Chapter 10

Design Methods

10.1 Distortions in Image Compression

One of the most successful applications of the wavelet transform is transform-based image compression (also called *image coding*). The coder in Figure 10.1 operates by transforming the data to remove redundancy, then quantizing the transform coefficients (a lossy step), and finally entropy coding the quantized output. Because of superior energy compaction and correspondence with the human visual system, wavelet compression has produced good results. Since the wavelet basis functions have short support for high frequencies and long support for low frequencies, large smooth areas of an image may be represented with very few bits. High frequency detail is added where it is needed.

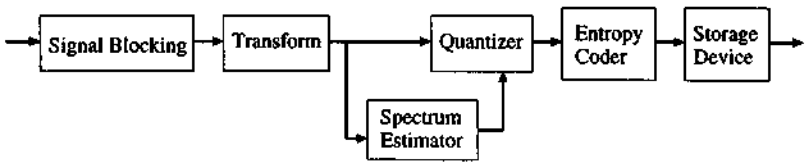

Figure 10.1: Transform-based image coder.

Figure 10.2 shows the original image and its decomposition (three levels) using the two-channel biorthogonal filter bank with $(9, 7)$ coefficients. The upper left block is a smoothed approximation of the original image; it comes from three iterations of the lowpass filter with downsampling. The other subbands contain details at various scales. Since the filters have four vanishing moments, the signal energy in those subbands is small. Consequently, they have to be amplified in order to see the details.

The ten subimages are processed by a spectrum estimator for bit allocation. The statistical properties of the subimages guide this quantization step. The lowpass subband (upper left) is allocated the most bits. One strategy is based on the rate distortion curve. Given B bits, how does one allocate b_k bits per pixel to the kth subimage such that the reconstructed image has smallest distortion? Important distortion measures are mean-square-error (MSE), peak signal-to-noise (PSNR), and maximum error.

At high and medium bit rate (low and medium compression ratio), there is a strong correlation between the image quality and these objective measures. A good compression algorithm

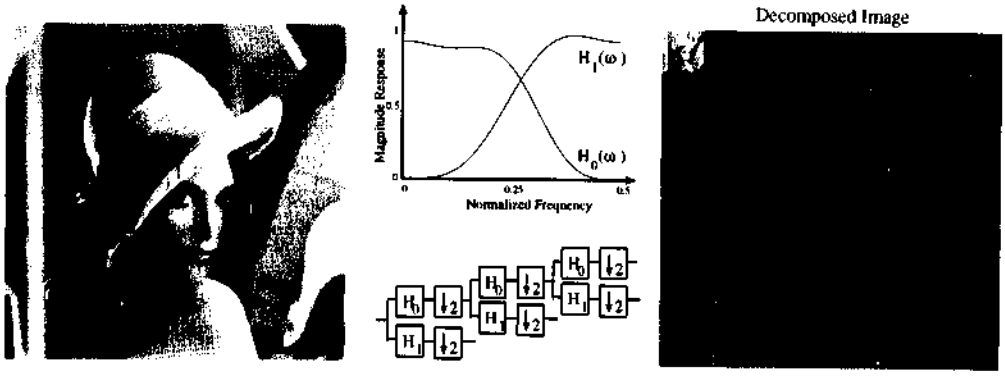

Figure 10.2: Original image "Lenna" and its multiresolution decomposition. The filter bank used is the (9, 7)-tap linear-phase filter bank.

would reconstruct the image with low MSE, high PSNR, and low Max error. *At low bit rate* (high compression), *the basis function characteristics have profound effects on the quality of the reconstructed image*. Typical distortions are ringing, blocking, and blurring artifacts.

Figure 10.3 shows the reconstructions for compression ratios of 32 : 1 (0.25 bits per pixel) and 64 : 1 (0.125 bpp) and 100 : 1 (0.08 bpp). The objective measures for 32 : 1 compression are MSE = 37.32 , PSNR = 32.41 dB, Max error = 62.66 . For 64 : 1 compression the error is larger: MSE = 61.36, PSNR = 30.25 dB, Max error = 68.67; and for 100 : 1 compression MSE = 104.3, PSNR = 27.95 dB, Max error = 121.3.

For medium bit rate (0.25 bpp), the objective measures are good indicators of the subjective quality of the image. At low bit rate, ***this is not always the case***! There are many artifacts in the reconstructed image at 0.08 bpp. The choice of filters is important, as well as bit allocation.

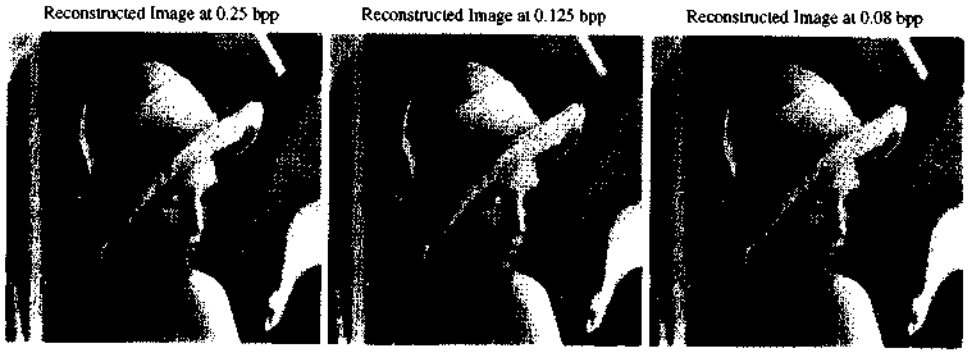

Figure 10.3: Reconstructed images for 32 : 1, 64 : 1, and 100 : 1 compression with (9, 7) filters.

With four levels and 13 subbands, there are several artifacts in the reconstructed images: blurring, border distortions, blocking, checkerboarding, and ringing. Each artifact will be elaborated below. They are the results of the filter choices, the bit allocation algorithms, and the convolutions.

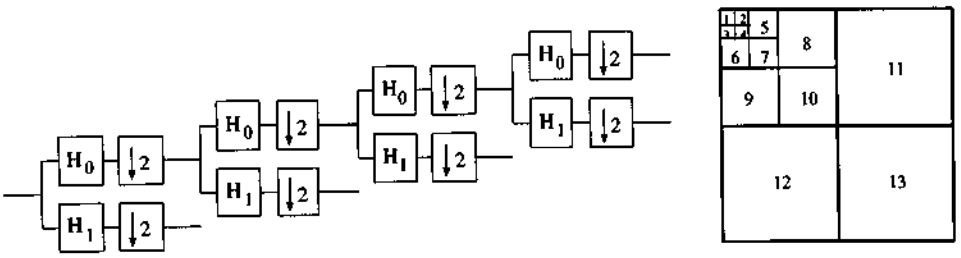

Blurring Artifacts

Blurring occurs when the algorithm does not assign enough bits to the higher subbands 5 to 13. In the extreme case, we allocate all bits to subbands 1 to 4 and none to subbands 5 to 13. The compression ratio is 48 to 1 (0.167 bpp). Figure 10.4 shows that the reconstructed image is essentially an interpolated version of subband 1. It suffers from blurring.

The blurring will decrease if one allocates more bits to the higher subbands. Figure 10.4(b) shows the reconstructed image when the bit allocations are

9.29 5.63 6.52 4.96 3.38 4.20 2.81 0.97 1.73 0.23 0.00 0.00 0.00

The resulting compression is also 48 to 1. This image is much sharper.

Figure 10.4: Blurring in image compression: Too few bits for the detailed subbands.

Border Distortions

The image and the filters have finite length. One has to be careful in computing their convolution. Border distortion is the result of choosing the wrong extension. Chapter 8 discusses the filtering operations for finite-length signals by zero-padding of the input signals, circular convolution,

and symmetric extension. In image compression at 32 to 1 (0.25 bpp), with the (9, 7) symmetric filters, we can compare the effects of these three extensions.

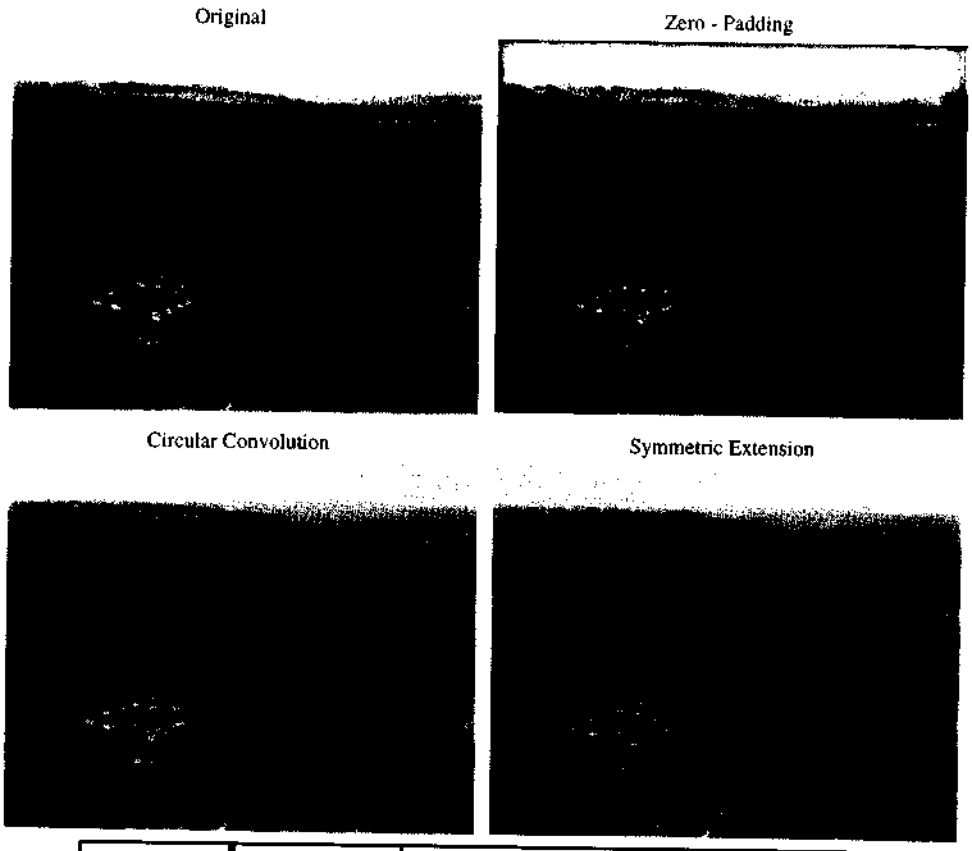

Figure 10.5	Zero-Padding	Circular Convolution	Symmetric Extension
MSE	278.40	96.10	88.67
PSNR (dB)	23.68	28.30	28.75
Max Error	168.40	54.43	55.26

The zero-padding method assumes that images are zero outside their borders. Therefore they are discontinuous signals. The pixel intensity at the reconstructed border is very low (dark). The large error of zero-padding occurs at the boundary.

