Theoretical and Experimental Rate Distortion Performance in Compression of Ambulatory ECG's

Patrick S. Hamilton and Willis J. Tompkins, *Senior Member, IEEE*

Abstract—**We compare ECG data compression algorithms based on signal entropy for a given mean-square-error (MSE) compression distortion. By defining the distortion in terms of the MSE and assuming the ECG signal to be a Gaussian process we are able to estimate theoretical rate distortion bounds from average ECG power spectra. These rate distortion bounds give estimates of the minimum bits per second (bps) required for storage of ECG data with a given MSE regardless of compression method.**

From average power spectra of the MIT/BIH arrhythmia database we have estimated rate distortion bounds for ambulatory ECG data, both before and after average beat subtraction. These rate distortion estimates indicate that, regardless of distortion, average beat subtraction reduces the theoretical minimum data rate required for ECG storage by approximately 100 bits per second (bps). Our estimates also indicate that practical ambulatory recording requires a compression distortion on the order of 11 μV rms.

We have compared the performance of common ECG compression algorithms on data from the MIT/BIH database. We sampled and quantized the data to give distortion levels of 2, 5, 8, 11, and 14 μV rms. These results indicate that, when sample rates and quantization levels are chosen for optimal rate distortion performance, minimum data rates can be achieved by average beat subtraction followed by first differencing of the residual signal. Achievable data rates approximate our theoretical estimates at low distortion levels and are within 60 bps at higher distortion levels.

I. Introduction

A significant amount of work has been done on compression of ECG data for storage and transmission of diagnostic ECG's [1]. Recently there has been interest in compression of ECG data for ambulatory monitoring [2], [3], and a number of digital Holter recorders are presently available. A solid-state Holter recorder offers the possibility of improved low frequency response for ST-segment monitoring. Such recorders can be smaller and eventually, less expensive than tape-based recorders.

Because Holter tape recorders generally record two leads for 24 h, and at present 4 Mbytes (MB) is about as much memory as a small battery-powered system can support for 24 h, a practical solid-state Holter recorder requires significant data compression to store 24 h of two-channel ECG in four MB of memory.

Manuscript received August 9, 1989; revised July 16, 1990. This work was supported in part by a grant from the Burdick Corporation.

P. S. Hamilton was with the Department of Electrical Engineering, University of Wisconsin-Madison, Madison, WI 53706. He is now with the Department of Electrical Engineering, Bucknell University, Lewisburg, PA 17837.

W. J. Tompkins is with the Department of Electrical Engineering, University of Wisconsin-Madison, Madison, WI 53706.

IEEE Log Number 9042177.

The majority of reported ECG compression algorithms have been written for compression and transmission of resting ECG's. These ECG compression algorithms generally fall into one of three categories: significant-point-extraction [2], [4]–[15], linear predictive coding [16]–[22], or orthogonal transform [23]–[26]. Entropy encoding is a compression technique that has been widely applied in conjunction with linear predictive coding, but is applicable to all compression algorithms [28]. Subtracting average beats from the original signal and storage of the resulting residual is another compression technique that has been used, with mixed results, in conjunction with significant-point-extraction [3], linear predictive coding [22], and orthogonal transform [23] compression.

From the literature, it is not clear which compression algorithm, or combination of algorithms, is best suited for compression of ambulatory ECG data. Direct comparison of ECG data-compression algorithms is difficult because different algorithms have been evaluated on different databases, generally derived from resting ECG's, and different methods have been used to report distortion and compression.

In this paper we estimate some theoretical limits on the compression of ambulatory ECG's, and evaluate the relative performance of some data compression algorithms that are suitable for ambulatory compression. To better estimate algorithm performance on ambulatory data, we have evaluated algorithms and estimated compression limits using data from the MIT/BIH arrhythmia database, a standard ambulatory ECG database composed of 48 half-hour Holter recordings [29].

We have evaluated the performance of different algorithms according to their resulting data rates as a functions of signal distortion. Data rates are reported as signal entropy in bits per second (bps) and distortions are reported in terms of mean squared error (MSE).

The entropy H of a signal X is calculated as

$$H(x) = \sum_{\text{all } x} p(x) \log_2 \left(\frac{1}{p(x)} \right)$$

where $p(x)$ represents the probability of signal value x. Signal entropy is the natural measure of compression because it represents the minimum average number of bits needed to noiselessly code a sample from a given signal [27].

