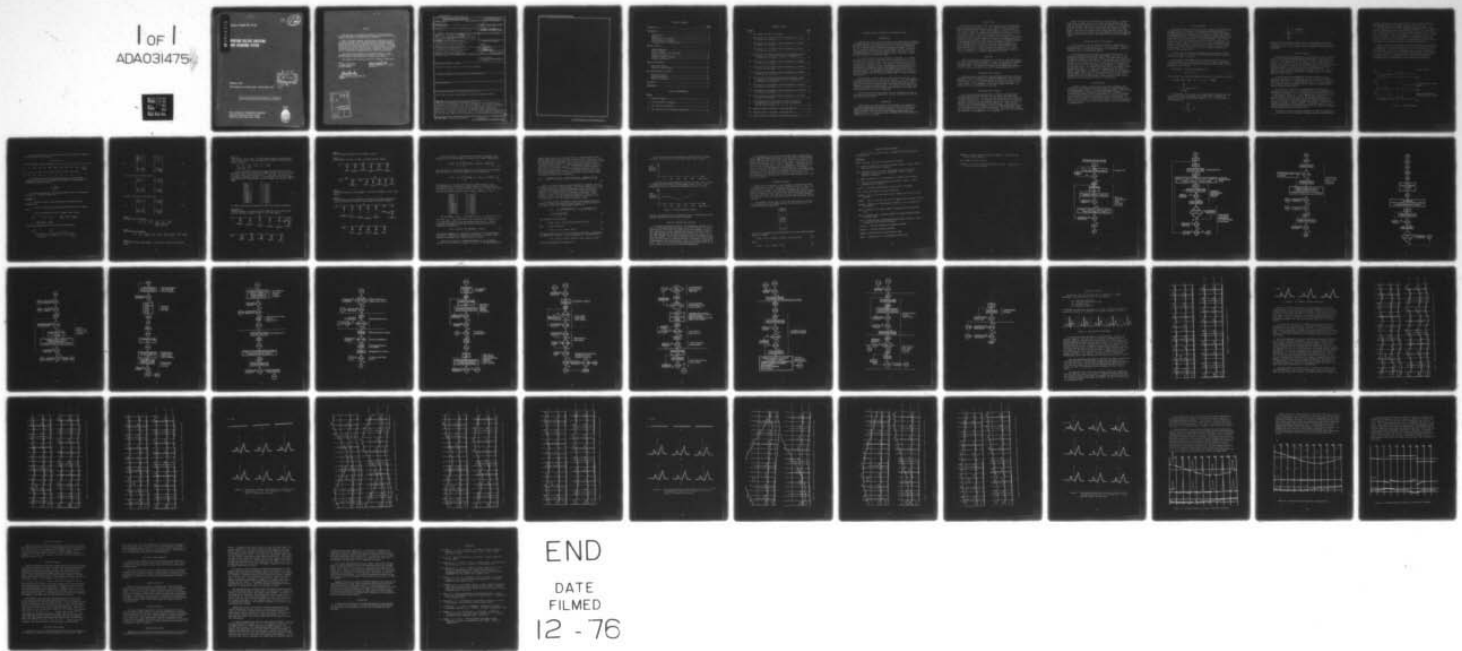
SCHOOL OF AEROSPACE MEDICINE BROOKS AFB TEX
USAFSAM ECG/VCG DIGITIZING AND AVERAGING SYSTEM. (U)
SEP 76 M E WOMBLE, H N KEISER
SAM-TR-76-31

F/G 9/2

UNCLASSIFIED

NL

1 of 1
ADA031475



END
DATE
FILMED
12 - 76

AD A031475

Report SAM-TR-76-31

FG

12

**USAFSAM ECG/VCG DIGITIZING
AND AVERAGING SYSTEM**

September 1976

Final Report for Period August 1972-February 1976

DDC
RECEIVED
NOV 2 1976
B

Approved for public release; distribution unlimited.

**USAF SCHOOL OF AEROSPACE MEDICINE
Aerospace Medical Division (AFSC)
Brooks Air Force Base, Texas 78235**



NOTICES

This final report was submitted by personnel of the Clinical Sciences Division, USAF School of Aerospace Medicine, Aerospace Medical Division, AFSC, Brooks Air Force Base, Texas, under job order 7755-01-02.

When U.S. Government drawings, specifications, or other data are used for any purpose other than a definitely related Government procurement operation, the Government thereby incurs no responsibility nor any obligation whatsoever; and the fact that the Government may have formulated, furnished, or in any way supplied the said drawings, specifications, or other data is not to be regarded by implication or otherwise, as in any manner licensing the holder or any other person or corporation, or conveying any rights or permission to manufacture, use, or sell any patented invention that may in any way be related thereto.

This report has been reviewed by the Information Office (OI) and is releasable to the National Technical Information Service (NTIS). At NTIS, it will be available to the general public, including foreign nations.

This technical report has been reviewed and is approved for publication.

M Edward Womble
M. EDWARD WOMBLE, Ph.D.
Project Scientist

Malcolm C. Lancaster
MALCOLM C. LANCASTER, Col, USAF, MC
Supervisor

Robert G. McIver
ROBERT G. MCIVER, Colonel, USAF, MC
Commander

ACCESSION FOR	
NTIS	White Section <input checked="" type="checkbox"/>
DIC	Buff Section <input type="checkbox"/>
DATA CHANGED	<input type="checkbox"/>
JUSTIFICATION	
BY	
DISTRIBUTION/AVAILABILITY CODES	
Dist. Avail. and/or SPECIAL	
A	

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 14 SAM-TR-76-31 ✓	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER 9
4. TITLE (and Subtitle) USAFSAM ECG/VCG DIGITIZING AND AVERAGING SYSTEM		5. TYPE OF REPORT & PERIOD COVERED Final Rept. Aug 1972 - Feb 1976
7. AUTHOR 10 M. Edward Womble, Ph.D.; Harold N. Keiser M.S.; Charles R. Meyer, Capt, USAF; John H. Triebwasser, Col, USAF, MC; and Malcolm C. Lan- caster, Col, USAF, MC		6. PERFORMING ORG. REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS USAF School of Aerospace Medicine (NG) ✓ Aerospace Medical Division (AFSC) Brooks Air Force Base, Texas 78235		8. CONTRACT OR GRANT NUMBER(s)
11. CONTROLLING OFFICE NAME AND ADDRESS USAF School of Aerospace Medicine (NG) Aerospace Medical Division (AFSC) Brooks Air Force Base, Texas 78235		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 62202F 7755-01-12
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE 11 September 1976
		13. NUMBER OF PAGES 55 (12) 58 p.
		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Digitizing, filtering, least-squares, multiplex, computers, ECG/VCG		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) A system for digitizing and filtering ECG/VCGs has been developed and is operational on a Data General Nova 1200 minicomputer. The theory and implementations of various algorithms for amplifying, digitizing, filtering, and calibrating VCGs are discussed along with the hardware needed to make them operational. Flowcharts are given of the digital filter algorithms along with the results of extensive verification tests. The system operation procedure and the two possible formats of the digitized and filtered ECG/VCGs are discussed.		

DD FORM 1473

1 JAN 73

EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

317 000

482

TABLE OF CONTENTS

	<u>Page</u>
INTRODUCTION.....	3
CALIBRATION.....	3
Analog Filter.....	4
Amplifier.....	4
Multiplexer and Digitizer.....	4
Data-Acquisition Software.....	4
Calibration-Factor Determination.....	5
DIGITAL FILTER.....	5
Slope Estimation.....	6
Baseline Removal.....	7
Fiducial Detection and Averaging.....	15
Bad-Data Rejection.....	16
Computer Program Flowchart.....	17
Program Verification.....	33
OUTPUT DATA FORMATS.....	51
RDOS File Formats.....	51
IBM 360/65 Tape Formats.....	51
DIGITIZING-SYSTEM OPERATION.....	52
Required Programs.....	52
Required Hardware.....	52
Operating Procedures.....	52
CONCLUSION.....	54
REFERENCES.....	55

LIST OF ILLUSTRATIONS

Figure

1. Baseline removal.....	8
2. The fine baseline estimate.....	15
3. The reference cardiogram.....	33
4. The results for the reference signal only.....	34

CONTENTS (Cont'd)

<u>Figure</u>	<u>Page</u>
5. The average reference cardiogram.....	35
6. Test results for reference signal corrupted by a 0.5-Hz, 0.5656-mV rms sine wave.....	36
7. Test results for reference signal corrupted by a 0.5-Hz, 0.2828-mV rms sine wave.....	37
8. Test results for reference signal corrupted by a 0.5-Hz, 0.1414-mV rms sine wave.....	38
9. The average waveforms calculated for the reference signal corrupted by a 0.5-Hz and 0.5656-mV, 0.2828-mV, and 0.1414-mV rms sine wave.....	39
10. Test results for reference signal corrupted by a 0.158-Hz, 2.121-mV rms sine wave.....	40
11. Test results for reference signal corrupted by a 0.158-Hz, 1.061-mV rms sine wave.....	41
12. Test results for reference signal corrupted by a 0.158-Hz, 0.55146-mV rms sine wave.....	42
13. The average waveforms calculated for the reference signal corrupted by a 0.158-Hz and 2.121-mV, 1.061-mV, and 0.55146-mV rms sine wave.....	43
14. Test results for reference signal corrupted by a 0.079-Hz, 4.242-mV rms sine wave.....	44
15. Test results for reference signal corrupted by a 0.079-Hz, 2.121-mV rms sine wave.....	45
16. Test results for reference signal corrupted by a 0.079-Hz, 1.06-mV rms sine wave.....	46
17. The average waveform calculated for reference signal corrupted by a 0.079-Hz and 4.242-mV, 2.121-mV, and 1.06-mV rms sine wave.....	47
18. Baseline-removal evaluation, triangular QRS wave.....	48
19. Baseline-removal evaluation, square QRS wave.....	49
20. Baseline-removal evaluation, step baseline changes.....	50

USAFSAM ECG/VCG DIGITIZING AND AVERAGING SYSTEM

INTRODUCTION

Computers are being used more and more to aid physicians in analyzing electrocardiograms. This aid ranges from filtering the cardiograms and amplifying them to enhance their presentation to the cardiologist all the way to analyzing the cardiograms and making diagnoses. The first step in using the computer to aid the cardiologist entails converting the cardiogram analog signals into digital form. Next, prior to any automated analyses, the digitized signals must be filtered, to remove noise, and stored in the correct format.

At the USAF School of Aerospace Medicine (USAFSAM), it is common practice in the Clinical Sciences Division to adjust the gain of the instrumentation amplifiers and apply a d.c. bias to maximize the range of the signal recorded on paper. This allows a larger picture for the cardiologist to examine. In the Calibration section, we discuss the equipment used in acquiring data, the equipment's characteristics, and the procedure used for determining the actual voltage level of the digitized signal. The procedure is based on passing an a.c. reference signal of known amplitude through the system after the amplifier gains are adjusted.

The procedures used for removing noises from the digitized cardiograms are discussed in the Digital Filter section. First, straightening the baseline and shifting to zero removes low-frequency noise. Afterward, the noise which is uncorrelated with itself from beat-to-beat is removed by averaging successive beats. The Digital Filter section also contains the results of verification tests that were run to prove that the algorithms were performing adequately.

Later sections provide the formats used in outputting the data, a detailed set of operating instructions, and the conclusions derived from this work for future work.

CALIBRATION

ECG/VCG data to be digitized and later processed are supplied to the signal conditioning and digitizing system from an analog instrumentation magnetic-FM-tape recorder/reproducer. Prior to 1976, the signal-to-noise ratios from ECG signals reproduced from tape were limited to 40 dB. New equipment being installed should boost signal-to-noise ratios to above 50 dB.

Analog Filter

Each of three channels of ECG data (assumed to be three orthogonal leads) from the tape reproducer is first passed through an analog 4th-order Butterworth low-pass filter, which has a d.c. gain of zero and a 100-Hz cutoff frequency. Within the passband the filters are maximally flat, while outside the passband their attenuation increases at a rate of 80 dB per decade. These filter characteristics were chosen to improve the signal-to-noise ratio from the tape reproducer and guarantee satisfaction of the Nyquist sampling criterion for sampling rates as low as 200 Hz. While tapes recorded from our VCG sources (Hewlett-Packard, VCG Programmer, model 1507C) may contain information in the 0- to 200-Hz band because of the cutoff frequencies of the preamplifiers, tapes recorded from the treadmill ECG source (Hewlett-Packard, Pre-amplifier, model 8811) contain information in only the 0- to 100-Hz band. In addition, it is doubtful that surface ECG data contain significant information in frequency bands above 100 Hz [5].

Amplifier

Each filtered output is amplified by a factor of 3 and presented to the input of the analog multiplexer. A gain of 3 was chosen to maximize the quantization accuracy without driving the multiplexer into its nonlinear region. The range of the output of the analog tape recorder is ± 1.4 V, while the range of the input to the multiplexer is ± 5 V.