Circular convolution assumes that the image is periodic. This assumption is not valid for our original image. If the subband images are not quantized, the reconstructed image has no distortion since the filter bank is a PR system. But quantization yields errors at the border when using circular convolution. The periodic assumption results in a discontinuous intensity level at the border.

Symmetric extension extends the images in a continuous way at the border. Although the objective performance of circular convolution is close, the subjective (perceptual) quality from

circular convolution is not as good as from symmetric extension.

Blocking Artifacts

One major disadvantage of the DCT in image coding is that it cannot be used for high compression because of its blocking artifacts. The basis functions typically have short length $M = 8$ for the DCT block transform with no overlap. Blocking will show up at the output. One way to improve is to use longer basis functions, which overlap neighboring blocks. Figure 10.5 also shows the reconstructed image using GenLOT (length 48). Most of the blocking artifacts are now unnoticeable.

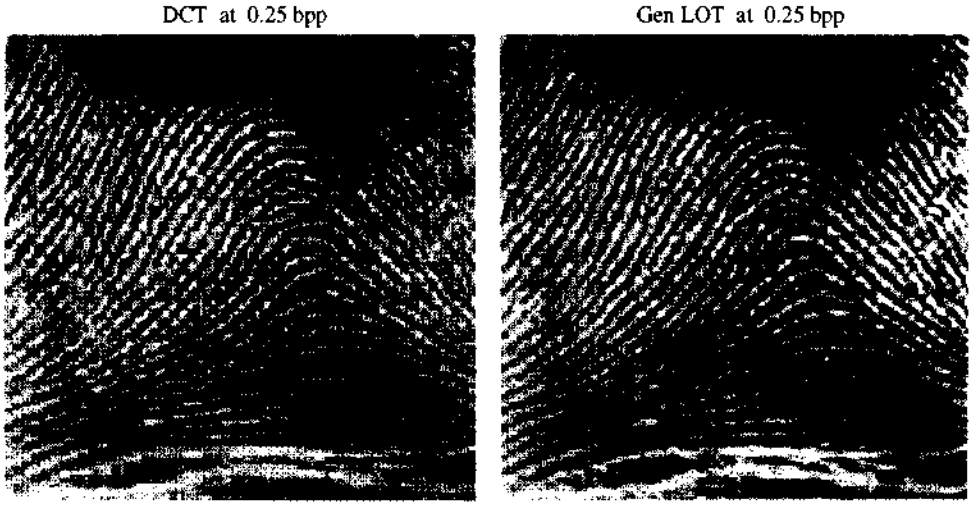

Figure 10.5: Blocking artifacts in image compression using DCT and lapped transforms.

High Frequency Distortion: Ringing

The reconstructed image is a linear combination of the scaling function and wavelets of the ***synthesis bank***. Since symmetric filters are not orthogonal, the analysis and synthesis functions are not the same. The scaling function of the synthesis filter should be smooth and the wavelet should be short. The error is larger at higher subbands, assigning fewer bits to the wavelet coefficients. When the wavelet has length N, the error at the pixel (K, L) affects its neighbors $(K \pm k, L \pm \ell)$, where $k, \ell \leq N/2$.

Figure 10.6 shows the original image and reconstructed image at 0.5 bpp. The lengths of the synthesis filters are 11 and 17. Even at medium bit rate, the ringing effect is quite severe. This is especially true at strong edges.

Nonsmooth Functions: Checkerboard Artifacts

In all these compression examples, the filters are linear-phase factors of a maximally-flat halfband filter. All scaling functions and wavelets in the synthesis bank are smooth, which is essential for image reconstruction. When the functions are not smooth, what is the effect on reconstruction?

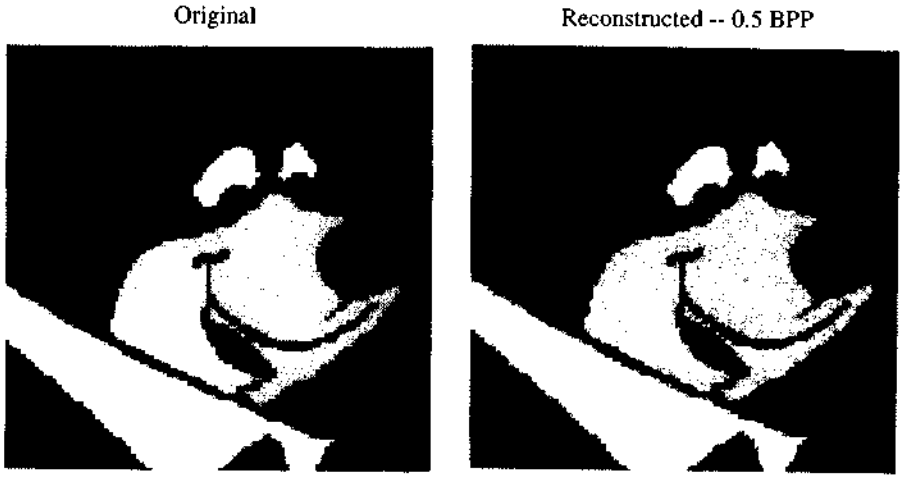

Figure 10.6: Ringing artifacts are clearly visible at edges.

The lowpass responses $H_0(z)$ and $F_0(z)$ in Figure 10.7 are designed such that the stopband error is equiripple (a desirable property in conventional digital signal processing). *The scaling functions are not smooth.* The reconstructed image suffers from the the ***checkerboard effect.*** This can be explained by observing that the scaling function from $F_0(z)$ has a large and narrow peak. The magnitudes near this narrow peak are about 10% below it. This explains the bright horizontal and vertical lines in the reconstructed image. The stronger intensity produces the checkerboard effect. This artifact can be attributed to the "roughness" of $\phi(t)$.

Lowpass synthesis:	***Long and smooth to avoid blocking and checkerboarding***
Highpass synthesis:	***Short to avoid ringing***

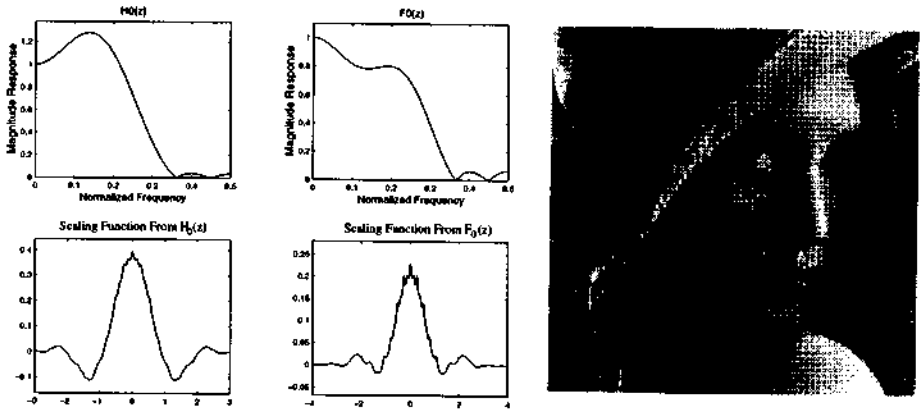

Figure 10.7: These lowpass equiripple filters are linear-phase factors of a halfband equiripple filter. The lengths are 9/7 and the rate is 0.25 bpp. The nonsmooth scaling functions create a checkerboard.

10.2 Design Methods—General Perspective

In all of our examples, one important factor for image quality is the choice of coefficients in the filter banks. For an objective function of coding gain or degree of smoothness or stopband attenuation, *one optimizes the filter coefficients.* This section will present several methods to design and optimize filter banks and wavelets. These methods are based on spectral factorizations, lattice structures and time-domain formulations. We will present several common choices for the objective function Φ, and its parameters, and the constraints to impose:

Design Problem: Minimize Φ (parameters) subject to constraints. (10.1)

The design parameters could be the lattice coefficients or the filter coefficients. Recall how a lattice structure can impose the PR property and also linear phase, pairwise-mirror-image or cosine modulation. The lattice structure of a two-channel paraunitary filter bank includes $\boldsymbol{R}_k =$ rotation by θ_k. There is a one-to-one correspondence between lattice coefficients $\theta_0, \theta_1, \ldots, \theta_N$ and filter coefficients. The advantage of the lattice coefficients is that the quantized filters are still PR. The rotations are still orthogonal when the angles are rounded off.

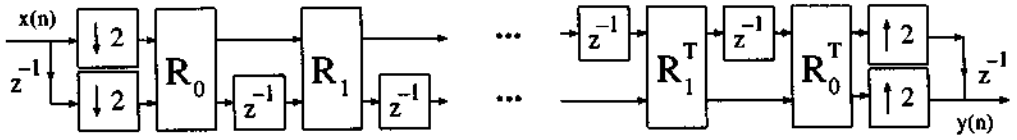

Objective Function Φ

Useful objective functions are the *stopband attenuation, coding gain,* and *degree of smoothness.* Stopband attenuation is very important in audio compression, whereas coding gain and smoothness are more important in image compression.

Stopband Attenuation: Φ measures the passband and stopband errors of $H(z)$:

$$\Phi = \begin{cases} \displaystyle\int_{\Omega} W(e^{j\omega})\,|H_d(e^{j\omega}) - H(e^{j\omega})|^2\,d\omega, & L_2 \text{ norm (energy)} \\ \text{Max}\ \left[W(e^{j\omega})\,|H_d(e^{j\omega}) - H(e^{j\omega})|\right], & L_\infty \text{ norm (maximum).} \end{cases} \tag{10.2}$$

The frequency region Ω consists of the passband $0 \le \omega \le \omega_p$ and stopband $\omega_s \le \omega \le \pi$. $H_d(e^{j\omega})$ is the desired response of the lowpass filter. $W(e^{j\omega})$ is a positive weighting function. Examples of typical $H_d(e^{j\omega})$ and $W(e^{j\omega})$ are:

$$H_d(e^{j\omega}) = \begin{cases} e^{-jL\omega}, & 0 \le \omega \le \omega_p \\ 0, & \omega_s \le \omega \le \pi \end{cases} \qquad W(e^{j\omega}) = \begin{cases} a, & 0 \le \omega \le \omega_p \\ b, & \omega_s \le \omega \le \pi. \end{cases} \tag{10.3}$$

If one increases the ratio b/a, the stopband error δ_s is much smaller than the passband error δ_p. In a filter bank with M channels, Φ has a term from each filter $H_k(e^{j\omega})$. The bandpass filters have two stopbands enclosing one passband.

The analysis filters in a paraunitary bank satisfy $\sum |H_k(e^{j\omega})|^2 = 1$. Consequently, it is sufficient to consider only the stopbands. With three filters, the figure shows $\frac{\pi}{3} \le \omega \le \frac{2\pi}{3}$ as a

stopband of H_0 and H_2 and as the passband of H_1. In its passband, $H_1(e^{j\omega})$ must approximate unity. The passband error $\delta_{p1} = 1 - \sqrt{1 - \delta_{s0}^2 - \delta_{s2}^2}$ is very small, if both stopband errors δ_{s0} and δ_{s2} are small. Filter banks with high stopband attenuation are essential in audio compression and signal detection.