Measuring data rate as a function of MSE distortion allows us to estimate the optimal rate distortion function for an ambulatory ECG compression algorithm. This optimal rate distortion function indicates the theoretical minimum data rates required to store an ambulatory ECG over a range of MSE distortions, regardless of compression technique. Our rate distortion measure also leads to a method for selection of sample rate

0018-9294/91/0300-0260$01.00 © 1991 IEEE

and quantization level to achieve minimum data rates for given distortion levels.

In the first section of this paper the concept of a rate distortion function is reviewed, and from our database we estimate the rate distortion function for an ambulatory ECG. We are also able to quantitatively estimate the improvement in compression resulting from average beat subtraction. In the second section we investigate the rate distortion performance for a number of compression algorithms.

Estimation of Rate Distortion Bounds for Ambulatory ECG's

Theoretically the minimum data rate required for coding a given signal will be a function of the distortion in the compressed signal. More distortion must be accepted as the signal is coded with fewer bits per second. The statistics of the signal will determine the exact nature of this rate distortion function. Berger [30] has shown that for a discrete-time stationary Gaussian source the minimum MSE distortion resulting from coding is given parametrically as a function of θ by the equation:

$$D_\theta = \frac{1}{f_0}\int_{-f_0/2}^{f_0/2} \min\,[\theta, \Phi(f)]\; df.$$

The rate is given by the parametric expression

$$R(D_\theta) = \frac{1}{2f_0}\int_{-f_0/2}^{f_0/2} \max\left[0, \log_2\left(\frac{\Phi(f)}{\theta}\right)\right] df$$

where f_0 is the sampling frequency and $\Phi(f)$ is the signal power as a function of frequency. With the logarithm calculated in base 2, the resulting rate represents the minimum number of bits per sample required to code the signal with an MSE distortion of D_θ.

Fig. 1 shows a graphical representation of the rate-distortion calculation. The curve represents the spectral density of the original bandlimited signal. The θ parameter represents the level of the horizontal line drawn across the power spectrum. The area of the lower shaded region represents spectral density of the minimum compression error given θ. The upper shaded area from $-f_c$ to f_c represents the spectral density of the coded signal. With $\Phi(f)$ plotted on a log scale, the area of the upper shaded region represents the minimum number of bits per sample required to code the signal with the distortion level D_θ.

The rate-distortion calculations as they have been presented assume that the signal to be coded is Gaussian and stationary. However, the ECG signal is neither Gaussian nor stationary. Berger has shown that the rate for Gaussian signals represents an upper bound on the rate for nonGaussian signals [30]. Thus, the rate calculated from the spectral density of the ECG signal represents an upper limit on the minimum rate required to code the ECG. It is theoretically possible to code a nonGaussian signal with a given power spectrum more efficiently than a Gaussian signal having the same power spectrum. If the ECG signal distribution is approximately Gaussian, these rate-distortion calculations should give reasonable estimates for the bounds on ECG data compression.

Estimation of Rate Distortion Curves for Ambulatory ECG's

In calculating the rate-distortion curve for our ECG data we assume that the ECG signal can be perfectly represented with a bandwidth of 100 Hz. With this assumption, we can estimate the rate distortion curve from the power spectrum of the ECG signal between dc and 100 Hz. For a given data rate, our estimates will give the minimum MSE distortion relative to an ECG signal sampled with infinite precision at 200 samples per second (sps).

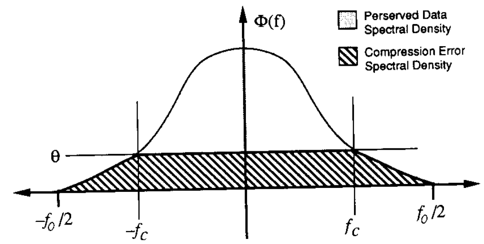

Fig. 1. Graphical representation of rate-distortion calculation for a given θ.

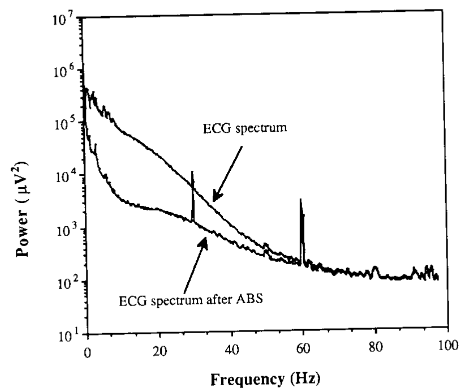

Fig. 2. Average ECG power spectrum calculated from the MIT/BIH arrhythmia database.

We calculated the average power spectrum for one channel of each of the 48 half-hour tapes from the MIT/BIH arrhythmia database. We digitally recorded ECG data from analog tapes of the MIT/BIH database with a sample rate of 200 sps and an effective resolution of 4.88 μV/LSB. We estimated the power spectrum for each tape by averaging the power spectra from consecutive 2048-point data segments.

