Six-band sub-band coder on ECG waveforms

S. C. Tai

Institute of Electrical Engineering, National Cheng-Kung University, Tainan, Taiwan, Republic of China

Abstract—*An ECG sampled at a rate of 500 samples s^{-1} or more produces a large amount of redundant data that are difficult to store and transmit. A process is therefore required to represent the signals with clinically acceptable fidelity and with the least code bits possible. In the paper, an efficient sub-band coding method for encoding ECG waveforms is presented. Although sub-band coding has been successfully applied to speech signals, it is the first time that this technique has been applied to the encoding of ECG waveforms. A frequency band decomposition of an ECG waveform is carried out by means of quadrature mirror filters (QMF), which split the ECG spectrum into six bands of unequal width. In the lower frequency bands, which contain most of the ECG spectrum energy, a larger number of bits per sample is used, whereas in upper frequency bands, which contain noise-like signals, fewer bits per sample and the run length coding method are used. The simulation results are presented in terms of bit rates and the quality of the reconstructed waveforms. The results show that a reproduction with an average signal-to-noise ratio (SNR) of 29·97 dB can be achieved even at an average bit rate of 0·81 bits per sample.*

Keywords—*Data compression, Huffman code, Quadrature mirror filter, Run length coding, Sub-band coding*

Med. & Biol. Eng. & Comput., 1992, **30**, 187–192

1 Introduction

DIGITAL PROCESSING of electrocardiography (ECG) has been widely used for many years (RIPLEY, 1976–1989). It has the benefits of low noise and low distortion and the disadvantage of a large amount of data that are difficult to transmit and store. For example, let us consider the case in the Veterans General Hospital in Taipei. For each patient, 2·5 s of ECG data are sampled for each of the 12 leads. The sample rate is 500 samples s^{-1} and each sample takes 12 bits (KAO *et al.*, 1983). Thus, 22·5 K bytes of memory are needed for each patient's ECG data. For some special cases, such as arrhythmia, 30 s more rhythm strip is taken and the amount of storage needed is increased to 90 K bytes (CHEN *et al.*, 1984). Note that this is just for one patient. In a big hospital like the Veterans General Hospital in Taipei, 1000 patients per month is not exceptional. This produces a large quantity of data that are difficult to store and transmit. A process is required that reduces the amount of data with acceptable fidelity, and this is the object of any data compression algorithm.

Reducing the number of code bits permits

(*a*) individual signals to be transmitted faster
(*b*) more parallel channels to be transmitted through a communication link
(*c*) a reduction of transmitter power
(*d*) more compact signal storage.

First received 26th July 1990 and in final form 5th April 1991

To obtain a representation of signals with acceptable fidelity and with the least code bits possible, we need to reduce the amount of redundant information. By redundancy, we mean that data are highly correlated, *i.e.* data are in fact not mutually independent. A compressor must therefore 'break' the highly correlated data into less correlated ones, if it is not feasible to obtain completely uncorrelated data. An ECG signal in its digital form is one of the natural signals that are highly correlated. It has been shown by MEAD *et al.* (1979) that 95 per cent of ECG spectrum energy is concentrated between 0·25 and 35 Hz, *i.e.* ECG signals are low frequency oriented and are therefore highly correlated.

2 Sub-band coding

One way to eliminate the correlation of ECG data is to decompose it in the frequency domain into several bands such that the data from each sub-band contain poorly correlated data. This is the basic concept of sub-band coding, which has become quite popular for median source encoding of speech (CROCHIERE *et al.*, 1976; JAYANT and NOLL, 1984; CROCHIERE and RABINER, 1983).

The basic idea of sub-band coding (SBC) (CROCHIERE *et al.*, 1976; JAYANT and NOLL, 1984; CROCHIERE and RABINER, 1983) is to decompose the input signal into narrow sub-bands. Each band is then sampled (or resampled) at its Nyquist rate (twice the width of the band) and digitally encoded with a PCM or other method. In this process, each sub-band can be encoded according to perceptual criteria that are specific to that band. On reconstruction, the sub-band signals are decoded and summed to give a close replica of the original signal.

Encoding in sub-bands offers several advantages. By appropriately allocating the bits in different bands, the number of quantiser levels, and hence the reconstruction error variance, can be separately controlled in each band, and the shape of the overall reconstruction error spectrum can be controlled as a function of frequency. In the lower frequency bands, which contain most of the ECG spectrum energy, a larger number of bits per sample can be used, whereas in upper frequency bands, which contain noise-like signals, fewer bits per sample can be used. Further, quantisation noise can be contained within bands to prevent the masking of a low-level input frequency range by quantising noise in another frequency range (Jayant and Noll, 1984).