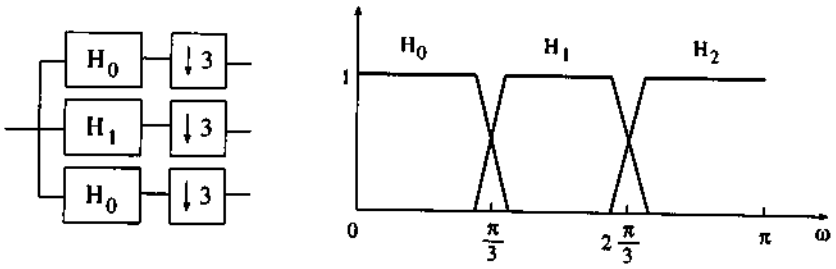

Example 10.1. Consider the design of a linear-phase filter bank with (9, 7) taps based on maximum error in the stopband. Figure 10.8(a) shows the equiripple stopband responses of $H_0(z)$ and $F_0(z)$. They are the factors of an equiripple halfband filter of length 15. There is no zero at $\omega = \pi$ and the scaling function and wavelets are *not smooth*. This pair of filters was used earlier to demonstrate checkerboard artifacts in image compression.

Figure 10.8(b) shows the frequency responses of a 9/7 maximally-flat halfband filter. *The stopband attenuation of both filters is bad.* However, each filter has four zeros at $\omega = \pi$. The scaling functions and wavelets are very smooth. They are the filters used in the FBI fingerprint image compression standard.

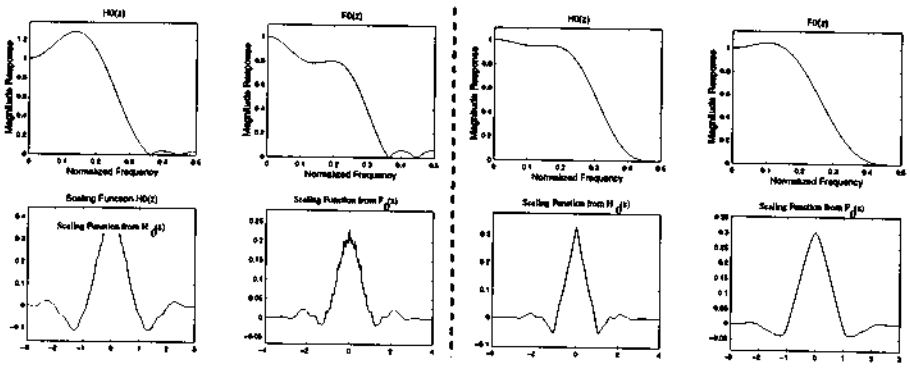

Figure 10.8: Example of equiripple design (left) and maximally-flat design (right).

Example 10.2. Figure 10.9 shows the filter responses of a cosine-modulated filter bank (8 channels, filter lengths 128). The objective function is $\int_{\omega_r}^{\pi} |H(e^{j\omega})|^2 \, d\omega$, the stopband energy of the prototype filter.

Coding Gain: A uniform quantizer has an equivalent noise model. The input signal $x(n)$ is quantized to b bits to obtain $v(n)$. For large b, one can assume that the quantization error $q(n) = v(n) - x(n)$ is a random variable, uniformly distributed in $-2^b \leq q \leq 2^b$. The error variance of $q(n)$ is $\frac{1}{3}2^{-2b}$.

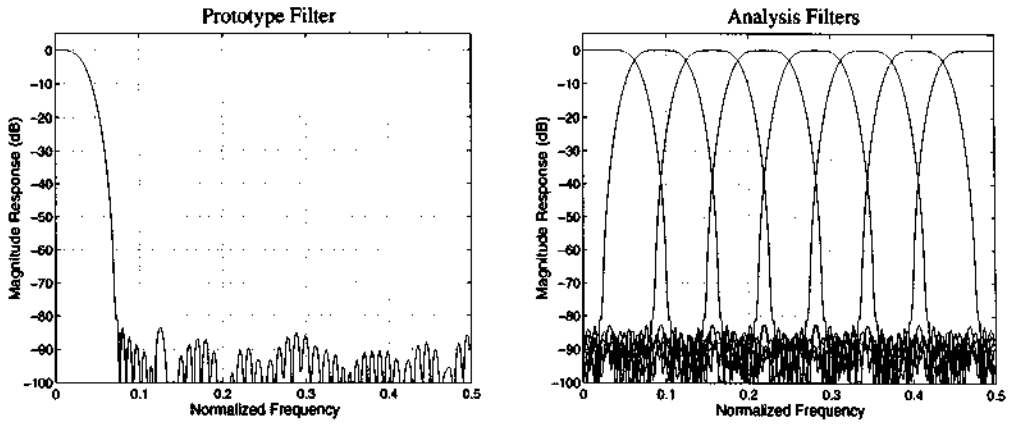

Figure 10.9: A cosine-modulated filter bank with stopband attenuation as objective Φ.

Instead of quantizing the input directly, one can transform it into another domain. The input signal is blocked into length N and transformed to length M by $\boldsymbol{P}^T$. The subband signal $y_k(n)$ is then quantized to produce $\widehat{y}_k(n)$. The transform should be designed so that the energy of $\boldsymbol{x}$ is concentrated in the first few coefficients of y_k. The number of bits in the quantizers Q_k should be proportional to the energy of $y_k(n)$.

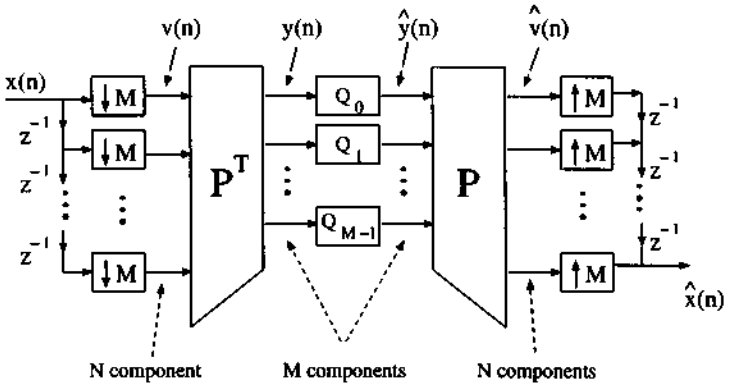

The coding gain measures the energy compaction from the transform. An orthogonal transform with high coding gain yields reconstructed images with high quality at low bit rate. Chapter 11 will discuss coding gain in more detail. For *biorthogonal* systems, [Katto-Yasuda] derived a unified formula when the input signal $\boldsymbol{x}(n)$ is modeled as a Markov process. Its inter-sample correlation factor is ρ. The coding gain is G when the sampling ratios are α_k:

$$G(\rho) = \prod_{k=0}^{M-1} (A_k B_k)^{-\alpha_k} \quad \text{where} \quad \left\{ \begin{array}{lcl} A_k & = & \sum_n \sum_m h_k(n)\, h_k(m)\, \rho^{|m-n|} \\ B_k & = & \sum_n (f_k)^2(n) \end{array} \right.$$

For example, a 3-level wavelet transform has four analysis filters:

$$H_0(z)H_0(z^2)H_0(z^4) \quad H_0(z)H_0(z^2)H_1(z^4) \quad H_0(z)H_1(z^2) \quad H_1(z). \tag{10.4}$$

The coefficients in these combinations are used in the formula for coding gain. Assuming

$\rho = 0.95$, one can compute G and improve the performance. The sampling ratios are 1/8, 1/8, 1/4 and 1/2.

Example 10.3. Figure 10.10 shows the frequency responses of an 8-channel GenLOT with length 48, for maximum stopband attenuation and also for maximum coding gain. Note that the responses are different. These GenLOTs are compared at 50 : 1 compression. The objective performances with high stopband attenuation are MSE = 135, PSNR = 26.84 dB and Maximum error = 96. The corresponding numbers for the GenLOT with high coding gain are 130, 26.99 dB and 84. Even though these are comparable, maximizing the attenuation produced blocking artifact. This is not due to the block length 48, but it is due to the DC leakage. The bandpass and highpass filter responses are not zero at $\omega = 0$, so the DC term has leaked out of the lowpass.

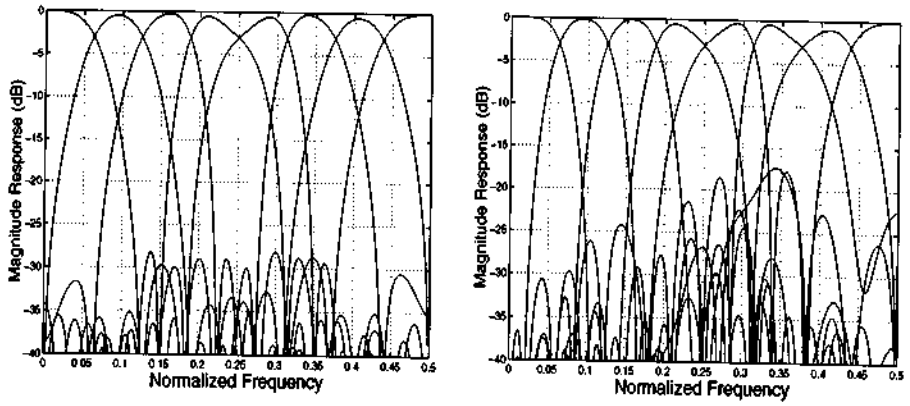

Stopband Attenuation Coding Gain

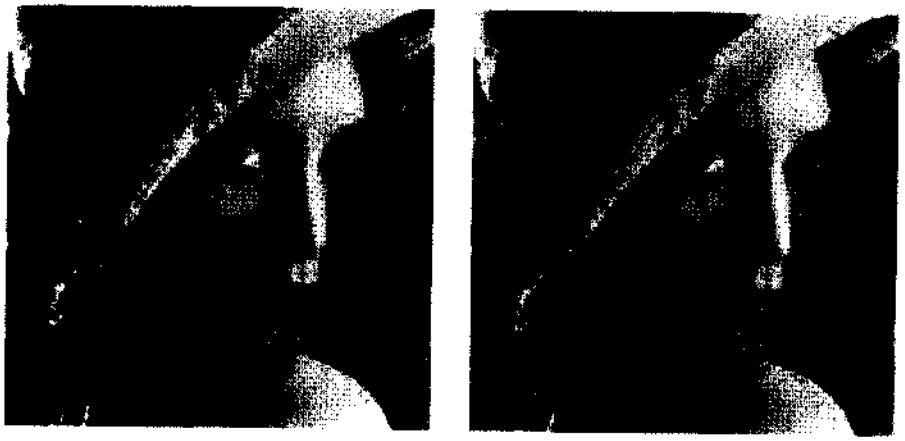

Figure 10.10: GenLOT designs optimized for stopband attenuation (left) and coding gain (right).

