

Efficient Design of Oversampled NPR GDFT Filter Banks*

Matthew R. Wilbur, Timothy N. Davidson and James P. Reilly

Technical Report,
Department of Electrical and Computer Engineering,
McMaster University,
Hamilton, Ontario, L8S 4K1, Canada.

Email: davidson@mcmaster.ca

July 2003

Abstract

In this paper, we present a flexible, efficient design technique for the prototype filter of an oversampled near perfect reconstruction (NPR) generalized Discrete Fourier Transform (GDFT) filter bank. Such filter banks have several desirable properties for subband processing systems which are sensitive to aliasing, such as subband adaptive filters. Our design criteria for the prototype filter are explicit bounds (derived herein) on the aliased components in the subbands and the output, the distortion induced by the filter bank, and the imaged subband errors in the output. It is shown that the design of an optimal prototype filter can be transformed into a convex optimization problem which can be efficiently solved. Our design technique provides an efficient and effective tool for exploring many of the inherent trade-offs in the design of the prototype filter, including the trade-off between aliasing in the subbands and the distortion induced by the filter bank. In our examples we calculate several of these trade-offs and demonstrate that our method can generate filters with significantly better performance than filters obtained using current design methods.

*A condensed version of this report has been accepted, subject to minor revisions, for publication in the *IEEE Transactions on Signal Processing*

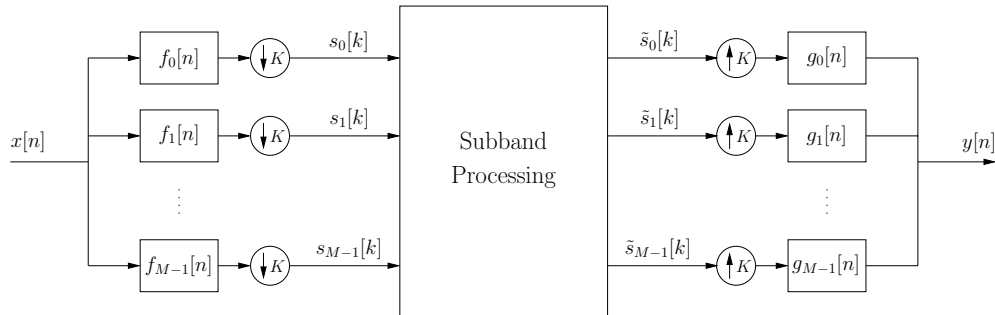


Figure 1: An M -channel uniform subband signal processing system.

1 Introduction

Uniform multi-rate filter banks form the basic unit of many multi-rate signal processing systems in a diverse set of applications that includes audio and image compression, denoising, feature detection and extraction, and adaptive filtering [1–4]. A typical example of such a system is illustrated in Figure 1. The standard design techniques for uniform filter banks are based on (approximating) the perfect reconstruction condition that in the absence of any subband processing, the output signal is simply a scaled and delayed version of the input [1–4]. It is now well known that perfect reconstruction can be (exactly) achieved with finite impulse response (FIR) filters, in both the critically-sampled ($K = M$) and oversampled ($K < M$) arrangements [2–6]. Classes of efficiently implementable ‘modulated’ filter banks are also available [1–4, 7–10]. However, it is becoming apparent that perfect reconstruction filter banks do not necessarily provide optimal performance of the subband signal processing system as a whole; e.g., [11]. Designs based on the perfect reconstruction condition typically allow considerable aliasing in the subband signals, $s_m[k]$, but structure these aliased components so that in the absence of any subband processing they are cancelled by the synthesis filter bank. This characteristic may be undesirable if the subband processing block is sensitive to aliasing in the subband signals, or if the subband processing block distorts the aliased components in the subband signals in a way that reduces the effectiveness of alias cancellation. An application which is particularly sensitive to aliasing in the subbands is subband adaptive filtering [12–17]. The subband processing block of a subband adaptive filtering system typically consists of a diagonal matrix of adaptive filters, each of which operates on one of the subband signals, and operates independently of the other adaptive filters.¹ The objective of the subband processing block is, quite naturally, to filter the subband signals. However, by doing so it may distort the aliased components in the subband signals and hence reduce the effectiveness of alias cancellation.

For subband adaptive filtering, and other applications with similar characteristics, near perfect reconstruction (oversampled) filter banks which *suppress* aliasing in the subbands and ‘imaging’ [2]

¹The use of off-diagonal ‘cross’ filters [13] can reduce the sensitivity to aliasing, but may adversely affect the convergence properties of the adaptive filter, and increases the computational cost of the system.

of the subband errors in the output, rather than relying on *cancellation*, offer the potential for improved performance. In particular, the performance of systems based on the class of oversampled generalized Discrete Fourier Transform (GDFT) filter banks is quite encouraging [16–22]. These filter banks are better able to suppress aliasing in the subbands than (uniform) oversampled cosine-modulated filter banks [18, 19, 21], and can be efficiently implemented using the GDFT [1]. In this paper we provide a flexible, efficient design technique for the prototype filter of an oversampled near perfect reconstruction (NPR) GDFT filter bank. The design criteria are explicit bounds (derived herein) on the aliased components in the subbands (and the output), the imaged subband error components in the output, and the distortion induced by the filter bank. These bounds rigorously amalgamate several intuitively developed design criteria in the current literature [15, 18–22], and subsume the criteria derived in [23] and [16, Section IV–B]. Our design criteria generate familiar constraints on the prototype filter: the aliasing criteria result in bounds on the stop-band energy and the maximum stop-band level, the imaging criterion results in an additional bound on the transition-band energy, and the distortion criterion results in a measure of the distance between the prototype filter and a ‘self-orthogonal’ filter.

In their direct form, these constraints generate a non-convex feasible set. Therefore, careful detection and management of locally optimal solutions may be required in order to obtain a filter which achieves an objective value sufficiently close to that of a globally optimal prototype.² The key step in obtaining our efficient design technique is to show that the design criteria can be (precisely) transformed into convex functions of the (deterministic) autocorrelation of the prototype filter, and hence that a (globally) optimal prototype filter can be obtained from the solution of a convex optimization problem that can be efficiently solved. A feature of our proposed method is that the semi-infinite constraints generated by the maximum stop-band level constraint and the constraint that the autocorrelation can be spectrally factorized are transformed into (finite) linear matrix inequalities.

Our convex formulation not only provides an efficient algorithm for finding an optimal prototype filter, but by doing so it provides an efficient method for determining the inherent trade-offs between competing prototype design criteria. Of particular interest is the inherent trade-off between aliasing in the subbands and the distortion induced by the filter bank. In our examples we calculate several of these trade-offs and demonstrate that filters designed via our formulation can provide significantly better performance than filters designed using current methods.

The paper is arranged as follows: In Section 2 we describe the class of GDFT filter banks and discuss the relationships between our design method and current techniques in greater detail. In Section 3 we derive our design criteria and in Section 4 we provide the transformation to the convex design problem. In Section 5 we demonstrate how our method generates filters which achieve many of the inherent design trade-offs, and in Section 6 we illustrate how this leads to improved

²Rigorous methods (e.g., [24]) are effective but tend to be computationally expensive. Simpler techniques, such as running a standard local optimization algorithm from multiple starting points, are less expensive but may not be as effective.

performance over filters designed by two competing methods.

2 GDFT Filter Banks

The generalized Discrete Fourier Transform (GDFT) filter bank [1] we will consider is of the form in Figure 1, with the analysis and synthesis filters consisting of exponentially modulated versions of a single real-valued FIR prototype filter $p[n]$:

$$f_m[n] = p[n]e^{j2\pi(m+m_0)(n+n_0)/M}, \quad (1a)$$

$$g_m[n] = f_m[-n]^* = p[-n]e^{j2\pi(m+m_0)(n-n_0)/M}, \quad (1b)$$

where m_0 and n_0 are frequency-shift and phase-shift constants, respectively. Such filter banks can be efficiently implemented using the GDFT [1]. It is well known that the only critically-sampled filter banks of the form in (1) with perfect reconstruction (PR) are those generated by a prototype filter which is a length M rectangular window [25]. Such a prototype filter leads to substantial aliasing in the subbands. Moreover, critically-sampled near perfect reconstruction (NPR) filter banks of the form in (1) have approximately the same prototype. In the oversampled case, there are longer prototypes which generate PR and NPR filter banks of the form in (1), and hence prototypes which generate much lower aliasing in the subbands can be obtained. The focus of this paper is on NPR filter banks, because the effectiveness of the alias cancellation in PR filter banks may be compromised by the subband processing, and because improved alias suppression in the subbands and image suppression in the output can be achieved by relaxing the PR constraint. The fact that we employ the same prototype filter in the analysis and synthesis banks means that our NPR filter banks have polyphase matrices [2] which are nearly paraunitary, and hence have favourable noise robustness and numerical properties.

One possible approach to the design of oversampled GDFT filter banks is to use the fact that any prototype filter for an L -times oversampled PR cosine-modulated filter bank is a prototype for a $2L$ -times oversampled PR DFT filter bank [9, 10]. The design of prototype filters for NPR and PR cosine-modulated filter banks is usually expressed as a constrained optimization problem for the filter coefficients [26, 27], or as an unconstrained optimization problem over some alternative variables (in some cases called lattice coefficients) which are related to the filter coefficients in such a way that PR is automatically satisfied [29]. Unfortunately, the constraints in the former method and the objective in the latter are non-convex functions of the design parameters, and hence these formulations can be quite awkward to solve. Furthermore, exploiting the relationship with cosine-modulated filter banks generates a subclass of the (G)DFT filter banks, but not the whole class [9, 10]. As one might expect, direct formulation of the design of a GDFT prototype (over the whole available class) also leads to non-convex optimization problems [16, 18, 19, 22]. (Some simple, but ad-hoc, prototype design methods have also been proposed [15, 20].) The approach

we outline in this paper involves the transformation³ of the design into a convex optimization problem in which the autocorrelation coefficients of the prototype are the design variables. This convex optimization problem can be efficiently solved for the (globally) optimal autocorrelation, from which an optimal prototype can be obtained via spectral factorization. A number of methods of spectral factorization are described in [30,35]. The simplicity of autocorrelation-based prototype design has been recognized before [36], in the context of cosine-modulated filter banks. However, in [36] the transformation of the design criteria into functions of the autocorrelation sequence involves approximations, and these approximations manifest themselves in undesirable properties of the prototype. In contrast, our design transformation is precise, and since our design criteria are explicit bounds on the aliasing, imaging and distortion energies, we obtain prototype filters with many desirable properties.

3 Derivation of Design Criteria

In this section we establish design criteria for the prototype filter which enable us to control the aliased components which appear in the subband signals $s_m[k]$, the imaged components which appear in the output, and the distortion induced by the filter bank. As discussed in the Introduction, and studied in greater detail in [16, 17], the performance of a GDFT-filter-bank-based subband adaptive filter is critically dependent on these terms.

If $X(z) = \sum_n x[n]z^{-n}$ denotes the z -transform of the input signal, then the m -th subband signal in Figure 1 has a z -transform

$$S_m(z) = \frac{1}{K}F_m(z^{1/K})X(z^{1/K}) + \frac{1}{K} \sum_{k=1}^{K-1} F_m(z^{1/K}W_K^k)X(z^{1/K}W_K^k), \quad (2)$$

where $W_K = e^{-j2\pi/K}$. The first term on the right hand side of (2) represents the desired component of the subband signal, and the second term represents the aliased components.

The nature of the output of the subband processing system in Figure 1 clearly depends on the nature of the subband processing block. In subband adaptive filtering applications, the subband processing block typically consists of a diagonal matrix of adaptive filters. If the adaptive filter in the m th subband has converged to $H_m(z)$ [and is ‘frozen’ from that point], then the subband

³Similar transformations have also led to convex formulations of some other FIR filter design problems [30–34], but the design problems considered herein are specially tailored to the design of prototype filters for oversampled NPR GDFT filter banks.

outputs are $\tilde{S}_m(z) = H_m(z)S_m(s)$, and the output of the system in Figure 1 is

$$\begin{aligned} Y(z) &= \sum_{m=0}^{M-1} G_m(z)\tilde{S}_m(z) \\ &= \frac{1}{K} \sum_{m=0}^{M-1} G_m(z)H_m(z^K)F_m(z)X(z) + \frac{1}{K} \sum_{m=0}^{M-1} G_m(z)H_m(z^K) \sum_{k=1}^{K-1} F_m(zW_K^k)X(zW_K^k). \end{aligned} \quad (3)$$

If there is no subband processing, then $H_m(z) = 1$ and (3) simplifies to the standard expression for the output of a filter bank [1–4]. The transfer function $D(z) = (1/K) \sum_{m=0}^{M-1} G_m(z)H_m(z^K)F_m(z)$ in (3) describes the processing of the input signal $X(z)$, while the second term in (3) represents the aliased components which appear at the output. In the following subsections we will analyze the aliased components in (2) and (3), the undesired distortion induced by the filter bank, and the imaged components of errors in $\tilde{S}_m(z)$ that appear in the output. We will show how the energy of each term can be bounded by natural properties of the prototype filter. For simplicity, our analysis is based on a deterministic input signal, $x[n]$, but a complementary analysis based on a stochastic model for $x[n]$ can also be performed.

