A Vector-Perturbation Technique for Near-Capacity Multi-Antenna Multi-User Communication—Part II: Perturbation

Bertrand M. Hochwald, Christian B. Peel, and A. Lee Swindlehurst

Abstract

Recent theoretical results describing the sum-capacity when using multiple antennas to communicate with multiple users in a known rich scattering environment have not yet been followed with practical transmission schemes that achieve this capacity. We introduce a simple encoding algorithm that achieves near-capacity at sum-rates of tens of bits/channel-use. The algorithm is a variation on channel inversion that regularizes the inverse and uses a "sphere encoder" to perturb the data to reduce the energy of the transmitted signal.

The paper comprises two parts. In this second part we show that, after the regularization of the channel inverse introduced in the first part, a certain perturbation of the data using a "sphere encoder" can be chosen to further reduce the energy of the transmitted signal. The performance difference with and without this perturbation is shown to be dramatic. With the perturbation, we achieve excellent performance at all SNR's. The results of both uncoded and turbo-coded simulations are presented.

Index Terms

B. M. Hochwald (*hochwald@lucent.com*) is with the Mathematics of Communications Research Department, Bell Laboratories, Murray Hill, NJ. C. B. Peel (*chris.peel@ieee.org*) is with the Swiss Federal Institute of Technology, Zurich, Switzerland. A. L. Swindlehurst (*swindle@ee.byu.edu*) is with Brigham Young University, Provo, UT. This work was supported in part by the National Science Foundation under Wireless Initiative Grant CCR-9979452 and Information Technology Grant CCR-0081476.

Broadcast channel, multiple-antenna multiple-user wireless, MIMO, Tomlinson-Harashima precoding, dirty-paper coding, spatial equalization, sphere encoder.

I. Introduction

When an access point with multiple antennas is used to communicate with many users, each with one antenna, the communication problem is complicated by the fact that the each user must decode his/her signal independently from the remaining users. With K antennas at the access point and K users, simple inversion of the square matrix channel at the transmitter allows independent signals to be directed to the various users. However, in the first part of this paper [1], we show that plain channel inversion performs poorly at all SNR's and for any number of users. Regularization of the channel, where a scaled identity matrix is added before the inverse is taken, improves performance substantially but still leaves a gap to capacity.

In this second part of the paper, instead of further modifying the inverse, we modify the data that is transmitted by judiciously adding an integer vector offset. An important data-modifying technique originally developed for the intersymbol interference channel is Tomlinson-Harashima precoding [2], [3]. This technique applies a scalar integer offset at the transmitter that allows cancellation of the interference after application of a modulo function at the receiver. Other researchers have recently examined the information-theoretic aspects [4] of Tomlinson-Harashima precoding, applied it to space-time coding [5] and used it to overcome inter-chip interference in CDMA [6], [7].

We show that a technique related to both Tomlinson-Harashima precoding and channel inversion can achieve near-sum-capacity even at high SNR, with each user receiving $1/K$th of the sum-capacity. Our technique does not require explicit dirty paper techniques. In fact, while the technique requires the transmitter to know the channel, each receiver needs to know only a single prearranged scalar related to the SNR of the channel.

Our method requires the joint selection of a *vector* perturbation of the signal to be transmitted to all the receivers. We show that sequential application of standard scalar Tomlinson-Harashima

precoding to choose this offset does not perform nearly as well as our joint selection. The perturbation algorithm has the simple interpretation of placing the largest signal components along the smallest singular values of the inverse channel, and the smallest signal components along the largest singular values of the inverse channel. In general, techniques such as the Fincke-Pohst algorithm [8], [9] (which in our context we label "sphere encoding"), can aid in selecting the desired vector perturbation. In all cases, however, the processing at the receiver is simple.

II. Model Summary

We briefly review the model introduced in [1]. A general model for the forward link of a multi-user system includes an access point with M transmit antennas and K users, each with one receive antenna. The received data at the K users can be written using a vector equation

$$\mathbf{y} = H\mathbf{x} + \mathbf{w} \,, \tag{1}$$

where $\mathbf{y} = [y_1, \ldots, y_K]^T$ contains the data received at each user. The transmitted signal is $\mathbf{x}$ which contains an element for each of the M transmit antennas; the receiver noise vector $\mathbf{w}$ has an element for each user, and the $K \times M$ matrix H has $h_{k,m}$ as elements, indicating the channel gain between transmit antenna m and user k. The power constraint $\mathrm{E}\,||\mathbf{x}||^2 = 1$ is imposed, with $\mathrm{E}\,\mathbf{w}\mathbf{w}^* = K \cdot I_K$. We construct an unnormalized signal $\mathbf{s}$ such that

