4. Practical QMF Filterbank Design

4.1 Overview

The ideal QMF filters ($H0(z)$ and $H1(z)$) in section 3.2.1 consist of perfect rectangular frequency domain windows and consequently have infinite latency, order and computational cost. This assumption has utility in exploring the general properties of subband hierarchy but for a practical implementation, major compromises are necessary. A primary limitation is that latency for real-time operation, as specified in section 1.2.2, has an upper bound of $T_{max} \cong 20$ms. Additionally, transferring the computational burden from oscillator bank to filterbank has the risk, in the extreme, of resulting in worse efficiency than classical AS (see section 3.1.7). Another impetus for reducing prototype filterbank complexity is to maximise the number of filterbank instantiations that may be multiplexed through the same hardware in order to provide a satisfactory number of independent output streams as discussed in section 1.2.1. A related topic, discussed in section 4.2 onwards, is in the modifications required for functional transparency.

The initial topic of this chapter is an investigation of efficient standard QMF filterbank designs in the context of their application to MAS. Both FIR and IIR forms are possible and each has attendant advantages and disadvantages. Practical QMF stage design is parameterised by the classical interpolator design variables from section 3.1.1. which are Δ_f (transition width) and δ_p and δ_s (passband and stopband ripple). It is appropriate to set out their tolerances prior to documenting their impact on QMF design. δ_s controls the relative power of quantisation noise transmitted during interpolation and, ultimately, interpolated oscillator SNR; for high-fidelity synthesis δ_s= -80dB which is comparable to that used in TOB's of recent years (Jansen, 1991). The constraint upon δ_p is quite subjective; an oscillator placed at a passband ripple maximum should have the same perceived power as at a minimum. However, given δ_s, varying δ_p has little effect on filterbank cost. Therefore the single design variable of interest is Δ_f .

4.2 FIR QMF Filterbank Design

4.2.1 Optimal Structure for a FIR QMF Stage

Finite Impulse Response (FIR) digital filter designs for QMF stages are well documented in the literature (Crochiere and Rabiner, 1993) and hence a brief tutorial will suffice. Invariably, symmetrical impulse responses are used which have a latency of precisely M samples where the filter length is $N=2M+1$. Applying an inverse Fourier transform to the ideal rectangular "brick wall" frequency response of $H0(z)$, which is unity within the bounds $\pm\pi/2$ and zero elsewhere creates the well-known time series for $h0[m]$ described by eqn. (4.1), known as a *sinc* function. $H1(z)$ is $H0(z)$ displaced in frequency by $f_s/4$ which has the time series $(-1)^m$. When multiplied by $h0[m]$, the desired impulse response for $h1[m]$ is obtained in eqn. (4.2). m is bounded such $M \le m \le M$ and M is odd so that coefficients at the extremes are non-zero to eliminate redundant latency. Example functions are plotted in Fig. 4.1 for $M=11$ ($N=23$): note that $h0[m]$ and $h1[m]$ are non-zero for all odd m. Truncation of the infinite time series gives rise to non-ideal filter behaviour which can be minimised by the techniques discussed in the next section.

$$h0[m] = 2\left(\frac{\sin(m\pi/2)}{m\pi}\right)$$
$$h1[m] = (-1)^m h0[m] \qquad (4.1, 4.2)$$