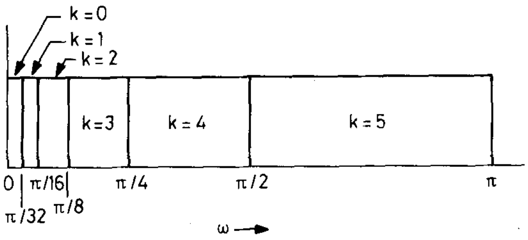

Fig. 1 *Octave-band stacking arrangement for filter bank channels*

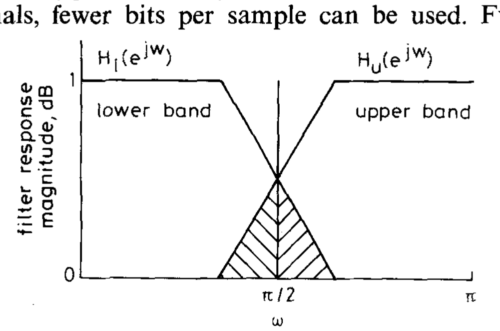

Fig. 2 *Quadrature mirror filtering for splitting an input into two equal-width sub-bands: a qualitative illustration of a filter-bank response that provides alias-image cancellation (Esteban and Galand, 1977)*

In general, the sub-band coding system consists of two basic components

(i) the analysis/synthesis subsystem, which is composed of a filter bank along with decimators/interpolators which are used for filtering, down-sampling and up-sampling operations
(ii) the coding subsystem, which generally employs some form of quantisation and bit allocation strategy.

Naturally, both components must be carefully designed to achieve good performance.

The filter bank that is used in our system is of the octave-spaced design (Nelson *et al.*, 1972; Barabell and Crochiere, 1979) with six bands, as shown in Fig. 1. The six-band octave band has a stacking arrangement in which each band k is twice the width of the proceeding band $k - 1$ (except for the $k = 0$, 1 bands). Experiments show that this six-band octave band successfully decomposes the ECG waveforms into six very poorly correlated data.

It is well known that in practical implementation, aliasing, frequency distortion and phase distortion resulting from decimation and filtering may occur if the filter bank is not properly designed. This effect is found to be unacceptable (Crochiere *et al.*, 1976; Crochiere and Rabiner, 1983; Jayant and Noll, 1984) in speech processing and needs to be removed. Yet, this problem is very elegantly tackled in the quadrature mirror filter (QMF) bank approach of Fig. 2 (Esteban and Garland, 1977). This figure shows the division of a full-band signal of a maximum radian frequency π into two signals of equal width using a constrained pair of low-pass and high-pass filters. By repeated subdivisions of resulting sub-bands using the QMF filter bank, an SBC filter bank with K given by a power of two can be realised. Values of K that are not powers of two can also be realised by simply ignoring appropriate sub-band branches in the QMF tree. For example, Fig. 3 shows the tree-structured realisation of a $K = 6$ band filter bank that is used in our system. The software implementation of the QMF bank we use is a 32-tap filter $h(n)$ designated as 32C in Crochiere and

