Non-Uniform Filterbank Design Using The Running Fourier Transform

M S Murphy * & F J Owens **

* Marconi Underwater Systems
Watford
UK

** Dept of Electrical & Electronic Eng
University of Ulster
UK

RESUME

Cet article décrit une méthode pour construire une collection de filtres uniformes et non-uniformes, qui est fondée sur la transformée de Fourier courante (TFC). La TFC est accomplie en convoluant le signal avec l'un des membres d'une famille de fonctions fenêtres $h(nT) = (nT)^k e^{-\alpha nT}$ avec k entier. Les valeurs de k et de α peuvent être choisies pour spécifier respectivement l'ordre et la largeur de bande de chaque filtre d'analyse. La TFC est meilleure que la méthode plus traditionnelle de la transformée de Fourier à temps discret (TFD), parce que des intervalles de filtre inégales à une largeur de bande d'analyse inconstante sont permis et parce qu'on peut contrôler la fuite spectrale en travers des canaux.

En traitant l'opération de la TFC comme des filtres bande-pass, on a calculé une équation de la réponse d'impulsion composée. On peut employer cette réponse pour rendre optimale la réponse d'amplitude composée de la TFC. On peut construire un groupe de filtres à l'intervalle et à la largeur de bande non-uniformes, en divisant la bande de fréquence en un nombre de sections uniformes et en rendant optimale la réponse d'impulsion composée de chaque section séparément. Enfin on a presenté une méthode cepstrale modifiée, qui rend lisse un spectre non-uniforme. On a illustré aussi que cette technique est meilleure que la méthode de filtre traditionnelle.

SUMMARY

This paper describes a technique for designing uniform and non-uniform filterbanks based on the Running Fourier Transform (RFT). The RFT is implemented by convolving the input signal with one of a family of windows, $h(nT) = (nT)^k e^{-\alpha nT}$, where k and α may be chosen to specify the order and bandwidth, respectively, of each analysing filter. The RFT is superior to the more traditional Discrete Fourier Transform (DFT) in that non-uniform channel spacing with variable analysing bandwidth is permissable and, also, spectral leakage across channels can be controlled.

By viewing the operation of the RFT filterbank in terms of bandpass filtering, an expression for the equivalent composite impulse response has been derived. This response can be used to optimise the composite amplitude response of the RFT filterbank. Filterbanks with non-uniform channel spacing and bandwidth can be designed by splitting the frequency band of interest into a number of uniform sections and then optimising the equivalent composite impulse response of each section independently. Finally a modified cepstral smoothing technique for non-uniform spectra is presented and shown to be superior to conventional bi-pass filtering.

1. Introduction

The Running Fourier Transform [Flanagan, 1972] is a method for deriving the spectral content, $F(\omega,nT)$, of a quasi-stationary signal at any instant in time. For a discrete-time signal, it is expressed as

$$F(\omega,nT) = \sum_{r=-\infty}^{n} f(rT).h(nT-rT)e^{-j\omega rT} = a(\omega,nT) - jb(\omega,nT), \quad (1)$$

where

$$a(\omega,nT) = [f(nT).\cos(\omega nT)]*h(nT),$$
$$b(\omega,nT) = [f(nT).\sin(\omega nT)]*h(nT), \quad (2)$$

$h(nT)$ is the impulse response of a lowpass filter and * denotes discrete convolution. By definition, the magnitude spectrum, $|F(\omega,nT)|$, is given by

$$|F(\omega,nT)| = [\overline{F}(\omega,nT).F(\omega,nT)]^{1/2} = [a^2(\omega,nT) + b^2(\omega,nT)]^{1/2} \quad (3)$$

where $\overline{F}(\omega,nT)$ denotes the complex conjugate of $F(\omega,nT)$. Fig 1 shows a system, using real operations only, for carrying out the spectral measurement indicated by equation (1).