3.1 Aliasing in the subbands

Our analysis of the aliased components in the m th subband signal, $s_m[k]$, will be expedited by considering an up-sampled version of $s_m[k]$ which has the same sample rate as $x[n]$. This signal will be denoted by $v_m[n]$, where $v_m[kK] = s_m[k]$ and $v_m[kK + i] = 0$, $i = 1, 2, \dots, K - 1$. Since up-sampling does not change the energy of the signal, and since the spectrum of $v_m[n]$ is periodic (in frequency) with period $2\pi/K$, the energy of $s_m[k]$ can be computed from $v_m[n]$ as follows:

$$E_{s_m} = \frac{1}{2\pi} \int_{-\pi}^{\pi} |S_m(e^{j\omega})|^2 d\omega = \frac{1}{2\pi} \int_{\Omega} |V_m(e^{j\omega})|^2 d\omega, \quad (4)$$

where Ω is an interval of width $2\pi/K$. For our analysis, a convenient interval is

$$\Omega_M^m = \left\{ \omega : \frac{2\pi(m+m_0)}{M} - \frac{\pi}{K} \leq \omega < \frac{2\pi(m+m_0)}{M} + \frac{\pi}{K} \right\}, \quad (5)$$

which is the principal spectral support of the m -th analysis filter in a GDFT filter bank. To simplify some of our expressions, we define

$$\alpha_m = 2\pi(m + m_0)/M \quad \text{and} \quad \psi_k = 2\pi k/K. \quad (6)$$

Using (4) and (2), the energy of the aliased components in the m -th subband signal for a deterministic signal $x[n]$ is

$$E_{A_m} = \frac{1}{2\pi} \int_{\omega \in \Omega_M^m} \left| \frac{1}{K} \sum_{k=1}^{K-1} F_m(e^{j(\omega-\psi_k)})X(e^{j(\omega-\psi_k)}) \right|^2 d\omega. \quad (7)$$

Setting $\phi = \omega - \alpha_m$ and using the fact (1) that $F_m(e^{j\omega}) = e^{j\alpha_m n_0} P(e^{j(\omega - \alpha_m)})$, where $P(e^{j\omega})$ is the frequency response of the prototype filter, we have that

$$\begin{aligned} E_{A_m} &= \frac{1}{2\pi} \int_{-\pi/K}^{\pi/K} \left| \frac{1}{K} \sum_{k=1}^{K-1} P(e^{j(\phi - \psi_k)}) X(e^{j(\phi + \alpha_m - \psi_k)}) \right|^2 d\phi \\ &= \frac{1}{2\pi K^2} \sum_{k,\ell=1}^{K-1} \int_{-\pi/K}^{\pi/K} P(e^{j(\phi - \psi_k)}) P^*(e^{j(\phi - \psi_\ell)}) X(e^{j(\phi + \alpha_m - \psi_k)}) X^*(e^{j(\phi + \alpha_m - \psi_\ell)}) d\phi. \end{aligned} \quad (8)$$

We will now develop bounds for E_{A_m} under the assumption that $\sum_n |x[n]|$ is finite. This assumption ensures that the signal energy $E_x = \sum_n |x[n]|^2$ and

$$U_X \triangleq \max_{\omega} |X(e^{j\omega})| \quad (9)$$

are finite, and also that

$$J_X(\mu) \triangleq \frac{1}{2\pi} \int_{-\pi}^{\pi} |X(e^{j\omega}) X^*(e^{j(\omega + \mu)})| d\omega \quad (10)$$

is finite for all μ . (The term $J_X(\mu)$ is closely related to the spectral correlation [37] of $x[n]$ for a frequency separation of μ .) Note that $J_X(0) = E_x$ and that $J_X(-\mu) = J_X(\mu)$. Since $J_X(\mu)$ is finite, the following portion of $J_X(\mu)$ is also finite:

$$\tilde{J}_X(\mu_1, \mu_2, \alpha) \triangleq \frac{1}{2\pi} \int_{-\pi/K + \alpha}^{\pi/K + \alpha} |X(e^{j(\omega - \mu_1)}) X^*(e^{j(\omega - \mu_2)})| d\omega \quad (11)$$

By taking the absolute value of the integrand in (8), extracting the terms dependent on the prototype filter we obtain the following bound on E_{A_m} :

$$E_{A_m} \leq \frac{1}{2\pi K^2} \sum_{k,\ell=1}^{K-1} \max_{\phi \in [-\pi/K, \pi/K]} |P(e^{j(\phi - \psi_k)}) P^*(e^{j(\phi - \psi_\ell)})| \int_{-\pi/K + \alpha_m}^{\pi/K + \alpha_m} |X(e^{j(\lambda - \psi_k)}) X^*(e^{j(\lambda - \psi_\ell)})| d\lambda \quad (12)$$

$$= \frac{1}{K^2} \sum_{k,\ell=1}^{K-1} \max_{\phi \in [-\pi/K, \pi/K]} |P(e^{j(\phi - \psi_k)}) P^*(e^{j(\phi - \psi_\ell)})| \tilde{J}_X(\psi_k, \psi_\ell, \alpha_m). \quad (13)$$

Let $U_P \triangleq \max_{\omega} |P(e^{j\omega})|$ denote the maximal spectral component of the prototype, and let $U_{P,\text{sb}} \triangleq \max_{\omega \in [\pi/K, 2\pi - \pi/K]} |P(e^{j\omega})|$, with $U_{P,\text{sb}} \leq U_P$, denote the maximum sidelobe level of the prototype. That is, let the prototype filter satisfy a spectral mask of the form shown in Figure 2. Since $k, \ell \geq 1$ we have that

$$\begin{aligned} \max_{\phi \in [-\pi/K, \pi/K]} |P(e^{j(\phi - \psi_k)}) P^*(e^{j(\phi - \psi_\ell)})| &\leq \max_{\phi \in [-\pi/K, \pi/K]} |P(e^{j(\phi - \psi_k)})| \max_{\phi \in [-\pi/K, \pi/K]} |P^*(e^{j(\phi - \psi_\ell)})| \\ &= U_{P,\text{sb}}^2. \end{aligned} \quad (14)$$




Figure 2: The mask on the magnitude spectrum of the prototype filter.

Hence,

$$E_{A_m} \leq \frac{U_{P,\text{sb}}^2}{K^2} \sum_{k,\ell=1}^{K-1} \tilde{J}_X(\psi_k, \psi_\ell, \alpha_m). \quad (15)$$

Equation (15) is an explicit bound for the energy of the aliased components in each subband signal and is a simple multiple of the square of the maximum stop-band level of the prototype filter.

An alternative bound on E_{A_m} can be obtained by writing $E_{A_m} = E_{A_{m,a}} + E_{A_{m,b}}$, where $E_{A_{m,a}}$ contains the terms in (8) with $\ell = k$, and $E_{A_{m,b}}$ contains the remaining terms. Using the technique used to obtain (15), $E_{A_{m,b}}$ can be bounded as $E_{A_{m,b}} \leq (U_{P,\text{sb}}^2/K^2) \sum_{k=1}^{K-1} \sum_{\ell=1, \ell \neq k}^{K-1} \tilde{J}_X(\psi_k, \psi_\ell, \alpha_m)$. The remaining term is

$$\begin{aligned} E_{A_{m,a}} &= \frac{1}{2\pi K^2} \sum_{k=1}^{K-1} \int_{-\pi/K}^{\pi/K} |P(e^{j(\phi-\psi_k)})X(e^{j(\phi+\alpha_m-\psi_k)})|^2 d\phi \\ &= \frac{1}{2\pi K^2} \int_{-2\pi+\pi/K}^{-\pi/K} |P(e^{j\phi})|^2 |X(e^{j(\phi+\alpha_m)})|^2 d\phi \\ &\leq E_{P,\text{sb}} \frac{U_X^2}{K^2}, \end{aligned} \quad (16)$$

where the last step is obtained by applying the Hölder inequality [38] and observing that

$$\max_{\theta \in [-2\pi+\pi/K, -\pi/K]} |X(e^{j(\theta+\alpha_m)})| \leq U_X.$$

The term $E_{P,\text{sb}}$ is the stop-band energy of the filter,

$$E_{P,\text{sb}} = \frac{1}{2\pi} \int_{\pi/K}^{2\pi-\pi/K} |P(e^{j\theta})|^2 d\theta. \quad (17)$$

Hence, we have the following alternative bound to that in (15),

$$E_{A_m} \leq E_{P,\text{sb}} \frac{U_X^2}{K^2} + \frac{U_{P,\text{sb}}^2}{K^2} \sum_{k=1}^{K-1} \sum_{\substack{\ell=1 \\ \ell \neq k}}^{K-1} \tilde{J}_X(\psi_k, \psi_\ell, \alpha_m), \quad (18)$$

which is a linear combination of the stop-band energy and the square of the maximum stop-band level of the prototype filter.

3.2 Aliasing in the output

The analysis of the aliased components in the output of the subband processing system is more difficult than that of the aliased components in the subbands, because it depends on the nature of the subband processing. However, for the adaptive filtering scenario described at the beginning of Section 3, we can determine the energy of the aliased components in the output by evaluating the energy of the second term on the right hand side of (3). That is,

$$E_{A_{\text{out}}} = \frac{1}{2\pi} \int_{-\pi}^{\pi} \left| \frac{1}{K} \sum_{m=0}^{M-1} \sum_{k=1}^{K-1} G_m(e^{j\omega}) H_m(e^{jK\omega}) F_m(e^{j(\omega-\psi_k)}) X(e^{j(\omega-\psi_k)}) \right|^2 d\omega. \quad (19)$$

Using techniques akin to those used to derive (18) (the details of which have been omitted for brevity) we can obtain the following bound for $E_{A_{\text{out}}}$:

$$E_{A_{\text{out}}} \leq \frac{U_P^2}{K^2} \left(E_{P,\text{sb}} K U_X^2 \left(\sum_m U_{H_m}^2 \right) + U_{P,\text{sb}}^2 \sum_{k=1}^{K-1} \sum_{\substack{\ell=1 \\ \ell \neq k}}^{K-1} \sum_{m=0}^{M-1} \sum_{\substack{n=0 \\ n \neq m}}^{M-1} \Gamma_X(m, n, k, \ell) \right), \quad (20)$$

where $U_{H_m} \triangleq \max_{\omega} |H_m(e^{j\omega})|$ and

$$\Gamma_X(m, n, k, \ell) = \frac{1}{2\pi} \int_{-\pi}^{\pi} |H_m(e^{jK\omega}) H_n^*(e^{jK\omega}) X(e^{j(\omega-\psi_k)}) X^*(e^{j(\omega-\psi_\ell)})| d\omega. \quad (21)$$

This bound consists of simple functions of the maximum stop-band level, the stop-band energy, and the maximum component of the spectrum of the prototype filter. If the adaptive filters are such that $U_{H_m} \leq 1$, then $\Gamma_X(m, n, k, \ell) \leq J_X(\psi_{k-\ell})$, where $J_X(\mu)$ was defined in (10), and we can obtain the following simplified bound:

$$E_{A_{\text{out}}} \leq U_P^2 \frac{M}{K^2} \left(E_{P,\text{sb}} K U_X^2 + U_{P,\text{sb}}^2 (M-1) \sum_{k=1}^{K-1} \sum_{\substack{\ell=1 \\ \ell \neq k}}^{K-1} J_X(\psi_{k-\ell}) \right). \quad (22)$$