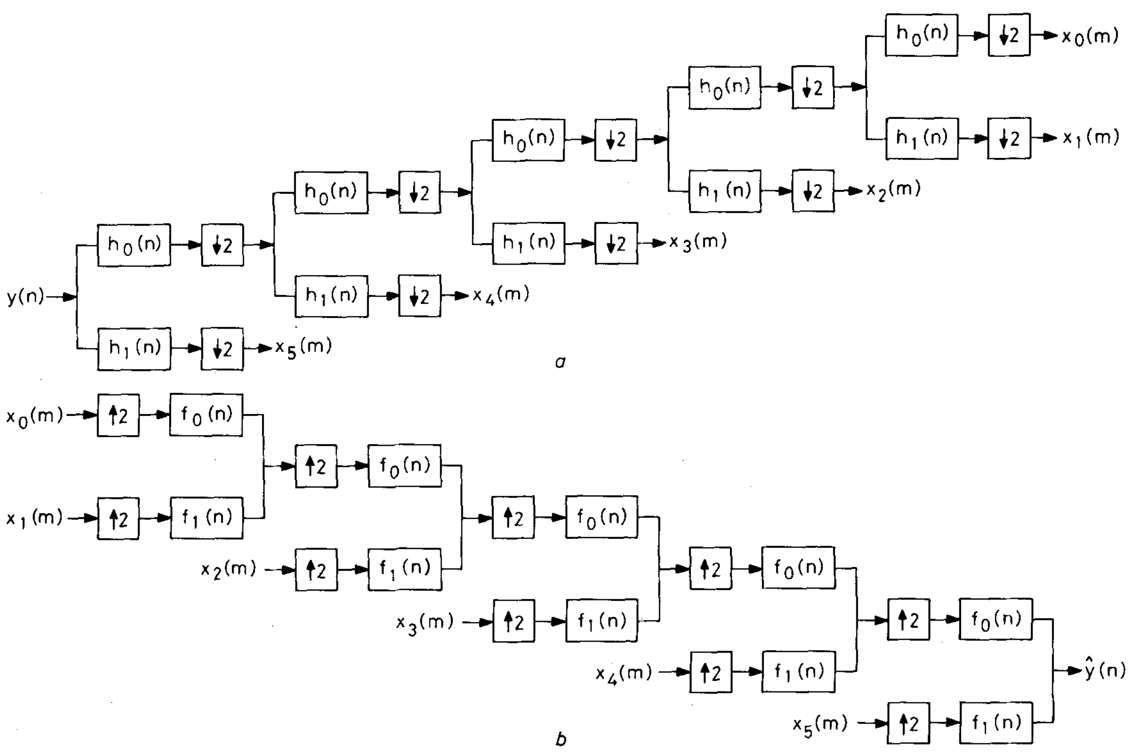

Fig. 3 *Tree-structured realisation for a six-band filter bank: (a) analysis filter bank; (b) synthesis filter bank*

Rabiner (1983), while in Fig. 3,

$$\begin{cases} h_0(n) = h(n) & \text{for all } n \\ h_1(n) = (-1)^n h(n) & \text{for all } n \end{cases}$$

$$\begin{cases} f_0(n) = h(n) & \text{for all } n \\ f_1(n) = -(-1)^n h(n) & \text{for all } n \end{cases}$$

3 Sources of ECG waveforms used in the research

For the compression tests, we used the ECG waveforms from the MIT/BIH database and the Veterans General Hospital's clinical data. The MIT/BIH database contains twin-channel ECG waveforms of 30 min each from 48 patients (two leads per patient). As for the Veterans

