Efficient Design of Robust Pulse Shapes for Communications using Average Performance Criteria

Timothy N. Davidson*
Department of Electrical and Computer Engineering,
McMaster University,
Hamilton, Ontario, L8S 4K1, Canada.

Abstract

The design of a pulse shaping filter which provides optimal worst-case performance in an unknown frequency-selective channel, has recently been formulated as a convex optimization problem from which an optimal filter can be efficiently obtained. However, the resulting design can be quite conservative in an average performance sense if the worst-case scenario rarely occurs. In this paper, we develop an alternative design problem based on the average performance over a given statistical model for the channel uncertainty. We show that this problem can also be formulated as a convex optimization problem, and that the design can be made robust to inaccuracies in the assumed statistical model for the channel uncertainty. The convex formulation is used to efficiently design pulse shapes which provide superior performance to that selected as the 'chip waveform' in a recent standard for digital mobile telephony.

1 Introduction

In digital communications, waveform coding is often performed by linear pulse amplitude modulation (PAM) of translated versions of a given pulse [1]. The choice of pulse shape critically impacts many system performance criteria, including spectral efficiency and mitigation of expected channel distortion. In applications in which an accurate channel model is available, there are several established techniques by which a pulse shape can be designed [1]. However, in some wireless applications the transmission environment may undergo substantial variations and it might not be possible to obtain an accurate channel model. In that case, one ought to design pulses which provide robust performance in the presence of channel uncertainty. Since that design problem is non-convex in its direct form, effective techniques for designing robust pulses have been rather scarce. Recently it was shown [2, 3] that the design of a pulse shape which provides optimal worst-case performance over a class of bounded channel uncertainties can be posed as a convex optimization problem in the autocorrelation of the pulse shape. Furthermore, bounds on the performance of the pulse shape in an ideal channel and constraints on the spectral occupation of the PAM scheme can be easily incorporated into this convex problem. Therefore, effective robust pulse shapes can be efficiently obtained. However, if the worst-case scenario rarely occurs, the resulting design can be quite conservative in an average performance sense. In this paper, we develop an alternative design problem based on the average performance over a given statistical model for the channel uncertainty, and show that this problem can also be transformed into a convex optimization problem. In addition, we show that the design can be made robust to inaccuracies in the statistical model for the channel uncertainty. Such design problems are particularly

* Email: davidson@mcmaster.ca

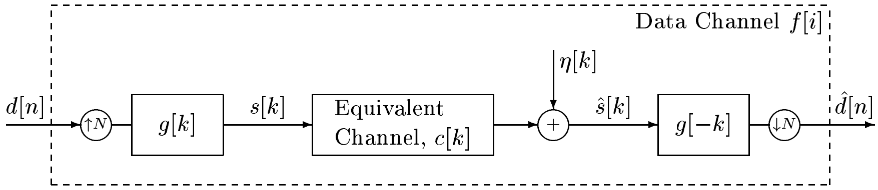

Figure 1: Discrete-time model of baseband PAM.

important in PAM applications in which spectral efficiency is required and the channel may vary (slowly, with respect to the pulse duration), but for which equalization is deemed to be too expensive. We will use the convex formulation to efficiently design pulse shapes which provide superior performance to that selected as the 'chip waveform' in the IS95 standard [4] for code division multiple access (CDMA) digital mobile telephony.

2 Robust Pulse Amplitude Modulation

Consider the discrete-time baseband PAM scheme illustrated in Fig. 1, where the equivalent channel includes conversion to and from a continuous-time signal, carrier modulation and demodulation, and the physical frequency-selective (fading) channel. If the equivalent channel does not vary significantly (in time) over the duration of the waveform, the received data estimate $\hat{d}[n]$ is

$$\hat{d}[n] = f[0]d[n] + u[n] + \sum_k g[k-Nn]\eta[k], \tag{1}$$

where $u[n] = \sum_{i\neq 0} f[i]d[n-i]$ is the intersymbol interference (ISI), $f[i] = \sum_k c[k]r_g[k-Ni]$, $r_g[m] = \sum_k g[k]g[k+m]$ is the autocorrelation function of the 'pulse-shaping' filter $g[k]$, and $\eta[k]$ models the additive noise. For notational convenience, we have normalized the time indices in (1) and hence the equivalent channel may not be causal. For simplicity, we will consider real-valued systems with white Gaussian noise, but the techniques described herein have natural extensions which allow uncertainty in the 'colour' of the noise [3], complex-valued waveforms, and explicit modelling of delays. When threshold detection is used to detect $d[n]$ from $\hat{d}[n]$, the worst-case value of the ISI over all data sequences is a key performance indicator [1,3]. Using the Cauchy-Schwarz inequality we have that