Multiplexer and Digitizer

The multiplexer, model MP4110, and the analog-to-digital converter, model MP2812, are both products of Analogic Corporation. The converter provides 12 bits of digital output within an absolute accuracy of $\pm 0.02\%$ of full scale. Total conversion time for one channel is 4 μ sec. Word output format is 2's-complement to match the internal arithmetic conventions of the Nova minicomputer/controller.

Data-Acquisition Software

Data acquisition was performed with a FORTRAN callable assembly language subroutine which utilizes the Data General Real-Time Disk Operating System (RDOS). The data were digitized and stored temporarily on a mass-storage device. The sampling rate and duration, the number of channels to digitize, and the identification of the output file are all defined in the calling sequence of this subroutine. These features allow the use of different sampling rates and acquisition-period lengths for the calibration and vectorcardiogram data. The calibration signal was digitized at 125 samples/sec for 6 seconds, while the vectorcardiogram was digitized at 500 samples/sec for 12 seconds.

Another feature of this subroutine is limit checking. To detect overscaling in either direction, a window slightly smaller than the range of the analog-to-digital converters was defined. Each digitized value was checked. When a value was outside this window, an element of an array defined in the calling sequence for this purpose was incremented. This array is assumed to two elements for each channel digitized. The first of the pair counted the number of values exceeding the positive limit; and the second, the number of values below the negative limit.

Calibration-Factor Determination

After the biases and amplifier gains are set, a pulse train is passed through the system and recorded. The peak-to-peak variations of the pulses are 0.5 mV, and the pulse rate is 30 pulses/min. These pulses are used to calibrate the signals.

The calibration process requires two passes over the data. The first pass defines the midrange for each channel as the mean of the maximum and minimum values for the channel. The lower value of the calibration waveform is determined by the second pass and is defined as the mean of all values below midrange. Because the analog filters cause overshoot on the leading edge of the calibration pulse and because of the low duty cycle, this technique would not produce a reliable value for the upper limit. To determine the upper value, the distance between midrange and maximum is divided into 50 segments and the number of points in each segment is determined. These frequencies are examined, and the upper limit is defined as the midvalue of the three consecutive segments whose sum frequency is the largest for each channel. The calibration factor for each channel is then defined as twice the difference of the upper and lower limits for that channel.

DIGITAL FILTER

Jointly, USAFSAM and the Charles Stark Draper Laboratory have developed a system which filters and averages electrocardiograms and vectorcardiograms (ECG/VCG). The system has been coded in ANSI Fortran and implemented on a Data General Nova 1230 digital computer. It accepts three separate signals, each sampled at a rate of 500 samples/sec, straightens and shifts the baseline to zero, and averages successive beats to produce a representative beat for each signal. The slope of the baseline is continuously monitored, and beats corresponding to a baseline slope greater than 2 mV/sec are not used in determining the average beat.

Slope Estimation

At several stages in the digital filter, curves are approximated by a straight line. For instance, in the Baseline Removal section, a straight line will be fitted to 11 selected points, each 80 msec apart. The midpoint of the straight line will be used to estimate the baseline of the incoming signal. In the Fiducial Detection and Averaging section, a straight line will be fitted to 7 consecutive samples, each 2 msec apart. The slope of the straight line will be used as an estimate of the derivative of the incoming signal at the time corresponding to the midpoint of the 7 samples. Finally, in the Computer Program Flowchart section, a straight line will be fitted to 11 consecutive points, each 80 msec apart. The slope of the line will be used to determine the quality of the incoming data.

In all of the cases mentioned, straight lines are being fitted to more than two points; therefore, since two points determine a straight line, the fit cannot be perfect.

The curves are approximated by the straight line that minimizes the sum of the square of the error in the fits at each point (least-square error [9]). Since the same approach will be used several times, a brief description of it follows.

The problem is to find the straight line

$$\hat{x}(k) = b + ks \quad (1)$$

which minimizes the sum of the square of the error in the fits at N points

$$x(k), \left[k = -\frac{N-1}{2}, -\frac{N-1}{2} + 1, \dots, -1, 0, 1, \dots, \frac{N-1}{2} \right]$$
$$J = \sum_{k=-\frac{N-1}{2}}^{\frac{N-1}{2}} [\hat{x}(k) - x(k)]^2 \quad (2)$$

The reasons for the peculiar and clumsy-appearing indexing of time are to simplify the calculations and to make b the estimate of the straight line at the time-interval midpoint ($k = 0$). The solution to this problem is [9]

$$b = \frac{1}{N} \sum_{k=-\frac{N-1}{2}}^{\frac{N-1}{2}} x(k) \quad (3)$$

$$s = \frac{\sum_{k=1}^{N-1} k^2 \frac{x(k) - x(-k)}{2k}}{\sum_{k=1}^{N-1} k^2} \quad (4)$$

Observe that the slope of straight line s is the weighted sum of the differences in the last and the first points, the next-to-last and the second points, etc. Also, observe that b is simply the average value of the points.

Baseline Removal

Precisely defining the baseline is a very difficult task. Some say that it is the voltage measured between either the T and P waves or, if a U wave is detected, between the U and the P waves. Others define the baseline to be somewhere on the P-R interval, which in many cases is sloped. The problem is further complicated by the fact that the electric potential measured at either of the points fluctuates as a result of both physiological events and the electrical instrumentation. In addition, while recording cardiograms, we normally add a d.c. bias and adjust the amplifier gain so that the cardiogram fills up the allowed space on a strip chart.

In making amplitude measurements, we use the baseline as the reference; therefore, the computer must locate the baseline to make amplitude measurements. The approach used with the digital filter is to define the baseline as a weighted average of the P-R and T-P intervals, and to make the baseline potential be 0 mV, which will simplify subsequent amplitude measurements.

The traditional technique for straightening and shifting the baseline to zero has been to use a high-pass filter, which removes the low-frequency baseline (<0.2 Hz) while passing the heart signal (0.2 Hz < heart signal < 100 Hz) [7]. This approach has worked fairly well in the past when the subsequent analysis was performed by physicians; however, filters are not perfect, and the resulting distortion of the heart signal itself is easily detectable when either the cardiogram display is enlarged or computer analysis is performed. Because of this problem and because of the availability of a digital computer, a more sophisticated approach is used that straightens and removes the baseline without changing the heart signal itself.

The baseline-removal strategy is based on the idea of removing the P, T, and QRS waves from the signal and approximating the baseline as a

straight line fitted to the remaining data. The baseline estimate is then subtracted from the original signal, as shown in Figure 1. The P, T, and QRS waves are eliminated from the signal as follows. The signal is partitioned into 80-msec segments (40 samples at 500 samples/sec), and the instantaneous slope is checked at each point. If at any point in an 80-msec segment the instantaneous slope is greater than 10 mV/sec, then that segment of data is eliminated. If it is not, the value of the signal at the end of the 80-msec segment is recorded and used to calculate the baseline. For time segments which contain either QRS complexes or T and P waves, the value of the signal saved for the previous time segment is recorded and used in estimating the baseline.

After 11 values of the signal are recorded (800 msec), the baseline is estimated at the midpoint of the 800-msec interval of data by fitting a straight line to the 11 recorded points. The oldest data point is then discarded, a new 80-msec segment of data is checked, and either a new point is recorded or the last point is shifted down. The estimate of the baseline at the midpoint of the shifted segment of data is then calculated and the process repeated. For the 80 msec between the estimates, linear interpolation is used to estimate the baseline.

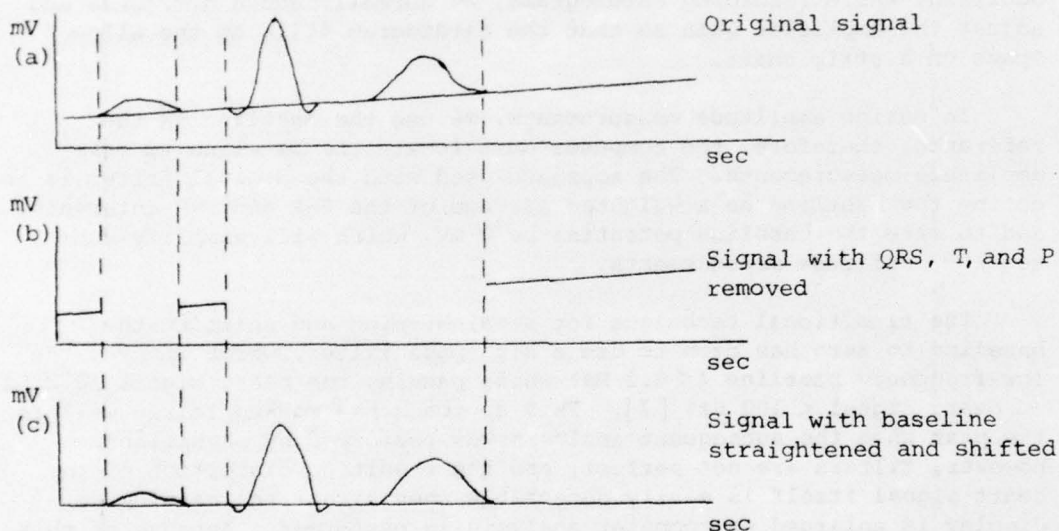



Figure 1. Baseline removal.

The straight line that is fitted to the 11 points (80-msec segments, 800-msec total) is of the form

$$x(k) = b + ks$$

where k refers to the 11 samples (k = -5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5).




As indicated in the Slope Estimation section, this particular indexing of time makes the estimate of the baseline at the midpoint of the 800-msec window of data be

$$b = \sum_{k=-5}^5 x(k) \quad (5)$$


To better understand how the baseline-removal logic is implemented, consider the following typical case.

Situation 1

- a. Array W2 contains data with baseline removed and is 600 samples long.
- b. Array ST contains the latest data received and is 200 samples long.
- c. Array SAV contains the last 40 samples read into W2.



In the above diagram, W2 is an array of 600 data points



$$\begin{array}{l}
 \text{W2} \\
 = \\
 \left[\begin{array}{c}
 \text{W2 (0)} \\
 \text{W2 (1)} \\
 \text{W2 (2)} \\
 \cdot \\
 \cdot \\
 \cdot \\
 \text{W2 (598)} \\
 \text{W2 (599)}
 \end{array} \right] \\
 = \\
 \left[\begin{array}{c}
 \text{x (t}_0\text{)} \\
 \text{x (t}_1\text{)} \\
 \text{x (t}_2\text{)} \\
 \cdot \\
 \cdot \\
 \cdot \\
 \text{x (t}_{598}\text{)} \\
 \text{x (t}_{599}\text{)}
 \end{array} \right]
 \end{array} \quad (6)$$

ST is an array of 200 data points


$$\begin{array}{l}
 \text{ST} \\
 = \\
 \left[\begin{array}{c}
 \text{ST (0)} \\
 \text{ST (1)} \\
 \cdot \\
 \cdot \\
 \cdot \\
 \cdot \\
 \text{ST (198)} \\
 \text{ST (199)}
 \end{array} \right] \\
 = \\
 \left[\begin{array}{c}
 \text{x (t}_{600}\text{)} \\
 \text{x (t}_{601}\text{)} \\
 \cdot \\
 \cdot \\
 \cdot \\
 \cdot \\
 \text{x (t}_{798}\text{)} \\
 \text{x (t}_{799}\text{)}
 \end{array} \right]
 \end{array} \quad (7)$$

and SAV is an array of 40 data points

$$\begin{array}{l}
 \text{SAV} \\
 = \\
 \left[\begin{array}{c}
 \text{SAV (0)} \\
 \text{SAV (1)} \\
 \cdot \\
 \cdot \\
 \cdot \\
 \cdot \\
 \text{SAV (38)} \\
 \text{SAV (39)}
 \end{array} \right] \\
 = \\
 \left[\begin{array}{c}
 \text{x (t}_{560}\text{)} \\
 \text{x (t}_{561}\text{)} \\
 \cdot \\
 \cdot \\
 \cdot \\
 \cdot \\
 \text{x (t}_{598}\text{)} \\
 \text{x (t}_{599}\text{)}
 \end{array} \right]
 \end{array} \quad (8)$$


Step 1

Move SAV ahead 40 samples:



Step 2

Move ST ahead 40 samples:



Step 3

Using data through 839th sample, calculate the baseline at the 639th sample.

Step 4

Using present and previous (at 599th sample) baseline estimates and linear interpolation, remove the baseline from data in SAV and move W2 ahead 40 samples

40	W2	639	640	ST	839
600	SAV	639			



The 11 data points needed to estimate the baseline at the 639th sample are stored in the arrays WYR, ZYREM, and SAV. ZYREM is an array of 6 data points, and WYR an array of 5. The latter two arrays are operated as a pushdown register. At the beginning of the last example, every 40th point, beginning with the 399th sample, was stored in WYR and ZYREM.

$$\begin{array}{|l}
 \text{WYR(5)} \\
 \text{WYR(4)} \\
 \text{WYR(3)} \\
 \text{WYR(2)} \\
 \text{WYR(1)} \\
 \text{ZYREM(6)} \\
 \text{ZYREM(5)} \\
 \text{ZYREM(4)} \\
 \text{ZYREM(3)} \\
 \text{ZYREM(2)} \\
 \text{ZYREM(1)}
 \end{array}
 =
 \begin{array}{|l}
 x(t_{799}) \\
 x(t_{759}) \\
 x(t_{719}) \\
 x(t_{679}) \\
 x(t_{639}) \\
 x(t_{599}) \\
 x(t_{559}) \\
 x(t_{519}) \\
 x(t_{479}) \\
 x(t_{439}) \\
 x(t_{399})
 \end{array}
 \tag{9}$$

In terms of the previous diagrams, we have the following situation.

