Reduction of Peak-to-Average Power Ratio in Transform Domain Communication Systems

Richard K. Martin and Marshall Haker

***Abstract*—Transform Domain Communication Systems (TDCS) are spectrum-scavenging systems which modulate data by a waveform whose magnitude is non-zero only in unused frequency bins, and whose phase is pseudo-random. This creates a noise-like waveform suitable for secure, low probability of intercept (LPI) communications, as well as spectrum-sharing applications. However, it also creates a high peak-to-average power ratio (PAPR). We investigate PAPR reduction in TDCS by modifying the projection onto complex sets approach for multicarrier systems. Simulations show that adequate PAPR reduction is achieved, transmitter-receiver waveform mismatches are not exacerbated by PAPR reduction, and the resulting waveform is still LPI.**

***Index Terms*—Peak-to-average power ratio, interference suppression.**

I. Introduction

TRANSFORM domain communication systems (TDCS) [1], [2] employ dynamic spectrum access by sensing the available spectrum, and only transmitting on unused frequency bins. This is done by forming a frequency-selective pulse shape that will be used for pulse amplitude modulation (PAM). The magnitude of the pulse shape is determined by the available spectrum, and a pseudo-random phase is assigned to each frequency bin. By using a low transmit power and the pseudo-random phases, the waveform appears noise-like and allows for undetected coexistence. This enables covert communications for the military, hence TDCS have been investigated by the Air Force Research Labs and the Air Force Institute of Technology for some time [1], [3]. Since it transmits in "spectrum white space," TDCS can also be used for dynamic spectrum access in a cognitive radio, as a spectrum overlay secondary user [4]; and since it transmits at low power, it can be used on the control channel in cognitive radios [2].

In a multicarrier system, each time domain sample is a weighted sum (i.e. an IFFT) of many independent random variables, which are often drawn from a QAM constellation. In a TDCS, each time domain sample is a weighted sum of many independent random variables, which have a distribution of a unit magnitude and a (pseudo) random phase. By the central limit theorem, the time-domain samples are approximately Gaussian, yeilding a high peak-to-average power ratio (PAPR) [5], [6], [7]. This is a problem for RF amplifiers, which usually operate in a small linear region. However, we actually desire for the signal to look Gaussian, since for a given transmit power, the source distribution with the highest entropy (and thus with the least apparent stucture to be used for detection) is a Gaussian. (Note that the IFFT does not change the entropy of the entire length-N waveform, but the entropy of an individual sample can be maximized.)

PAPR reduction has been thoroughly studied in the context of multicarrier systems [8]. However, the problem is somewhat different for TDCS, since the resources that can be manipulated are more limited. In [5], the frequency domain block is partitioned into pieces, each piece is weighted, and they are recombined. Similarly, in [9], [10] each block is partitioned, phase rotations are applied, and the order of the pieces is permuted, either within a symbol [9] or across antennas [10]. However, for these methods, side information (the weights) must be known; and in TDCS, if the initial set of used tones differ slightly at the transmitter and receiver (due to imperfect spectrum estimates), the sets of used tones cannot be partitioned in a consistent manner and the final transmitted waveform and receiver matched filter may differ greatly.

In [6], [7], [11], [12] "tone reservation" was used, wherein small values were transmitted on the otherwise unused null tones in a multicarrier system. In [6], the reserved tones were set by alternating clipping in the time domain and restoration of the data tones in the frequency domain; in [11], [12] they were set using a gradient method; and in [7] they were set by first balancing the two largest time-domain spikes, then decreasing them to balance with the third spike, and so on. Similarly, "active constellation extension," [13] injects additional power into the *data* tones. The method of PAPR reduction we will use in this paper is similar to tone reservation or constellation extension, with the main difference being that we will adjust the phases rather than the magnitudes of the signal in the frequency domain. We will adapt the method of [6], although in principle the methods of [7], [11], [12] could also be used. In contrast to multicarrier systems, we have a magnitude constraint (low power on used tones, no power on currently occupied tones) which is why we do not *directly* apply the methods of [6], [7], [11], [12]; and PAPR reduction cannot exacerbate potential waveform mismatches between the transmitter and the receiver, which is why we do not apply the methods of [5], [9], [10]. The contributions of this paper are a modification of [6] suitable for TDCS and a study of the tradeoff of PAPR reduction and the desired

Manuscript received January 14, 2008; revised April 17, 2008 and June 27, 2008; accepted August 24, 2008. The associate editor coordinating the review of this letter and approving it for publication was M. Morelli.

R. Martin (corresponding author) is with the Dept. of Electrical and Computer Engineering, The Air Force Inst. of Technology, Wright-Patterson AFB, Ohio (e-mail: richard.martin@afit.edu).

M. Haker is with the Air Force Research Labs, Sensors Directorate.

R. Martin is funded in part by the Air Force Research Labs, Sensors Directorate. The views expressed in this paper 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.

Digital Object Identifier 10.1109/TWC.2009.080057

properties of TDCS.

