

This Provisional PDF corresponds to the article as it appeared upon acceptance. Fully formatted PDF and full text (HTML) versions will be made available soon.

Sparsening filter design for iterative soft-input soft-output detectors

EURASIP Journal on Wireless Communications and Networking 2012,
2012:72 doi:10.1186/1687-1499-2012-72

Raquel G Machado (raquel@wpi.edu)
Andrew G Klein (klein@ece.wpi.edu)
Richard K Martin (Richard.Martin@afit.edu)

ISSN 1687-1499

Article type Research

Submission date 15 May 2011

Acceptance date 29 February 2012

Publication date 29 February 2012

Article URL <http://jwcn.eurasipjournals.com/content/2012/1/72>

This peer-reviewed article was published immediately upon acceptance. It can be downloaded, printed and distributed freely for any purposes (see copyright notice below).

For information about publishing your research in *EURASIP WCN* go to

<http://jwcn.eurasipjournals.com/authors/instructions/>

For information about other SpringerOpen publications go to

<http://www.springeropen.com>

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 29 FEB 2012		2. REPORT TYPE		3. DATES COVERED 00-00-2012 to 00-00-2012	
4. TITLE AND SUBTITLE Sparsening Filter Design For Iterative Soft-Input Soft-Output Detectors				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Worcester Polytechnic Institute, Department of Electrical and Computer Engineering, 100 Institute Rd, Worcester, MA, 01609				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES EURASIP Journal on Wireless Communications and Networking , 29 February 2012, Government or Federal Purpose Rights License.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 37	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Sparsening filter design for iterative soft-input soft-output detectors

Raquel G Machado^{*1}, Andrew G Klein^{*1} and Richard K Martin²

¹Department of Electrical and Computer Engineering, Worcester Polytechnic Institute, 100 Institute Rd, Worcester, MA, 01609, USA

²Department of Electrical and Computer Engineering, Air Force Institute of Technology (AFIT), Wright-Patterson AFB, OH

^{*}Corresponding authors: raquel@wpi.edu; klein@ece.wpi.edu

Email address:

RKM: richard.martin@afit.edu

Abstract

A large body of research exists around the idea of channel shortening, where a prefilter is designed to reduce the effective channel impulse response to some smaller number of contiguous taps. This idea was originally conceived to reduce the complexity of Viterbi-based maximum-likelihood equalizers. Here, we consider a generalization of channel shortening which we term “channel sparsening”. In this case, a prefilter is designed to reduce the effective channel to a small number of nonzero taps which do not need to be contiguous. When used in combination with belief-propagation-based maximum *a posteriori* (MAP) detectors, an analogous complexity reduction can be realized. We address the design aspects of sparsening filters, including several approaches to minimize the bit error rate of MAP detectors. We devote attention to the interaction of the sparsening filter and detector, and demonstrate the performance gains through simulation.

Keywords: belief propagation; turbo equalization; channel sparsening; channel shortening.

1 Introduction

Intersymbol interference (ISI) caused by frequency selective channels is one of the chief impairments faced by modern, high data-rate communication receivers. The issue of compensating for ISI has been studied at length over the past five decades, and a wide range of strategies are available for use by communication system designers. For optimal performance, a maximum *a posteriori* (MAP) or maximum-likelihood (ML) sequence estimator may be implemented using the Bahl-Cocke Jelinek-Raviv (BCJR) or Viterbi algorithm, respectively. These optimal approaches, however, are exponentially complex in the number of channel coefficients, and consequently suboptimal ISI compensation techniques are used in most applications.

Sparse impulse responses are characterized as having only a small fraction of nonzero coefficients. This behavior can arise, for example, in underwater acoustic communication channels or in terrestrial communication channels over hilly terrain. Compensation of sparse ISI channels is considerably challenging since these channels can often have very long delay spreads, and optimal approaches like BCJR and Viterbi are therefore infeasible. Recently, a MAP detector employing belief propagation (BP) was proposed [1] for ISI compensation in sparse channels.

The proposed scheme is attractive because it permits near-optimal performance with complexity that depends only on the number of nonzero coefficients. The complexity of this algorithm is *exponential* in the number of nonzero channel coefficients, however, so it may still be prohibitively complex for the majority of applications. A hybrid version of this detector was proposed [2] which uses a linear prefilter in the receiver just before the BP-based MAP detector. By designing the prefilter so that the combined response of the sparse channel and prefilter has a reduced, limited number of nonzero coefficients, the use of the BP-based detector becomes feasible in a wider range of applications.

The utility of the hybrid structure in [2] has been demonstrated through an extensive simulation study that showed significant error-rate improvement in sparse channels when compared with competing approaches that employ decision-feedback equalizers. While the simulation results are encouraging, relatively little attention is paid to the interaction of the prefilter and the BP-based detector. For example, the prefilter is arbitrarily designed so that the taps of the combined response of the sparse channel and prefilter coincide with the dominant taps in the original channel, yet no motivation is provided for this choice.

In this work, we focus on the design of *sparsening* prefilters for use with soft-input soft-output MAP detectors of the form considered in [1, 2]. While [1, 2] primarily focused on the case where the original channel is sparse, we note that even non-sparse channels can be sparsened with a simple linear, finite impulse response (FIR) filter. Consequently, our work can be applied in general situations, even where the original channel is not sparse. We address the issue of sparsening filter design with the goal of minimizing the detector bit error rate (BER). We consider the interaction of the sparsening filter and BP detector, and develop a practically-implementable sparsening filter design method.

