An Optimized Transmitter Precoding Scheme for Synchronous DS-CDMA

Erik S. Hons, Amir K. Khandani and W. Tong

Department of Electrical and Computer Engineering

University of Waterloo

Waterloo, Ontario, Canada, N2L 3G1

Technical Report UW-E&CE 02-08

May 30, 2002

Abstract

This article[1] presents a technique to reduce the multiple access interference in the forward-link of a conventional DS-CDMA system by applying an energy constrained transformation to the transmitter output. For each symbol period, a new transformation is selected which minimizes the mean-squared error at the output of a bank of matched-filter detectors. This selection is subject to a constraint on total transmitted energy. It is shown that the proposed method results in substantial advantages over earlier similar techniques. The proposed algorithm can be implemented efficiently at the transmitter with existing optimization techniques that solve the quadratic trust-region subproblem.

Keywords

code division multiple access, interference cancelation, transmitter precoding, multiple access channels, trust region, quadratic optimization

I. Introduction

THE conventional DS-CDMA detector assumes an independent additive Gaussian noise model for the Multiple Access Interference (MAI). In reality, the MAI term is highly structured which makes this model invalid. Multiuser detection exploits the structure of the MAI to improve performance at the cost of additional processing overhead at the receiver which is usually quite large. This overhead is especially problematic in the forward link of a cellular mobile network where the receiver is usually a highly resource constrained mobile unit.

Recently, approaches have been proposed which transfer some or all of the processing burden from the receiver to the transmitter (see [2], [3], [4], [5]). These methods work by applying a linear transformation at either the transmitter side [2],[4], [5], or at both the transmitter and receiver sides [3]. By using the minimum mean-squared error criterion, the transformations can be selected to minimize the bit error rate due to the MAI. Reference [4] studies a transmitter precoding scheme tailored to exploit multi-path components, as well as a second approach based on joint optimization of spreading sequences.

The main framework for the methods based on *processing solely on the transmitter side* is the so-called *transmitter precoding* of [2]. In this paper, we follow a structure similar to [2] which is used as a baseline for comparison.

It is shown in [2] that transmitter precoding results in complete decorrelation of the transmitted symbols. In a noiseless channel, this is equivalent to the well known decorrelating detector [6]. However, by applying the decorrelation operation at the transmitter before noise is added to the signal,

[1] A more detailed version of this work is reported in [1].

transmitter precoding avoids the problem of noise enhancement. Two methods are considered in [2] to handle the increased transmitted energy of the precoder: so called *unconstrained* and *constrained* transmitter precoding. The unconstrained method simply scales the transmitter output to the appropriate average energy. In contrast, the constrained method imposes an average energy constraint on the selection of the precoding transformation. It is found that the raw bit error rate for both methods is similar and thus the higher overhead of the constrained method is not justified.

In this work, we revisit constrained transmitter precoding for synchronous DS-CDMA with a new formulation in which the transmitted energy is capped for each symbol period. The MAI will be minimized in an MMSE sense subject to this constraint, thus we refer to the technique as the "*optimizing precoder*". The focus will be on a channel with Forward Error Correction (FEC) where it is demonstrated that constrained transmitter precoding performs significantly better than in the uncoded case and in fact outperforms unconstrained transmitter precoding (this is different from the conclusion reached in [2] for the uncoded case). It will also be shown that the optimizing precoder not only outperforms the constrained precoder reported in the earlier works, but also is easily solved with efficient algorithms from non-linear optimization.

We first present a model for the target system in section II and the formulation for the various precoders in section III. Solution algorithms for the optimizing precoders are presented in section IV. Finally, results and conclusions are presented in sections V and VI, respectively.

II. SYSTEM MODEL

The transmitted signal, $x(t)$, in the forward link of a symbol and chip synchronous DS-CDMA system with K active users and symbol duration T_b is,

$$x(t) = \sum_{n=1}^{K} A_n s_n(t) b_n, \qquad 0 \le t \le T_b, \qquad 1 \le n \le K, \tag{1}$$

where A_n is the transmitter gain for the nth user and b_n, $s_n(t)$ are the n^{th} user antipodal modulated data symbol and signature waveform, respectively. We assume that the data symbols are binary and equi-probable. Additionally, we assume that the signature waveforms are linearly independent, zero outside the range $[0, T_b]$, and normalized to have unit energy so that,