Situation 2

Same as situation 1 except the arrays ZYREM and WYR have been added. Also the baseline estimate is known at the 599th sample.




Step 1

Move SAV ahead 40 samples (see situation, step 1):

Step 2

Store ZYREM(1) and move ST, WYR, and ZYREM ahead 40 samples.




Step 3

Using data through the 839th sample, calculate the baseline at the 639th sample.

Step 4

Using present and previous baseline estimates and linear interpolation, remove the baseline from data in SAV and move W2 ahead 40 samples.



At the third step, the data needed to estimate the baseline (see equation 5) at the 639th sample are stored in the arrays ZYREM and WYR:

$$b(639) = \frac{1}{11} \left\{ \sum_{i=1}^5 [ZYREM(i) + WYR(i)] + ZYREM(6) \right\} \quad (10)$$

The data needed to estimate the baseline at the 599th sample (40 samples ago) were stored in the arrays ZYREM and WYR (see equation 9)

$$b(599) = \frac{1}{11} \left\{ \sum_{i=1}^5 [ZYREM_{-1}(i) + WYR_{-1}(i)] + ZYREM_{-1}(6) \right\} \quad (11)$$

The subscript "-1" indicates the values of ZYREM and WYR at step 1. In going from step 1 to step 2, the data points in ZYREM and WYR are pushed down, $x(t_{399})$ is stored as ZSAV1, and $x(t_{839})$ is added; now ZYREM and WYR contain the following points (compare equations 12 and 9):

$$\begin{bmatrix} WYR(5) \\ WYR(4) \\ WYR(3) \\ WYR(2) \\ WYR(1) \\ ZYREM(6) \\ ZYREM(5) \\ ZYREM(4) \\ ZYREM(3) \\ ZYREM(2) \\ ZYREM(1) \end{bmatrix} = \begin{bmatrix} x(t_{839}) \\ x(t_{799}) \\ x(t_{759}) \\ x(t_{719}) \\ x(t_{679}) \\ x(t_{639}) \\ x(t_{599}) \\ x(t_{559}) \\ x(t_{519}) \\ x(t_{479}) \\ x(t_{439}) \end{bmatrix} \quad (12)$$

The new sample, $x(t_{839})$, is stored in WYR(5) if, and only if, the data from t_{799} to t_{839} pass certain tests indicating that neither a P, T, nor QRS wave occurs during that 80-msec time interval. Finally, equations 10 through 12 yield the following recursive formula for estimating the baseline:

$$b(k) = b(k-40) + \frac{1}{11} (ZYREM(1) - ZSAV1) \quad (13)$$

The baseline estimate is initialized via equation 10, and determined by equation 13. Equation 10 requires 10 additions and 1 division, while equation 13 requires only 2 additions and 1 division.

There are two tests for determining whether or not an 80-msec interval of data contains a P, QRS, or T wave: (1) The maximum negative

slope of the R wave in the x lead is used as a fiducial point in performing signal averaging (see Digital Filter section). Therefore, since the fiducial point is available, S-T intervals are eliminated by rejecting from the baseline estimation all data occurring during the 100 msec following a fiducial point. (2) The instantaneous slope is checked at every point in the 80-msec interval. If that slope exceeds 6 mV/sec at any point in the interval, it is concluded that either a P, QRS, or T wave occurred during that interval. The 6-mV/sec threshold was determined experimentally and is valid only with the 7-point slope estimate of equation 4, which can be rewritten as

$$s(k) = .643 \left[\frac{x(k+3) - x(k-3)}{6} \right] + .286 \left[\frac{x(k+2) - x(k-2)}{4} \right] + .071 \left[\frac{x(k+1) - x(k-1)}{2} \right] \quad (14)$$

Each of the three terms enclosed by brackets in equation 14 is an estimate of the slope of the waveform. The actual estimate of the slope at the kth data point is merely the weighted average of the three estimates, with the weightings proportional to the time intervals involved. For instance, the first difference, stretching over six samples (12 msec), is given the heaviest weighting (.643). These weightings coupled with the 6-mV/sec threshold are important in distinguishing the P, QRS, and T waves from signal noise.

The actual version of equation 4 used in the computer program is recursive in nature and requires fewer additions and multiplications. The version used follows from comparing equation 15

$$s(k+1) = .643 \left[\frac{x(k+4) - x(k-2)}{6} \right] + .286 \left[\frac{x(k+3) - x(k-1)}{4} \right] + .071 \left[\frac{x(k+2) - x(k)}{2} \right] \quad (15)$$

with equation 14, which yields

$$s(k+1) = s(k) - \left\{ \Sigma(k) + 3 [x(k+4) + x(k-3)] \right\} / 28 \quad (16)$$

where Σ is the solution of


$$\Sigma(k+1) = \Sigma(k) + x(k+4) - x(k-2) \quad (17)$$

The slope estimate is initiated via equation 14 and then is propagated via equation 16. The sum Σ , used in equation 16, is initiated via

$$\Sigma(k) = x(k+3) + x(k+2) + x(k+1) + x(k) + x(k-1) + x(k-2)$$

and is then propagated via equation 17.

Up to this point, we have presented an algorithm which estimates the values of the baseline at every 40th sample (every 80 msec).



To subtract the nonzero baseline from the original signal, we need an estimate of the baseline at every sample (every 2 msec). This is generated by fitting a straight line between each pair of points.




Figure 2. The fine baseline estimate.

Finally, the signal baseline is shifted to zero by subtracting the fine baseline estimate from the original cardiogram.

Fiducial Detection and Averaging

To reduce the quantity of data to be subsequently analyzed and to remove the noncardiac-generated noise that is uncorrelated with itself from one beat to the next, the digital filter will average successive beats and produce a single representative beat for subsequent analysis. This approach has been successfully used by others [1, 2, 4, 8, 10]. The steps required are to locate each beat, align them temporally, add the beats together, and divide the resulting sum by the number of beats used. The point used in locating each beat is the point on the x lead corresponding to the maximum absolute value of its slope or time derivative. This fiducial point has been successfully used by others [1, 4] and is relatively insensitive to baseline disturbances.

The instantaneous time derivative estimate used in the baseline-removal algorithm (equations 16 and 17) is also used to detect fiducial points. The algorithm for detecting these points first determines the maximum absolute value of the slope of the x lead in the first 1.6 sec of data. This maximum value is used to determine a threshold, which is set at 60% of the max value. After the threshold is initialized, the slope is calculated and compared to the threshold as the data are received. After the threshold is exceeded, the point of maximum slope is used as the fiducial point; the threshold is then reset to 60% of the maximum value just determined, and the fiducial point detector is deactivated for the next 400 msec to prevent double-peaked R waves from becoming a problem.

Bad-Data Rejection

Under certain conditions, it is impossible to extract the actual cardiograms from the noise. Although it does not cover all cases, the criterion used to reject or accept data is based on the slope of a pseudobaseline estimate. This estimate is determined as was the actual baseline estimate with the exception that points occurring during the P, QRS, and T waves are used. Any time that the slope of this straight line exceeds 2 mV/sec, the data are not used in calculating the average waveform.

The pushdown register used to store the data needed to calculate the pseudobaseline (11 successive points, each 80 msec apart) is composed of the arrays YREM and ST.

YREM(1)
YREM(2)
.
.
.
YREM(6)
ST(39)
ST(79)
.
.
.
ST(199)

With $N = 11$, equation 4 can be manipulated (as in the Baseline-Removal section) to obtain

$$\tilde{s}(k+1) = \tilde{s}(k) - \left\{ \tilde{\Sigma}(k)+5 [x(k+6) + x(k-5)] \right\} / 110 \quad (18)$$

where

$$\tilde{\Sigma}(k+1) = \tilde{\Sigma}(k) + x(k+6) - (k-4) \quad (19)$$

Computer Program Flowchart

A list of the key variables and a flowchart of the digital filter program follow.

Variables

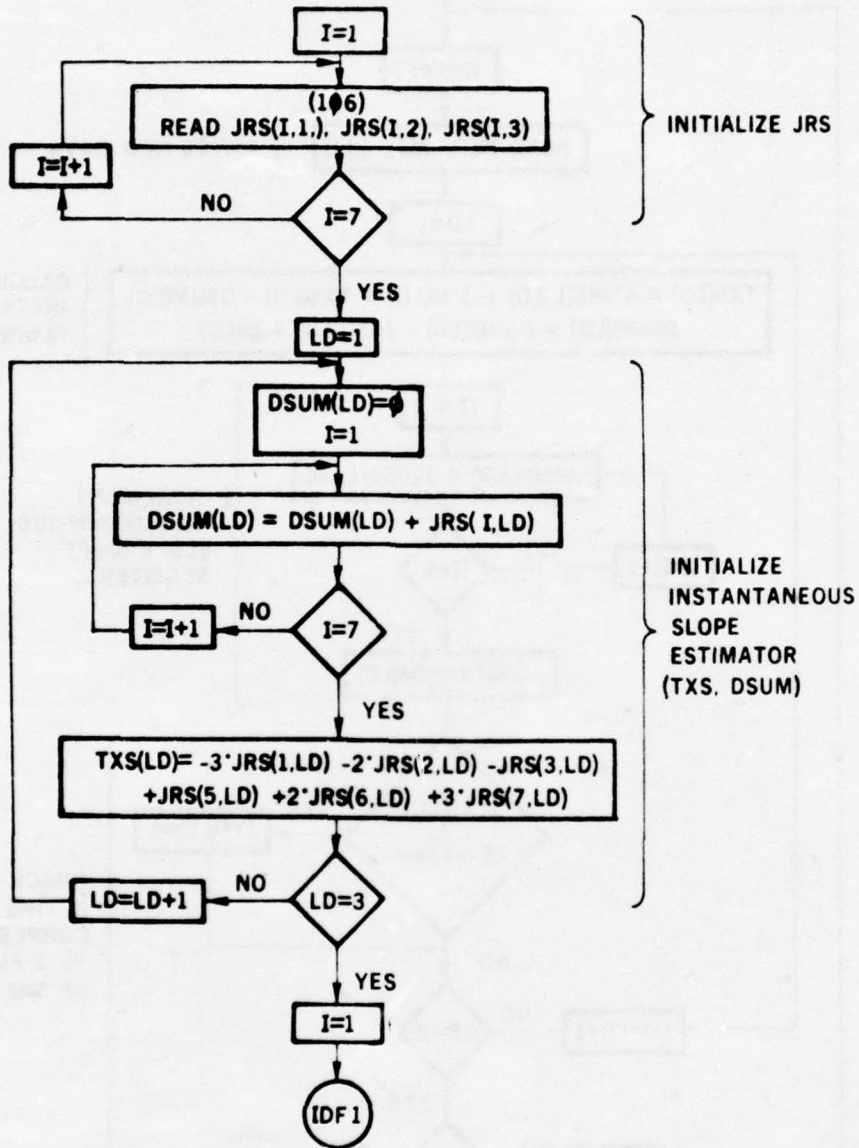
- TREM - Variable used with the peak-detection scheme.
- LC - Counter activated after fiducial has been located. Used to disable fiducial detection for 65 samples.
- ITR - Flag which is set to 1 when instantaneous slope is less than threshold (fiducial detector is enabled). Reset to 0 when fiducial is detected.
- K - Dummy variable used in initialization of average slope estimator.
- JF - Counter activated when average slope exceeds 2 mV/sec. Stops averaging for 100 samples.
- IR(10,3) - Counters activated when fiducial point is detected.
- IRC - Counts the number of fiducials detected.
- ZYSUM(3) - Sum of 11 points, once every 20 samples and with QRS points excluded.
- YSUM(3) - Sum of 11 points, once every 20 samples with no points excluded.
- TYSUM(3) - Numerator of least-squares estimate of slope, extends over 200 samples.
- NW(3) - Counter used to record number of beats used to form average beat.
- NR(3) - Counters used to determine number of samples between average waveform update.
- JYF(3) - Flags set when instantaneous slope exceeds 10 mV/sec. Used to exclude data from baseline estimate.
- W2(299,3) - Contains waveforms with baseline removed.
- AV(299,3) - Contains averaged waveforms.
- THRESH - Threshold used in peak-detection logic.
- TXS(3) - Instantaneous V slope estimate for each lead.