We note that channel sparsening filters can be seen as a generalization of so-called *channel shortening filters* proposed in [3–7]. Given an FIR channel $\mathbf{h}$ of length L_h , the channel shortening problem roughly amounts to designing a filter $\mathbf{w}$ so that the energy in the combined response $\mathbf{h} \star \mathbf{w}$ is concentrated in $\mu < L_h$ *contiguous* taps. Channel sparsening is nearly the same, though the μ taps which contain the majority of the energy are not constrained to be contiguous. Furthermore, while much of the recent interest in channel shortening has been for application to multicarrier systems, the original idea of channel shortening [3] was proposed for a reduced-complexity hybrid prefilter/ML detector which bears some resemblance to the one considered here. More recent works such as [8] have considered channel shortening in conjunction with iterative MAP detectors. Again, however, these works impose a constraint that the taps in the combined channel/filter response must be contiguous. One very recent work [9] has considered use of matching pursuit to find a sparse, non-contiguous target impulse response (TIR), and it is shown to yield a lower mean squared error (MSE) compared to the conventional contiguous approach.

In Section 2 we present the basic system model, and in Section 3 we describe the basic operation of the BP detector. Sections 4 through 6 focus on the design aspects of the channel sparsening filter (CSF), while Section 7 presents simulation results which compare various sparsening filter design methods. Finally, Section 8 concludes the article.

2 System model

We consider the system model shown in Figure 1. A sequence of symbols $x[k]$ drawn from an M -ary alphabet is transmitted through an intersymbol interference channel denoted $\mathbf{h}$ —which may or may not be sparse—and additive white Gaussian noise (AWGN) $n[k]$ with variance σ_n^2 is added. At the receiver, we employ a detector which consists of the cascade of a CSF which we denote $\mathbf{w}$, followed by a belief-propagation-based detector [1]. As mentioned, the BP detector is exponentially complex in the number of nonzero channel taps. Consequently, the purpose of the CSF is to reduce the number of nonzero coefficients in the effective channel to some specified quantity μ so that use of the BP detector becomes practically feasible.

Let $\mathbf{c} = \mathbf{h} \star \mathbf{w}$ be the effective channel (or combined response) where the $\star$ operator denotes convolution. We assume that the channel, CSF, and effective channel are modeled as FIR filters of lengths L_h , L_w , and $L_c = L_h + L_w - 1$, respectively. Thus, the received data is given by

$$y[k] = \sum_{l=0}^{L_h-1} h[l]x[k-l] + n[k].$$

After filtering by the CSF, the data is

$$\begin{aligned} z[k] &= \sum_{l=0}^{L_w-1} w[l]y[k-l] \\ &= \sum_{l=0}^{L_c-1} c[l]x[k-l] + v[k]. \end{aligned} \tag{1}$$

where the CSF output noise, which is colored, is given by $v[k] = n[k] \star w[k]$. Finally, the output of the CSF is passed to the soft-input soft-output BP detector which outputs likelihood values that can be used to make decisions as to what was transmitted.

In this work, we focus our attention on the optimal design of the sparsening filter. As such, we make the simplifying assumption that the channel $\mathbf{h}$ is known perfectly to the receiver. It is rather straightforward, however, to extend our proposed design method to adaptive implementations which can be employed in situations where the channel is unknown and/or slowly time-varying.

3 Belief propagation detector

Before discussing the sparsening filter design in detail, we first provide some details about the BP detector. The BP algorithm used in the detector is in the class of message passing algorithms, and is sometimes called the *sum-product algorithm* [1]. By representing the ISI channel as a factor graph, the BP algorithm can be used to implement MAP detection, thereby finding the sequence $\mathbf{x}$ which maximizes the joint *a posteriori* probability density function $P(\mathbf{x}|\mathbf{y})$. The BP algorithm proceeds iteratively, and computes log likelihood ratios of the transmitted bits which become more reliable with each iteration. After a sufficient number of iterations, the log likelihood ratios can be used to make bit decisions.

To compute the likelihood ratios, the BP detector needs to know the effective channel impulse response. Given a finite-length filter $\mathbf{h}$, it is not possible in general to find a finite-length filter $\mathbf{w}$ such that the combined response is exactly equal to some prescribed FIR response $\mathbf{c}$ since the problem is overdetermined.^a In other words, in designing the CSF filter $\mathbf{w}$, we must accept that it is not possible to perfectly sparsen the channel so that the effective channel $\mathbf{c}$ consists of only μ nonzero taps. Consequently, the remaining $L_c - \mu$ taps of $\mathbf{c}$ will not be exactly equal to zero in general unless the CSF is chosen to have infinite length. Nevertheless, to keep computational complexity at the level prescribed by the choice of μ , we only use the largest μ taps of $\mathbf{c}$ in the computation of the likelihood ratios used inside the BP detector. As such the residual ISI contribution from the smallest $L_c - \mu$ taps of $\mathbf{c}$ in (1) will be treated as noise by the BP detector. A sufficiently large choice of CSF length L_w , however, can ensure arbitrarily small additional ISI.