$$u[n]^2 \leq C_d^2 \sum_{i\neq 0} f[i]^2, \tag{2}$$

where $C_d^2 = \max_{d[k]} \sum_{i\neq 0, i\in \mathrm{supp}\, f} d[n-i]^2$ and supp f denotes the support of f. (Note that for standard systems, we will need $f[i]$, and therefore $c[k]$, to be FIR to ensure that C_d is finite.) Therefore, a natural design criterion would be to minimize $\sum_{i\neq 0} f[i]^2$. Unfortunately, in many applications, the channel coefficients, $c[k]$ are unknown, and hence $f[i]$ cannot be computed in advance. In previous work, we have addressed that problem by minimizing the maximum value of a measure of the ISI over a bounded class of channels [2,3]. The alternative approach developed in the next section seeks to minimize $E\left\{\sum_{i\neq 0} f[i]^2\right\}$, where $E\{\cdot\}$ denotes the expectation over an assumed statistical model for the channel coefficients. Before that development begins, we discuss a standard constraint on pulse shape design: the spectral occupation of the resulting PAM scheme.

The spectral occupation of a PAM scheme is usually measured in terms of the (time-averaged) power spectrum of the transmitted signal. For stationary white data with zero mean and variance v_d, the power spectrum of $s[k]$ is $S_s(e^{j\omega}) = v_d\left|G(e^{j\omega})\right|^2$, where $G(e^{j\omega}) = \sum_k g[k]e^{-j\omega k}$ is the discrete-time Fourier Transform of $g[k]$. The spectral occupation can be constrained in many ways (e.g., [5]), but in common with many communications standards, we will enforce a relative spectral mask on the transmitted signal. Using standard techniques, this mask can be transformed (e.g., [3, 5]) into a spectral mask on $s[k]$ of the form: $M_\ell(e^{j\omega}) \leq S_s(e^{j\omega})/\zeta \leq M_u(e^{j\omega})$, where $\zeta > 0$ is a reference value and $M_\ell(e^{j\omega})$ and $M_u(e^{j\omega})$ denote, respectively, the lower and upper bounds of the mask. An observation that is a key element in our design method is that $R_g(e^{j\omega}) = |G(e^{j\omega})|^2$. Hence $S_s(e^{j\omega})$ is a linear function of $r_g[m]$, but is, in general, a quadratic function of $g[k]$.

3 Design Technique

Simple communication systems of the form in Fig. 1 are most appropriate for scenarios in which the expected channel is ideal (i.e., $E\{c[k]\} \propto \delta[k]$), because the receiver structure resembles that of the optimal receiver for the expected channel. Channels with this characteristic are common in mobile communication systems in which there is a 'line of sight' propagation path, and we will restrict our attention to such channels. In more general channels, block-based PAM schemes with equalization may be more appropriate than the scheme in Fig. 1, and extensions of the principles of this paper to such schemes are being explored [6]. For the scheme in Fig. 1 and for a given constellation, the worst case ISI over all possible data patterns is bounded by a constant multiple of $\sum_{i\neq 0} f[i]^2$; see (2), and given a statistical model of the channel coefficients, the average value of this ISI bound is proportional to $E\left\{\sum_{i\neq 0} f[i]^2\right\}$. Therefore, a natural design objective is to ensure that the averaged bound on the worst-case ISI is small with respect to the expected gain of the desired symbol, $E\{f[0]\}$, subject to spectral occupation constraints. An instance of that problem can be stated more formally as: *For a given statistical model for $c[k]$ with $E\{c[k]\} = \delta[k]$, a relative spectral mask specified by $M_\ell(e^{j\omega}) \geq 0$ and $M_u(e^{j\omega})$, and for given N and L_g, find a filter of length L_g achieving $\min_{g[k]} E\left\{\sum_{i\neq 0} f[i]^2\right\}$, subject to $E\{f[0]^2\} = \left(\sum_{k=0}^{L_g-1} g[k]^2\right)^2 = 1$, and $\zeta M_\ell(e^{j\omega}) \leq \left|G(e^{j\omega})\right|^2 \leq \zeta M_u(e^{j\omega})$ for all $\omega \in [0,\pi]$ and some $\zeta > 0$, or show that none exist.*

