

Acoustic Signal Compression with Wavelet Packets

MLADEN VICTOR WICKERHAUSER

Numerical Algorithms Research Group

Department of Mathematics

Yale University

New Haven, CT 06520

Email: victor@math.yale.edu

13 August 1989

Abstract. The wavelet transform generalizes to produce a library of orthonormal bases of modulated wavelet packets. Each basis comes with a fast transform; these bases are similar to adapted windowed Fourier transforms. There is a notion of the “best basis” for a signal, given a cost function. This paper discusses some early results in acoustic signal compression using a simple counting cost function.

By generalizing the method of multiresolution decomposition, it is possible to construct orthonormal wavelet packets which provide a family of orthonormal bases for $L^2(\mathbf{R})$. These wavelet packets are well localized in both space and frequency. Correlating digitally sampled signals with these wavelet packets and discarding small coefficients can substantially compress the signals. The wavelet packets can be adapted to optimize fast data transformation, or to facilitate signal recognition. This paper presents results of computer experiments with the author’s signal compression algorithms on digitized speech signals.

Correlations with the entire family of orthonormal bases may be rapidly calculated for discrete functions of \mathbf{R}^N , where N is a positive integer power of 2. The algorithm used is based on the conjugate quadrature mirror filter method described by Mallat [M], with generalizations described in the Appendix. In the generalized case, it has the same complexity as the discrete “fast” Fourier transform. This complexity is $O(N \log_2 N)$ to obtain more than 2^N orthonormal bases all at once. These bases are especially well suited for signal processing, and the availability of choices suggests adaptive algorithms which optimize for a given problem. Orthogonality insures that signal analysis and resynthesis are computationally equivalent.

If the basis vectors of the family are arranged in an array of dimensions N columns by $n = \log_2 N$ rows, then any admissible subset of the entries will be an orthonormal basis of \mathbf{R}^N . Roughly speaking, an admissible subset is a choice of elements with 3 properties:

- (1) every column contains exactly one element,
- (2) elements in a single row appear in contiguous blocks of 2^k elements, where $0 \leq k \leq n$ is an integer, and
- (3) row blocks begin an integer multiple of their length from the left edge of the array.

This notion will be made precise in the Appendix.

The redundancy of the bases has benefits for both numerical algorithms and signal processing. For example, it is possible to choose the basis (i.e., the admissible subset of the array) which has the fewest large coefficients. The choice algorithm, as described in the appendix, can be implemented in $O(N \log_2 N)$ steps. The far smaller number of surviving coefficients (after thresholding) reduces the complexity of subsequent processing. Such an algorithm may also be used to compress audio signals. Applications to 2-dimensional signal compression may be found in [CMQW].

Reconstructing the signal corresponding to an array of zeros with a single 1 gives a particular tone or tone-burst, corresponding to a note produced by a musical instrument. Graphs of individual wave packets bear a striking resemblance to bursts of sound, suggesting that these bases are well suited to representing acoustic signals. This provides a method for digital music synthesis.

Another example is the recognition of acoustic signatures. By displaying an array of grey scales proportional to the absolute values of the corresponding coefficients, one obtains a picture of the signal together with a family of Fourier-like transformations. Nicolas [N] has build an expert system which can read this picture and detect characteristic features of a particular signal with high probability.

Perhaps the most important application is the rapid transformation of digital data. The same wavelet packet bases that compress signals will conjugate a large class of (discretized) operators into sparse or even band-diagonal form. Thus data in compressed form

is already prepared for fast and useful numerical algorithms. Already wavelet transform methods have produced fast numerical algorithms for matrix inversion and calculation of potentials [BCR]. Such a marriage of compression and transformation could have far reaching technical and economic consequences.

Synthesis of discrete wavelets and wavelet packets. For a longer exposition on the construction of wavelets, see Daubechies [D]. The starting point is a summable sequence $h = \{h_k\} \in l^1(\mathbf{R})$ with the following three properties:

- (1) $\sum_k h_{2k} = \sum_k h_{2k+1} = 2^{-1/2}$,
- (2) For all integers $l \neq 0$, $\sum_k h_k h_{k+2l} = 0$, and
- (3) There exists $\epsilon > 0$ such that $\sum_k |h_k| |k|^\epsilon < \infty$.

Finite sequences obviously satisfy (3). Complete characterizations of the finite sequences $\{h_k\}_{-M \leq k < M}$ which satisfy conditions (1) and (2) may be seen in [D]. Call such a finite sequence h a *summing filter*, and define the associated *differencing filter* $g = \{g_k\}$ by $g_k = (-1)^k h_{-k+1}$. The pair (h, g) are called *conjugate mirror filters* of length $r = 2M$. They may be used to construct orthonormal bases in two ways.

First, there is a continuous, compactly supported real-valued function ϕ on \mathbf{R} solving the functional equation

$$(1) \quad \phi(x) = 2^{1/2} \sum_k h_k \phi(2x - k).$$

This function may be approximated very efficiently by iteration. Define an associated function ψ by

$$(2) \quad \psi(x) = 2^{1/2} \sum_k g_k \phi(2x - k).$$

This function is also compactly supported—in fact, its support is contained in an interval of length r . Its dyadic dilates and integer translates generate a Hilbert basis for $L^2(\mathbf{R})$. Namely, writing ψ_{jk} for the function $\psi_{jk}(x) = 2^{-j/2} \psi(2^{-j}x - k)$, there is the following result of Daubechies and Meyer:

THEOREM 1. *The collection $\{\psi_{jk}, j, k \in \mathbf{Z}\}$ is a complete orthonormal basis for $L^2(\mathbf{R})$. ■*

Second, conjugate mirror filters provide a basis for $l^2(N)$, where N shall be shorthand for the finite set $\{1, 2, \dots, 2^n\}$. To construct this, introduce the summing and differencing operators H and G on $l^2(\mathbf{Z})$, respectively defined by

$$(3) \quad Hf(i) = \sum_k h_{k-2i}f(k), \quad Gf(i) = \sum_k g_{k-2i}f(k).$$

These two operators have adjoints H^* and G^* defined by

$$(4) \quad H^*f(k) = \sum_i h_{k-2i}f(i), \quad G^*f(k) = \sum_i g_{k-2i}f(i).$$

