Wavelet Analysis and Signal Processing

Ronald R. Coifman, Yves Meyer, and Victor Wickerhauser
Yale University, New-Haven, Ct 06520, USA

Wavelet Analysis consists of a versatile collection of tools for the analysis and manipulation of signals such as sound and images,as well as more general digital data sets. The user is provided with a collection of standard libraries of waveforms , which can be chosen to fit specific classes of signals . These libraries come equipped with fast numerical algorithms enabling realtime implementation of a variety of signal processing tasks, such as data compression, extraction of parameters for recognition and diagnostics, transformation and manipulation of data. The process of analysis of data is usually started by comparing acquired segments of data with stored waveforms.

As can be seen at the top portion Figure 1, representing a segment of a recording of the word armadillo, voice signals consist of modulated oscillations of small duration and varying frequencies and intensity. Figures (3-16) represent a variety of waveforms selected from different libraries, as well as illustrations of analysis tasks performed on them as described in the figure captions.

A general signal (Figure 1 or Figure 2 for example) is a superposition of different structures occurring on different time scales at different times (or spacial scales at different locations). One purpose of analysis is to separate and sort these structures. The example of music (or voice) can be used to illustrate some of these ideas. A musical note can be described by four basic parameters, intensity (or amplitude), frequency, time duration, time position. Wavelet packets or trigonometric wave forms are indexed by the same parameters, plus others corresponding to choice of library (we can think of a library as a musical instrument, i.e. the recipe used to generate all the waveforms, notes, in the library).

Figure 1

The first 1024 samples ($\frac{1}{8}$ second) of the word armadillo, are plotted on the top part. The library of local sine waveforms is then used to select the combination of windows of highest efficiency (lowest entropy). Expansion coefficients are then ordered by window in decreasing order. The top 5% are plotted in the center and used to reconstruct a compressed form of the signal which is plotted below.

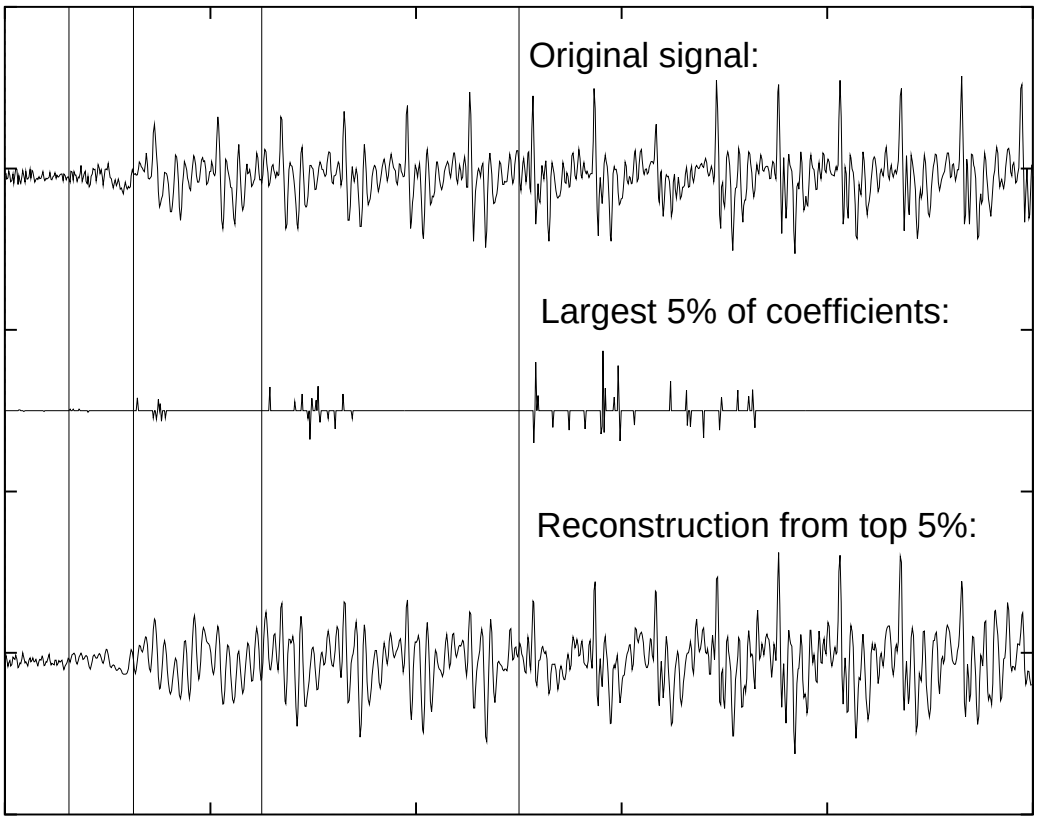

Figure 2

Automatic segmentation of a voice recording (armadillo) by using least entropy windowing in the local sine library. The windows are selected to obtain optimal efficiency in representing the signal. We see different patterns fall into distinct windows.

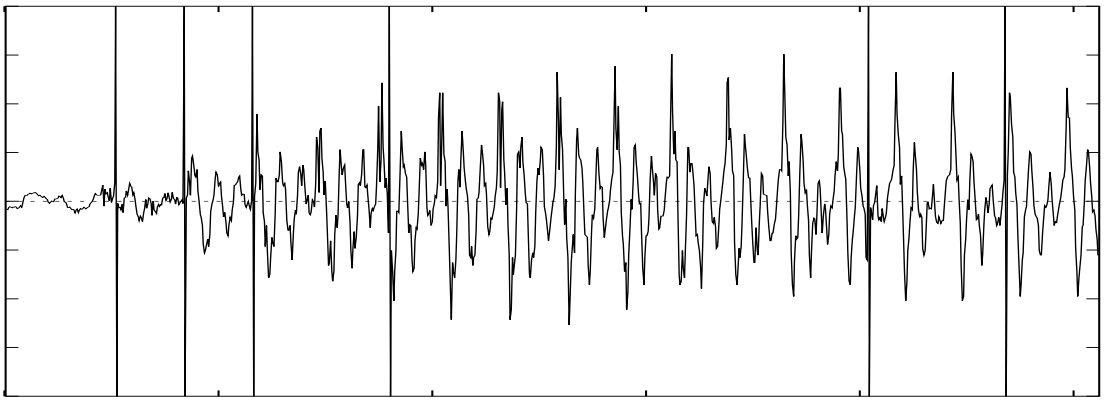

The process of analysis compares a sound (or other signals) with all elements of a given library, picks up large correlations (or notes which are closest to segments of the signal) and uses these to rebuild the signal with a minimal number of waveforms. The result provides an economical transcription, which if ordered by decreasing intensity sorts the main features in order of importance. (In Figure 1 the signal has been segmented in windows, the top 5% of the expansion kept, and used to rebuild the signal in the bottom half).[1]

This realization of the signal in terms of the "best basis" providing efficient superpositon in terms of oscillatory modes on different time scales, can be used to compress signals for digital transmission and storage.

Of more importance for applications is the ability to compute and manipulate data in compressed parameters. This feature is particularly important for recognition and diagnostic purposes.

As an illustration consider a hypothetical diagnostic device for heartbeats, in which fifty

[1]See Figures 8-15 for examples of analysis of various typical examples.

consecutive beats are recorded. We would like to use this data as a statistical foundation for detection of significant changes in the next batch of beats. Theoretically this can be done by factor analysis (or Karhunen-Loéve bases), unfortunately the computation involving raw data is too large to be useful. The ability to efficiently compress the recorded data in terms of a single statistical best basis (or the representation of the data in terms of an adapted efficient coordinate system) enables us to perform a factor analysis (if needed) and to compute the deviation of the next few heartbeats from their predecessors, thereby detecting significant changes on the fly.

This procedure, in which we first compress a large data set of measurements, in order to compute with the compressed parameters, can reduce dramatically the time needed to compute and manipulate data, it generalizes the usual transform methods (like FFT) by building an adapted fast transform for various classes of data or of operations on that data (as an example the data could consist of a three dimensional atmospheric pressure map, and the computation would involve the evolution of the pressure. In this case it is natural to break the computation as a sum of interactions on different scales, and some limited interaction between adjacent scales, such breakup is automatic if the pressure map is expressed in the wavelet basis, which in this case is also the natural choice for compression of the data.) Demo software is available on anonymous ftp from Yale [**6**].

Definitions of Modulated Waveform Libraries. We now introduce the concept of a "Library of orthonormal bases". For the sake of exposition we restrict our attention to two classes of numerically useful waveforms, introduced recently [**2**][**3**].

We start with trigonometric waveform libraries. These are localized sine transforms LST associated to covering by intervals of $\mathbf{R}$ (more generally, of a manifold).

We consider a cover $\mathbf{R} = \bigcup_{-\infty}^{\infty} I_i$ $I = [\alpha_i \alpha_{i+1})$ $\alpha_i < \alpha_{i+1}$, write $\ell_i = \alpha_{i+1} - \alpha_i = |I_i|$ and let $p_i(x)$ be a window function supported in $[\alpha_i - \ell_{i-1}/2, \alpha_{i+1} + \ell_{i+1}/2]$ such that

$$\sum_{-\infty}^{\infty} p_i^2(x) = 1$$

and

$$p_i^2(x) = 1 - p_i^2(2\alpha_{i+1} - x) \quad \text{for} \quad x \quad \text{near} \quad \alpha_{i+1}$$

then the functions

$$S_{i,k}(x) = \frac{2}{\sqrt{2\ell i}} p_i(x) \sin[(2k+1)\frac{\pi}{2\ell_i}(x - \alpha_i)]$$

form an orghonormal basis of $L^2(\mathbf{R})$ subordinate to the partition p_i. The collection of such bases forms a library of orthonormal bases. See Figure 3.