Since the BP detector is typically implemented in the log domain, the majority of its complexity is

due to the many summation operations which must be performed [1]. If the BP algorithm proceeds over N total iterations, the total complexity requires on the order of $N(\mu + 1)M^{\mu+1}$ summations, where M is the size of the source alphabet and μ is the number of significant effective channel taps used in the detection. As such, the complexity of the BP is exponential in μ , and so the system designer can specify the total complexity by appropriate choice of μ .

We note that the BP detector performance only truly coincides with the MAP detector when two conditions are met: (1) there are no cycles in the factor graph corresponding to the channel, and (2) the additive noise is white and Gaussian. In general, the first of these conditions is never satisfied. In practice cycles have been shown to be of little concern since they are a low probability event (in the case of potentially detrimental length 4 cycles) [10], or the cycles themselves do not pose a noticeable performance penalty [1]. The second condition on the noise, however, is more serious for this hybrid structure. Since the AWGN gets colored by the CSF, the noise at the input of the BP detector is no longer white.^b We will address this issue in the sequel.

We emphasize that the CSF does not change the operation of the BP detector. As the CSF changes the effective channel taps, however, and passes the μ largest effective channel taps to the BP detector, the CSF obviously affects the behavior and performance of the combined filter/detector structure. Since the BP detector itself is unaltered from [1], it can accommodate a system employing channel codes such as LDPC encoding considered in [1], or can readily be extended to the MIMO case with, for example, space-time coding as in [2, 8]. Since our focus is on the design of the CSF, we consider an uncoded system.

4 Channel sparsening

In the design of the CSF $\mathbf{w}$, the goal is for the number of significant (nonzero) taps of $\mathbf{c}$ to be μ or less, regardless of where they lie in $\mathbf{c}$ or whether they are contiguous or not. We note that $\mu \in \{1, 2, \dots, L_h\}$ is a parameter chosen by the system designer. If $\mu = 1$, then the detector coincides with traditional linear equalization since the goal of the CSF design is to make the effective channel be a single spike. At the

other extreme, the choice $\mu = L_h$ corresponds to “pure” BP detection as in [1] since the CSF need not do any sparsening and can be a simple unity gain filter. Larger choices of μ will result in an exponentially more complicated BP detector, but will also result in better BER performance.

Ideally, we would like to choose $\mathbf{w}$ to minimize the system BER. As no closed form expression for the BER exists, direct minimization of BER is intractable. Consequently, we consider choosing $\mathbf{w}$ to maximize a proxy for the BER which we term the Sparse Shortening SNR (SSSNR). In the sequel, we will assess the validity of this metric by measuring the proximity of the SSSNR-optimal filter to the BER-optimal filter for some low-dimensional examples. Accordingly, we define the SSSNR as the ratio of the useful signal power coming out of the μ large taps of $\mathbf{c}$ over the total power of the received signal plus noise,

$$\begin{aligned} J_S(\mathbf{w}) &= \frac{\sigma_x^2 \|\text{large taps of } \mathbf{c}\|^2}{\sigma_x^2 \|\text{large taps of } \mathbf{c}\|^2 + \sigma_x^2 \|\text{small taps of } \mathbf{c}\|^2 + \sigma_v^2} \\ &= \frac{\sigma_x^2 \sum_{k \in \mathcal{S}} |c_k|^2}{\sigma_x^2 \sum_k |c_k|^2 + \mathbb{E} \left\{ |v[k]|^2 \right\}} \end{aligned} \quad (2)$$

where $\mathcal{S}$ is the set of the locations of the desired largest μ taps in $\mathbf{c}$. The numerator of (2) is the signal power scaled by the power of the μ significant taps in $\mathbf{c}$, and the denominator contains the total received signal power. Ideally, the energy in the significant taps, $\sum_{k \in \mathcal{S}} |c_k|^2$, will make up almost all of the energy in the channel, $\sum_k |c_k|^2$, since we want all other taps to be as close to zero as possible. Ignoring noise for a moment, $0 \leq J_S \leq 1$, and the only way to force $J_S \rightarrow 1$ it to make all but μ taps in $\mathbf{c}$ go to zero. Adding in the noise term to the denominator ensures that the residual self-interference is weighted comparably to the noise, so that the excess taps are not made small at the expense of noise amplification.

The SSSNR in (2) is analogous to Melsa’s Shortening SNR (SSNR) [4], with a few distinctions: the set of desired taps $\mathcal{S}$ is not contiguous, the denominator includes the noise power, and the denominator includes both the desired and undesired taps (rather than just the latter). The last distinction is for numerical reasons, and it can be shown that keeping or omitting the desired taps in the denominator leads to the same solution in this type of problem [11, III.B].