Constraints

PR constraints: The equivalent conditions for M-channel perfect reconstruction are

$$\begin{cases} \sum_n h_k(n) f_\ell(2m-n) = \delta(m)\delta(k-\ell) \\ \boldsymbol{F}_p(z)\,\boldsymbol{H}_p(z) = z^{-n_0}\boldsymbol{I}. \end{cases} \tag{10.5}$$

How does one design $H_k(z)$ and $F_k(z)$ that satisfy (10.5)? One way is to design $h_k(n)$ first and then solve for $f_k(n)$. For arbitrary analysis filters, synthesis filters might not exist. The time domain approach in [Nayebi1] presents an iterative method to solve for $h_k(n)$ and $f_k(n)$.

For FIR filters, (10.5) is true only if the determinant of $\boldsymbol{H}_p(z)$ is a delay. The special case of *paraunitary* $\boldsymbol{H}_p(z)$ can be parametrized by a set of rotation angles or Householder matrices. This is the lattice structure. The PR constraints are *automatically satisfied.*

For two channels, a simple PR condition exists for FIR systems: $H_0(z)F_0(z)$ is a halfband filter. This is the result when aliasing is cancelled. Such simplicity does not exist for systems with more channels.

Smoothness constraints: From the image compression examples, it is clear that the synthesis functions should be smooth. This relates directly to the number p of zeros at $\omega = \pi$ of $F_0(z)$. The smoothness property is also known as the DC-leakage property in traditional image processing.

Example 10.4. The lowpass Daubechies orthogonal filter $H_0(z)$ is the minimum-phase spectral factor of a maximally-flat halfband filter $P(z)$. Since $P(z)$ has $2p$ zeros at $\omega = \pi$, $H_0(z)$ and $F_0(z)$ have p zeros. Note that the functions in Figure 10.11 of D_{24} are considerably smoother than D_4, since D_{24} has 12 zeros at $\omega = \pi$.

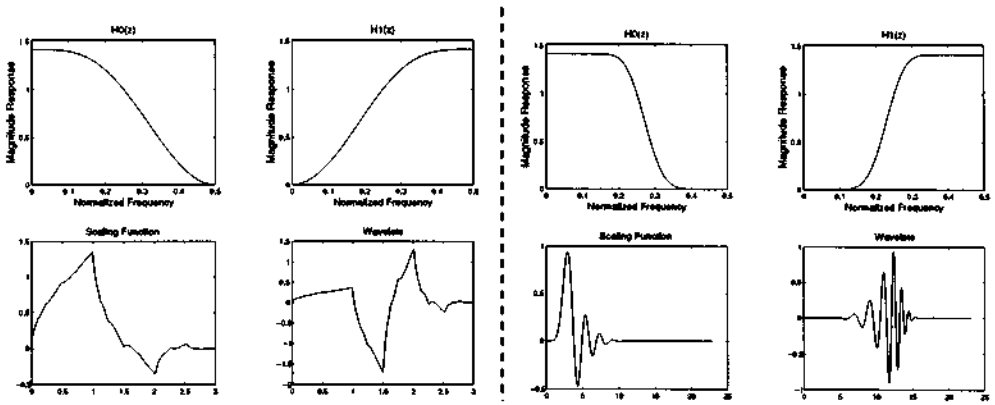

Figure 10.11: Magnitude response, scaling function and Daubechies wavelets D_4 and D_{24}.

10.3 Design of Perfect Reconstruction Filter Banks

The three main design methods extend in different ways to M channels:

Spectral factorization	factor an M-th band filter
Lattice structure method	minimize Φ (lattice coefficients)
Time-domain method	minimize Φ (filter coefficients) subject to PR.

The first approach uses spectral factorization. We find roots or factors $F_0(z)H_0(z)$ of a linear-phase polynomial $P(z)$. For applications that require large stopband attenuation, $P(z)$ has many

zeros on or near the unit circle. Its factorization becomes difficult. A special class that is very useful in function approximation is the halfband maxflat Daubechies filter (Section 5.5). One can distribute its zeros in several ways to $H_0(z)$ and $F_0(z)$ to obtain:

$$\begin{cases} \text{Paraunitary filter banks} & -- & \text{Orthogonal Daubechies wavelets} \\ \text{Type A linear phase filter banks} & -- & \text{Symmetric even-length wavelets} \\ \text{Type B linear phase filter banks} & -- & \text{Symmetric odd-length wavelets.} \end{cases} \tag{10.6}$$

These wavelets have excellent approximation from the zeros at $\omega = \pi$.

The second approach is based on the lattice structure implementation of the filter bank. The unknown parameters are the lattice coefficients k_i. Since all PR filter banks are characterized by an appropriate set of k_i, the design problem reduces to minimizing the objective function $\Phi(k_i)$. Then the corresponding filters can be computed. The minimization is a challenging problem because of the *nonlinear relation* between the lattice and filter coefficients.

The time-domain approach expresses the PR conditions directly on the filter coefficients. Nonlinear quadratic constrained optimization yields good filters. One optimization method formulates the PR conditions in terms of matrices and the filter coefficients [Nayebi1]. The algorithm iteratively solves (10.5). An alternative is to formulate the objective function and constraints in quadratic forms:

$$\text{Minimize } \boldsymbol{h}^T \boldsymbol{P} \boldsymbol{h} \text{ subject to } \boldsymbol{h}^T \boldsymbol{Q}_k \boldsymbol{h} = c_k. \tag{10.7}$$

We may use Lagrange multipliers (explicitly or implicitly). The advantage of this constrained quadratic approach is that the derivatives and the Hessian of both the objective function and the PR conditions can be computed exactly.

Spectral Factorization Method

Given a polynomial of the form $P(z) = F_0(z)H_0(z)$, the objective here is to find the two unknown polynomials $F_0(z)$ and $H_0(z)$. We assume that all polynomials have real coefficients. Note that $F_0(z) = z^{-N}H_0(z^{-1})$ in a paraunitary filter bank, and the problem becomes a spectral factorization. For a low-degree polynomial $P(z)$, one can find its zeros and distribute them to $F_0(z)$ and $H_0(z)$. Since all polynomials are real-valued, the zeros have to be in complex-conjugate pairs.

For example, consider the maximally-flat halfband filter $P(z)$ of length 15. Its zeros are shown in Figure 10.12 with 8 zeros at $\omega = \pi$. Given the 14 zeros of $P(z)$, there are many ways of distributing its zeros to $F_0(z)$ and $H_0(z)$. One can assign half the zeros to $H_0(z)$ and half to $F_0(z)$ in such a way that $F_0(z) = z^{-7}H_0(z^{-1})$. This means that z_0 is a zero of $H_0(z)$ when z_0^{-1} is a zero of $F_0(z)$. The factorization assigns half of the zeros at π to $H_0(z)$, and the remaining half to $F_0(z)$. For minimum phase, $H_0(z)$ has all other zeros inside the unit circle, whereas $F_0(z)$ has all the zeros outside the unit circle. Both filters have length 8. This factorization yields a paraunitary filter bank and the Daubechies D_8 orthogonal wavelets.

Part b shows an alternate zero distribution. The zeros are chosen such that the resulting filters have odd length and linear phase. Note that $H_0(z)$ and $F_0(z)$ have the same number of zeros at $\omega = \pi$. This (9, 7) pair is used in the FBI fingerprint image compression standard.

Part c shows another zero distribution that yields linear phase with even length. Note that the zeros at π for $H_0(z)$ and $F_0(z)$ are not the same. The scaling function and wavelet from the

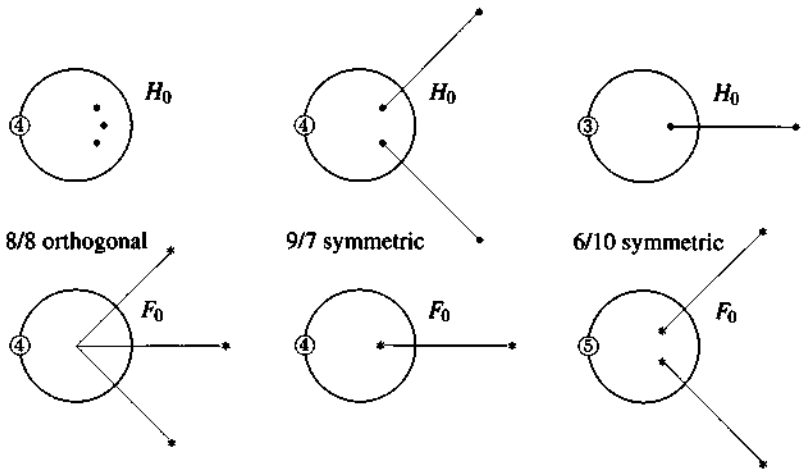

Figure 10.12: The zero distributions of $F_0(z)$ and $H_0(z)$. (a) Paraunitary filter bank (b) Linear phase filter bank (Type B, odd length) and (c) Linear phase filter bank (Type A, even length).

synthesis bank is smoother (it has more zeros at π). This is the (6, 10)-tap filter pair. Its performance is very good in image compression. The three factorizations demonstrate the essential ideas in designing two-channel paraunitary filter banks (orthogonal wavelets) and two-channel linear phase filter banks (symmetric wavelets).

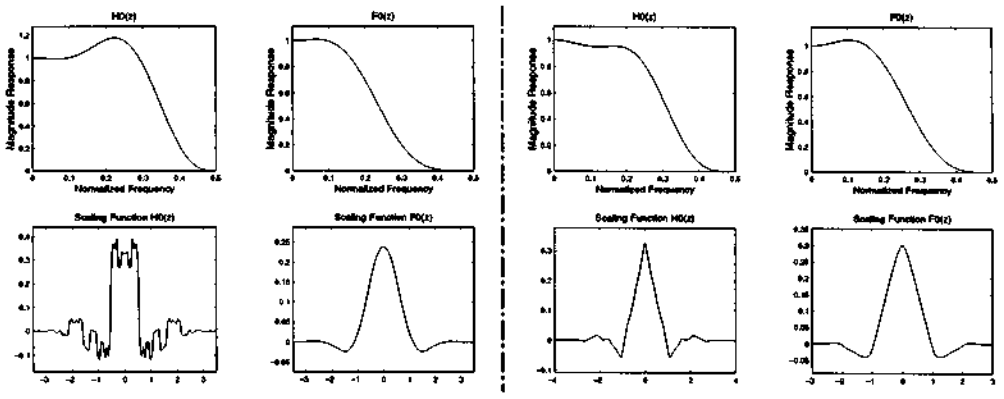

Spectral factorization applies to the design of two-channel filter banks and Near-Perfect-Reconstruction cosine-modulated filter banks. In the two-channel filter bank case, $F_0(z)$ and $H_0(z)$ are the factors of $P(z)$. In the NPR cosine-modulated case, the prototype filter $H(z)$ is a spectral factor of a $2M$-th band filter. For an arbitrary filter bank, such simple PR/NPR conditions do not exist. There is no optimization process involved in spectral factorization. In general, an optimal filter $P(z)$ does not guarantee optimal filters $H_0(z)$ and $F_0(z)$.