Figure 3

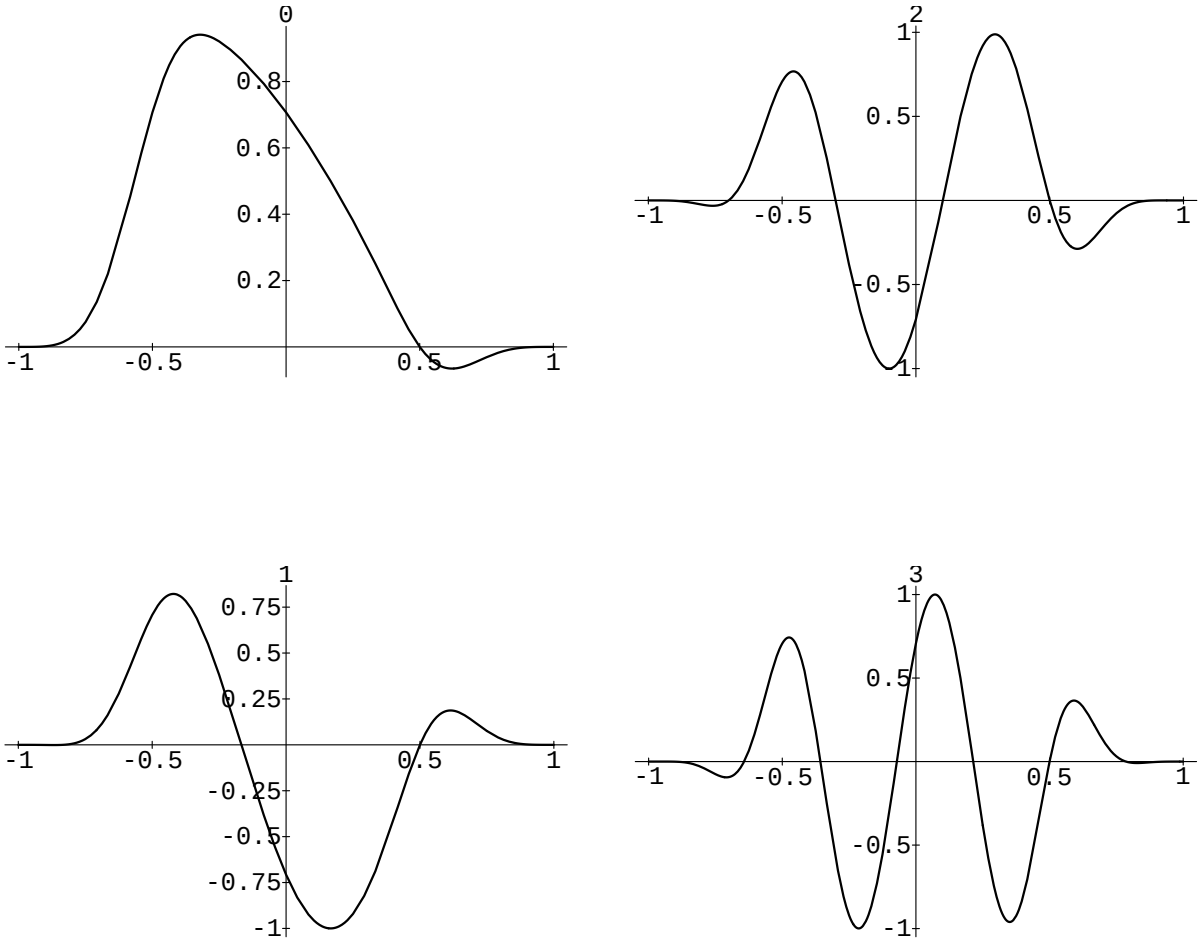

Figure 4

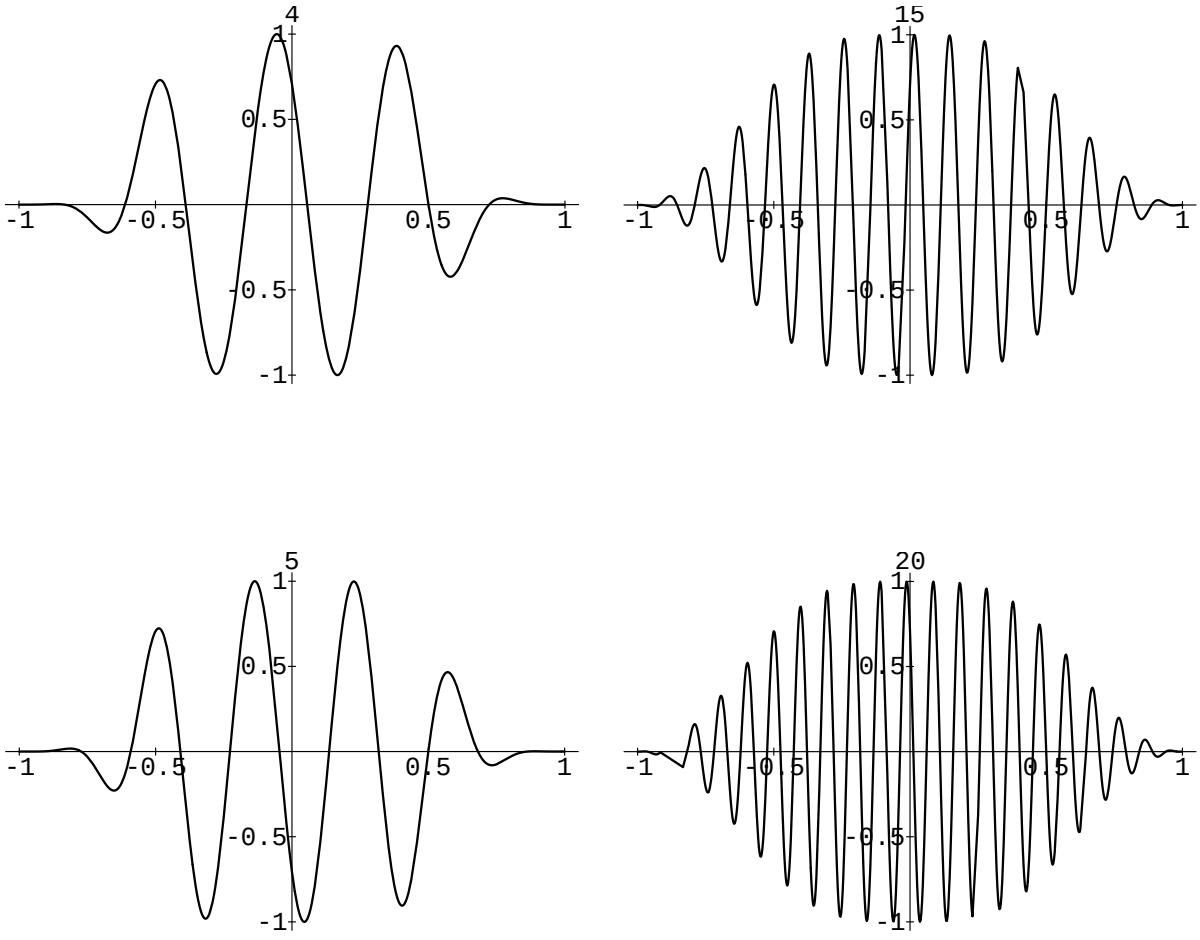

It is easy to check that if H_{I_i} denotes the space of functions spanned by $S_{i,k}$ $k = 0, 1, 2, ...$ then $H_{I_i} + H_{I_{i+1}}$ is spanned by the functions

$$P(x)\frac{1}{\sqrt{2(\ell_i + \ell_{i+1})}}\sin[(2k+1)\frac{\pi}{2(\ell_i + \ell_{i+1})}(x - \alpha_i)]$$

where

$$P^2 = p_i^2(x) + p_{i+1}^2(x)$$

is a "window" function covering the interval $I_i \cup I_{i+1}$.

Another new library of orthonormal bases called the Wavelet packet library can be constructed. This collection of modulated wave forms, corresponds roughly to a covering

of "frequency" space. This library contains the wavelet basis, Walsh functions, and smooth versions of Walsh functions called wavelet packets. See Figure 5.

Figure 5

A few wavelet packets, $W_0(x) \cdots W_7(x)$ from C-18 by using the relations (1.2). These waveforms are mutually orthogonal, moreover, each of them is orthogonal to all of its integer translates and dyadic rescaled versions. The full collection of these wavelet packets (including translates and rescaled versions) provides us with a library of "templates" or "notes" which are matched "efficiently" to signals for analysis and synthesis.

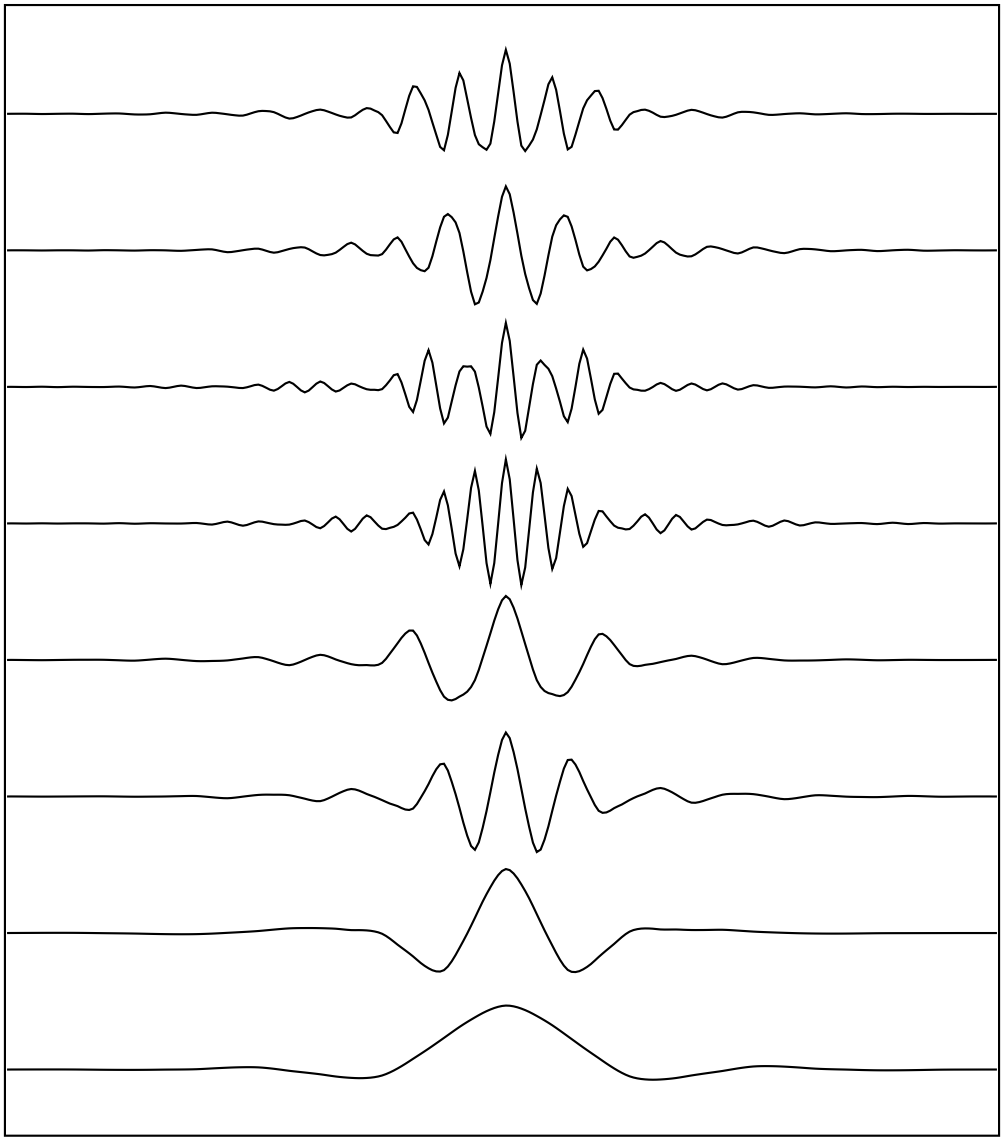

We'll use the notation and terminology of [**4**], whose results we shall assume.