The original 3 conditions imposed upon h guarantee that $H^*H + G^*G = I$, that $HG^* = 0$, and that $l^2 = H^*l^2 \oplus G^*l^2$. Repeating this last decomposition n times yields

$$(5) \quad l^2 = \oplus \sum_{j=0}^{n-1} (H^*)^j G^* l^2 \oplus (H^*)^n l^2.$$

Now, it is possible to define $H, G : l^2(2^m) \rightarrow l^2(2^{m-1})$, and $H^*, G^* : l^2(2^{m-1}) \rightarrow l^2(2^m)$. First replace h, g with periodized filters \tilde{h}_m, \tilde{g}_m defined by

$$\tilde{h}_m(k) = \sum_{l \equiv k \pmod{2^m}} h(l), \quad \tilde{g}_m(k) = \sum_{l \equiv k \pmod{2^m}} g(l).$$

Then restrict the ranges of i and k in Eqs.(3,4) as follows:

$$(6) \quad Hf(i) = \sum_{k=1}^{2^m} \tilde{h}_m(k-2i)f(k), \quad Gf(i) = \sum_{k=1}^{2^m} \tilde{g}_m(k-2i)f(k).$$

$$(7) \quad H^*f(k) = \sum_{i=1}^{2^{m-1}} \tilde{h}_m(k-2i)f(i), \quad G^*f(k) = \sum_{i=1}^{2^{m-1}} \tilde{g}_m(k-2i)f(i).$$

For each $m > 0$, the periodized filters \tilde{h}_m, \tilde{g}_m satisfy the same orthogonality conditions as h, g , even though they differ (in general) at different levels m . Hence, these operators H, G satisfy the relations

$$(8) \quad l^2(2^m) = H^*l^2(2^{m-1}) \oplus G^*l^2(2^{m-1}), \quad HG^* = GH^* = 0, \quad \text{and} \quad H^*H + G^*G = I_m,$$

where I_m is shorthand for the identity operator on $l^2(2^m)$.

Counting dimensions in this finite-dimensional case shows that Eq.(5) may be rewritten as

$$(9) \quad l^2(N) = H^* l^2(N/2) \oplus G^* l^2(N/2) = \oplus \sum_{j=0}^{n-1} (H^*)^j G^* l^2(N/2^{j+1}) \oplus (H^*)^n l^2(1).$$

Taking the standard bases within each of the direct summands gives the so-called *wavelet basis* for the finite-dimensional space $l^2(N)$. Computing the expansion of a vector in this basis may be done recursively in $O(Nr)$ operations, where r is the length of the filters h and g , i.e., the cost of applying H or G .

But of course, each of the direct summands $l^2(N/2^j)$ may itself be expanded in the wavelet basis, and its summands recursively expanded still further, until all the summands are 1-dimensional. This yields the formula

$$(10) \quad l^2(N) = \oplus \sum_{j=0}^{N-1} F_{\epsilon_n(j)}^* \dots F_{\epsilon_1(j)}^* l^2(1), \quad \text{where } F_\epsilon^* = \begin{cases} G^*, & \text{if } \epsilon = 0, \\ H^*, & \text{if } \epsilon = 1, \end{cases}$$

and $\epsilon_i(j)$ is the i th digit in the binary expansion of j . Taking one coefficient for each of the 1-dimensional summands in Eq.(10) gives the lowest-level *wavelet packet basis* expansion of a vector. These functions are related to the classical Walsh functions, which are obtained as above by taking filters of length 2. One obtains the bit-reversed and sequency-ordered Walsh function bases by respectively composing bit-reversal or Gray-encoding with the function ϵ . Longer filters give smooth generalizations of Walsh functions, which are totally new objects.

Individual wavelet packets may be constructed as follows. Fix $N = 2^n$ and define the j th wavelet packet $w_j^n \equiv w_j \in l^2(N)$ from Eq.(10) above, namely, by

$$(11) \quad w_j = F_{\epsilon_n(j)}^* \dots F_{\epsilon_1(j)}^* 1.$$

This is a unit vector in $l^2(N)$ by an easy induction argument.

LEMMA 2. $\langle v, w_j \rangle = F_{\epsilon_1(j)} \dots F_{\epsilon_n(j)} v$, where F_ϵ is the adjoint of F_ϵ^* .

PROOF: Starting from Eq.(11), the operators F_ϵ^* may be transposed onto v , giving

$$\langle v, w_j \rangle = \langle v, F_{\epsilon_n(j)}^* \dots F_{\epsilon_1(j)}^* 1 \rangle = \langle F_{\epsilon_1(j)} \dots F_{\epsilon_n(j)} v, 1 \rangle.$$

But the right-hand side is an inner product in \mathbf{R}^1 . ■

Other wavelet packets. The periodized lowest-level wavelet packet w_j^n is not localized. It stretches throughout the entire index interval $\{0, 1, \dots, 2^n - 1\}$, as is easily shown by induction on the steps of the reconstruction algorithm. Since the filter length is at least 2, the number of non-zero coefficients at least doubles with each application of F_ϵ^* . But there are n filter applications. This property distinguishes lowest-level wavelet packets from wavelets, which are supported in subintervals of length proportional to their scale. Since it contains only frequency, no position information, the inner product of a vector with $w_j^n, j = 0, 1, \dots, N-1$ is analogous to the Fourier transform. The lowest-level wavelet packets themselves are the analogs of pure tones.

The reconstruction algorithm may be generalized to provide wavelet packets supported in subintervals. For $m = 1, 2, \dots, n$, define

$$(12) \quad w_j^m = F_{\epsilon_n(j)}^* \dots F_{\epsilon_{n-m+1}(j)}^* e_{\hat{j}},$$

where $e_{\hat{j}}$ is the unit vector in $\mathbf{R}^{2^{n-m}}$ which has a 1 in the \hat{j} position and zeros everywhere else, with $\hat{j} = j \bmod 2^{n-m}$. In other words, $\epsilon_k(j) = \epsilon_k(\hat{j})$ for $k = 1, 2, \dots, n-m$. Let w_j^0 be the standard basis vector in \mathbf{R}^N which has a 1 in the j th position and zeros elsewhere. The component of $v \in \mathbf{R}^N$ in the w_j^m direction is given by the \hat{j} coefficient in the vector $F_{\epsilon_{n-m+1}(j)} \dots F_{\epsilon_n(j)} v$. The proof is nearly identical with that of Lemma 2:

LEMMA 3. $\langle v, w_j^m \rangle = \langle e_{\hat{j}}, F_{\epsilon_{n-m+1}(j)} \dots F_{\epsilon_n(j)} v \rangle$. ■

The width of the support of w_j^m is at most $2^m + (r-2)(2^m - 1)$, where r is the length of the filters h, g . For small m and r , this can be much smaller than 2^n . Thus inner products with $w_j^m, j = 0, 1, \dots, N-1$ contain some position information, like windowed Fourier transforms. The narrower wavelet packets themselves correspond to tonebursts.