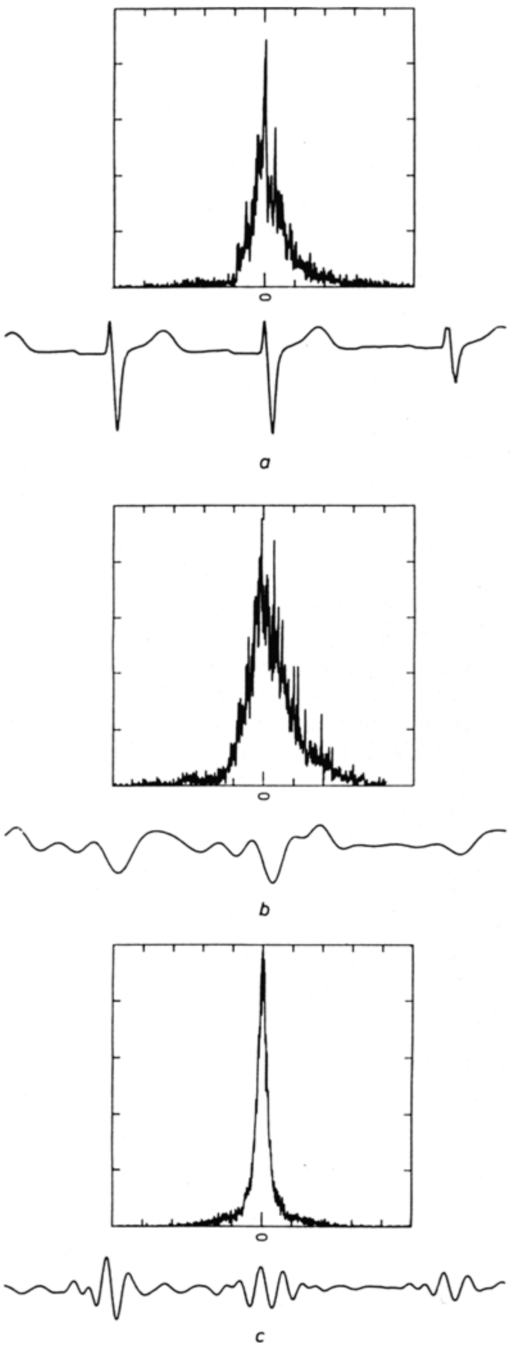

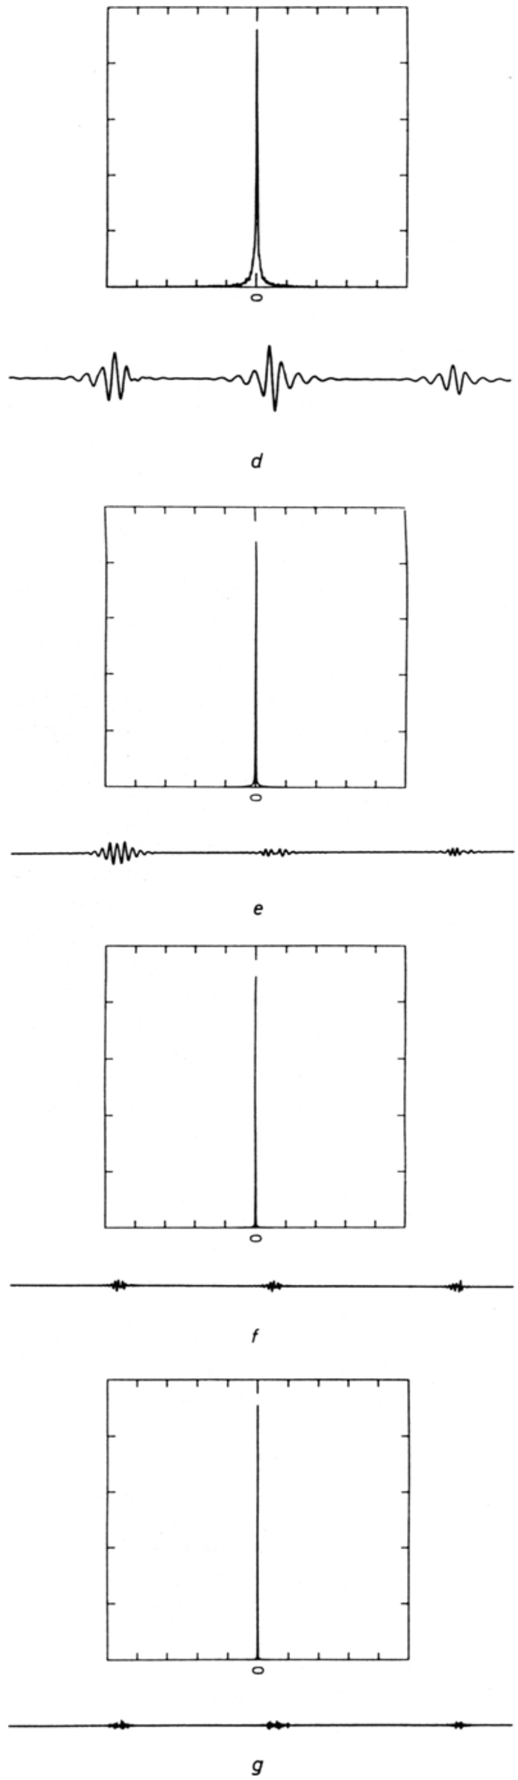

Fig. 4 *Some results of the six-band sub-band decomposition for an ECG waveform. Each figure shows a sub-band signal and its quantiser level distribution: (a) original ECG; (b) the zeroth band signal; (c) the first band signal; (d) the second band signal; (e) the third band signal; (f) the fourth band signal; (g) the fifth band signal. The third, fourth and fifth sub-band signals have been magnified by a factor of two, two and four, respectively*

General Hospital's clinical data, it contains single-channel ECG waveforms of 30 s each from 63 patients (one lead per patient). Without loss of generality, and for convenience, we shall use the ECG waveforms from the Veterans General Hospital's clinical data to depict our system throughout this paper. The same idea, with slight modification, applies to the MIT/BIH database.

4 Encoding of the sub-bands

4.1 *Overview*

In the sub-band analysis process, we have used the tree-structured QMF bank shown in Fig. 3 to decompose the ECG signal into six sub-bands. As each sub-band signal is sampled at twice the width of that band, the total samples of the original signal equal the sum of the samples of each sub-band signal. It seems that the sub-band analysis does not result in any data compression. However, it yields a desirable separation of the input signal. Consider Fig. 4, which shows some results of the six-band decomposition. As most of the ECG spectrum energy is concentrated around the low-frequency components (MEAD *et al.*, 1979), the lower bands are vital for later reconstruction and therefore should be very carefully encoded. The higher bands contain noise-like waveforms that can be considered as a 'refinement' of the waveform synthesised from the lower band signals. Consider the quantiser levels distribution shown in Fig. 4. It is interesting to observe that the higher bands have the quantiser level distributions highly concentrated at or around zero and the variances highly reduced compared with those of the input signal. This interesting phenomenon tells us that higher bands contain less information and so we can use fewer bits per sample to encode them.