We also estimated the power spectra of the residual ECG signal produced by average beat subtraction. The residual ECG signals were generated with an operator-assisted program that identified the *QRS* complexes, calculated average complexes for each type of beat in a given record, and subtracted the average beats from the original ECG data. Fig. 2 shows power spectra produced by averaging the power spectra for all of the MIT/BIH tapes. The power spectrum in Fig. 2 is graphed with a logarithmic amplitude scale to show the relative contribution of different spectral components to the minimum data rate. The area between the two curves is indicative of the reduction in data rate that can be achieved with average beat subtraction. Fig. 3 shows the average rate-distortion curves calculated from the average spectra in Fig. 2.

Storage of 24 h of two-channel data requires a data rate of approximately 190 bps per channel. Fig. 3 indicates that this data rate requires a coding distortion on the order of 15 μV rms. If average beat subtraction is applied to the signal before compression, our calculations indicate that, at data rates near

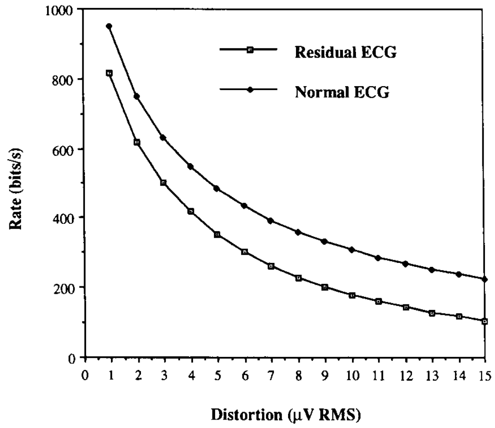

Fig. 3. Comparison of average rate-distortion curves for the ECG signal with and without average beat subtraction.

190 bps, the residual ECG signal can be stored with distortion levels on the order of 11 μV rms. Fig. 3 also shows that regardless of the distortion it should be possible to code the residual ECG with data rates approximately 100 bps less than the normal ECG signal.

A Comparison of Data Compression Techniques

Though all the algorithms that have been described in the literature can be implemented in real time with present technology, computationally efficient algorithms can be implemented on more economical hardware. Computationally efficient algorithms can also run more slowly, decreasing power consumption. Because computationally efficient algorithms are better suited for ambulatory ECG compression, we have restricted our study to the more computationally efficient significant-point-extraction and linear predictive coding algorithms.

We compared ECG data-compression schemes by dividing the processing into three steps: 1) sampling and quantization, 2) compression, and 3) entropy encoding. Fig. 4 shows a block diagram of these steps. We take compression to be the process which maps the set of sampled and quantized ECG values to a new set of values which can be coded with fewer bits. This may be accomplished by reducing the entropy of the points to be coded, as with linear predictive coding, and/or reducing the number of points to be coded, as with the Fan. In the final step the ECG is coded in an effort to produce a signal with an average bit count that is near the entropy of the compressed data.

In our experimental investigation, distortion results only from quantization and sampling. No additional distortion is introduced during compression. We select the quantization level and sample rate, based on rate-distortion principles, to approximate optimal compression for a given distortion. Different compression reduction schemes are then evaluated with the ECG signal sampled and quantized to give a range of distortion levels. The resulting signal entropys are used to compare the effectiveness of the different compression schemes.

Selection of Sample Rate and Quantization Level

Previous papers on ECG data compression have experimentally investigated data compression schemes over different ranges of sample rates and quantization levels [18], [19] with

Fig. 4. ECG data compression steps.

no theoretical criteria for choosing the sample rate and quantization level. Berger's method for calculation of the rate-distortion curve implies that sample rate and quantization level may be selected as a function of the desired MSE distortion level [30].

The method for optimal coding of a signal can intuitively be seen from the graph in Fig. 1. The signal should only be coded in portions of the spectrum where the signal level exceeds the parametric distortion θ. In this way the distortion for frequencies that are not coded is limited to the signal power over those frequencies. In the area of the spectrum that is coded, the distortion is equal to the coding distortion.

For a signal with a monotonically decreasing power spectrum there is a cutoff frequency f_c where

$$\Phi(f_c) = \theta.$$

For optimal coding, signal frequencies below f_c are coded and the signal frequencies above f_c are discarded. When the power spectrum decreases with increasing frequency, as the average ECG power spectrum does, the signal can be low-pass filtered to eliminate the high frequencies and sampled at $2f_c$. The filtered signal should then be coded with a distortion equal to θ times the ratio of the new bandwidth to the original bandwidth to give a coding distortion of $f_c\theta$. The filtered, sampled signal should then be codeable with the minimum data rate for the distortion level.