Write w^m for the orthonormal basis $\{w_j^m : j = 0, 1, \dots, N-1\}$. Expanding a vector into w^n generates all the expansions into w^1, \dots, w^n simultaneously, suggesting the following picture. Denote by w^0 the standard basis in \mathbf{R}^N . A vector in the w^0 basis may be expanded in the $\log_2 N$ wavelet packet bases w^1, \dots, w^n with the results placed row-by-row into a

rectangular array. Each location in this array corresponds to a wavelet packet orthogonal to other wavelet packets in that row, of support width increasing with row number. The array thus contains $1 + \log_2 N$ representations of a vector, including analogs of windowed Walsh (or Fourier) transforms at all dyadic window-widths.

Likewise, if row m has a single 1 in column j with all its other entries 0, the rest of the array is determined, and the standard representation of w_j^m will appear in row 0.

Any wavelet packet is extensible to a smooth periodic function, of the same degree of smoothness as the solution ϕ of the functional equation (1). (In the finite-dimensional case, *numerical smoothness of degree k* will mean rapid decrease of the first k successive differences.)

Functions underlying the discrete algorithm. Wavelet packet coefficients can be interpreted as inner products of an underlying function of \mathbf{R} with a family of smooth orthonormal elements of $L^2(\mathbf{R})$. Let ϕ, ψ be the continuous functions in $L^2(\mathbf{R})$ defined by equations (1) and (2), and suppose that $S \in L^2(\mathbf{R})$ is a function to be approximated with one value in each subinterval of length δ . One possibility is to calculate an approximate value $s(k) = \delta^{-1} \int_{\mathbf{R}} S(x) \phi(\delta^{-1}x - k) dx$ at each integer k , relying on the good localization of the bump function ϕ . In fact, if we arrange that higher moments of ϕ vanish, this approximation is very good for smooth functions S :

LEMMA 4. *If $\int_{\mathbf{R}} x^d \phi(x) dx = 0$ for all $0 < d < D$, and $S \in C^D(\mathbf{R})$, then for x near k , $S(x) = s(k) + O(\delta^D)$.*

PROOF: Expand S in its Taylor series about k . The inner product ϕ evaluates S at $x = k$ and kills all lower order terms than D . ■

It can be arranged that any fixed number of moments of ϕ vanish. The number of vanishing moments is at most half the filter length, however, so longer finite filters must be used.

Now define a family of wave functions $W_n \in L^2(\mathbf{R})$, $n = 0, 1, \dots$, recursively from ϕ and

ψ by the following: $W_0(x) = \phi(x)$, $W_1(x) = \psi(x)$, and

$$(13) \quad W_{2n}(x) = 2^{1/2} \sum_k h_k W_n(2x - k), \quad W_{2n+1}(x) = 2^{1/2} \sum_k g_k W_n(2x - k).$$

These are orthogonal transformations in L^2 , so that W_n is a unit vector for every n . From these definitions, we see that any sequence s that approximates a function S will have as its wave packet coefficients the inner products of S with the dyadic dilates and integer translates of W_n for various n . The following formula, noted by Coifman and Meyer, describes this correspondence:

PROPOSITION 5. *For any integer k , and nonnegative integers m, n , the relationship between wavelet packet coefficients and underlying inner products is given by the formula:*

$$2^{-m/2} \int_{\mathbf{R}} S(x) W_n(2^{-m}x - k) dx = \langle e_k, F_{\epsilon_1(n)} \cdots F_{\epsilon_m(n)} s \rangle,$$

where $\epsilon_j(n)$ is the j th binary digit of n , padded with leading zeros if necessary, and e_k is the unit sequence in l^2 which has a single 1 in the k th position and zeros elsewhere.

PROOF: The inner product with e_k merely picks out one coefficient. The rest of the formula may be demonstrated by considering the effect of a filter convolution on s . But if $\sigma = \sigma(k)$ is a partial result, we have:

$$\begin{aligned} F_\epsilon \sigma(k) &= 2^{-q/2} \int_{\mathbf{R}} S(x) \left(\sum_l f_\epsilon(l - 2k) W_p(2^{-q}x - l) \right) dx \\ &= 2^{-q/2} \int_{\mathbf{R}} S(x) \left(\sum_l f_\epsilon(l) W_p(2^{-q}x + 2k - l) \right) dx \\ &= \begin{cases} 2^{-q/2} \int_{\mathbf{R}} S(x) 2^{-1/2} W_{2p}(2^{-q-1}x - k) dx, & \text{if } \epsilon = 0, \\ 2^{-q/2} \int_{\mathbf{R}} S(x) 2^{-1/2} W_{2p+1}(2^{-q-1}x - k) dx, & \text{if } \epsilon = 1. \end{cases} \end{aligned}$$

Induction on q up to m completes the proof. ■

Wavelets correspond to $n \equiv 1$, $m = 0, 1, 2, \dots$, and $k \in \mathbf{Z}$. Wave packets correspond to a fixed level $m \equiv l > 0$, $n = 0, 1, 2, \dots$, and $k \in \mathbf{Z}$.