3.3 Distortion in the output

The third performance measure for the filter bank is the distortion which it induces on the signal. In order to isolate the (undesirable) distortion induced by the filter bank from the (desired) processing performed by the subband processing block, we will analyze the distortion in the absence of any subband processing. That is, with $H_m(z) = 1$ in (3). If the filter bank is distortion free then $(1/K) \sum_{m=0}^{M-1} G_m(z) F_m(z) = cz^{-d}$, where c is a constant and d is an integer [2–4]. In that case,

the first term in (3) becomes,

$$\frac{1}{K} \sum_{m=0}^{M-1} G_m(z) F_m(z) X(z) = cX(z)z^{-d}.$$

Without loss of generality in the design criteria for the prototype filter, we can restrict our attention to the case where the output is neither scaled nor delayed; i.e., where $c = 1$ and $d = 0$. In that case, the energy of the distortion induced by the filter bank is

$$E_D = \frac{1}{2\pi} \int_{-\pi}^{\pi} \left| \frac{1}{K} \sum_{m=0}^{M-1} G_m(e^{j\omega}) F_m(e^{j\omega}) X(e^{j\omega}) - X(e^{j\omega}) \right|^2 d\omega \quad (23)$$

$$= \frac{1}{2\pi} \int_{-\pi}^{\pi} \left| \left(\frac{1}{K} \sum_{m=0}^{M-1} |P(e^{j(\omega-\alpha_m)})|^2 - 1 \right) X(e^{j\omega}) \right|^2 d\omega. \quad (24)$$

This may be bounded in the following way:

$$E_D \leq \frac{U_X^2}{2\pi} \int_{-\pi}^{\pi} \left| \frac{1}{K} \sum_{m=0}^{M-1} |P(e^{j(\omega-\alpha_m)})|^2 - 1 \right|^2 d\omega \quad (25)$$

$$= U_X^2 \sum_n \left| \sum_{\ell} \frac{M}{K} p[\ell] p[\ell - nM] - \delta[n] \right|^2, \quad (26)$$

where $\delta[n]$ is the Kronecker delta, and we have used Parseval's Relation. (The derivation of (26) is analogous to that of (44) in the Appendix for which more details are provided.) If we normalize the energy of the prototype filter so that $\sum_{\ell} p[\ell]^2 = K/M$, then

$$E_D \leq \gamma_p^2 \frac{MU_X^2}{K}, \quad (27)$$

where

$$\gamma_p^2 = \sum_{\substack{n \\ n \neq 0}} \left| \sum_{\ell} p[\ell] p[\ell - nM] \right|^2. \quad (28)$$

As shown in the Appendix, for a normalized prototype the energy of the distortion in the output can also be bounded by

$$E_D \leq \tilde{\gamma}_p^2 \frac{ME_x}{K}, \quad (29)$$

where $\tilde{\gamma}_p = \sum_{n, n \neq 0} \sum_{\ell} |p[\ell] p[\ell - nM]|$ and, as previously defined, $E_x = \sum_n x[n]^2$ is the energy of the input signal. In this paper we will focus on the bound on E_D in (27), but previous work in a different application [34] suggests that using the bound in (29) will generate qualitatively similar results. This is because both γ_p and $\tilde{\gamma}_p$ are measures of the 'distance' between $p[\ell]$ and a filter $q[\ell]$ which satisfies $\sum_{\ell} q[\ell] q[\ell - nM] = K/M \delta[n]$. (Such $q[\ell]$ are often said to be self-orthogonal [3].) The term γ_p is a two-norm measure, and $\tilde{\gamma}_p$ is a one-norm measure.

3.4 Imaging of subband errors in the output

The remaining performance measure of the filter bank captures the extent to which errors between the desired and actual outputs of the m th subband adaptive filter corrupt the filter bank output in other subbands. To be more precise, let $\tilde{y}_m(z)$ denote the (z -transform of the) desired output of the m th subband adaptive filter. The error signal which drives that adaptive filter is $\tilde{E}_m(z) = \tilde{y}_m(z) - \tilde{S}_m(z)$; see also Figure 11 in Section 6. Such errors naturally manifest themselves in the m th subband of the output. However, images of $\tilde{E}_m(z)$ also appear outside the m th subband of the output. If we let $\tilde{\Omega}_M^m = \{\omega : \alpha_m - \pi/M \leq \omega \leq \alpha_m + \pi/M\}$ denote the m th subband of the output, then the energy of the imaged components of $\tilde{E}_m(z)$ is

$$E_{I_m} = \frac{1}{2\pi} \int_{(-\pi, \pi] \setminus \tilde{\Omega}_M^m} |G_m(e^{j\omega}) \tilde{E}_m(e^{jK\omega})|^2 d\omega.$$

If we let $U_{\tilde{E}_m} \triangleq \max_{\omega} |\tilde{E}_m(e^{j\omega})|$, then

$$E_{I_m} \leq \frac{U_{\tilde{E}_m}^2}{2\pi} \int_{\pi/M}^{2\pi - \pi/M} |P(e^{j\phi})|^2 d\phi = U_{\tilde{E}_m}^2 (E_{P,\text{tb}} + E_{P,\text{sb}}),$$

where $E_{P,\text{tb}} = (1/\pi) \int_{\pi/M}^{\pi/K} |P(e^{j\omega})|^2 d\omega$ is the energy in the transition band of the prototype filter, and $E_{P,\text{sb}}$ is the stop-band energy.

4 Prototype design

As we have argued in the Introduction, the performance of subband adaptive filtering systems is sensitive to the error mechanisms discussed in Section 3. Now that the extent of these error mechanisms for oversampled GDFT filter banks has been bounded by simple functions of the prototype filter, we can design filters which optimize these bounds, so that optimized subband signal processing systems can be realized. Using the analysis in Section 3, it is clear that:

1. For a given normalization, small values of the maximum stop-band level, $U_{P,\text{sb}}$, the stop-band energy, $E_{P,\text{sb}}$, and the maximum spectral component, U_P , of the prototype filter will guarantee that the energy of the aliased components in the subbands and the output are small.
2. A small value of γ_p in (28) (and/or $\tilde{\gamma}_p$) will guarantee that the energy of the (amplitude and phase) distortion in the output is small.
3. Small values of the transition-band energy, $E_{P,\text{tb}}$, and the stop-band energy, $E_{P,\text{sb}}$, will guarantee that the energy of imaged subband errors in the output are small.

Although various combinations of some of these criteria have been employed by other authors (on a somewhat ad-hoc basis), we have shown how they explicitly bound the energy of the aliased

components in the subbands and the output, the energy of the distortion in the output, and the energy of the imaged subband errors in the output. Natural design criteria for the prototype can be obtained by minimizing a (linear) combination of $U_{P,\text{sb}}$, $E_{P,\text{sb}}$, U_P , γ_p and $E_{P,\text{tb}}$, subject to bounds on their individual values. For example, we might wish to find the length L prototype filter which minimizes the stop-band energy, subject to fixed bounds on the maximum stop-band level, the distortion coefficient, the transition-band energy and the maximum spectral component of the filter, and subject to the filter being normalized. That is, we might seek the solution of the following optimization problem:

$$\min_{\substack{p[\ell] \\ 0 \leq \ell \leq L-1}} \frac{1}{2\pi} \int_{\pi/K}^{2\pi-\pi/K} |P(e^{j\omega})|^2 d\omega \quad (30a)$$

$$\text{subject to } |P(e^{j\omega})| \leq \epsilon_{\text{sb}} \quad \forall \omega \in [\pi/K, \pi], \quad (30b)$$

$$\sum_{\substack{n \\ n \neq 0}} \left| \sum_{\ell} p[\ell] p[\ell - Mn] \right|^2 \leq \epsilon_{\gamma}^2, \quad (30c)$$

$$\frac{1}{\pi} \int_{\pi/M}^{\pi/K} |P(e^{j\omega})|^2 d\omega \leq \epsilon_{\text{tbe}}, \quad (30d)$$

$$|P(e^{j\omega})| \leq B \quad \forall \omega, \quad (30e)$$

$$\sum_{\ell} p[\ell]^2 = K/M, \quad (30f)$$

where ϵ_{sb} , ϵ_{γ}^2 , ϵ_{tbe} and B are fixed constants. Note that in contrast to standard peak-constrained least-squares filter design [39], our design criteria do not include an explicit lower bound on the magnitude spectrum in the pass-band of the prototype filter. However, by generalizing Nyquist's first criterion for intersymbol-interference-free pulse amplitude modulation, it can be shown that the distortion and maximum stop-band level constraints [(30c) and (30b), respectively] implicitly control the pass-band 'ripple' [34, Appendix A].

The integrals in (30a) and (30d) can be analytically evaluated, resulting in a convex quadratic objective and a convex quadratic constraint, respectively, and by squaring both sides of (30b) and (30e) we obtain two infinite sets of convex quadratic constraints (each expression generates one constraint for each relevant frequency). These infinite sets of constraints can be approximated by discretization [40]. (A simple discretization scheme is discussed in (32) below.) Therefore, in the absence of (30c) the problem in (30) could be efficiently solved, using, for example, second order cone programming techniques [41]. Unfortunately, the distortion constraint in (30c) is, in general, a non-convex quadratic function of $p[\ell]$. Hence, the problem in (30) is a non-convex optimization problem which may require careful (and computationally expensive) management of locally optimal solutions in order to obtain a filter whose performance is sufficiently close to that of a globally optimal filter. This is important because the objective and the constraints in (30) are competing criteria. For example, it is well known that there is a trade-off between the maximum stop-band level and the stop-band energy (e.g., [39]). Achieving a good design involves an exploration of the

trade-offs between these criteria, followed by the design of a filter which achieves a desired position on the trade-off surface. The non-convexity of (30) can make it quite awkward to get an accurate description of the trade-off surface, and hence quite difficult to determine how far a given filter is from providing an optimal trade-off. Furthermore, it can be quite awkward to determine when the constraints in (30) conflict so that there is no filter of the given length which satisfies all the constraints; i.e., reliable detection of infeasibility of (30) can be difficult to achieve.

Given the difficulties involved in solving (30), some authors have developed interesting approximations to the distortion measure (30c), such as considering only amplitude distortion and ignoring phase distortion [22]. Others have simplified the design by iteratively linearizing $\sum_{\ell} p[\ell]p[\ell - Mn]$ and employing the iterative least squares technique to determine (locally) optimal filters for the simplified formulation [19]. (The method in [19] will be discussed in more detail in Section 6.) In contrast, our method solves (30) without approximation or additional constraints, and allows efficient and accurate calculation of the design trade-offs and filters which achieve them.

The key observation in the development of our design method is that the objective and the constraints in (30) are all convex functions of the autocorrelation of the filter coefficients,

$$r_p[n] = \sum_{\ell} p[\ell]p[\ell - n].$$

Using the fact that $r_p[-n] = r_p[n]$ and $p(e^{j\omega}) = |P(e^{j\omega})|^2 = r_p[0] + 2\sum_{n \geq 1} r_p[n] \cos(\omega n)$, the integrals in (30a) and (30d) can be analytically evaluated. They are equal to $\sum_{n \geq 0} b_s[n]r_p[n]$ and $\sum_{n \geq 0} b_t[n]r_p[n]$, respectively, where $b_s[0] = 1 - 1/K$, $b_s[n] = -2\sin(\pi n/K)/(\pi n)$ for $n \geq 1$, $b_t[0] = 1/K - 1/M$ and $b_t[n] = 2(\sin(\pi n/K) - \sin(\pi n/M))/(\pi n)$ for $n \geq 1$. Therefore, the design problem in (30) can be transformed into the following optimization problem in $r_p[n]$, $n \geq 0$:

$$\min_{\substack{r_p[n] \\ 0 \leq n \leq L-1}} \sum_{n \geq 0} b_s[n]r_p[n] \quad (31a)$$

$$\text{subject to} \quad p(e^{j\omega}) \leq \epsilon_{\text{sb}}^2 \quad \forall \omega \in [\pi/K, \pi], \quad (31b)$$

$$\sum_{i \geq 1} r_p[Mi]^2 \leq \epsilon_{\gamma}^2/2, \quad (31c)$$

$$\sum_{n \geq 0} b_t[n]r_p[n] \leq \epsilon_{\text{tbe}} \quad (31d)$$

$$p(e^{j\omega}) \leq B^2 \quad \forall \omega, \quad (31e)$$

$$r_p[0] = K/M, \quad (31f)$$

$$p(e^{j\omega}) \geq 0 \quad \forall \omega. \quad (31g)$$

This change of variables must be handled carefully because not all sequences $r_p[n]$ are the autocorrelation coefficients of some filter. A necessary and sufficient condition for $r_p[n]$ to correspond to the autocorrelation coefficients of a filter is (31g); e.g., [42]. Given a sequence $r_p[n]$ which solves (31), filter coefficients $p[\ell]$ which generate this autocorrelation can be found using standard

spectral factorization techniques [30, 35].

