# EFFICIENT DESIGN OF OVERSAMPLED NPR GDFT FILTER BANKS

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

Department of Electrical and Computer Engineering,
McMaster University, Hamilton, Ontario, Canada

**ABSTRACT**

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 that are sensitive to aliasing; e.g., subband adaptive filters. Our design criteria for the prototype filter are explicit bounds on the aliased components in the subbands, the aliased components in the output, and the distortion induced by the filter bank. It is shown that the design of an optimal prototype filter can be transformed into a convex optimization problem that 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.

## 1. INTRODUCTION

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]. While many classes of perfect reconstruction filter banks are available, 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., [2]. Designs based on the perfect reconstruction condition typically allow considerable aliasing in the subband signals, 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 processing applied to the subband signals is sensitive to aliasing, or if that processing distorts the aliased components in a way that reduces the effectiveness of alias cancellation.

For applications which are quite sensitive to aliasing, such as subband adaptive filtering, oversampled near perfect reconstruction (NPR) filter banks that suppress aliased components, rather than relying on alias cancellation, offer the potential for improved performance. In particular, the performance of systems based on oversampled generalized Discrete Fourier Transform (GDFT) filter banks is quite encouraging [3–6]. 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 in [7]) on the aliased components in the subbands, and the reconstruction error. These bounds rigorously amalgamate several intuitively developed design criteria in the current literature [3,4], and subsume the criteria derived in [5, 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, 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, delicate management of locally optimal solutions may be required in order to obtain a filter whose performance is "good enough". 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 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. (Similar transformations have also led to convex formulations of some other FIR filter design problems; e.g., [8, 9].) 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 demonstrate the value of these trade-off curves in filter design, and show that filters designed via our formulation can provide significantly better performance than filters designed using current methods.

## 2. GDFT FILTER BANKS AND DESIGN CRITERIA

The generalized Discrete Fourier Transform (GDFT) filter bank we will consider has $M$ subbands, each of which is downsampled by $K < M$. The analysis and synthesis filters for the $m$th subband, $f_m[n]$ and $g_m[n]$ respectively, are exponentially modulated versions of a single real-valued FIR prototype filter $p[n]$:

$$f_m[n] = p[n]e^{j2\pi(m+m_0)(n+n_0)/M}, \quad 0 \le m \le M-1, \quad (1)$$

and $g_m[n] = f_m[-n]^*$. This modulated structure facilitates efficient implementation. Since we employ the same prototype filter in the analysis and synthesis banks, our NPR filter banks have polyphase matrices [1] which are nearly paraunitary, and hence have favourable noise robustness and numerical properties.

To describe the operation of the filter bank we let $X(z) = \sum_n x[n]z^{-n}$ denote the $z$-transform of $x[n]$, the input signal. The $m$-th subband signal has a $z$-transform

$$S_m(z) = \frac{1}{K}\sum_{k=0}^{K-1} F_m(z^{1/K}W_K^k)X(z^{1/K}W_K^k), \quad (2)$$

where $W_K = e^{-j2\pi/K}$. The $k = 0$ term on the right hand side of (2) represents the desired component of the subband signal, while

the remaining terms represent the aliased components. In subband adaptive filtering applications, the subband signals are typically filtered individually. If the filter in the $m$th subband has converged and is denoted by $H_m(z)$, then the output of the whole subband adaptive filtering system is

$$Y(z) = \frac{1}{K} \sum_{k=0}^{K-1} \sum_{m=0}^{M-1} G_m(z) H_m(z^K) F_m(z W_K^k) X(z W_K^k). \tag{3}$$

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]. The transfer function $\sum_{m=0}^{M-1} G_m(z) H_m(z^K) F_m(z)$ in (3) distorts the input signal $X(z)$, while the remaining terms represent the aliased components which appear at the output.

In the full version of this paper [7], we showed that the aliased components in (2) and (3) and the undesired distortion induced by the filter bank, can be bounded by natural properties of the prototype filter. To state those bounds succinctly, let $E_x = \sum_n |x[n]|^2$ denote the energy of a signal $x[n]$, and define

$$J(\mu) \triangleq \frac{1}{2\pi} \int_{-\pi}^{\pi} \left| X(e^{j\omega}) X^*(e^{j(\omega+\mu)}) \right| d\omega,$$