The periodized case can be interpreted similarly. One can easily produced a series of graphs depicting approximations of these functions W_n . Filters with regular limits ϕ, ψ

produce a library of wavelet packets with a remarkable resemblance to sound bursts or smoothly modulated notes. Expansion in the wavelet packet basis correlates a signal to these notes, and their appearance suggests that the library is suitable for acoustic signal analysis. Some examples of these graphs may be seen below:

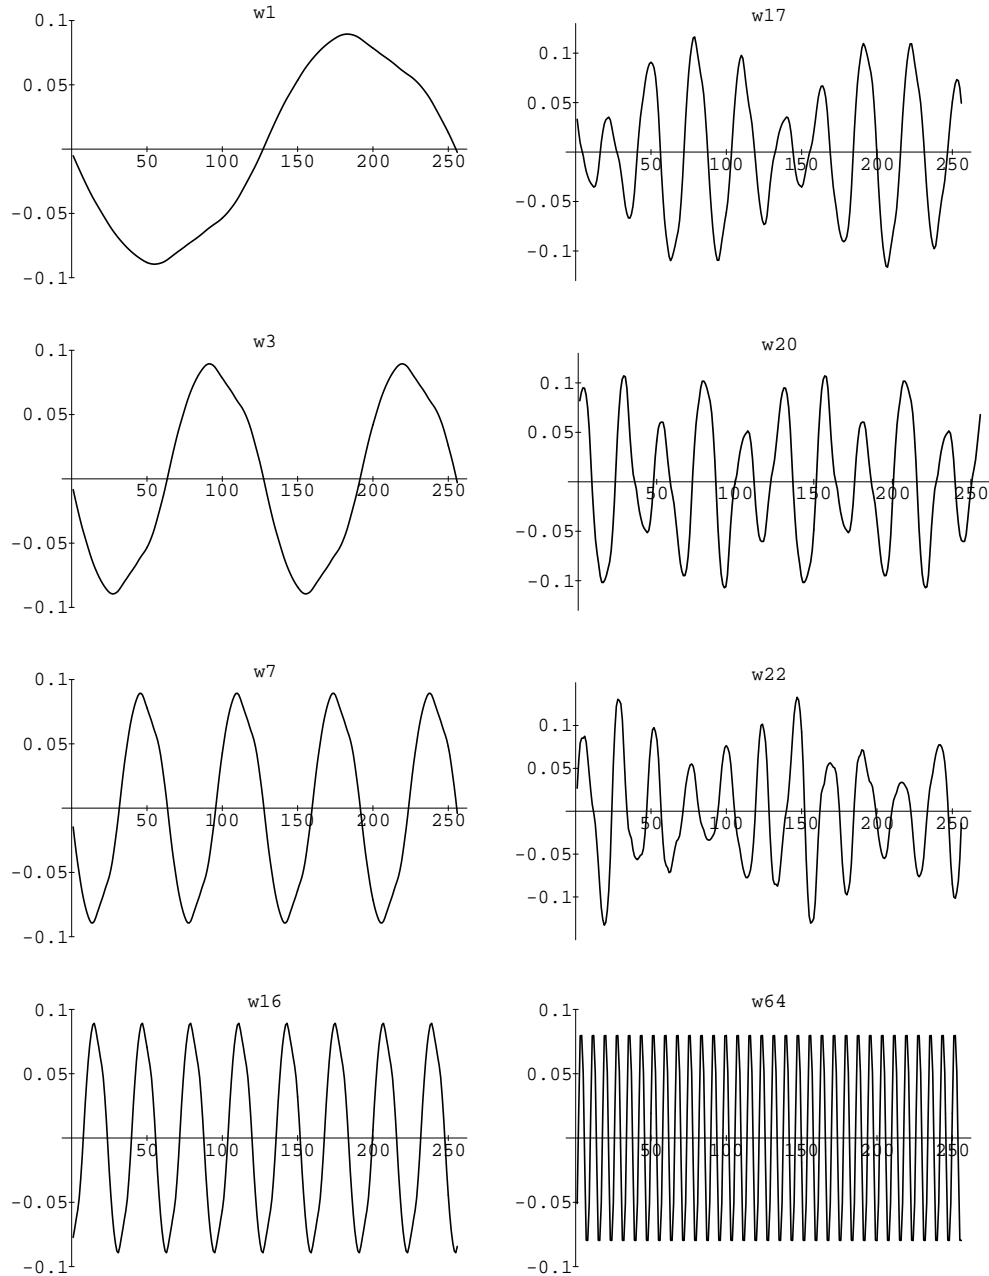


Figure 1. Periodized wavelet packets

Wavelet packets as musical tones and tonebursts. It is amusing to listen to wavelet packets. They approximate continuous periodic functions on \mathbf{R} , and can be played back as tones through a digital-to-analog converter and loudspeaker. The curious reader who wishes to hear these tones may obtain from the author a recording of a few examples.

The speed of the wavelet packet algorithm suggests that wavelet packet generators could be useful in sound synthesis. A single wavelet packet generator may replace a large number of oscillators. Through experimentation, a musician could determine combinations of wave packets that produce especially interesting sounds. Having an orthonormal basis offers the musician the widest possible range of sounds.

Also, consider approximating the sound of a musical instrument. A sample of the notes produced by the instrument may be decomposed into its wavelet packet coefficients, using a suitable filter length. Reproducing the note requires reloading those coefficients into a wavelet packet generator and playing back the result. Properties such as attack and decay may be controlled separately (e.g., with envelope generators), or by using longer wavelet packets and encoding those properties as well into each note. Any of these processes may be controlled in real time, for example by a keyboard.

Of course, the musical instrument could just as well be a human voice, and the notes words or phonemes.

Conjugate mirror filters of length greater than 2 result in numerically smooth wavelet packets, which are presumably well suited to representing digitally-sampled smooth analog signals. Results with the two sample signals studied in the next section support this belief: longer filters result in less measurable and audible distortion.

Wavelet packet coefficients, like wavelet coefficients, are mostly very small for digital samples of smooth signals. The differencing filters remove correlations between adjacent, nearly identical samples. Discarding coefficients below some predetermined (or dynamically determined) cutoff introduces only small errors while effecting data compression for smooth signals. In effect, this means that a wave packet-based music synthesizer could store many complex sounds efficiently. Similarly, a wavelet packet-based speech synthesizer could be used to reconstruct highly compressed speech signals.

Data compression with wavelet packets. Let F be a field, finite or infinite, let $v \in F^N$ be a vector, and consider the problem of representing v with fewer than N coefficients. If F is finite (i.e., an alphabet), then there are numerous lossless data compression algorithms that can replace v with a substitution table and symbol list, removing some of the redundancy of information in the standard representation of v . Such algorithms rely on searching through v for repeated strings and tabulating some fraction of them. The best implementations (for small alphabets F) yield compression ratios between 2 and 10, i.e., the compressed vector is 2 to 10 times shorter than the original.