4.2 *Encoding strategies for each sub-band signal*

As the lower bands contain most of the ECG spectrum energy, a larger number of bits is used. This implies that more quantiser levels are used. Consider the ECG waveforms from the Veterans General Hospital's clinical data which is shown in Fig. 4*a*. Originally, each sample contains 12 bits. After being decomposed by the tree-structured filter bank, the lowest two bands, i.e. the zeroth and the first bands, contain only 1/32 of the amount of samples in the input due to the decimation process. Therefore, the sample-to-sample correlations of these two band signals are no longer as strong as the original signal. This phenomenon also occurs in the higher sub-band signals. The original highly correlated data have now been successfully 'broken' into six poorly correlated data. Our simulation results show that the correlation of each sub-band signal is so reduced that, applying predictive coding techniques, such as DPCM, to encode each sub-band signal with allowance of acceptable distortion, will in fact need more bits per sample than using PCM. Therefore, in the following paragraphs, we show the techniques that are used in our coding subsystem.

The signal from each sub-band is first requantised with a different number of quantiser levels. The lower sub-band signals are requantised with more quantiser levels so that they may be more accurately preserved. The number of quantiser levels for the signal from each sub-band are listed below, where the zero quantiser level for the fifth sub-band signal means that zero bits per sample are allocated for this sub-band signal because it contains so little spectrum energy that it can be totally ignored, as depicted in Fig. 4.

(i) zeroth sub-band: 2048 quantiser levels
(ii) first sub-band: 2048 quantiser levels
(iii) second sub-band: 1024 quantiser levels
(iv) third sub-band: 512 quantiser levels
(v) fourth sub-band: 256 quantiser levels
(vi) fifth sub-band: zero quantiser levels.

The quantiser levels of each sub-band signal are Huffman coded, where the Huffman codes for the quantiser levels of each sub-band are obtained by considering the 63 clinical data from the Veterans General Hospital (every ECG waveform from the MIT/BIH database) as the training set. We compute the average number of bits needed for each sample of its corresponding sub-band. They are listed below:

(i) zeroth sub-band: 7·44 bits per sample
(ii) first sub-band: 5·73 bits per sample
(iii) second sub-band: 3·66 bits per sample
(iv) third sub-band: 1·58 bits per sample
(v) fourth sub-band: 1·07 bits per sample.

As can be seen, the average number of bits per sample is less for higher sub-bands. This is due to the number of quantiser levels used for each sub-band, as we described in the previous paragraph and the fact that the quantiser level distributions become much more concentrated at or around zero and the variances decrease in the higher bands, as depicted in Fig. 4.

The average number of bits for each sample in the third and fourth sub-bands, as shown in (iv) and (v), respectively, is further reduced by run length encoding. The strategies are described in the next two paragraphs.

As can be seen from Fig. 4, the quantiser level zero dominates for higher bands. Experiments show that a long run of zeros occurs frequently for higher bands, in particular, the third and fourth bands. Thus, the quantiser level zero is specially encoded by run length coding, *i.e.* it is coded by the Huffman code for zero, the length of the run of zeros. As we need extra bits to represent the length of a run of zeros, and the length of a run of zeros is related to the sub-band to which it belongs, we also have a program to select the number of bits to represent the lengths of the runs for band three and band four, respectively. The values are listed below:

(i) third sub-band: six bits (runs of length 1-64)
(ii) fourth sub-band: eight bits (runs of length 1-256).

For the third and fourth bands, a sequence of samples with quantiser levels not equal to zero and shorter than or equal to two is ignored and the sample(s) is (are) rescaled to zero. Observations show that these short sequences contribute little to the subjective quality of the reconstructed waveforms, and the subsequent rescaling to zero will enlarge the length of a run of zeros and therefore makes the run length coding strategy more effective.