Let $\mathbf{H}$ be the $L_c \times L_w$ tall convolution matrix of $\mathbf{h}$, i.e., a Toeplitz matrix with first column $[\mathbf{h}^T, \mathbf{0}_{1 \times L_w - 1}]^T$ and first row $[h(0), \mathbf{0}_{1 \times L_w - 1}]$. Then let $\mathbf{H}_\mathcal{S}$ be the $\mu \times L_w$ matrix obtained by extracting the μ rows of $\mathbf{H}$ corresponding to the desired nonzero tap locations, $k \in \mathcal{S}$. As an example, for $L_w = 3$, $L_h = 3$, $\mu = 2$, and $\mathcal{S} = \{1, 3\}$,

$$\mathbf{H} = \begin{bmatrix} h_0 & 0 & 0 \\ h_1 & h_0 & 0 \\ h_2 & h_1 & h_0 \\ 0 & h_2 & h_1 \\ 0 & 0 & h_2 \end{bmatrix}, \quad \mathbf{H}_\mathcal{S} = \begin{bmatrix} h_0 & 0 & 0 \\ h_2 & h_1 & h_0 \end{bmatrix}$$

With these definitions, the combined response is $\mathbf{c} = \mathbf{H} \cdot \mathbf{w}$, and the values of the desired nonzero taps are contained in the vector $\mathbf{c}_\mathcal{S} = \mathbf{H}_\mathcal{S} \cdot \mathbf{w}$. Then we have

$$\sum_{k \in \mathcal{S}} |c_k|^2 = \|\mathbf{H}_\mathcal{S} \cdot \mathbf{w}\|^2$$

$$\sum_k |c_k|^2 = \|\mathbf{H} \cdot \mathbf{w}\|^2,$$

and the SSNR from (2) becomes

$$J_S(\mathbf{w}) = \frac{\mathbf{w}^H \mathbf{B}_\mathcal{S} \mathbf{w}}{\mathbf{w}^H \mathbf{C} \mathbf{w}} \tag{3}$$

$$\mathbf{B}_\mathcal{S} = \sigma_x^2 \mathbf{H}_\mathcal{S}^H \mathbf{H}_\mathcal{S}$$

$$\mathbf{C} = \sigma_x^2 \mathbf{H}^H \mathbf{H} + \sigma_n^2 \mathbf{I}.$$

The SSNR expression in (3) is a generalized Rayleigh quotient. The value of $\mathbf{w}$ that maximizes this quantity, i.e., the SSNR-optimal CSF for a given set $\mathcal{S}$, is computed by finding the generalized eigenvector of the matrix pair $(\mathbf{B}_\mathcal{S}, \mathbf{C})$ corresponding to the largest generalized eigenvalue. An algorithm for this general problem is given in [12, Section 8.7.2].

The method used to compute the tap values in [2], which is based on [8], is mathematically similar to our approach, with two key differences. Most importantly, the set $\mathcal{S}$ is fixed in [2]. Second, [2] uses the concept of a TIR. The optimal CSF is written as a function of the TIR, and then the TIR is optimized.

(This is implicit within [2, Equation (25)].) The choice of CSF is “optimal” in the sense that it minimizes the MSE between the outputs of the CSF and TIR, and the TIR is optimal in the sense that it maximizes the signal-to-noise ratio (SNR) at the CSF output. Similar to the channel shortening literature where the minimum MSE and the maximum SSNR channel shorteners are equivalent [13, Section 5], the approach in [2] is mathematically equivalent to our approach with the exception of the fixed sparse coefficient locations. However, the minimum MSE approach is more convoluted to implement, as two filters must be designed rather than one.

The CSF that maximizes (3) is only optimal for a given choice of $\mathcal{S}$. As such, the design of $\mathbf{w}$ involves two issues: picking the best locations for the μ nonzero taps in $\mathbf{c}$ (i.e., choosing the set $\mathcal{S}$), and picking the values of the filter coefficients so that (3) is maximized. The first issue is related to the problem of choosing the optimal delay in linear minimum mean squared error (LMMSE) equalization, which is known to be nontrivial since there is no known expression for the optimal delay [14]. In the classical equalization problem, it is feasible to conduct a brute-force search over the L_c possible delays. Here, however, the problem is considerably more challenging since there are $\binom{L_c}{\mu} = \frac{L_c!}{(L_c-\mu)!\mu!}$ possible choices of $\mathcal{S}$. In this article, we consider three methods of choosing the set $\mathcal{S}$.

- Use the indices of the μ largest magnitude taps of $\mathbf{h}$, as in [2]. This will be referred to as Roy’s tap selection method.
- Try all of the possible combinations. This will be referred to as the combinatorial tap selection method, and it is optimal (though expensive).
- Try the heuristic approach outlined below, which will be referred to as the greedy tap selection method.

The greedy method is as follows. First, set $\bar{\mu} = 1$, and find the location $\mathcal{S}_1$ of a single tap that maximizes the SSSNR. This involves computing $\mathbf{w}$ for all L_c possible tap choices. Next, set $\bar{\mu} = 2$ and $\mathcal{S} = \{\mathcal{S}_1, \mathcal{S}_2\}$. Keep $\mathcal{S}_1$ from the prior step, and find the location $\mathcal{S}_2$ such that the best two-tap channel is produced.

This involves computing $\mathbf{w}$ for each of $L_c - 1$ values. Continue adding one tap at a time until μ locations have been chosen.

Roy's method requires designing a single CSF, although the tap locations are likely far from optimal (as will be demonstrated in Section 7). The combinatorial method requires designing $\frac{L_c!}{(L_c-\mu)!\mu!}$ filters. Finally, the greedy method requires designing $\frac{1}{2}\mu(2L_c - \mu + 1)$ filters. It is far cheaper than the combinatorial method, although its performance approaches that of the combinatorial method, as will be demonstrated in Section 7. For example, with $L_c = 20$ and $\mu = 2$, the greedy method is 4.9 times cheaper; and with $L_c = 25$ and $\mu = 5$, the greedy method is 460 times cheaper than the combinatorial method.