$$\int_0^{T_b} s_n^2(t)\, dt = 1, \qquad 1 \le n \le K. \tag{2}$$

In the additive white Gaussian noise (AWGN) channel, the received waveform is,

$$r(t) = x(t) + n(t), \qquad 0 \le t \le T_b, \tag{3}$$

where $n(t)$ is a Gaussian process with zero mean and power spectral density of $\sigma^2 = N_0/2$. The matched-filter output for user n is the scalar y_n where,

$$y_n = \int_0^{T_b} r(t) s_n(t) dt, \qquad 1 \le n \le K. \tag{4}$$

If we combine the matched-filter outputs to form the vector $\mathbf{y} = [y_1, \cdots, y_K]^t$, then the set of outputs can be described in matrix notation as,

$$\mathbf{y} = \mathbf{RAb} + \mathbf{n}, \tag{5}$$

where $\mathbf{b} = [b_1, \ldots, b_K]^t$ is the vector of modulated data symbols, $\mathbf{A} = \mathrm{diag}\{A_1, \ldots, A_K\}$ is the set of amplitudes, $\mathbf{R} = [R_{i,j}]$ is the $K \times K$ cross-correlation matrix for the current set of signature waveforms where,

$$R_{i,j} = \int_0^{T_b} s_i(t) s_j(t) dt. \tag{6}$$

and $\mathbf{n} = [n_1, \cdots, n_K]^t$ is a Gaussian noise vector whose elements have zero mean and covariance matrix $\mathbf{R}$. The detection strategy which gives the minimum bit error rate selects the vector $\mathbf{b}$ with maximum-likelihood given the set of observations $\mathbf{y}$. If the off-diagonal elements of $\mathbf{R}$ are zero (implying the spreading codes are orthogonal), and the data symbols are equi-probable (as assumed), then it is well known that the optimal decision rule is $\mathrm{sgn}(\mathbf{y})$ which is equivalent to a zero threshold decision rule.

For simplicity, we consider a binary phase-shift keyed (BPSK) system for which the signature waveforms are composed of square waveforms called "chips" which have uniform duration T_c, where T_c is typically much smaller than T_b. With this structure, the normalized signature waveforms $s_n(t)$ can be represented as the binary code vectors $\mathbf{s}_n \in \{\frac{-1}{\sqrt{L}}, \frac{1}{\sqrt{L}}\}^L$, where $L = \frac{T_b}{T_c}$, and the collection of codes as the matrix,

$$\mathbf{M} = \begin{bmatrix} \mathbf{s}_1 \\ \vdots \\ \mathbf{s}_K \end{bmatrix}. \tag{7}$$

Note that $\mathbf{R} = \mathbf{M}^t\mathbf{M}$ where $\mathbf{M}^t$ denotes the transpose of $\mathbf{M}$ and that $\mathbf{R}$ is symmetric and positive definite. We assume that the signature codes are selected randomly. Note that a system with random codes will closely match the system that uses long pseudo-random sequences to generate its signature codes.

III. Problem Formulation

Our goal is to improve the bit error-rate (BER) of a DS-CDMA system by preprocessing the transmitted signal. We use the minimum mean-squared error as the criterion in the design of the preprocessor. To estimate a (BPSK modulated) data vector $\mathbf{b}$ based on the set of observations $\mathbf{y}$, we use the function $\hat{\mathbf{b}}()$ which minimizes the mean-squared error, i.e.,

$$E[(\mathbf{b} - \hat{\mathbf{b}}(\mathbf{y}))^2]. \tag{8}$$

One option for preprocessing is transmitter precoding which applies a linear transformation $\mathbf{T}$ to the transmitted vector so that,

$$\mathbf{y} = \mathbf{RTAb} + \mathbf{n}. \tag{9}$$

It was shown in [2] that when the MMSE criterion is used to solve for $\mathbf{T}$ and expectation is taken with respect to $\mathbf{b}$ and $\mathbf{n}$ the result is $\mathbf{T} = \mathbf{R}^{-1}$.