We use $\mathbf{1}_{m,n}$ and $\mathbf{0}_{m,n}$ to indicate $m \times n$ matrices of ones or zeros. The vector operations $\| \cdot \|$, $\langle \cdot \rangle$, and $\odot$ are the 2-norm, the sample average, and the Hadamard (element-wise) product; and $(\cdot)^T$, $(\cdot)^*$, and $(\cdot)^H$ indicate transpose, conjugate, and Hermitian transpose. Elements of a vector are given in brackets. $\mathcal{F}$ and $\mathcal{F}^{-1}$ indicate the fast Fourier transform (FFT) and its inverse.

II. System Model

The system model is depicted in Fig. 1. All processing is performed at an oversampling factor of N compared to the symbol rate. The transmitter and receiver independently monitor the spectrum use, and create a spectral mask $\mathbf{m}$ of L ones and $N - L$ zeros (not necessarily contiguous) indicating which of the N frequency bins are available for use. This is usually done by comparing the spectrum occupancy to a threshold, which is set by determining how much interference can be tolerated or detected by the primary user of the spectrum. L is determined by the number of unoccupied bins. Next, a pseudo-random (PR) sequence generator is used to create an $N \times 1$ vector of phases $\boldsymbol{\theta}$ that are distributed across $[0, 2\pi)$, which are applied element-wise to the spectral mask. The resulting vector is passed through an inverse fast Fourier transform (IFFT) and scaled to the desired power level. The PR phases make this waveform appear Gaussian. Thus, the TDCS waveform in the frequency-domain and time-domain (respectively) are the $N \times 1$ vectors

$$\mathbf{X} = \sigma_x \, \mathbf{m} \odot e^{j\boldsymbol{\theta}}, \tag{1}$$

$$\mathbf{x} = \mathcal{F}^{-1}\mathbf{X}, \tag{2}$$

where σ_x^2 is the transmitted power per frequency bin. The vector $\mathbf{x}$ will pass through a digital-to-analog converter to create the continuous-time frequency-selective pulse shape, although we perform an approximate analysis by considering the discrete-time waveform $\mathbf{x}$. Then $\mathbf{x}$ is used as the pulse shape for PAM. (Cyclic shift keying (CSK) is often used in TDCS systems, but since CSK and PAM waveforms with the same pulse shape have the same PAPR, it is sufficient to consider PAM in this paper for simplicity.) Finally, $\mathbf{x}$ is passed through a high power amplifier (we use that of (7) in [14] with nonlinearity parameter $p = 2$), and then upconverted to passband using a single carrier frequency.

This process is performed separately at the transmitter and receiver, with the assumption that the spectral occupancy appears the same to both and that both use the same PR phase sequence. If there are differing spectrum estimates, perhaps because multipath channels from the primary user of the spectrum induce different fading on the band edges at the TDCS transmitter and receiver, there may be a slight mismatch between the transmit waveform and the matched filter used by the receiver. In this case, $\mathbf{x}$, $\mathbf{m}$, $\boldsymbol{\theta}$, and L will have subscripts "tx" and "rx" as needed; whereas a lack of such a subscript indicates agreement between the transmitter and the receiver. The PR sequence is always the same at the transmitter and receiver.

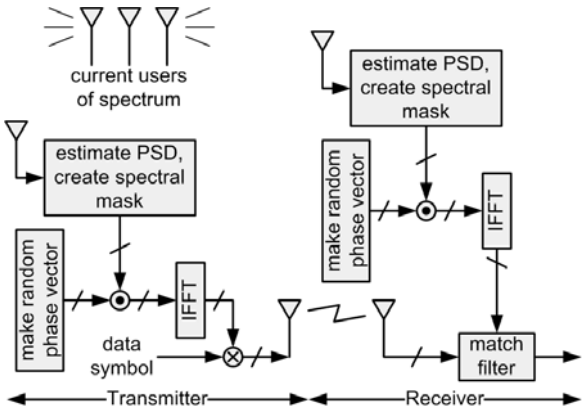

Fig. 1. Generation of a TDCS pulse shape waveform at the transmitter and (ideally) identical matched filter at the receiver. Inverted triangles are antennae, an arrow with a slash indicates a vector of data, and $\odot$ indicates the Hadamard product. Digital to analog conversion and modulation/demodulation are not depicted, to focus attention on the generation of the pulse shape.

III. PAPR Reduction

Gatherer and Polley [6] used a method called "projection onto complex sets (POCS)" which exploits the structure of multicarrier systems. We will first review POCS in the context of multicarrier systems, then modify it for use in TDCS. This PAPR reduction must be performed at the transmitter and receiver, in order to keep the waveform at both ends similar. Many multicarrier systems have tones that carry no data, called "null tones." They may be used as a guard band, or they may have insufficient SNR to effectively carry data. In IEEE 802.11a wireless LANs, for example, 12 of the 64 tones are null tones; 11 of which form a guard band, and one of which is the DC tone. The transmitted signals on these tones thus have no importance to the receiver, and they can be modified freely in an effort to reduce the PAPR.