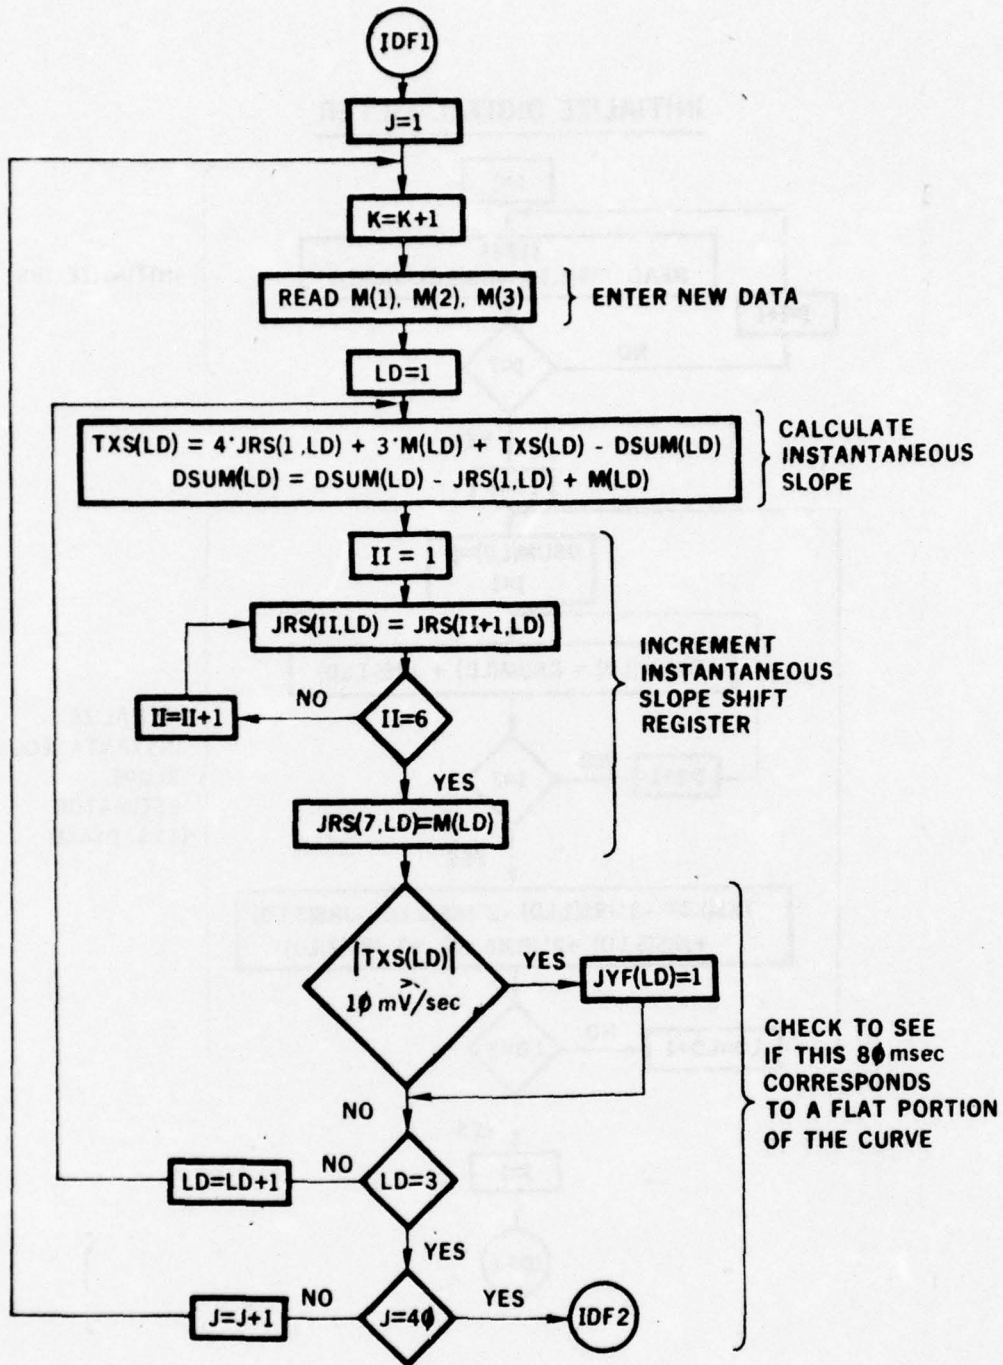
ZYREM(6,3) - Array used with baseline estimator. Points with bad baselines are removed.

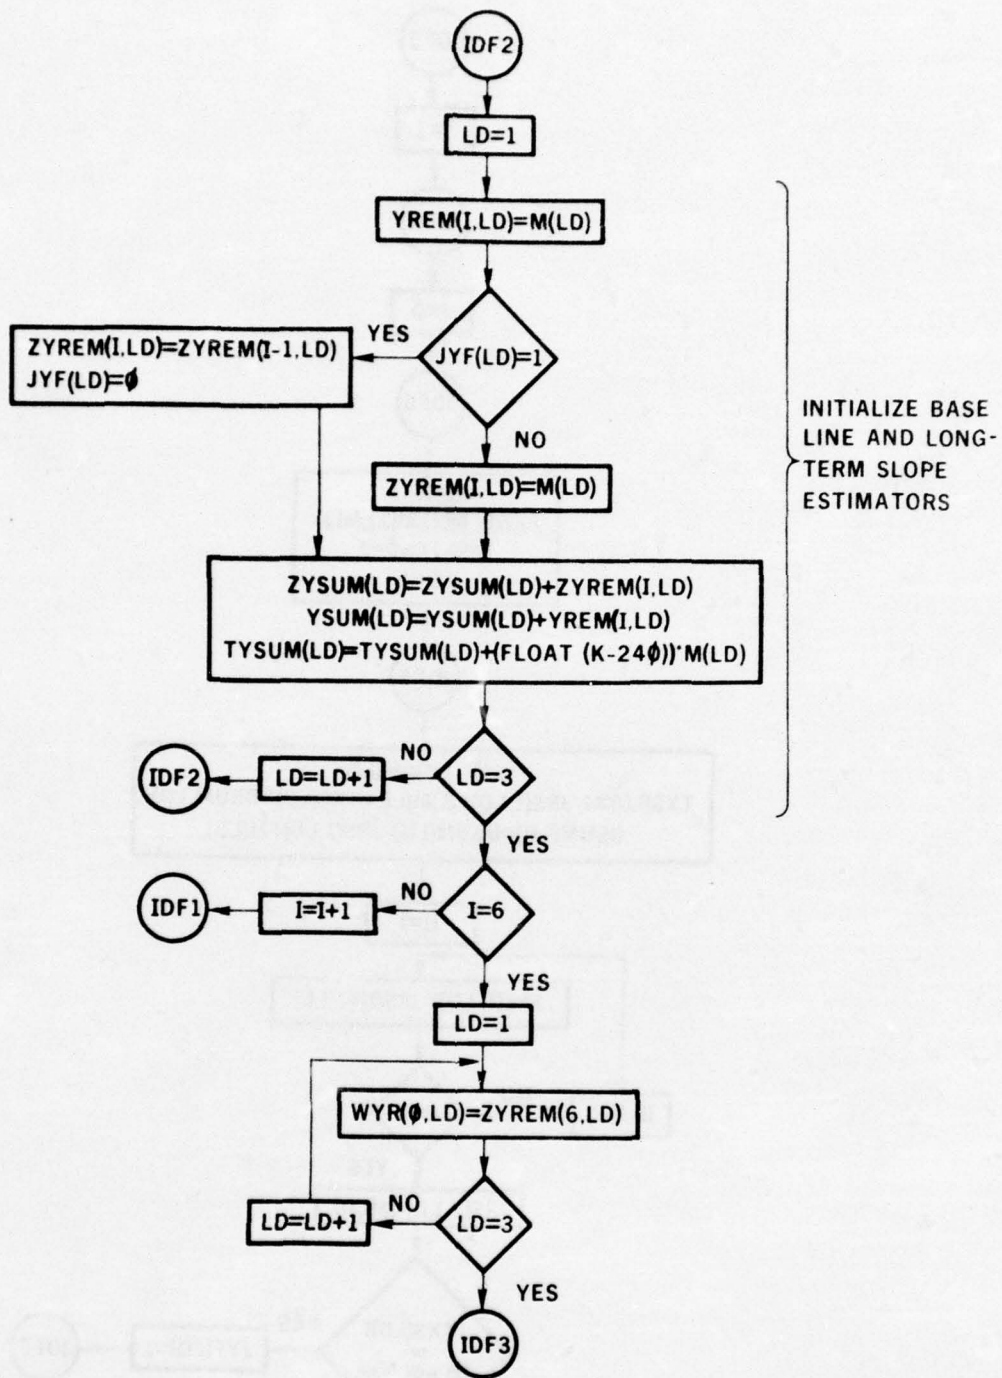
IRC - Number of beats detected.

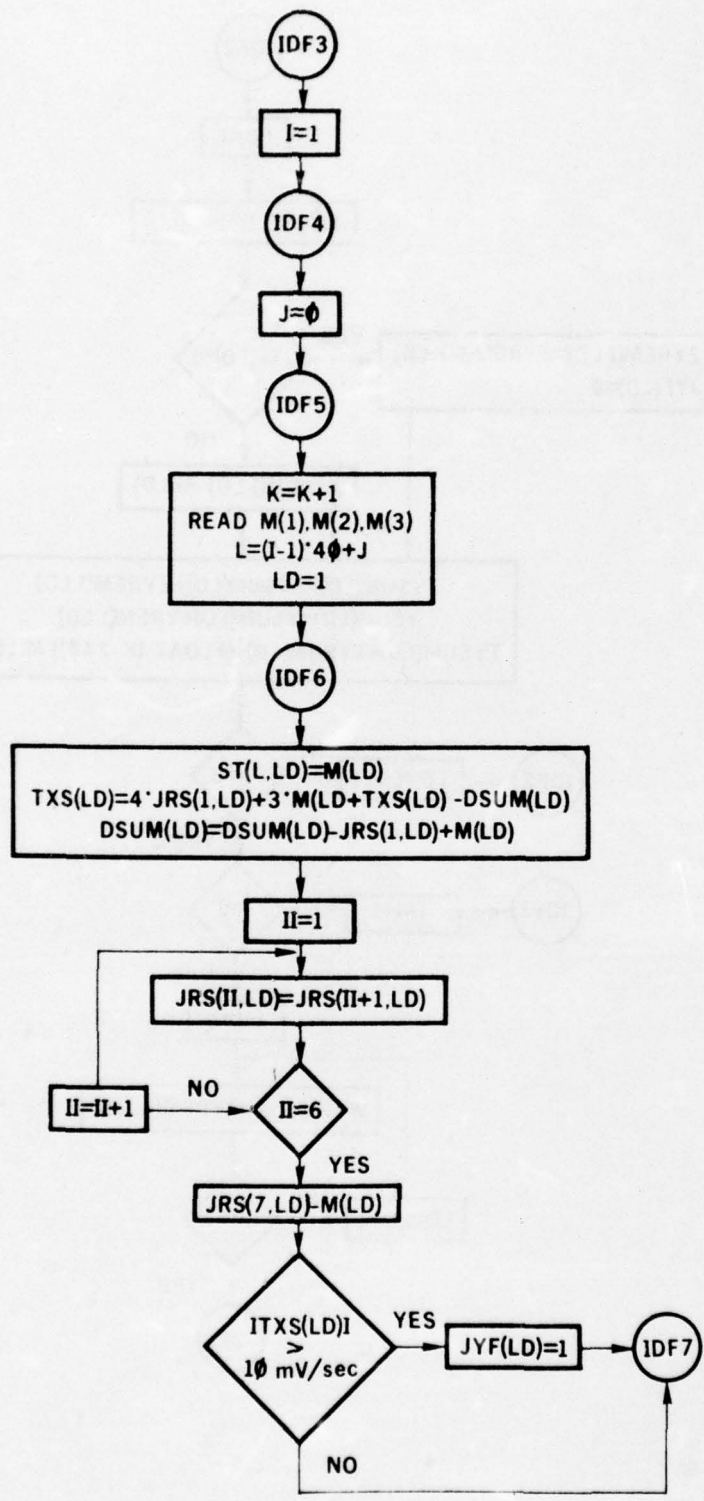
YREM(6,3) - Array used with average-slope estimator. Contains all points--none are rejected.