Similarly, we can consider a more general (possibly non-linear) transformation on $\mathbf{b}$ which gives the real-valued vector $\mathbf{b}'$ so that,

$$\mathbf{y}' = \mathbf{R}\mathbf{b}' + \mathbf{n}, \tag{10}$$

where we have incorporated $\mathbf{A}$ into $\mathbf{b}'$. Using the MMSE criterion to select $\mathbf{b}'$, we obtain mean-square error ρ where,

$$\rho = E[\|\mathbf{Ab} - \mathbf{y}'\|^2], \tag{11}$$

$$= E[\|\mathbf{Ab} - (\mathbf{R}\mathbf{b}' + \mathbf{n})\|^2], \tag{12}$$

which, after taking expectation with respect to $\mathbf{n}$, gives

$$\rho = \|\mathbf{Ab} - \mathbf{R}\mathbf{b}'\|^2. \tag{13}$$

As a result, the optimization problem can be expressed as,

$$\underset{\mathbf{b}'}{\operatorname{argmin}} \|\mathbf{A}\mathbf{b} - \mathbf{R}\mathbf{b}'\|^2, \tag{14}$$

where the solution $\mathbf{b}' = \mathbf{R}^{-1}\mathbf{A}\mathbf{b}$ is the same as for transmitter precoding.

Similar to the decorrelating detector, using the optimal $\mathbf{T}$ completely eliminates MAI. However, because the decorrelating operation occurs before transmission, this method (unlike the decorrelating detector) does not suffer from noise enhancement.

A side effect of precoding is increased transmitter power. It is natural to consider imposing an energy constraint at the transmitter, say $\|\mathbf{M}\mathbf{b}'\|^2 \leq r$. Replacing $\mathbf{R} = \mathbf{M}^t\mathbf{M}$ in (14), we can reformulate the optimization as:

$$\begin{aligned} \min_{\mathbf{b}'} \quad & \|\mathbf{A}\mathbf{b} - \mathbf{M}^t\mathbf{M}\mathbf{b}'\|^2 \\ \text{subject to:} \quad & \|\mathbf{M}\mathbf{b}'\|^2 \leq r. \end{aligned} \tag{15}$$

We refer to this method as the *optimizing precoder*. From the set of $\mathbf{b}'$ with transmitted power less than or equal to r, it selects the element which results in the minimum squared error due to MAI. The parameter r is held constant from symbol to symbol so instantaneous transmitter power is capped by a fixed value.

It is crucial to realize that the method proposed here acts on the amplitude of the transmitted signal, and consequently, does not modify its spectral property. This can be considered as a form of fast power control where the power of each user is readjusted at each bit interval (taking into account the bit values transmitted to other users). This is unlike the known alternatives, say [2], which are based on applying a linear transformation to the transmitted signal which obviously will affect the spectral properties of the transmitted signal. We note that having a flat spectrum is a desirable property of many CDMA systems and the method proposed here does not interfere with this property.

By constraining transmitter power, we force the system to include MAI in situations where eliminating it would require inappropriate power levels. For this reason, the optimizing precoder can be expected to have worse raw performance than unconstrained precoding due to allocating different level of MAI at different bit positions. The situation changes, however, when FEC is used because FEC coders achieve a form of averaging of instantaneous signal-to-noise ratio over several bit periods, compensating for the aforementioned effect of variable MAI. This conclusion will be borne out by simulation results.

A constrained transmitter precoding scheme is also presented in [2]. However, there are some critical differences between that formulation and (15), the most significant of which (as already discussed) is that our method is based on a scaling of the transmitted signals, and consequently, preserves their spectral properties. Another significant difference is the scope of the precoder design. In (15), a new transformation is generated for each symbol period, with expectation of MSE taken with respect to environmental noise only. In [2], expectation of MSE is taken with respect to $\mathbf{b}$ as well as $\mathbf{n}$ so that the current data vector is not a distinct influence on the design of the precoding transformation. Furthermore, (15) caps instantaneous power at each interval while in [2], *average* power is constrained. The latter method allows for undesirable spikes in instantaneous power.

IV. Solution Algorithm

Numerical optimization techniques are required to solve (15). The problem specifies minimization of a convex quadratic function (because $\mathbf{M}\mathbf{M}^t$ is positive semi-definite) over an ellipsoid (a convex region) whose surface is the set of signal vectors which have a transmitted power level of r. Note that due to the convexity of the problem, any local optimum solution will be also globally optimum. To simplify the problem, we apply the change of variable $\mathbf{x} = \mathbf{M}\mathbf{b}'$ which transforms (15) into minimization over a sphere, i.e.,

$$\begin{aligned} &\min_{\mathbf{x}} \quad \|\mathbf{A}\mathbf{b} - \mathbf{M}^t\mathbf{x}\|^2 \\ &\text{subject to:} \quad \|\mathbf{x}\|^2 \leq r. \end{aligned} \tag{16}$$