For orthogonal filters, the cepstral algorithm of [Mian-Nainer] is widely used. It computes the minimum phase spectral factor without finding the zeros. Section 5.4 describes this (additive) splitting of the logarithm of $P(e^{j\omega})$.

Lattice Structure

Most PR filter banks can be implemented based on lattice structure. For any set of lattice coefficients, the paraunitary, linear-phase, and cosine-modulation properties are structurally imposed. This means that those properties (and also PR) are enforced independent of the lattice coefficient values. The objective function could take many different forms depending on the applications. Examples are stopband attenuation, coding gain, and interference energy. The lattice method is used for two-channel paraunitary filter bank [VaidHoang], the two-channel biorthogonal linear-phase filter bank [V], the M-channel linear-phase paraunitary filter bank [Soman], and the paraunitary cosine-modulated filter bank [Koil2].

Another role for lattices begins in the time domain. Optimization yields a *Near Perfect Reconstruction (NPR)* solution:

$$\sum_n h_k(n) f_\ell(2m-n) < \epsilon < 10^{-10} \quad \text{for } k \neq \ell, m \neq 0. \tag{10.8}$$

From these NPR filters, one can find the lattice coefficients. *Those give PR filters* $h_k^{PR}(n)$ *and* $f_k^{PR}(n)$. If the reconstruction error ϵ is small, the NPR filters and PR filters are close and $\Phi^{PR} \approx \Phi$. If ϵ is not small, the filters are not close and Φ could change drastically.

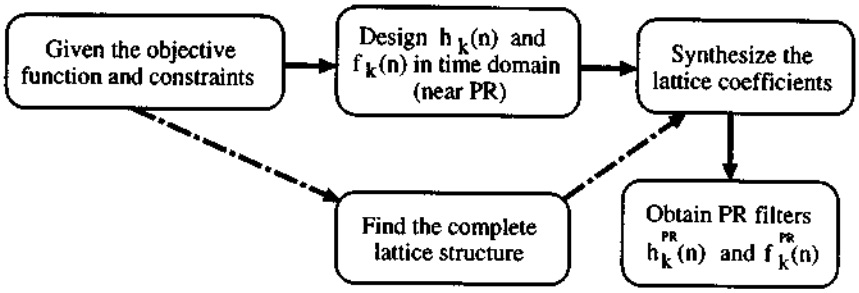

There are two key points to ensure success:

- Existence of *complete* and *minimal lattice structure*. The *completeness* is crucial here since the lattice structure must cover all filters with the given properties. *Minimal* is desirable to ensure efficient implementation.
- The reconstruction error ϵ must be sufficiently small to guarantee that $\Phi^{PR} \approx \Phi$. A typical objective function is stopband attenuation.

Example 10.5. A two-channel linear phase system of length 32 is shown in Figure 10.13(a) (solid lines). These filters are designed with $\epsilon = 10^{-4}$. They are used in the above lattice synthesis procedure to obtain a PR filter bank (broken lines). The attenuation has changed from -42 dB to -25 dB. Apparently, 10^{-4} is not small enough. Figure 10.13(b) shows the frequency response of the NPR solution with $\epsilon = 10^{-14}$ (solid line) and the PR solution (broken line). The NPR filters are very close to PR. Practically, there are no changes in the frequency responses.

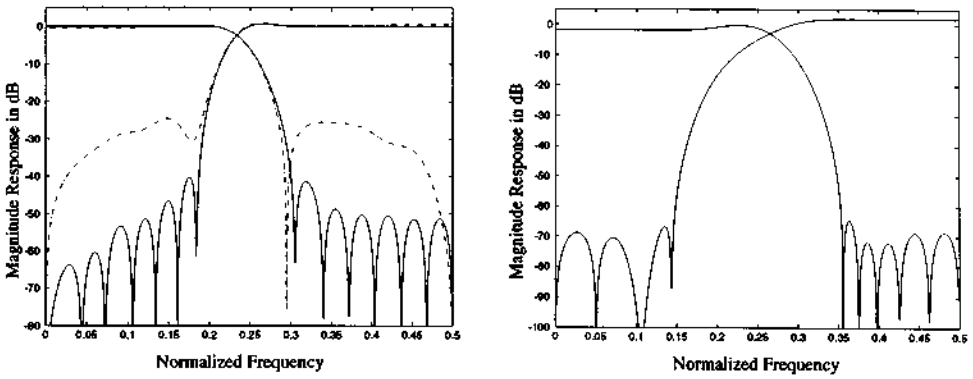

Figure 10.13: NPR and PR filters obtained by lattice synthesis for $\epsilon = 10^{-4}$ (left) and $\epsilon = 10^{-14}$ (right).

What went wrong in the lattice synthesis that changes the filter response drastically when ϵ is not small? Note the subtle difference between PR and NPR lattice structures in Figure 10.14. The delays are no longer z^{-2}, but z^{-1}. There are additional blocks $\widehat{\Gamma}_k = \begin{bmatrix} 1 & \widehat{\gamma}_k \\ \widehat{\gamma}_k & 1 \end{bmatrix}$ between Γ_k. A large ϵ will yield nonzero $\widehat{\gamma}_k$, and setting them to zero will change the filters dramatically. Since $\widehat{\gamma}_k$ are very small if $\epsilon < 10^{-14}$, setting $\widehat{\gamma}_k = 0$ does not change the frequency responses or the objective function.

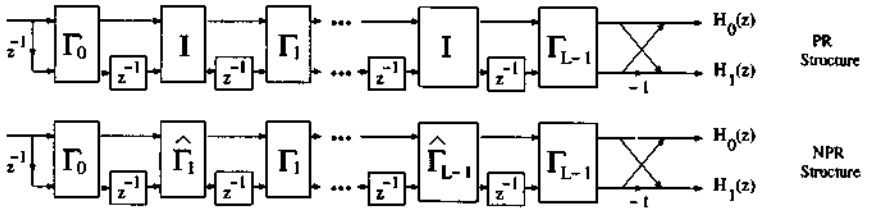

Figure 10.14: Lattice structures for PR and NPR Type-A linear phase filter banks.

Time-Domain Methods

Block matrix approach: This method is suggested by [Nayebi1] to solve the PR equations when the filter length is $N = LM$:

$$\underbrace{\begin{bmatrix} \boldsymbol{P}_0^T & \boldsymbol{0} & \cdots & \boldsymbol{0} \\ \boldsymbol{P}_1^T & \boldsymbol{P}_0^T & \cdots & \boldsymbol{0} \\ \vdots & \vdots & \ddots & \vdots \\ \boldsymbol{P}_{L-1}^T & \boldsymbol{P}_{L-2}^T & \cdots & \boldsymbol{P}_0^T \\ \boldsymbol{0} & \boldsymbol{P}_{L-1}^T & \cdots & \boldsymbol{P}_1^T \\ \vdots & \vdots & \ddots & \vdots \\ \boldsymbol{0} & \boldsymbol{0} & \cdots & \boldsymbol{P}_{L-1}^T \end{bmatrix}}_{\boldsymbol{A}} \underbrace{\begin{bmatrix} \boldsymbol{Q}_0^T \\ \boldsymbol{Q}_1^T \\ \vdots \\ \boldsymbol{Q}_{L-1}^T \end{bmatrix}}_{\boldsymbol{Q}} = \underbrace{\begin{bmatrix} \boldsymbol{0} \\ \vdots \\ \boldsymbol{0} \\ \boldsymbol{J} \\ \boldsymbol{0} \\ \vdots \\ \boldsymbol{0} \end{bmatrix}}_{\boldsymbol{B}} \tag{10.9}$$

J is the exchange matrix (reverse identity). Its location determines the delay of the system. The entries of the filter matrices are $P_{k\ell} = h_k(\ell)$ and $Q_{k\ell} = f_k(\ell)$. The $M \times M$ submatrices P_j and Q_j are defined by:

$$P = \begin{bmatrix} P_0 & P_1 & \cdots & P_{L-1} \end{bmatrix} \quad \text{and} \quad Q = \begin{bmatrix} Q_0 & Q_1 & \cdots & Q_{L-1} \end{bmatrix}. \tag{10.10}$$

Chapter 9 used this formulation for LOT ($N = 2M$ and $L = 2$), where (10.10) reduces to $P_1^T P_1 + P_0^T P_0 = I$ and $P_1^T P_0 = 0$.

In summary, the time domain PR constraints $AQ = B$ are general. Given the analysis filters $h_k(n)$ (and thus A), this is not always solvable. One can find "optimal" synthesis filters $Q = (A^T A)^{-1} A^T B$ by minimizing $\|A Q - B\|$. Then fix those filters $f_k(n)$ and reverse the equation to find optimal $h_k(n)$. This is the iterative block matrix approach.

Quadratic Constrained Least Squares method (QCLS): For an M-channel paraunitary filter bank, the PR condition requires no distortion ($k = 0$) and no aliasing ($k > 0$):

$$T_k(z) = \frac{1}{M} \sum_{\ell=0}^{M-1} z^{-N} H_\ell(z^{-1}) H_\ell(z W^k) = \delta(k)\, z^{-\ell}.$$

Let h be a column vector consisting of all the filter coefficients $h_k(n)$:

$$h = (h_0(0) \cdots h_0(N) \cdots h_{M-1}(0) \cdots h_{M-1}(N))^T.$$

The PR conditions can be written in the form $h^T Q_k h = 0$ and $h^T S_k h = 1$. Furthermore, the objective function for stopband attenuation can be formulated as a quadratic $\Phi = h^T P h$, where P is symmetric and positive definite. The optimized filter is precisely h_{opt} such that

$$h_{\text{opt}} \quad \text{minimizes} \quad h^T P h \quad \text{subject to} \quad \begin{cases} h^T Q_k h = 0, \\ h^T S_k h = 1. \end{cases} \tag{10.11}$$

Since $\mathbf{Q}_k$ is normally not positive definite, the solution can be difficult. However, there are effective optimization procedures that linearize the quadratic constraints [Schittkowski]. These procedures will yield an approximate solution (the constraints are not satisfied exactly). The errors are very small and can be ignored in most practical cases. The QCLS method is used to design the two-channel PR linear-phase filter bank [QCLS], the NPR Pseudo-QMF bank [Nguyen1], the PR cosine-modulated filter bank [QCLS], and the M-channel paraunitary linear-phase filter bank [Nguyen2].

10.4 Design of Two-Channel Filter Banks

This book has emphasized the construction of $H_0(z)$ and $F_0(z)$ as factors of a halfband filter $P_0(z)$. This method is conceptually simple, and good in practice. The other two approaches, based on lattice structure and quadratically constrained optimization, look awkward by comparison.