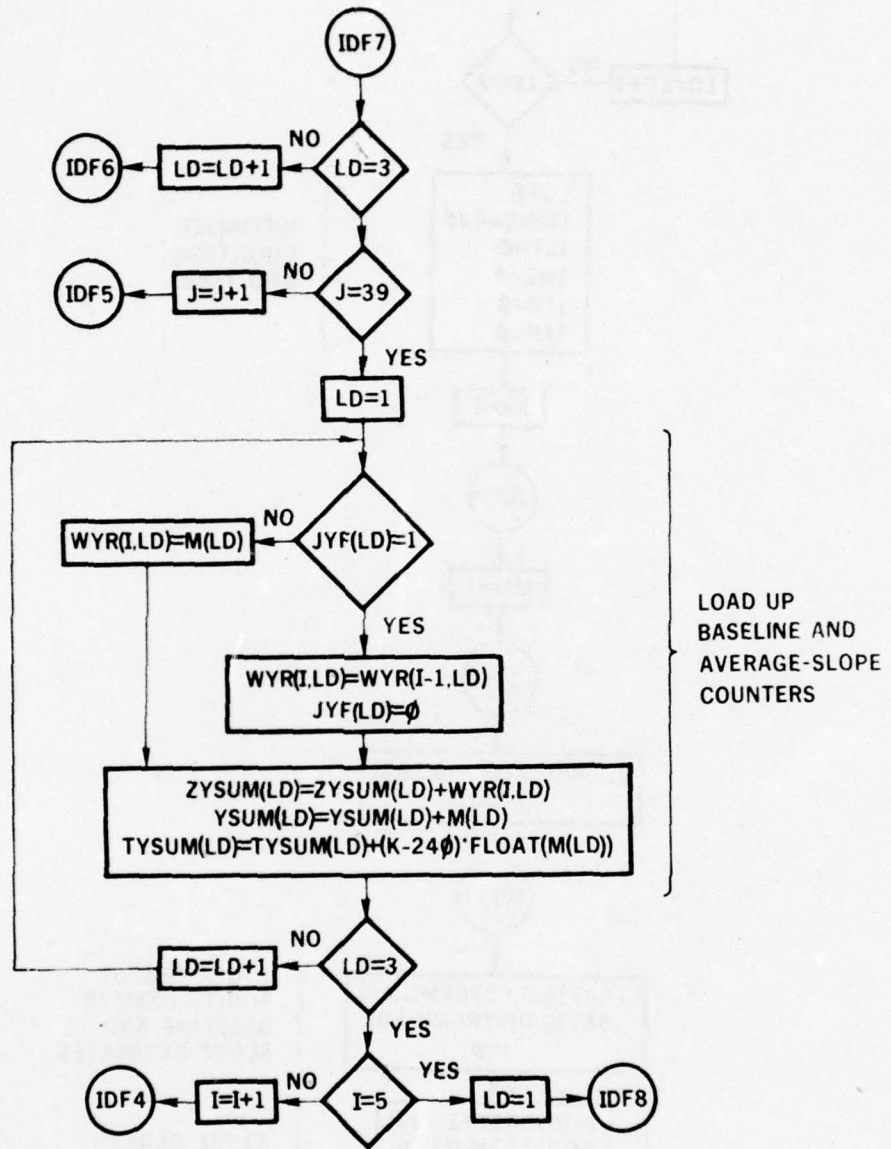
INITIALIZE DIGITAL FILTER

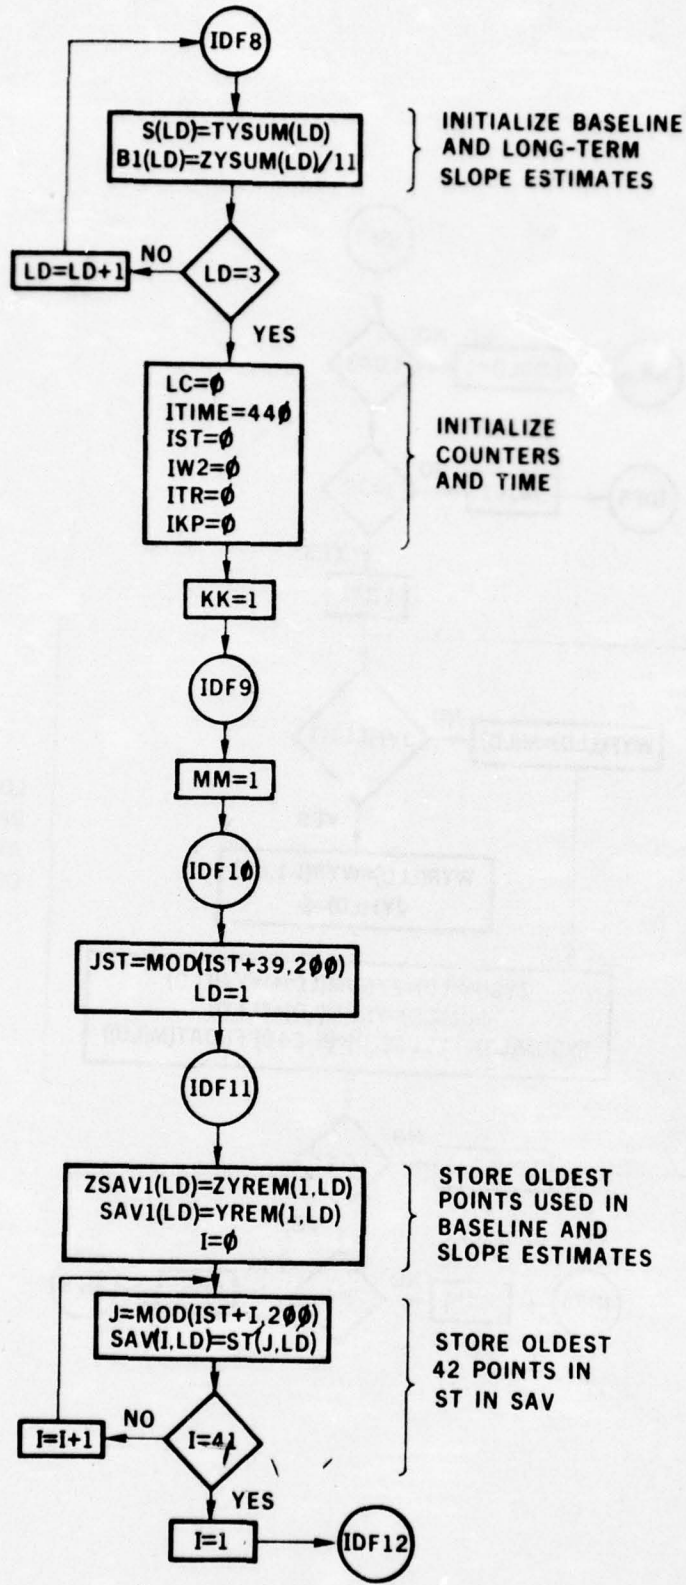


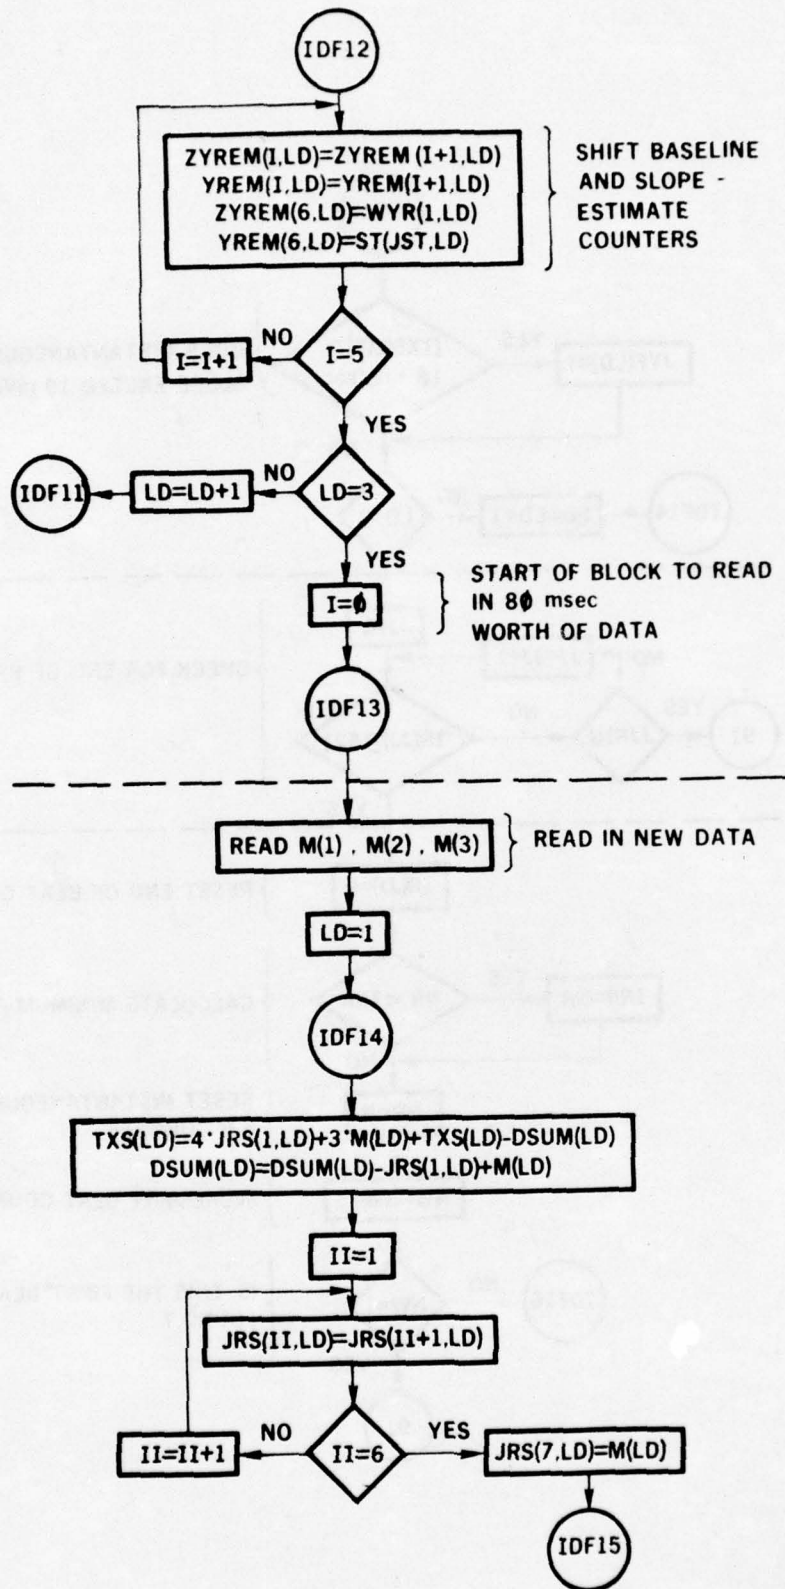


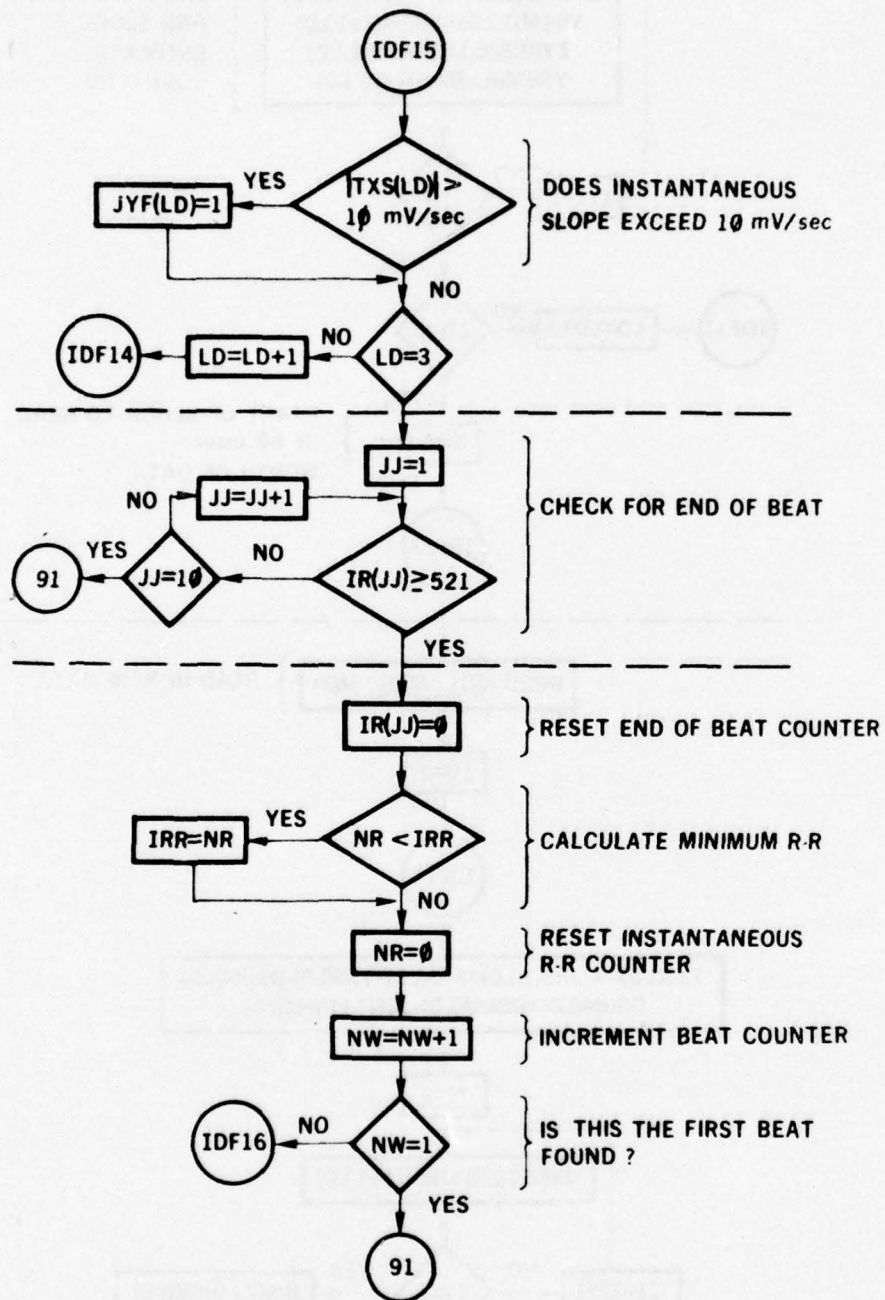


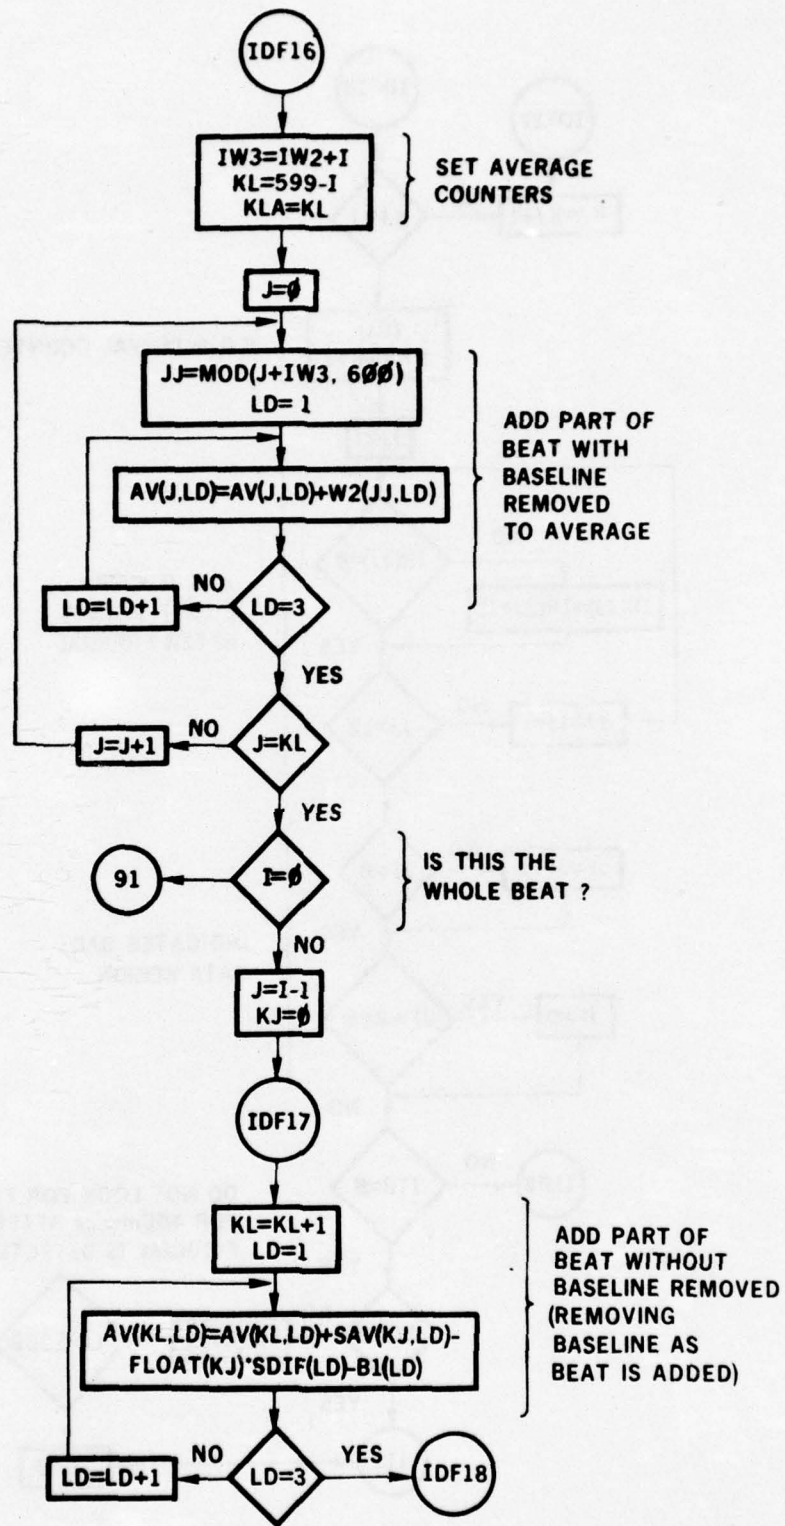


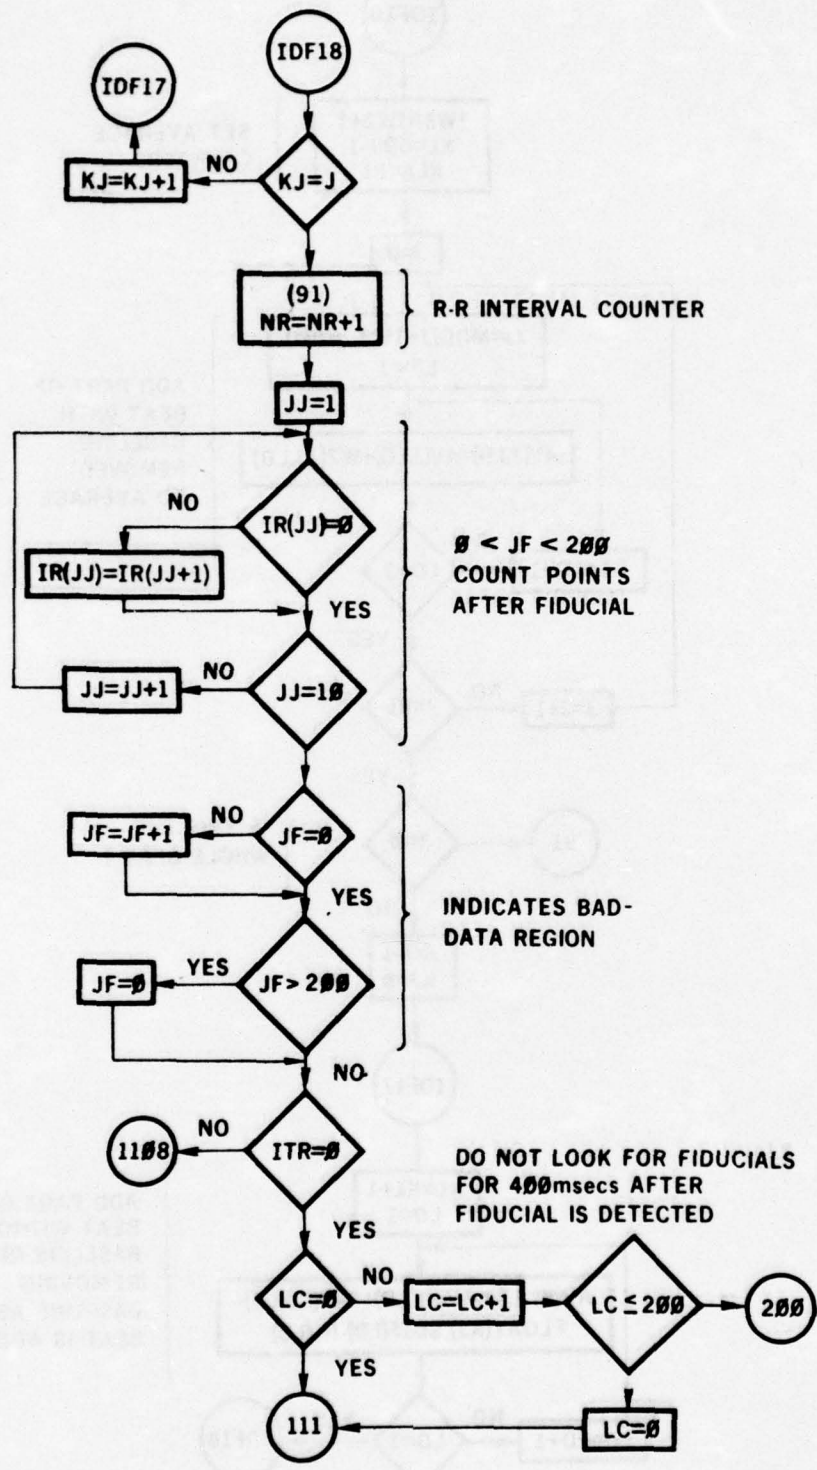


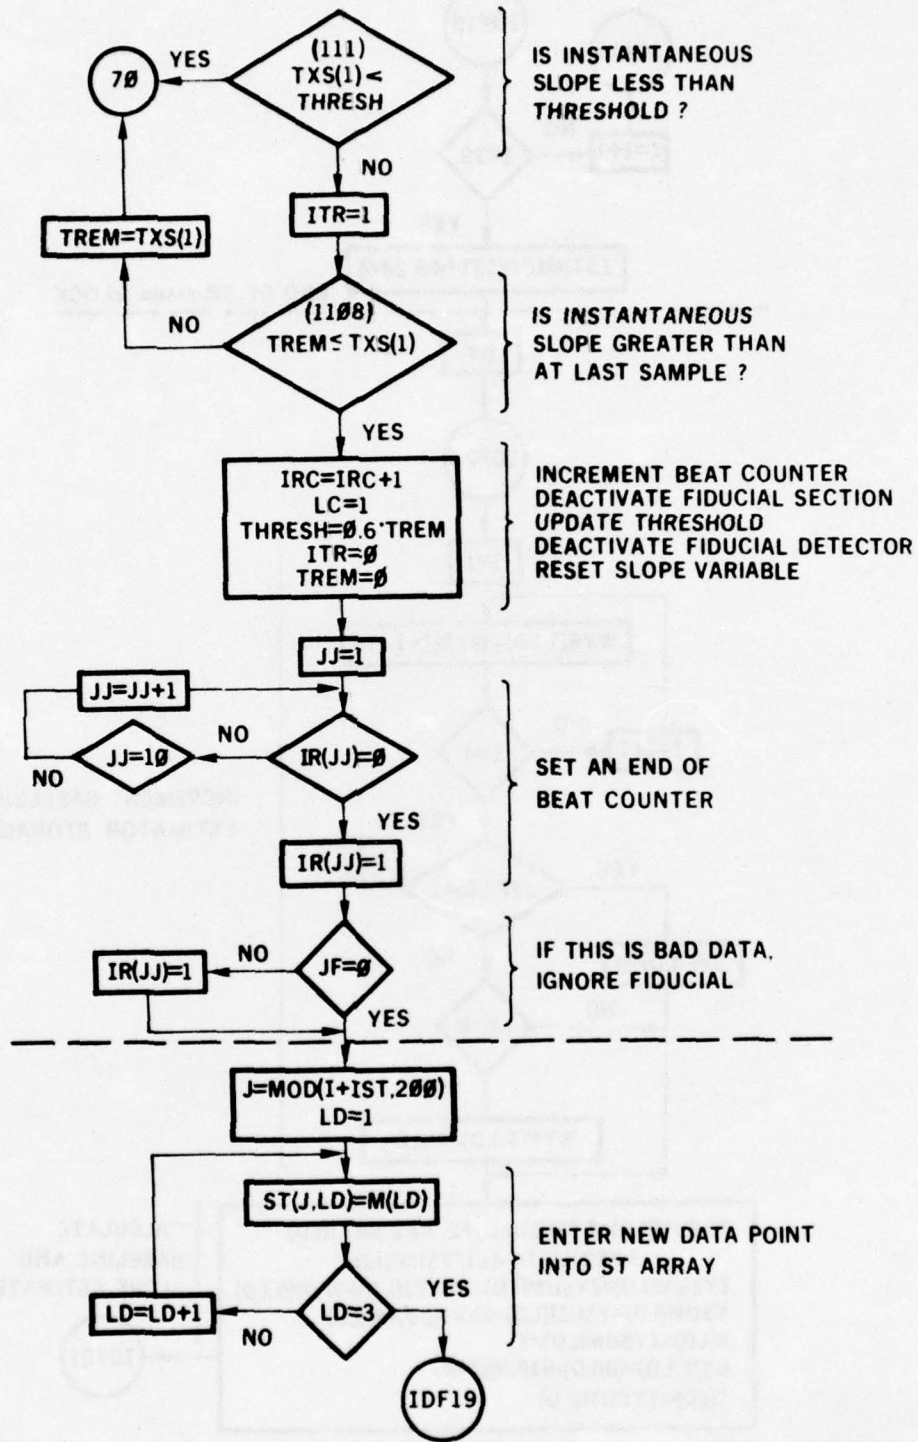


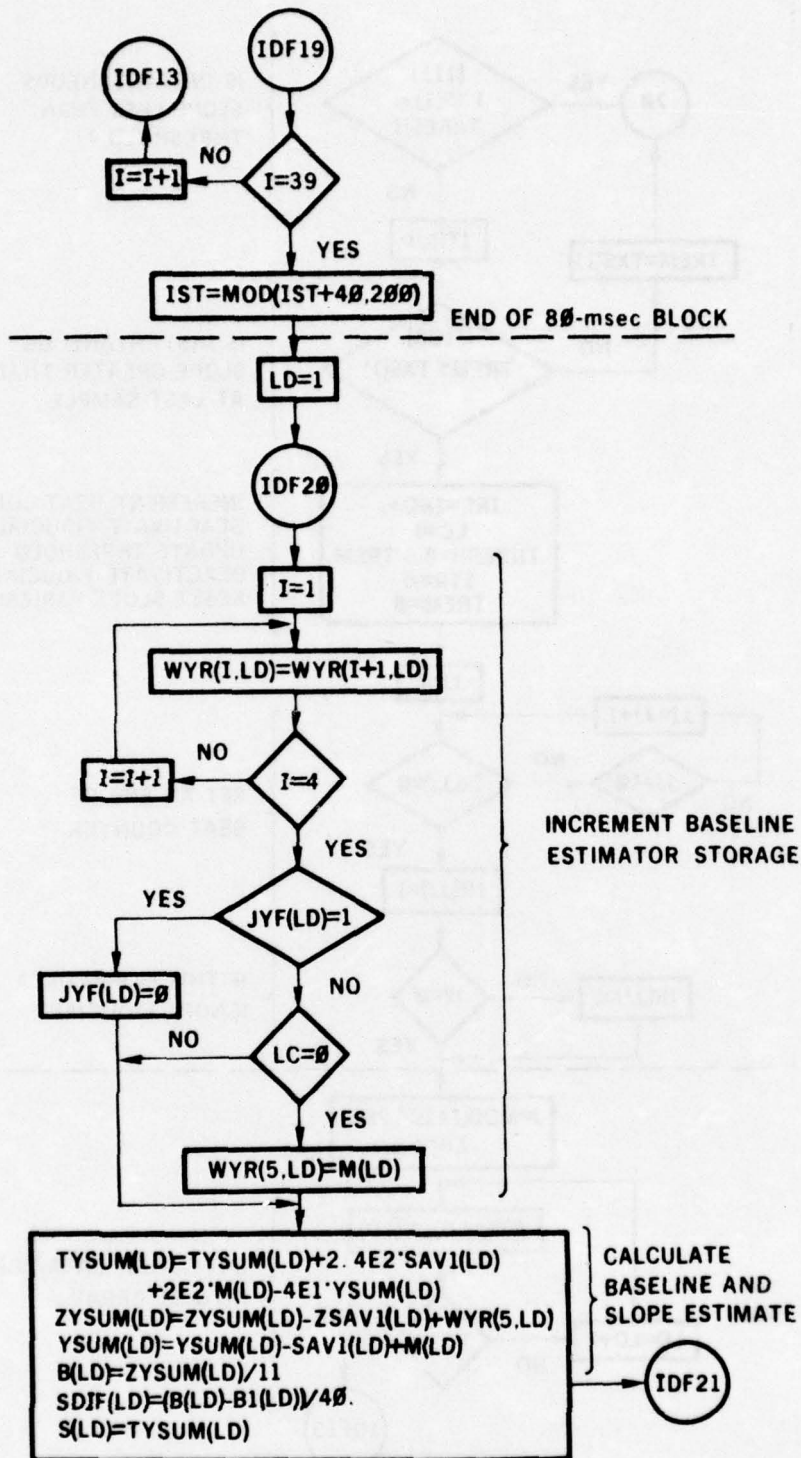


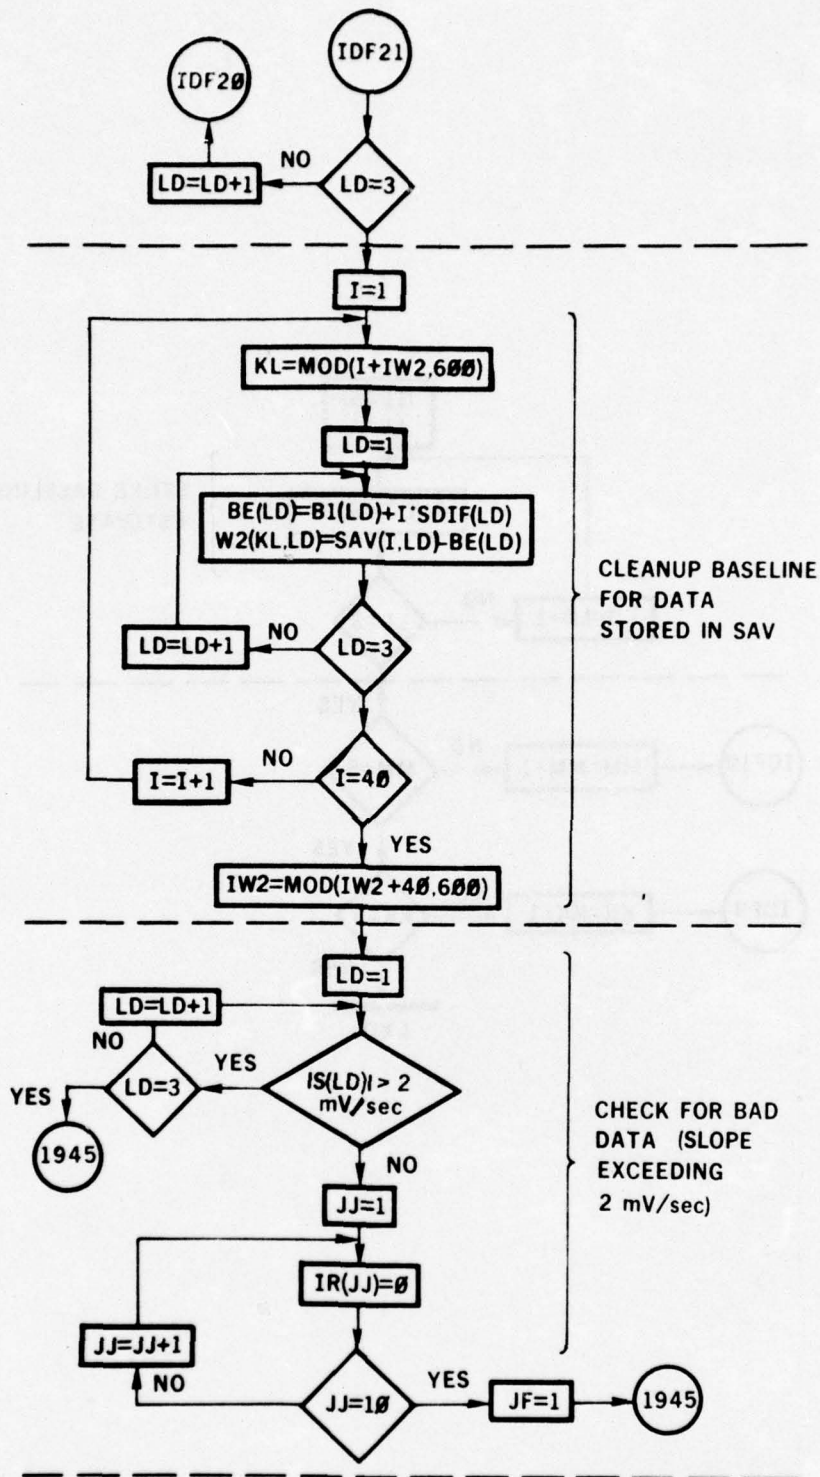


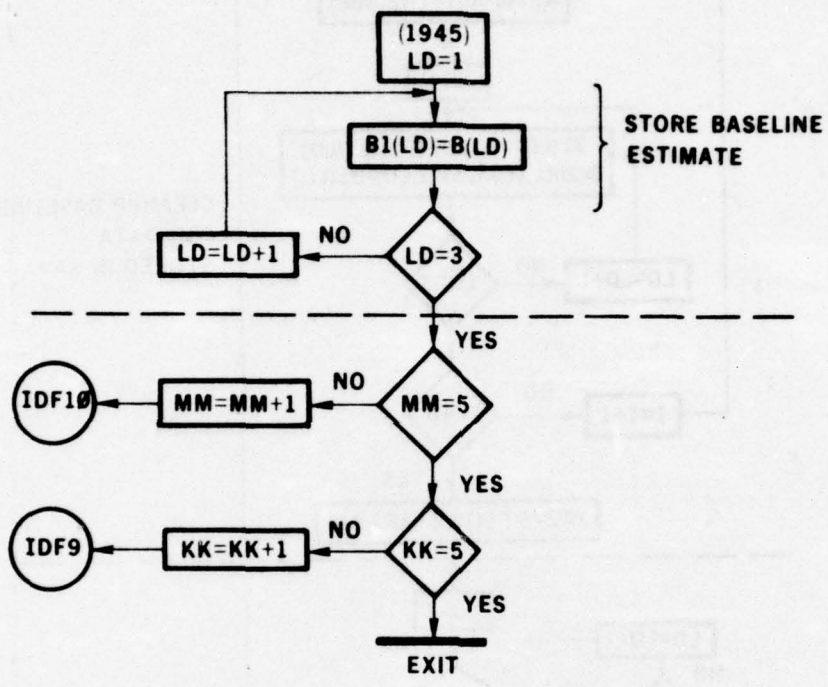












Program Verification

The program logic was verified both at USAFSAM and at Draper Laboratory. The verification tests can be grouped as

- (1) Baseline-removal test
- (2) Fiducial-point-detection test
- (3) Averaging test
- (4) Data-good, Data-bad

At USAFSAM, the tests were performed as follows. A nurses' training aid was used to generate the reference heart signal represented in Figure 3.




Figure 3. The reference cardiogram.

The reference signal was then fed into the system. Results are shown in Figure 4: The top trace (a) of Figure 4 corresponds to the original reference cardiogram; the middle trace (b) is the output of the baseline-removal algorithm; and the bottom trace (c) is the difference between the output of the algorithm and its input. The slight d.c. bias has been removed so that no unwanted signals have been added. A 3-in (7.62-cm) vertical line is printed to signify where fiducial points have been detected. The averaging logic writes a 0.31-in (0.79-cm) vertical line on the middle trace, immediately after a complete beat has been located and added to the average. In this test, all 17 beats were detected and used to form the averaged beat that is displayed in Figure 5.