Consider S_1, the set of time domain vectors $\mathbf{x}$ corresponding to input vectors $\mathbf{X}$ with the original data values in place but possibly non-zero values on the null tones; and S_2, the set of time domain vectors $\mathbf{x}$ whose maxima are below a clipping threshold c. POCS involves alternating projections onto the sets S_1 and S_2 [6]. If the sets truly are convex as is the case for S_1 and S_2, POCS is globally convergent to a time domain vector $\mathbf{x}$ in the intersection of S_1 and S_2 if the intersection is non-empty. The choice of c dictates whether or not a solution exists.

The POCS PAPR reduction algorithm for multicarrier systems is as follows:

1. IFFT the initial vector $\mathbf{X}$ to get a point $\mathbf{x}$ in S_1.
2. Clip $\mathbf{x}$ using the threshold c. If $\mathbf{x}$ changes, go to step 3. Otherwise, terminate.
3. FFT $\mathbf{x}$ and restore the data elements (but not the null tones) of $\mathbf{X}$ to their original values.
4. IFFT to get a new $\mathbf{x}$, then go to step 2.

Steps 2 and 3 each implement a projection. Each iteration requires an FFT and an IFFT, and the algorithm must be run in full for each transmitted block. The clipping threshold is related to a threshold on the PAPR via

$$\text{PAPR}_o = \frac{Nc^2}{L\sigma_x^2}. \tag{3}$$

To modify the POCS method for use in TDCS, note that there are no null tones in TDCS, so set S_1 must be redefined. The parameters at our disposal are the phases on the used tones, and the constraint is that the magnitude of the spectral mask must be maintained. Enforcing unit magnitude on the used tones is the simplest method of enforcing a flat power spectrum on the used tones. This also keeps any single block from having spikes in frequency that would be more likely to be detected or cause interference. Consider the set

$$\hat{S}_1 = \left\{ \mathbf{x} \left| \mathcal{F}\mathbf{x} = \begin{bmatrix} \mathbf{1}_{L,1} \odot e^{j\boldsymbol{\theta}} \\ \mathbf{0}_{N-L,1} \end{bmatrix} \right. \right\}, \tag{4}$$

where $\boldsymbol{\theta}$ is any phase vector and for simplicity of notation, the used frequency bins are assumed to be contiguous. Now POCS can be used by simply substituting $\hat{S}_1$ for S_1, and by replacing step 3 of the POCS algorithm with

$\hat{3}$. FFT $\mathbf{x}$ and restore the magnitudes (but not the phases) of $\mathbf{X}$ to their original values.

However, there is a caveat: as opposed to S_1 and S_2, the set $\hat{S}_1$ is not actually convex, hence the method is not *guaranteed* to find a solution. To see this, note that S is said to be convex if and only if for all $\mathbf{x}_1, \mathbf{x}_2 \in S$ and for all $\alpha \in [0,1]$, the vector $\mathbf{x}_{avg} = \alpha\mathbf{x}_1 + (1-\alpha)\mathbf{x}_2 \in S$. In the context of TDCS, using (4) we have

$$\mathcal{F}\mathbf{x}_{avg} = \begin{bmatrix} \mathbf{1}_{L,1} \odot \left(\alpha e^{j\boldsymbol{\theta}_1} + (1-\alpha) e^{j\boldsymbol{\theta}_2} \right) \\ \mathbf{0}_{N-L,1} \end{bmatrix}, \tag{5}$$

so $\mathbf{x}_{avg} \notin \hat{S}_1$ and $\hat{S}_1$ is not convex. Thus, we refer to the algorithm in this section as "Approximate POCS (APOCS)," since the convexity is only approximate. Although convergence of APOCS is not guaranteed, our simulations never failed to converge, and the next section gives a heuristic justification for this convergence.

To visualize APOCS, consider a toy problem with $N = 4$, $L = 3$, and

$$\mathbf{x} = \mathcal{F}^{-1} \left[e^{j\theta_1}, e^{j\theta_2}, e^{j\theta_3}, 0 \right]^T \tag{6}$$

$$= e^{j\theta_3} \mathcal{F}^{-1} \left[e^{j(\theta_1-\theta_3)}, e^{j(\theta_2-\theta_3)}, 1, 0 \right]^T. \tag{7}$$

The PAPR of $\mathbf{x}$ depends on the magnitude of its elements, hence it only depends on two phase differences rather than three phases. Fig. 2 shows contours of the PAPR as a function of the phase differences, as well as two sample trajectories of APOCS for a target PAPR of 1.5. Observe that APOCS does not approximate a gradient descent algorithm, since the trajectories are not quite perpendicular to the contours. However, each iteration does drop to a lower contour.