Expanding the new objective function results in,

$$\mathbf{b}^t\mathbf{A}^t\mathbf{A}\mathbf{b} - 2\,\mathbf{x}^t\mathbf{M}\mathbf{A}\mathbf{b} + \mathbf{x}^t\mathbf{M}\mathbf{M}^t\mathbf{x}.$$

Let $\mathbf{Q} = 2\mathbf{M}\mathbf{M}^t$, $\mathbf{c} = -2\mathbf{M}\mathbf{A}\mathbf{b}$ and delete the constant term to obtain,

$$\begin{aligned} &\min_{\mathbf{x}} \quad \tfrac{1}{2}\mathbf{x}^t\mathbf{Q}\mathbf{x} + \mathbf{x}^t\mathbf{c} \\ &\text{subject to} \quad \|\mathbf{x}\|^2 \leq r. \end{aligned} \tag{17}$$

This type of problem occurs frequently in nonlinear optimization as the *trust-region subproblem*. Its solution has a standard treatment which is given in [7]. A more convenient solution algorithm is given in [8]. We give an overview of that algorithm here.

At a point $\mathbf{x}$, the gradient to the objective function is given by $\mathbf{Q}\mathbf{x} + \mathbf{c}$ while the gradient to the constraint is given by $-2I\mathbf{x}$ where I is the identity matrix. If the *Karush-Kuhn-Tucker* (KKT) conditions for local[2] optimum point $\mathbf{x}^*$ are satisfied with Lagrange multiplier λ^*, then the following holds:

$$(\mathbf{Q} + \lambda^*\mathbf{I})\mathbf{x}^* + \mathbf{c} = 0 \tag{18a}$$

$$\sqrt{r} - \|\mathbf{x}^*\| \geq 0 \tag{18b}$$

$$\lambda^* \geq 0 \tag{18c}$$

$$\lambda^*(\sqrt{r} - \|\mathbf{x}^*\|) = 0, \tag{18d}$$

where the factor 2 has been absorbed into λ^* for notational convenience.

To establish whether a particular $\mathbf{x}^*$ is a minimum, we check the second order condition which requires positive curvature of the objective function at $\mathbf{x}^*$. Because the matrix $\mathbf{Q}$ is positive semi-definite, we are assured that the second order condition will be satisfied for any point in the constraint region. This observation, coupled with a non-zero constraint gradient at all points other than the origin, allows us to conclude that the KKT conditions are satisfied with a unique λ^* and that the associated point $\mathbf{x}^*$ is a global optimum solution.

For matrix $\mathbf{A}$, we use $\underline{\psi}(\mathbf{A})$ to denote the least eigenvalue of $\mathbf{A}$. The range of possible values for λ^* is already lower bounded by zero in (18c). We can obtain an upper bound for λ^* by rearranging and taking the norm of each side of (18a). After some manipulation, we have,

$$\|(\mathbf{Q} + \lambda^*\mathbf{I})^{-1}\mathbf{c}\| = \|\mathbf{x}\|,$$

which implies,

$$\frac{1}{\underline{\psi}(\mathbf{Q}) + \lambda^*}\|\mathbf{c}\| \geq \|\mathbf{x}\|, \tag{19}$$

and thus,

$$\underline{\psi}(\mathbf{Q}) + \lambda^* \leq \frac{\|\mathbf{c}\|}{\|\mathbf{x}\|}, \tag{20}$$

finally,

$$\lambda^* \leq \frac{\|\mathbf{c}\|}{\|\mathbf{x}\|} - \underline{\psi}(\mathbf{Q}), \tag{21}$$

[2] Note that due to the convexity of the problem, any local optimum solution will be also globally optimum.

as shown in [8]. If a non-zero λ^* satisfies the KKT conditions, then the corresponding $\mathbf{x}^*$ must have $\|\mathbf{x}^*\| = \sqrt{r}$. If we do not have access to the eigenvalues of $\mathbf{Q}$, then we can drop the $\underline{\psi}(\mathbf{Q})$ term from (21) which gives the relaxed upper bound $\lambda^* \leq \frac{\|\mathbf{c}\|}{\sqrt{r}}$.