But those methods have their own advantages. The lattice structure is efficient to implement, once it is found. The optimization can yield "best" filters rather than "good" filters. This brief section will record the details for the two cases of greatest significance: *orthogonal filter banks and linear phase filter banks.*

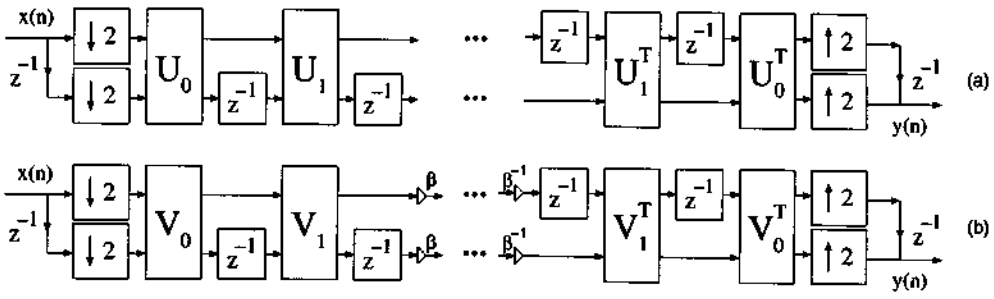

Figure 10.15: Lattice structures for the two-channel paraunitary filter bank. The first structure uses rotation angles as parameters and the second structure uses lattice coefficients.

Lattice Structure Method

Figure 10.15 shows two lattice structures for a paraunitary filter bank and orthogonal wavelets. The objective is to find the right set of angles θ_k in U_k (or lattice coefficients α_k in V_k) such that an objective function is minimized. The objective function considered by [VaidHo] is the stopband error (L_2 norm). There are several tables compiled for orthogonal filters with various frequency specifications.

Figure 10.16 shows the two lattice structures for linear phase filter banks. Recall that Type A yields filters with even length (one is symmetric and the other is antisymmetric). Type B yields symmetric filters with odd length. In Type A, the parameters are the lattice coefficients in Γ_k. [NgVa1] presents design examples for both Type A and B systems, designed using the nonlinear optimization where the parameters are the lattice coefficients.

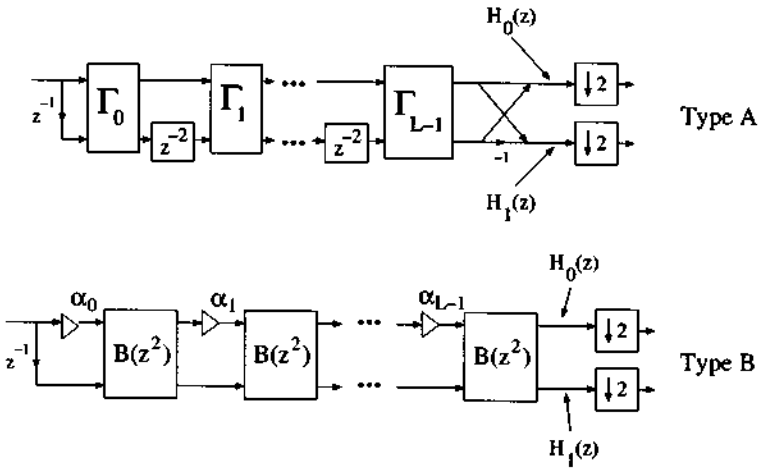

Figure 10.16: Lattice structures for symmetric filter banks, even and odd lengths.

Time-Domain Method—QCLS Formulation of Orthogonality

The orthogonality condition is that $P(z) = z^{-N} H_0(z^{-1}) H_0(z)$ is a halfband filter. We will show how to formulate this in quadratic form $\boldsymbol{h}^T \boldsymbol{Q}_k \boldsymbol{h} = c_k$ using

$$\boldsymbol{h} = [h_0(0)\ h_0(1)\ \cdots\ h_0(N)]^T \quad \text{and} \quad \boldsymbol{e}(z) = \left[1\ z^{-1}\ \cdots\ z^{-N}\right]^T \quad \text{and} \quad H_0(z) = \boldsymbol{e}^T(z)\boldsymbol{h}.$$

Then $z^{-N} H_0(z^{-1}) = \boldsymbol{h}^T \boldsymbol{J} \boldsymbol{e}(z)$ where $\boldsymbol{J}$ is the reverse identity matrix (antidiagonal). Therefore $P(z) = \boldsymbol{h}^T \boldsymbol{J} \boldsymbol{e}(z) \boldsymbol{e}^T(z) \boldsymbol{h}$. Let

$$\boldsymbol{e}(z)\boldsymbol{e}^T(z) = \sum z^{-k} \boldsymbol{D}_k \quad \text{where} \quad [\boldsymbol{D}_k]_{i,j} = \begin{cases} 1; & i + j = k \\ 0; & \text{otherwise.} \end{cases} \tag{10.12}$$

Then $P(z)$ can be rewritten as

$$P(z) = \sum z^{-k} \boldsymbol{h}^T \boldsymbol{J} \boldsymbol{D}_k \boldsymbol{h} = \sum p(k) z^{-k}. \tag{10.13}$$

In other words, the kth coefficient of $P(z)$ is $p(k) = \boldsymbol{h}^T \boldsymbol{J} \boldsymbol{D}_k \boldsymbol{h}$. The halfband condition on $P(z)$ gives the constraints in the design problem:

$$\min_{\boldsymbol{h}} \ \boldsymbol{h}^T \boldsymbol{P} \boldsymbol{h} \quad \text{subject to} \quad \boldsymbol{h}^T \boldsymbol{J} \boldsymbol{D}_{N-1-2n} \boldsymbol{h} = \delta(n).$$

Here $\boldsymbol{h}^T \boldsymbol{P} \boldsymbol{h}$ denotes the stopband energy. See [Nguyen3] for more examples.

Lagrange Multiplier Method

With the choice of $F_0(z) = H_1(-z)$ and $F_1(z) = -H_0(z)$ to cancel aliasing, the filters $H_0(z)$ and $H_1(z)$ must satisfy the PR condition $H_0(z)H_1(-z) - H_1(z)H_0(-z) = z^{-\ell}$. This is equivalent to $\boldsymbol{C}\boldsymbol{h}_1 = \boldsymbol{m}$, where $\boldsymbol{C}$ is a matrix whose elements come from $\boldsymbol{h}_0(n)$. The highpass filter is $\boldsymbol{h}_1$ and the k-th element of $\boldsymbol{m}$ is $\delta(k - l)$. Given $\boldsymbol{h}_0(n)$, the minimization problem becomes:

$$\text{Minimize} \quad \Phi \quad \text{subject to} \quad \boldsymbol{C}\boldsymbol{h}_1 = \boldsymbol{m}. \tag{10.14}$$

The minimization constructs a Lagrangian function

$$L(\boldsymbol{h}_1, \lambda) = \Phi - \lambda^T (\boldsymbol{C}\boldsymbol{h}_1 - \boldsymbol{m}). \tag{10.15}$$

The optimality conditions $\partial L / \partial \boldsymbol{h}_1 = 0$ and $\partial L / \partial \lambda = 0$ are

$$\frac{\partial \Phi}{\partial \boldsymbol{h}_1} - \boldsymbol{C}^T \lambda = \boldsymbol{0} \quad \text{and} \quad \boldsymbol{C}\boldsymbol{h}_1 - \boldsymbol{m} = \boldsymbol{0}. \tag{10.16}$$

In some cases Φ can be expressed explicitly in terms of $\boldsymbol{h}_1$. When stopband error is the objective function, Φ has the form $\frac{1}{2}\boldsymbol{h}_1^T \boldsymbol{S}_1 \boldsymbol{h}_1 + \boldsymbol{S}_2^T \boldsymbol{h}_1 + s_3$. The optimality conditions become

$$\begin{bmatrix} -\boldsymbol{S}_1 & \boldsymbol{C}^T \\ \boldsymbol{C} & \boldsymbol{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{h}_1 \\ \lambda \end{bmatrix} = \begin{bmatrix} \boldsymbol{S}_2 \\ \boldsymbol{m} \end{bmatrix}. \tag{10.17}$$

QCLS formulation for Type A: The filter length is $2m$ and the PR condition is

$$T(z) = \sum_{n=0}^{4m-2} t(n)z^{-n} = -H_0(z)H_1(-z) + H_0(-z)H_1(z) = z^{-(2m-1)}. \tag{10.18}$$

Substituting $(-z)$ for z in (10.18), one obtains $t(n) = 0$ for even n. Consequently we have m conditions on the odd-numbered coefficients:

$$t(2k+1) = \begin{cases} 0, & 0 \le k \le m-2 \\ 1, & k = m-1. \end{cases} \tag{10.19}$$

Let $\boldsymbol{h} = (\,\boldsymbol{h}_0(0) \;\; \cdots \;\; \boldsymbol{h}_0(m-1) \;\; \boldsymbol{h}_1(0) \;\; \cdots \;\; \boldsymbol{h}_1(m-1)\,)^T$. Our objective is to express the m conditions in (10.19) in terms of $\boldsymbol{h}$. The polynomials $H_0(z)$, $H_1(z)$, $H_0(-z)$ and $H_1(-z)$ can be written using $\boldsymbol{e}(z) = (1 \;\; z^{-1} \;\; \cdots \;\; z^{-(m-1)})^T$:

$$\begin{cases} H_0(z) = \boldsymbol{h}^T \begin{pmatrix} \boldsymbol{e}(z) + z^{-m}\boldsymbol{J}\boldsymbol{e}(z) \\ \boldsymbol{0} \end{pmatrix}, & H_0(-z) = \boldsymbol{h}^T \begin{pmatrix} \boldsymbol{U}\boldsymbol{e}(z) + (-1)^{m/2}z^{-m}\boldsymbol{J}\boldsymbol{U}\boldsymbol{e}(z) \\ \boldsymbol{0} \end{pmatrix} \\ H_1(z) = \boldsymbol{h}^T \begin{pmatrix} \boldsymbol{0} \\ \boldsymbol{e}(z) - z^{-m}\boldsymbol{J}\boldsymbol{e}(z) \end{pmatrix}, & H_1(-z) = \boldsymbol{h}^T \begin{pmatrix} \boldsymbol{0} \\ \boldsymbol{U}\boldsymbol{e}(z) - (-1)^{m/2}z^{-m}\boldsymbol{J}\boldsymbol{U}\boldsymbol{e}(z) \end{pmatrix} \end{cases}$$