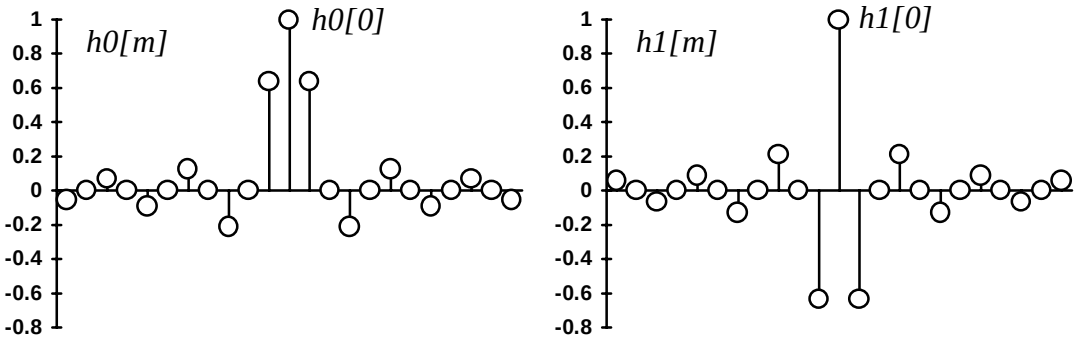

Figure 4.1 Unwindowed FIR Responses for h0[m] and h1[m] for M=11, N=23

Upsampling by $I=\uparrow 2$ is equivalent to inserting a zero-valued sample between each sample of $x0[n]$ and $x1[n]$ to create $x0[m]$ and $x1[m]$. This creates a modulo-2 activation pattern of the coefficients in $h0[m]$ and $h1[m]$. When a sample is shifted into the first stage of either filter, all samples at even values of m are zero and therefore only the coefficients at odd values of m are required in the computation of the output. Conversely, when a zero-valued sample is shifted into either filter, coefficients at even values of m are activated. Fortunately, the only non-zero coefficient for even m is $h0[0]=h1[0]=1$. The filters are therefore realised by applying an input sample to both a delay line of length $(M+1)/2$ and a FIR of length $M+1$ comprising the coefficients at odd values of m, and generating two output samples by commutating between, in sequence, the FIR and delay line: a simple form of a polyphase structure as discussed in section 3.1.2.

Taking advantage of FIR symmetry reduces computation by "folding" the FIR such that $(x[-m]+x[m])\times h0[m]$ replaces $x[-m]\times h0[-m]+x[m]\times h0[m]$ because $h0[-m]=h0[m]$. The only structural difference between a realisation of $h0[m]$ and $h1[m]$ is that the FIR coefficients are inverted in sign in the latter. Therefore $x0[n]+x1[n]$ is input to the delay line while $x0[n]-x1[n]$ is input to the FIR which has coefficients from $h0[m]$: the subtraction of $x1[n]$ is functionally equivalent to an inversion of coefficient signs. By this means, the signals $x0[n]$ and $x1[n]$ are upsampled by $\uparrow 2$, interpolated by their respective impulse responses $h0[m]$ and $h1[m]$ and summed in an efficient structure which minimises computation and storage as illustrated in Fig. 4.2. A mean of $(M+1)/4$ multiplies are required per $y[m]$.

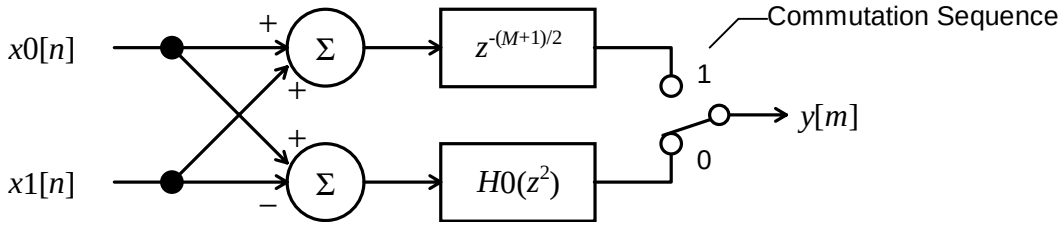

Figure 4.2 Optimal Realisation of FIR QMF Synthesis Stage

4.2.2 FIR Frequency Response Optimisation

A truncated *sinc* function for $H0(z)$ is equivalent to a rectangular window which has poor frequency domain performance. A solution is to employ the Kaiser window FIR design technique which minimises filter length by shaping the frequency response such that it is closely bounded by the tolerance margins Δ_f and $\delta=\delta_s=\delta_p$ from which M is determined by eqn. (4.3). The formulae for the time-domain window is based on a zeroth-order Bessel function and is in the literature (Crochiere and Rabiner, 1993). Note that the same ripple is specified for passband and stopband. Kaiser filters are sub-optimal because, in practice, the height of the first stopband sidelobe is closest to δ and the rest fall monotonically below this value, constituting excess stopband gain (see Fig. 4.3). However, it is simple to compute and hence a popular design technique where strict optimality is not required.

$$M = \frac{-20\log_{10}\delta - 7.95}{28.72\Delta_f} \tag{4.3}$$

An alternative to windowing is the Parks and McClellan (1972) “equiripple” design algorithm. This takes the more specific parameters of $\{f_{pass}, f_{stop}, \delta_p, \delta_s\}$and applies an iterative technique, the Remez exchange algorithm, to adjust the sidelobe amplitudes in the passband and stopband to conform to the equiripple condition (i.e. of equal height). For a given specification, it achieves a design with minimum M. Empirical formulae for M are given in eqns. (4.4, 4.5, 4.6) which are quoted from (Crochiere and Rabiner, 1993). It differs from the Kaiser technique in that the condition $\delta_p=\delta_s$ is not necessary. Specifying $\delta_p\neq \delta_s$ results in lower M (and therefore lower latency) but also generates non-zero coefficients at even-numbered values of m making it more expensive to compute than a windowed *sinc* design. However, setting $\delta_p= \delta_s$ and $f_{pass}=\pi-f_{stop}$ results in an optimally windowed *sinc* function with zero coefficients at even-numbered m, of less M than the Kaiser window. Fig. 4.3. provides a visual comparison of the performance of the two FIR design methods for a typical MAS specification as indicated by section 6.4.4.