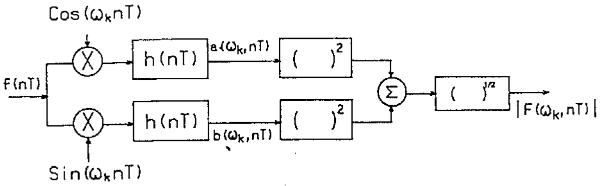

Fig 1 : Implementation of Runnning Fourier Transform

From a frequency-domain perspective, Fig 1 shows that the input signal is heterodyned to base-band and then filtered by a lowpass filter with impulse response h(nT). The result is effectively a bandpass filter with centre frequency ω and a passband extending from $(\omega - \omega_c)$ to $(\omega + \omega_c)$, where ω_c is the radian bandwidth of the lowpass filter. By carrying out the measurement indicated by Fig 1 at a number of different analysing frequencies ω_k, a running spectrum can be obtained. The advantage of this method of spectral derivation over, for example, the Discrete Fourier Transform (DFT) is that analysing bandwidth and spectral resolution are independently variable and therefore spectral leakage across channels can be controlled.

An appropriate family of lowpass filter impulse responses, h(nT), for use with the Running Fourier Transform are defined by [Owens, 1988]

$$h(nT) = (nT)^k . e^{-\alpha nT}, \text{ k integer.} \quad (4)$$

The above family of lowpass functions have order k+1 and a 3dB radian cut-off frequency ω_c given by

$$\omega_c = \alpha.(\ 2^{1/(k+1)} - 1\ )^{1/2}. \quad (5)$$

Using the impulse-invariant transformation, the z-domain transfer function, H(z), of this family of filters is

$$H(z) = T^k [\sum_{n=1}^{k} a_n (e^{-\alpha T} z^{-1})^n] / [1 + \sum_{m=1}^{k+1} b_m (e^{-\alpha T} z^{-1})^m]. \quad (6)$$

A wide variety of filterbank configurations can therefore be realised by appropriate choice of ω, α and k.

2. Non-Uniform Filterbank Design

A desirable feature in filterbank design is that the composite frequency response of the bank should be flat with linear phase. This is not too difficult to achieve if uniform channel spacing and constant bandwidth suffice. However if non-uniform spectral resolution and definition are required then the problem of obtaining a flat composite respnse becomes extremely difficult. This paper will consider composite amplitude response only since phase response is often of secondary importance in spectral analysis.

The composite amplitude response of the filterbank can be computed by summing the response of each channel to a swept-frequency input signal. This can be efficiently computed from an expression for the frequency response of each channel. Replacing the dummy summation variable r in equation (1) by n-r gives

$$F(\omega,nT) = \sum_{r=0}^{\infty} f(nT-rT).h(rT).e^{-j\omega(n-r)T} \quad (7)$$

$$= e^{-j\omega nT}[a_1(\omega,nT) + jb_1(\omega,nT)]$$

where

$$a_1(\omega,nT) = f(nT)*h(nT)\cos(\omega,nT)$$
$$b_1(\omega,nT) = f(nT)*h(nT)\sin(\omega,nT) \quad (8)$$

The magnitude spectrum, $|F(\omega,nT)|$, is now given by

$$|F(\omega,nT)| = [a_1^2 + b_1^2]^{1/2} \quad (9)$$

The above equations show that the RFT may be alternatively viewed as filtering the input signal with bandpass filters with impulse responses $h(nT)\cos(\omega nT)$ and $h(nT)\sin(\omega nT)$, squaring and adding the magnitude of each filter output and then taking the square-root. By deriving the z-domain transfer function of each bandpass filter using the impulse invariant transformation, squaring and adding the magnitude responses and finally taking the square-root, an expression for the overall amplitude response of each channel can be derived.

Consider first the case of uniform channel spacing. As previously indicated, the filtering action of the system in Fig 1 can be viewed in terms of bandpass filtering, though its impulse response is the window function h(nT). A true bandpass filter would have an impulse response $h(nT).\cos(\omega_k nT)$. Assuming no channel centred on dc, the composite impulse response, p(nT), of an M-channel bank of bandpass filters is

$$p(nT) = h(nT).d(nT), \quad d(nT) = \sum_{k=1}^{M} \cos(\omega_k nT) \quad (10)$$