Lossless data compression is also possible with wavelet packets. If more of the wavelet packet coefficients are small, then the entire expansion has lower entropy than the original signal. Fewer bits need to be transmitted for the small coefficients. As a simple test of this idea, one sample signal was compressed by Huffman coding, yielding a compression ratio of 2. For comparison, its complete best-level wave packet expansion was compressed by the same Huffman coding algorithm, yielding a compression ratio of 3. The reconstructed signal was identical to the original in both cases.

Such methods lose efficiency when F is very large or infinite, such as when it is taken to be \mathbf{R} . The alternative is a method that introduces some acceptable losses or distortions in exchange for greater efficiency.

Loss or distortion introduced in reconstruction shall be defined as the root-mean-square or l^2 norm of the difference between v and its approximation \hat{v} . For simplicity of comparison, this shall be normalized into the relative difference $\|v - \hat{v}\|/\|v\|$. Distortions of up to 5% are considered acceptable in some speech processing applications. An optimal compression method for a speech signal is one that yields no more than 5% distortion while maximizing the compression ratio. While it is too much to expect that a single method will be optimal for all signals, it is not unreasonable to expect that a method which is close to optimal for one signal is close to optimal for similar signals—similar bandwidths, dynamic ranges, etc.

The class of algorithms from which the optimal compression method is to be chosen will be the following: A fixed signal of length $N = 2^n$ will be decomposed into the n wavelet

packet bases w^1, \dots, w^n . The entries in each row which are smaller than a predetermined cutoff (depending only on the norm of the signal) will be set to 0. Then the admissible subsets of the array will be searched to find that set which has the fewest nonzero coefficients. For filter length r , this algorithm takes $O(rN \log_2 N)$ steps: $rN \log_2 N$ operations to expand the signal, $N \log_2 N$ to cut off small coefficients, and $2N \log_2 N$ to find the optimal admissible subset. In addition, the compressed coefficients must have a header appended which defines their locations within the admissible subset. Enlarging the class of admissible subsets (which improves the compression ratio) causes this header to grow, forcing a tradeoff.

In the extreme case where the definition of admissible is conditions (1,2,3) of the introduction, the number of admissible subsets is very large. Let A_n be the number of admissible subsets for a signal of length $N = 2^n$. Then $A_1 = 2$ (the standard basis, or {sum,difference}), and there is a recurrence relation deducible from the recursive construction of wavelet packets. Namely, the admissible subset is either w^0 or a choice of admissible subsets from the wave packet expansions of Gl^2 and Hl^2 . This generalizes to provide the inductive step,

$$(14) \quad A_{n+1} = 1 + A_n^2.$$

The constant 1 quickly becomes negligible as n increases, and an obvious simplification gives the estimate $A_{n+1} \geq 2^{2^n} = 2^N$. In this case the header attached to each coefficient must state which level it comes from, together with the position in that level. This requires at least $\log_2 N + \log_2 \log_2 N$ bits to transmit. In the appendix is a formula which evaluates the effect of this requirement upon the compression ratio.

At the other extreme, always choosing the same admissible set—say, the lowest-level wavelet packets w^n —shortens the header to $\log_2 N$ bits per coefficient.

An intermediate case is to fix a segment length N and then allow within each segment only complete rows w^m , for $0 < m \leq n$, as the admissible subsets. This permits some adaptability to the signal while adding a header only $\log_2 \log_2 N$ bits long to an entire segment of coefficients, quite negligible compared to N as $N \rightarrow \infty$.

Experiments with speech signal compression. A speaker recorded the phrase “Joe brought a young girl” into computer memory, using a 22kHz sampling rate and a resolution of 8 bits per sample. There was no pre-emphasis, and the analog-to-digital converter was assumed to be reasonably linear. This resulted in a signal v containing slightly fewer than 2^{15} samples, which was then padded at the end with dither to a length of 2^{15} . Hence $N = 32768$ and $n = 15$ for this example.

This signal was then transformed into a list of single-precision floating point numbers. For calculating the L^2 norm, the signal length was taken to be 1 with subinterval width $1/N$.

Three kinds of compression were tried on this list of floating-point numbers. Method 1 involved expanding v in the wavelet basis, using filters of length 2, 10, and 18, cutting off all terms less than ϵ , and then reconstructing the signal. Hardware limitations forced the signals to be divided into 2 segments of 16384 samples each before transforming to the wavelet basis. Table 1 shows the results obtained with this method.

For the next two methods, the signal was divided into 8 segments of $4096 = 2^{12}$ samples, and each segment expanded into the bases w^1, \dots, w^{12} , using filters of length 2, 10, and 18. In method 2, each segment was reconstructed from its w^{12} representation after those coefficients less than ϵ were discarded. Results using this method are presented in Table 2.

In method 3, each segment was reconstructed from its “best” representation w^b , $1 \leq b \leq 12$, where b was chosen for each segment to maximize the number of coefficients smaller than ϵ . This gave the best results, shown in Table 3.

In all three methods, the values $\epsilon = 0.5, 2.0, 8.0$ were used, corresponding to approximately 0.6%, 2.5%, and 10% of the maximum value of the signal. This is expressed in dB below peak in Table 4, using the formula $(\text{dB below peak}) = -20 \log_{10}(\epsilon/\text{Peak})$.

The experiment was repeated with a male speaker and a female speaker. The dynamic range and rate of speech for the two samples were adjusted to be as similar as practicable, so that bandwidth was the principle difference between the signals.

While it is necessary to listen to the reconstructed signals in order to judge their in-

telligibility, some observations are possible from the estimates of errors introduced after reconstruction. First, the distortion seems to increase linearly with compression ratio. It becomes noticeable at a value of 5%, which occurs for compression ratios around 5 or 10. Second, at large cutoff values the longest filter gives the least distortion and the highest compression ratio. This is not surprising, given the smoothness of speech signals. Third, distortions become audible as superimposed hiss, while the speaker remains recognizable.

For contrast, Table 5 lists the distortions and compressions introduced by simply squelching the original signal at the same cutoff levels used in methods 1, 2, and 3. Even allowing large distortions, this method gives substantially smaller compression ratios.