We are given an exact quadrature mirror filter $h(n)$ satisfying the conditions of Theorem (3.6) in [**4**], p. 964, i.e.

$$\sum_n h(n-2k)h(n-2\ell) = \delta_{k,\ell},\ \sum_n h(n) = \sqrt{2}.$$

We let $g_k = h_{l-k}(-1)^k$ and define the operations F_i on $\ell^2(\mathbf{Z})$ into "$\ell^2(2\mathbf{Z})$"

$$F_0\{s_k\}(i) = 2\sum s_k h_{k-2i} \tag{1.0}$$
$$F_1\{s_k\}(i) = 2\sum s_k g_{k-2i}.$$

The map $\mathbf{F}(s_k) = F_0(s_k) \oplus F_1(s_k) \in \ell^2(2\mathbf{Z}) \oplus \ell^2(2\mathbf{Z})$ is orthogonal and

$$F_0^* F_0 + F_1^* F_1 = I \tag{1.1}$$

We now define the following sequence of functions.

$$\begin{cases} W_{2n}(x) = \sqrt{2}\sum h_k W_n(2x-k) \\ W_{2n+1}(x) = \sqrt{2}\sum g_k W_n(2x-k). \end{cases} \tag{1.2}$$

Clearly the function $W_0(x)$ can be identified with the scaling function φ in [**4**] and W_1 with the basic wavelet ψ.

Let us define $m_0(\xi) = \frac{1}{\sqrt{2}}\sum h_k e^{-ik\xi}$ and

$$m_1(\xi) = -e^{i\xi}\bar{m}_0(\xi+\pi) = \frac{1}{\sqrt{2}}\sum g_k e^{ik\xi}$$

Remark. The quadrature mirror condition on the operation $\mathbf{F} = (F_0, F_1)$ is equivalent to the unitarity of the matrix

$$\mathcal{M} = \begin{bmatrix} m_0(\xi) & m_1(\xi) \\ m_0(\xi+\pi) & m_1(\xi+\pi) \end{bmatrix}$$

Taking the Fourier transform of (1.2) when $n = 0$ we get

$$\hat{W}_0(\xi) = m_0(\xi/2)\hat{W}_0(\xi/2)$$

i.e.,

$$\hat{W}_0(\xi) = \prod_{j=1}^{\infty} m_0(\xi/2^j)$$

and

$$\hat{W}_1(\xi) = m_1(\xi/2)\hat{W}_0(\xi/2) = m_1(\xi/2)m_0(\xi/4)m_0(\xi/2^3)\cdots$$

More generally, the relations (1.2) are equivalent to

$$\hat{W}_n(\xi) = \prod_{j=1}^{\infty} m_{\varepsilon_j}(\xi/2j) \tag{1.3}$$

and $n = \sum_{j=1}^{\infty} \varepsilon_j 2^{j-1} (\varepsilon_j = 0 \text{ or } 1)$.

The functions $W_n(x-k)$ form an orthonormal basis of $L^2(\mathbf{R}^n)$.

We define a library of wavelet packets to be the collection of functions of the form $W_n(2^\ell x - k)$ where $\ell, k \in \mathbf{Z}, n \in N$. Here, each element of the library is determined by a scaling parameter ℓ, a localization parameter k and an oscillation parameter n. (The function $W_n(2^\ell x - k)$ is roughly centered at $2^{-\ell}k$, has support of size $\approx 2^{-\ell}$ and oscillates $\approx n$ times).

We have the following simple characterization of subsets forming orthonormal bases.

PROPOSITION. *Any collection of indices (ℓ, n) such that the intervals $[2^\ell n, 2^\ell n + 1)$ form a disjoint cover of $[0, \infty)$ gives rise to an orthonormal basis of L^2*[2].

Motivated by ideas from signal processing and communication theory we were led to measure the "distance" between a basis and a function in terms of the Shannon entropy of the expansion. More generally, let H be a Hilbert space.

Let $v \in H$, $\|v\| = 1$ and assume

$$H = \oplus \sum H_i$$

an orthogonal direct sum. We define

$$\varepsilon^2(v, \{H_i\}) = -\sum \|v_i\|^2 \ell n \|v_i\|^2$$

[2]We can think of this cover as an even covering of frequency space by windows roughly localized over the corresponding intervals.

as a measure of distance between v and the orthogonal decomposition.

ε^2 is characterized by the Shannon equation which is a version of Pythagoras' theorem.

Let

$$H = \oplus(\sum H^i) \oplus (\sum H_j)$$
$$= H_+ \oplus H_-$$

H^i and H_j give orthogonal decompositions $H_+ = \sum H^i, H_- = \sum H_j$. Then

$$\varepsilon^2(v; \{H^i, H_j\}) = \varepsilon^2(v, \{H+, H_-\} + \|v_+\|^2 \varepsilon^2 \left(\frac{v_+}{\|v_+\|}, \{H^i\}\right) + \|v_-\|^2 \varepsilon^2 \left(\frac{v_-}{\|v_-\|}, \{H_j\}\right)$$

This is Shannon's equation for entropy (if we interpret as in quantum mechanics $\|P_{H_+} v\|^2$ as the "probability" of v to be in the subspace H_+).

This equation enables us to search for a smallest entropy space decomposition of a given vector.

In fact, for the example of the first library restricted to covering by dyadic intervals we can start by calculating the entropy of an expansion relative to a local trigonometric basis for intervals of length one, then compare the entropy of an adjacent pair of intervals to the entropy of an expansion on their union. Pick the expansion of minimal entropy and continue until a minimum entropy expansion is achieved (see Figures 1 and 2).

In practice, discrete versions of this scheme can be implemented in $CN \ \log \ N$ computations (where N is the number of discrete samples $N = 2^L$.)

For voice signals and images this procedure leads to remarkable compression algorithms see below.

Of course, while entropy is a good measure of concentration or efficiency of an expansion, various other information cost functions are possible, permitting discrimination and choice between various special function expansion.

Other possible libraries can be constructed. The space of frequencies can be decomposed into pairs of symmetric windows around the origin ,on which a smooth partition of unity is

constructed. This and other constructions were obtained by one of our students E. Laeng [**L**].

Higher dimensional libraries can also be easily constructed,(as well as libraries on manifolds) leading to new and direct analysis methods for linear transformations.

We will describe an algorithm to produce a rectangle in which coefficients are grouped by frequency, since this is simpler and since the transformation to the other form is evident. For definiteness, consider a function defined at 8 points $\{x_1, \ldots, x_8\}$, i.e., a vector in $\mathbf{R}^8$. We may develop the (periodized) wavelet packet coefficients of this function by filling out the following rectangle:

x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8
s_1	s_2	s_3	s_4	d_1	d_2	d_3	d_4
ss_1	ss_2	ds_1	ds_2	sd_1	sd_2	dd_1	dd_2
sss_1	dss_1	sds_1	dds_1	ssd_1	dsd_1	sdd_1	ddd_1

A rectangle of wavelet packet coefficients.

Figure 6

Wavelet packets are generated by the same algorithm in which different QMF are selected, generating the shapes in Figure 5. Wavelet packets from the deepest level have the largest scale. Coefficients computed in the shaded box represent a correlation between the initial signal $x_1, \ldots, x_8$ and the waveform described on the left. Here we take the Haar QMF, i.e. $s_1 = \frac{1}{\sqrt{2}}(x_1 + x_2) \quad d_1 = \frac{1}{\sqrt{2}}(x_1 - x_2), \ldots$ etc. The waveforms produced are the classical Walsh functions.

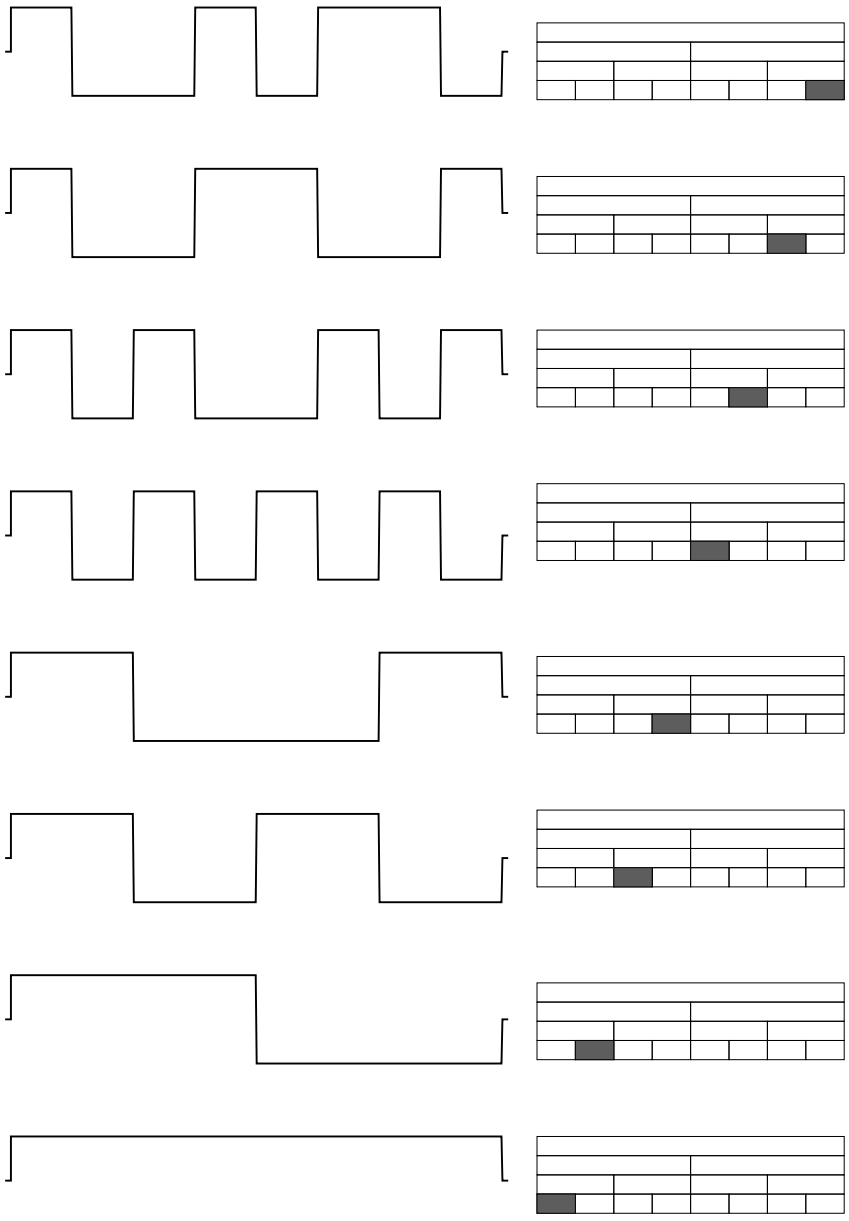

Figure 7

Wavelet packets from intermediate levels have a shorter time duration than the Walsh functions.