$$D_\infty(\delta_p,\delta_s) = \log_{10}(\delta_s(0.005309(\log_{10}\delta_p)^2 + 0.07114\log_{10}\delta_p - 0.4761)) \\ +(-0.00266(\log_{10}\delta_p)^2 - 0.5941\log_{10}\delta_p - 0.4278) \tag{4.4}$$

$$f(\delta_p,\delta_s) = 11.012 + 0.512(\log_{10}\delta_p - \log_{10}\delta_s),\ \text{for}\ |\delta_s| \le |\delta_p| \tag{4.5}$$

$$M = \frac{D_\infty(\delta_p,\delta_s)}{2\Delta_f} - \frac{f(\delta_p,\delta_s)\Delta_f}{2} \tag{4.6}$$

Kaiser Window *M*=25, *N*=51

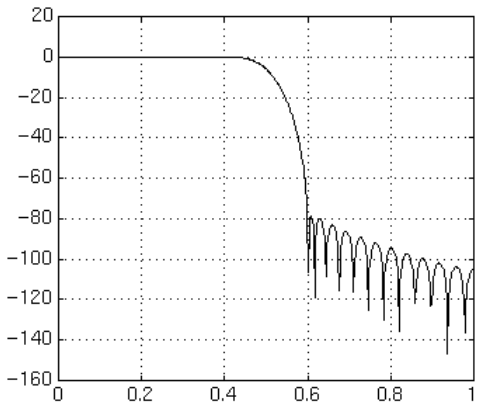

Gain (dB) vs Frequency (Nyquist=1.0)

Parks-McClellan *M*=23, *N*=47

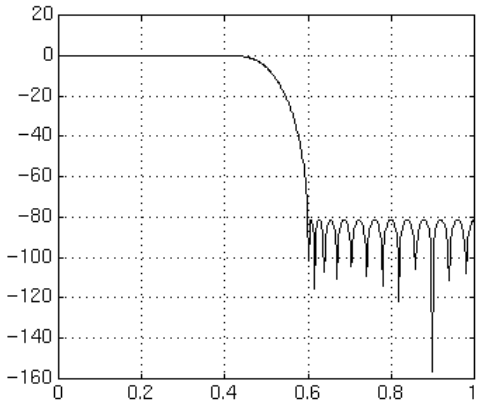

Gain (dB) vs Frequency (Nyquist=1.0)

Figure 4.3 FIR designs for H0(z) with Δ_f=0.1, δ_p=δ_s=-80dB

4.3 MAS Normalisation Techniques for QMF FIR Filterbanks

The objective of achieving functional transparency requires consideration. For MAS to integrate with existing spectral modelling techniques, it should appear to the high-level SME, and ultimately the musician, as a TOB (granted the need for *a priori* data). However, fragmentation of a note partial series across different subbands in a subband hierarchy corrupts (i) time, (ii) frequency and (iii) phase relationships between partials. Normalisation techniques which address each of these problems in an FIR QMF filterbank are presented in the following sub-sections, and form the basis of those used in the IIR QMF implementation of choice, discussed in section 4.4.

4.3.1 Latency Normalisation

$$T_{fb}(K) = \frac{M}{fs}\sum_{i=1}^{K-1} 2^{i-1} = \frac{M(2^{K-1}-1)}{fs} \tag{4.7}$$

Assume the case of a note partial series allocated across a constant-Q series of subbands as illustrated in section 3.2.3. High frequency partials in the fullband (k=0) are not subject to any filterbank delay. Lower frequency partials are allocated deeper in the hierarchy and for allocation level k are subject to a latency of $T_{fb}(k)$ according to eqn. (4.7) which exponentiates with k. Therefore at note onset, high frequencies emerge before low frequencies. Though this phenomena is a reflection of some natural processes, it is desirable to align envelope features between partials for accurate resynthesis: feature relationships in PWL envelope data encapsulate the spectral characteristics of the analysed source tone and should remain undistorted by the filterbank. A simple solution is to include delay lines for inputs to subbands $k<K$ such that latency from any level in the filterbank is normalised to the maximum of $T_{fb}(K)$. This is possible because FIR QMF stages exhibit a frequency-independent group delay of M samples. The length of delay line $d(K,k)$ required for level k is expressed by eqn. (4.8). The deepest subbands at level K do not require delay-lines.