Unfortunately, $E\left\{\sum_{i\neq 0} f[i]^2\right\}$ is a quartic function of $g[k]$, and the lower bound constraint on the power spectrum generates non-convex quadratic constraints on $g[k]$. Therefore, any direct design algorithm for the optimal $g[k]$ is complicated by the intricacies of dealing with potential local minima. In contrast, by defining $\boldsymbol{r}_g$ such that $[\boldsymbol{r}_g]_m = r_g[m]$, $0 \leq m \leq L_g - 1$ and $\tilde{\boldsymbol{r}}_g = \boldsymbol{T}\boldsymbol{r}_g$, where

$$\boldsymbol{T} = \begin{bmatrix} \boldsymbol{0} & \boldsymbol{J}_{L_g-1} \\ 1 & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{I}_{L_g-1}, \end{bmatrix} \tag{3}$$

with $\boldsymbol{I}_K$ being the $K \times K$ identity matrix and $\boldsymbol{J}_K$ being the $K \times K$ matrix consisting of ones on the anti-diagonal and zeros elsewhere, we can write $f[i] = \boldsymbol{c}_i^T \tilde{\boldsymbol{r}}_g$, where $[\boldsymbol{c}_i]_\ell = c[\ell + Ni]$. Hence, $E\left\{\sum_{i\neq 0} f[i]^2\right\} = \boldsymbol{r}_g^T \boldsymbol{Q}_c \boldsymbol{r}_g$, where $\boldsymbol{Q}_c = \boldsymbol{T}^T \sum_{i\neq 0} E\{\boldsymbol{c}_i \boldsymbol{c}_i^T\}\boldsymbol{T}$, and hence $E\left\{\sum_{i\neq 0} f[i]^2\right\}$ is a convex quadratic function of $r_g[m]$. Furthermore, $\left|G(e^{j\omega})\right|^2$ is a linear function of $r_g[m]$. To complete the re-formulation of the design in terms of $r_g[m]$ instead of $g[k]$, we must add the additional constraint $R_g(e^{j\omega}) \geq 0$ for all $\omega \in [0,\pi]$, which is a necessary and sufficient condition for $r_g[m]$ to be factorizable in the form $r_g[m] = \sum_k g[k]g[k+m]$, (by the Féjer-Riesz Theorem). However, this constraint is implicit in the spectral mask because $M_\ell(e^{j\omega}) \geq 0$. By performing this re-formulation we obtain the following convex optimization problem for the optimal $\boldsymbol{r}_g$:

Problem 1 *Given $M_\ell(e^{j\omega}) \geq 0$, $M_u(e^{j\omega})$, $\boldsymbol{Q}_c$, N and L_g, find a filter of length L_g achieving $\min \boldsymbol{r}_g^T \boldsymbol{Q}_c \boldsymbol{r}_g$ over $\boldsymbol{r}_g$ and $\zeta > 0$, subject to $r_g[0] = 1$ and the spectral mask*

$$\zeta M_\ell(e^{j\omega}) \leq R_g(e^{j\omega}) \leq \zeta M_u(e^{j\omega}), \qquad \textit{for all } \omega \in [0,\pi], \tag{4}$$

or show that none exist.