The objective in (31a) and the constraints in (31b), (31d), (31e), (31f), (31g) are linear, and hence convex, in $r_p[n]$, $0 \leq n \leq L - 1$, and (31c) is a convex quadratic constraint.⁴ Therefore, the trade-offs between these competing prototype design criteria can be efficiently evaluated and an optimal autocorrelation efficiently found using convex optimization techniques. Furthermore, infeasibility of (31) can be reliably detected.

Unfortunately, the constraints in (31b), (31e) and (31g) each generate an infinite number of linear constraints on $r_p[n]$, one for each relevant frequency, and it may appear that these could be awkward to handle in practice. One possible approach is to approximate the constraints by discretization [40]. A simple scheme for discretizing (31b) is to choose a set of frequencies, $\pi/K \leq \omega_1 \leq \omega_2 \leq \dots \leq \omega_N \leq \pi$, often uniformly spaced, and approximate (31b) with N ordinary inequality constraints

$$p(e^{j\omega_k}) \leq \epsilon_{\text{sb}}^2 - \epsilon_N, \quad k = 1, 2, \dots, N, \quad (32)$$

where ϵ_N is chosen heuristically so that satisfaction of (32) guarantees satisfaction of (31b). In this way, the design problem can be approximated by a finite dimensional problem with linear and convex quadratic constraints—a problem which is efficiently solvable, using, for example, second-order cone programming techniques [41]. As the number of discretization points, N , is increased, ϵ_N can be reduced and hence the quality of the approximation improves. However, this also increases the number of constraints in the optimization problem, which may result in longer solution times and may expose the problem to numerical difficulties. A ‘rule of thumb’ is that $N \geq 15L$ will provide a sufficiently accurate approximation [30].

An elegant, precise and finite-dimensional alternative to discretizing these linear constraints is to transform them [31] into linear matrix inequalities (LMIs), which can be efficiently enforced using semidefinite programming (SDP) techniques [43]. In the case of (31g), we can exploit the Positive Real Lemma which states (e.g., [31]) that $p(e^{j\omega}) \geq 0, \forall \omega$ if and only if there exists a symmetric positive semidefinite matrix $\mathbf{X}_1 \in \mathbb{R}^{L \times L}$ such that

$$\sum_{k=0}^{L-1-n} [\mathbf{X}_1]_{k, k+n} = r[n], \quad n = 0, 1, \dots, L - 1, \quad (33)$$

where $[\cdot]_{ij}$ denotes the (i, j) th element of a matrix. (A symmetric matrix is said to be positive semidefinite if all its eigenvalues are non-negative. This will be denoted by $\mathbf{X}_1 \geq \mathbf{0}$.) Similarly, we can use the Bounded Real Lemma (e.g., [31]) to enforce (31e): $p(e^{j\omega}) \leq B^2, \forall \omega$ if and only if there exists a symmetric positive semidefinite matrix $\mathbf{X}_2 \in \mathbb{R}^{L \times L}$ such that

$$\sum_{k=0}^{L-1-n} [\mathbf{X}_2]_{k, k+n} = B^2 \delta[n] - r_p[n], \quad n = 0, 1, \dots, L - 1. \quad (34)$$

⁴For reasons of numerical accuracy, one may wish to replace (31c) by $(\sum_{i \geq 1} r_p[Mi]^2)^{1/2} \leq \epsilon_\gamma / \sqrt{2}$, which is a (convex) second-order cone constraint [41].

Equations (33) and (34) are finite sets of linear constraints on the semidefinite matrix \mathbf{X}_i , and are (exactly) equivalent to the infinite set of linear constraints on $r_p[n]$ generated by (31g) and (31e), respectively. Similarly, for the maximum stop-band level constraint, $r_p(e^{j\omega}) \leq \epsilon_{\text{sb}}^2, \forall \omega \in [\pi/K, \pi]$ if and only if [31] there exist positive semidefinite matrices $\mathbf{X}_3 \in \mathbb{R}^{L \times L}$ and $\mathbf{Z} \in \mathbb{R}^{(L-1) \times (L-1)}$, such that

$$(2 - \delta[n])x_{\text{sd}}[n] + q[n] = (\epsilon_{\text{sb}}^2 + r[0])\delta[n] - 2r_p[n], \quad (35)$$

where

$$x_{\text{sd}}[n] = \sum_{k=0}^{L-1-n} [\mathbf{X}_3]_{k,k+n}, \quad n = 0, \dots, L-1, \quad (36)$$

$$q[n] = \begin{cases} d_0 z_{\text{sd}}[0] + d_1 z_{\text{sd}}[1] & n = 0 \\ 2d_0 z_{\text{sd}}[n] + d_1 z_{\text{sd}}[n-1] + d_1 z_{\text{sd}}[n+1] & n = 2, \dots, L-3 \\ 2d_0 z_{\text{sd}}[L-2] + d_1 z_{\text{sd}}[L-3] & n = L-2 \\ d_1 z_{\text{sd}}[L-2] & n = L-1, \end{cases}, \quad (37)$$

$$z_{\text{sd}}[n] = \sum_{k=0}^{L-2-n} [\mathbf{Z}]_{k,k+n}, \quad n = 0, \dots, L-2, \quad (38)$$

$d_0 = 2(\cos(\pi/K) - \cos^2(\pi/K))$ and $d_1 = 2(\cos(\pi/K) - 1)$. This LMI approach has the advantage that the transformation to the finite problem is precise, and hence we avoid having to select N and ϵ_N in (32). The resulting optimization problems can be expressed as

$$\min_{r_p[n], \mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \mathbf{Z}} \sum_{n \geq 0} b_s[n] r_p[n] \quad (39a)$$

$$\text{subject to} \quad \mathbf{X}_1 \geq \mathbf{0}, \mathbf{X}_2 \geq \mathbf{0}, \mathbf{X}_3 \geq \mathbf{0}, \mathbf{Z} \geq \mathbf{0}, \quad (39b)$$

$$(33), (34), (35), (31c), (31d) \text{ and } (31f). \quad (39c)$$

The problem in (39) is a (convex) cone programme with a combination of linear, second-order and semidefinite cones and can be efficiently solved using general purpose solvers for such problems (e.g., [44]). Furthermore, there are some early indications (e.g., [45]) that the inherent structure in (39) can be exploited using specially designed algorithms that are substantially more efficient than general purpose methods.

5 Design Trade-offs

In this section we illustrate how our convex formulation (39) can be used to efficiently evaluate the trade-offs in prototype filter design for oversampled NPR GDFT filter banks. As discussed in the Introduction, an important trade-off is that between aliasing in the subbands and the distortion induced by the filter bank. In Section 3 we showed that the energy of the aliased components in the

subbands can be bounded by linear functions of the maximum stop-band level and the stop-band energy of the filter, and that these two quantities along with the maximum spectral component of the filter bound the energy of the aliased components at the output. In addition, we showed that the energy of the distortion error induced by the filter bank can be bounded by a multiple of the distortion coefficient, γ_p^2 [see (27)]. Therefore, a natural trade-off to explore is how the minimum achievable stop-band energy varies with the bound on the distortion coefficient, ϵ_γ^2 , for a given bound on the maximum stop-band level (and a fixed bound on the maximum spectral component and the transition-band energy). This trade-off can be efficiently obtained by solving (39) [or a discretized version of (31)] for different values of ϵ_γ^2 and fixed values of ϵ_{sb}^2 , B^2 and ϵ_{tbe} . The resulting trade-off curve generates considerable insight into the design of the prototype filter, as we illustrate in the following example.

Example 1 Consider a GDFT filter bank with $M = 8$ subbands and a down-sampling factor of $K = 6$. As in Section 3.3, we normalize the prototype to have energy $\sum_\ell p[\ell]^2 = K/M$. To make an appropriate choice for the upper bound on the maximum spectral component, B , we observe that if the stop-band suppression is substantial, then most of the filter’s energy lies in the pass band, $\omega \in [0, \pi/M]$. Using Parseval’s relation, in order for the prototype to have energy K/M , the average value of $|P(e^{j\omega})|^2$ over this band must be around K . Since $K = 6$ in this example, this is around 7.8 dB. To allow for some variation of the spectrum over the pass-band, we set $B^2 = 10^{0.1}K$; i.e., around 8.8 dB. The trade-off between the minimum achievable stop-band energy and the bound on the distortion coefficient for certain maximum allowable stop-band levels for filters of length 48 is given in Figure 3.⁵ (As discussed below, the constraint on the transition-band energy was not activated in this example.) These curves represent the *inherent* trade-off between the stop-band energy and the distortion coefficient because all points on or above (and to the right of) the curves can be achieved with a length 48 filter, and no length 48 filter can achieve any point below (and to the left of) the curves. Each point on these curves was found by solving (39) using the `SeDuMi` toolbox [44] for MATLAB. This task required between three and six seconds of CPU time on a 1.6 GHz Pentium IV workstation. The power spectra of three representative filters are given in Figure 4, along with the corresponding masks. From Figure 3 we can see that when considerable distortion is allowed, the maximum stop-band level constraints we have chosen are inactive. (The trade-off curves coalesce.) However, as the distortion constraint is made more stringent, the maximum stop-band level constraints become active. In fact, when the maximum stop-band level constraint is 34.5 or 36 dB below the maximum spectral component constraint, there is no length 48 filter which satisfies this constraint and has a vanishingly small distortion coefficient. (The achievable regions for these constraints, indicated by the dashed and dash-dot curves respectively, are bounded on the left.) Note, also, that a stop-band energy ‘floor’ is encountered when considerable distortion is allowed. This floor is due to the maximum spectral

⁵The filter length is chosen to be an integer multiple of $M/2$ for compatibility with the competing interpolation-based method described in Section 6. Our design method does not place any constraints on the filter length. Length 49 filters for this example were designed in [46].




Figure 3: Trade-off between the fraction of the total filter energy that is in the stop band, $E_{P, sb}/(K/M)$, and the (normalized) bound on the distortion coefficient, $\epsilon_\gamma^2/(K/M)$ for length 48 filters with maximum stop-band levels 0 dB (solid), 30 dB (dotted with Δ), 33 dB (dotted with \square), 34.5 dB (dashed), and 36 dB (dash-dot) below the maximum spectral component constraint. The symbols \circ , Δ and \square denote the trade-offs achieved by the filters in Figure 4(a), (b) and (c), respectively. The point denoted by the $+$ is used in Example 4.

component constraint (31e).

Given the proximity of π/M and π/K in this example, compatible constraints on the transition-band energy have only marginal effects on Figure 3. More precisely, the range of transition-band energy constraints which are active and result in the optimization problem having a non-empty feasible set is rather narrow. Hence, the corresponding curves have been omitted from Figure 3 for clarity. (They have also been omitted from the remaining figures in this section.) However, in Section 6 we will show that constraints on the transition-band energy have a considerable impact when the filter bank has a downsampling factor of $K = 4$ rather than $K = 6$. \square




Figure 4: Power spectra (in decibels) of filters which achieve the stop-band energy versus distortion coefficient trade-off curve in Figure 3 for a distortion coefficient of $10^{-6}K/M$ and maximum stop-band level constraints 0, 30, and 33 dB below the maximum spectral component constraint. For clarity the spectral mask imposed by the maximum spectral component and maximum stop-band level constraints is shown by the dashed line. In part (a) we have also indicated the stop-band edge, $1/(2K)$, with a dotted line. In parts (b) and (c) this edge is clear from the mask.

A distinct advantage of the design method developed in this paper is its flexibility. Many other design trade-offs can be explored using simple modifications to the design problem in (39). For example, we can determine how the minimum achievable maximum stop-band level varies with the bound on the distortion coefficient for a given bound on the stop-band energy by simply replacing the objective in (39) by ϵ_{sb}^2 and by introducing the additional linear constraint $\sum_{n \geq 0} b_s[n]r_p[n] \leq \epsilon_{\text{sbe}}$. The resulting problem remains convex and can be efficiently solved using similar techniques. This trade-off curve generates further insight into the design of the prototype filter, as we now show.