Next, the reference cardiogram was corrupted by adding an additional sine wave baseline. Three different frequencies were used (0.079 Hz, 0.158 Hz, and 0.5 Hz), and for each frequency three different amplitudes were used (the lowest, not to trip the data-good data-bad logic; the middle, to trip it occasionally; and the highest, to trip it every time.)

The results for the 0.5-Hz sine wave are given first. Figures 6 and 9(a) show the results of choosing a 0.5656-mV rms amplitude to make the 2-mV/sec data-bad limit always be exceeded; no beats were averaged. All fiducial points were detected and labeled, and the 0.5-Hz sine wave was attenuated $\approx 20\%$.





Figure 4. The results for the reference signal only.




Figure 5. The average reference cardiogram.

Figures 7 and 9(b) show the system's performance with the same frequency noise (0.5 Hz), but at an amplitude (0.2828-mV rms) carefully chosen to make the data-bad limit be exceeded at times and not exceeded at others. As expected, some beats were used in calculating the average beat, while others were not; also, the fiducial point was detected for every beat. Finally, although the 0.5-Hz contamination was slightly attenuated (possibly 20%), the averaged beat (Fig. 9(b)) was slightly tilted.

To finish the 0.5-Hz test, an amplitude was chosen that would not exceed the 2-mV/sec data-bad logic. The results for this amplitude, 0.1414-mV rms, are shown in Figures 8 and 9(c). All fiducials were detected, and as expected, all beats were used to calculate the averaged beat. The contaminating sine wave was attenuated 20% by the filter, and because of the lower amplitude of the sine wave, the averaged beat was less tilted.