Problem 1 consists of a convex quadratic objective, subject to a linear equality constraint and linear inequality constraints (4), and hence it is a quadratic programme. If $\boldsymbol{Q}_c = \boldsymbol{L}_c \boldsymbol{L}_c^T$ is a factorization of $\boldsymbol{Q}_c$ (e.g., Cholesky), then by introducing an additional variable θ_{av}, Problem 1 can be re-written as: min θ_{av}, *subject to* $\|\boldsymbol{L}_c^T \boldsymbol{r}_g\|^2 \leq \theta_{av}$, $r_g[0] = 1$ *and* (4)*;* which is a 'rotated' second-order cone programme (SOCP) [7]. (We will use this alternative form in our implementation.) The spectral mask (4) generates an infinite set of linear inequality constraints (two for each frequency ω), and this set must be rendered finite before Problem 1 can be efficiently solved. One standard technique is to approximate the infinite constraint set via discretization techniques.[1] The resulting quadratic programme has a finite number of linear constraints, and hence the globally optimal autocorrelation can be efficiently found (e.g., using interior point methods). However, in many communications applications, we encounter spectral masks which are piecewise constant. For these masks (and for piecewise trigonometric polynomial masks) we can avoid the heuristic approximation induced by discretization and *precisely* transform the semi-infinite spectral mask constraint into a finite set of linear constraints on a finite set of positive semidefinite matrices [9, 10]. The resulting optimization problem is a cone programme with a concatenation of second-order and semidefinite cones and can also be efficiently solved using interior point methods. Once the optimal autocorrelation has been found by solving Problem 1 (using either mask representation), an optimal pulse shaping filter can be found by spectral factorization [8, 11].

While solutions to Problem 1 generate PAM schemes which perform well under the assumed statistical model for the channel, one may wish to modify that problem to ensure that the resulting PAM scheme performs well even if the statistical channel model is inaccurate. To do so, we let $\boldsymbol{Q}_c = \bar{\boldsymbol{Q}}_c + \boldsymbol{\Delta}_{\boldsymbol{Q}}$, where $\bar{\boldsymbol{Q}}_c$ represents the matrix $\boldsymbol{Q}_c$ for the nominal statistical model, and the symmetric matrix $\boldsymbol{\Delta}_{\boldsymbol{Q}}$ represents the uncertainty in $\boldsymbol{Q}_c$ due to inaccuracies in the statistical model. We model this uncertainty by constraining $\boldsymbol{\Delta}_{\boldsymbol{Q}}$ to lie in the parameterized 'admissible set' $\Omega_\upsilon = \{\boldsymbol{\Delta}_{\boldsymbol{Q}} \,|\, \boldsymbol{\Delta}_{\boldsymbol{Q}} = \boldsymbol{\Delta}_{\boldsymbol{Q}}^T, -\upsilon \boldsymbol{I} \leq \boldsymbol{\Delta}_{\boldsymbol{Q}} \leq \upsilon \boldsymbol{I}\}$, where υ represents the 'size' of the uncertainty. This admissible set is the set of all symmetric matrices with all their eigenvalues bounded by $\pm\upsilon$, and as such $\boldsymbol{\Delta}_{\boldsymbol{Q}} \in \Omega_\upsilon$ incorporates uncertainties in both the variance of individual channel coefficients $c[k]$ and the correlation between coefficients. The design approach we will take here is to minimize the worst-case value of $\boldsymbol{r}_g^T \boldsymbol{Q}_c \boldsymbol{r}_g$ over all $\boldsymbol{\Delta}_{\boldsymbol{Q}}$'s in the admissible set, subject to a performance degradation under the nominal statistical model of at most $100\nu\%$. That problem can be simplified by observing that $\max_{\boldsymbol{\Delta}_{\boldsymbol{Q}} \in \Omega_\upsilon} \boldsymbol{r}_g^T \boldsymbol{Q}_c \boldsymbol{r}_g = \boldsymbol{r}_g^T \bar{\boldsymbol{Q}}_c \boldsymbol{r}_g + \upsilon \boldsymbol{r}_g^T \boldsymbol{r}_g$, and that the performance degradation constraint can be written as $\boldsymbol{r}_g^T \bar{\boldsymbol{Q}}_c \boldsymbol{r}_g \leq (1+\nu)\theta_{\text{av}}^\star$, where $\theta_{\text{av}}^\star$ is the optimal value of the objective in Problem 1 with $\boldsymbol{Q}_c = \bar{\boldsymbol{Q}}_c$. Therefore, the design problem reduces to the following quadratic programme:

Problem 2 *Given* $M_\ell(e^{j\omega}) \geq 0$, $M_u(e^{j\omega})$, $\bar{\boldsymbol{Q}}_c$, υ, N L_g, $\theta_{\text{av}}^\star$ *and* ν, *find a filter of length* L_g *achieving* min $\boldsymbol{r}_g^T \bar{\boldsymbol{Q}}_c \boldsymbol{r}_g + \upsilon \boldsymbol{r}_g^T \boldsymbol{r}_g$ *over* $\boldsymbol{r}_g$ *and* $\zeta > 0$, *subject to* $r_g[0] = 1$, $\boldsymbol{r}_g^T \bar{\boldsymbol{Q}}_c \boldsymbol{r}_g \leq (1+\nu)\theta_{\text{av}}^\star$ *and the spectral mask in Eq.* (4)*, or show that none exist.*

Observe that as either υ, the 'size' of the uncertainty in $\boldsymbol{Q}_c$, or ν, the nominal performance degradation factor, approach zero, the solution to Problem 2 approaches that of Problem 1. (Using similar techniques to those described above, Problem 2 can be re-cast as an SOCP, and we will use that version of the problem in our implementation.) Since the objective in Problem 2 can be written as $\boldsymbol{r}_g^T(\bar{\boldsymbol{Q}}_c + \upsilon \boldsymbol{I})\boldsymbol{r}_g$, minimizing the worst case $\boldsymbol{r}_g^T \boldsymbol{Q}_c \boldsymbol{r}_g$ over all the statistical models in the admissible set subject to $r_g[0] = 1$ can be shown to be equivalent to minimizing $\boldsymbol{r}_g^T \check{\boldsymbol{Q}}_c \boldsymbol{r}_g$ subject to $r_g[0] = 1$, where $\check{\boldsymbol{Q}}_c = \bar{\boldsymbol{Q}}_c + \upsilon \boldsymbol{I}$ is the corresponding matrix generated by a modified version of the nominal statistical model for the channel. The modified statistical model is obtained from the nominal one by adding an independent component of variance $\upsilon/2$ to each channel coefficient. This 'diagonal loading' of a nominal kernel matrix of a quadratic form to obtain robustness is common to many 'regularization' approaches to robust (adaptive) signal processing,

[1] A 'rule of thumb' is to choose $15L_g$ uniformly spaced discretization points [8], plus any 'corner' frequencies of the mask.

but the explicit constraint on the performance degradation under the nominal statistical model, $\boldsymbol{r}_g^T \bar{\boldsymbol{Q}}_c \boldsymbol{r}_g \leq (1+\nu)\theta_{\mathrm{av}}^\star$, appears less frequently.

An interesting connection between the design method for deterministically bounded uncertainty in earlier work [2, 3] and those for statistically modelled uncertainty in this paper can be obtained by analyzing a limiting case of the modified design in Problem 2. As υ, the size of the uncertainty in $\boldsymbol{Q}_c$, grows, the objective of Problem 2 is dominated by the term $\boldsymbol{r}_g^T \boldsymbol{r}_g$, and as the variance of the nominal statistical model approaches zero, $\boldsymbol{r}_g^T \bar{\boldsymbol{Q}}_c \boldsymbol{r}_g$ approaches $2\sum_{i\geq 1} r_g[Ni]^2$. In this limiting case, and with the normalization $r_g[0]=1$, Problem 2 reduces to the minimization of $\boldsymbol{r}_g^T \boldsymbol{r}_g$ subject to a bound on $\sum_{i\geq 1} r_g[Ni]^2$, and the spectral mask. This is essentially the same as the design problem derived from the minimization of the worst case ISI over a set of bounded channel coefficients in [2, 3].

4 Chip Waveform Design for IS95

We now demonstrate the effectiveness of pulse shape design via Problems 1 and 2 by designing pulse shapes with improved performance over that of the ‘chip’ waveform in the IS95 standard for mobile telephony [4]. The filter specified for the synthesis of the chip waveform in IS95 [4] has $N=4$ and $L_g=48$. While that filter satisfies the spectral mask specified in the standard, it generates a rather large amount of ISI, even in an ideal channel. To determine whether this shape can be improved upon, we will use Problems 1 and 2 to design pulse shapes for a slowly-varying Rician-like channel with additive white Gaussian noise. (This type of channel is typical of mobile communication channels which involve a ‘line of sight’ propagation path.) The actual statistical model in which the designs will be evaluated is a real FIR channel of length 33 (i.e., eight chips) with $c[0]=1$ and the remaining $c[k]$ being independent zero-mean Gaussian random variables with standard deviation 0.05. The optimal autocorrelation for this statistical channel model, subject to the IS95 spectral mask was computed by solving[2] Problem 1 and is shown in Fig. 2, along with that of the IS95 filter. The improved ‘zero-crossing’ behaviour of the designed filter is evident from that figure. The power spectra of the IS95 and optimal filters are shown in Figs 3(a) and (b), respectively, from which the improved flatness in the pass band of the optimal filter is clear. It is also clear from that figure that the IS95 filter satisfies the specified spectral mask by a considerable margin. The power spectrum of an optimal filter for spectral mask *achieved* by the IS95 filter is shown in Fig. 3(c). (The autocorrelation of this new optimal filter is indistinguishable from that of the optimal filter in Fig. 3(b) at the scale of Fig. 2.)