Compression ratios are somewhat misleading out of context. In particular, if the sampling rate is more than twice the signal bandwidth, then oversampling contributes redundancy. Furthermore, silences between words or at the beginning and end of the phrase pad the signal length. Any compression scheme should remove these sources of redundancy. What remains is the Nyquist information content, and to judge a compression method it is important to measure the distortion introduced as the signal length is compressed below the Nyquist length. Estimating this length for the test signal above requires counting only those samples larger than some squelch level, and determining from a sonogram the essential bandwidth of the signal.

For the two speech samples studied, “silences” were considered to be any samples with values -1 , 0 , or 1 from the range $\{-128, \dots, 127\}$. With the maximum amplitude of the male sample being 82, this represented a squelch level of about -38 dB. The following output is from a computer program that counted the squelched samples:

```
Please name the input file of bytes: jbayg.male
Set the squelch level (1,2,...,127): 1
Squelched: 7671 of 32768, or 23.4% of signal. (True length 25097)
```

With the maximum amplitude of the female sample being 116, the squelch level represented -41 dB. Corresponding output for that signal is

```
Please name the input file of bytes: jbayg.female
```

Set the squelch level (1,2,...,127): 1

Squelched: 7653 of 32768, or 23.4% of signal. (True length 25115)

The speech samples were produced by subjects intent on speaking at the same rate. Higher squelch levels result in diverging “True lengths,” providing some justification for the rather arbitrary choice of 1.

Bandwidth determinations are less precise. One method is to find a frequency f such that $1 - \epsilon$ of the signal energy, averaged over time, is at frequencies below f . It then remains to determine ϵ . Under the assumption that 5% rms errors are acceptable, set $\epsilon = 0.05$. Then the bandwidth for the male sample is 5.5 kHz, while for the female sample it is 7.5 kHz. Thus, the male’s speech is oversampled by a factor of 2, while the female’s is oversampled by a factor of 1.5.

Removing these redundancies would result in signals of Nyquist length. These are tabulated below:

SIGNAL	BANDWIDTH	NYQUIST LENGTH	COMPRESSION
Male	5.5 kHz	12549 (38%)	2.6
Female	7.5 kHz	16743 (51%)	2.0

Only compression ratios greater than the compression to the Nyquist length will be independent of the sampling method. The absolute compression ratio divided by the Nyquist compression will be called *compression below the Nyquist length* or the *invariant compression ratio*. For confirmation of the validity of this estimate, notice that long filters and negligible cutoffs result in compression to this Nyquist length, regardless of method.

Whenever the invariant compression ratio exceeds 1, there will be distortion. This distortion might be acceptable if it is small, or if the essential features of the signal are still intelligible. In fact, all of the reconstructed compressed signals retained recognizability as well as intelligibility: the speakers’ friends could easily identify the speakers. Other compression methods, which exploit specific features of human speech to obtain much higher compression ratios, will be considered in a subsequent article.

<u>SIGNAL</u>	<u>FILTER</u>	<u>CUTOFF</u> (dB)	<u>COMPRESSION</u>	<u>INV. COMPRESSION</u>	<u>Δ_{RMS}</u> (%)
Male	2	44	1.4	0.5	0.0
		32	2.7	1.0	4.7
		20	7.9	3.0	14
	10	44	2.5	1.0	0.7
		32	6.3	2.4	4.0
		20	13.9	5.3	9.4
	18	44	2.6	1.0	0.7
		32	6.7	2.6	4.0
		20	14.3	5.5	9.1
Female	2	47	1.3	0.7	0.1
		35	2.1	1.1	3.2
		23	5.2	2.6	12
	10	47	2.1	1.1	0.4
		35	4.3	2.1	3.0
		23	8.7	4.3	8.0
	18	47	2.2	1.1	0.9
		35	4.6	2.3	2.9
		23	8.9	4.5	7.5

Table 1: Discard small wavelet components.

<u>SIGNAL</u>	<u>FILTER</u>	<u>CUTOFF</u> (dB)	<u>COMPRESSION</u>	<u>INV. COMPRESSION</u>	<u>Δ_{RMS}</u> (%)
Male	2	44	1.4	0.5	0.4
		32	2.8	1.1	5.0
		20	7.2	2.8	15
	10	44	2.3	0.9	0.7
		32	5.8	2.2	4.1
		20	14.0	5.4	10
	18	44	2.4	0.9	0.7
		32	6.1	2.3	4.0
		20	15.3	5.9	10
Female	2	47	1.3	0.7	0.3
		35	2.0	1.0	3.5
		23	4.6	2.3	12
	10	47	2.0	1.0	0.4
		35	4.0	2.0	3.0
		23	8.9	4.5	8.6
	18	47	2.1	1.1	0.5
		35	4.4	2.2	3.0
		23	9.7	4.9	8.2

Table 2: Discard small lowest-level wavelet packet components.

<u>SIGNAL</u>	<u>FILTER</u>	<u>CUTOFF</u> (dB)	<u>COMPRESSION</u>	<u>INV. COMPRESSION</u>	<u>Δ_{RMS}</u> (%)
Male	2	44	1.7	0.7	1.2
		32	3.3	1.3	4.9
		20	8.8	3.4	14
	10	44	2.6	1.0	0.8
		32	6.8	2.6	4.0
		20	16.6	6.4	9.7
	18	44	2.7	1.0	0.8
		32	7.3	2.8	4.0
		20	17.9	6.9	9.8
Female	2	47	1.5	0.7	0.7
		35	2.4	1.2	3.4
		23	5.6	2.8	11
	10	47	2.3	1.1	0.5
		35	5.0	2.5	3.0
		23	10.8	5.4	8.0
	18	47	2.4	1.2	0.5
		35	5.5	2.7	2.9
		23	12.4	6.2	7.7

Table 3: Choose best-level wavelet packets.

<u>SIGNAL</u>	<u>PEAK</u>	<u>ϵ</u>	<u>dB BELOW PEAK</u>
Male	82	0.5	44
		2.0	32
		8.0	20
Female	116	0.5	47
		2.0	35
		8.0	23

Table 4: Relative attenuation of certain cutoffs.

<u>SIGNAL</u>	<u>CUTOFF (dB)</u>	<u>COMPRESSION</u>	<u>INV. COMPRESSION</u>	<u>Δ_{RMS} (%)</u>
Male	44	1.1	0.4	0.0
	32	1.3	0.5	2.4
	20	2.4	0.9	17
Female	47	1.1	0.5	0.0
	35	1.3	0.7	1.7
	23	2.0	1.0	11

Table 5: Squelch small samples.