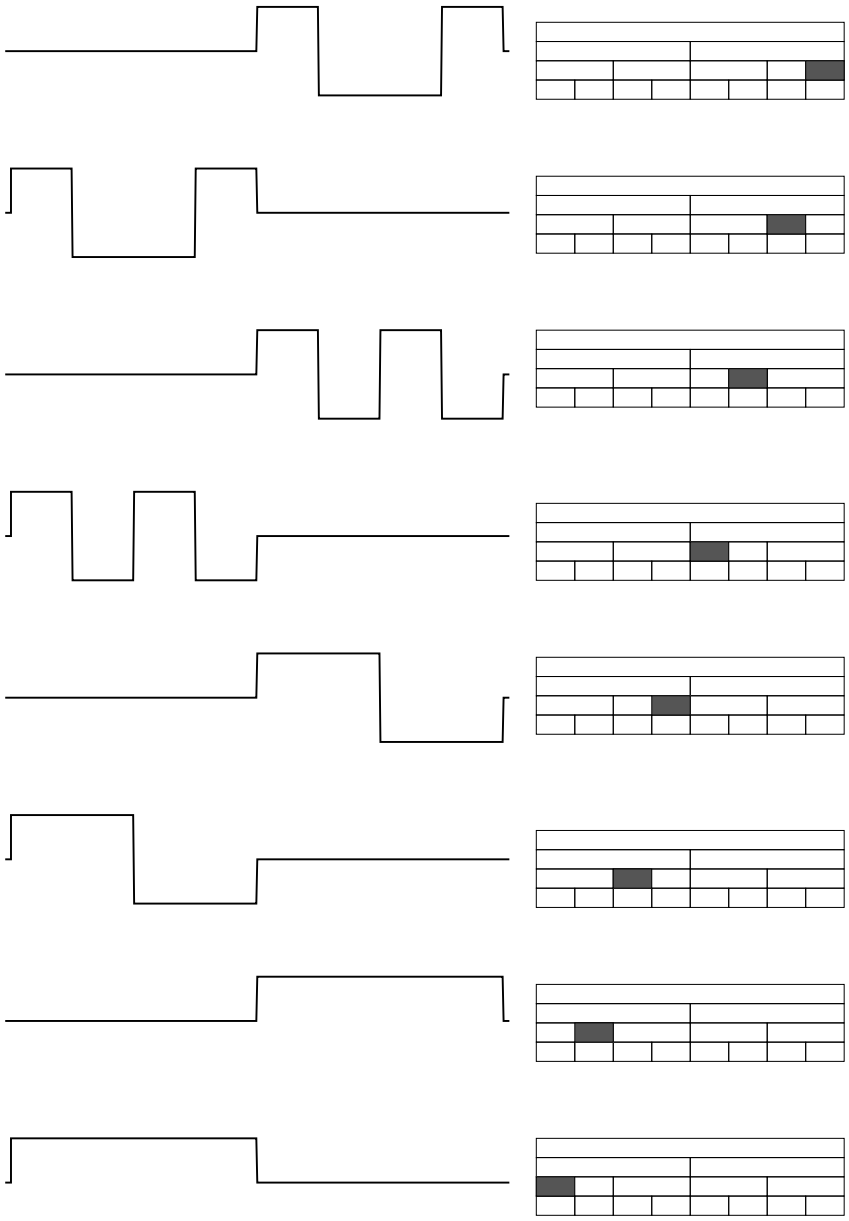

Figure 8
Wavelet packets from the first level have the smallest time duration.

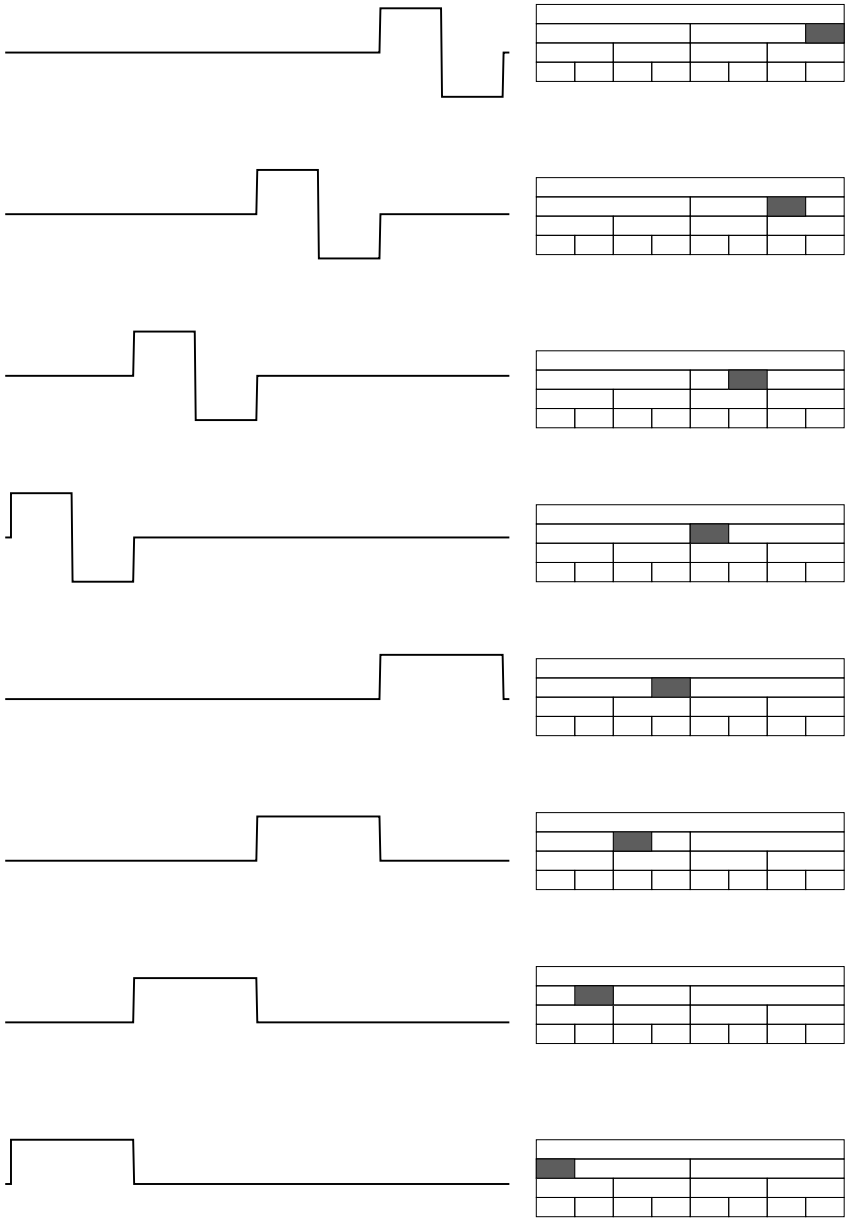

Each row is computed from the row above it by one application of either F_0 or F_1, which we think of as "summing" (s) or "differencing" (d) operations, respectively. Thus, for example the subblock $\{ss_1, ss_2\}$ is obtained by convolution-decimation of $\{s_1, s_2, s_3, s_4\}$ with F_0, while $\{ds_1, ds_2\}$ comes from similar convolution-decimation with F_1. In the simplest case, where we use the Haar filters $h = \{\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}\}$ and $g = \{\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}}\}$, we have in particular $ss_1 = \frac{1}{\sqrt{2}}(s_1 + s_2)$, $ss_2 = \frac{1}{\sqrt{2}}(s_3 + s_4)$, $ds_1 = \frac{1}{\sqrt{2}}(s_1 - s_2)$, and $ds_2 = \frac{1}{\sqrt{2}}(s_3 - s_4)$. The two daughter s and d subblocks on the $n+1$st row are determined by their mutual parent on the nth row, which conversely is determined by them through the adjoint anticonvolution.

Reconstructing the nth row from the $n+1$st row consists of applying F_0^* to the left daughter and F_1^* to the right daughter, then summing the images into the parent. In this manner, we generated the graphs of the functions which are included in fig 3. We used a rectangle of size 1024×10 to obtain 1024-point approximations. We filled the rectangle with 0's except for a single 1, then applied the deconvolutions F_0^* and F_1^* up to 10 times in various orders, so as to generate a vector of length 1024. This vector approximates one of the 10240 wavelet packets in R^{1024}. The details of this reconstruction determine the frequency, scale, and location parameters.

From this rectangle, we may choose subsets of N coefficients which correspond to orthonormal bases for $\mathbf{R}^N$. For example, the subset corresponding to the labelled boxes in the figure below is the wavelet basis.

<table>
<tr><td colspan="8"></td></tr>
<tr><td colspan="4"></td><td>d_1</td><td>d_2</td><td>d_3</td><td>d_4</td></tr>
<tr><td colspan="2"></td><td>ds_1</td><td>ds_2</td><td colspan="2"></td><td colspan="2"></td></tr>
<tr><td>sss_1</td><td>dss_1</td><td></td><td></td><td></td><td></td><td></td><td></td></tr>
</table>

The wavelet basis.

The two figures below give other orthonormal basis subsets.

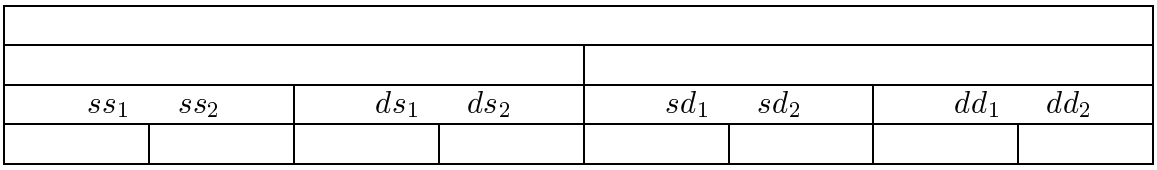

A subband basis.