$$\mathbf{x} = \frac{\mathbf{s}}{\sqrt{\gamma}}, \qquad \text{or} \qquad \mathbf{x} = \frac{\mathbf{s}}{\sqrt{\mathrm{E}\,\gamma}}. \tag{2}$$

where $\gamma = ||\mathbf{s}||^2$. The ergodic sum-capacity for the model (1) is

$$C_{\mathrm{sum}} = \mathrm{E} \sup_{D \in \mathcal{A}} \log |I_M + \rho H^* D H| \,, \tag{3}$$

where I_M is the $M \times M$ identity matrix, $\mathcal{A}$ is the set of $K \times K$ diagonal matrices D with non-negative elements that sum to one ($\mathrm{tr}\, D = 1$), and we define $\rho = 1/\sigma^2$. The Hermitian transpose of H is indicated by H^*. In this paper we consider only the case $K = M$; other cases are considered in Part I of this paper [1] and elsewhere [10].

III. Perturbing the Data

In [1] we argue that many of the problems with inverting the channel when $K = M$ are due to the normalization constant γ, which is often very large because of the large singular values in the inverse of the channel matrix H. One way to help H is to regularize its inverse, as described in [1]. Another way is to make sure the transmitted data does not lie along the singular vectors associated with the large singular values of H^{-1}. In this section, we present a way to "perturb" the data in a data-dependent way (unknown to the receivers) so that the data vector is approximately orthogonal to the right singular vectors associated with the large singular values. Our goal therefore is to form a $\tilde{\mathbf{u}}$ from the data vector $\mathbf{u}$ such that

$$\mathbf{s} = H^{-1}\tilde{\mathbf{u}} \tag{4}$$

has norm (much) smaller than $H^{-1}\mathbf{u}$, but the entries of $\tilde{\mathbf{u}}$ can still be decoded individually at the receivers.

We cannot, in general, perturb $\mathbf{u}$ by an arbitrary complex vector because this perturbation is not known to the receivers and would therefore cause decoding errors. We can, however, use an idea derived from Tomlinson-Harashima precoding [2], [3] where we allow each element of $\mathbf{u}$ to be perturbed by an integer. In the simplest case, we set

$$\tilde{\mathbf{u}} = \mathbf{u} + \tau \boldsymbol{l} \tag{5}$$

where τ is a positive real number and $\boldsymbol{l}$ is a K-dimensional complex vector $a + ib$, where a and b are integers. The scalar $\gamma = \|\mathbf{s}\|^2$ is computed as before and the transmitted signal is

$$\mathbf{x} = \frac{1}{\sqrt{\gamma}} H^{-1}\tilde{\mathbf{u}}. \tag{6}$$

The scalar τ is chosen large enough so that the receivers may apply the modulo function

$$f_\tau(y) = y - \left\lfloor \frac{y + \tau/2}{\tau} \right\rfloor \tau, \tag{7}$$

where the function $\lfloor \cdot \rfloor$ is the largest integer less than or equal to its argument. The function (7) removes the effect of the integer multiple of τ. (The function $f_\tau(y)$ is applied separately to the

real and imaginary components of a complex y.) We have more to say about the choice of τ shortly. After passing through the channel H, the transmitted signal $\mathbf{x}$ in (6) appears at receiver k as

$$y_k = \frac{1}{\sqrt{\gamma}}\tilde{u}_k + w_k.$$

If we ignore for the moment the effect of w_k, and assume that $\gamma = 1$, then

$$f_\tau(y_k) = f_\tau(u_k + \tau l_k) = u_k,$$

and we recover the transmitted symbol. The receivers know γ and therefore may compensate for $\gamma \neq 1$ by dividing τ by $\sqrt{\gamma}$. As we note in Section II, the transmitter may instead divide by $\sqrt{\mathrm{E}\,\gamma}$; Figure 6 shows that the performance difference is not significant. Our other figures assume that the transmitter divides by $\sqrt{\gamma}$. An error is made at the receiver if the additive channel noise pushes the received signal across the standard symbol decoding boundaries or across the nonlinear boundaries of $f_\tau(y)$ at $\pm\tau/2$.