5 Noise coloration

As mentioned previously, even if the noise $n[k]$ is white, the CSF outputs colored noise. To see this, we let $\mathbf{W}$ be a wide Tœplitz convolution matrix corresponding to the filter $\mathbf{w}^T$, and compute the covariance matrix of the noise observed by the BP detector as

$$\begin{aligned} \mathbb{E} \{ \mathbf{v} \mathbf{v}^H \} &= \mathbb{E} \{ (\mathbf{W} \mathbf{n}) (\mathbf{W} \mathbf{n})^H \} = \mathbf{W} \mathbb{E} \{ \mathbf{n} \mathbf{n}^H \} \mathbf{W}^H \\ &= \sigma_n^2 \mathbf{W} \mathbf{W}^H \neq \sigma_n^2 \mathbf{I}. \end{aligned}$$

This is a problem because the BP algorithm assumes white noise. The coloration in the noise will harm the BP performance, potentially making it worse than a classical LMMSE equalizer followed by a simple slicer. Thus, to avoid this pitfall, we consider a penalty term based on the squared autocorrelation of the

output noise, or equivalently of the CSF,

$$\begin{aligned}
J_A(\mathbf{w}) &= \frac{1}{\|\mathbf{w}\|^4 \sigma_n^4} \sum_{l=1}^{L_w-1} |\mathbb{E} \{v_k v_{k-l}^*\}|^2 \\
&= \frac{1}{\|\mathbf{w}\|^4} \sum_{l=1}^{L_w-1} \left| \sum_{m=0}^{L_w-1} w_m w_{m-l}^* \right|^2 \\
&= \frac{1}{\|\mathbf{w}\|^4} \sum_{l=1}^{L_w-1} \sum_{m,n=0}^{L_w-1} w_m w_n^* w_{m-l}^* w_{n-l}
\end{aligned}$$

It can be shown that J_A is equivalent to

$$J_A = \int_0^1 \left(\frac{|W(f)|^2}{\int_0^1 |W(f')|^2 df'} - 1 \right)^2 df + 1,$$

where $2\pi f = \omega$. Thus, J_A penalizes non-flatness of the spectrum of $\mathbf{w}$, since J_A drops to its minimum value of 1 as the spectrum $W(\omega)$ approaches any constant value $\forall f$.

In order to combine J_S , which should be maximized with J_A , which should be minimized, we invert the former and include a relative weighting term,

$$J = J_S^{-1} + \beta J_A. \quad (4)$$

The weight β can be chosen to try to force the minimum of J to be in the proximity of the BER cost surface, J_E . In the next section, we look at the surfaces J_S , J , and J_E in order to visualize the effect of β .

The value of β can be set several ways. The simplest is to try various values of β and get a sense of which values lead to good results for the class of parameter values of interest. For example, for the parameters in our simulations, $\beta \in [0.1, 0.5]$ seems to yield good results. Alternatively, β can be included in the optimization problem. One could search the objective function of (4) for a new value of $\mathbf{w}$ (but without changing β), then occasionally adjust β (but not $\mathbf{w}$) to improve the BER, and repeat. If β is updated on a much slower time scale than $\mathbf{w}$, then the computationally-intensive BER does not have to be evaluated very often during the search.

For a given value of β , (4) can be minimized over $\mathbf{w}$ by any method of unconstrained non-linear optimization. We chose to use the simplex method of [15], since it was already available in Matlab, via the “fminsearch” function.

6 Cost surfaces

To visualize the cost surfaces, consider the following example. The channel is $\mathbf{h} = [1, 0.5, 0.9, 0.3]$, the target number of taps is $\mu = 2$, the SNR is 8 dB, the CSF $\mathbf{w}$ has three taps so $L_c = 6$, and we use the unit norm constraint $\|\mathbf{w}\| = 1$. With this constraint, the CSF lies on a unit sphere, and can be represented in spherical coordinates: $w_0 \triangleq w_x = \cos(\theta) \sin(\phi)$, $w_1 \triangleq w_z = \cos(\phi)$, $w_2 \triangleq w_y = \sin(\theta) \sin(\phi)$. A contour plot of the SSSNR (inverted) is shown in Figure 2; note that this is an amalgamated surface (i.e., maximized across all choices of $\mathcal{S}$). There is symmetry of a sort due to the fact that $\mathbf{w}$ and $-\mathbf{w}$ have the same SSSNR.

The impulse response magnitudes of the channel, SSSNR-optimal CSF, and effective channel are shown in Figure 3, and the $\mu = 2$ significant taps of $\mathbf{c}$ are filled in. While the other taps are small, they are not exactly equal to zero, and will contribute some residual ISI that is left uncorrected by the BP detector. Note that the impulse responses as shown have been normalized to have unit infinity norm. In Figure 4, the i th subplot (counting across the first row then the second row) shows the regions in which that tap is used in the final design, with “dark” indicating that the tap is used. The axes are as in Figure 2 and we see, for example, that the sixth tap is never selected to be one of the nonzero taps in the effective channel. Furthermore, the amalgamated cost surface is highly multimodal, with significant shape changes at the boundaries of each tap’s usage region.