In practice, the statistical model used in the design of the pulse shape may be incorrect, and hence we developed the modified design method in Problem 2. To demonstrate its effectiveness, we now design chip waveforms intended for the above Rician-like channel using an inaccurate statistical model of the channel which has $c[0]=1$ and the remaining $c[k]$ being from a first-order autoregressive process (of the same standard deviation) generated by passing an independent and identically distributed zero-mean Gaussian sequence with standard deviation 0.05 through the causal filter with transfer function $\sqrt{1-z_p^2}/(1-z_p z^{-1})$. This filter has a single real pole at $z=z_p$ and no zeros, and as z_p approaches zero, the statistical model used in the design approaches the true statistical model. In Fig. 4, the variation of $\boldsymbol{r}_g^T \boldsymbol{Q}_c \boldsymbol{r}_g$ in the actual statistical model is plotted against the pole position of the inaccurate statistical model used in the design of $\boldsymbol{r}_g$, for both the original and modified designs; Problems 1 and 2, respectively. (For visual clarity we only report results for the specified mask. Those for the achieved mask are qualitatively similar.) It is evident from Fig. 4 that when the design model is substantially inaccurate, the modified design provides improved performance. The (small) performance degradation (controlled by ν) that the modified design incurs for reasonably accurate

[2]For each of the design problems in this paper, the spectral masks were precisely enforced using the linear matrix inequality technique in [9, 10]. Each instance of the resulting convex cone programme was solved in around 25 seconds on a 400 MHz PENTIUM II workstation using a MATLAB-based general purpose semidefinite program solver called `SeDuMi` [12].

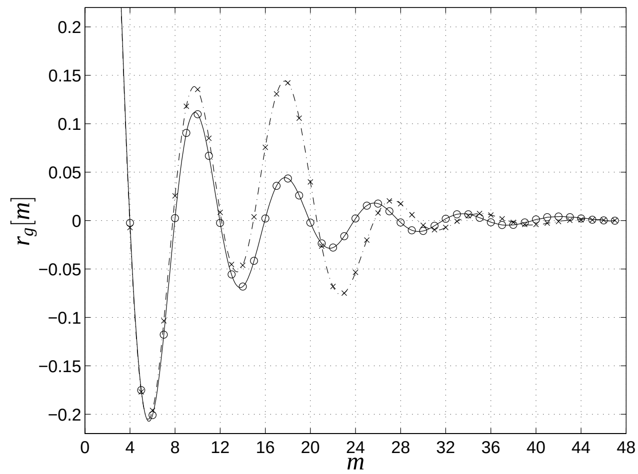

Figure 2: Detail of the autocorrelation sequences for the designed ('○') and IS95 ('×') filters. For visual clarity, the sequences have been interpolated using an ideal ('sinc-function') interpolator.