Fig. 5 shows the optimal sampling rate as a function of distortion calculated from the average ECG power spectrum in Fig. 2. While θ is less than the minimum value of the power spectrum, the sample rate is equal to twice the bandwidth of the original signal. The optimal sample rate drops drastically for coding with a distortion larger than 8.5 μV rms.

The quantization level is chosen to give an MSE distortion equal to θ times the ratio of the sampled signal bandwidth to the original signal bandwidth. Because the power spectrum of the quantized signal is not equivalent to that of the original signal, coding the bandlimited, quantized signal only approximates the ideal coding process [30], but it is a reasonable approximation when the variance of the quantization error is small compared to $\Phi(f)$.

Quantization noise can be modeled as additive white noise with a variance given by

$$\sigma^2 = \frac{q^2}{12}$$

where q is the quantization step size. Because the quantized signal may be smoothed to give a reconstructed signal with a reduced distortion, the actual distortion resulting from quantization will be slightly less than the quantization error variance. We chose quantization levels that resulted in the desired MSE distortion after optimal smoothing for minimum MSE.

Fig. 6 graphically shows the quantization levels and sample rates calculated for each of 48 tapes from the MIT/BIH databases at the 2, 5, 8, 11, and 14 μV distortion levels. Although the spread of appropriate sample rates increases with increasing distortion level, for optimal compression, the bulk of the ECG records required quantization levels that increased at roughly 3

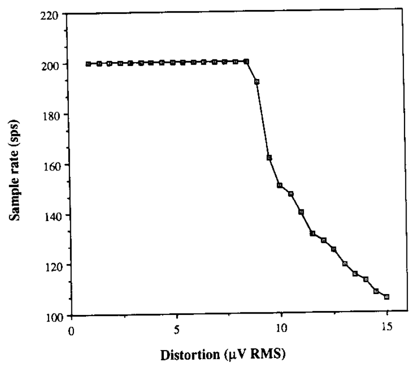

Fig. 5. Optimal sample rate as a function of signal distortion as calculated from the average ECG power spectrum.

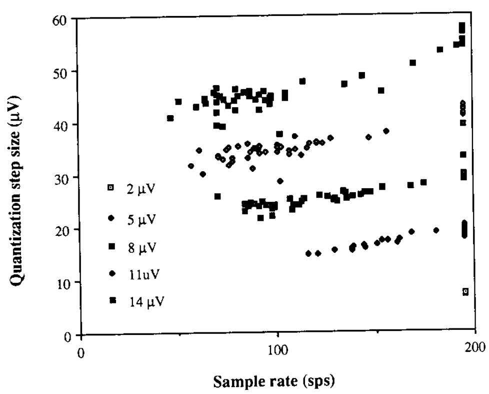

Fig. 6. ECG quantization level as a function of sample rate for efficient compression of ECG data with distortion levels of 2, 5, 8, 11, and 14 μV rms.

μV per μV rms of distortion. The corresponding sample rates seem to decrease at about 10 sps per μV rms of distortion. These relations may serve as guides in selecting sample rates and quantization levels for storing ambulatory ECG data.

Compression Techniques

We compared the performance of six compression techniques: first difference, second difference, minimum mean squared error (MMSE) linear prediction, least mean square (LMS) adaptive linear prediction, sign-adaptive linear prediction, and Fan adaptive sampling. These techniques were applied to the sampled and quantized ECG data with and without average beat subtraction. Average beats were removed from the sampled and quantized signals as described in the section on estimation of rate distortion bounds. The resulting signal entropies were used as a measure of compression efficiency.

Differencing and MMSE Linear Prediction

First and second differenced ECG signals were tested, as well as error signals produced using linear predictors of orders 2, 3, 5, and 10 with predictor coefficients estimated from the average correlations over the length of each half-hour tape.

Adaptive Linear Prediction

We tested second- and third-order LMS and sign-adaptive predictors. The LMS algorithm updated predictor coefficients according to the equation:

$$a_i(t) = a_i(t - 1) + \beta x(t - i)\, e(t).$$

The predictor coefficients in the sign-adaptive algorithm were updated according to the equation:

$$a_i(t) = a_i(t - 1) + \Delta x(t - i) \operatorname{sgn}(e(t)).$$

We chose the parameters β and Δ experimentally by monitoring the magnitude of the error signal. If in the course of a test, the error signal increased dramatically, we decreased the β or Δ parameter by an order of magnitude and restarted the test. After a stable test, we decreased the adaptation parameters again and ran a second test. With quantization values on the order of 10 μV, we found that stable values of β were generally less than 0.00001, and stable values of Δ were generally less than 0.01.