Example 2 Consider the GDFT filter bank in Example 1 with the same energy normalization and maximum spectral component constraint. The trade-off between the minimum achievable maximum stop-band level and the bound on the distortion coefficient for certain bounds on the stop-band energy is provided in Figure 5. Again, all points on or above the curves are achievable with a length 48 filter, and no length 48 filter can achieve any point below them. Each point was found by solving the above-mentioned modified version of (39) using `SeDuMi` toolbox [44] and was obtained in less than eight seconds. The power spectra of three representative filters are given in Figure 6, from which the differences in the stop-band structure are clear. From Figure 5 it is clear that when considerable distortion is allowed, the stop-band energy constraints that we have chosen become inactive. (The trade-off curves coalesce.) However, as the distortion constraint becomes more stringent, the choice of the bound on the stop-band energy has a considerable impact. In particular, there is no length 48 filter which has less than 0.02% of its energy in the stop band and a vanishingly small distortion coefficient. (The achievable regions indicated by the dotted curves are bounded on the left.) □




Figure 5: Trade-off between the maximum relative stop-band level (below the maximum spectral component constraint) and the (normalized) bound on the distortion coefficient, $\epsilon_\gamma^2/(K/M)$, for stop-band energies at most $\rho K/M$. The filters are of length 48 and $\rho = 1$ (solid), 5×10^{-4} (dashed), 3×10^{-4} (dash-dot), 2×10^{-4} (left/lower dotted curve), 1×10^{-4} (right/upper dotted curve). The symbols $*$, \triangleright and \diamond denote the trade-offs achieved by the filters in Figure 6(a), (b) and (c), respectively. For reference, the point denoted by the Δ is achieved by the same filter that achieves the point denoted by the Δ in Figure 3. This point will be used in Example 4.




Figure 6: Power spectra (in decibels) of filters which achieve the maximum stop-band level versus distortion coefficient trade-off curves in Figure 5 for a distortion coefficient of $10^{-6} K/M$ and stop-band energy constraints as indicated. [In (a) the stop-band energy constraint is inactive.] For clarity the spectral mask is shown by the dashed line.

Another interesting trade-off which can be efficiently evaluated using our method is that between the minimum achievable stop-band energy and the maximum allowable stop-band level for a given bound on the distortion coefficient. Curves of this type are often encountered in peak-constrained least-squares filter design [39], but usually without the distortion constraint. We can determine this trade-off without modifying (39). We simply solve (39) for a range of values of ϵ_{sb}^2 given fixed values of ϵ_{γ}^2 and B^2 . (As mentioned in Example 1, the transition-band energy constraint was activated in this section.) As we illustrate below, these trade-off curves provide additional insight into the design of the prototype.

Example 3 Once again, we consider the GDFT filter bank from Example 1, with the same energy normalization and maximum spectral component constraint. The inherent trade-off between the minimum achievable stop-band energy and the maximum allowable stop-band level for certain bounds on the distortion coefficient is provided in Figure 7. As is clear from the figure, relatively mild constraints on the distortion coefficient have a considerable impact on the inherent trade-off. (The achievable region bounded by the dotted-with- \star curve is substantially smaller than that bounded by the solid-without-symbols curve.) However, as the distortion constraint becomes more stringent, the impact of changes in this constraint is far lower. (The achievable regions bounded by the dotted-without-symbols and solid-with-symbols curves are almost the same.) Filters which achieve the trade-offs indicated by the symbols \star , \triangleleft , \blacksquare and \square in Figure 7 have the same maximum allowable stop-band level, and hence their gross spectral features will be similar at the scale of Figure 4. (The power spectrum of the filter which achieves the point denoted by \square in Figure 7 is given in Figure 4(c).) However, these four filters have substantially different distortion coefficients. The impact of these different distortion coefficients on the pass-band spectra of the filters is illustrated in Figure 8. The filters with smaller distortion coefficients have substantially flatter pass-band characteristics. In fact, the relationship between the distortion coefficient and the flatness of the pass-band response can be made explicit using a recent generalization of Nyquist’s first condition for intersymbol-interference free pulse transmission [34, Appendix A]. \square

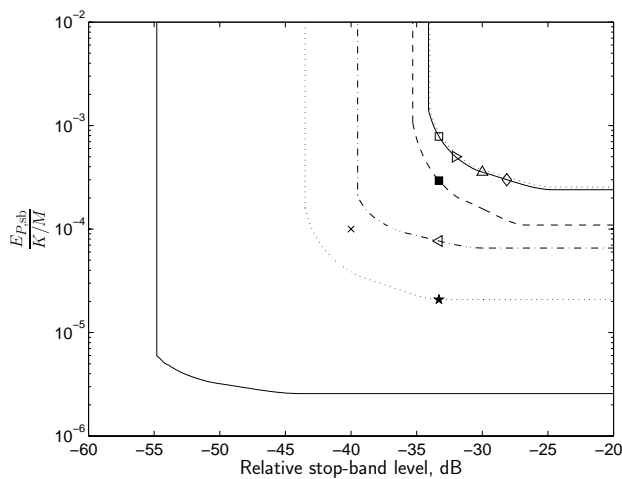


Figure 7: Trade-off between the (normalized) stop-band energy, $E_{P, sb}/(K/M)$, and the maximum relative stop-band level (below the maximum spectral component constraint) for distortion coefficients, γ_p^2 , at most $\alpha K/M$. The filters are of length 48 and $\alpha = 1$ (solid, no symbols), 10^{-2} (dotted with \star), 10^{-3} (dash-dot), 10^{-4} (dashed), 10^{-6} (solid with symbols), 10^{-8} (dotted, no symbol). The points denoted by the symbols \square and \triangle are achieved by filters which achieve the corresponding points in Figure 3. Similarly, the points denoted by the \triangleright and \diamond are achieved by filters which achieve the corresponding points in Figure 5. The symbols \star , \triangleleft , \blacksquare and \square denote the trade-offs achieved by the filters in Figure 8. The point denoted by the \times is used in Example 4.




Figure 8: Details of the pass-band power spectra (in decibels) for filters achieving the stop-band energy versus maximum stop-band level trade-off in Figure 7 for a stop-band level at least 33 dB below the maximum spectral component constraint and distortion coefficients of at most $\alpha K/M$, where $\alpha = 10^{-2}$ (dotted with \star), $\alpha = 10^{-3}$ (dash-dot with \triangleleft), $\alpha = 10^{-4}$ (dashed with \blacksquare), $\alpha = 10^{-6}$ (solid with \square).

An auxiliary problem that arises in filter bank design is the choice of the length of the prototype filter. Typically, for complexity reasons, we would like to use the shortest prototype which satisfies a given performance level; i.e., satisfies given bounds on the stop-band energy, the maximum stop-band level, the distortion coefficient, and the transition-band energy. Finding this minimum length prototype can be quite awkward using the non-convex formulation in (30) because it can be quite difficult to determine whether or not a given set of constraints generates a non-empty feasible set. In contrast, infeasibility of our convex formulation [(31) or (39)] can be reliably detected. Therefore, the shortest length which achieves a given specification can be efficiently determined using a bisection based search (on the filter length) for the feasibility boundary of a modified version of (31) [or (39)] in which the objective is removed and the linear constraint $\sum_{n \geq 0} b_s[n]r_p[n] \leq \epsilon_{sbe}$ is added. (An analogous procedure in a different application appears in [47].) We illustrate that procedure in the following example.

Example 4 In Examples 1–3 we examined the design trade-offs for prototype filters of length 48. Using the above-mentioned bisection search method, we found that the shortest filters which achieve the trade-off denoted by the + in Figure 3 for maximum stop-band levels 0, 30, 33, 34.5 and 36 dB below the maximum spectral component constraint are of lengths 35, 39, 43, 44 and 45, respectively. Similarly, we found that the shortest filters which achieve the trade-off denoted by the Δ in Figure 5 for stop-band energies at most $\rho K/M$, where $\rho = 1, 5 \times 10^{-4}, 3 \times 10^{-4}, 2 \times 10^{-4}$, and 1×10^{-4} are 44, 47, 49, 50 and 55, respectively. Notice that the last three cases require filters longer than 48. This is what we would expect, because the Δ is outside the achievable region for length 48 filters in these cases. Finally, the shortest filters which achieve the point denoted by the \times in Figure 7 for distortion coefficients of at most $\alpha K/M$, where $\alpha = 1, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-6}$ and 10^{-8} are 36, 44, 52, 54, 59, and 62, respectively. \square

As we have seen in this section, a feature of our design approach is that it is universal, in the sense that the trade-off curves can be obtained without knowledge of the application nor the input signal. However, selection of an appropriate point on the resulting trade-off surface depends on the relative importance of our three error mechanisms (aliasing in the subbands and the output, distortion, and imaging in the output) on the application at hand, and on the properties of the input signal. While the selection of that point will require a certain amount of application-specific empirical performance evaluation, the bounds derived in Section 3 provide some guidance as to how the properties of the input signal affect the relative importance of a small stop-band energy, a small maximum stop-band level, a small distortion coefficient, and a small transition-band energy. For example, the signal dependent coefficients of the bound on the aliasing in the subbands in (15) and the bound on the distortion in (27) are quite different. If these two error mechanisms are of equal importance in a given application, then for a given signal, the relative sizes of the signal dependent coefficients will guide the designer toward a point on the maximum stop-band level versus distortion coefficient trade-off curve that is appropriate for that signal.

6 Performance comparison

Now that we have provided examples of the inherent trade-offs in prototype filter design and have shown that our design method provides filters which achieve these trade-offs, we will compare our design method with two competing methods [19, 20]. Both methods are of interest because, like our method, they generate filter banks with nearly paraunitary polyphase matrices, and they tackle the distortion induced by the filter bank directly. In contrast, many other methods employ approximations of the distortion, such as considering only amplitude distortion [22].

The method proposed by Harteneck *et al.* [19] (see also [16]) involves minimizing a linear combination of the stop-band energy, $E_{P,\text{sb}}$, the distortion coefficient, γ_p^2 , and any normalization error, subject to the prototype filter having linear phase. In the notation of the present paper, this can be written as

$$\min_{\substack{p[\ell] \\ 0 \leq \ell \leq L-1}} \lambda E_{P,\text{sb}} + \gamma_p^2 + (\sum_{\ell} p[\ell]^2 - K/M)^2 \quad (40a)$$

$$\text{subject to } P(e^{j\omega}) \text{ having linear phase,} \quad (40b)$$

where $\lambda \geq 0$ is a chosen weighting. For symmetric filters, imposing phase linearity is equivalent to requiring $p[\ell] = p[L-1-\ell]$. (The alternative objective $\mu E_{P,\text{sb}} + (1-\mu)(\gamma_p^2 + (\sum_{\ell} p[\ell]^2 - K/M)^2)$ for some $\mu \in [0, 1]$ could also be used, but (40a) was chosen by Harteneck *et al.* [19].) The problem in (40) is not convex, and hence may require delicate management of local minima. However, given a linear-phase filter as a ‘starting point’, a local minimum can be found quite efficiently using an iterative least-squares (ILS) technique [27]. In our implementation we ran the ILS algorithm from several systematic starting points and several random starting points, and then chose the locally optimal filter with the lowest objective as our solution. Our systematic starting points included the length L truncation of the square-root raised cosine filter (e.g., [28, p. 496]) with cut-off frequency π/M and roll-off parameter $(M/K) - 1$.

Although our analysis (see Section 3) and some simulations (below) suggest that our standard formulation [(30), (31) or (39)] is more appropriate for many subband processing applications, we can obtain a convex problem similar to that in (40) by simply dropping the maximum stop-band level, transition-band energy, and maximum spectral component constraints [(31b),(31d) and (31e), respectively] from (31) and lifting the distortion constraint (31c) into the objective:

$$\min_{r_p[n], \tau_1, \tau_2} \lambda \left(\sum_{n \geq 0} b_s[n] r_p[n] \right) + \tau_1 + \tau_2 \quad (41a)$$

$$\text{subject to } \sum_{i \geq 1} r_p[Mi]^2 \leq \tau_1/2, \quad (r_p[0] - K/M)^2 \leq \tau_2 \quad \text{and} \quad p_p(e^{j\omega}) \geq 0, \forall \omega. \quad (41b)$$

As discussed in Section 4, the second constraint in (41b) can be handled via discretization or via transformation into an LMI. Notice, however, that we do not impose the phase linearity constraint in (41), or in any of our other formulations.