Next, the frequency of the noise signal was decreased to 0.158 Hz, and the same sequence of amplitudes (high, to always trigger data-bad test; medium, to trigger it some time but not always; low, to never trigger data-bad test) were used to generate baseline noises. The results obtained with the high-amplitude (2.121-mV rms) noise are given in Figures 10 and 13(a). All fiducials were identified, and as expected, the data-bad criterion was always exceeded; the average waveform was not calculated. The 0.158-Hz contamination was attenuated 80%. Then the amplitude of the sine wave contamination was decreased to 1.061-mV rms, a value that should make the data-bad limit be exceeded at times and not at others; this was the case, as shown in Figures 11 and 13(b). Finally, the amplitude of the sine wave contamination was decreased to 0.55146-mV rms, a level that should make all the data acceptable. As indicated by Figures 12 and 13(c), all fiducials were detected and all beats were used in determining the averaged beat.

The same procedure was repeated with a 0.079-Hz sine wave (a more realistic baseline). The results (using 4.242-, 2.121-, and 1.06-mV rms), which were the same as before, are shown in Figures 14 through 17.




Figure 6. Test results for reference signal corrupted by a 0.5-Hz, 0.5656-mV rms sine wave.




Figure 7. Test results for reference signal corrupted by a 0.5-Hz, 0.2828-mV rms sine wave.




Figure 8. Test results for reference signal corrupted by a 0.5-Hz, 0.1414-mV rms sine wave.

— 1 mV




Figure 9. The average waveforms calculated for the reference signal corrupted by a 0.5-Hz: (a) 0.5656-mV; (b) 0.2828-mV; and (c) 0.1414-mV rms sine wave.




Figure 10. Test results for reference signal corrupted by a 0.158-Hz, 2.121-mV rms sine wave.




Figure 11. Test results for reference signal corrupted by a 0.158-Hz, 1.061-mV rms sine wave.




Figure 12. Test results for reference signal corrupted by a 0.158-Hz, 0.55146-mV rms sine wave.