We now have a range of possible values for λ^*. We start by checking $\lambda^* = 0$. If the norm of the solution is less than $\sqrt{r}$ then we terminate with a solution inside the constraint. Otherwise, we can test an arbitrary λ for optimality by solving (18a) for $\mathbf{x}$ and then checking its norm. If $\|\mathbf{x}\| = \sqrt{r}$, then $\mathbf{x}$ is the optimal solution to the problem. This suggests a bisection algorithm for solving (17). Assume that we wish to find λ^* to within ϵ of the exact solution. The bisection algorithm, introduced in [8], proceeds as follows:

1. set $\lambda_{\text{low}} = 0$ and $\lambda_{\text{high}} = \frac{\|\mathbf{c}\|}{\sqrt{r}}$.
2. set $\lambda = \frac{1}{2}(\lambda_{\text{low}} + \lambda_{\text{high}})$.
3. if $(\lambda_{\text{high}} - \lambda_{\text{low}}) \leq \epsilon$ exit procedure, otherwise go to step 4.
4. solve $(\mathbf{Q} + \lambda\mathbf{I})\mathbf{x} = -\mathbf{c}$ for $\mathbf{x}$
5. if $\|\mathbf{x}\| > \sqrt{r}$, then set $\lambda_{\text{low}} = \lambda$, otherwise set $\lambda_{\text{high}} = \lambda$.
6. goto step 2

A change of variables was applied to obtain the formulation in (17). The solution to the original problem can be obtained by back-solving through the symmetric positive-definite system $\mathbf{M}\mathbf{b}' = \mathbf{x}$ once $\mathbf{x}$ is found.

The complexity of this algorithm is determined by the desired accuracy and the cost of the linear system solution in step 4. At least $\log_2 \frac{\|\mathbf{c}\|}{\sqrt{r}} - \log_2 \epsilon$ iterations are required to converge to precision ϵ. If the eigenvalues of $\mathbf{Q}$ are known and non-zero, the number of iterations can be reduced. The structure of the linear system in step 4 is helpful in reducing the complexity. Because $\lambda > 0$ during each iteration, $(\mathbf{Q} + \lambda\mathbf{I})$ must be symmetric positive definite which implies fast Cholesky factorizations can be used to solve the system.

Simulation will show that two iterations are enough to obtain solutions with near-optimal performance. Each of these iterations requires the factorization of an $L \times L$ matrix. Thus, complexity depends on the system chip rate, but not on the number of active users.

Alternatively, another solution algorithm has been proposed in [9],[10] tailored for large scale systems for which forming factorizations is infeasible. The basis of this technique is the Lanczos method which can approximate eigenvalues using only matrix-vector products. On specialized processors, this technique may provide a performance gain over the method described above.

V. Simulation Results

In this section, we analyze the performance of the optimizing precoder in a simulated DS-CDMA environment. Simulation results will be compared with those from conventional systems as well as with both the constrained and unconstrained transmitter precoding methods of [2]. It will be shown that, in environments utilizing error-correcting codes, energy-constrained methods out-perform all other methods considered here (this is unlike the conclusion reached in [2] for the uncoded case). Furthermore, it will be shown that the energy-constrained method proposed in this article out-performs the energy-constrained method of [2].

A. Simulation Setup

The channel model used for simulation matches the forward link of a single-cell synchronous DS-CDMA system using BPSK modulation as specified in section II. The system chip rate is set to $L = 32$ in all cases. Both coded and uncoded systems are considered with coded systems using the rate $1/3$ convolutional code from the IS-95 cellular communication standard. Receivers use a standard Viterbi decoder which is supplied with soft outputs from the matched-filter detector.

The simulator generates independent pseudo-random data streams for each user which are independently coded when FEC is used. Another independent pseudo-random bit stream, shared by all users, is used to generate the spreading codes. After modulation and spreading, AWGN is added to the combined signals. A bank of independent detectors then produces soft-outputs from the received signals which are either fed to a Viterbi decoder or used to make hard decisions for the original data bits.

B. Optimizer Setup

With the structure of the simulation system in place, the optimal value of the power constraint r for the proposed method can be determined. With no closed form solution for the optimal r in (15), we must use simulation to find it. By holding the system parameters and SNR constant while varying r, we hope to find a value which gives the minimum BER for each system configuration. Fortunately, BER curve is smooth and continuous with respect to changes in r which aids in the search. Figure 1 shows the effect of r on the optimizing precoder performance when the system configuration and SNR level are held constant. In all four curves, performance reaches a single global optimum with respect to r. This is a result of the optimization step which results in a different set of optimizer solutions for each r. One such set will have the minimum total MAI, implying that r must have a minimum value.