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

Let $U_X \triangleq \max_\omega \ |X(e^{j\omega})|$ denote the maximal spectral component $x[n]$. Let $U_{P,\text{sb}} \triangleq \max_{\omega \in [\pi/K, \pi]} \ |P(e^{j\omega})|$ denote the maximum stop-band level of the prototype filter, and let $E_{P,\text{sb}} = (1/\pi) \int_{\pi/K}^{\pi} |P(e^{j\theta})|^2 \, d\theta$ denote its stop-band energy. Finally, let $\alpha_m = 2\pi(m + m_0)/M$ and $\psi_k = 2\pi k/K$. Using various forms of the Hölder inequality, the energy of the aliased components in (2), denoted by $E_{A_m}$, can be bounded by [7]

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

The energy of the aliased components in the output, $E_{A_{\text{out}}}$, can also be bounded by simple functions of the stop-band energy and the maximum stop-band level [7]. If the subband adaptive filters are scaled so that $\max_\omega \left| H_m(e^{j\omega}) \right| \leq 1$, then [7]

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

In order to isolate the (undesirable) distortion induced by the filter bank from the (desired) processing performed by the subband processing system, we will analyze the distortion in the absence of any subband processing. That is, with $H_m(z) = 1$ in (3). If we normalize the energy of the prototype filter so that $\sum_\ell p[\ell]^2 = K/M$, then the energy of the distortion induced by the filter bank, $E_D$, can be bounded by [7]

$$E_D \leq \gamma_p^2 \ \frac{M U_X^2}{K}, \tag{4}$$

where $\gamma_p^2 = \sum_{n \neq 0} \left| \sum_\ell p[\ell] p[\ell - nM] \right|^2$, measures the distance from the prototype filter to the nearest "self-orthogonal" filter.

## 3. DESIGN METHOD

Using the bounds in Section 2, 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 $\gamma_p$ will guarantee that the energy of the (amplitude and phase) distortion in the output is small.

Although various combinations of some of these criteria have been employed by other authors (on a somewhat ad-hoc basis), we have shown (in [7]) how they explicitly bound the energy of the aliased components in the subbands and the output, and the energy of the distortion in the output. Natural design criteria for the prototype can be obtained by minimizing a (linear) combination of $U_{P,\text{sb}}$, $E_{P,\text{sb}}$, $U_P$ and $\gamma_p$, 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, 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:

$$\min_{p[\ell]} \quad \frac{1}{\pi} \int_{\pi/K}^{\pi} \left| P(e^{j\omega}) \right|^2 d\omega \tag{5a}$$

$$\text{subject to} \quad \left| P(e^{j\omega}) \right| \leq \epsilon_{\text{sb}} \quad \forall \omega \in [\pi/K, \pi], \tag{5b}$$

$$\sum_{n \neq 0} \left| \sum_\ell p[\ell] p[\ell - Mn] \right|^2 \leq \epsilon_\gamma^2, \tag{5c}$$

$$\left| P(e^{j\omega}) \right| \leq B \quad \forall \omega, \tag{5d}$$

$$\sum_\ell p[\ell]^2 = K/M, \tag{5e}$$

where $\epsilon_{\text{sb}}$, $B$ and $\epsilon_\gamma^2$ are fixed constants. Although (5a), (5b) and (5d) can be expressed as convex functions of $p[\ell]$, the distortion constraint in (5c) is a non-convex quadratic function of $p[\ell]$. Hence, the problem in (5) is a non-convex optimization problem which may require careful (and computationally expensive) management of locally optimal solutions. This is important because the objective and the constraints in (5) are competing criteria. The non-convexity of (5) can make it quite awkward to get an accurate description of the trade-offs between these criteria, and to determine when the constraints in (5) conflict so that there is no filter of the given length which satisfies all the constraints.

The key observation in the development of our efficient design method is that the objective and the constraints in (5) 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 $R_p(e^{j\omega}) = |P(e^{j\omega})|^2 = r_p[0] + 2 \sum_{n \geq 1} r_p[n] \cos(\omega n)$, the integral in (5a) can be analytically evaluated. It is equal to $\sum_{n \geq 0} b[n] r_p[n]$, where $b[0] = 1 - 1/K$, and for $n \geq 1$, $b[n] = -2 \sin(\pi n/K)/(\pi n)$. Therefore, the design problem in (5) can be transformed into the following optimization problem in $r_p[n]$:

$$\min_{r_p[n]} \quad \sum_{n \geq 0} b[n] r_p[n] \tag{6a}$$

$$\text{subject to} \quad R_p(e^{j\omega}) \leq \epsilon_{\text{sb}}^2 \quad \forall \omega \in [\pi/K, \pi], \tag{6b}$$