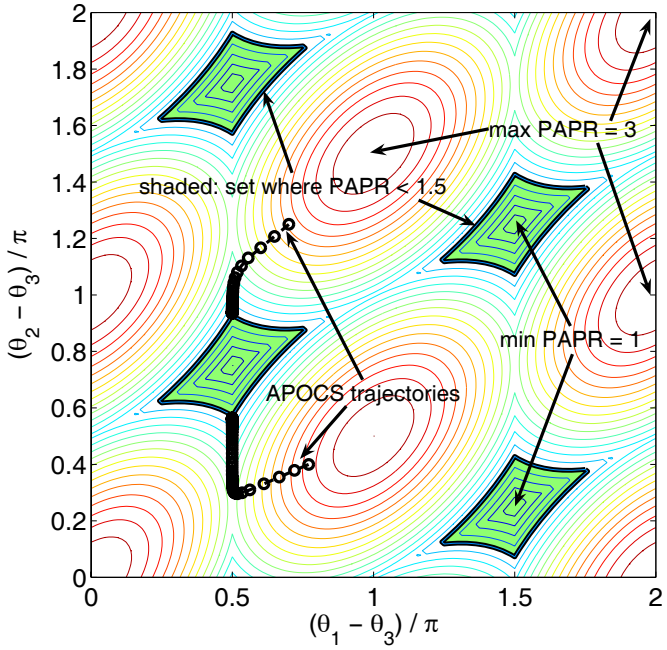

Fig. 2. Contours of the PAPR of $\mathbf{x}$, where $\mathcal{F}\mathbf{x} = \left[e^{j\theta_1}, e^{j\theta_2}, e^{j\theta_3}, 0 \right]^T$. There are also two sample trajectories of the APOCS algorithm. Note that one of the phases, e.g. θ_3, can be factored out of $\mathcal{F}\mathbf{x}$, hence the signal power depends only on two phase differences rather than three phases.

IV. Enforcing Convexity

In this section, we show that APOCS could be slightly modified to be globally convergent, heuristically implying that APOCS is often globally convergent. The intent is not to actually use the new algorithm, but to use it to justify the convergence of APOCS.

Consider a single complex-valued element $X[k]$ of the initial vector $\mathbf{X}$. Let $X_\perp[k]$ be a unit magnitude vector geometrically (not statistically) orthogonal to $X[k]$ in the complex plane,

$$X_\perp[k] = \frac{X_{\mathcal{I}}[k] - jX_{\mathcal{R}}[k]}{|X[k]|}, \tag{8}$$

where subscripts $\mathcal{R}$ and $\mathcal{I}$ denote real and imaginary parts. Let $\mathbf{X}_\perp$ be the vector with elements $X_\perp[k]$, and let $\mathbf{b}$ be a vector of real-valued weights. Consider the set

$$\tilde{S}_1(\mathbf{X}) = \{ \mathbf{x}' \,|\, \mathcal{F}\mathbf{x}' = \mathbf{X} + \mathbf{b} \odot \mathbf{X}_\perp, |b[k]| \le \epsilon \ \forall k \}. \tag{9}$$

For small ϵ, this approximates letting the phase of each $X[k]$ vary by a small amount, although there will be a very small amount of magnitude distortion as well. The set $\tilde{S}_1$ will depend on the original waveform $\mathbf{X}$. To show that $\tilde{S}_1$ is convex, let $\mathbf{x}_1, \mathbf{x}_2 \in \tilde{S}_1$ and $\mathbf{x}_{avg} = \alpha\mathbf{x}_1 + (1-\alpha)\mathbf{x}_2$. Taking the Fourier transform of $\mathbf{x}_{avg}$,

$$\begin{aligned} \mathbf{X}_{avg} &= \alpha(\mathbf{X} + \mathbf{b}_1 \odot \mathbf{X}_\perp) + (1-\alpha)(\mathbf{X} + \mathbf{b}_2 \odot \mathbf{X}_\perp) \\ &= \mathbf{X} + \underbrace{(\alpha\mathbf{b}_1 + (1-\alpha)\mathbf{b}_2)}_{\mathbf{b}_{avg}} \odot \mathbf{X}_\perp. \end{aligned} \tag{10}$$

Since $\mathbf{x}_1, \mathbf{x}_2 \in \tilde{S}_1$, we have $|b_1[k]| \le \epsilon$ and $|b_2[k]| \le \epsilon$ for all k. From the triangle inequality,

$$\begin{aligned} |b_{avg}[k]| &\le \alpha |b_1[k]| + (1-\alpha) |b_2[k]| \\ &\le \alpha\epsilon + (1-\alpha)\epsilon \ = \epsilon. \end{aligned} \tag{11}$$

Thus, $\mathcal{F}\mathbf{x}_{avg} = \mathbf{X} + b_{avg} \odot \mathbf{X}_\perp$ with $|b_{avg}[k]| \le \epsilon \ \forall k$, so $\mathbf{x}_{avg} \in \tilde{S}_1$ and $\tilde{S}_1$ is convex.