If the channel spacing is fixed, then the function d(nT) depends only on the number of channels M and the channel spacing $\delta\omega$ and can be expressed in closed form as

$$d(nT) = [\sin(N\delta\omega nT/2)]/[\sin(\delta\omega nT/2)] \quad (11)$$

For an odd number of channels covering the double-sided frequency band of interest, N = 2M + 1 and for an even number, N = 2M. In general, d(nT) is a sequence of pulses, of maximum amplitude N, occuring at intervals of $2\pi/\delta\omega$. Such a sequence is shown in Fig 2 for M = 100, $\delta\omega$ = 40Hz, f_s = 10kHz and a frequency range of dc to 4020 Hz. When the double-sided band of interest covers the range $-\pi/T$ to $+\pi/T$, then d(nT) becomes an impulse train of amplitude N.

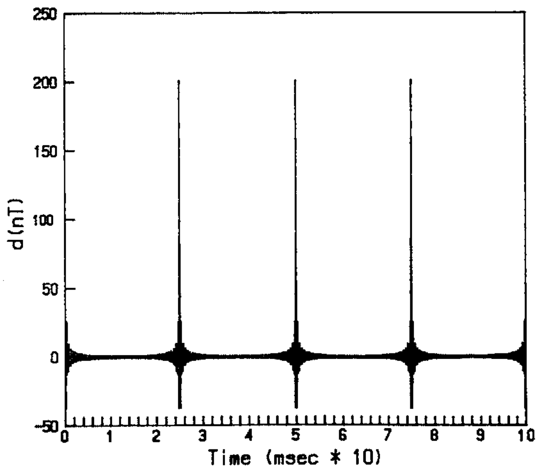

Fig 2 : An Example of the Function d(nT)

In the composite impulse response of the filterbank (Eqn. 10), the function d(nT) is weighted by the impulse response of the lowpass filter, h(nT). By iteratively adjusting the channel spacing and bandwidth, a composite impulse response can be obtained which approaches the ideal, that is a single delayed impulse. This is illustrated in Fig 3. Note that an RFT filterbank measures the spectral content of a signal by repeatedly heterodyning and lowpass filtering, whereas the composite impulse response shown is for an equivalent bank of bandpass filters.

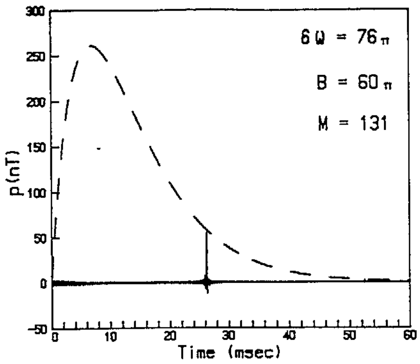

Fig 3 : Conceptual Composite Impulse Response of RFT Filterbank.

It is possible to extend the above method of design to obtain a degree of non uniformity in the channel spacing and bandwidth of the filterbank. This is achieved by dividing the frequency band of interest into a number of sections. Each section must have uniform resolution and definition but the values of each can vary from section to section. The equivalent composite impulse response can be computed for each section and the design procedure applied to each in order to optimise the overall composite frequency response.

To ensure that the relative gains of each section are equal, a scaling factor is required for each section such that the amplitude level measured for a unit-amplitude sinusoid, falling exactly on the centre frequency of any channel, will always be unity. This is readily achieved by ensuring that the gain of each filter is unity at dc. From equation (6), the scaling factor S is given by

$$S = [1 + \sum_{m=1}^{k+1} b_m e^{-\alpha mT}]/[T^k . \sum_{n=1}^{k} a_n e^{-\alpha nT}] \qquad (12)$$

The a_n coefficients are always positive but the b_m coefficients are negative for m odd. To ensure that S is neither negative or zero, it has been shown [Murphy, 1988] that for lowpass filter orders of 2, 3, 4, 5 and 6, analysing bandwidths less than 1, 11, 44, 90 and 160Hz respectively are non-usable when used in a filterbank with varying filter order. This limitation on analysing bandwidth is a result of preserving the desired impulse response of $(nT)^k.e^{-\alpha nT}$.