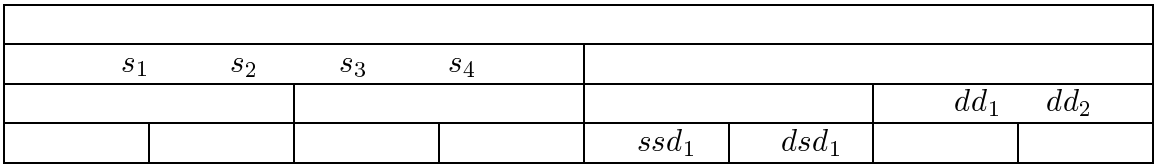

An orthonormal basis subset.

The boxes of coefficients in the rectangle have a natural binary tree structure. Each box is a direct sum of its two children. Call a subset of the rectangle a *graph* if it contains only whole boxes and each column of the rectangle has exactly one element. We have the following relation between dyadic coverings and o.n bases.

PROPOSITION. *Every graph is an orthonormal basis subset.*

The number of graphs may be counted by induction. If $N = 2^L$, let A_L be the number of graphs in the coefficient rectangle of N columns and L rows. Then $A_0 = 1$ and we have the relation $A_{L+1} = 1 + A_L^2$, which implies that $A_{L+1} > 2^{2^L} = 2^N$.

This last algorithm is beautifully suited for a best basis selection algorithm. By comparing the information cost of two "children" with their parent box we can, starting from the bottom of the rectangle, replace each node of the tree by the least costly combination.

If entropy is taken as information cost, the Shannon equation quarantees that we will end up with a basis with minimum entropy.

A simple variant on this selection algorithm permits the construction of a statistical best basis. Here we start with a collection of vectors X_n $n = 1, \ldots, N$ in $\mathbf{R}^d$ (for example, recording of successive distinct heartbeats). We construct the average vector $\bar{X} = \frac{1}{N} \sum_{n=1}^{N} X_n$

and would like to select in our library a basis which is most efficient, on the average, in compressing all vectors $\tilde{X}_n = X_n - \bar{X}$.

This is easily achieved by repeating the preceding search where, in each node of the tree, we compare the total cost (or entropy) of the node to the cost of its children. (Where by total cost we mean the sum of entropies of all vectors contributing to the node, or some other measure of information).

Of course this procedure is related to the Karhunen-Loeve expansion, in which we find in $\mathbf{R}^d$ the most efficient basis by diagonalizing the autocovariance matrix

$$M_{ij} = \frac{1}{N} \sum_{n=1}^{N} \tilde{X}_n(i) \tilde{X}_n(j) \quad i = 1, \ldots, d \; j = 1, \ldots, d.$$

Intuitively we think of the various sample vectors as forming an ellipsoidal cloud centered at $\bar{X}$, with the principal axis of this ellipsoid being the eigenvectors of M, pointing in the direction of maximum variance. This Karhunen-Loeve basis is the most efficient basis for capturing on the average most of the energy of a random sample. Ideally, to analyze the fluctuations of $\tilde{X}_n$ we should compute this basis and expand each sample in it. Numerically this task is too expensive when dealing with raw data. It is advisable to first find the statistical best basis within a library thereby compressing the data. In this new coordinate system one can compute the Karhunen-Loeve basis and compare its entropy to that of the best basis. If the two entropies are close, we have an indication that we selected the correct library for compression of the collection $\tilde{X}_n$.

The reduction in complexity achieved by finding a best basis provides a significant speedup in computation time, not only for finding the $K - L$ basis, but also for obtaining a fast algorithm to compute correlations with the $K - L$ basis. Normally to compute an expansion in $K - L$ will take d^2 computations (if $N \simeq d$) by proceeding through the best basis. This can be reduced to $Cd \log d$.

Returning to the specific example of heartbeats, we could design a diagnostic tool as follows. We compress fifty consecutive heartbeats (to a desired accuracy) in their statistical

best basis. Next, we compute a $K-L$ basis which we use to analyze the variance of the next batch (of course, it may be simpler and also sufficient to skip the $K-L$ construction).

We should also mention, that the $K-L$ basis, while optimal in the L^2 sense, may be quite lacking in efficiency in other norms. This fact for the case of trigonometric expansions has led us to pick wavelets for more local questions. We are therefore proposing the best basis selection with different norm criteria as a tool for obtaining more flexible coding and compressions.

Best basis for numerical computations:

As seen in the preceding paragraph, expressing the raw data in a statistical best basis provides a transform permitting faster manipulation and computation with the data. The paper on wavelets and numerical algorithms [**B.C.R**] in this volume provides illustrations of this fact where, rather than diagonalize a matrix (generally, a numerically expensive procedure), one chooses to express it in an appropriate wavelet basis obtaining a banded version and fast computation. In that case two procedures arise. The first is a mere coordinate change to a well chosen basis. The second, or so-called nonstandard form, consists in compressing a matrix $[a_{ij}]$ as if it were an image by finding a best 2-dimensional basis to compress it. Since the two dimensional bases of wavelet-packets or trigonometric waveforms can be obtained as a product of one-dimensional versions, this procedure amounts to a separation of variables and translates efficient compression into fast computation of linear transformations.

Figure 9

Wavelet packet analysis of a segment of "armadillo." The shaded rectangles measure the correlation with selected elements of the best basis, they are located above the portion of the signal with which they correlate. They have a base corresponding to duration and a height centered at the main frequency. See Appendix I for a description of the screen display

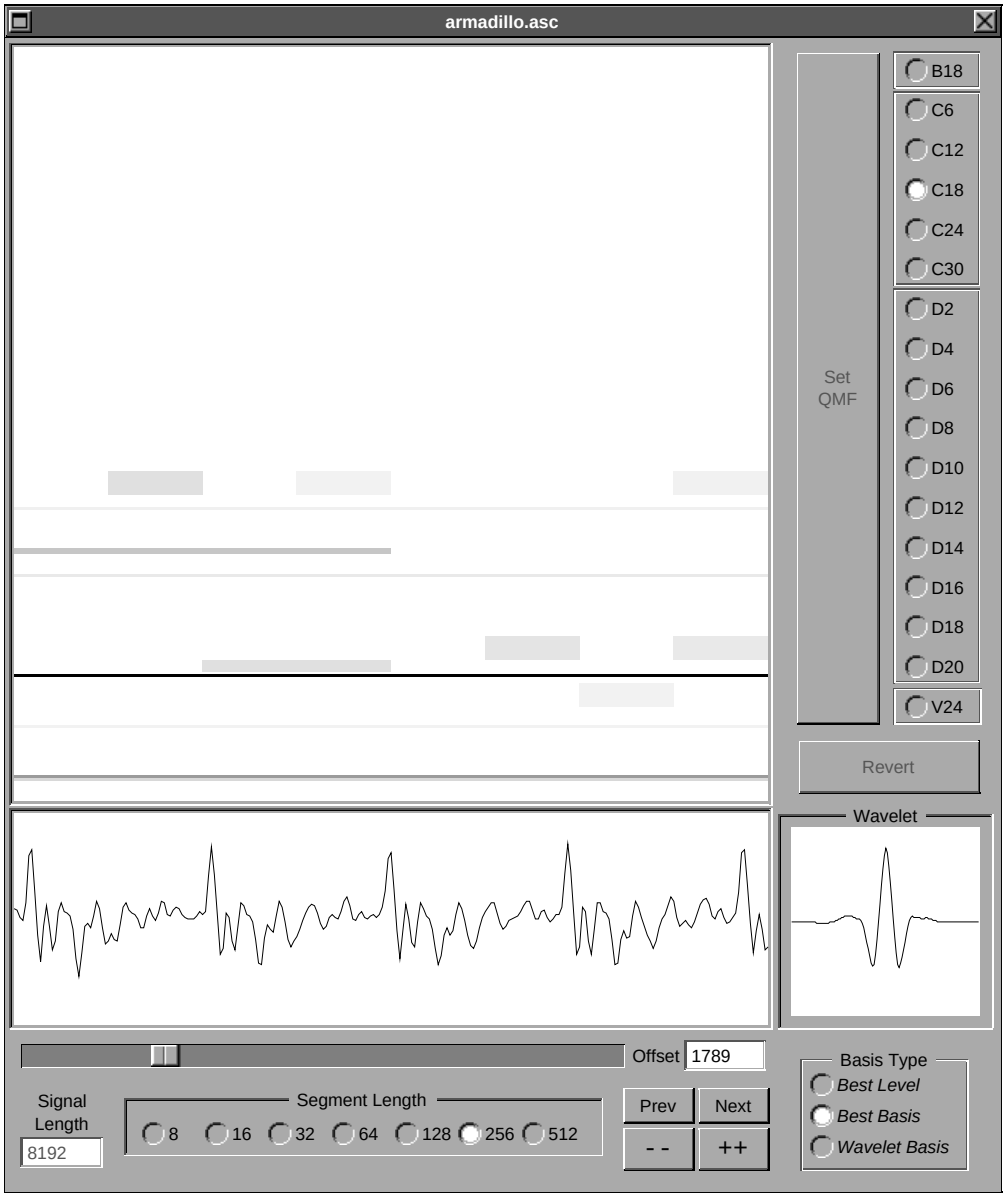

Figure 10

Best basis analysis, using the library V24 of a whistle. There are essentially two basic frequencies along the whole signal.