To project the output of the clipping step, $\mathbf{x}_{clip}$, onto this set, we must find the vector in $\tilde{S}_1$ that is closest in Euclidean distance, in either the time or frequency domain. Let the FFT of the clipped waveform be $\mathbf{X}_{clip} = \mathbf{X} + \boldsymbol{\Delta}$, where $\boldsymbol{\Delta} =$

$\mathbf{e} \odot \mathbf{X} + \mathbf{b}_c \odot \mathbf{X}_\perp$ for some $\mathbf{e}$ and $\mathbf{b}_c$; and let the projection (to be determined) be $\mathbf{X}_{pro} = \mathbf{X} + \mathbf{b} \odot \mathbf{X}_\perp$. The distance is minimized via

$$\begin{aligned} \mathbf{b} &= \arg \min_{\mathbf{b}, |b[k]| \le \epsilon} \|\mathbf{X}_{clip} - \mathbf{X}_{pro}\|^2 \\ &= \arg \min_{\mathbf{b}, |b[k]| \le \epsilon} \sum_{k=1}^{N} \left((b_{clip}[k] - b[k]) X_\perp[k] + e[k] X[k] \right)^2 . \end{aligned} \tag{12}$$

The squared terms can be minimized separately for each value of k, via

$$\begin{aligned} b[k] &= \arg \min_{b[k], |b[k]| \le \epsilon} \left((b_{clip}[k] - b[k]) X_\perp[k] + e[k] X[k] \right)^2 \\ &= \begin{cases} b_{clip}[k], & |b_{clip}[k]| \le \epsilon \\ \epsilon \ \text{sign}\left(b_{clip}[k]\right), & |b_{clip}[k]| > \epsilon \end{cases} . \end{aligned} \tag{13}$$

Thus, the projection first restricts each $X[k]$ to lie in the subspace given by $b[k]X_\perp[k]$ for *any* value of $b[k]$, then it restricts the magnitude of $b[k]$ to be at most ϵ. We call the resulting algorithm "Enforced POCS (EPOCS)." It involves substituting $\tilde{S}_1$ for S_1, and by replacing step 3 of the POCS algorithm with

3̃. FFT $\mathbf{x}$. For each element k, restore the component in the $X[k]$ direction to its original value and clip the component in the $X_\perp[k]$ direction to length ϵ.

$\tilde{S}_1$ is very similar to $\hat{S}_1$ for small ϵ, since $\mathbf{X}_\perp$ leads primarily to phase rotations of elements of $\mathbf{X}$ with very little magnitude distortion; hence APOCS and EPOCS will lead to similar waveforms. Since $\tilde{S}_1$ is convex, EPOCS is globally convergent, and we infer that its approximation (APOCS) is usually globally convergent. Since EPOCS induces slight violation of the magnitude constraints, is slightly more computationally complex than APOCS, and performs slightly worse than APOCS (as shown in Section VI), it is preferable to implement APOCS rather than EPOCS.

V. Performance Metrics

Our goals are (i) reduce the PAPR, (ii) maintain correlation between the transmit and receive waveforms, and (iii) maintain the noise-like quality of the waveform. For (i), PAPR will be assessed via the complimentary cumulative distribution function (CCDF),

$$\text{CCDF}\left(\text{PAPR}_o\right) = P\left[\text{PAPR} > \text{PAPR}_o\right]. \tag{14}$$

For (ii), there may be some mismatch between the transmit and receive spectral masks, so the algorithm starts at two different points. The final waveforms could differ significantly, disrupting the matched filter at the receiver. We measure this via the correlation coefficient,

$$\rho = \frac{\mathbf{x}_{tx}^H \mathbf{x}_{rx}}{\|\mathbf{x}_{tx}\| \|\mathbf{x}_{rx}\|}, \tag{15}$$

The effective SNR at the receiver is reduced by a factor of ρ, increasing the bit error rate.

Even if the transmitter and receiver use different tones, the initial phases on the used tones match. This is easily enforced by always creating N phases per symbol, even though some are not used. Before PAPR reduction, if the transmitter uses L_{tx} out of N frequency bins and the receiver uses L_{rx} out of N frequency bins, with L bins in common, then the initial value is

$$\rho_{init} = \frac{L\sigma_x^2}{\sqrt{L_{tx}}\sigma_x \cdot \sqrt{L_{rx}}\sigma_x} = \frac{L}{\sqrt{L_{tx}L_{rx}}} \le 1. \tag{16}$$

After PAPR reduction, the transmit and receive phases will differ slightly, yielding

$$\rho_{fin} = \frac{\sum_{i \in \mathcal{M}} e^{j(\theta_{i,rx} - \theta_{i,tx})}}{\sqrt{L_{tx}L_{rx}}}, \tag{17}$$

where $\mathcal{M}$ is the set of L tones that are used by both the transmitter and receiver. Define $\widetilde{\theta}_i = \theta_{i,rx} - \theta_{i,tx}$, which is zero on average. PAPR reduction decreases ρ by a factor of

$$\eta = \frac{\rho_{fin}}{\rho_{init}} = \frac{1}{L} \sum_{i \in \mathcal{M}} e^{j\widetilde{\theta}_i} \approx 1 - \frac{1}{2} \left\langle \widetilde{\theta}_i^{\,2} \right\rangle \tag{18}$$

by a Taylor series expansion. Recall that $\langle \cdot \rangle$ denotes sample average. The more the PAPR is reduced, the larger the variance of the phase discrepancy is likely to be, and the more ρ is decreased, as shown in Section VI.