Only 16 user results are shown. However, the optimum value of r was consistently lower for coded transmission than for uncoded for all numbers of users. Moreover, the optimal r for FEC was much less affected by the number of users and changing channel conditions. Likely, this is due to fact that by averaging over several symbol periods, the effect of high instantaneous MAI is reduced and the benefit of tightly constrained transmitted power levels is enhanced.

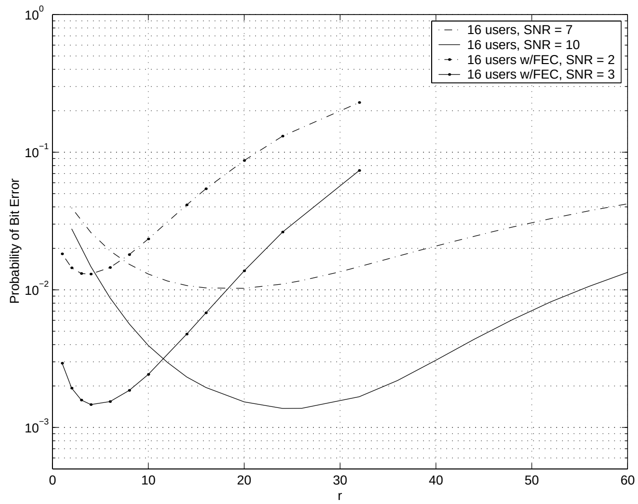

Fig. 1. BER vs r: 16 users

Despite the effects of FEC, the optimum r was not constant for different system configurations or even for different SNR levels. The system is required to adjust r constantly for changing conditions to achieve optimum performance. However, the set of optimum values for r need only be calculated once and stored in a look-up table. Alternatively, r can be chosen to optimize performance at a certain SNR with only a small loss at other SNR levels due to the smoothness of the BER curve. Note that a conventional detector does not require knowledge of r to produce soft output metrics for this type of signal so that the modulator is free to adjust r as necessary.

C. Channel Simulation

Each simulated “system” consists of a modulator/detector pair and produces a single curve on each graph. Two conventional systems are considered (all systems simulated in a similar environment): a system with a standard modulator and a standard matched-filter detector, and a system with a decorrelating detector. The former is labeled “Conventional System” and provides a baseline performance curve with no performance enhancing techniques. The latter is labeled the “Decorrelating Detec-