A. Choice of $\boldsymbol{l}$ *and* τ

An obvious choice of $\boldsymbol{l}$ at the transmitter minimizes $\gamma = \|\mathbf{s}\|^2$,

$$\boldsymbol{l} = \arg\min_{\boldsymbol{l}'}(\mathbf{u} + \tau\boldsymbol{l}')^*(HH^*)^{-1}(\mathbf{u} + \tau\boldsymbol{l}') \tag{8}$$

This is a K-dimensional integer-lattice least-squares problem for which there is a large selection of exact and approximate algorithms. See, for example, algorithms by Fincke and Pohst [8] or Kannan [11]. A review of algorithms can be found in [9], and recent developments may be found in [12]–[14]. The Fincke-Pohst algorithm is used for space-time demodulation in [15] where it is called a *sphere decoder*. Because we are using this algorithm for encoding data to be transmitted, we refer to it as the sphere *encoder*. We leave the details of this algorithm to the references, but mention that the algorithm avoids an exhaustive search over all possible integers in the lattice by limiting the search space to a sphere of some given radius centered around a

starting point. In our case the center is the vector $\mathbf{u}$. Generally, the sphere encoder works on real lattices, so we assume that a complex version is used [16], or that the equation (8) has been converted to a $2K$-dimensional real lattice problem.

The scalar $\tau > 0$ is a design parameter that may be chosen to provide a symmetric decoding region around (the real or imaginary part of) every signal constellation point, whence

$$\tau = 2(|c|_{\max} + \Delta/2) \tag{9}$$

where $|c|_{\max}$ is the absolute value of the constellation symbol(s) with largest magnitude and Δ is the spacing between constellation points. If we want to reduce the effects of the perturbation vector $\boldsymbol{l}$, we may increase τ, thereby increasing the decoding region at the upper and lower extremes of the constellation. While this improves error performance in these decoding regions, the γ that results is typically also larger, possibly reducing total error performance. If τ is made too large, the minimization in (8) yields $\boldsymbol{l} = 0$ independently of $\mathbf{u}$, and the perturbation technique reduces to simple channel inversion. If τ is made smaller than $2|c|_{\max}$ then error-free decoding becomes impossible even in the absence of channel noise. We find that choosing τ as in (9) often works well.

B. Analysis of vector perturbation technique

As shown in [1], plain channel inversion performs poorly because $(HH^*)^{-1}$ has a badly behaved large eigenvalue. In this section we provide a brief theoretical discussion of why using the perturbation vector $\tau\boldsymbol{l}$ improves the performance significantly, especially for large K. We confine our discussion to large K, where the analysis is most tractable. Since the analysis of integer algorithms is generally difficult, we are forced to relax the constraints of rigor in this section. The reader who would prefer more details about the implementation of the algorithm may skip this discussion and proceed directly to Section III-C.

The vector $\boldsymbol{l}$ is chosen to minimize the norm of $\mathbf{s} = H^{-1}\tilde{\mathbf{u}}$ in (4) (using the cost function

(8)). Using the eigendecomposition $(HH^*)^{-1} = Q\Lambda^{-1}Q^*$, we can express the cost function as

$$\gamma = \|\mathbf{s}\|^2 = \sum_{k=1}^{K} \mu_k \delta_k^2, \tag{10}$$

where $\mu_k = 1/\lambda_k$, $\delta_k = |\mathbf{q}_k^* \tilde{\mathbf{u}}|$, and $\mathbf{q}_k$ is the kth column of Q. We assume that $\mu_1 > \mu_2 > \ldots > \mu_K$. The vector perturbation algorithm minimizes (10) over $\tilde{\mathbf{u}}$, where $\tilde{\mathbf{u}} = \mathbf{u} + \tau \boldsymbol{l}$ and we search over the integer vector $\boldsymbol{l}$. We would like to examine the behavior of $\mathrm{E}\,\gamma$ as a function of K. In [1] it is shown that $\mathrm{E}\,\gamma = \infty$ for plain (without the perturbation) channel inversion. We argue that, with the perturbation, $\mathrm{E}\,\gamma$ is approximately constant with K, and therefore the sum-rate for our method grows linearly with K.

Recall that τ is chosen large enough so that no element of $\tilde{\mathbf{u}}$ can be made zero. In fact, with our choice of τ, the norm of $\tilde{\mathbf{u}}$ is minimized by choosing $\boldsymbol{l} = 0$. Thus, although a nonzero $\boldsymbol{l}$ increases the norm of $\tilde{\mathbf{u}}$, the norm of $\mathbf{s} = H^{-1}\tilde{\mathbf{u}}$ is decreased in the process. There is no norm constraint on $\boldsymbol{l}$ and hence the choice of possible points $\tilde{\mathbf{u}}$ form an infinite lattice.