Fan

It is generally difficult to predict the distortion which will result from Fan coding because the points which are saved have a distortion related to the original quantization level of the signal, and the reconstructed points that are calculated from these points have an average distortion which is related to the error tolerance used for compression. It is difficult to predict the MSE distortion because it is difficult to predict what fraction of the original number of points will be saved.

No optimal method has been proposed for choosing the tolerance parameter in the Fan encoding algorithm with respect to the original quantization level of the signal. We have implemented a Fan compression according to the centerline criterion reported by Ishijima *et al.* [8]. By choosing our tolerance level to be half the signal's quantization level, we may perfectly reconstruct the original, regularly sampled, uniformly quantized signal from the points stored by the Fan algorithm. This allows us to directly compare the compression produced by a Fan algorithm and the compression resulting from linear prediction because both techniques produce compressed signals with the same MSE distortion.

Fan compression gives a set of points which are stored as vectors representing amplitude and time. To further compress the data produced by the Fan algorithm, we first differenced both the amplitude and time coordinates of each point and then calculated the entropy of the compressed data as the sum of the entropies for the differenced time and amplitude values.

Results

Tables I and II list the average entropy in bps produced by our six data compression techniques both with and without average beat subtraction. The entropies are the average for all 48 of the MIT/BIH tapes recorded with distortion levels of 2, 5, 8, 11, and 14 μV rms. The tables also list for comparison the theoretical values calculated in our previous section for these five distortion levels.

For nearly all distortion levels, both with an without average beat subtraction, first differencing produced residuals with the lowest entropy. With average beat subtraction and a 2 μV rms distortion, compression based on a tenth-order MMSE predictor

TABLE I
AVERAGE ENTROPY (bits/s) FROM COMPRESSION WITHOUT AVERAGE BEAT SUBTRACTION

Distortion (rms):	2 μV	5 μV	8 μV	11 μV	14 μV
1st Diff.	710.8	452.7	306.2	239.4	197.2
2nd Diff.	752.1	503.8	354.1	286.1	241.8
MMSE 2	773.0	485.9	340.5	270.0	225.0
MMSE 3	738.0	489.5	340.5	270.9	225.6
MMSE 5	733.2	486.1	340.7	270.5	225.0
MMSE 10	729.6	483.6	340.7	270.2	225.5
LMS 2	723.1	475.7	326.3	258.7	214.8
LMS 3	723.5	479.7	330.5	263.7	218.3
Sign 2	727.5	479.2	321.7	249.7	205.1
Sign 3	730.9	479.3	325.7	253.8	208.3
SAPA	753.5	517.7	356.1	282.4	234.4
Theoretical	749.93	485.6	358.0	285.6	237.5

TABLE II
AVERAGE ENTROPY (bits/s) RESULTING FROM COMPRESSION WITH AVERAGE BEAT SUBTRACTION

Distortion (rms):	2 μV	5 μV	8 μV	11 μV	14 μV
1st Diff.	621.7	387.4	263.7	208.5	174.4
2nd Diff.	718.8	488.7	344.7	279.6	238.2
MMSE 2	630.7	400.9	272.3	215.1	181.3
MMSE 3	629.1	399.6	271.6	215.5	181.2
MMSE 5	624.1	396.9	271.1	214.3	179.8
MMSE 10	620.6	393.5	269.7	213.2	178.6
LMS 2	629.9	390.3	266.5	211.2	177.0
LMS 3	626.3	391.5	267.1	211.5	177.2
Sign 2	640.4	400.0	270.5	213.6	178.4
Sign 3	645.8	404.7	272.5	214.4	179.1
SAPA	668.2	464.6	319.4	256.7	215.9
Theoretical	617.1	352.8	228.1	159.7	115.8

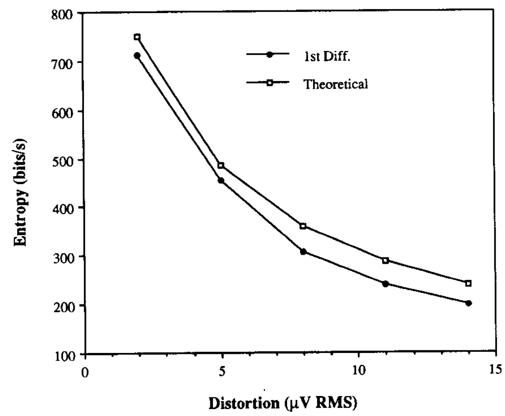

Fig. 7. Comparison of theoretical rate-distortion curve and experimental rate/distortion curve produced by first differencing the ECG signal.