In contrast to our optimization-based method and that of Harteneck *et al.* [19], Liu *et al.* [20] proposed a simple, but *ad-hoc*, method based on interpolated filters [48, 49]. To obtain a ‘good’ prototype filter Liu *et al.* [20] suggest that one can simply interpolate the low-pass filter from an orthonormal two-channel filter bank by a factor of $M/2$. There are many well-known constructions of such filter banks [2–4, 50], and many choices for the interpolation algorithm. As suggested by Liu *et al.* [20] we will use MATLAB’s `interp` function to perform the interpolation. (While that choice does provide ‘good’ interpolation, the interpolation depends on the filter from the two-channel filter bank.) We will restrict attention to interpolations of low-pass filters from the standard Daubechies family [50] of orthonormal two-channel filter banks, the symlet family [50], and those obtained from ‘windowed’ ideal autocorrelations [51, 52]. We will choose the minimum phase spectral factor for the window-based designs, and will denote them by “MSB”. After evaluating the performance of the standard data analysis windows [53], we found that the ‘Blackmann-Harris’ window [53, p. 65] is an appropriate choice for the application at hand.

In the following example, we demonstrate how our formulations can provide prototype filters with significantly better design trade-offs than those generated by Harteneck’s method, and those generated by interpolation.

Example 5 In this example we again consider a filter bank with $M = 8$, $K = 6$, and length 48 filters. In Figure 9 we provide (with a dashed curve) the locus of normalized distortion coefficient–stop-band energy pairs $(\gamma_p^2/E_p, E_{P, \text{sb}}/E_p)$, achieved by Harteneck’s method [19] for different values of λ in (40). Here, $E_p = \sum_{\ell} p[\ell]^2$ is the energy of the filter. For clarity, we have indicated the points achieved by Harteneck’s method for specific values of λ by the symbols on the dashed curve in Figure 9. As one would expect, increasing λ moves the achieved trade-off towards the bottom-right corner of the figure. Figure 9 also contains the trade-off points achieved by interpolation of filters from certain standard orthonormal two-channel filter banks, and the inherent trade-off (achieved by our method) between the stop-band energy and the distortion coefficient (indicated by the solid curve). To obtain the inherent trade-off we solved an LMI version of the problem:

$$\min_{r_p[n]} \sum_{n \geq 0} b_s[n] r_p[n], \quad \text{subject to } \gamma_p^2 \leq \alpha K/M, r_p[0] = K/M \text{ and } p(e^{j\omega}) \geq 0, \forall \omega, \quad (42)$$

for a range of values of α . Over the range of distortion coefficients in Figure 9, the solid curve coincides with the solid curve in Figure 3. (Note that for clarity the scale of Figure 9 is slightly different from that in Figure 3.) However, when greater distortion is allowed, the solid curve in Figure 9 does not exhibit the floor effect seen on the right of Figure 3, because (42) does not contain a constraint on the maximum spectral component. For reference, we have also incorporated the dotted curves from Figure 3 into Figure 9. These curves correspond to the stop-band energy versus distortion trade-off subject to bounds on the maximum spectral component and the maximum relative stop-band level and were obtained by solving (39). The power spectra of representative filters are provided in Figure 10. In addition, the power spectra of filters which achieve the trade-offs denoted by the \circ , \triangle and \square in Figure 9 are given in Figure 4.




Figure 9: Trade-offs between the normalized stop-band energy, $E_{P, sb}/E_p$, and the normalized distortion coefficient, γ_p^2/E_p , for Example 5. Solid: the inherent trade-off [achieved by our method, (42)]; Dashed: trade-off achieved by Harteneck's method [19]; Dotted: trade-off achieved by our method with additional constraints on the maximum spectral component and the maximum relative stop-band level, (39). Dotted with \boxtimes , \triangle and \blacktriangle : relative stop-band level is -30 dB; Dotted with \boxplus , \square and \blacksquare : relative stop-band level is -33 dB. (The solid and dotted curves coincide with the corresponding curves in Figure 3.) The symbols \blacktriangleright , \blacktriangleleft , \blacktriangledown , \blacklozenge indicate the trade-offs achieved by Harteneck's method with $\lambda = 8.86 \times 10^{-3}$, 3.91×10^{-2} , 1.01, and 2.09, respectively. The symbols $+$, \times and $*$ denote the trade-offs achieved by interpolating length-12 Daubechies, symlet and MSB filters, respectively.

It is clear from Figure 9 that by optimizing the prototype and avoiding the phase linearity constraint, our method provides a significantly better trade-off between stop-band energy and distortion.⁶ In particular, the dotted curves in Figure 9 denote the stop-band-energy versus distortion coefficient trade-off achieved by filters which must also satisfy rather stringent spectral masks. Despite the fact that these filters satisfy this additional constraint, for a given distortion coefficient they (almost always) achieve a lower stop-band energy than the filter designed using Harteneck's method. \square

⁶Recall that our method actually achieves the *inherent* trade-off in the sense that no length 48 filter can achieve any point below the (solid) curve generated by our method.




Figure 10: Power spectra (in decibels) of filters which achieve the stop-band energy versus distortion coefficient trade-offs indicated by the ■, ▼ and × in Figure 9. For clarity the stop-band edge is indicated by the dotted line.

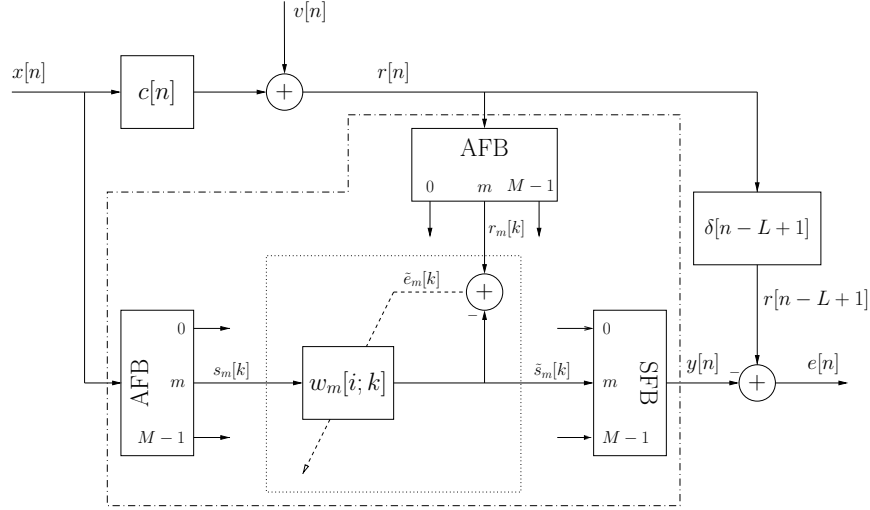


Figure 11: The subband adaptive filtering system (enclosed by the dash-dot polygon) used in Example 6. For clarity only one element of the (diagonal) subband processing block is shown (in the dotted box). The symbols AFB and SFB denote the analysis and synthesis filter banks, respectively, $c[n]$ is the impulse response of the unknown system and $w_m[i; k]$ is the impulse response of the adaptive filter in the m th subband at the k th instant.

To verify that the improved trade-offs achieved by our method in Example 5 (see Figure 9) can generate significant performance gains for the subband signal processing system as a whole, we now examine the performance of the simple subband adaptive filtering system illustrated in Figure 11. The purpose of the subband adaptive filter (bounded by the dash-dot polygon in Figure 11) is to approximate the impulse response of an unknown system, $c[n]$. This configuration is often called a ‘system identification’ configuration as is commonly used in acoustic echo cancellation (AEC) applications [13–17, 20–22]. The subband adaptive filter operates by passing the (known) input signal, $x[n]$, and noisy measurements, $r[n]$, of the output of the unknown system through separate analysis filter banks. (The sequence $v[n]$ denotes the noise.) The subband processing block is a diagonal matrix of standard adaptive filters, each of which operates (independently of the other filters) on one of the subband signals. (The elements which process the m th subband are enclosed in the dotted box in Figure 11, where the dashed line indicates that the error signal $e_m[k]$ drives the adaptation of $w_m[i; k]$.) The filtered subband signals are then reconstructed, and as the filter converges, $y[n]$ approximates a delayed version of $r[n]$. Although Figure 11 indicates that adaptive filtering is performed in every subband, if $x[n]$ and $c[n]$ are real, M is even, and if we insert a real part operator and a gain of two between the synthesis filter bank and the summer (i.e., if $e[n] = r[n - L + 1] - 2 \operatorname{Re}(y[n])$), then we can perform adaptive filtering on only $M/2$ of the subbands and obtain the same performance [16, 17, 19]. If we choose $m_0 = 1/2$ in (1) we need only filter subbands $0 \leq m \leq M/2 - 1$. The resulting reduction in computational load has a significant impact in practice, and we will use this reduced complexity system in the following simulations.

Example 6 In this example we examine the performance of the subband adaptive filter in Figure 11

equipped with the GDFT filter banks designed in Example 5 (see Figure 9) in a synthetic acoustic echo cancellation (AEC) environment. Recall that $M = 8$, $K = 6$ and $L = 48$ in Example 5. The prototype filters designed by Liu’s method are obtained directly, and those designed by Harteneck’s method are obtained directly from the output of the optimization routine. However, our method generates the optimal autocorrelation sequence, from which an optimal filter can be obtained using standard spectral factorization techniques [30, 35]. For simplicity, we only report results for the minimum phase spectral factor, but other experiments have indicated that in the scenario of this example, the performance variation over the different spectral factors tends to be small compared to the performance variations in Table 1. In synthesizing the filter banks from the prototype filter using (1), we chose $m_0 = 1/2$ so that we need only implement the adaptive filters in subbands $0 \leq m \leq 3$, and we chose $n_0 = -(L - 1)/2 = 24$, so that if $p[n]$ has linear phase (as it does in Harteneck’s designs) then all the filters $f_m[k]$ in the filter bank also have linear phase.

We evaluate the average performance of the subband adaptive filter over a class of randomly generated unknown systems $c[n]$ of length 60 whose impulse response coefficients tend to decay exponentially. More specifically, $c[n] \sim \mathcal{N}(0, e^{-n/10})$, $0 \leq n \leq 59$. This class of unknown systems shares many of the characteristics of the acoustic impulse response encountered in practical AEC applications. The adaptive filters, $w_m[i; k]$, have length 10 and are adapted using the normalized least-mean square (NLMS) algorithm [54], with step-size coefficient $\tilde{\mu} = 0.8$. The input signal, $x[n]$, is a (real) zero-mean white Gaussian signal of unit variance, and in order to isolate the performance of the filter bank, no noise was injected into the measured signal $r[n]$; i.e., $v[n] = 0$.

The mean square values of the error, $e[n]$, averaged over 10,000 realizations of the unknown system and the input signal, for the three prototype filters from Figure 10 (see Example 5) are provided in Figure 12. It is clear from that figure that the improved stop-band energy versus distortion trade-off achieved by our design method results in a substantial improvement (around 3.4 dB) in the steady-state error over the interpolation method, and a significant improvement (around 1.6 dB) over Harteneck’s method. To illustrate the influence of the properties of the prototype on the performance of the subband adaptive filter in greater detail, we have provided in Table 1 the steady-state mean square error (again averaged over 10,000 realizations of the unknown system and the input signal) for systems based on filters which achieve the marked trade-offs in Figure 9. For convenience we have listed in that table, the (normalized) distortion coefficient, the relative stop-band level (below the maximum spectral component constraint, $B^2 = 10^{0.1}K$), and the (normalized) stop-band and transition-band energies of each filter. (Note that when $\gamma_p^2/E_p = 10^{-3}$, the dotted-with- Δ curve and the solid curve in Figure 9 coincide and hence generate the same prototype filters. Therefore, only three filters are listed for this distortion coefficient.) From Table 1 it is clear that for a given distortion coefficient, the lower stop-band and transition-band energies achieved by our simplest formulation [which has no spectral mask, (42)] result in a significant improvement in the steady-state error over that achieved by the corresponding filter designed by Harteneck’s method. Table 1 also shows that our standard formulation [(31) or (39)] with a mild constraint on the maximum stop-band level results in further performance improvement. However,




Figure 12: Mean square error (averaged over 10,000 realizations) for the subband adaptive filtering system in Figure 11 using filter banks designed in Example 5; see also Figure 9. Legend—■: our method; ▼: Harteneck’s method; ×: interpolated symplet.

if the maximum stop-band level constraint is too stringent, and the resulting increase in the stop-band and transition-band energies is too large, then the steady-state error begins to increase. It can also be observed from Table 1 that as the distortion constraint is relaxed from $10^{-8}E_p$, the performance of the filters from each design method improves, but as this constraint becomes rather loose, the performance begins to degrade. Finally, Table 1 justifies our earlier decision (in Section 5) to leave the transition-band energy constraint inactive for this scenario. Even without a specific transition-band energy constraint, our method achieves a consistently lower transition-band energy than the competing methods. □

Table 1: Steady-state mean square error (SS-MSE) for the adaptive filtering system in Example 6 equipped with prototype filters from Example 5 (see Figure 9). Also included are the normalized distortion coefficient (γ_p^2/E_p), the maximum relative stop-band level below $B^2 = 10^{0.1}K$ (rel. SBL), and the normalized stop-band ($E_{P,\text{sb}}/E_p$) and transition-band ($E_{P,\text{tb}}/E_p$) energies.