Define

$$\nu = \mathrm{E}\left(\prod_{k=1}^{K} \frac{\delta_k^2}{c^2}\right)^{1/K}, \tag{11}$$

where

$$c^2 = (1/K)\mathrm{E}\,\|\tilde{\mathbf{u}}\|^2, \tag{12}$$

and the expectation is over Q and $\mathbf{u}$. We take as an empirical axiom that ν is positive and approximately independent of K as $K \to \infty$. Equation (11) is the expected geometric mean of $\delta_1, \ldots, \delta_K$. The fact that $\nu > 0$ is a consequence of the constraints on $\tilde{\mathbf{u}}$, for if $\tilde{\mathbf{u}}$ is unconstrained, then $\nu = 0$ (the minimizer of (10) is parallel to $\mathbf{q}_K$ and therefore obeys $\delta_1 = \ldots = \delta_{K-1} = 0$, $\delta_K = \|\tilde{\mathbf{u}}\|$). We contend that forcing $\boldsymbol{l}$ to have integer components when minimizing (10) does not generally permit $\tilde{\mathbf{u}}$ be chosen exactly parallel to $\mathbf{q}_K$ (an axis in a random coordinate system), and thus the $\boldsymbol{l}$ that minimizes (8) generates a $\tilde{\mathbf{u}}$ that can only be coarsely oriented in the coordinate system defined by $\mathbf{q}_1, \ldots, \mathbf{q}_K$. The orientations with respect to $\mathbf{q}_1, \ldots, \mathbf{q}_K$

do not change significantly with K, and hence the expected geometric mean of $\delta_1, \ldots, \delta_K$ is approximately independent of K. A similar statement can be made for the expected arithmetic mean of $\delta_1, \ldots, \delta_K$. Since the columns of Q form an orthonormal basis,

$$\sum_{k=1}^{K} \delta_k^2 = \|\tilde{\mathbf{u}}\|^2. \tag{13}$$

Combining (13) and (12), we see that c^2 can also be rewritten

$$c^2 = \frac{1}{K} \mathrm{E} \sum_{k=1}^{K} \delta_k^2 \tag{14}$$

and is approximately independent of K.

We apply the arithmetic-geometric-mean inequality

$$\mathrm{E}\,\gamma = \mathrm{E} \sum_{k=1}^{K} \mu_k \delta_k^2 \geq K \mathrm{E} \left[\left(\prod_{k=1}^{K} \mu_k \right)^{1/K} \left(\prod_{k=1}^{K} \delta_k^2 \right)^{1/K} \right]. \tag{15}$$

It can be shown that $K(\prod_{k=1}^{K} \mu_k)^{1/K} \to e$ as $K \to \infty$ [17] (observe that no expectation is needed here). Equation (11) implies that $\mathrm{E}\,(\prod_{k=1}^{K} \delta_k^2)^{1/K} = \nu c^2$. Therefore, (15) becomes

$$\mathrm{E}\,\gamma \geq e \nu c^2. \tag{16}$$

By combining (16) with (14), we also conclude that

$$\mathrm{E}\,\gamma \geq e \nu \frac{\mathrm{E}\,\|\tilde{\mathbf{u}}\|^2}{K}. \tag{17}$$

Equality in (16) and (17) is achieved if

$$\mathrm{E}\,\mu_1 \delta_1^2 = \ldots = \mathrm{E}\,\mu_K \delta_K^2 = \frac{e \nu c^2}{K}. \tag{18}$$

Thus, a way to minimize $\mathrm{E}\,\gamma$ is to have the optimum $\tilde{\mathbf{u}}$ orient itself towards each eigenvalue in inverse proportion to the eigenvalue (on average). The values of c and ν are determined by simulation. Observe that the lower bound (16), if achievable, suggests that γ is approximately independent of K as $K \to \infty$.

It turns out that the vector perturbation algorithm minimizing (8) nearly achieves the lower bound (17). Figure 1 shows a numerical example of minimizing (8) for $K = 10$. Plotted is

the average δ_k^2 and $\mu_k \delta_k^2$ for $k = 1, \ldots, 10$, where $\mathbf{u}$ was chosen randomly from a 16-QAM constellation and $\tau = 2.5298$. We computed $\mathrm{E}\,\gamma = 2.34$ and $\mathrm{E}\,\|\tilde{\mathbf{u}}\|^2 = 23.7$ and $\nu = 0.24$, thus nearly achieving the lower bound $\mathrm{E}\,\gamma \geq 1.55$ in (17). Finally, we see that $\mathrm{E}\,\mu_1 \delta_1^2 \approx \ldots \approx \mathrm{E}\,\mu_{10} \delta_{10}^2 \approx \mathrm{E}\,\gamma/10$, thus obeying (18) closely.