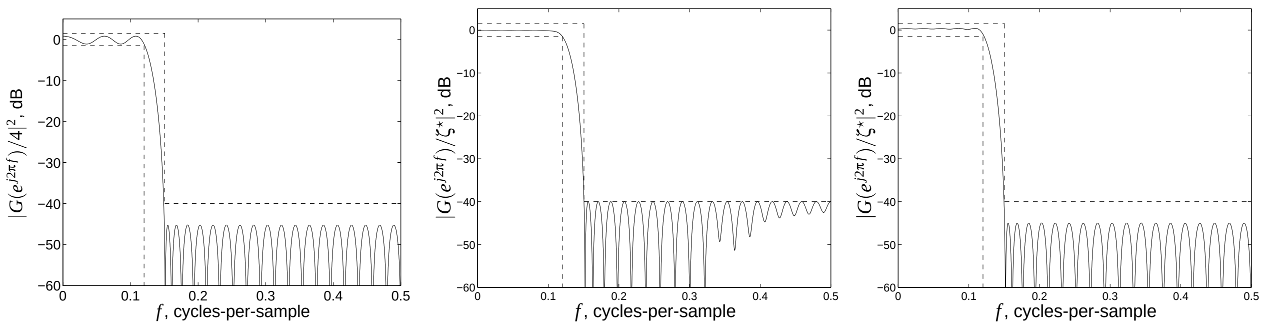

(a) IS95 (b) Robust, specified mask (c) Robust, achieved mask.

Figure 3: Relative power spectra (in decibels) of the filters with the IS95 spectral mask. Here, $\zeta^\star$ is the optimal value of ζ from Problem 1.

models of the channel statistics (in order to obtain robustness) is evident from the gap between the curves as the pole approaches the origin. The chip error rates (CERs) evaluated in the actual statistical channel model are shown in Fig. 5, for designs based on accurate and inaccurate statistical channel models. The improved robustness of the modified design and the (small) performance degradation it incurs for accurate channel models are also evident from these CER curves.

5 Conclusions

In this paper we have shown that the design of pulse shapes which provide robust performance in an average sense can be formulated as a convex optimization problem in the autocorrelation of the pulse shape. Furthermore, we have shown that robustness to uncertainties in the assumed statistical model for the channel can be easily incorporated into the design. We used this convex formulation to efficiently design pulse shapes with superior performance to that used as the 'chip' waveform in the IS95 standard for mobile wireless communications.

In closing, it is pointed out that the goal of this work was to obtain efficient design algorithms for waveforms which provide robustness to uncertain, but linear and time-invariant, channels. An interesting

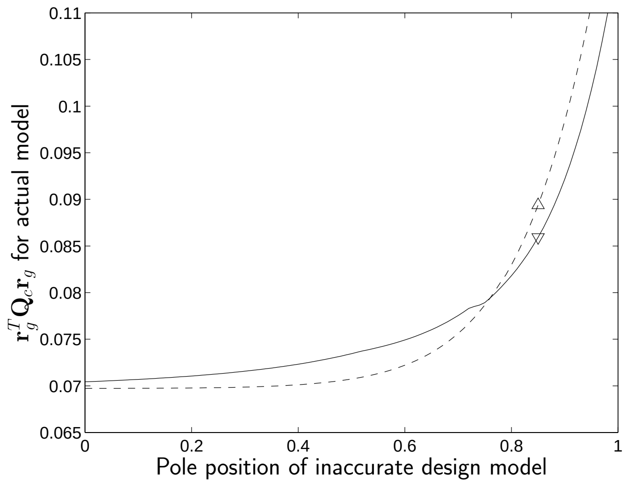

Figure 4: Value of $\boldsymbol{r}_g^T \boldsymbol{Q}_c \boldsymbol{r}_g$ for the actual statistical channel model against the pole position of the inaccurate statistical model used in the design of $\boldsymbol{r}_g$. (When the pole position is zero the design model is accurate.) Legend—Dashed: original design using Problem 1; Solid: modified design using Problem 2 with $\upsilon = 0.2$ and $\nu = 0.01$; $\triangle$, $\triangledown$: designs for the inaccurate model used in Fig. 5. Note that for the actual statistical channel model the IS95 filter generates $\boldsymbol{r}_g^T \boldsymbol{Q}_c \boldsymbol{r}_g = 0.0948$.