$$d(K,k) = M(2^{K-k}-1) \text{ for } k < K \tag{4.8}$$

4.3.2 Frequency Normalisation

Once an oscillator x is allocated to the optimal subband $s_{k,l}$, its PWL frequency envelope $F_x[m]$ - from eqn. (1.1) - requires translation into the context of $s_{k,l}$ such that the desired frequency appears at the filterbank output. This is achieved via the linear relationship of eqn. (4.9) which generates the required digital frequency envelope $\Omega_x[m]$. If $s_{k,l}$ is inverted in frequency (by interpolation via an odd number of H1's, see section 3.2.2), $\Omega_x[m]$ must be pre-inverted to cancel the effect. The computational overhead in a PWL context is trivial as only normalisation of breakpoints is necessary, and the divisor for

level k is 2^{-k} when frequencies are prenormalised to the k=0 fullband as occurs in the TOB (Jansen 1991).

$$\Omega_x[m] = \begin{Bmatrix} \pi a, \text{ non-inverted sub-band} \\ \pi(1-a), \text{ inverted sub-band} \end{Bmatrix} \text{ where } a = \left(\frac{F_x[m] - f_{min}(s_{k,l})}{f_{max}(s_{k,l}) - f_{min}(s_{k,l})} \right) \qquad (4.9)$$

In a MAS paradigm of depth K=3, $\Omega_x[m]$ is decimated by a factor of $D \in \{2, 4, 8\}$ or undecimated for the fullband case of k=0. The decimated frequency envelope is denoted by $\Omega_x[n]$ where $m=Dn$. Now consider the allocation of x somewhere in the fully-overlapping octave spaced subband series $s_{k,1.}$ It is desirable that the corresponding decimated phase accumulator $\Phi_x[n]$ should track the parabolic trajectory of its undecimated equivalent $\Phi_x[m]$ so that errors introduced by decimated phase accumulation are eliminated. This is achieved by a trapezoidal integration over the decimation interval of the 'missing' $\Omega_x[m]$ envelope samples as expressed in eqn. (4.10): $\phi'_x[n]$ is the phase-normalisation term described in the next section. The 2^{-k} term is necessary to denormalise $\Omega_x[n]$ back to the context of $\Omega_x[m]$ in which the integration takes place. However, since $D=2^k$, the multiplier $2^{-k}D$ resolves to unity. For clarity, Fig. 4.4 illustrates an example for D=8 (k=3) in which the missing $\Omega_x[m]$ samples are assumed to lie on a line segment linking $\{\Omega_x[n], \Omega_x[n+1]\}$.

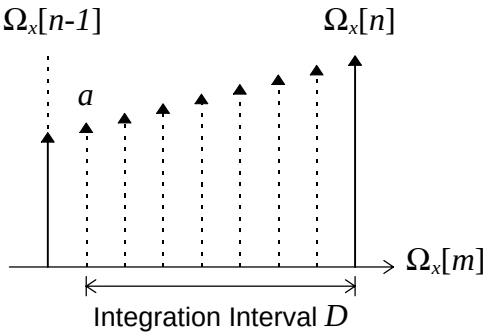

Figure 4.4 Trapezoidal Integration of Decimated Frequency Envelope

$$\Phi_x[n] = \Phi_x[n-1] + 2^{-k} D\left(\frac{a + \Omega_x[n]}{2}\right) + \phi'_x[n]$$
$$\text{where } a = \Omega_x[n] + \frac{\Omega_x[n] - \Omega_x[n-1]}{D} \tag{4.10}$$

4.3.3 Phase Normalisation

At the output of the filterbank, it is desirable that each sinusoid should have correct relative phase: a property usefully termed as “phase transparency”. To analyse the phase characteristics of the QMF FIR synthesis filterbank, consider the properties of a single stage. A stationary sinusoid is applied to each input, both lowpass ($x0[n]$) and highpass ($x1[n]$) according to eqns. (4.11, 4.12). At n=0, $x0[n]$ and $x1[n]$ have instantaneous phase ϕ_0 and ϕ_1 respectively. After a delay of M samples at the output, the interpolated baseband of $x0[n]$ preserves its phase of ϕ_0 at $y[m=M]$ as expressed in eqn. (4.13). The only transformation is a halving of relative digital frequency Ω_0 independent of ϕ_0.

$$\begin{aligned} x0[n] &= A_0 \sin(\Omega_0 n + \phi_0) \\ x1[n] &= A_1 \sin(\Omega_1 n + \phi_1) \end{aligned} \tag{4.11, 4.12}$$

$$y[m - M] = A_0 \sin(\frac{\Omega_0}{2} m + \phi_0) + A_1(\sin\frac{2\pi - \Omega_1}{2} m + \frac{2\pi - \Omega_1}{\Omega_1}\phi_1) \tag{4.13}$$