We conclude that optimizing (8) tends to generate a $\tilde{\mathbf{u}}$ that, on average, is oriented towards each eigenvalue of $(HH^*)^{-1}$ in inverse proportion to the eigenvalue as in (18). The value of γ that results nearly achieves the lower bound (17), and is approximately independent of K.

C. A successive algorithm for generating an integer offset vector

For comparison with the vector-perturbation technique, we briefly present a method for generating an integer offset vector by repeated application of scalar Tomlinson-Harashima precoding. The method uses an ordered decomposition of the channel matrix H such that the last user sees no interference, but the kth user sees interference from users $k+1, \ldots, K$. The transmitter compensates for this interference by using its knowledge of $s_{k+1}, \ldots, s_K$ to generate s_k from u_k, starting with $k = K - 1$. Methods based on the QR decomposition have been explored for use with dirty-paper codes in [18]. QR-based algorithms have been used for crosstalk cancellation in digital subscriber lines [19] and for CDMA transmission to distributed receivers [6]. The achievable capacity of a greedy-ordering form of the QR decomposition is analyzed in [20], where it is shown to be close to the sum-capacity when used with perfect dirty-paper interference cancellation codes. Since we do not have such codes available to us, we consider a VBLAST-based ordering for the QR decomposition.

The VBLAST algorithm, originally designed to simplify receiver processing in a point-to-point link [21], provides an ordering of the users' signals that is determined by their channel quality. In [22] the algorithm is adapted to transmitter precoding, and we briefly summarize it here. Let $H = P^* B W^{-1}$, where P is a permuation matrix, B is an upper triangular matrix with ones on the diagonal, and W has orthogonal columns. Let $\eta = \rho/(1+\rho)$ be Costa's "α"

parameter [23], renamed here to avoid confusion with the α used for regularization in [1]. We first generate the signal $\tilde{\mathbf{s}}$, and then form the transmitted signal $\mathbf{s} = W\tilde{\mathbf{s}}$. A successive technique is used to generate $\tilde{\mathbf{s}}$:

$$\begin{aligned} \tilde{s}_K &= u_K \\ \tilde{s}_{K-1} &= f_\tau \left(u_{K-1} - \eta b_{(K-1),K} \tilde{s}_K \right) \\ &\vdots \\ \tilde{s}_1 &= f_\tau \left(u_1 - \eta \sum_{\ell=2}^{K} b_{1,\ell} \tilde{s}_\ell \right) , \end{aligned} \tag{19}$$

where $b_{k,\ell}$ is the (k,ℓ)th entry of the matrix B. We may write this equation in terms of the vector of integers $\mathbf{m}$ that the modulo function effectively adds to the signal:

$$\tilde{\mathbf{s}} = \left((1-\eta)I + \eta B \right)^{-1} \left(\mathbf{u} + \tau \mathbf{m} \right) . \tag{20}$$

The signal $\mathbf{s} = W\tilde{\mathbf{s}}$ is formed, normalized, and then sent through the channel. The K users receive

$$\mathbf{y} = \frac{1}{\sqrt{\gamma}} H\mathbf{s} + \mathbf{w} = \frac{1}{\sqrt{\gamma}} P^* B \left((1-\eta)I + \eta B \right)^{-1} \left(\mathbf{u} + \tau \mathbf{m} \right) + \mathbf{w}.$$

The parameter η has a similar effect as α in [1] since it boosts the SINR at each user. The matrix P^* permutes the order of the users according to the VBLAST criterion. Decoding occurs at user k based on [18], [24]

$$y'_k = f_{\tilde{\tau}_k}(\eta y_k) ,$$

where $\tilde{\tau}_k = \frac{\tau}{\sqrt{\gamma}}$. Each receiver models the received data as

$$y'_k = \frac{\eta}{\sqrt{\gamma}} u_k + w'_k ,$$

where w'_k combines the additive receiver noise w_k and the interference. We do not analyze the algorithm but simply mention that at high ρ (where $\eta \to 1$),

$$\mathbf{s} = WB^{-1} \left(\mathbf{u} + \tau \mathbf{m} \right) . \tag{21}$$

The algorithm described here differs slightly from [22] in our use of the parameter η. However, the use of η is well-known in the dirty-paper coding literature [18], [24], where its use improves performance at low power levels.