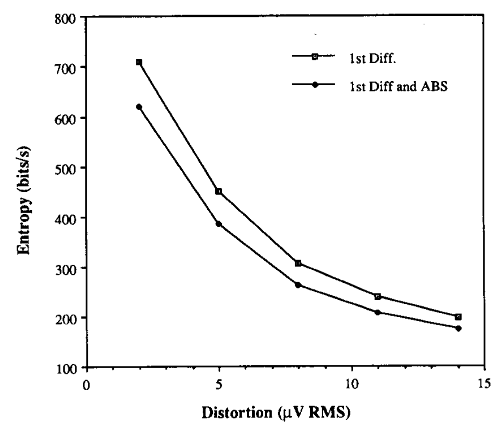

Fig. 8. Comparison of ECG residual signal entropies resulting from first differencing with and without average beat subtraction.

produced a slightly lower signal entropy than did first differencing. The adaptive linear predictors performed slightly worse than first differencing, but slightly better than the MMSE linear predictors of equal order. Fan and second differencing produced the least amount of compression, with the Fan algorithm performing slightly better than second differencing.

Fig. 7 compares the signal entropy resulting from first differencing of the ECG signal with the rate-distortion curve calculated from the average ECG power spectrum. The experimental rate-distortion curve parallels the theoretical rate-distortion curve, but first differencing produces rates which are 30 to 50 bps lower than the estimated theoretical rate-distortion curve. This large discrepancy between theoretical rate distortion bounds and experimental rate distortion performance is most likely due to the nonGaussian distribution of the ECG signal. Estimation of the rate distortion bound assumed a Gaussian distribution.

Average beat subtraction reduced the ECG residual signal entropy when used with all the compression methods we tested. Fig. 8 compares the ECG residual signal entropy resulting from first differencing of the ECG signal both with and without average beat subtraction. The resulting signal entropy is lower for all distortion levels when average beat subtraction is used, but the improvement resulting from average beat subtraction decreases with increasing signal distortion. This decrease in the effectiveness of average beat subtraction with increasing quantization and decreasing sample rate is due to misalignment of average beats and roundoff error in calculation of the average beat.

It should also be noted that, even at low distortion levels, average beat subtraction does not reduce the compressed signal entropy as much as our theoretical estimates would lead us to suspect. This may result because the distribution of the residual ECG is closer to that of a Gaussian distribution than the distribution of the raw ECG. Thus, the actual lower bound for the residual rate distortion curve will be closer to our calculated bound.

DISCUSSION

We have investigated the relationship between data rates and signal distortion in compression of ambulatory ECG signals. Our rate-distortion calculations indicate that, for ECG compression with an average data rate of 200 bps, which would allow 24-h Holter recording with 4 MBytes of memory, we must tolerate distortion in the compressed signal of almost 15 μV rms. Our calculations indicate that the use of average beat subtraction should provide data rates near 200 bps with an rms distortion of about 11 μV. We have demonstrated that achievable data rates without average beat subtraction are somewhat lower than theoretical estimates for ECG compression without average beat subtraction. Table I shows that first difference compression results in data rates that are roughly 30–40 bps lower than theoretical estimates. As shown in Table II, when average beat subtraction is employed at lower distortion levels achievable data rates are close to theoretical estimates, but at higher distortion levels achievable data rates are as much as 60 bps higher than theoretical estimates.

When we evaluated data compression algorithms on ambulatory ECG data that were sampled and quantized to give a specified rms distortion we determined that minimum data rates are achieved with average beat subtraction and first differencing of the residual signal. This would seem to contradict the findings of Ruttiman and Pipberger [18] and Pahlm *et al.* [19] who concluded that second differencing produces better compression than first differencing. However, Ruttiman and Pipberger evaluated compression algorithms on ECG's that were taken at rest, and Pahlm *et al.* filtered their ambulatory recordings with a low-pass cutoff of 20 Hz. High frequency muscle noise in ambulatory ECGs that are not severely filtered reduces the effectiveness of compression with second differencing. Our findings have recently been supported by two studies which also used the MIT/BIH arrhythmia database. Moody *et al.* noted that the results of compression with first and second differences were comparable when sampling at lower rates with relatively coarse quantization [2]. Bertinelli *et al.* [31] reported better performance with first differencing than with second differencing on data sampled at 100 Hz with 10-bit quantization.