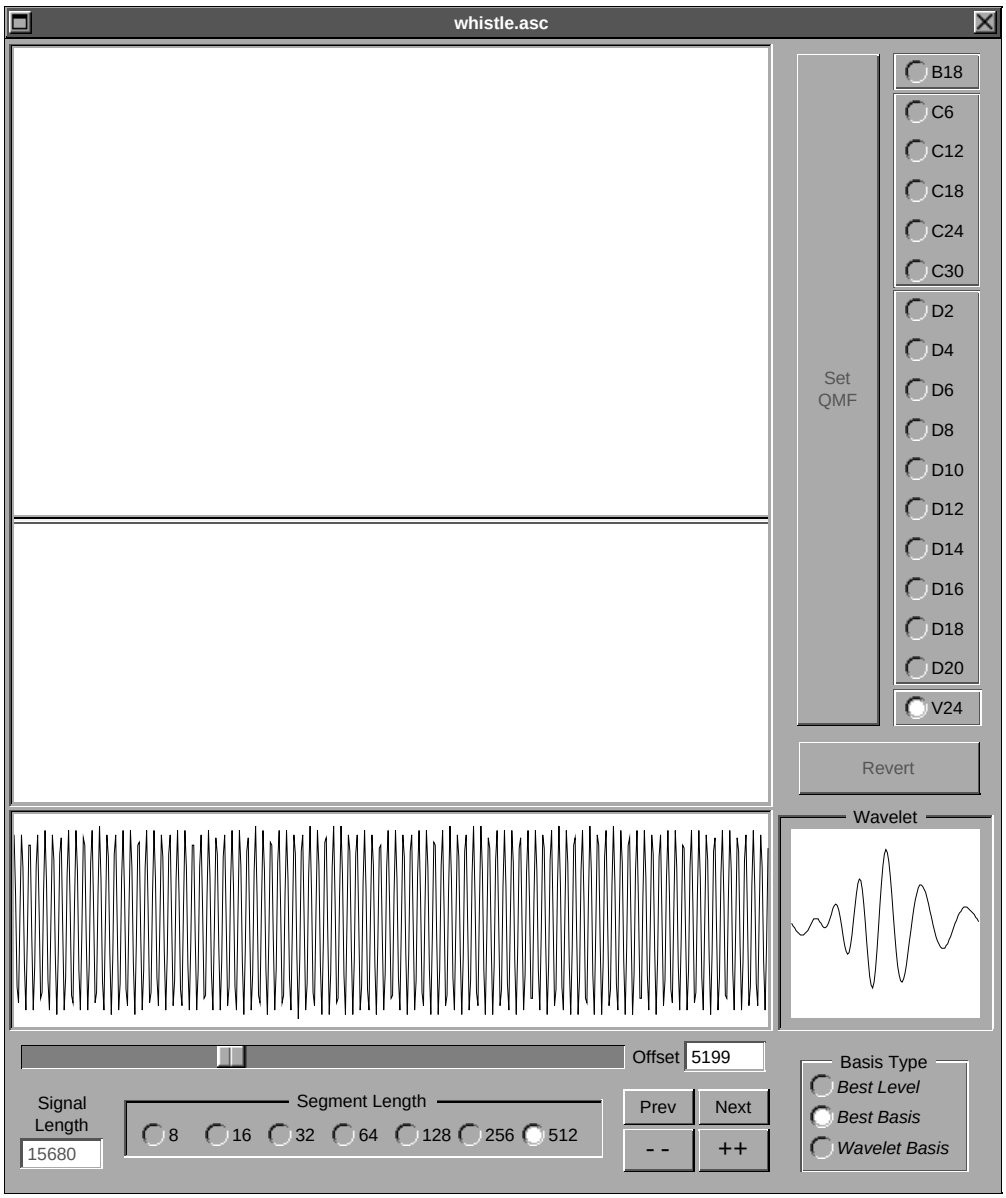

Figure 11

This is a plot of $\sin(250\ \pi x^2)$, a chirp up to the Nyquist frequency.

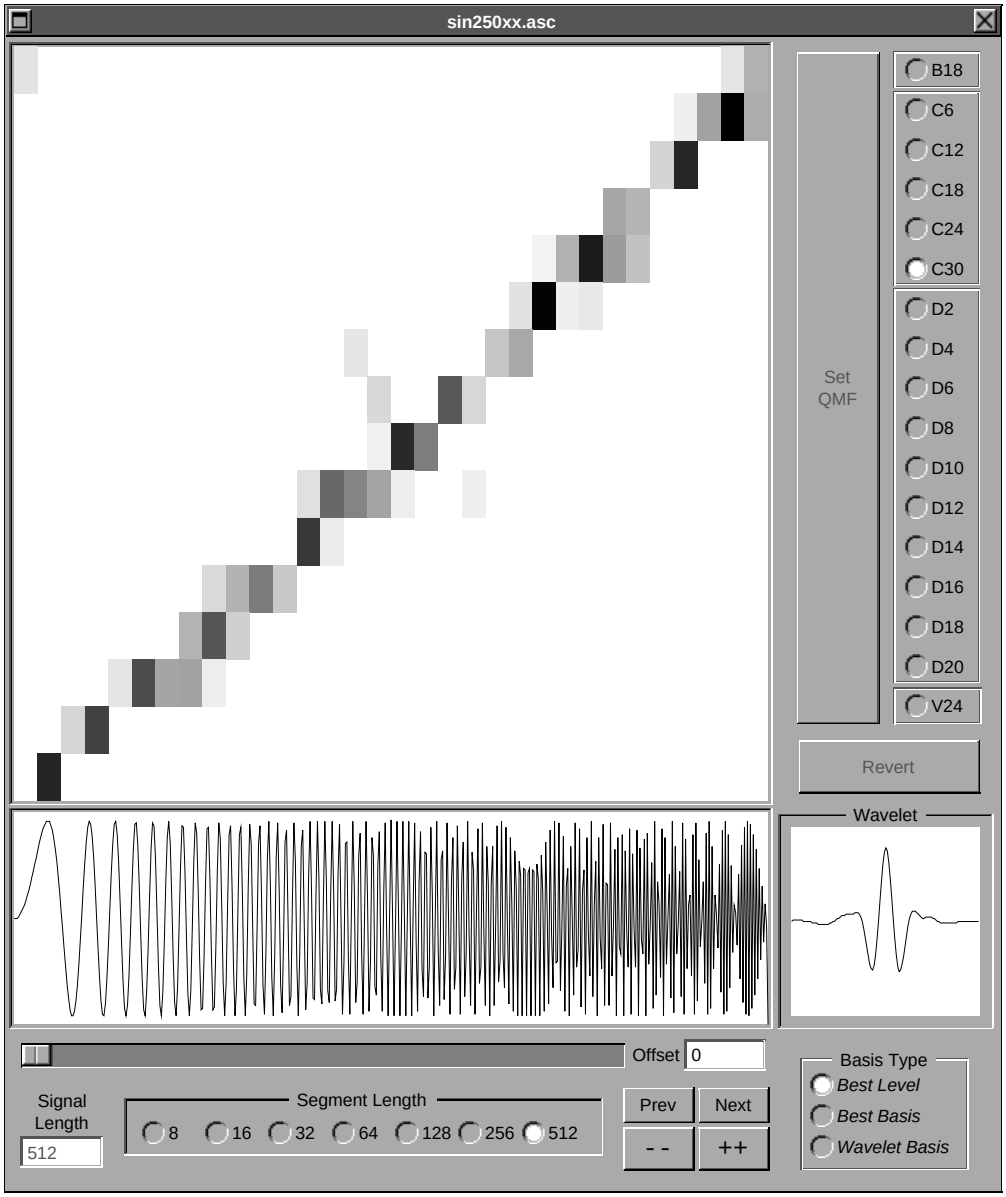

Figure 12

This is a plot of $\sin(750\ \pi x^2)$, showing aliasing between 1 and 3 times the Nyquist frequency.

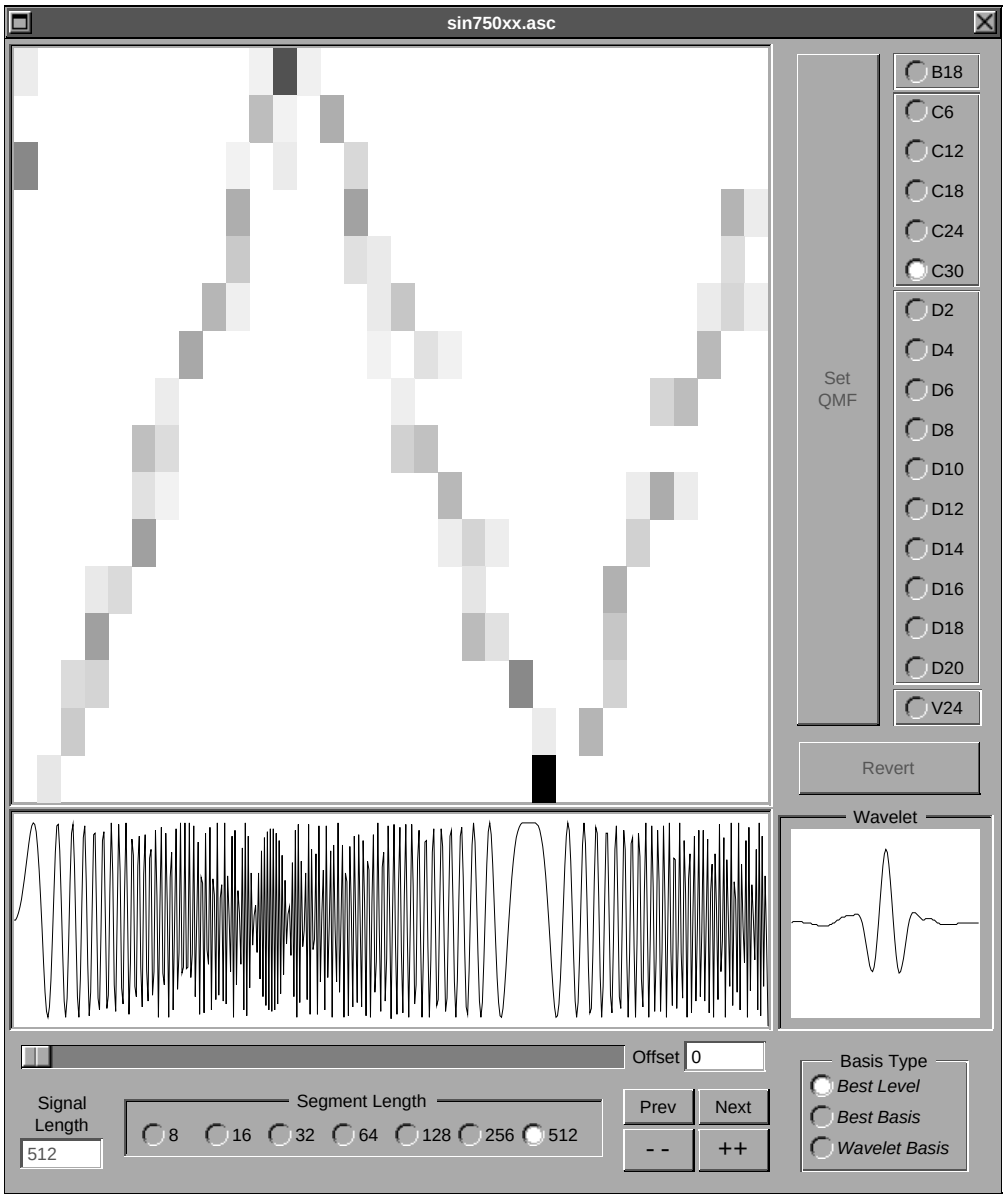

Figure 13
This is a plot of $\sin(250\ \pi x^3)$, a chirp whose frequency increases like a parabola.