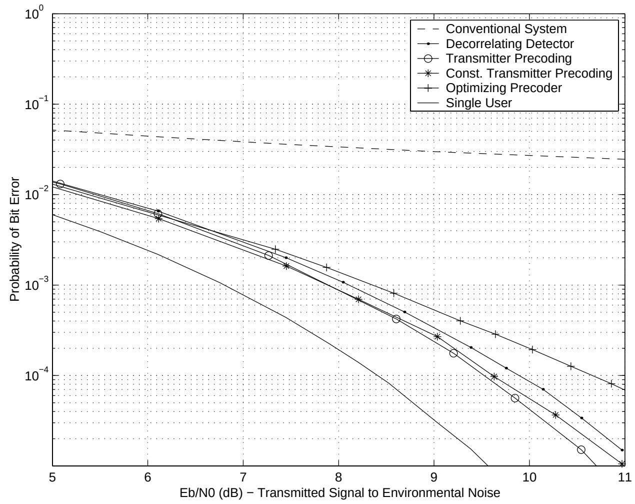

Fig. 2. BER: 8 Users in $L = 32$ Chips, Not Coded

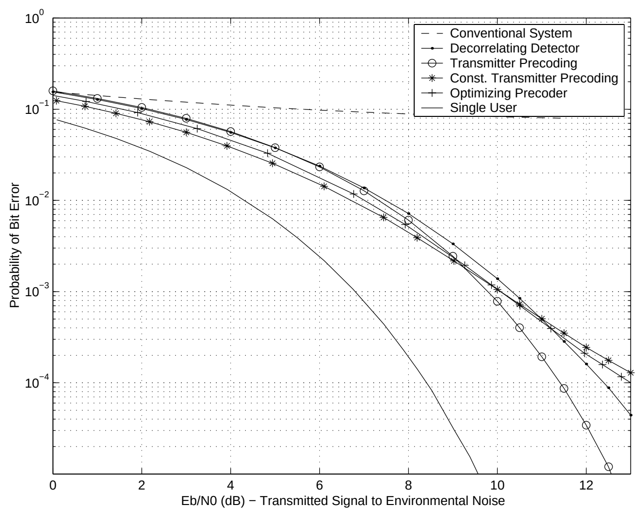

Fig. 3. BER: 16 Users in $L = 32$ Chips, Not Coded

tor" and gives an example of a standard receiver-based performance enhancement technique. The two methods from [2] are labeled "Transmitter Precoding" and "Constrained Transmitter Precoding". These are appropriate generalizations of the method proposed in [2] to a case comparable to ours (with FEC and power constraint added at the transmitter side [3]). The optimizing precoder proposed in this paper is labeled "Optimizing Precoder" and is paired with a standard detector. The single user case is also given for comparison.

In Figures 2 and 3, uncoded data is being transmitted to $8, 16$ users. In all cases, transmitter

[3] Reference [2] does study the case of imposing a constraint on the power at the transmitter side, however, as it does not use FEC, the conclusion has been that imposing such a constraint does not improve the performance. This agrees with our results for the uncoded case, however, we have shown that imposing a constraint on power indeed improves the performance for the coded case.

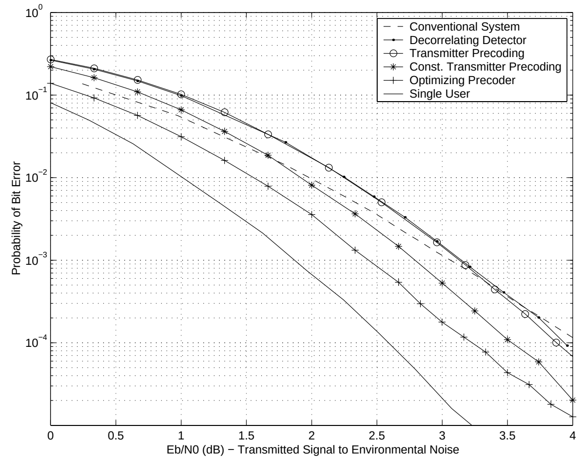

Fig. 4. BER: 8 Users in $L = 32$ Chips, Coded

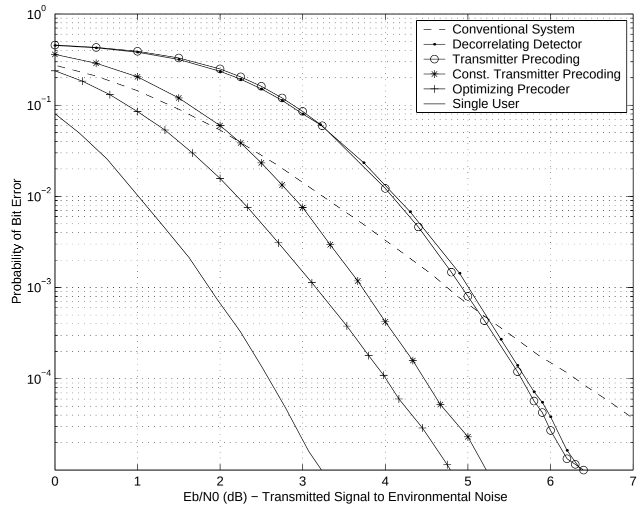

Fig. 5. BER: 16 Users in $L = 32$ Chips, Coded

precoding gives the best performance at practical system error rates of 10^{-3} and below. As stated in [2], constrained transmitter precoding performs better for low SNR, but eventually crosses the unconstrained method as SNR rises. The optimizing precoder from this paper generally performs worse than constrained transmitter precoding in this setting. However, this situation changes in the presence of FEC.

In Figures 4 and 5, data is independently coded and then transmitted to $8, 16$ users. Clearly, the energy-constrained methods significantly out-perform unconstrained methods in this setting. For 16 users, the constrained transmitter precoding has a gain of more than 1.0dB over unconstrained transmitter precoding. The optimizing precoder performs even better, achieving a further gain of 0.5dB over constrained precoding. Simulations for the 8 user case show this gain to be 0.25dB.