Here $\boldsymbol{J}$ is the exchange matrix and $\boldsymbol{U}$ is diagonal with $\boldsymbol{U}_{kk} = (-1)^k$. Substituting into (10.18), $T(z)$ is simplified to

$$T(z) = \boldsymbol{h}^T \begin{pmatrix} \boldsymbol{0} & \Gamma(z) \\ \boldsymbol{0} & \boldsymbol{0} \end{pmatrix} \boldsymbol{h} \quad \text{where} \quad \boldsymbol{E}(z) = \boldsymbol{e}(z)\boldsymbol{e}^T(z) \quad \text{and} \tag{10.20}$$

$$\Gamma(z) = [\boldsymbol{U}\boldsymbol{E}(z) - \boldsymbol{E}(z)\boldsymbol{U}] + z^{-m}\Big[(-1)^{m/2}\big(\boldsymbol{J}\boldsymbol{U}\boldsymbol{E}(z) + \boldsymbol{E}(z)\boldsymbol{U}\boldsymbol{J}\big) - \big(\boldsymbol{U}\boldsymbol{E}(z)\boldsymbol{J} + \boldsymbol{J}\boldsymbol{E}(z)\boldsymbol{U}\big)\Big] + z^{-2m}(-1)^{m/2}\,[\boldsymbol{J}\boldsymbol{E}(z)\boldsymbol{U}\boldsymbol{J} - \boldsymbol{J}\boldsymbol{U}\boldsymbol{E}(z)\boldsymbol{J}].$$

Substituting (10.12) for $\boldsymbol{E}(z)$ produces a polynomial $\sum_{k=0}^{4m-2} z^{-k}\boldsymbol{h}^T \begin{pmatrix} \boldsymbol{0} & \Gamma_k \\ \boldsymbol{0} & \boldsymbol{0} \end{pmatrix} \boldsymbol{h}$, where the Γ_k are constant matrices $\boldsymbol{D}_k, \boldsymbol{J}$ and $\boldsymbol{U}$. Comparing term by term in (10.20), (10.19) becomes

$$\begin{cases} \boldsymbol{h}^T\boldsymbol{Q}_{2k+1}\boldsymbol{h} = 0; & 0 \le k \le m-2 \\ \boldsymbol{h}^T\boldsymbol{Q}_{2m-1}\boldsymbol{h} = 1; \end{cases} \quad \text{where} \quad \boldsymbol{Q}_n = \begin{pmatrix} \boldsymbol{0} & \Gamma_n \\ \boldsymbol{0} & \boldsymbol{0} \end{pmatrix} \quad \text{and} \tag{10.21}$$

$$\Gamma_n = \begin{cases} \boldsymbol{U}\boldsymbol{D}_n - \boldsymbol{D}_n\boldsymbol{U}; & 0 \le n \le m-1 \\ \boldsymbol{U}\boldsymbol{D}_n - \boldsymbol{D}_n\boldsymbol{U} + (-1)^{m/2}\,(\boldsymbol{J}\boldsymbol{U}\boldsymbol{D}_{n-m} + \boldsymbol{D}_{n-m}\boldsymbol{U}\boldsymbol{J}) - \big(\boldsymbol{U}\boldsymbol{D}_{n-m}\boldsymbol{J} + \boldsymbol{J}\boldsymbol{D}_{n-m}\boldsymbol{U}\big); & m \le n \le 2m-2 \\ (-1)^{m/2}\,(\boldsymbol{J}\boldsymbol{U}\boldsymbol{D}_{m-1} + \boldsymbol{D}_{m-1}\boldsymbol{U}\boldsymbol{J}) - \big(\boldsymbol{U}\boldsymbol{D}_{m-1}\boldsymbol{J} + \boldsymbol{J}\boldsymbol{D}_{m-1}\boldsymbol{U}\big); & n = 2m-1. \end{cases}$$

In summary, the PR condition in (10.19) is rewritten as m quadratic constraints on $\boldsymbol{h}$.

Example 10.6. The QCLS formulation is used in the design of a Type A bank with filter lengths 30. The objective function Φ is

$$\int_0^{\omega_p} [|H_0(e^{j0}) - H_0(e^{j\omega})|^2 + |H_1(e^{j\omega})|^2]\,d\omega + \int_{\omega_s}^{\pi} [|H_0(e^{j\omega})|^2 + |H_1(e^{j\pi}) - H_1(e^{j\omega})|^2]\,d\omega$$

Note that both passband and stopband errors are included since this is not a power complementary system. This is an NPR filter bank with $\epsilon = 1.5 \times 10^{-15}$. The figure below shows the frequency responses of the filters.

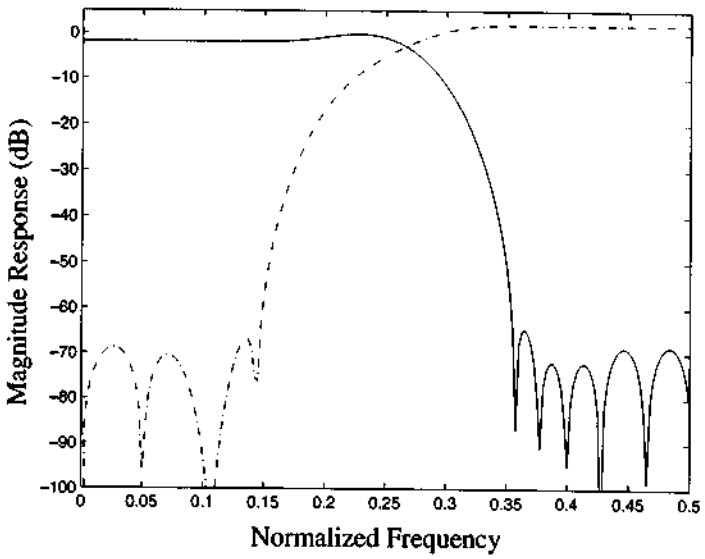

We close by noting four other design methods for two channels:

- Biorthogonal filter banks with integer coefficients divided by 2^k [AH].
- LMS-based algorithm [Roy].
- Biorthogonal linear-phase filter bank [Phoong1, Kiya].
- Alias-free two-channel allpass-based IIR filter bank [V, Vaid1, Nguyen4, EkPr].

10.5 Design of Cosine-modulated Filter Banks

A cosine-modulated filter bank is very efficient. It can be implemented using *one prototype filter* $H(z)$ and a cosine-modulation block. The M filters $H_k(z)$ are cosine-modulated versions of $H(z)$. As a result, a 32-channel bank with length 512 has only 256 unknowns (assuming linear phase), instead of 16,384 unknowns in the general case. This helps the design algorithm tremendously in terms of convergence speed and results.

There are many design algorithms and they all use $H(z)$. What properties does $H(z)$ satisfy to obtain NPR solution or PR solution?

- ***Nearly PR*** If $H(z)$ is a spectral factor of a $2M$-th band filter, amplitude and phase distortions are eliminated. Aliasing is suppressed by sufficient stopband attenuation.
- ***Exactly PR*** Let $H(z)$ be linear phase. The $2M$ polyphase components $G_k(z)$ of a paraunitary cosine-modulated filter bank have to satisfy

$$G_k(z^{-1})\, G_k(z) + G_{M+k}(z^{-1})\, G_{M+k}(z) = \frac{1}{2M}. \tag{10.22}$$

We begin with NPR methods that are based on spectral factorization. Lattice structure and time-domain approaches will also be discussed, for PR designs. We will see that NPR methods tend to give better stopband attenuation.

Traditional Pseudo-QMF Bank Design

Pseudo-QMF Banks are NPR and cosine-modulated. They were proposed for communication and audio processing applications. Aliasing of the adjacent bands is cancelled by appropriate phase terms in the modulation. The nonadjacent aliasing terms are minimized by designing a prototype with good stopband attenuation δ_s. The aliasing level is comparable to δ_s, and the reconstruction error consists of distortion (Φ_d) and aliasing (Φ_a):

$$\Phi = \Phi_d + \Phi_a = (1-\alpha)\int_0^{\pi} \left|T_0(e^{j\omega}) - e^{-j\ell\omega}\right|^2 d\omega + \alpha \int_{\omega_s}^{\pi} \left|H(e^{j\omega})\right|^2 d\omega.$$

The weighting factor has $0 \leq \alpha \leq 1$, and ℓ is the system delay. There is typically a tradeoff between distortion and aliasing. Distortion is eliminated if $H(z)$ is restricted to be a spectral factor of a $2M$-th band filter $P(z)$. The objective function becomes simply Φ_a.

Spectral Factorization Evaluating $H(z)H(z^{-1})$ on the unit circle $z = e^{j\omega}$ yields a nonnegative function $P(e^{j\omega}) = |H(e^{j\omega})|^2$. How does one design $P(z)$ where $P(e^{j\omega}) \geq 0$? The design procedure is as follows:

- Design an equiripple $2M$-th band filter with $\widehat{p}(N) = \frac{1}{2M}$ and $\widehat{p}(N - 2M\ell) = 0$ for $\ell \neq 0$. N is large since $\widehat{P}(z)$ must also have good stopband attenuation δ_s.
- Form another $2M$-th band filter $P(z)$ (nonnegative!) by $P(z) = \delta_s\, z^{-N} + \widehat{P}(z)$. Factor $P(z)$.

The lowpass filter $H_0(z)$ and the highpass filter $H_{M-1}(z)$ are most affected by the choice of factor (minimum phase or near linear phase). Their passbands are combinations of two passbands and they have bumps or dips [Koil3] if the phase of $H(z)$ is not linear.

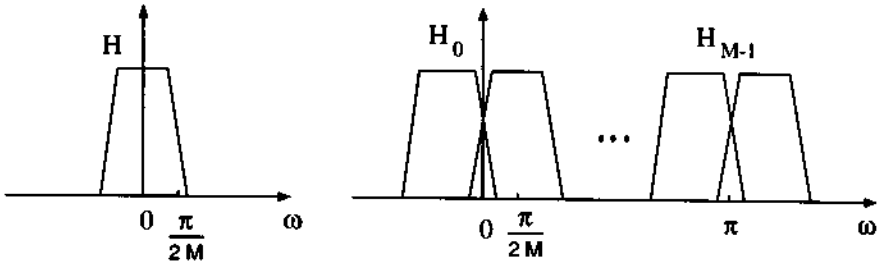

[Nguyen1] uses QCLS for linear-phase factorization. The resulting cosine-modulated filter bank has no distortion, small aliasing (comparable to stopband attenuation), and flat passband responses for H_0 and H_{M-1}. A 32-channel system of length 512 is shown, with distortion 5×10^{-11} and aliasing -83 dB.