Figure 13. The average waveforms calculated for the reference signal corrupted by a 0.158-Hz: (a) 2.121-mV; (b) 1.061-mV; and (c) 0.55146-mV rms sine wave.




Figure 14. Test results for reference signal corrupted by a 0.079-Hz, 4.242-mV rms sine wave.




Figure 15. Test results for reference signal corrupted by a 0.079-Hz, 2.121-mV rms sine wave.




Figure 16. Test results for reference signal corrupted by a 0.079-Hz, 1.06-mV rms sine wave.




Figure 17. The average waveform calculated for reference signal corrupted by a 0.079-Hz: (a) 4.242-mV; (b) 2.121-mV; and (c) 1.06-mV rms sine wave.

Several tests were made at the Charles Stark Draper Laboratory to precisely examine the response of the baseline-removal algorithm to idealized high-energy waveforms. The objectives were to determine if any computer-generated artifacts occurred and to determine the operational limits of the algorithm. Three types of waveforms were used in the tests: (1) triangular waves, (2) square waves, and (3) pure steps.

Results obtained using a very-high-energy triangle waveform are seen in Figure 18. The triangles have a 4-mV amplitude and a sinusoidal baseline with a 2-mV amplitude and 0.16-Hz frequency. With these values, the maximum baseline slope is 2.01 mV/sec, just exceeding the 2-mV/sec threshold for bad-data rejection; however, no bad-data tests were used here--the objective was to evaluate the baseline-removal algorithm. The underlying heart rate is 100 beats/sec, which is borderline tachycardia. Higher heart rates were used to insure that QRS complexes in close proximity would not lead to problems in reconstruction. Figure 18 demonstrates good waveform reconstruction. The only computer-generated artifact is a small residual baseline which tends to average out in time. Slope estimate is affected by the large triangular waveforms, but is centered around the correct value and exhibits no phase shift.




Figure 18. Baseline removal evaluation, triangular QRS wave.

Next, a series of square waves were made using amplitudes of up to 4 mV and durations of up to 160 msec. In Figure 19, the results obtained using a 4-mV square wave of 80-msec duration, superimposed on the same sinusoidal baseline, are shown. The baseline removal is excellent, with only a small residual artifact that depends upon the baseline remaining. The slope estimate, while affected by the square waves, is correctly centered and exhibits zero phase shift. Square waves of duration greater than 80 msec could not be handled. If duration is greater, two samples may occur within the square wave, indicating zero slope. Although provision could easily be made for handling wider square waves, this does not appear necessary since waves are nonphysiological in origin.




Figure 19. Baseline-removal evaluation, square QRS wave.

Finally, tests were made to determine the response to sudden step changes in the baseline. Figure 20 shows the results of a typical test. First a jump of 0 to 1 mV was used, with a subsequent jump of -2 to -1 mV. The corresponding baseline and slope estimates are shown in the figure. The effects of the step on the estimates are spread out evenly on both sides of the jump, owing to the symmetry of the moving window. There is a baseline error within 220 samples (440 msec) of the jump, and a slope error within 200 samples (400 msec) of the jump. The maximum baseline error is 0.5 mV/sec per mV, and the maximum slope error is about 2 mV/sec per mV. The results of using the idealized waveforms of Figures 18-20 indicate that the algorithm operates satisfactorily in these cases and introduces only small artifacts that tend to average out with time.




Figure 20. Baseline-removal evaluation, step baseline changes.

OUTPUT DATA FORMATS

Three types of data output are generated by the processing system. All files are on 0.5-in (1.27-cm) 9-track magnetic tape recorded at 800 BPI. Two of these are the standard RDOS formatted files, and the third is written for input to an IBM 360/65 for further analysis of the vectors and amplitudes which cardiologists use to identify coronary problems [11]. Logs of all tapes are created on the line printer as the data are stored on tape.

RDOS File Formats

Both of the RDOS files use the standard system of 512 bytes of data and 2 bytes of file identification. These files are buffered and may be referenced an element at a time. The data tape produced by the data acquisition or digitizing program consists of a 6-byte identifier, normally the vectorcardiogram number, right-justified with leading zeros. This is followed by three 16-bit integers containing the age, height, and weight of the subject. Then follows 2250 16-bit integers of calibration pulses, 750 per channel, written $X_1, Y_1, Z_1, X_2, Y_2, Z_2, \dots, X_{750}, Y_{750}, Z_{750}$. Finally 18,000 16-bit integers of vectorcardiogram data are written in the same manner as the calibration data. This is followed by an end-of-file mark. A segment of data is checked for overscaling before it is transferred to magnetic tape. If overscaling of calibration data has occurred, a message is written on the line printer and the control terminal and a restart occurs; if not, the data is written on tape. If overscaling occurs on the vectorcardiogram, the message is written again and the operator can restart.

The output of the processing program described in the Digital Filter section contains the same identification, age, height, and weight information as the digitizing program. This is followed by three real values which are the scale factors for X, Y, and Z leads respectively. To transform the filtered vectorcardiogram data to millivolts, each value is divided by the scale factor for that lead. After the scale factors, come 16,800 16-bit integers of the filtered vectorcardiogram in $X_1, Y_1, Z_1, X_2, Y_2, Z_2, \dots, X_{5600}, Y_{5600}, Z_{5600}$ sequence. These data are followed by a sentinel integer of 32767 for each lead. Then comes one real value which is the heart rate for this segment. Next, 1800 real values, representing the averaged and calibrated waveforms, are written in $X_1, Y_1, Z_1, X_2, Y_2, Z_2, \dots, X_{600}, Y_{600}, Z_{600}$ sequence. The patient record is closed with an end-of-file mark.

IBM 360/65 Tape Formats

The data for input to the IBM 360/65 are the same as the output of the processing program but have been written in six blocks of 7200

bytes each. The first block contains the identification information in ASCII characters: the age, height, weight, scale factors, and 3588 vectorcardiogram values. Blocks 2, 3, 4, and 5 contain the remainder of the vectorcardiogram, the sentinel, and the heart rate. Block 6 contains the 1800 real values of the averaged waveform.

DIGITIZING-SYSTEM OPERATION

To operate the programs used to acquire and process cardiograms, the operator must be familiar with the Data General RDOS. This system is described in Data General's manual entitled "RDOS, Real-Time Disk Operating System."

Once familiar with this system, the operator can use the programs to obtain data and request actions through a series of questions and messages on the operator's console. Certain pieces of hardware must be operational and identified to the system for the programs to operate properly.

Required Programs

Three programs are required to complete the process from data acquisition to IBM-360-compatible tape generation. The first program, DRADIG, acquires the data and creates a magnetic tape file for each subject. The second program, F5DRA, processes the output of the first program by determining calibration factors, filtering the vectorcardiogram data, and producing, for each lead, an average beat calibrated to millivolts. This program also produces an output file for each program. The third program, TO360, converts the output of F5DRA to IBM-360-compatible tape.

Required Hardware

The basic hardware requirements for the current system are two 9-track magnetic tapes (800 BPI), a fixed-head NOVADISK for temporary data storage while digitizing, a graphics terminal connected to channel 1 of the multiplexer, and at least three channels of analog-to-digital (A/D) converters. Also required are a 32K NOVA minicomputer with floating-point processor and hardware multiply and divide, and an analog magnetic tape transport with the X channel connected to A/D channel 0, Y to channel 1, and Z to channel 2.

Operating Procedures

DRADIG--The data acquisition program, DRADIG, obtains all operator-provided information and control instructions through the operator's

console. A magnetic tape must be mounted on unit zero and loaded. The tape is identified to the system by the operator typing INIT MTO and RETURN. The program can then be started by typing DRADIG and RETURN. The program then asks NEW TAPE, YES OR NO, and the operator types the proper response and a return. If he types NO, the program assumes that the tape contains some data and asks NUMBER OF FILES USED. The operator responds with the proper file number and a return. After positioning the tape to the proper file, the program asks CASE NUMBER. The operator responds with six characters, using leading zeros, and a return. The program then asks for age, height, and weight. Response to each must be an integer followed by a return.

The program message DIGITIZING CALIBRATION--TYPE RETURN TO START appears; and, when the analog magnetic tape is positioned at the calibration data and running at the proper speed, the operator types RETURN. Six seconds of calibration data are digitized by the program, examined for overscale, and, if no overscaling is detected, spooled to magnetic tape. If overscale is detected, the bell on the console rings and numbers representing the number of samples out of range appear on the console and line printer. The message PAUSE then appears, and the program waits for a return. After the operator types RETURN, the program asks for another case number and continues as before.

The program message DIGITIZING DATA--TYPE RETURN TO START appears; and, when the analog tape is positioned to the data segment required and running at the proper speed, the operator types RETURN. These data are also examined for overscale, and transferred to tape if none occurs. If overscale occurs, the same procedure applies as for calibration overscale. After the data have been transferred to tape, the program asks for another case number. If the operator wishes to terminate the run, he types DONE and RETURN. If he wishes to continue, he supplies a new case number and a return.

F5DRA--The digital filtering and averaging program obtains all control information from the operator's console and does any graphic output on the graphic terminal connected to multiplexer channel 1. Under normal automatic processing, a data tape must be mounted on one tape transport and loaded, and an output tape mounted on a second transport and loaded. Both must be identified to the system with the INIT instruction.

To load the program, the operator types F5DRA and RETURN. When the program message PLOT, 0-NONE, 1-X, 2-Y, 3-Z appears, the operator types the proper value to eliminate plotting or selects the lead to be plotted, and a return. The program message MODE; AUTOMATIC OR SINGLE appears; and the operator types either AUTO or SINGLE, and RETURN. A response of SINGLE will cause the program to terminate after processing only one file; AUTO will cause the program to process sequential files from the input magnetic tape until end of tape is encountered. If files from the disk are to be processed, the SINGLE mode is selected; and the

program then asks FIRST INPUT FILE. If AUTO mode is selected, the response is of the form "MTN:mm"--where N is the unit number of the input tape transport, and mm is the number of the first file for input. Next the program asks FIRST OUTPUT FILE. For AUTO mode, the same type of response is given as was given for the input file in AUTO. For SINGLE mode, each response is the 6-character file name.

The program then proceeds to filter, scale, and average the data, using the graphic terminal to plot the selected lead and the averaged waves. If the hard-copy unit is on, it also makes hard copy of each screen full of data. The filtered and averaged data are written to the output file. Individual R-R intervals, along with ratios of successive intervals, are recorded on the line printer. The ratios are used to detect rhythm disturbances. The averaged waveforms are always plotted. The running time of the program is nearly doubled if plotting of a lead is requested.

TO360--The program to generate IBM-360-compatible tape requires two tapes for operation. Both must be loaded and identified to the operating system with the INIT instruction. To load the program, the operator types TO360 and RETURN. The queries and responses are the same as those of the processing program for mode, output and input files. The program sequentially processes an input tape and creates multiple files on the output tape for further data processing on an IBM 360/65 at San Antonio Data Services Center.

CONCLUSION

A system for digitizing and filtering ECG/VCGs has been designed and is operational at USAFSAM. Over 900 VCGs have been digitized and are being used in ongoing research work, both at USAFSAM and under contract.

REFERENCES

1. Brandan, C. W., and D. A. Brady. A hardware trigger for temporal indexing of the electrocardiographic signal. *Comput Biomed Res* 3:47-57 (1970).
2. Cox, J. R. Special purpose digital computers. *Comput Biomed Res* 2:79-87 (1965).
3. Frederick, D. K., and A. B. Carlson. *Linear systems in communication and control*, pp. 129-133. New York: John Wiley, 1971.
4. Froelicher, V. F., et al. Treadmill exercise testing at the USAF School of Aerospace Medicine: Physiological responses in aircrewmembers and the detection of latent coronary artery disease. *AGARDOGRAPH* No. 210, AGARD, Seine, France, 1976.
5. Golden, D. P., Jr., R. A. Wolthuis, and G. W. Hoffler. A spectral analysis of the normal resting electrocardiogram. *IEEE Trans Biomed Eng* 20(5):366-372 (1973).
6. Gustafson, D. E., A. S. Willsky, and J. Y. Wang. Cardiac arrhythmia detection and classification through signal analysis. Charles Stark Draper Laboratory Final Report on contract F41609-75-C-0013, Cambridge, Mass., July 1975.
7. Lamb, L. E. *Electrocardiography and vectorcardiography: Instrumentation, fundamentals, and clinical applications*. Philadelphia: W. B. Saunders, 1965.
8. Rautaharju, P. M. Deterministic type waveform analysis in electrocardiography. *Ann NY Acad Sci* 128(3):939-954 (1966).
9. Sokolnikoff, I. S., and R. M. Redheffer. *Mathematics of physics and modern engineering*, pp. 702-711. New York: McGraw-Hill, 1958.
10. Womble, M. E., S. J. Mitter, and J. S. Halliday. Classifying vectorcardiograms via signal pattern recognition. C. S. Draper Lab Report R-766, Cambridge, Mass., Oct 1973.
11. Womble, M. E., et al. A VCG measurement and display system. Abstract to appear in *Proceedings of 1976 ACEMB*. (Paper in preparation).