Goal (iii) can be assessed via the Lilliefors [15] and Jarque-Bera [16] tests for Normality. The Lilliefors test computes the empirical CCDFs of the data (samples of $\mathbf{x}$) and of a Gaussian whose mean and variance match the data. Then the test determines whether the maximum discrepancy between the two CCDF's is large enough to be statistically significant. The Jarque-Bera test computes the skewness S and kurtosis K (normalized third and fourth central moments of the data, respectively), and computes the test statistic

$$JB = \frac{N}{6} \left(S^2 + \frac{(K-3)^2}{4} \right), \tag{19}$$

where N is the number of observations (the FFT size in our case). For a Gaussian, S and $K - 3$ should be zero. Thus, increasing values of JB indicate increasing departure from Normality.

VI. Simulations

We now simulate APOCS and quantify the resulting PAPR distribution. We also examine whether mismatches between the transmitter's and receiver's spectrum estimates are magnified by APOCS, and we examine the probability that the transmitted PAPR-reduced waveform still passes the Lilliefors and Jarque-Bera tests for normality. The TDCS has a spectral mask of

$$\mathbf{m} = \left[\mathbf{1}_{1,39}, \mathbf{0}_{1,20}, \mathbf{1}_{1,10}, \mathbf{0}_{1,20}, \mathbf{1}_{1,39}\right]^T, \tag{20}$$

hence $L = 88$ out of $N = 128$ frequency bins are used. From (3), since PAPR ≥ 1, for this frequency allocation the clipping threshold $c \ge 0.829$. We tested $c \in \{1, 1.25, 1.5\}$, i.e. $\text{PAPR}_o = 1.45c^2 \in \{1.45, 2.27, 3.27\}$. In Fig. 3, the CCDF is shown for three thresholds, for $i = 4$ and $i = 10$ iterations each. 10^5 different waveform realizations were used. With 10 or more iterations, the PAPR is very nearly hard-limited to the threshold, and with as few as four iterations, the majority of the PAPR reduction is obtained. For two of the curves, we also simulated EPOCS, obtaining almost identical results. We used

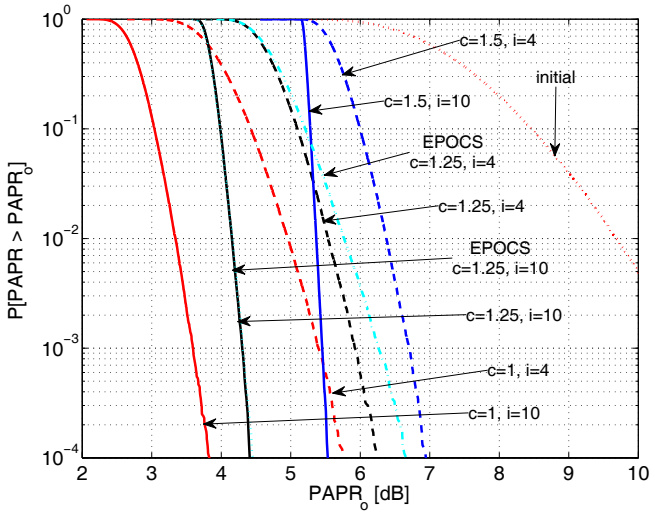

Fig. 3. Complimentary CDF of the PAPR for three different thresholds c and two different numbers of iterations i. All curves are for APOCS except for the two explicitly labelled "EPOCS."

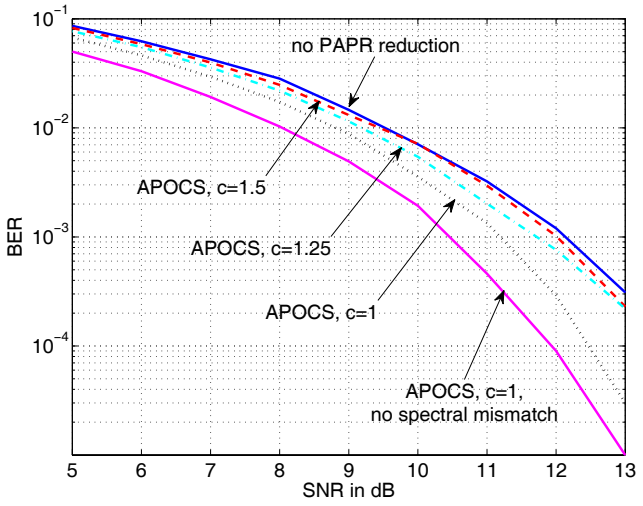

Fig. 5. BER for no PAPR reduction; for PAPR reduction with $c = 1.5$, $c = 1.25$, and $c = 1$ in the presence of spectral mismatch; and PAPR reduction with $c = 1$ with no spectral mismatch.