5 Simulation results

Owing to strong electromagnetic fields of 50 or 60 Hz present in the modern hospital and home environment, ECG signals under test are first filtered by fast digital filters to subtract this interference (LYNN, 1977; TAYLOR and MACFARLANE, 1974; LEVKOV *et al.*, 1984). The ECG waveforms of the 63 clinical data from the Veterans General Hospital (and both channels of all the 48 patients from the MIT/BIH database) in their original form are then compressed by our sub-band coding system. Fig. 5 shows some ECG signals in their original form and their reconstructed signals. The subjective quality of each reconstructed ECG waveform as well as the bit rate and SNR (signal-to-noise ratio) are shown, where the SNR of the reconstructed ECG waveform $\hat{y}(n)$ with respect to the original ECG waveform $y(n)$ is defined as

$$\text{SNR} = 10 \log \frac{\sum_n y(n)^2}{\sum_n \{\hat{y}(n) - y(n)\}^2}$$

We list in Table 1 some statistics of the results. As can be seen from this table, an average of 0·81 bits per sample can be achieved with the high SNR value 29·97 dB. Signal fidelities are very well preserved. Note that the data rate

Table 1 Some statistics

	Maximum bits per sample	Minimum bits per sample	Average bits per sample
	1·131042	0·649722	0·813992
SNR (dB)	34·410	29·130	29·968

for each ECG signal depends on the 'complexity' of the signal, i.e. the data complexity. For example, the worst case bit rate and the best case bit rate are 1·13 bits per sample and 0·65 bits per sample, respectively. We consider this to be a normal phenomenon as a good data compressor uses more bits for a complicated signal and fewer bits for a simple signal so that the quality of the compressed signal is kept constant.

6 Discussions and future research directions

Sub-band coding has been successfully applied to speech signals (Jayant and Noll, 1984). As ECG signals have properties similar to those of speech signals, we can apply this coding strategy to ECG data. Experimental results show that the sub-band coding technique is also a promising coding technique for ECG waveforms.

The analysis filter bank decomposes the original ECG waveform into six sub-band signals. Will the performance be better if we decompose the original signals into more sub-band signals? Based on the following considerations, we did not use more bands:

(*a*) Using more bands increases the hardware complexity. As the overall system is to be implemented by hardware, we must take this complexity into consideration

(*b*) The sub-band signals achieved are already poorly correlated. Thus, it is not worth increasing the number of sub-bands which increases the hardware complexity.