$$\sum_{i\geq 1} r_p[Mi]^2 \leq \epsilon_\gamma^2/2, \tag{6c}$$

$$R_p(e^{j\omega}) \leq B^2 \quad \forall\omega, \tag{6d}$$

$$r_p[0] = K/M, \tag{6e}$$

$$R_p(e^{j\omega}) \geq 0 \quad \forall\omega. \tag{6f}$$

The additional constraint in (6f) is a necessary and sufficient condition for $r_p[n]$ to correspond to the autocorrelation coefficients of a filter. Given a sequence $r_p[n]$ which solves (6), a corresponding filter $p[\ell]$ can be found using standard spectral factorization techniques; e.g., [8]. The objective in (6a) and the constraints in (6b), (6d), (6e), (6f) are linear, and hence convex, in $r_p[n]$, and (6c) is a convex quadratic constraint. 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 (6) can be reliably detected. However, the constraints in (6b), (6d) and (6f) 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. These constraints can be approximated by discretization, but a precise alternative is to transform them [9] into linear matrix inequalities (LMIs), which can be efficiently enforced using semidefinite programming (SDP) techniques. We will use the latter technique in our designs.

## 4. PERFORMANCE COMPARISON

We compare our design method with a method recently proposed by Harteneck *et al.* [3]; see also [5]. This method is of interest because, like our method, it generates filter banks with nearly paraunitary polyphase matrices, and it tackles the distortion induced by the filter bank directly. In the notation of the present paper, Harteneck's formulation can be written as

$$\min_{p[\ell]} \qquad \lambda E_{P,\mathrm{sb}} + \gamma_p^2 + \left(\textstyle\sum_\ell p[\ell]^2 - K/M\right) \tag{7a}$$

$$\text{subject to} \qquad P(e^{j\omega}) \text{ having linear phase,} \tag{7b}$$

where $\lambda \geq 0$ is a chosen weighting. For odd length symmetric filters, imposing phase linearity is equivalent to requiring $p[\ell] = p[L-1-\ell]$. The problem in (7) is not convex, but local minima can be found quite efficiently using an iterative least-squares technique. 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.

**Example 1** In this example we design length 49 prototype filters for a GDFT filter bank with $M = 8$ subbands and a downsampling factor of $K = 6$. In Fig. 1 we provide the tradeoffs between the normalized distortion coefficient, $\gamma_p^2/E_p$, and the normalized stop-band energy, $E_{P,\mathrm{sb}}/E_p$, achieved by different design methods. Here, $E_p = \sum_\ell p[\ell]^2$ is the energy of the filter. The dashed curve in Fig. 1 is the trade-off achieved by Harteneck's method [3]. We have indicated the points achieved for specific values of $\lambda$ by the symbols on that curve. The solid curve in Fig. 1 is the inherent tradeoff between the distortion coefficient and the stop-band energy, in the sense that no length 49 filter can achieve any point below the curve. Filters achieving this trade-off can be efficiently obtained by solving (6) in the absence of the spectral mask imposed by (6b) and (6d). The dotted curves in Fig. 1 are the inherent trade-off curves for filters which must also satisfy a spectral mask. For both dotted curves we chose $B^2 = 10^{0.1}K$ so that the maximal spectral component was at most 1 dB above the natural pass-band level induced by the energy normalization in (6e). The maximum stop-band levels were chosen to be 30 dB (the "taut" mask) and 33 dB (the "tighter" mask) below $B^2$, respectively. Despite having to satisfy these additional mask constraints, our method still generates a better distortion/stop-band energy trade-off than Harteneck's method. The power spectra of representative filters from the trade-off curves are provided in Fig. 2. Each filter was found in under 5 s by solving an LMI version of (6) using `SeDuMi` [10] on a 1.6 GHz Pentium IV workstation. □