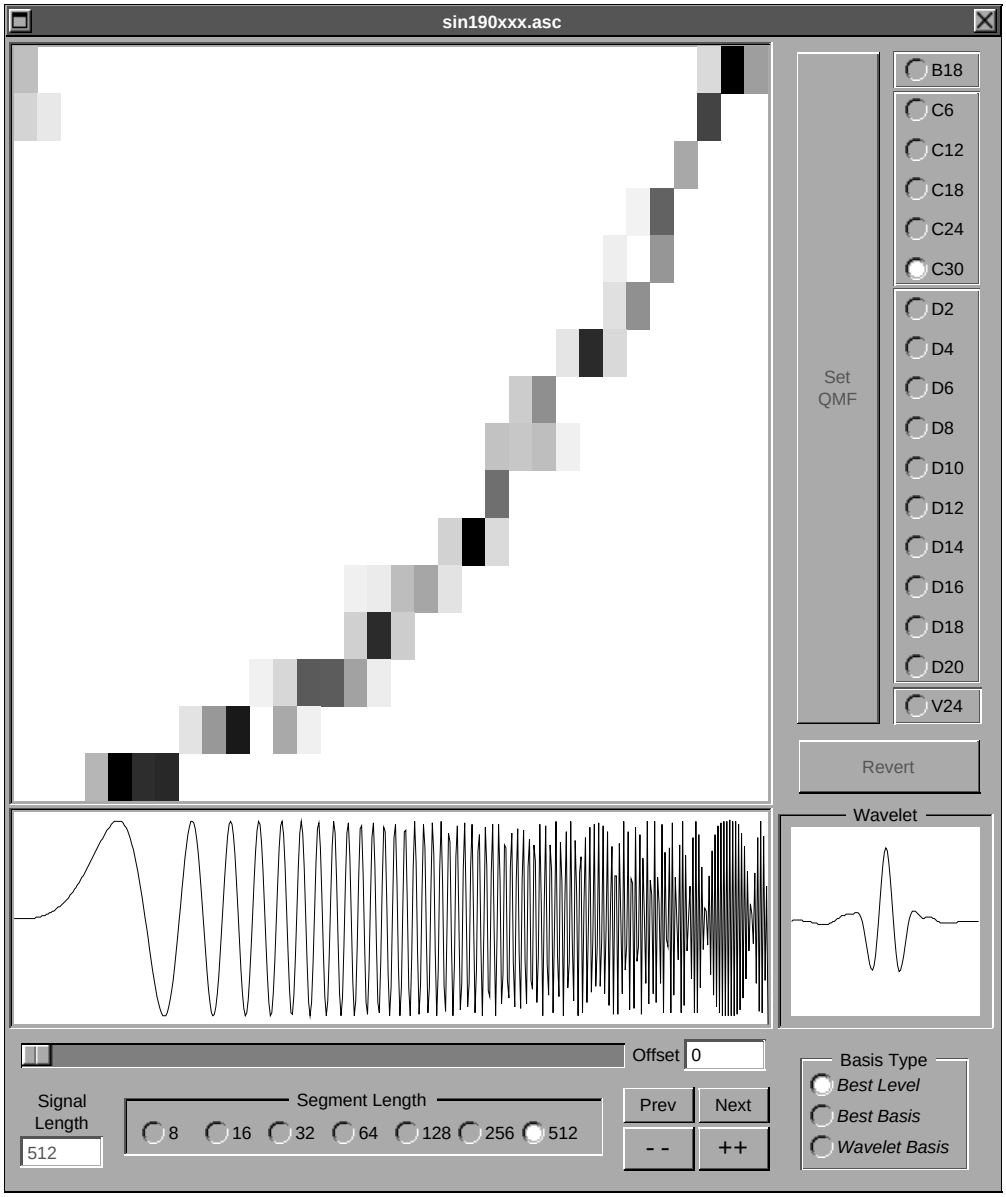

Figure 14

This plot shows $\sin(250\ \pi x^2) + \sin(80\ \pi x^2)$, two chirps whose frequencies increase at different rates.

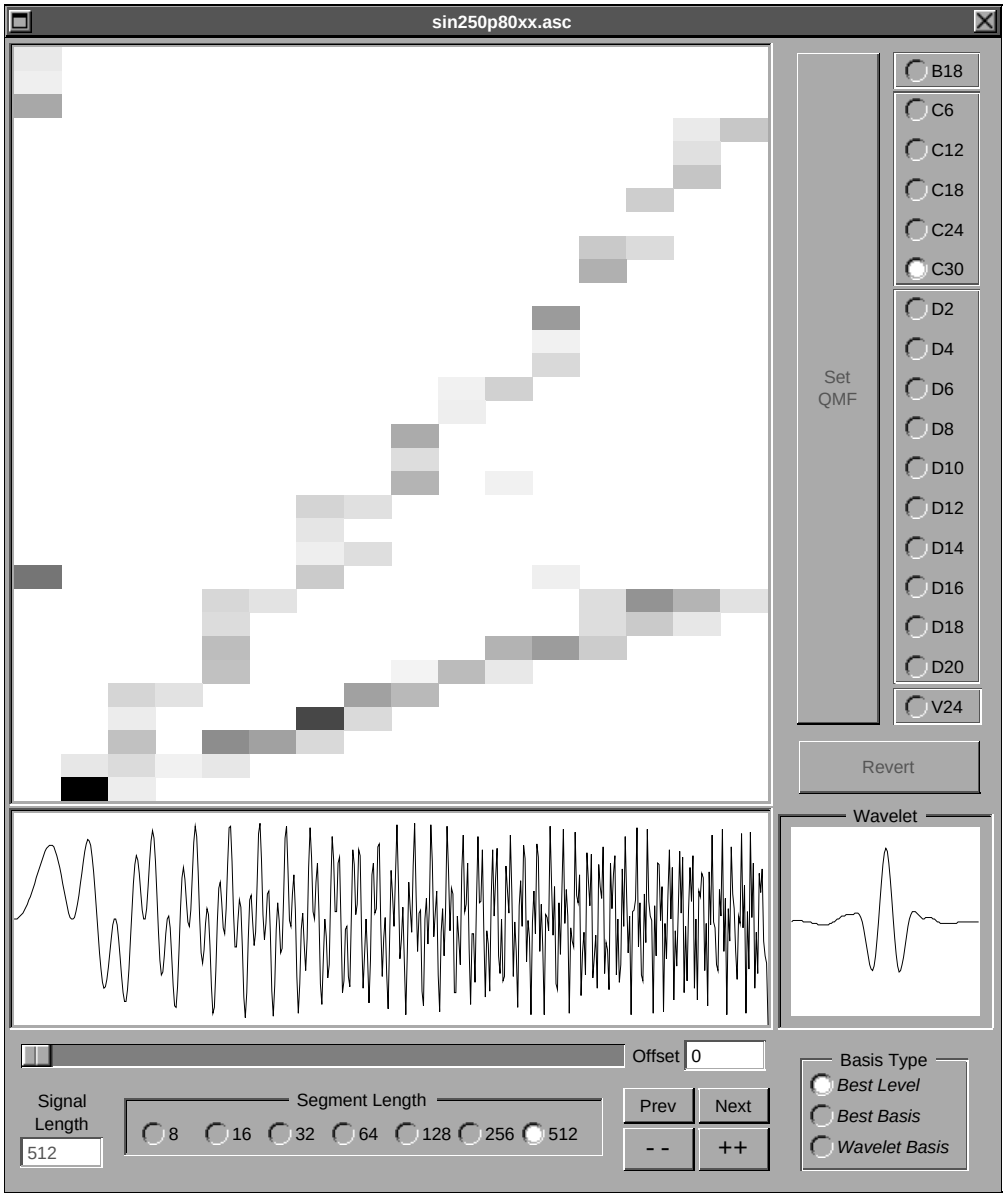

Figure 15

Let $y = x - 0.5$ and $z = x - 0.8$. This plot shows the superposition $\sin(250\ \pi x^2) + \sin(250\ \pi y^2) + \sin(250\ \pi z^2)$, three chirps of parallel increasing frequency which are 0.5 and 0.8 intervals out of phase.