However, in the context of the output sample rate, highpass filter $H1$ suppresses the baseband signal ($\Omega_1/2$) of $x1[n]$ and passes the sideband at π–($\Omega_1/2$). As phase is the integral of frequency, the phase of the $x1[n]$ sideband at $y[m=M]$ (the product of the baseband phase ϕ_1 and the frequency ratio of sideband to baseband) is distorted from its desired value ϕ_1 and is, furthermore, a function of oscillator frequency Ω_1. However, substituting a frequency-dependent phase factor in eqn. (4.12) to produce eqn. (4.14) cancels out this phase distorting term in eqn. (4.13) such that the interpolated sideband of $x1[n]$ has the desired phase ϕ_1 at $y[m=M]$.

$$x1[n] = A_1 \sin(\Omega_1 n + \frac{\Omega_1}{2\pi - \Omega_1}\phi_1) \tag{4.14}$$

$$\phi'_x[n] = c[n]\phi_x[n] \text{ where } c[n] = \frac{\text{base - band freq}}{\text{side - band freq}} = \frac{f_s 2^{-k} \Omega_x[n]}{2\pi F_x[m]} \quad (4.15)$$

where $m = nD + d(K,0)$ via eqn.(4.8) and $D = 2^k$ as in section 4.3.2

This principle may be extended to a QMF filterbank (e.g. the complete binary tree of Fig. 3.7) which includes delay lines for latency normalisation. According to eqn. (4.8), latency is a constant $d(K,0)$ samples irrespective of allocation depth and therefore phase relationships are preserved during interpolation. Phase normalisation can therefore be interpreted as the adjustment of baseband phase to give the required sideband phase. An oscillator x allocated to subband $s_{k,l}$ has actual baseband frequency $f_s 2^{-k}(\Omega_x[n]/2\pi)$ via eqn. (4.9) which, after suppression of unwanted side-bands in the filterbank, corresponds to the desired frequency of sideband $F_x[m]$ at the output. To ensure $F_x[m]$ has the desired phase $\phi_x[m]$ requires substitution of the normalising eqn. (4.15) in eqn. (4.10).

For a stationary sinusoid, eqn (4.15) need only be computed once. However, for sinusoids in AS with time-varying frequency, eqn. (4.15) must be computed each sample period to maintain precise phase relationships. As the computation involves a division, this constitutes a large potential overhead. It is redundant, however, in the case of the subband series $s_{k,l}$:{k=0..K,l=1} which comprises a cascade of $H0$ lowpass filters where the phase normalising term in eqn. (4.5) is unity: the baseband image itself is passed by the filterbank. This cascade is functionally equivalent to the fully-overlapping octave spaced subband decomposition discussed in section 3.1.6 and is an advantageous property of the FIR QMF filterbank (see section 4.5).

4.4 IIR QMF Design

Restricting the subject of Infinite Impulse Response (IIR) filter theory to that of QMF stage design singles out the family of *polyphase all-pass* (PA-IIR) filters that have low computational cost, and narrow tolerances for Δf, δ_s and δ_p. The building block of a PA-IIR QMF is a single stage allpass filter (Oppenheim and Schafer, 1989), which has a unity gain, but non-linear phase response governed by a single coefficient α (the z-transform consists of a pole at $-1/\alpha$ and a zero at $-\alpha$) with a difference equation of

$y[n]=\alpha(x[n]-y[n-2])+x[n-2]$. Prime advantages are that it is cheap to compute, unconditionally stable, and because it has unity gain, may exploit the full dynamic range of fixed-point arithmetic (in contrast to other IIR building-blocks such as 2nd order sections). The phase response of a cascade of n all-pass filters with response $G(z)$ is the sum of the individual phases dictated by $\alpha_i:\{1\leq i\leq n\}$as expressed in eqn. (4.16).

$$\angle G(z)=\sum_{i=1}^{n}\angle\left(\frac{\alpha_i+z^{-1}}{1+\alpha_i z^{-1}}\right) \tag{4.16}$$

Fig. 4.5. illustrates the phase response of a state-of-the-art high-performance IIR polyphase allpass filter as published by (Hart et al, 1993); an example which is most suitable for closer investigation. Two parallel all-pass cascades are used; $G0$ and $G1$ ($G1$ is in series with a unit delay z^{-1}). $G0$ has length n=6 with $\alpha_i\in$ {0.984964, 0.798278, 0.914901, 0.593341, 0.297311, 0.040409} and $G1$ has length n=5 with $\alpha_i\in$ {0.865132, 0.953132, 0.708912, 0.452729, 0.149350}. Coefficients are derived by an iterative design procedure. A signal is input into both cascades and at the output it can be seen that the relative phase difference (bold line) is zero for $0\leq z\leq\pi/2$ (i.e. $G0$ and $G1$ in-phase) and $-\pi$ radians for $\pi/2\leq z\leq\pi$ (i.e. $G0$ and $G1$ in antiphase) with a narrow transition width. Clearly, summing the output creates a half-band filter with gain $\times 2$ for $0\leq z\leq\pi/2$ and $\times\cong 0$ for $\pi/2\leq z\leq\pi$, equivalent to the $H0$ filter in a QMF stage.