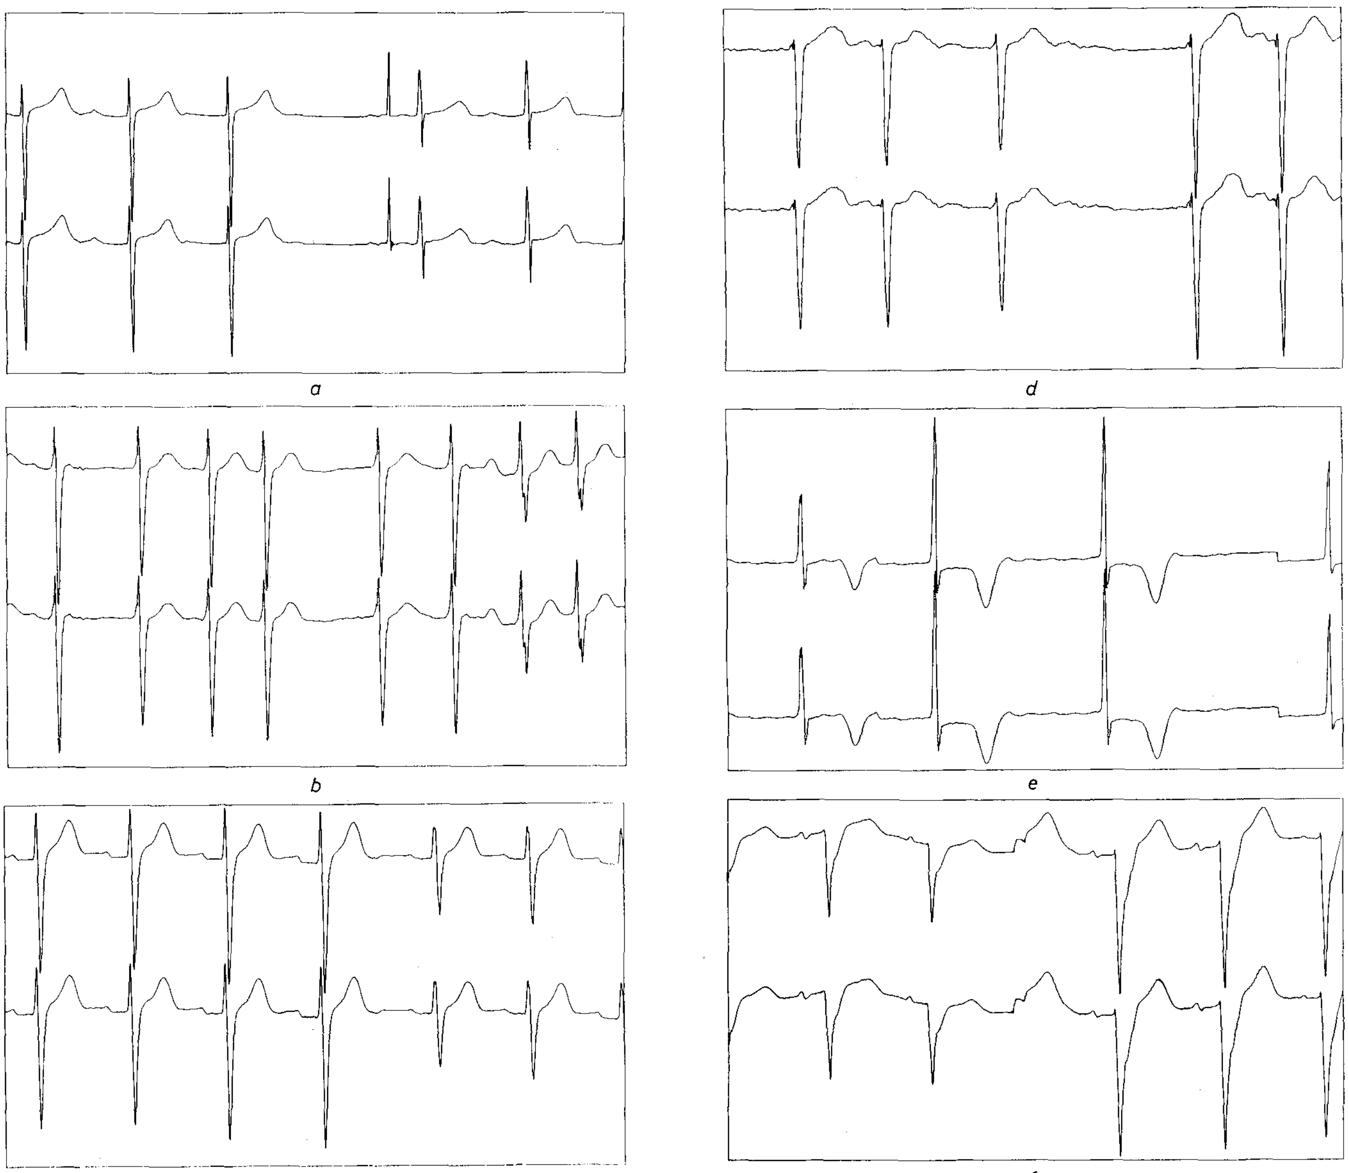

Fig. 5 *Some results for the reconstructed signal. In each block, the top waveform is the original signal and the lower one is the reconstructed signal: (a) 0·93 bits per sample (SNR = 27·78 dB); (b) 0·95 bits per sample (SNR = 29·84 dB); (c) 0·90 bits per sample (SNR = 31·13 dB); (d) 0·98 bits per sample (SNR = 29·80 dB); (e) 0·73 bits per sample (SNR = 29.93 dB); (f) 0·78 bits per sample (SNR = 32·85 dB)*

As the lower band signals contain most of the spectrum energy, a larger number of bits are allocated to these sub-band signals. The sub-band signals from the higher bands contain noise-like signals which are not vital to our later reconstruction and therefore are encoded with less bits. In the coding system, samples from each sub-band are requantised, Huffman coded and run length coded if they belong to higher bands and are of quantiser level zero. The encoded sub-band signals are then synthesised by the synthesis filter bank. We observe that signal fidelities are very well preserved even for a very low bit rate. This is because of our very conservative encoding of each sub-band signal, especially of the lower bands. In our future research, we shall improve the coding strategies for each sub-band signal so that a much lower bit rate can be achieved without lowering the fidelities.

The signal from the fifth sub-band is totally ignored in our system. This is proper for most cases. However, for a very special case, it may introduce some artefact which is not desired. For example, consider the fourth QRS complex in Fig. 5*a* which shows additional waves which look like 'ripples'. These appear because this QRS complex changes abruptly, just like a square wave. This is a rather special case of a QRS complex. High frequency components are indispensable for such a waveform. In our future research, the coding subsystem will be improved to be more flexible. The encoding strategies for the higher bands, instead of being fixed, will depend on the power contained in those bands.