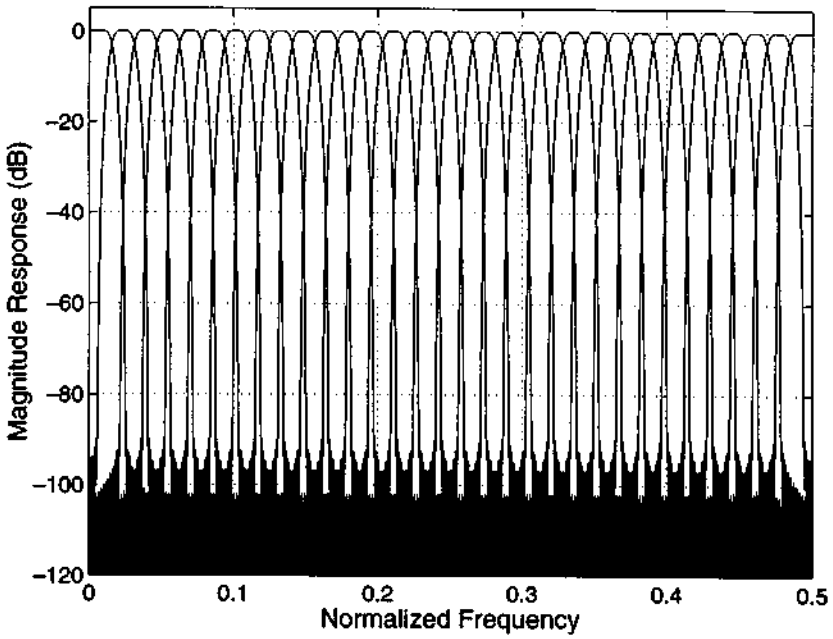

QCLS in NPR Pseudo-QMF Design

The linear phase relation $H(z) = z^{-N}H(z^{-1})$ expresses the coefficients $\boldsymbol{p}(n)$ as $\boldsymbol{h}^T\boldsymbol{D}(n)\boldsymbol{h}$. The matrix $\boldsymbol{D}(n)$ depends on the index n and the filter length [Nguyen1]. The design becomes a quadratic-constrained minimization:

Minimize the stopband error Φ_a subject to $\boldsymbol{p}(n) = \boldsymbol{h}^T\boldsymbol{D}(n)\boldsymbol{h} = 2M$-th band filter.

The quadratic forms have simple derivatives and Hessians, which makes the optimization search very fast. A high-order system design like the example above takes about 4 minutes on a Sun Workstation 10.

Lattice Structure Method

In a PR cosine-modulated filter bank of length $N = 2mM$, each polyphase component pair $(G_k(z), G_{M+k}(z))$ is paraunitary. The design procedure based on lattice structure is:

- Initialize the angles θ_k. Compute $H_k(z)$ and $\Phi = \int_{\omega_s}^{\pi} \left|H(e^{j\omega})\right|^2 d\omega$.
- Iterate on θ_k to minimize Φ.

This method works well for short filters. One needs a time domain approach to design long PR systems with high attenuation. We show the frequency responses of two eight-channel PR cosine-modulated filter banks with lengths 64 and 80. The stopband attenuation improves from -46 dB to -52 dB (no distortion or aliasing). The attenuation using the QCLS is -58 dB, a significant improvement from -52 dB.

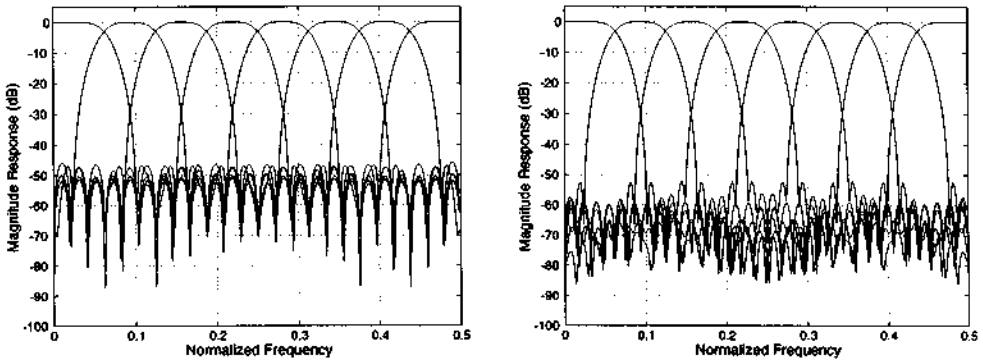

Time Domain Method in PR Design—QCLS Formulation

The objective is to design the linear phase prototype filter $H(z)$ such that its $2M$ polyphase components $G_k(z)$ satisfy the PR conditions (in pairs k and $M+k$). Let $\boldsymbol{h}$ be a vector of length mM consisting of the first half of $\boldsymbol{h}(n)$. Let $\boldsymbol{e}$ be a delay vector of length m with components z^{-k}. Then the polyphase components $G_k(z)$ can be written as $\boldsymbol{h}^T\boldsymbol{V}_k\boldsymbol{e}$, where $[\boldsymbol{V}_k]_{i,j} = 1$ only for $i = k + 2jM$. The PR condition (10.22) on the pair $G_k(z)$ and $G_{M+k}(z)$ becomes

$$\boldsymbol{h}^T\left[\boldsymbol{V}_k\boldsymbol{J}\boldsymbol{e}\boldsymbol{e}^T\boldsymbol{V}_k^T + \boldsymbol{V}_{M+k}\boldsymbol{J}\boldsymbol{e}\boldsymbol{e}^T\boldsymbol{V}_{M+k}^T\right]\boldsymbol{h} = \frac{1}{2M}z^{-(m-1)}. \tag{10.23}$$

Substituting $\boldsymbol{e}\boldsymbol{e}^T = \sum_{n=0}^{2m-2} z^{-n}\boldsymbol{D}_n$ where $[\boldsymbol{D}_n]_{i,j} = 1$ only for $i + j = n$, we obtain the

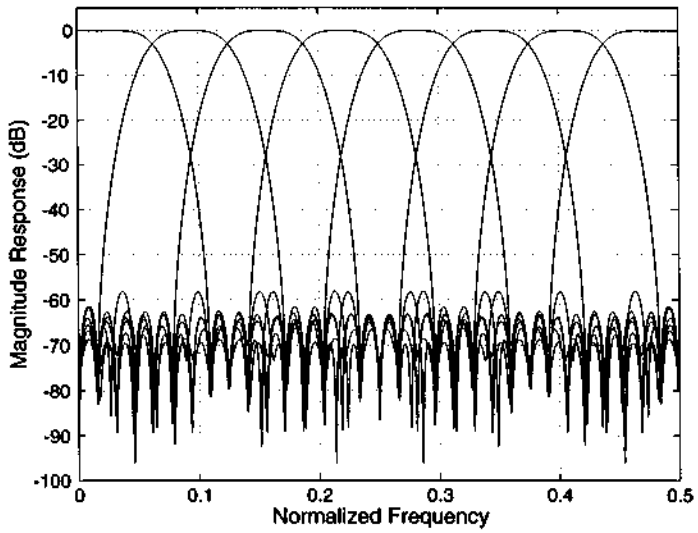

Example of an 8-channel cosine-modulated filter bank with length 80.

PR condition in terms of $\boldsymbol{h}(n)$:

$$\boldsymbol{h}^T\left[\boldsymbol{V}_k\boldsymbol{J}\boldsymbol{D}_n\boldsymbol{V}_k^T + \boldsymbol{V}_{M+k}\boldsymbol{J}\boldsymbol{D}_n\boldsymbol{V}_{M+k}^T\right]\boldsymbol{h} = \frac{1}{2M}\,\delta(n-m+1). \tag{10.24}$$

The design problem has these *quadratic constraints* [QCLS]. The objective function Φ is also quadratic in $\boldsymbol{h}$. As an example, we design an 8-channel cosine-modulated filter bank with length 80. The frequency responses of the analysis filters are shown above.

Final Design Example: *M* Channels with Linear Phase (GenLOT)

Linear phase filter banks are not cosine-modulated. For $M = 2$, they are constructed by factorization of the Daubechies polynomial or any symmetric halfband $P(z)$. They can also be designed by optimizing the lattice coefficients in the lattice structures [NgVaid1]. For $M > 2$, there is no simple spectral factorization method and one has to rely on lattice structure or time-domain approaches.

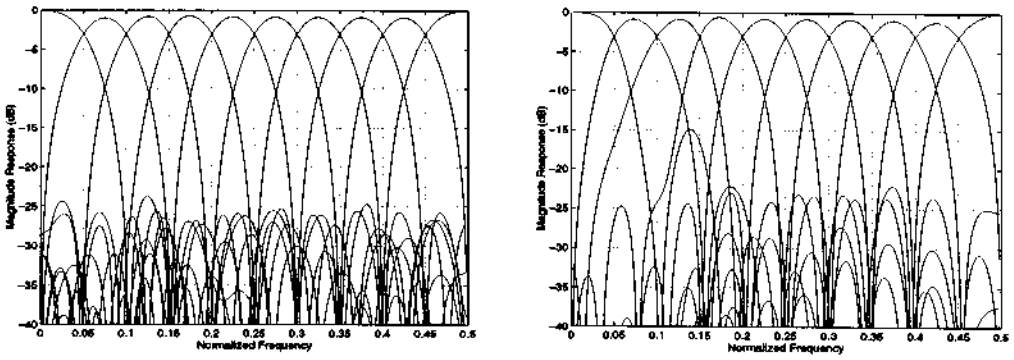

Frequency responses of a 10-channel linear-phase system with length 40.

Lattice structures exist for even M, when the length is a multiple of M (see Chapter 9 and [Soman, deQueiroz]). The figure above shows the frequency responses of a 10-channel linear-phase system with length 40. These are designed by optimizing the lattice coefficients for stopband attenuation (left) and coding gain (right).

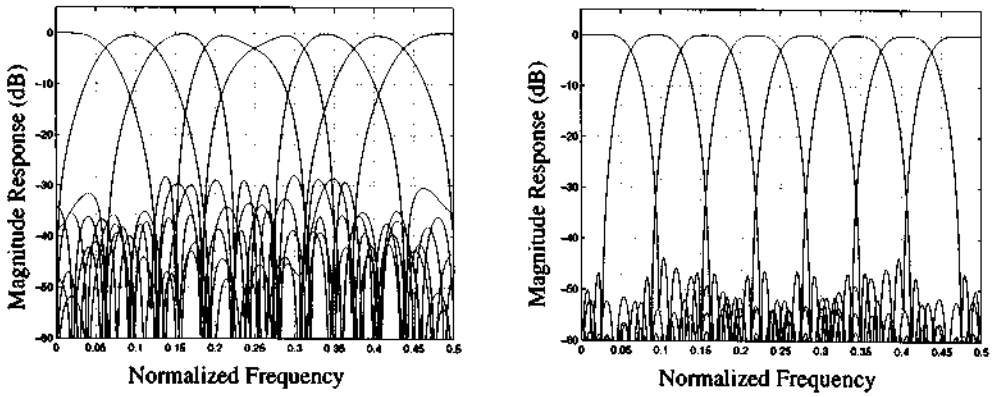

An NPR design (also linear phase) gave attenuation of 43.7dB.

One notices from those GenLOT designs that the stopband attenuation does not improve significantly as the filter length increases. One way to obtain high attenuation is to ***trade off the reconstruction error and stopband attenuation*** [Nguyen2]. An 8-channel PR design gave us attenuation of 28dB using the lattice structure. On the other hand, an NPR design (also linear phase) gave attenuation of 43.7dB in the last graph shown. The price was amplitude distortion of 0.0365 and aliasing of -47dB. The acceptability of a near PR design depends on the applications!