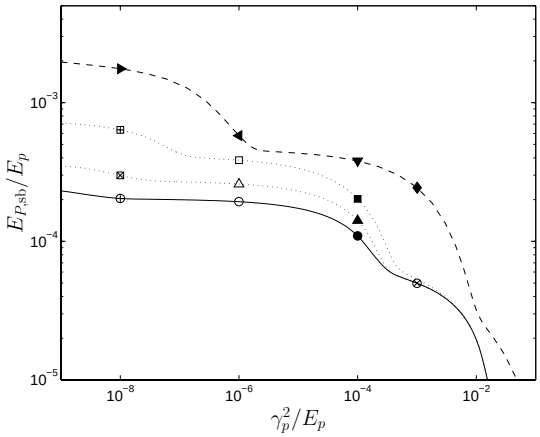


**Fig. 1**. Trade-offs between the stop-band energy and the distortion coefficient for Ex. 1. Solid: the inherent trade-off [our method, no mask]; Dashed: Harteneck's method [3]; Dotted: our method with additional spectral mask constraints; Dotted with △,▲: relative stop-band level is -30 dB; Dotted with □,■: relative stop-band level is -33 dB. The symbols ▶, ◀, ▼, ♦ indicate the trade-offs achieved by Harteneck's method with $\lambda = 6.75 \times 10^{-3}$, $4.34 \times 10^{-2}$, 1.19, and 2.40, respectively.

**Example 2** To verify that the improved trade-offs achieved by our method (see Fig. 1) can generate significant performance gains for the subband signal processing system as a whole, we examined the performance of a simple subband adaptive filtering system equipped with the GDFT filter banks designed in Ex. 1 in a synthetic acoustic echo cancellation (AEC) environment. In synthesizing the filter banks from the prototype filter using (1), we chose $m_0 = 1/2$, and $n_0 = -(L-1)/2$. The latter choice ensures that if $p[n]$ has linear phase (as it does in Harteneck's designs) then all the filters in the filter bank also have linear phase.

We evaluated the average performance of the subband adaptive filter over a class of randomly generated echo paths of length 60 with impulse response $c[n] \sim \mathcal{N}(0, e^{-n/10})$. This class of echo paths shares many of the characteristics of the acoustic impulse response encountered in practical AEC applications. The adaptive filters had length 10 and were adapted using the normalized least-mean square (NLMS) algorithm, with step-size coefficient $\tilde{\mu} = 0.8$. The input signal was 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. To illustrate the influence of the properties of the prototype on the performance of the subband adaptive filter, we have provided in Table 1 the steady-state mean square residual echo (averaged over

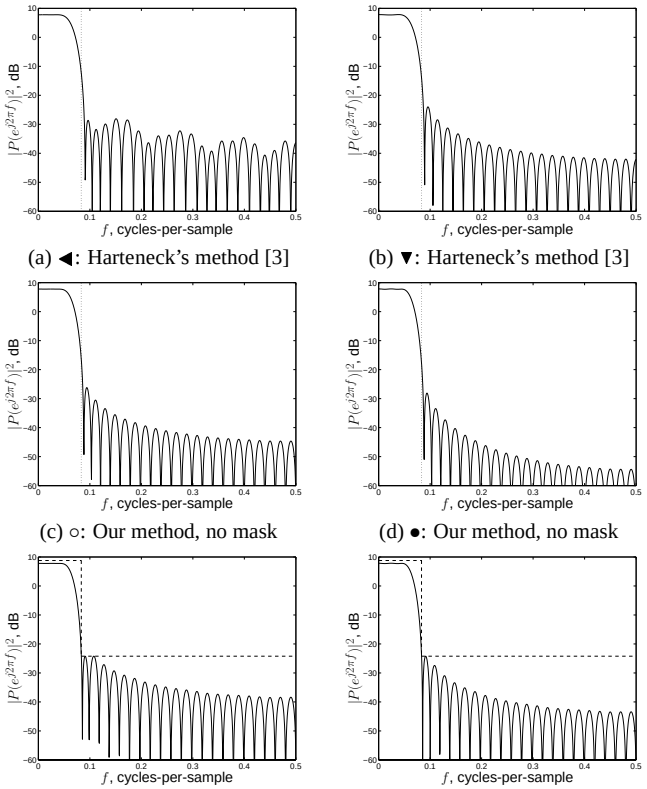


(a) ◀: Harteneck's method [3] (b) ▼: Harteneck's method [3]

(c) ○: Our method, no mask (d) ●: Our method, no mask

(e) □: Our method, tighter mask (f) ■: Our method, tighter mask

**Fig. 2**. Power spectra of filters which achieve the stop-band energy versus distortion coefficient trade-offs indicated by the symbols in Fig. 1. The dotted and dashed lines indicate the stop-band edge and spectral mask, respectively. In the left column $\gamma_p^2/E_p = 10^{-6}$, whereas in the right column $\gamma_p^2/E_p = 10^{-4}$.