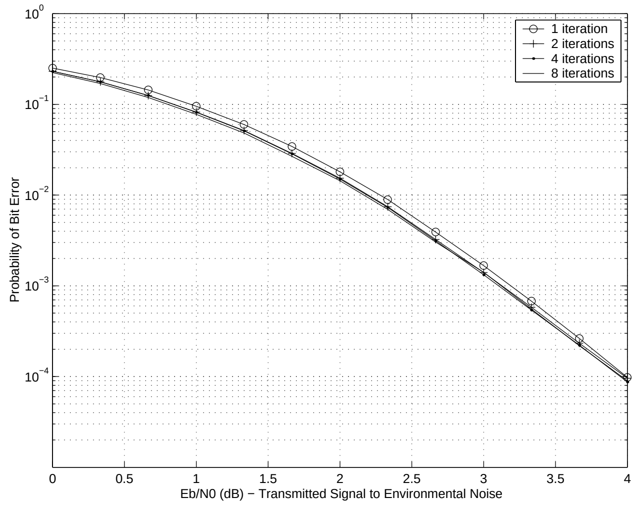

Fig. 6. Complexity: 16 Users in $L = 32$ Chips, Coded

D. Complexity Analysis

The computational complexity of the standard solution algorithm for both optimizing precoders is determined by the desired accuracy of the solution. Accuracy translates directly into iterations of a loop which requires the Cholesky factorization of a matrix. Figure 6 show the result of limiting the number of iterations for the 16 user case. From this graph it is clear that that near-optimal performance is achieved as soon as the second iteration, which requires only two matrix factorizations. This result dramatically improves the feasibility of these methods for practical systems.

In order notation, the Cholesky factorization is $O(L^3)$ while two $O(L^2)$ triangular system back-solves are required to obtain the final solution. The parameter L (chip rate) is usually fixed in practical systems, so that these costs are also fixed. These costs are not prohibitive in the forward link of a mobile cellular system because the processing is performed at the base station where power and processing resources are readily available. The chief benefit of this, and all other precoding methods, is the simplicity of the receiver, which is only a matched-filter detector.

VI. Conclusions

A new transmitter precoding method has been demonstrated in this article which, for practical systems which delivers gains of: (i) 0.75dB to 4dB over existing unconstrained precoding methods, (ii) 0.25dB to 0.5dB over existing constrained precoding methods, and (iii) requires only 2 to 4 fast Cholesky factorizations of an order L matrix to achieve those gains

REFERENCES

[1] E. S. Hons, “Constrained quadratic optimization: Theory and application for wireless communication systems,” M.S. thesis, University of Waterloo, Waterloo, Canada, 2001.

[2] B. R. Vojčić and W. M. Jang, “Transmitter precoding in synchronous multiuser communications,” *IEEE Transactions on Communications*, vol. 46, no. 10, pp. 1346–1355, Oct. 1998.

[3] W. M. Jang, B. R. Vojčić, and R. L. Pickholtz, “Joint transmitter-recevier optimization in synchronous multiuser communication over multipath channels,” *IEEE Transactions on Communications*, vol. 46, no. 2, pp. 269–278, Feb. 1998.

[4] M. Brandt-Pearce and A. Dharap, “Transmitter-based multiuser interference rejection for the down-link of a wireless CDMA system in a multipath environment,” *IEEE Journal on Selected Areas of Communications*, vol. 18, no. 3, pp. 407–417, Mar. 2000.

[5] Z. Tang and S. Cheng, “Interference cancellation for DS-CDMA systems over flat fading channels though pre-decorrelating,” in *PIMRC ’94*, The Hague, The Netherlands, 1994.

[6] S. Verdu, *Multiuser Detection*, Cambridge University Press, 1998.

[7] J. Nocedal and S. J. Wright, *Numerical Optimization*, Springer Series in Oerations Research. Springer–Verlag, New York, 1999.

[8] Y. Ye, “On affine scaling algorithms for nonconvex quadratic programming,” *Mathematical Programming*, vol. 56, pp. 285–300, 1992.

[9] D. C. Sorensen, “Minimization of a large-scale quadratic function subject to a spherical constraint,” *SIAM Journal on Optimization*, vol. 7, no. 1, pp. 141–161, 1997.

[10] M. Rojas, S. A. Santos, and D. C. Sorensen, “A new matrix-free algorithm for the large-scale traust-region subproblem,” *SIAM Journal on Optimization*, vol. 11, no. 3, pp. 611–646, 2000.