Figure 2 provides an uncoded symbol probability of error plot for plain channel inversion, regularized inversion, vector perturbation, the successive algorithm, as well as a regularized version of the sphere encoder that is presented in Section IV. The probability of error is shown for 16-QAM signaling with $K = 10$ transmit antennas and users, as a function of ρ. Although the vector perturbation technique does not do as well as the successive algorithm or regularized inversion for low ρ, it achieves a significant gain in performance for high ρ. The regularized perturbation technique described next in Section IV performs well for all ρ.

Figure 2 shows that the beneficial effect of regularization is generally a gain in ρ, with little effect on the high-ρ slope (or "diversity") of the error curve. The linear inversion-based methods have the lowest diversity. The VBLAST-based successive method seems to have a steep slope at first, but ultimately has high-ρ diversity that is similar to the linear inversion-based methods. Only the vector-perturbation method retains a high diversity at high ρ.

IV. Regularized Perturbation

We can marry the regularized inversion method of [1] with the vector perturbation technique of Section III to reduce γ more than either method could alone over a wide range of ρ. The choice of the integer vector $\boldsymbol{l}$ that minimizes γ is made with the modified cost function

$$\boldsymbol{l} = \arg\min_{\boldsymbol{l}'} \left|\left| H^*(HH^* + \alpha I_K)^{-1}(\mathbf{u} + \tau \boldsymbol{l}') \right|\right|^2 . \tag{22}$$

Unfortunately, the analysis of the combined method appears to be difficult. In [1], α is chosen to maximize an approximation to the SINR. We do not know how to compute the average SINR after the minimization (22), and α_{opt} is generally no longer K/ρ when regularization is combined with perturbation. Because γ is significantly smaller in (22) than with regularization alone, $\alpha = K/\rho$ is too large and gives too much crosstalk from the other users. The optimum α is generally

significantly smaller. For example, probability-of-error simulations show that $\alpha_{\rm opt} \approx 1/(5\rho)$ for $K = 4$, and $\alpha_{\rm opt} \approx 1/\rho$ for $K = 10$. We do not have a good explanation for these choices of α and leave this as an open problem.

Despite the difficulty of analyzing $\alpha > 0$, we note that much of the qualitative behavior of the minimizing solution described in Section III-B for $\alpha = 0$ still applies. For example, we find that $\mathrm{E}\,\mu_1\delta_1^2 \approx \ldots \approx \mathrm{E}\,\mu_K\delta_K^2$, as in Figure 1, except that $\mu_1, \ldots, \mu_K$ tend to be smaller with $\alpha > 0$.

A. Simulation of a Complete System

To check our distance from capacity, we simulated a complete system for $M = K = 4$ antennas/users and $M = K = 10$ antennas/users. The transmitted signal is

$$\mathbf{x} = \frac{1}{\sqrt{\gamma}} H^*(HH^* + \alpha I)^{-1}(\mathbf{u} + \tau \boldsymbol{l}), \tag{23}$$

The receivers know τ, but not $\boldsymbol{l}$. The K users receive (as a vector)

$$\mathbf{y} = \frac{1}{\sqrt{\gamma}} HH^*(HH^* + \alpha I)^{-1}(\mathbf{u} + \tau \boldsymbol{l}) + \mathbf{w}. \tag{24}$$

User k models its received signal as

$$y_k = \frac{1}{\sqrt{\gamma}}(u_k + \tau l_k) + w_k', \tag{25}$$

where w_k' contains not only the receiver noise w_k, but also the crosstalk from other users introduced by α. Each user then passes this signal through a modulo function that removes the effects of the unknown l_k and uses a turbo decoder to decode its intended data u_k. Since we are making comparisons with the ergodic sum-capacity (3), we allow the channel matrix H to be chosen randomly with every use. This randomly-chosen effect is obtained on a smoothly varying channel by using an interleaver over a long block of many consecutive channel uses.

To compare our results with the sum-capacity, we first present our operating points. We examine both $M = K = 4$ and $M = K = 10$, using 16-QAM constellations with either rate

$r = 1/2$ (2 bits/user) and rate $r = 3/4$ (3 bits/user) codes. The sum rate is therefore

$$R_{\text{sum}} = 4rK \quad \text{bits/channel use.} \tag{26}$$

The possible sum-rates for $r = 1/2$ are therefore $R_{\text{sum}} = 8$ bits/channel-use and $R_{\text{sum}} = 20$ for $K = 4$ and $K = 10$, respectively. The sum-rates for $r = 3/4$ are $R_{\text{sum}} = 12$ and $R_{\text{sum}} = 30$. To find the receiver operating points that correspond to these sum rates, we turn to Figure 3 which shows the sum capacity for $K = 4$ and $K = 10$ systems as a function of $\rho = 1/\sigma^2$. These sum-capacity curves are computed by evaluating (3) numerically (we omit the details). The operating point for $r = 1/2$ is approximately $\rho = 7$ dB for either $K = 4$ or $K = 10$, and the operating point for $r = 3/4$ is approximately $\rho = 11.2$ dB for either K.