APPENDIX A: FAST ALGORITHMS FOR SIGNAL PROCESSING

All of the algorithms mentioned above and used to perform the experiments have been coded by the author in the C programming language. The experiments were performed on a Sun 3 minicomputer at the Yale University Department of Mathematics, using samples recorded on a Macintosh II personal computer.

Transformation to wavelet bases and back. The fast wavelet transformation and its inverse are described in a preprint of Beylkin, Coifman, and Rokhlin [BCR]. There they discuss fast multiplication by singular integral operators of the vectors thus encoded. A future paper will explore the applications of such operators to problems of noise suppression, speech acceleration, and speech recognition.

Note from Tables 1 and 3 that compression with wavelet packets is 25% to 38% tighter than compression with wavelets. This advantage may be offset by the greater computation time required.

Transformation to wavelet packet bases and back. Efficiently transforming to periodized wavelet packets requires organizing the applications of F_ϵ into a pyramid scheme. Suppose $v = v^0 = \sum_{j=0}^{N-1} v_j^0 w_j^0$ is a vector in $l^2(N)$ written in the standard basis w^0 . Then $v = \sum_{j=0}^{N-1} v_j^1 w_j^1 = \cdots = \sum_{j=0}^{N-1} v_j^n w_j^n$ are the various representations of v in the n wavelet packet bases as $v^m \equiv (v_0^m, \dots, v_{N-1}^m)$, $m = 1, 2, \dots, n$. It will be shown that going from v^m to v^{m+1} requires Nr operations, where r is the filter length. Thus filling the entire rectangle $v^0 \mapsto v^1 \mapsto \cdots \mapsto v^n$ costs nrN operations.

But $\{0, 1, \dots, 2^n - 1\}$ has a “dyadic decomposition” into subsets of 2^m consecutive indices, $m = 0, 1, \dots, n$. Denote these subsets by $\Delta_i^m = \{i2^m, i2^m + 1, \dots, i2^m + 2^m - 1\}$, for $i = 0, 1, \dots, 2^{n-m} - 1$. Call m the *scale* of the interval. Clearly $\Delta_i^m \cap \Delta_j^m = \emptyset$ if $i \neq j$. Each dyadic interval of scale m contains two disjoint (left and right) *daughters* of scale $m-1$, and is contained in a unique *parent* of scale $m+1$. In fact, $\Delta_{2i}^{m-1} \cup \Delta_{2i+1}^{m-1} \subset \Delta_i^m \subset \Delta_{\lfloor i/2 \rfloor}^{m+1}$, where as usual $\lfloor x \rfloor$ denotes the greatest integer less than or equal to x .

The convolutions H, G can be arranged so that their images are indexed by the daughters of the domain interval. Denote by V_i^m , $m = 0, 1, \dots, n$, $i = 0, 1, \dots, 2^{n-m} - 1$, the 2^m -

dimensional subspaces of \mathbf{R}^N spanned by the standard basis vectors $\{e_j : j \in \Delta_i^m\}$. Then $\mathbf{R}^N = \oplus \sum_{i=0}^{2^{n-m}-1} V_i^m$, and composing with the appropriate injections gives, for each m ,

$$\begin{aligned} H : V_i^m &\rightarrow V_{2i}^{m-1}, & \text{for } i = 0, 1, \dots, 2^{n-m}-1, \\ G : V_i^m &\rightarrow V_{2i+1}^{m-1}, & \text{for } i = 0, 1, \dots, 2^{n-m}-1. \end{aligned}$$

Applying H, G successively to V_0^m, V_1^m, \dots produces N new coordinates. Each operator requires $rN/2$ multiplications, giving a total of rN operations for this step. It remains to show that these steps successively develop the coefficients v^1, v^2, \dots, v^n , but this is a straightforward induction. By Lemma 3,

$$v_j^{m+1} = \langle e_{\hat{j}}, F_{\epsilon_{n-m}(j)} F_{\epsilon_{n-m+1}(j)} \dots F_{\epsilon_n(j)} v \rangle = \langle e_{\hat{j}}, F_{\epsilon_{n-m}(j)} v^m | \Delta_j^m \rangle,$$

where $j \in \Delta_j^m$ determines \hat{j} uniquely,

$$= \begin{cases} (H v^m | \Delta_j^m)_{\hat{j}}, & \text{if } j \text{ is in the left daughter of } \Delta_j^m, \\ (G v^m | \Delta_j^m)_{\hat{j}}, & \text{if } j \text{ is in the right daughter of } \Delta_j^m. \end{cases}$$

Now let m range from 0 to $n-1$.

To reconstruct vectors from their wavelet packet representations, this pyramid scheme must be reversed. Composing with the appropriate injections gives

$$\begin{aligned} H^* : V_{2i}^{m-1} &\rightarrow V_i^m, & \text{for } i = 0, 1, \dots, 2^{n-m}-1, \\ G^* : V_{2i+1}^{m-1} &\rightarrow V_i^m, & \text{for } i = 0, 1, \dots, 2^{n-m}-1. \end{aligned}$$

Since the ranges of H^* and G^* are orthogonal by Eq.(8), as are the subspaces $V_i^m, i = 0, \dots, 2^{n-m}-1$, the previous wavelet packet coefficients can be recursively computed:

$$v_j^{m-1} = (H^* v^m | \Delta_{L(j)}^m + G^* v^m | \Delta_{R(j)}^m)_{\hat{j}},$$

where $\hat{j} \equiv j \pmod{2^{n-m+1}}$, and $L(j), R(j)$ index the left and right daughters of the unique dyadic subset at level $m-1$ which contains j .

Each adjoint filter convolution requires $r2^m/2$ multiplications on each of the 2^{n-m} subspaces V_i^m , plus the results must be added at the end. The operation count is therefore $O(rN)$. Letting m range from m_0 to 0 rebuilds the signal from any complete level m_0 of coefficients.