Our results also indicate that first differencing produces better compression than either adaptive linear predictive coders or linear predictive coders based on signal statistics. Ruttiman and Pipberger [18] and Pahlm *et al.* [19] already demonstrated that second differencing produces data rates that are nearly the same as those produced by linear predictors, based on signal statistics from ECG data taken at rest and filtered ECG data. We have extended these results to include minimally filtered ambulatory data. We have also shown that adaptive compressors, as proposed by Cohn and Melsa [21], [22] do not represent an improvement over simple first differencing. We did note that MMSE linear predictors and adaptive linear predictors produced residual signals with lower variances than the residual signals produced by first differencing, but first differencing resulted in lower signal entropies. This indicates that because different predictors produce error signals with different probability distributions, the resulting residual entropies are not related in the same way to the residual variances.

Data compression of the Fan algorithm was inferior to all other compression techniques except second differencing. Much has recently been written about the performance of the Fan algorithm for ECG compression [10]–[13], but these studies only considered the reduction in sample numbers. In contrast, we examined the entropy of the resulting sample values and sample times and the implied data rate required for storage. Our results indicated that, when evaluated with respect to MSE distortion, the reduction in sample number resulting from Fan compression is not sufficient to compensate for the increased number of bits required for storage. Better compression, with equal distortion, can be achieved with methods that store prediction errors.

Our results for the Fan algorithm imply that similar methods which report low rates, such as the TRIM compression algorithm [2], probably produce high MSE in the reconstructed signal. What is difficult to judge is the importance of MSE as a measure of distortion in the ECG signal. Algorithms which reconstruct the compressed signal from significant points are good at retaining the general shape of the beat complexes, but the high MSE may reduce the accuracy of measurements such as ST level, which require signal averaging.

As a final point we have confirmed, both experimentally and theoretically, that average beat subtraction can be used to improve ECG compression. We have also quantified the improvement in compression that can be expected when average beat subtraction is used. Although this is intuitive, and is standard practice in compression of resting ECG's, other researchers have not always been able to improve compression with average beat subtraction [3], [24]. Our theoretical calculations indicate that average beat subtraction should lower residual signal entropies by about 100 bits/s, but our experimental results produced reductions in entropy that ranged from only 90 bits/s for signals with low distortion to slightly more than 20 bits/s for signals with high distortion.

References

[1] S. M. S. Jalaleddine, C. G. Hutchens, R. D. Strattan, and W. A. Coberly, "ECG data compression techniques—A unified approach," *IEEE Trans. Biomed. Eng.*, vol. BME-37, pp. 329–343, 1990.

[2] G. B. Moody, K. Soroushian, and R. G. Mark, "ECG data compression for tapeless ambulatory monitors," *Comput. Cardiol.*, pp. 467–470, 1988.

[3] S. M. S. Jalaleddine, C. G. Hutchens, W. A. Coberly, and R. D. Strattan, "Compression of Holter ECG data," *Biomed. Sci. Instrument.*, vol. 24, pp. 35–45, 1988.

[4] J. R. Cox, F. M. Nolle, H. A. Fozzard, and G. C. Oliver, Jr., "AZTEC: a preprocessing program for real-time ECG rhythm analysis," *IEEE Trans. Biomed. Eng.*, vol. BME-15, pp. 128–129, 1968.

[5] B. Furht and A. Perez, "An adaptive real-time ECG compression algorithm with variable threshold," *IEEE Trans. Biomed. Eng.*, vol. BME-35, pp. 489–494, 1988.

[6] L. W. Gardenhire, "Data redundancy reduction for biomedical telemetry," in *Biomedical Telemetry*, C. A. Caceres, Ed. New York: Academic, 1965, pp. 255–298.

[7] —, "Redundancy reduction, the key to adaptive telemetry," in *Proc. 1964 Natl. Telemeter Conf.*, Los Angeles, CA, 1964, Sec. 1-5, pp. 155–298.

[8] M. Ishijima, S. Shin, G. H. Hostetter, and J. Sklansky, "Scan-along polygonal approximation for data compression of electrocardiograms," *IEEE Trans. Biomed. Eng.*, vol. BME-30, pp. 723–729, 1983.

[9] R. C. Barr, S. M. Blanchard, and D. A. Dipersio, "SAPA-2 is the Fan," *IEEE Trans. Biomed. Eng.*, vol. BME-32, p. 337, 1985.

[10] D. A. DiPersio and R. C. Barr, "Evaluation of the Fan method of adaptive sampling on human electrocardiograms," *Med. Biol. Eng. Comput.*, vol. 23, pp. 401–410, 1985.