To illustrate the design procedure, consider a filterbank, covering the range dc to 4500Hz, which is split into three 1500Hz sections with the channel spacing and bandwidth increasing by a factor of 2 from section to section. The complete initial specification for this filterbank (FB1) is given in Table 1. To minimise distortion in the composite amplitude response due to step increases in bandwidth at section boundaries, the filter orders also increase from section to section. Fig 4 gives its composite amplitude response. It is clear that the response deviates significantly from the ideal.

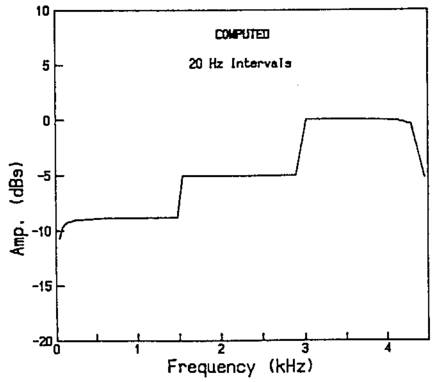

Fig 4 : Composite Amplitude Response of RFT Filterbank FB1.

By iteratively adjusting the channel bandwidth, the number of channels, the channel spacing and filter order, in conjunction with viewing the resulting composite impulse response, p(nT), for each section, a composite amplitude response which approaches the ideal flat characteristic can be achieved. By further iteratively adjusting the bandwidth in each section and observing the resulting composite amplitude response a certain degree of fine-tuning of the filterbank design can be achieved. Fig 5 shows the composite amplitude response for a filterbank (FB2) with a specification given in Table 1, which was designed by optimising the specification of FB1 in the above way.

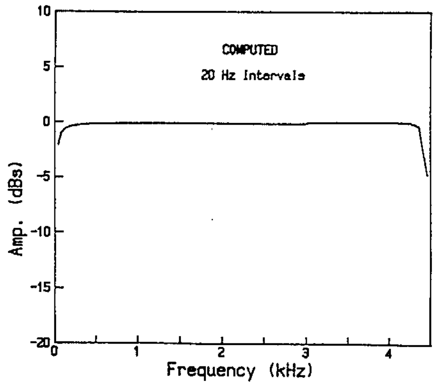

Fig 5 : Composite Amplitude Response of Optimised Filterbank FB2

Parameter	Section 1		Section 2		Section 3	
	FB1	FB2	FB1	FB2	FB1	FB2
Channels	37	37	18	25	10	15
1^{st}Chan.(Hz)	40	40	1540	1540	3020	3060
Spacing (Hz)	40	40	80	60	160	100
Filter Order	2	2	3	3	5	5
Bandwidth(Hz)	32	35	64	44	128	72

Table 1 Filterbank Specifications.

3. Spectral Smoothing.

One way of computing the envelope of a discrete spectrum is to treat the spectral components as a time series and smooth them using a bipass (forward and reverse time) filter [Kormylo, 1974]. However, the non-uniform sampling interval that results from using non-uniform channel spacing must be removed by interpolation. The result is a spectrum with non-uniform resolution but with linear definition. Such a spectrum can be smoothed not only by lowpass filtering but also by homomorphic filtering.

Fig 6 shows the unsmoothed short-time spectrum of a synthetic /a/ vowel, as computed using the optimised filterbank, FB2, in Table 1, overplotted with the transmission of the model used to synthesise the signal. Fig 7 shows the computed spectrum after linear interpolation and lowpass filtering with a variable, second-order, Butterworth lowpass filter. Different cut-off frequencies were used for each filterbank section. Only over the first two sections is there a reasonable spectral match. To remove all of the harmonic detail from section 3 requires an excessive level of smoothing.

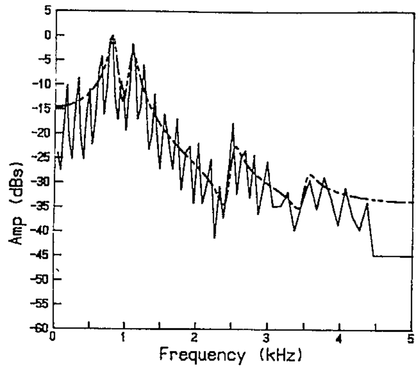

Fig 6 : Unsmoothed Short-Time Spectrum of Synthetic /a/ Vowel.