The corresponding amalgamated BER surface, shown in Figure 5 is also highly multimodal. This plot was generated via Monte-Carlo simulation by exhaustively computing the BER at the output of the BP detector for each value of $\mathbf{w}$. Here, there are two minima with nearly equivalent BER of 0.0065 (or, four equivalent minima if one counts symmetric pairs due to the fact that $\mathbf{w}$ and $-\mathbf{w}$ also yield the same

BER). By comparing Figures 2 and 5, we see that the SSSNR and BER surfaces look quite different. We note, however, that there are minima in nearly identical locations. This provides some evidence that, at least for this low dimensional example, the SSSNR is a good proxy for BER.

Figures 6, 7, and 8 consider the added squared autocorrelation penalizing term, and show the amalgamated cost surface of (4) with $\beta = 1$, as well as various optimal CSFs. These include the two local optima of the BER cost surface, the global optimum of the SSSNR, the near-optimum of the SSSNR with heuristically chosen sparse tap locations, and the weighted SSSNR solution with $\beta = 1$. We note that in Figure 6 there are two solutions with nearly identical BER; the heuristic solution is near one whereas the SSSNR solution is near the other, and the two pairs of solutions have different active taps. The goal is to determine which optimal solution is the best proxy for the minimum BER solution. Adding the squared autocorrelation penalizing term moves the SSSNR heuristic solution past the minimum BER solution, and it appears that β should be small, say in the range of 0 to 0.2. To further investigate the affect of β , Figure 9 shows the BER performance of the weighted SSSNR scheme as a function of β at 8 dB SNR, and we see that the optimal value of β occurs at 0.1.

7 Simulations

Having demonstrated that SSSNR is a good proxy for BER, we now compare the BER of the various CSF design approaches. We consider a longer channel than in the low-dimensional example of the previous sections, and we compare SSSNR and computational time among the different design metrics and tap selection methods. Second, we evaluate the BER for two channels employing different sparsening filters in conjunction with BP.

In the first example, we consider the channel $\mathbf{h}_1 = [0.0722, 0, 0, 0.7217, 0.6495, 0, 0, 0.2165, 0, 0.0722]$. We design the CSF to sparsen the channel to $\mu = 2$ taps, we let $L_w = 25$, we transmit uncoded BPSK symbols, and we use ten iterations in the BP detector.

The BER results are shown in Figure 10 for the three tap selection methods outlined at the end

of Section 4, and for comparison we include the performance of a 25-tap classical LMMSE equalizer with a memoryless slicer. The three methods employing a BP detector handily outperform the LMMSE equalizer, and the greedy tap selection approach is able to attain performance nearly equal to the optimal combinatorial method. In addition, we note that those methods which attempt to pick the set of nonzero taps $\mathcal{S}$ to maximize the SSSNR outperform Roy’s method by approximately 1.5 dB for this example.

The complexity, taps selected, and the SSNR of each tap-selection method is displayed in Table 1. The combinatorial approach achieved the best SSSNR performance, but requires the design of 561 filters. On the other hand, the greedy approach achieved almost the same SSSNR performance with much fewer filter designs, being more efficient than the combinatorial method. Roy’s tap selection method needs to design only one filter, but its SSSNR result is inferior to the other two.

In conducting simulations, we noticed that occasionally the hybrid CSF/BP structure yielded BER performance which was inferior to that of a simple linear equalizer with a memoryless slicer. Upon further investigation, it became clear that the performance degradation in such cases was due to noise coloring by the CSF, as addressed in Section 5. We now consider such an example, and show that the use of the modified cost function given in (4) results in flatter sparsening filters, and improves the BER performance. We now consider the channel $\mathbf{h}_2 = [-0.21, -0.36, 0.72, 0.5, 0.21]$, and we again design the CSF to sparsen the channel to $\mu = 2$ taps with a CSF and equalizer of length $L_w = 25$. As before, we transmit uncoded BPSK symbols, and use ten iterations in the BP detector. We also add the squared autocorrelation penalizing term to the combinatorial and greedy SSSNR CSF design metrics. To choose the β value, we performed a grid search with 11 values between 0.0 and 1.0 at a 14 dB SNR and the best value obtained was $\beta = 0.1$.

From the simulation results in Figure 11, it is apparent that approaches without the addition of the squared autocorrelation penalizing term perform worse than a simple linear equalizer, at least at low to moderate SNR. Overall, the addition of the squared autocorrelation penalizing term improves the BER performance for both the combinatorial and greedy approaches, by approximately 0.7 dB. Also, Roy’s tap-selection scheme performs significantly worse than the classical LMMSE equalizer for this example.

The motivation given for Roy’s tap-selection approach [2] was that by choosing the locations of the desired nonzero taps (i.e., the set $\mathcal{S}$) to match the locations of strong arrivals of the incoming signal, there is good “spectral matching” between the channel and the TIR which results in reduced noise enhancement. However, it is unclear that matching dominant tap *locations* between the channel and TIR (regardless of the tap values) results in spectra with a similar shape. In fact, in the example that intended to motivate the reduced noise enhancement of their approach [2, Figure 12], the CSF apparently amplifies the noise by about 40 dB in the region of a deep channel null. Thus, while extensive simulations in [2] have demonstrated significant error-rate improvement when compared with competing approaches that employ decision-feedback equalizers, we note that further BER improvement can be made by wiser design of the CSF.