1000 realizations of the unknown system and the input signal) for systems based on filters which achieve the marked trade-offs in Fig. 1. For convenience we have listed in that table, the distortion coefficient, the relative stop-band level (below $B^2 = 10^{0.1}K$), and the stop-band energy of each filter. From Table 1 it is clear that for a given distortion coefficient, the lower stop-band energies achieved by our formulations 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 the presence of a taut spectral mask results in improved performance, but that if this mask is too tight, then the performance degrades. Furthermore, it can be seen 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. □

While the performance differences in the simple setting in Ex. 2 are rather subtle, this example 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 stop-band energy, the maximum stop-band level, and the distortion coefficient; and second, to obtain optimized performance from the subband adaptive filter in a particular application, we should explore the trade-offs between these three quantities. Our convex formulation in (6) provides an efficient method for evaluating these trade-offs and should be a convenient tool for system designers.

**Table 1**. Steady-state mean square error (SS-MSE) for the adaptive filtering system in Ex. 2 equipped with prototype filters from Fig. 1. Also included are the distortion coefficient, the maximum relative stop-band level, and the stop-band energy.

| Design Method | Symbol in Fig. 1 | $\gamma_p^2/E_p$ | rel. SBL, dB | $E_{P,\mathrm{sb}}/E_p$, $\times 10^{-4}$ | SS-MSE, dB |
|---|---|---|---|---|---|
| Harteneck [3] | ▶ | $10^{-8}$ | -18.7 | 17.6 | -19.93 |
| Ours, no mask | ⊕ | $10^{-8}$ | -24.2 | 2.04 | -23.27 |
| Ours, taut mask | ⊠ | $10^{-8}$ | -30.0 | 2.99 | -23.61 |
| Ours, tighter mask | ⊞ | $10^{-8}$ | -33.0 | 6.36 | -22.98 |
| Harteneck [3] | ◀ | $10^{-6}$ | -20.3 | 5.78 | -21.98 |
| Ours, no mask | ○ | $10^{-6}$ | -24.5 | 1.93 | -23.35 |
| Ours, taut mask | △ | $10^{-6}$ | -30.0 | 2.59 | -23.69 |
| Ours, tighter mask | □ | $10^{-6}$ | -33.0 | 3.85 | -23.62 |
| Harteneck [3] | ▼ | $10^{-4}$ | -21.8 | 3.80 | -22.43 |
| Ours, no mask | ● | $10^{-4}$ | -25.7 | 1.09 | -23.77 |
| Ours, taut mask | ▲ | $10^{-4}$ | -30.0 | 1.41 | -24.00 |
| Ours, tighter mask | ■ | $10^{-4}$ | -33.0 | 2.02 | -24.07 |
| Harteneck [3] | ♦ | $10^{-3}$ | -23.3 | 2.44 | -22.01 |
| Ours, no mask | ⊗ | $10^{-3}$ | -30.0 | 0.50 | -23.21 |
| Ours, tighter mask | | $10^{-3}$ | -33.0 | 0.53 | -23.35 |

## 5. REFERENCES


[1] P. P. Vaidyanathan, *Multirate Systems and Filter Banks*, Prentice Hall, Englewood Cliffs, NJ, 1993.

[2] M. K. Mıhçak, P. Moulin, M. Anitescu, and K. Ramchandran, "Rate-distortion-optimal subband coding without perfect reconstruction constraints," *IEEE Trans. Signal Processing*, vol. 49, pp. 542–557, Mar. 2001.

[3] M. Harteneck, S. Weiss, and R. W. Stewart, "Design of near perfect reconstruction oversampled filter banks for subband adaptive filters," *IEEE Trans. Circuits Systems–II*, vol. 46, pp. 1081–1085, Aug. 1999.

[4] K. Eneman and M. Moonen, "DFT modulated filter bank design for oversampled subband systems," *Signal Processing*, vol. 81, pp. 1947–1973, Sep. 2001.

[5] S. Weiss, A. Stenger, R. W. Stewart, and R. Rabenstein, "Steady-state performance limitations of subband adaptive filters," *IEEE Trans. Signal Processing*, vol. 49, pp. 1982–1991, Sep. 2001.

[6] 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 Trans. Signal Processing*, vol. 50, pp. 2730–2743, Nov. 2002.

[7] M. R. Wilbur, T. N. Davidson, and J. P Reilly, "Efficient design of oversampled NPR GDFT filter banks," Aug. 2002, 30 pages. Submitted to the *IEEE Trans. Signal Processing*.

[8] 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. Birkhäuser, Boston, MA, 1997.

[9] 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 Trans. Signal Processing*, vol. 50, pp. 2702–2715, Nov. 2002.

[10] 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.