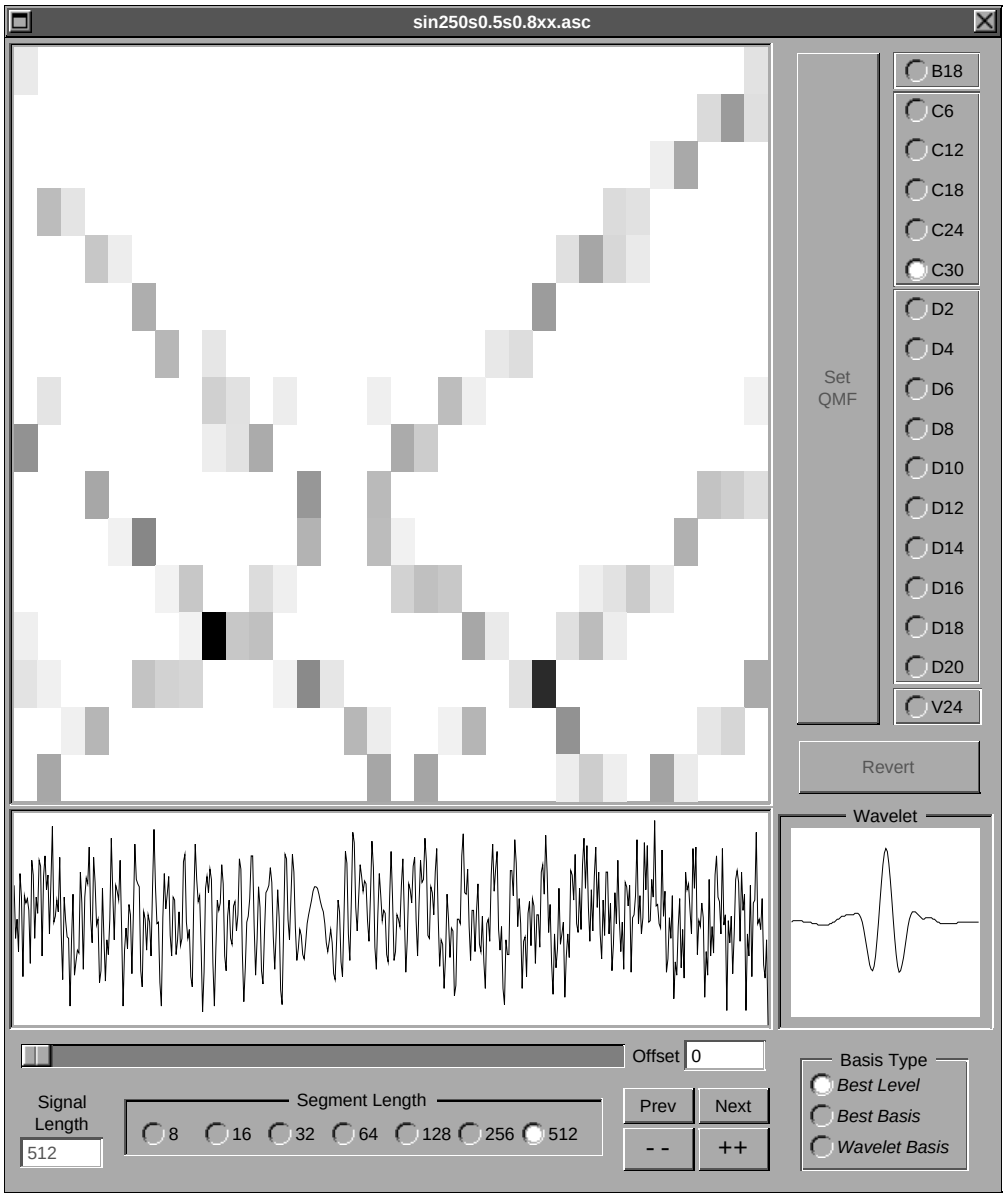

Figure 16

Let $y = x - 0.5$, and let $z = x - 0.8$. This plot of $\sin(250\,\pi x^2) + \sin(190\,\pi y^2) + \sin(120\,\pi z^2)$, showing three chirps increasing at different rates and 0.5 and 0.8 intervals out of phase, respectively.

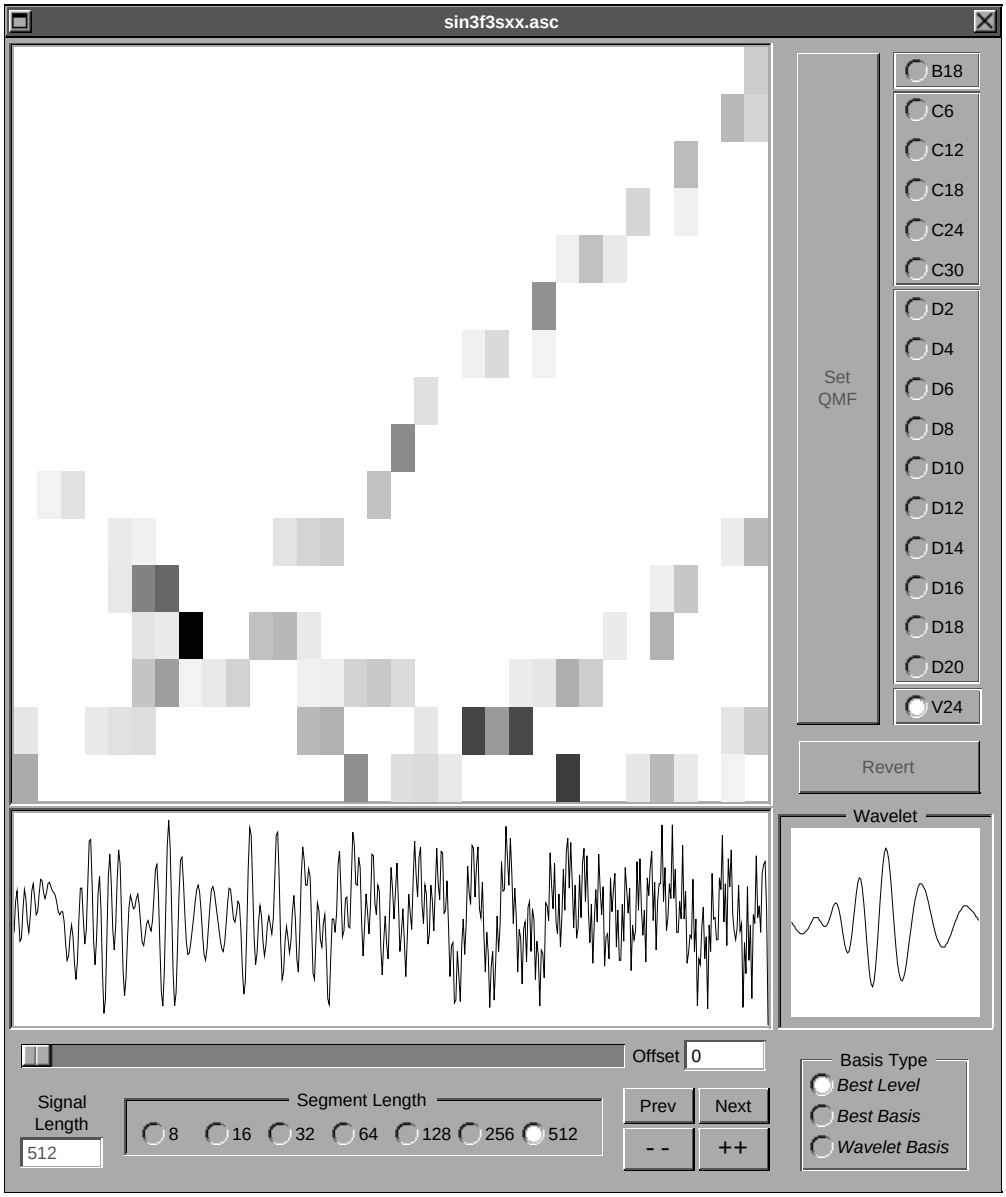

APPENDIX 1

USING THE WAVELET PACKET LABORATORY

Reading a Signal From a File

WPLab can read text files which have the extension ".asc" appended to their names, and which contain just ASCII floating-point numbers.

Selecting the "Open" item from the application's main menu brings up a browser panel which allows the user to select any single file with the proper extension. The entire file is read into a double-precision floating-point array in memory, and the array is padded with zeroes up to an integer multiple of the longest available segment length.

The number of samples in the signal is displayed in the "Signal Length" text field. The signal file name becomes the title for the main window, and also for the miniwindow upon miniaturization.

Displaying a Signal Segment

Segments of the signal file are plotted in the rectangular view near the bottom of the window.

WPLab can display segments whose lengths are a power of 2, and starting at arbitrary offsets. Use the "Segment Length" radio buttons to select this length, and any combination of the offset form, slider, and buttons to set the index of the first displayed sample.

The buttons "Prev", "Next", "++", and "–" adjust the offset. Their actions, respectively, are to subtract a segment length, add a segment length, add 1, and subtract 1. The program does the best it can given the signal length.

If a newly selected segment length is too long for the current offset, the offset is decreased to accommodate it.

Choosing a Quadrature Mirror Filter

There are 17 quadrature mirror filters (QMFs) available for wavelet packet analysis. They are identified by a letter ("B", "C", "D", or "V"), followed by (finite) impulse response length. For example, the Haar filters $(\sqrt{2}), (\sqrt{2}), (\sqrt{2}), -(\sqrt{2})$ are designated $D2$.

Preview a filter by clicking on its radio button. In the small "Wavelet" window will appear a plot of the mother wavelet associated to that QMF. This action enables the "Set QMF" and "Revert" buttons.

Roughly speaking, longer filters produce smoother wavelets and wavelet packets with better frequency resolution.

Click on the "Set QMF" button to register your choice and update the Phase Plane. This disables the "Set QMF" and "Revert" buttons.

Click on the "Revert" button to cancel previewing filters and return to the last registered QMFs. This will disable the "Set QMF" and "Revert" buttons.

The Phase Plane Representation of a Signal

The large square view contains that portion of the phase plane affected by the plotted segment. WPLab draws a rectangle in the phase plane for every modulated waveform in the basis chosen to represent the signal.

Each modulated waveform can be assigned 4 attributes: amplitude a, timescale s, frequency f, and position p. In a musical note, these correspond to loudness, duration, pitch, and the instant it is played.

Suppose that the signal segment has length $N = 2^n$. Coefficient (a, s, f, p) is displayed as the rectangle $[p2^s, (p+1)2^s] \times [f2^(n-s), (f+1)2^(n-s)]$, shaded in proportion to a^2.

Because of the Heisenberg uncertainty principle, position and frequency cannot both be specified to arbitrary precision. The uncertainty of the frequency is 2^s, and the uncertainty in position is $2^(n-s)$. Thus each rectangle or "Heisenberg box" has area $2^n = N$. Since the total area of the displayed section of the phase plane is N^2, there are exactly N Heisenberg boxes in a disjoint cover of the section.

A library consists of all possible Heisenberg boxes, and bases from the library consist of certain disjoint covers of the phase plane by such rectangles.

Choosing a Basis

Wavelet Basis: this forces a display of the wavelet basis constructed with the given

mother wavelet.

Best Level: this forces all of the Heisenberg boxes to have the same time scale. In particular, they must be congruent. There are (log N) such bases for a segment of length N, and the one displayed has minimum entropy.

Best Basis: this minimizes entropy over all bases corresponding to disjoint dyadic covers of the segment. There are more than 2^N such bases for a segment of length N.

Printing the Window

The entire contents of the key window may be printed at full scale with the "Print" menu item.

References

1. R. Coifman, *Adapted multiresolution analysis, computation, signal processing and operator theory*, ICM 90 (Kyoto).
2. R. Coifman, Y. Meyer, S. Quake and V. Wickerhauser, *Signal processing and compression with wavelet packets*, Numerical Algorithms Research Group, Yale University (1990).
3. R. Coifman and Y. Meyer, *Remarques sur l'analyse de Fourier à fenêtre*, série I, C. R. Acad. Sci. Paris **312** (1991), 259–261.
4. I. Daubechies, *Orthonormal bases of compactly supported wavelets*, Communications on Pure and Applied Mathematics **XLI** (1988), 909–996.
5. E. Laeng, *Une base orthonormale de $L^2(\mathbf{R})$, dont les éléments sont bien localisés dans l'espace de phase et leurs supports adaptés à toute partition symétrique de l'espace des fréquences*, série I, C. R. Acad. Sci. Paris **311** (1990), 677–680.
6. Available by anonymous ftp, *ceres.math.yale.edu*, InterNet address 130.132.23.22.