Examination of the frequency responses of the CSF and combined response for this second simulation scenario provides further evidence that does not support Roy’s reduced noise enhancement claim. Figure 12 depicts the frequency response of each sparsening filter chosen by the combinatorial approach, the combinatorial approach with addition of the penalizing term, and Roy’s method. Roy’s filter amplifies the input signal over the frequencies in the center of the band. Comparing the original channel $\mathbf{h}_2$ with the combined responses as shown in Figure 13 provides another picture showing that Roy’s chosen filter amplifies noise in the center of the band, which contributes to the degraded performance pictured in Figure 11.

Conversely, by adding the penalizing term to the combinatorial SSSNR approach, the resultant sparsening filter becomes flatter, producing an effective channel similar to the original one. This, in fact, reduces the noise enhancement. Moreover, the flatter frequency response shown by the CSF will also reduce noise coloring thereby improving BP detector BER performance. Both reasons explain the error performance improvement and motivate the incorporation of the squared penalizing term to the CSF design. The optimal choice of β however, remains an open issue and is likely to be channel-dependent.

In addition, to provide further evidence of the proposed method’s efficacy, we also considered the ITU Vehicular A channel [16] that has six paths arriving at [0, 310, 710, 1090, 1730, 2510] ns and a power-delay

profile of $[0, -1, -9, -10, -15, -20]$ dB. In our simulations we used a square-root raised cosine pulse and a symbol duration of $T = 80$ ns, which generally resulted in a sparse equivalent discrete channel with average length of 21 taps. Also, we transmit uncoded BPSK symbols, use ten iterations in the BP detector, let $\mu = 2$ non-zero taps, and use a CSF (and for comparison, an equalizer) of length $L_w = 32$. Again, to calculate the β value, we used a grid search at 20 dB SNR and the value obtained was $\beta = 0.2$.

Figure 14 shows the following tap-selection methods in comparison with the linear equalizer: Roy's, combinatorial, greedy and greedy with squared autocorrelation penalizing term. Once again the methods employing BP detectors outperformed the linear equalizer and the method with squared autocorrelation had the best performance. Moreover, at a 10^{-5} BER the greedy method with squared autocorrelation outperforms the combinatorial method by approximately 2 dB.

Finally, we again emphasize that this hybrid detector is quite flexible since its complexity can be adjusted by the system designer. While the complexity scales exponentially with μ , implementations are quite feasible on modern hardware in a wide range of applications [17]. While linear and decision feedback equalizer complexity scales only linearly with the channel length L_h , and are therefore attractive for applications where hardware simplicity is at a premium, the performance advantage offered by the hybrid BP detector (reported here and in [2]) may well be worth the additional complexity. Finally, when compared with traditional Viterbi and BCJR detectors which scale exponentially with L_h , the hybrid BP detector appears to have a considerable advantage in terms of complexity [1].

8 Conclusions

In this work we have considered the design of sparsening filters as a way to reduce the complexity of iterative soft-input soft-output MAP detectors. By designing the sparsening filter so that the combined response of the (possibly non-sparse) channel and filter has a sparse impulse response, i.e., a response with only a handful of significant taps, the use of a BP-based MAP detector becomes feasible for detecting the bits. We proposed a filter design metric called the SSSNR, and showed that maximizing this quantity

serves as a good proxy for minimizing BER. We developed a greedy algorithm for tap selection, and showed that this approach yields near-optimal performance with a significant reduction in complexity when compared to the optimal, combinatorial tap selection approach. In addition, we treated the issue of noise coloration introduced by the sparsening filter, and showed that the addition of a noise penalty term to the SSSNR results in solutions with a flatter frequency response, thereby limiting the amount of noise coloration. Numerical simulations compared our scheme with an existing approach due to Roy, and showed that significant performance gains can be had by intelligently choosing the tap locations.

Future work in this area will consider fractionally-spaced or adaptive versions of this approach, as well as the effects of sparsening filter length on performance. The authors would like to thank Yanjie Peng at WPI for providing the initial version of the simulation code to implement the BP detector.

Competing interests

The authors declare that they have no competing interests.

Acknowledgement

Martin is funded in part by the Air Force Research Labs, Sensors Directorate. The views expressed in this article are those of the authors, and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government. This document has been approved for public release; distribution unlimited.

Endnotes

^aNote that a SIMO system employing either multiple receive antennas or fractional sampling *can* perfectly sparsen the channel under certain conditions on subchannel roots [18]. ^bWhile it is possible that the noise observed at the receiver front-end is not white to begin with, we make the standard AWGN assumption throughout.