We briefly discuss a measure of bit-energy that we also include in our performance curves. The total transmitted signal power is normalized to unity, and the channel fading coefficients have unit variance. Therefore, for a transmitted signal that is independent of the channel, the signal received at each user also has unit power. The additive noise at each user has variance $1/\rho$. We may therefore define

$$\frac{E_b}{N_0} = \frac{\rho}{rb},$$

where b is the number of bits per constellation symbol ($b = 4$ for a 16-QAM constellation) and r is the channel code rate (rb is then the number of information bits per user). This measure of bit-energy is independent of K. Our transmitted signal is channel-dependent, thus potentially changing the power of the received signal.

The 16-QAM constellation is mapped to bits using a standard Gray mapping. Successful bit-level turbo decoding at the receiver requires accurate knowledge of the likelihood function of the transmitted bits. The signal received by user k is given by (25) where the additive w'_k is approximately Gaussian. Then y_k is passed through the function $f_{\tau/\sqrt{\gamma}}(\cdot)$, (see Section III), the output of which is fed to the turbo decoder. The modulo function $f_{\tau/\sqrt{\gamma}}(\cdot)$ operates on w'_k and the resulting likelihood function is no longer Gaussian [25]. We compute the new likelihood

function; for simplicity, we confine our analysis to the real parts of y_k, u_k, and w'_k. The imaginary parts are handled in an identical manner.

Let $y'_{k,\mathrm{r}} = f_{\tau/\sqrt{\gamma}}(y_{k,\mathrm{r}})$, where the "r" subscript denotes "real part". By definition, $f_{\tau/\sqrt{\gamma}}(y_{k,\mathrm{r}}) = f_{\tau/\sqrt{\gamma}}(y_{k,\mathrm{r}} + (\tau/\sqrt{\gamma})l)$ for any $y_{k,\mathrm{r}}$ and for any integer l. Therefore,

$$p(y'_{k,\mathrm{r}}|u_{k,\mathrm{r}}) = \sum_{m=-\infty}^{\infty} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\{-(y'_{k,\mathrm{r}} - (u_{k,\mathrm{r}} - m\tau)/\sqrt{\gamma})^2/2\sigma^2\}, \qquad y'_{k,\mathrm{r}} \in \left[-\frac{\tau}{2\sqrt{\gamma}}, \frac{\tau}{2\sqrt{\gamma}}\right). \tag{27}$$

In practice we may approximate the infinite sum in (27) by a sum of a few terms on both sides of $m = 0$. The bit-wise turbo decoder requires a likelihood ratio between having transmitted a binary "1" or "0" and we can use (27) to form this ratio. The difference between the density (27) and a standard Gaussian density can be dramatic especially at low SNR.

Figures 4 and 5 show the results of our trials. The turbo coders are formed from a UMTS standard parallel concatenated code with systematic component, feedforward polynomial $1 + D + D^3$, feedback polynomial $1 + D^2 + D^3$, and block length of 4000 bits. We assumed that the channel H is interleaved and is therefore independent but known from channel-use to channel-use, so as to make the comparison with the ergodic capacities in Figure 3 meaningful. We see that the combination of regularization and vector perturbation performs to within approximately 4 dB from capacity and is significantly better than regularization alone. The $K = 10$ user system performs better than $K = 4$. This is perhaps surprising since the total system sum rate for $K = 10$ is 2.5 times the sum rate for $K = 4$. Perhaps this is because the larger system showed less variability in the transmission scale factor γ.

V. Discussion

We are therefore transmitting at very high sum-rates (tens of bits/channel use) and we are reasonably close to capacity. There are ways to get closer that we have yet to explore: i) Match the turbo code carefully to the channel or increase its block length; ii) Transmit at higher rates to the users whose channels happen to be best, since the sum-capacity is not necessarily attained

by transmitting at equal rates to all of the users; iii) Compute and overcome the penalty for using the modulo operation at the receiver.