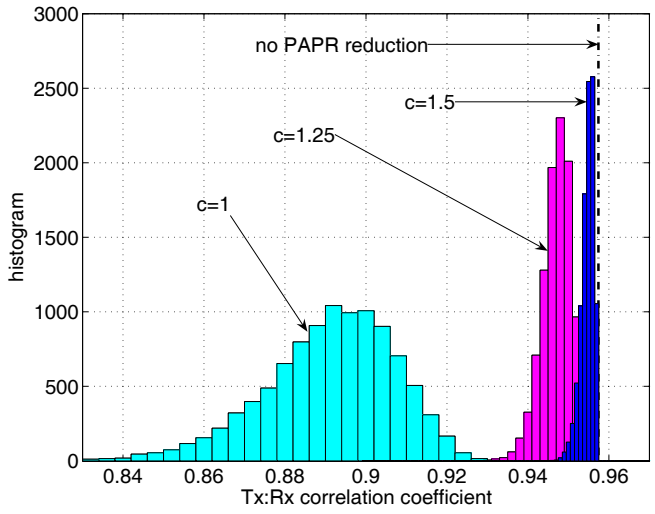

Fig. 4. Histogram of the magnitude of the correlation coefficient ρ between the transmitted and received waveforms, after APOCS. The dash-dot line is the correlation coefficient before PAPR reduction (note that in this case, *all* 10^4 samples fall into the same bin). The histograms have different size bins to keep their heights comparable. The effective SNR at the receiver is reduced by a factor of ρ, increasing the BER.

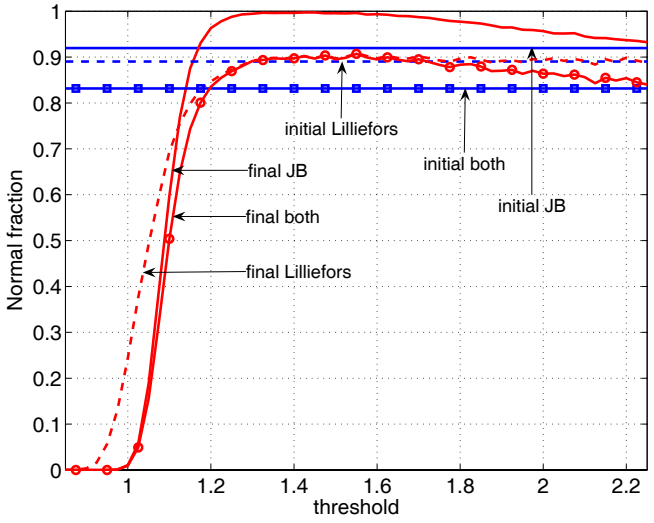

Fig. 6. Results of Lilliefors and Jarque-Bera testing for normality, for the initial TDCS waveforms and the outputs of the PAPR reduction for various clipping thresholds. "Both" indicates the fraction of symbols that tested as Normal under both tests at once.

$\epsilon = 0.1$, leading to a magnitude distortion of up to 0.5% on each tone. Note that EPOCS performs slightly worse in PAPR than APOCS and leads to slight violation of the magnitude constraints, which is why APOCS is preferred.

Now consider the effects of a spectral mismatch between the transmitter and receiver, where the transmitter uses the spectral mask of (20), and the receiver uses the spectral mask of

$$\mathbf{m}_{rx} = [\mathbf{1}_{1,41}, \mathbf{0}_{1,16}, \mathbf{1}_{1,14}, \mathbf{0}_{1,16}, \mathbf{1}_{1,41}]^T, \qquad (21)$$

although the PR phases are the same at both transmitter and receiver. This could occur if either the transmitter or receiver imprecisely identifies the band edges of occupied spectrum. Fig. 4 shows a histogram of $|\rho|$ of (15), computed over 10^4 independent trials. Before PAPR reduction, we had $|\rho| = 0.957$ due to the difference in the number of used tones. After 100 iterations of APOCS, the correlation only decreased by a small amount; although it dropped as low as 0.802 in some cases, when the PAPR reduction was greatest. Fig. 5 shows the bit error rate (BER) versus SNR with and without PAPR reduction, in the presence of both the spectral mismatch described above and a nonlinear amplifier, for an AWGN (flat-fading) channel. The PAPR reduction improves the BER by reducing distortion to the transmitted waveform, while the spectral mismatch degrades the BER; but overall PAPR reduction helps. However, note that the simulation parameters were chosen to be extreme, and in many cases the BER differences will be smaller than shown in Fig. 5. More common benefits of PAPR reduction include improving power efficiency and reducing out-of-band interference [8].