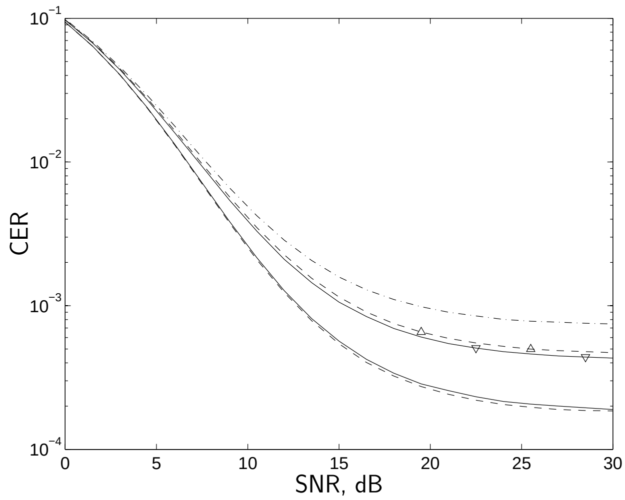

Figure 5: Simulated chip error rates (CERs) against SNR for the accurate and inaccurate design models. Legend—Dash-dot: IS95; Dashed: Problem 1 with the accurate design model; Solid: Problem 2 with $\upsilon = 0.2$, $\nu = 0.01$ and the accurate design model. Dashed with $\triangle$: Problem 1 with the inaccurate design model from Fig. 4; Solid with $\triangledown$: Problem 2 with the inaccurate design model from Fig. 4.

direction for further work is to examine ways in which robustness to uncertain time-varying non-linear channels can be incorporated into the current design framework. We also point out that we have used a sophisticated, but general purpose solver to solve our design problems. It remains to be seen whether an application specific implementation which exploits the special structure of our problem can be developed. It is expected that such a specialized implementation would be faster and would require less memory than our current implementation. The recent development of application specific implementations for related problems [13] suggests that these expectations could be realized.

References

[1] R. W. Lucky, J. Salz, and E. J. Weldon, Jr. *Principles of Data Communications*. McGraw Hill, 1968.

[2] T. N. Davidson. Efficient evaluation of design tradeoffs in waveform design for robust pulse amplitude modulation. In *Proc. Int. Conf. Acoust., Speech, Signal Processing*, Istanbul, Turkey, June 2000.

[3] T. N. Davidson. Efficient design of waveforms for robust pulse amplitude modulation. *IEEE Trans. Signal Processing*, December 2001. To appear.

[4] Proposed EIA/TIA interim standard. Wideband spread spectrum digital cellular system dual-mode mobile station-base station compatibility standard. Technical Report TR45.5, QUALCOMM, Inc., San Diego, CA, USA, April 1992.

[5] T. N. Davidson, Z.-Q. Luo, and K. M. Wong. Design of orthogonal pulse shapes for communications via semidefinite programming. *IEEE Trans. Signal Processing*, 48(5):1433–1445, May 2000.

[6] J. Milanović, T. N. Davidson, Z.-Q. Luo, and K. M. Wong. Design of robust redundant precoding filter banks with zero-forcing equalizers for unknown frequency-selective channels. In *Proc. Int. Conf. Acoust., Speech, Signal Processing*, Istanbul, Turkey, June 2000.

[7] M. S. Lobo, L. Vandenberghe, S. Boyd, and H. Lebret. Applications of second-order cone programming. *Lin. Algebra Applics*, 284(1–3):193–228, November 1998.

[8] S.-P. Wu, S. Boyd, and L. Vandenberghe. FIR filter design via spectral factorization and convex optimization. In B. Datta, editor, *Applied and Computational Control, Signals, and Circuits*, volume 1. Birkhauser, Boston, May 1999.

[9] T. N. Davidson, Z.-Q. Luo, and J. F. Sturm. Linear matrix inequality formulation of spectral mask constraints. In *Proc. Int. Conf. Acoust., Speech, Signal Processing*, Salt Lake City, UT, May 2001.

[10] T. N. Davidson, Z.-Q. Luo, and J. F. Sturm. Linear matrix inequality formulation of spectral mask constraints with appplications to FIR filter design, September 2000. Submitted to the *IEEE Trans. Signal Processing*. Revised October 2001.

[11] T. N. T. Goodman, C. A. Micchelli, G. Rodriguez, and S. Seatzu. Spectral factorization of Laurent polynomials. *Adv. in Comp. Math.*, 7(4):429–454, 1997.

[12] J. F. Sturm. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optim. Methods and Software*, 11–12:625–653, 1999. (See `http://fewcal.kub.nl/sturm/software/sedumi.html` for updates.).

[13] B. Alkire and L. Vandenberghe. Interior-point methods for magnitude filter design. In *Proc. Int. Conf. Acoust., Speech, Signal Processing*, Salt Lake City, UT, May 2001.