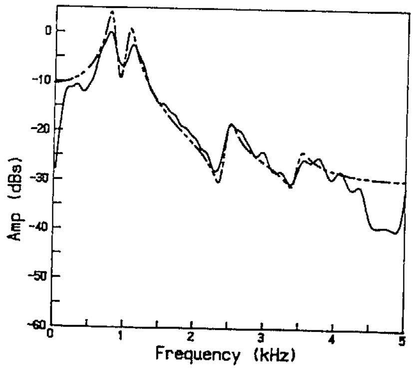

Fig 7 : Filter-Smoothed Short-Time Amplitude Spectrum of Vowel /a/.

In contrast to bipass filtering, homomorphic filtering results in the same level of smoothing being applied to all sections. However, multiple peaks of unequal amplitude occur in the cepstrum due to the different channel spacing in each section and, therefore, the spectrum of the high-time part of the cepstrum will not be flat (Fig 8). In cepstral smoothing, this spectrum is effectively subtracted from the original spectrum and so the smoothed spectrum will contain the spectral trend of the high-time part of the cepstrum. The amount of spectral trend can be determined by carrying out a linear regression on the peak amplitudes of the high-time harmonics in Fig 8 and this can then be removed from the smoothed spectrum. Fig 9 illustrates the result of applying this process to a synthetic /a/ vowel. A comparison of Fig 9 with Fig 7 highlights the superiority of using cepstral smoothing with trend removal over bi-pass filtering for deriving the spectral envelope of a non-uniform spectrum.

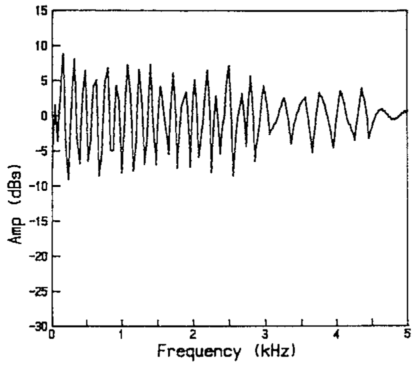

Fig 8 : High-Time Spectrum of Synthetic /a/ computed from RFT Filterbank (FB2).

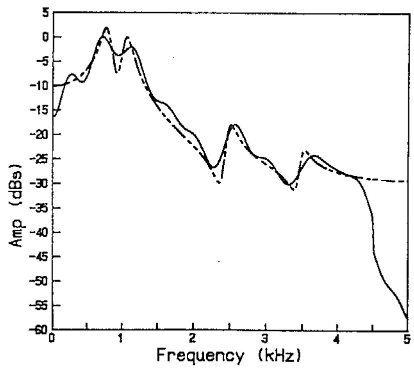

Fig 9 : Cepstrally Smoothed Filterbank (FB2) Spectrum of Synthetic /a/ with Spectral Trend Removed.

4. Conclusions

This paper has described a technique for designing uniform and non-uniform filterbanks based on the Running Fourier Transform (RFT). The RFT is implemented by convolving the input signal with one of a family of windows, $h(nT) = (nT)^k e^{-\alpha nT}$, where k and α may be chosen to specify the order and bandwidth, respectively, of each analysing filter.

Computation of an equivalent composite impulse response has been used to optimise the composite amplitude response of any RFT filterbank. Finally a modified cepstral smoothing technique for non-uniform spectra has been presented and has been shown to be superior to conventional bi-pass filtering.

References

Flanagan (1972), J L, Speech Analysis, Synthesis and Perception, Springer, Berlin, 2nd edition, 1972, p 142.

Kormylo (1974), J J & Jain, V K, Two-Pass Recursive Digital Filter with Zero Phase Shift, IEEE Trans, ASSP-22, No 5, 1974, pp 384-387.

Murphy (1988), M S, Speech Processing For A Spectral Pole-Zero Model, DPhil Thesis, University of Ulster, UK, 1988.

Owens [1988], F J & Murphy, M S, A Short-Time Fourier Transform, Signal Processing, North-Holland, Vol 14, No 1, 1988, pp 3-10.