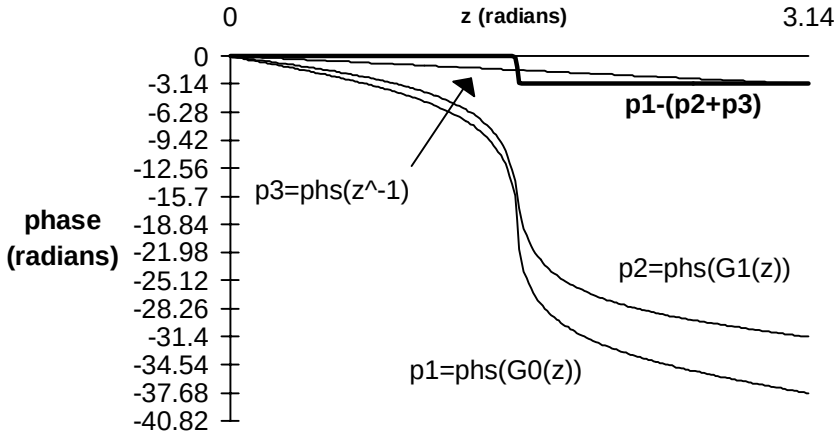

Figure 4.5 Phase Response of PA-IIR Filter

To translate $G0$, $G1$ into an equivalent QMF synthesis stage is straightforward. However for a polyphase implementation equivalent to the FIR QMF of Fig. 4.2, filters with the response $G0(z^2)$ and $G1(z^2)$ are required. This is achieved by decimating the difference equation by $\downarrow 2$ to $y[n]=\alpha(x[n]-y[n-1])+x[n-1]$ which permits filter computation at the lower sample rate with a net cost of 5.5 multiplies per $y[m]$. The commutation sequence implements the required z^{-1} series delay for $G1$. Fig. 4.6 illustrates the resulting structure. Half-band interpolation of the low-pass signal $x0[n]$ is performed as previously described giving the standard QMF response $H0(z)$. Interpolation of the high-pass signal $x1[n]$ by QMF response $H1(z)$ is realised by sending an inverted copy of $x1[n]$ to $G1$. This causes antiphase cancellation of the outputs of $G0$, $G1$ over $0 \le z \le \pi/2$ and in-phase superposition over $\pi/2 \le z \le \pi$ and therefore realises $H1(z)$. Gain ($\times 2$) from superposition is normalised to unity at the output. The actual normalised frequency response is illustrated in Fig. 4.7: note that δ_s =-100dB which is an improvement of 20dB over the specification in section 4.1 of δ_s=-80dB.

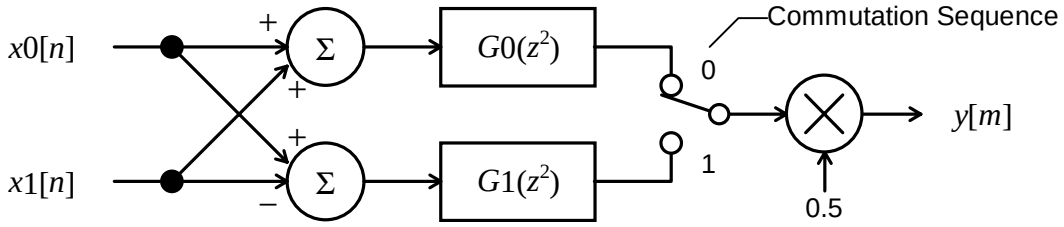

Figure 4.6 IIR Polyphase Allpass QMF Synthesis Stage

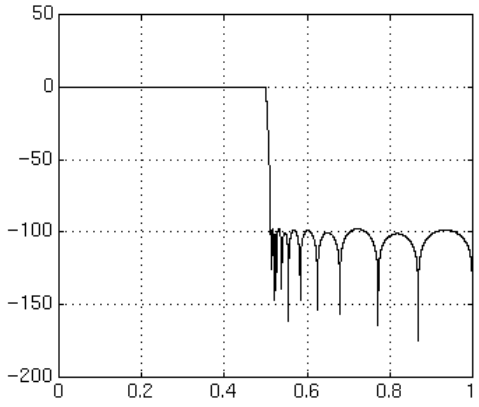

Gain (dB) vs. Frequency (Nyquist==1.0)

Figure 4.7 PA-IIR H0(z) Response

4.4.1 Normalisation with Non-linear Phase

Frequency normalisation in a QMF filterbank is independent of its phase response and hence the method in section 4.3.2 also applies to the PA-IIR design. However, the group delay of the PA-IIR stage is frequency-dependent in contrast to the fixed delay of M samples in a symmetric FIR filter. This means that phase relationships are distorted and envelope features are skewed by a PA-IIR filterbank during interpolation. To normalise phase therefore, requires the inclusion of the filterbank transfer function in eqn. (4.15). Latency normalisation by fixed-length delay lines (as proposed for the FIR QMF in section 4.3.1) is not possible as they function in the context of linear phase alone.