Design Method	Symbol in Fig. 9	γ_p^2/E_p ,	rel. SBL, dB	$E_{P,\text{sb}}/E_p$, $\times 10^{-4}$	$E_{P,\text{tb}}/E_p$ $\times 10^{-2}$	SS-MSE, dB
Harteneck [19]	►	10^{-8}	-20.0	17.4	6.24	-20.02
Ours, no mask, (42)	⊕	10^{-8}	-24.5	2.54	5.62	-23.05
Ours, (39)	⊠	10^{-8}	-30.0	3.76	5.26	-23.42
Ours, (39)	⊞	10^{-8}	-33.0	7.44	5.02	-22.74
Harteneck [19]	◄	10^{-6}	-20.6	9.06	6.20	-21.29
Ours, no mask, (42)	○	10^{-6}	-24.6	2.40	5.57	-23.11
Ours, (39)	△	10^{-6}	-30.0	3.55	5.23	-23.46
Ours, (39)	□	10^{-6}	-33.0	6.86	5.00	-22.88
Harteneck [19]	▼	10^{-4}	-22.01	4.66	5.88	-22.14
Ours, no mask, (42)	●	10^{-4}	-26.6	1.09	5.06	-23.75
Ours, (39)	▲	10^{-4}	-30.0	1.57	4.91	-23.97
Ours, (39)	■	10^{-4}	-33.0	2.68	4.78	-23.91
Harteneck [19]	◆	10^{-3}	-23.6	2.82	5.21	-21.89
Ours, no mask, (42)	⊗	10^{-3}	-30.0	0.65	4.27	-23.07
Ours, (39)		10^{-3}	-33.0	0.75	4.12	-23.27
Interp. [20], MSB [51, 52]	*	2.08×10^{-3}	-14.3	82.8	8.45	-15.52
Interp. [20], Daubech. [50]	+	1.02×10^{-4}	-16.9	26.6	6.96	-19.66
Interp. [20], symlet. [50]	×	2.14×10^{-5}	-15.9	21.0	7.04	-20.53

Table 2: Prototype filter characteristics and steady-state mean square error (SS-MSE) for the adaptive filtering system in Example 6 equipped with length 200 prototype filters

Design Method	γ_p^2/E_p ,	rel. SBL, dB	$E_{P,sb}/E_p$,	$E_{P,tb}/E_p$ $\times 10^{-2}$	SS-MSE, dB
Interp. [20], MSB [51, 52]	1.46×10^{-4}	-36.5	5.32×10^{-4}	1.73	-26.30
Interp. [20], symlet [50]	1.33×10^{-7}	-42.3	1.76×10^{-6}	3.58	-27.14
Ours, (39) [†]	1.51×10^{-9}	-114	2.26×10^{-13}	2.43	-29.68
Harteneck [19], $\lambda = 6.43^\ddagger$	1.51×10^{-9}	-88.0	1.88×10^{-11}	2.97	-28.63

[†] The constraints in (39), including that on the transition-band energy, were chosen to be the values achieved by the interpolated symlet.

[‡] The value of λ was chosen so that the same distortion as our method was achieved.

Although the steady-state mean square errors of the systems in Table 1 (see also Figure 12) might be regarded as being rather high, and although the performance differences are quite subtle in places, Example 6 has validated the major principles of our design approach. First, the performance of a GDFT-filter-bank-based subband adaptive filtering system depends on the properties of the prototype filter that we derived in Section 3, namely, the stop-band energy, the maximum stop-band level, the distortion coefficient, and the transition-band energy; and second, to obtain optimized performance from the subband adaptive filter in a particular application, we should explore the trade-offs between these quantities. Our standard formulation [(31) or (39)] provides an efficient method for evaluating these trade-offs and hence should be a convenient tool for system designers.

The performance of the system in Example 6 can be significantly improved by using longer prototype filters or by reducing the downsampling factor K , although both actions will increase the implementation complexity of the system. To indicate the extent of the potential improvements, Table 2 provides partial results for the scenario of Example 6 with length 200 filters, and in the following example we consider the same scenario with length 48 filters and a downsampling factor of $K = 4$.

Example 7 We consider the subband adaptive filtering system from Example 6, but with a downsampling factor of $K = 4$. The trade-offs between the (normalized) stop-band energy and distortion coefficient achieved by the three design methods are illustrated in Figure 13. For our method, we set the maximum spectral component constraint in the same was as in Example 1, and the maximum stop-band level constraint was chosen to be that achieved by the interpolated symlet. (This constraint is inactive in the figure, as suggested by Table 3.) The focus of this example is the effects of the transition-band energy on the system performance. The solid, dotted and dash-dot curves in Figure 13 indicate the inherent trade-offs (achieved by our method) under different constraints on the transition-band energy. It is clear from Figure 13 that for the filter bank considered in this example different constraints on the transition-band energy result in significantly different trade-offs between the stop-band energy and the distortion. In particular, Harteneck’s method [19], which does not explicitly control the transition-band energy, achieves a trade-off (indicated by the

Table 3: Prototype filter characteristics and steady-state mean square error (SS-MSE) for the adaptive filtering system in Example 7.

Design Method	Symbol in Fig. 13	γ_p^2/E_p ,	rel. SBL, dB	$E_{P,sb}/E_p$,	$E_{P,tb}/E_p$ $\times 10^{-2}$	SS-MSE, dB
Interp. [20], MSB [51, 52]	*	1.38×10^{-3}	-26.3	4.13×10^{-3}	8.86	-20.92
Interp. [20], Daubech. [50]	+	6.77×10^{-5}	-32.7	8.31×10^{-4}	7.14	-29.36
Interp. [20], symlet [50]	\times	1.42×10^{-5}	-43.6	3.62×10^{-5}	7.24	-38.83
Harteneck [19], $\lambda = 99.7$	\blacktriangledown	1.42×10^{-5}	-78.6	5.93×10^{-10}	8.86	-38.97
Ours, (39)	\bullet	1.42×10^{-5}	-85.7	1.27×10^{-10}	8.18	-39.58
Ours, (39)	\blacksquare	1.42×10^{-5}	-71.7	7.85×10^{-9}	7.24	-40.93
Ours, (39)	\blacktriangle	1.42×10^{-5}	-61.2	7.25×10^{-8}	5.79	-39.40
Harteneck [19], $\lambda = 0.642$	∇	10^{-8}	-66.5	9.19×10^{-9}	10.0	-38.76
Ours, (39)	\circ	10^{-8}	-83.2	3.12×10^{-10}	8.81	-40.35
Ours, (39)	\square	10^{-8}	-60.1	9.52×10^{-8}	7.24	-42.05
Ours, (39)	Δ	10^{-8}	-58.2	1.18×10^{-6}	5.79	-37.44

dashed curve) that is superior to that which can be achieved by any design method which imposes the transition-band energy constraints that we have chosen. (Recall that our method achieves the *inherent* trade-off for the given constraints, indicated by the dotted and dash-dot curves.) However, as illustrated in Table 3 and Figure 14, filters designed via Harteneck’s method may have rather large transition-band energies, and hence improved state-state mean square error (SS-MSE) adaptive filtering performance can be obtained using filters designed by our method. (Since $K = 4$ in this example, we have used length 15 adaptive filters in the subbands in place of the length 10 adaptive filters in Example 6.) In particular, Table 3 indicates that for the same distortion coefficient as the interpolated symlet, the filter designed by Harteneck’s method (symbol \blacktriangledown) performs only marginally better than the interpolated symlet (symbol \times). The filter designed by our method without constraints on the transition-band energy (symbol \bullet) performs a little better (a gain of 0.7 dB in the SS-MSE). Constraining the transition-band energy of our filter to be less than or equal to that of the interpolated symlet improves performance by more than 1.3 dB (for a total gain over the interpolated symlet and Harteneck’s method of around 2 dB, symbol \blacksquare), but over constraining the transition-band energy results in a degradation in performance (symbol \blacktriangle).

If the required normalized distortion coefficient is reduced to 10^{-8} , the same trends apply, but the performance of our filters improves and their performance advantages increase. (The performance of Harteneck’s filters actually degrades; symbol ∇ .) In particular, the filter indicated by the symbol \square has a SS-MSE which is around 3.2 dB lower than that achieved by the interpolated symlet (symbol \times), and 3.3 dB lower than that achieved by the corresponding filter designed by Harteneck’s method (symbol ∇). The convergence of the (averaged) output MSE of the subband adaptive filtering system for these and other representative filters is provided in Figure 14. \square




Figure 13: Trade-offs between the normalized stop-band energy, $E_{P, sb}/E_p$, and the normalized distortion coefficient, γ_p^2/E_p , for Example 7. Legend—Solid: our method, without a constraint on the transition-band energy (TBE); Dashed: Harteneck’s method [19]; Other curves: Our method with the TBE constrained to be less than ρ times that achieved by the interpolated symlet; Dotted: $\rho = 1$; Dash-dot: $\rho = 0.8$. The symbols $+$, \times and $*$ denote the trade-offs achieved by interpolated length-12 Daubechies, symlet and MSB filters, respectively. The characteristics and SS-MSE adaptive filtering performance of the filters which achieve the trade-offs indicated by the symbols are provided in Table 3.




Figure 14: Output mean square error (averaged over 1000 realizations) for Example 7. The symbols correspond to those in Table 3 and Figure 13. The two curves without symbols correspond to Harteneck’s method (marginally upper curve, symbol ∇) and the interpolated symlet (marginally lower curve, symbol \times).

7 Conclusions

In this paper we have derived explicit bounds on the aliasing in the subbands and the output, the distortion, and the imaged subband errors of an oversampled near perfect reconstruction (and near paraunitary) GDFT filter bank. We have shown that the design of a prototype filter which optimizes these bounds can be formulated as a convex optimization problem from which a globally optimal filter can be efficiently obtained. The key to developing the convex formulation was to show that the performance objectives can be written as linear and convex quadratic functions of the autocorrelation sequence of the filter. We showed that the convex formulation allows efficient and accurate exploration of the inherent trade-offs in the design, and generates filters with significantly improved performance over two current techniques. Since the intersection of convex sets is itself convex, our convex formulation is quite flexible and many other performance objectives can be incorporated into the design; e.g., a lower bound on the power spectrum of the filter in the pass band [55] (see also [30, 31, 34, 47]). Furthermore, it appears that our formulation may also be applicable to the design of windows for oversampled short-term Fourier Transforms which form a ‘snug’ frame [56], and the design of filtered multitone modulation schemes [57].

A Appendix: Derivation of (29)

Given (24),

$$E_D \leq E_x \max_{\omega} \left| \frac{1}{K} \sum_{m=0}^{M-1} |P(e^{j(\omega-\alpha_m)})|^2 - 1 \right|^2. \quad (43)$$

Now,

$$\begin{aligned} \frac{1}{K} \sum_{m=0}^{M-1} |P(e^{j(\omega-\alpha_m)})|^2 &= \frac{1}{K} \sum_{k,\ell} p[k]p[\ell] e^{-j(\omega-m_0/M)(k-\ell)} \sum_{m=0}^{M-1} e^{j2\pi m(k-\ell)/M} \\ &= \frac{M}{K} \sum_n \left(\sum_{\ell} p[\ell]p[\ell + Mn] \right) e^{-j(\omega M - 2\pi m_0)n}, \end{aligned} \quad (44)$$

where we have used the Poisson sum formula. Hence,

$$\begin{aligned} \max_{\omega} \left| \frac{1}{K} \sum_{m=0}^{M-1} |P(e^{j(\omega-\alpha_m)})|^2 - 1 \right|^2 &\leq \frac{M}{K} \left| \sum_n \left(\sum_{\ell} p[\ell]p[\ell + Mn] \right) - \delta[n] \right|^2 \\ &\leq \frac{M}{K} \left(\sum_n \left| \sum_{\ell} p[\ell]p[\ell + Mn] - \delta[n] \right| \right)^2. \end{aligned} \quad (45)$$

Equation (29) then follows from the normalization of the prototype filter, $\sum_{\ell} p[\ell]^2 = K/M$, and the definition of $\tilde{\gamma}_p$.