References

1. G Colavolpe, G Germe, On the application of factor graphs and the sum-product algorithm to ISI channels. *IEEE Trans. Commun.* **53**(5), 818–825 (2005)
2. S Roy, TM Duman, VK McDonald, Error rate improvement in underwater MIMO communications using sparse partial response equalization. *IEEE J. Ocean. Eng.* **34**(2), 181–201 (2009)
3. DD Falconer, FR Magee, Adaptive channel memory truncation for maximum likelihood sequence estimation. *Bell Sys. Tech. J* **52**, 1541–1562 (1973)
4. PJW Melsa, RC Younce, CE Rohrs, Impulse response shortening for discrete multitone transceivers. *IEEE Trans. Commun.* **44**, 1662–1672 (1996)
5. M Kallinger, A Mertins, Room impulse response shortening by channel shortening concepts, in *Proc. IEEE Asilomar Conf. on Signals, Systems and Comp.*, Pacific Grove, CA, 2005, pp. 898–902
6. G Arslan, BL Evans, S Kiaei, Equalization for discrete multitone receivers to maximize bit rate. *IEEE Trans. Signal Process.* **49**(12), 3123–3135 (2001)
7. K Vanbleu, G Ysebaert, G Cuypers, M Moonen, K Van Acker, Bitrate maximizing time-domain equalizer design for DMT-based systems. *IEEE Trans. Commun.* **52**(6), 871–876 (2004)
8. G Bauch, N Al-Dhahir, Reduced-complexity space-time turbo-equalization for frequency-selective MIMO channels. *IEEE Trans. Wirel. Commun.* **1**(4), 819–828 (2002)
9. A Gomaa, N Al-Dhahir, A new design framework for sparse FIR MIMO equalizers. *IEEE Trans. Commun.* **59**(8), 2132–2140 (2011)
10. M Kaynak, T Duman, E Kurtas, Belief propagation over frequency selective fading channels, in *Proc. IEEE Vehicular Technology Conf. (VTC'04)*, vol. 2, Los Angeles CA , 2004, pp. 1367–1371
11. RK Martin, JM Walsh, CR Johnson Jr., Low complexity MIMO blind, adaptive channel shortening. *IEEE Trans. Signal Process.* **53**(4), 1324–1334 (2005)

12. GH Golub, CF Van Loan, *Matrix Computations* (The Johns Hopkins University Press, Baltimore, 1996)
13. RK Martin, K Vanbleu, M Ding, G Ysebaert, M Milosevic, BL Evans, M Moonen, CR Johnson Jr., Unification and evaluation of equalization structures and design algorithms for discrete multitone modulation systems. *IEEE Trans. Signal Process.* **53**(10), 3880–3894 (2005)
14. L Szczecinski, Low-complexity search for optimal delay in linear FIR MMSE equalization. *IEEE Signal Process. Lett.* **12**(8), 549–552 (2005)
15. JC Lagarias, JA Reeds, MH Wright, PE Wright, Convergence properties of the nelder-mead simplex method in low dimensions. *SIAM J. Optim.* **9**, 112–147 (1998)
16. Guidelines for the evaluation of radio transmission technologies for IMT-2000, Recommendation ITU-R M.1225 (1997)
17. Y Peng, K Zhang, AG Klein, X Huang, Design and implementation of a belief propagation detector for sparse channels, in *Proc. IEEE Intl. Conf. on Application-specific Systems, Architectures and Processors (ASAP'11)*, Santa Monica, CA, 2011, pp. 259–262
18. L Tong, G Xu, T Kailath, Blind identification and equalization based on second-order statistics: a time domain approach. *IEEE Trans. Inf. Theory* **40**(2), 455–466 (1994)

Table 1. Computational complexity, taps selected, and SSSNR achieved at 8 dB SNR

Tap-selection method	Number of filters designed	Taps selected	SSSNR
Combinatorial	561	{17,18}	6.9344
Greedy	67	{18,19}	6.9336
Roy	1	{4,5}	5.7961

Figure 1. System model.

Figure 2. Contours of SSSNR for a three-tap unit norm filter, parameterized by two angles in spherical coordinates. The $\times$ indicates the filter with the highest SSSNR.

Figure 3. Channel, filter, and effective channel tap magnitudes, using the SSSNR-optimal 3-tap filter.

Figure 4. Usage maps of the six taps of the effective channel $\mathbf{c}$. Axes are identical to Figure 2. Dark areas indicate that the given tap is one of the μ largest taps in $\mathbf{c}$.

Figure 5. Contours of BER for a three-tap unit norm filter, parameterized by two angles in spherical coordinates. The ∇ indicates the filters with the lowest BER.

Figure 6. Comparison of different optimal CSFs, overlaid on contours of the surface (4) with $\beta = 1$.

Figure 7. Figure 6, zoomed in on the lower right section.

Figure 8. Figure 6, zoomed in on the middle section.

Figure 9. Bit error rate as a function of β for the weighted SSSNR scheme.

Figure 10. Bit error rate for different CSF design metrics on channel $\mathbf{h}_1$.

Figure 11. Bit error rate with $\mathbf{h}_2$ using same CSF design metrics shown in Figure 10 and addition of the squared autocorrelation penalizing term.

Figure 12. CSF frequency response for combinatorial, combinatorial with addition of penalizing term, and Roy's approach.

Figure 13. Frequency response of original channel and effective channel frequency responses for combinatorial, combinatorial with addition of penalizing term, and Roy's approach.

Figure 14. Bit error rate for the ITU vehicular A channel.




Figure 1




Figure 2






Figure 3







Figure 4




Figure 5





Figure 6






Figure 8




Figure 9




Figure 10




Figure 11




Figure 12




Figure 13




Figure 14