An interesting software alternative is to compute the latency T_x for a high-Q partial x at frequency $f(x)$ from the transfer function and "pre-skew" the $A_x[n]$, $F_x[n]$ envelopes by $T_{fb}(K)$-T_x, where $T_{fb}(K)$ is the PA-IIR filterbank worst-case latency. In a PWL envelope context, this may be done by attaching horizontal segments of length T_x to the start of the $A_x[n]$, $F_x[n]$ envelopes at note onset where T_x is computed from an LUT. However, the actual delay of a PA-IIR QMF stage is small, relative to the example PM-FIR QMF's discussed in section 4.2.2. As illustrated by Fig. 4.5 the greatest delay is -5π radians at z=π/2: that is about 10 samples.

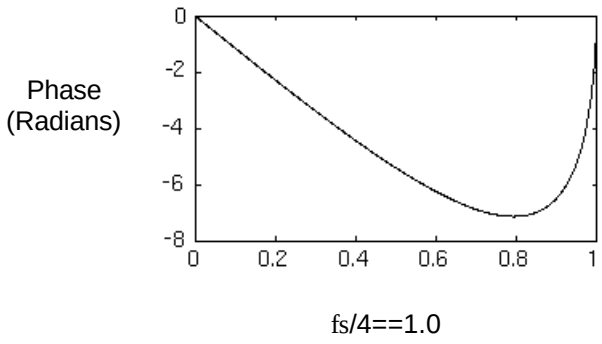

Figure 4.8 Required Passband Phase Response of P(z)

A structural solution to both of these problems is to incorporate an allpass phase-equaliser $P(z)$ ($|P(z)|$=1 for $\forall z$) into the prototype PA-IIR QMF stage output such that its net phase response from input to output is approximately linear. It is then it is functionally interchangeable with an FIR QMF stage and the algorithms of sections 4.2.1

and 4.2.2 are applicable. Modularity of the QMF filterbank structure is retained simplifying implementation. The alternative is to employ a single equaliser at the output, on the assumption that a single transfer equation characterises the entire filterbank.

Unfortunately, use of subband hierarchy in MAS means that at any particular value of absolute frequency, a signal can be sourced from one of a number of overlapping subbands, each with a unique transfer function and hence a single equaliser cannot be used. However, a problem with $P(z)$ is illustrated in Fig. 4.8 in that it has a phase response with a local minimum in the passband which cannot be realised by an allpass IIR filter which, by definition, has maximum phase. It is concluded that phase-linearity, and hence phase normalisation, is impractical for this PA-IIR QMF filterbank but that latency normalisation by envelope pre-skewing appears quite feasible and efficient.

4.5 Comparing PM-FIR and PA-IIR Performance for MAS

QMF Stage:	**Generic FIR**	**PM-FIR from Fig. 3.5 (*M*=23)**	**PA-IIR**
Storage (samples)	(M+1)	24	12
Mults per *y*[*m*]	(M+1)/4	6	5.5
Adds per *y*[*m*]	1+(M+1)/4	7	12

Table 4-1 Relative Complexity of FIR and PA-IIR QMF Stages

A comparison of the transition width characteristics between Figs. 4.3 and 4.7 demonstrates that the PA-IIR approach has superior performance: a feature that is a critical determinant in the choice of QMF filterbank for MAS as shall become clear in the next section. Additionally, δ_s=-100dB rather than δ_s=-80dB from section 4.1. The PM-FIR and PA-IIR QMF designs share the same general polyphase model (respectively Figs. 4.2 and 4.6) and differ only in the structure of their component phase shifters which may both function optimally in fixed-point arithmetic. Respective storage and computational requirements are tabulated in Table 4-1. Complexity of the PA-IIR stage

is thus comparable to the PM-FIR examples of Fig. 4.3. (which are good designs for typical MAS applications, see section 6.4.4)

If a non-linear phase response is tolerated, a PA-IIR implementation is most efficient in the light of its Δ_f performance. As discussed in section 4.3.3, phase normalisation in the PM-FIR QMF is complex to implement, but the subband series $s_{k,l}$:{k=0..K,l=1} offers phase transparency by default for subbands where the majority of low-frequency oscillators will be allocated in MAS. Sensitivity of the ear is most critical up to around 5kHz and diminishes with increasing frequency (Parsons, 1987). Phase has particular relevance in the context of stereo and sound spatialisation, but such a criterion does not apply so strongly to MAS as the filterbank is accumulating monoaural sources for subsequent spatialisation by an unspecified post-processor (see section 8.3.1). In the general case, altering the relative partial phases in a musical tone has minimal impact upon the perception of its timbre (Plomp, 1976), but exceptions to the rule can be contrived.