References

- [1] R. E. Crochiere and L. R. Rabiner, *Multirate Digital Signal Processing*. Englewood Cliffs, NJ: Prentice Hall, 1983.
- [2] P. Vaidyanathan, *Multirate Systems and Filter Banks*. Englewood Cliffs, NJ: Prentice Hall, 1993.
- [3] M. Vetterli and J. Kovačević, *Wavelets and Subband Coding*. Englewood Cliffs, NJ: Prentice Hall, 1995.
- [4] G. Strang and T. Nguyen, *Wavelets and Filter Banks*. Wellesley, MA: Wellesley-Cambridge Press, 1997.
- [5] Z. Cvetković and M. Vetterli, “Oversampled filter banks,” *IEEE Transactions on Signal Processing*, vol. 46, no. 5, pp. 1245–1255, 1998.
- [6] H. Bölcskei, F. Hlawatsch, and H. G. Feichtinger, “Frame-theoretic analysis of oversampled filter banks,” *IEEE Transactions on Signal Processing*, vol. 46, no. 12, pp. 3256–3268, 1998.
- [7] P. N. Heller, T. Karp, and T. Q. Nguyen, “A general formulation of modulated filter banks,” *IEEE Transactions on Signal Processing*, vol. 47, no. 4, pp. 986–1002, 1999.
- [8] Z. Cvetković and M. Vetterli, “Tight Weyl-Heisenberg frames in $\ell^2(\mathbb{Z})$,” *IEEE Transactions on Signal Processing*, vol. 46, no. 5, pp. 1256–1259, 1998.
- [9] J. Kliewer and A. Mertins, “Oversampled cosine-modulated filter banks with arbitrary system delay,” *IEEE Transactions on Signal Processing*, vol. 46, no. 4, pp. 941–955, 1998.
- [10] H. Bölcskei and F. Hlawatsch, “Oversampled cosine modulated filter banks with perfect reconstruction,” *IEEE Transactions on Circuits and Systems—II: Analog and Digital Signal Processing*, vol. 45, no. 8, pp. 1057–1071, 1998.
- [11] M. K. Mihçak, P. Moulin, M. Anitescu, and K. Ramchandran, “Rate-distortion-optimal subband coding without perfect reconstruction constraints,” *IEEE Transactions on Signal Processing*, vol. 49, no. 3, pp. 542–557, 2001.
- [12] J. J. Shynk, “Frequency-domain and multirate adaptive filtering,” *IEEE Signal Processing Magazine*, pp. 15–37, Jan. 1992.
- [13] A. Gilloire and M. Vetterli, “Adaptive filtering in subbands: Analysis, experiments, and application to acoustic echo cancellation,” *IEEE Transactions on Signal Processing*, vol. 40, no. 8, pp. 1862–1875, 1992.

- [14] Q. Jin, Z. Q. Luo, and K. M. Wong, "Optimum filter banks for signal decomposition and its application in adaptive echo cancelation," *IEEE Transactions on Signal Processing*, vol. 44, no. 7, pp. 1669–1680, 1996.
- [15] B. Farhang-Boroujeny and Z. Wang, "Adaptive filtering in subbands: Design issues and experimental results for acoustic echo cancellation," *Signal Processing*, vol. 61, pp. 213–223, 1997.
- [16] S. Weiss, A. Stenger, R. W. Stewart, and R. Rabenstein, "Steady-state performance limitations of subband adaptive filters," *IEEE Transactions on Signal Processing*, vol. 49, no. 9, pp. 1982–1991, 2001.
- [17] J. P. Reilly, M. Wilbur, M. Seibert, and N. Ahmadvand, "The complex subband decomposition and its application to the decimation of large adaptive filtering problems," *IEEE Transactions on Signal Processing*, vol. 50, no. 11, pp. 2730–2743, 2002.
- [18] M. Harteneck, J. M. P. Borrillo, and R. W. Stewart, "An oversampled subband adaptive filter without cross adaptive filters," *Signal Processing*, vol. 64, no. 1, pp. 93–101, 1998.
- [19] M. Harteneck, S. Weiss, and R. W. Stewart, "Design of near perfect reconstruction oversampled filter banks for subband adaptive filters," *IEEE Transactions on Circuits and Systems–II: Analog and Digital Signal Processing*, vol. 46, no. 8, pp. 1081–1085, 1999.
- [20] Q.-G. Liu, B. Champagne, and D. K. C. Ho, "Simple design of oversampled uniform DFT filter banks with applications to subband acoustic echo cancellation," *Signal Processing*, vol. 80, no. 5, pp. 831–847, 2000.
- [21] K. Eneman and M. Moonen, "Hybrid subband/frequency-domain adaptive systems," *Signal Processing*, vol. 81, pp. 117–136, 2001.
- [22] —, "DFT modulated filter bank design for oversampled subband systems," *Signal Processing*, vol. 81, pp. 1947–1973, 2001.
- [23] M. R. Petraglia and S. K. Mitra, "Performance analysis of adaptive filter structures based on subband decomposition," in *Proceedings of the International Symposium on Circuits and Systems*, vol. I, Chicago, IL, 1993, pp. 60–63.
- [24] R. Horst and P. M. Pardalos, Eds., *Handbook of Global Optimization*. Dordrecht, The Netherlands: Kluwer, 1995.
- [25] M. Vetterli, "Filter banks allowing perfect reconstruction," *Signal Processing*, vol. 10, pp. 219–244, 1986.
- [26] T. Q. Nguyen, "Digital filter bank design: Quadratic-constrained formulation," *IEEE Transactions on Signal Processing*, vol. 43, no. 9, pp. 2103–2105, 1995.

- [27] H. Xu, W.-S. Lu, and A. Antoniou, “Efficient iterative design method for cosine-modulated QMF banks,” *IEEE Transactions on Signal Processing*, vol. 44, no. 7, pp. 1657–1668, 1996.
- [28] M. C. Jeruchim, P. Balaban, and K. S. Shanmugan, *Simulation of Communication Systems*, 2nd ed. New York: Kluwer, 2000.
- [29] R. D. Koilpillai and P. P. Vaidyanathan, “Cosine-modulated FIR filter banks satisfying perfect reconstruction,” *IEEE Transactions on Signal Processing*, vol. 40, no. 4, pp. 770–783, 1992.
- [30] S.-P. Wu, S. Boyd, and L. Vandenberghe, “FIR filter design via spectral factorization and convex optimization,” in *Applied Computational Control, Signal and Communications*, B. Datta, Ed. Boston, MA: Birkhäuser, 1997, also available from <http://www-isl.stanford.edu/people/boyd>.
- [31] T. N. Davidson, Z.-Q. Luo, and J. F. Sturm, “Linear matrix inequality formulation of spectral mask constraints with applications to FIR filter design,” *IEEE Transactions on Signal Processing*, vol. 50, no. 11, pp. 2702–2715, 2002.
- [32] P. Moulin, M. Anitescu, K. O. Kortanek, and F. A. Potra, “The role of linear semi-infinite programming in signal adapted QMF bank design,” *IEEE Transactions on Signal Processing*, vol. 45, no. 9, pp. 2160–2174, 1997.
- [33] J. Tuqan and P. P. Vaidyanathan, “A state space approach to the design of globally optimal FIR energy compaction filters,” *IEEE Transactions on Signal Processing*, vol. 48, no. 10, pp. 2822–2838, 2000.
- [34] T. N. Davidson, “Efficient design of waveforms for robust pulse amplitude modulation,” *IEEE Transactions on Signal Processing*, vol. 49, no. 12, pp. 3098–3111, 2001.
- [35] T. N. T. Goodman, C. A. Micchelli, G. Rodriguez, and S. Seatzu, “Spectral factorization of Laurent polynomials,” *Advances in Computational Mathematics*, vol. 7, no. 4, pp. 429–454, 1997.
- [36] R. D. Koilpillai and P. P. Vaidyanathan, “A spectral factorization approach to psuedo-QMF design,” *IEEE Transactions on Signal Processing*, vol. 41, no. 1, pp. 82–92, 1993.
- [37] W. A. Gardner, *Statistical Spectral Analysis. A Nonprobabilistic Theory*. Englewood Cliffs, NJ: Prentice Hall, 1988.
- [38] W. Rudin, *Real and Complex Analysis*. New York, NY: McGraw-Hill, 1987.
- [39] J. W. Adams and J. L. Sullivan, “Peak-constrained least-squares optimization,” *IEEE Transactions on Signal Processing*, vol. 46, no. 2, pp. 306–321, 1998.
- [40] R. Hettich and K. O. Kortanek, “Semi-infinite programming: Theory, methods and applications,” *SIAM Review*, vol. 35, no. 3, pp. 380–429, 1993.

- [41] M. Lobo, L. Vandenberghe, S. Boyd, and H. Le Bret, “Applications of second-order cone programming,” *Linear Algebra and its Applications*, vol. 284, pp. 193–228, 1998.
- [42] A. Papoulis, *Signal Analysis*. New York, NY: McGraw-Hill, 1977.
- [43] L. Vandenberghe and S. Boyd, “Semidefinite programming,” *SIAM Review*, vol. 38, 1996.
- [44] J. F. Sturm, “Using SeDuMi 1.02, A Matlab toolbox for optimization over symmetric cones,” *Optimization Methods and Software*, vol. 11-12, pp. 625–653, 1999, see <http://www.fewcal.nl/sturm/software/sedumi.html> for updates.
- [45] B. Alkire and L. Vandenberghe, “Interior-point methods for magnitude filter design,” in *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, Salt Lake City, UT, May 2001, pp. 3821–3824.
- [46] M. R. Wilbur, T. N. Davidson, and J. P. Reilly, “Efficient design of oversampled NPR GDFT filter banks,” in *Proceedings of the 2003 International Conference on Acoustics, Speech and Signal Processing*, vol. VI, Hong Kong, Apr. 2003, pp. 489–492.
- [47] T. N. Davidson, Z.-Q. Luo, and K. M. Wong, “Design of orthogonal pulse shapes for communications via semidefinite programming,” *IEEE Transactions on Signal Processing*, vol. 48, no. 5, pp. 1433–1445, 2000.
- [48] Y. Neuvo, C. Y. Dong, and S. K. Mitra, “Interpolated finite impulse response filters,” *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. ASSP–32, no. 6, pp. 563–570, 1984.
- [49] R. Lyons, “Interpolated narrowband lowpass FIR filters,” *IEEE Signal Processing Magazine*, vol. 20, no. 1, pp. 50–57, 2003.
- [50] I. Daubechies, *Ten Lectures on Wavelets*. Philadelphia, PA: SIAM, 1992.
- [51] F. Mintzer, “Filters for distortion-free two-band multirate filter banks,” *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. ASSP–33, no. 3, pp. 626–630, 1985.
- [52] M. J. T. Smith and T. P. Barnwell, III, “Exact reconstruction techniques for tree-structured subband coders,” *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. ASSP–34, no. 3, pp. 434–441, 1986.
- [53] F. J. Harris, “On the use of windows for harmonic analysis with the discrete Fourier transform,” *Proceedings of the IEEE*, vol. 66, no. 1, pp. 51–83, 1978.
- [54] S. Haykin, *Adaptive Filter Theory*, 4th ed. Upper Saddle River, NJ: Prentice Hall, 2002.
- [55] M. R. Wilbur, “The decomposition of large blind equalization problems using GDFT filter banks,” Master’s thesis, McMaster University, 2000.

- [56] Z. Cvetković, “On discrete short-time Fourier analysis,” *IEEE Transactions on Signal Processing*, vol. 48, no. 9, pp. 2628–2640, 2000.
- [57] G. Cherubini, E. Eleftheriou, and S. Olcer, “Filtered multitone modulation for very high-speed digital subscriber lines,” *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 5, pp. 1016–1028, 2002.