It remains to show that any admissible subset of coefficients suffices to rebuild a complete level. The most general admissible subsets considered here are those of the next section, namely the disjoint dyadic covers. But if any member of the admissible subset is on the bottom level, it is there together with all the other members of its scale- n dyadic subset, and also with its twin sister on that level. The algorithm reconstructs each parent from its two daughters, which may then be discarded from the bottom, leaving another admissible subset with fewer levels. Clearly, the operation count to reconstruct some of the parents is at most the operation count to reconstruct all of them, or $O(rN)$. By induction on the number of levels, the algorithm after at most n steps yields an admissible subset on level 0, i.e., the original signal. The total operation count is $O(rN \log_2 N)$, since there are at most $n = \log_2 N$ levels.

Searching for optimal admissible subsets. Suppose that a vector $v \in \mathbf{R}^N$ has been expanded into all wave packet bases, with the coordinate vectors v^0, \dots, v^n forming rows in a rectangular array. An *admissible subset* Σ of this array, as defined in the introduction, corresponds to any disjoint cover of $\{0, 1, \dots, N-1\}$ by the dyadic subsets $\Delta_j^m, m = 0, 1, \dots, n, j = 0, 1, \dots, 2^{n-m+1}$.

The correspondence is given by: w_j^m is part of the basis $\iff \Delta_j^m \in \Sigma$, where $j \in \Delta_j^m$ determines j uniquely.

An *optimal basis* for a given vector v and threshold ϵ is an admissible subset for which as many indexed coordinates are smaller than ϵ (in absolute value) as for any other. There are more than 2^N admissible subsets, but the optimal basis may be found in $O(N \log_2 N)$ operations via the following algorithm.

For simplicity of description, suppose that the array to be searched contains only “ones” (large entries) and “zeros” (entries smaller than the threshold). Finding an optimal basis

means finding an admissible subset with the maximal number of zeros, or equivalently the minimal number of 1's. This search may be implemented recursively, and its properties proved by induction.

Consider a pair of dyadic subintervals Δ_{2i}^k and Δ_{2i+1}^k at level k , and suppose that each has been assigned the cost of “ones” present in the most efficient wavelet packet subbasis of itself. The sum of these costs may be compared to that of the joint parent Δ_i^{k-1} . If the parent is cheaper, it is marked “kept” and is assigned its own cost. If the two children sum to a lower cost, then this sum is assigned to the parent but the parent is not marked.

To start the induction, we mark as “kept” all the children at the bottom level n , and assign each of them their own cost in “ones.” The induction on k proceeds until $k = 0$. At each step, the chosen subset below each parent is the disjoint dyadic cover of that parent interval which indexes the maximal number of zeros. The union of these over all parents in that row is always an admissible subset. At the last step, when there is only one parent Δ_0^0 , the chosen subset will be an optimal basis for the original signal. Of course, it is not necessarily unique.

Advancing from row m to $m-1$ requires counting zeros among N numbers for size, in groups of 2^m , then comparing these counts against adjacent sums from a table of $2N/2^m$ previous counts. Altogether, advancing one step takes $O(N)$ operations. There are $n = \log_2 N$ steps, resulting in the $O(N \log_2 N)$ complexity.

The topmost “kept” intervals in the marked tree constitute the best basis. They are exactly those intervals which are better than any combination of their descendents, but which have no better ancestors. These “best” intervals can be found by a depth-first search whose complexity is $O(N)$, the size of the tree of dyadic subintervals down to level $\log_2 N$.

To estimate the storage requirements of this algorithm, note that passing from step m to step $m-1$ requires a table of $2N/2^m$ integers from the set $0, 1, \dots, 2^m-1$, which takes N bits of space. This may be reused. Recording the current best admissible subset at each row may be done with $2N$ bits. Order the family of $2N$ dyadic subintervals in any manner, labeling the result $\{\Delta(k), k = 1, 2, \dots, 2N-1\}$. Let D be the unique integer between 0 and $2^{2N}-1$ with the property that $\epsilon_k(D) = 1 \iff \Delta(k)$ is a maximal subinterval of

the admissible subset. Otherwise, put $\epsilon_k(D) = 0$. D is constructed during the search, and if the dyadic subsets are correctly ordered, the storage requirements for the fraction of D determined at step m are $N + N/2 + \dots + N/2^m$. After the last step, D is equivalent to an optimal basis.

There are more efficient ways of coding the basis description, of course, for example by using standard lossless compression methods to eliminate repetition.

Transmission of compressed signals. After expansion into wavelet or wavelet packet coefficients and compression by discarding the small ones, it is necessary to transmit the surviving values as well as their positions in the dyadic cover. Suppose that the signal consists of N samples each b bits long, thus requiring Nb bits of storage. When transformed into the optimal basis, say that only N_ϵ of the coefficients exceed the cutoff ϵ . Since no cheating is allowed, the coefficients can only contain b bits each. Specifying their positions in an N by $\log_2 N$ array takes $\log_2 N + \log_2 \log_2 N$ bits, or equivalently the level and position within that level of the coefficient. This adds up to $N_\epsilon(b + \log_2 N + \log_2 \log_2 N)$ bits. The compression ratio of the transmission method is then

$$\rho = \frac{N_\epsilon}{N} \times \left(1 + \frac{\log_2 N + \log_2 \log_2 N}{b}\right).$$

The first factor is the compression ratio computed above. The second factor may be controlled by breaking up the signal into segments of approximately the same length as the dynamic range 2^b .

REFERENCES

- [BCR] Gregory Beylkin, Ronald R. Coifman, and Vladimir Rokhlin, *Fast wavelet transforms and numerical algorithms I*, Yale University Research Report YALEU/DCS/RR-696 (August, 1989).
- [CMQW] Ronald R. Coifman, Yves Meyer, Steven Quake and M. Victor Wickerhauser, *Signal processing and compression with wave packets*, Proceedings of the conference on wavelets, Marseilles, Spring 1989.
- [D] Ingrid Daubechies, *Orthonormal bases of compactly supported wavelets*, Communications on Pure and Applied Mathematics **XLI** (1988), 909–996.
- [Ma] Stephane G. Mallat, *A Theory for Multiresolution Signal Decomposition: The Wavelet Decomposition*, IEEE Transactions on Pattern Analysis and Machine Intelligence **11** (1989), 674–693.
- [N] Nicolas, *Acoustic signature recognition with arborescent wavelets.*, Proceedings of the conference on wavelets, Marseilles, Spring 1989.

Permanent address: Department of Mathematics, Washington University, St. Louis, Missouri 63130