4.6 On Non-infinitesimal Δ_f in a Subband Hierarchy

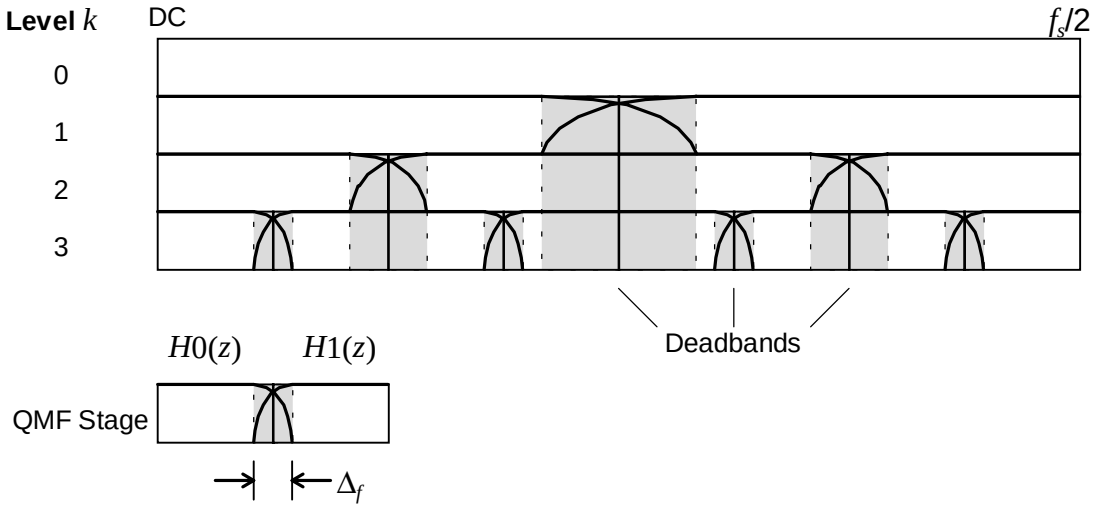

Figure 4.9 Deadbands in Spectral Hierarchy due to Non-infinitesimal Δ_f

The normalisation schemes outlined in this chapter are, as yet, insuffient to produce a functional MAS paradigm. A critical side-effect, as Fig. 4.9 illustrates, is that gradual passband rolloff appears in the transition widths of a practical QMF stage (due to non-

infinitesimal Δ_f) which causes attenuation towards the high frequency end of $H0$ and the low frequency end of $H1$. As discussed in section 3.1.1, an oscillator with time-varying frequency must not enter this region because it will suffer fading as a function of frequency, giving rise to a parasitic amplitude modulation effect that will affect the quality of perceived tone. Also, a sideband will also appear in the transition region, mirrored in $z=\pi/2$, which is boosted above the inaudible level of stop-band gain (δ_s) with the danger of becoming perceptible. Fortunately, this eventuality is prohibited by marking the transition widths as "deadbands", and restricting oscillator allocation to filter passbands which is permissible because time-invariant filterbanks are presumed with constant Δ_f.

The effect of deadbands is not isolated to individual QMF stages; sibling stages inherit the deadbands of their parents, falling under their "shadow" in absolute frequency terms. This is because an oscillator may be correctly interpolated through several passbands before entering the deadband of a stage further up the hierarchy due to cascaded, multistage interpolation. These deadband effects are eliminated by extending the allocation policy of section 3.2.3. Oscillators that are predicted according to their a priori parameters $\{f_{min}(x), f_{max}(x)\}$ to be in danger of falling into a deadband are promoted to the subband that lies immediately above the deadband thus logically excluding transition widths from the MAS paradigm: a formal algorithm is proposed in Chapter 5.

4.7 Review

The PM-FIR and PA-IIR designs outlines offer two efficient alternatives to QMF implementation. The former has poorer Δ_f performance but is linear-phase. The practical problems of achieving functional transparency are identified. Frequency normalisation requires trivial pre-processing of frequency envelope breakpoints after subband allocation. Latency normalisation by delay lines is practical if the QMF is linear-phase, but invalid for the phase non-linear PA-IIR, though envelope pre-skewing is a feasible solution. The most difficult issue is phase normalisation; intractable in a PA-FIR and only valid in a PM-FIR for the series $s_{k,l}$:$\{l=1\}$. Finally, non-ideal filter behaviour manifests

itself as an impinging of a non-infinitesimal Δ_f upon the idealised MAS subband hierarchy which creates deadbands where oscillators cannot be allocated.