Fig. 6 shows the results of the Lilliefors and Jarque-Bera

tests, before and after PAPR reduction. This indicates the fraction of the waveforms (out of 10^4) that tested as Normal, and the curve "both" indicates the fraction of the waveforms that simultaneously satisfied both tests. For thresholds above 1.25, almost all of the waveforms were normal, but below this value, the fraction dropped off quickly. Oddly, between thresholds of 1.2 to 2.2, PAPR reduction appears to actually increase the chance that the signal appears Gaussian. This is because reducing the PAPR reduces the kurtosis, but the JB test statistic of (19) depends on the square of the excess kurtosis. Thus, if the average kurtosis decreases monotonically with decreasing thresholds, the average JB test statistic need not increase monotonically.

VII. Conclusions

We have investigated the POCS method of PAPR reduction in TDCS. Despite the fact that the set of the desired waveforms is not convex, our modified APOCS algorithm significantly reduces the PAPR. When the transmitter and receiver have mismatched spectral estimates, the magnitude of the correlation coefficient between the transmitter's and receiver's waveforms drops, but not severely (from 0.96 to 0.80 in the worst case we have observed). By comparing the amount of PAPR reduction for various thresholds in Fig. 3, the amount of possible waveform decorrelation for various thresholds in Fig. 4, and the likelihood of retaining a Gaussian appearance for various thresholds in Fig. 6, a value of the threshold that best balances the desired criteria can be selected.

References

[1] V. D. Chakravarthy, A. K. Shaw, M. A. Temple, and J. P. Stephens, "Cognitive radio—an adaptive waveform with spectral sharing capability," in *Proc. Wireless Communications and Networking Conference*, vol. 2, Mar. 2005, pp. 724–729.

[2] C. Han, J. Wang, Y. Yang, and S. Li, "Addressing the control channel design problem: OFDM-based transform domain communication system in cognitive radio," *Elsevier Computer Networks J.*, vol. 52, no. 4, pp. 795–815, Mar. 2008.

[3] V. Chakravarthy, A. S. Nunez, J. P. Stephens, A. K. Shaw, and M. A. Temple, "TDCS, OFDM, and MC-CDMA: a brief tutorial," *IEEE Commun. Mag.*, vol. 43, no. 9, pp. S11–S16, Sept. 2005.

[4] Q. Zhao and A. Swami, "A survey of dynamic spectrum access: signal processing and networking perspectives," in *Proc. IEEE Int. Conf. on Acoustics Speech and Signal Processing*, vol. 4, Hononlulu, HI, Apr. 2007, pp. 1349–1352.

[5] R. W. Bauml, R. F. H. Fischer, and J. B. Huber, "Reducing the peak-to-average power ratio of multicarrier modulation by selected mapping," *Electron. Lett.*, vol. 32, pp. 2056–2057, Oct. 1996.

[6] A. Gatherer and M. Polley, "Controlling clipping probability in DMT transmission," in *Proc. Asilomar Conf. on Signals, Systems, and Computers*, Pacific Grove, CA, Nov. 1997.

[7] B. S. Krongold and D. L. Jones, "An active-set approach for OFDM PAR reduction via tone reservation," *IEEE Trans. Signal Processing*, vol. 52, no. 2, pp. 495–509, Feb. 2004.

[8] S. H. Han and J. H. Lee, "An overview of peak-to-average power ratio reduction techniques for multicarrier transmission," *IEEE Wireless Commun.*, vol. 12, no. 2, pp. 56–65, Apr. 2005.

[9] M. Tan and Y. Bar-Ness, "OFDM peak-to-average power ratio reduction by combined symbol rotation and inversion with limited complexity," in *Proc. IEEE Global Comm. Conf.*, Dec. 2003, pp. 605–610.

[10] M. Tan, Z. Latinović, and Y. Bar-Ness, "STBC MIMO-OFDM peak-to-average power ratio reduction by cross-antenna rotation and inversion," *IEEE Commun. Lett.*, vol. 9, no. 7, pp. 592–594, July 2005.

[11] J. Tellado, "Peak to average power reduction for multicarrier modulation," Ph.D. dissertation, Stanford University, 2000.

[12] A. T. Erdogan, "A low complexity multicarrier PAR reduction approach based on subgradient optimization," *Signal Processing*, vol. 86, pp. 3890–3903, Dec. 2006.

[13] B. S. Krongold and D. L. Jones, "PAR Reduction in OFDM via active constellation extension," in *Proc. Int. Conf. on Acoustics, Speech, and Signal Processing*, vol. 4, Hong Kong SAR, China, Apr. 2003, pp. 525–528.

[14] T. G. Pratt, N. Jones, L. Smee, and M. Torrey, "OFDM link performance with companding for PAPR reduction in the presence of non-linear amplification," *IEEE Trans. Broadcasting*, vol. 52, no. 2, pp. 261–267, June 2006.

[15] H. Lilliefors, "On the Kolmogorov-Smirnov test for normality with mean and variance unknown," *J. Amer. Statistical Association*, vol. 62, pp. 399–402, June 1967.

[16] C. M. Jarque and A. K. Bera, "Efficient tests for normality, homoscedasticity and serial independence of regression residuals," *Economics Lett.*, vol. 6, no. 3, pp. 255–259, 1980.