Consider the hardware needed to accomplish the overall system. Both the analysis filter bank and the synthesis filter bank can be implemented by pipeline architectures; each consisting of five stages. In each stage, there is one high-pass filter and one low-pass filter. Each filter is made by simple convolution (OPPENHEIM and SCHAFER, 1975), in which 32 real multiplications and 31 real additions are needed. As 2 ms is long enough for chips available off the shelf to accomplish 32 real multiplications and 31 real additions, the analysis/synthesis subsystem can thus be implemented in a real-time environment. The encoding (decoding) subsystem serves as the last (first) stage of the pipeline analysis (synthesis) filter. The overall system can be processed in real time if the hardware for the encoding (decoding) subsystem takes 2 ms. This is not at all difficult because it needs only simple operations. Thus, we conclude that the overall system can be implemented in a real-time environment by a pipeline architecture.

In conclusion, the first attempt of sub-band coding on ECG waveforms shows acceptable results. We believe that it is an interesting area that is worth further investigation. In our future research, we shall improve the coding subsystem such that the average compression result is less than 0·5 bits per sample.

Acknowledgment—This research work was partially supported by the National Science Council of the Republic of China under contract NSC81-0404-E-006-119.

References

BARABELL, A. J. and CROCHIERE, R. E. (1979) Sub-band coder design incorporating quadrature filters and pitch prediction. Proc. ICASSP, Washington, DC, 530–533.

CHEN, J. J., KAO, C. D., CHEN, C. Y. and YU, B. C. (1984) Computer diagnosis of electrocardiograms. II: 3-channel ECG data transmission system for arrhythmia detection. *Chinese Med. J.*, **33**, 159–165.

CROCHIERE, R. E., WEBBER, S. A. and FLANAGAN, J. L. (1976) Digital coding of speech in sub-bands. *Bell System Tech. J.*, **55**, 1069–1085.

CROCHIERE, R. E. and RABINER, L. R. (1983) *Multirate digital signal processing*. Englewood Cliffs, Prentice-Hall, New Jersey.

ESTEBAN, D. and GALAND, C. (1977) Application of quadrature mirror filters to split band voice coding schemes. Proc. ICASSP, Hartford, CT, 191–195.

JAYANT, N. S. and NOLL, P. (1984) *Digital coding of waveforms*. Englewood Cliffs, Prentice-Hall, New Jersey.

KAO, C. D., CHEN, J. J. and YU, B. C. (1983) A microcomputer-based electrocardiographic data transmission system. *Proc. N. S. C. ROC(A)*, **7**, 275–284.

LEVKOV, C., MICHOV, G., IVANOV, R. and DASKALOV, I. K. (1984) Subtraction of 50 Hz interference from the electrocardiogram. *Med. & Biol. Eng. & Comput.*, **22**, 371–373.

LYNN, P. A. (1977) Online digital filters for biological signals: some fast designs for a small computer. *Med. & Biol. Eng. & Comput.*, **15**, 534–540.

MEAD, C. N., CLARK, K. W., POTTER, S. J., MOORE, S. M. and THOMAS, L. J. (1979) Recent process in frequency-domain analysis of the ECG. *IEEE Comp. in Cardiol.*, 43–47.

NELSON, G. A., PFEIFER, L. L. and WOOD, R. C. (1972) High speed octave band digital filtering. *IEEE Trans.*, **AU-20**, (1), 8–65.

OPPENHEIM, A. V. and SCHAFER, R. W. (1975) *Digital signal processing*. Englewood Cliffs, Prentice-Hall, New Jersey.

RIPLEY, K. L. (1976–1989) *Computers in cardiology*. IEEE

TAYLOR, T. P. and MACFARLANE, P. W. (1974) Digital filtering of the ECG–a comparison of low-pass digital filters on a small computer. *Med. & Biol. Eng. & Comput.*, **12**, 493–502.

Author's biography

S. C. Tai was born in Hsinchu, Taiwan, in 1960. He received the BS and MS degrees in Electrical Engineering from the National Taiwan University, Taipei, Taiwan, in 1982 and 1986, respectively, and a Ph.D. degree in computer Science from the National Tsing Hua University, Hsinchu, Taiwan, in 1989. He is currently an Associate Professor of Electrical Engineering at the National Cheng Kung University, Tainan, Taiwan. His teaching and research interests include design automation of VLSI, data compression, biomedical engineering and computer algorithms.