[11] S. M. Blanchard and R. C. Barr, "Comparison of methods for adaptive sampling of cardiac electrograms and electrocardiograms," *Med. Biol. Eng. Comput.*, vol. 23, pp. 377–386, 1985.

[12] L. N. Bohs and R. C. Barr, "Real-time adaptive sampling with the Fan algorithm," *Med. Biol. Eng. Comput.*, vol. 26, pp. 565–573, 1988.

[13] —, "Prototype for real-time adaptive sampling using the Fan algorithm," *Med. Biol. Eng. Comput.*, vol. 26, pp. 574–583, 1988.

[14] W. C. Mueller, "Arrhythmia detection program for an ambulatory ECG monitor," *Biomed. Sci. Instrum.*, vol. 14, pp. 81–85, 1978.

[15] J. P. Abenstein and W. J. Tompkins, "A new data-reduction algorithm for real-time ECG analysis," *IEEE Trans. Biomed. Eng.*, vol. BME-29, pp. 43–48, 1982.

[16] K. L. Ripley and J. R. Cox, "A computer system for capturing transient electrocardiographic data," *Proc. Comput. Cardiol.*, pp. 439–445, 1976.

[17] D. Stewart and G. E. Dower, "An ECG compression code," *J. Electrocardiol.*, vol. 6, pp. 175–176, 1973.

[18] O. Pahlm, P. O. Borjesson, and O. Werner, "Compact digital storage of ECG's," *Comput. Prog. Biomed.*, vol. 9, pp. 263–300, 1979.

[19] U. E. Ruttiman and H. V. Pipberger, "Compression of ECG by prediction or interpolation and entropy encoding," *IEEE Trans. Biomed. Eng.*, vol. BME-26, pp. 613–623, 1979.

[20] U. E. Ruttiman, A. S. Berson, and H. V. Pipberger, "ECG data compression by linear prediction," *Proc. Comput. Cardiol.*, pp. 313-315, 1976.

[21] D. L. Cohn and J. L. Melsa, "Application of sequential adaptation to data rate reduction," in *Proc. Eighth Asilomar Conf. Circ., Syst. Comput.*, 1974, pp. 201-206.

[22] —, "Efficient digitization methods for electrocardiograms," in *Proc. Joint Automatic Control Conf.*, 1977, pp. 1386-1391.

[23] N. Ahmed, P. J. Milne, and S. G. Harris, "Electrocardiographic data compression via orthogonal transforms," *IEEE Trans. Biomed. Eng.*, vol. BME-22, pp. 484-487, 1975.

[24] M. E. Womble, J. S. Halliday, S. K. Mitter, M. C. Lancaster, and J. H. Triebwasser, "Data compression for storing and transmitting ECG's/VCG's," *Proc. IEEE*, vol. 65, pp. 702-706, 1977.

[25] W. S. Kuklinski, "Fast Walsh transform data-compression algorithm: e.c.g. applications," *Med. Biol. Eng. Comput.*, vol. 21, pp. 465-472, 1983.

[26] B. R. S. Reddy and I. S. N. Murthy, "ECG data compression using Fourier descriptors," *IEEE Trans. Biomed. Eng.*, vol. BME-33, pp. 428-434, 1986.

[27] C. E. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J.*, vol. 27, pp. 379-423 (pt. I), pp. 623-656 (pt. II), 1948.

[28] D. A. Huffman, "A method for construction of minimum-redundancy codes," *Proc. IRE*, vol. 40, pp. 1098-1101, 1952.

[29] R. G. Mark, P. S. Shulter, G. B. Moody, P. Devlin, and D. Chernoff, "An annotated ECG database for evaluating arrhythmia detectors," *IEEE Front. Eng. Med. Health Care*, pp. 205-210, 1982; database is available from Joe Mietus, Beth Israel Hospital, Biomed. Eng. Division, KB-26, 330 Brookline Ave., Boston, MA 02215.

[30] T. Berger, *Rate Distortion Theory; A Mathematical Basis for Data Compression.* Englewood Cliffs, NJ: Prentice-Hall, 1971.

[31] M. Bertinelli, A. Castelli, C. Combi, and F. Pinciroli, "Some experiments on ECG data compression in the presence of arrhythmias," *Comput. Cardiol.*, pp. 473-476, 1988.

Patrick S. Hamilton, for a photograph and biography, see this issue, p. 259.

Willis J. Tompkins, (S'61-M'66-SM'77), for a photograph and biography, see this issue, p. 259.