We have almost no analysis of the combination of regularization and perturbation, nor do we have any information-theoretic limit for the basic perturbation method. Although not rigorous, our analysis of γ in Section III-B predicts the approximate behavior of the integer minimization (8) and suggests that the basic perturbation algorithm should work for any K, limited only by the complexity of the minimization (8). The sphere encoder allows us to handle up to $K = 15$ with relative ease. The transmitter power normalization constant γ seems to go to a limiting constant as $K \to \infty$, implying that we may be a fixed distance from the sum-capacity for any K. We would like a theory that predicts the limiting value of γ as $K \to \infty$.

One way to achieve capacity on "dirty-paper" and "dirty-tape" interference channels is to encode n consecutive symbols using an n-dimensional lattice [26]. At the receiver, a modulo operation is performed against the lattice, removing the influence of the interference. This technique may be applied to our problem as well; a multi-user signal may be encoded with a temporal lattice and applied to the inverse channel H^{-1} to possibly reduce γ even further. A lattice technique was shown to achieve the diversity-vs-multiplexing tradeoff for point-to-point multiple-antenna channels in [27]. Perhaps lattice techniques could be used successfully on our problem.

Another area we treated only superficially is computing the exact effect on performance of normalizing at the transmitter with $\sqrt{\gamma}$ versus normalizing with $\sqrt{\mathrm{E}\,\gamma}$ (equation (2)). The most practical choice is $\sqrt{\mathrm{E}\,\gamma}$ because the receivers then do not need to know γ. We, however, chose $\sqrt{\gamma}$ in our paper for three reasons: i) $\mathrm{E}\,\gamma$ does not need to exist (this is important in Part I of the paper); ii) It is simpler in our simulations to instantly compute $\sqrt{\gamma}$ rather than to compute $\mathrm{E}\,\gamma$; iii) We found the performance difference to be very small. For example, we show in Figure 6 the bit probability of error for rate $r = 3/4$ turbo-encoded data using 16-QAM symbols (same scenario as in Figure 4 for $M = K = 10$) when normalizing by $\sqrt{\gamma}$ (instantaneous) and by

$\sqrt{\mathrm{E}\,\gamma}$ (average). We see that the performance is actually improved very slightly by normalizing by $\sqrt{\mathrm{E}\,\gamma}$.

When the transmitter normalizes by $\sqrt{\mathrm{E}\,\gamma}$, the receivers do not need to anything at all about the channel for our techniques to work. Perhaps we should be comparing our results with the channel capacity that is attained when only the transmitter knows the channel. Unfortunately, this capacity is apparently not as easy to compute as when both transmitter and receivers know the channel.

We mention that we have not considered how to handle multiple receive antennas at the terminals, or how to find α_{opt} analytically in the cost function (22). We have also not analyzed the optimum τ to choose; for example, increasing τ reduces decoding errors due to the mod-operation but increases γ. Finally, we have also not discussed how to handle users with differing average received signal power. This extension would be particularly useful for systems where there are many users and some are much nearer to the access point than others.

Other possible applications of our technique include (single or multiple-antenna) CDMA systems, where a transmitter communicates with multiples users through spreading sequences that are nonorthogonal because of intersymbol interference. Our algorithm provides a method for compensating for the cross-coupling in the matrix of effective spreading codes H. We leave the analysis of this possible application for future work.

Since the appearance of a preliminary version of this work [28], Shi and Schubert [29] have considered applying our methods with unequal SINR's at each user, and Windpassinger *et al.* [30] have used the LLL algorithm [31] to lower the complexity of minimizing the cost function (22).

REFERENCES

[1] C. B. Peel, B. M. Hochwald, and A. L. Swindlehurst, "A Vector-Perturbation Technique for Near-Capacity Multi-Antenna Multi-User Communication—Part I: Channel Inversion and Regularization," *Submitted to IEEE Trans. Comm.*, Sept. 2003. Available from http://mars.bell-labs.com/.

[2] M. Tomlinson, "New automatic equaliser employing modulo arithmetic," *Electronics Letters*, vol. 7, pp. 138–139, March 1971.

[3] H. Harashima and H. Miyakawa, "Matched-transmission technique for channels with intersymbol interference," *IEEE Trans. Communication*, vol. 20, pp. 774–780, August 1972.

[4] R. D. Wesel and J. M. Cioffi, "Achievable rates for Tomlinson-Harashima precoding," *IEEE Trans. Information Theory*, vol. 44, pp. 824–830, March 1998.

[5] R. F. H. Fischer, *Precoding and Signal Shaping for Digital Transmission*. IEEE-Wiley, 2002.

[6] R. Fischer, C. Windpassinger, A. Lampe, and J. Huber, "MIMO precoding for decentralized receivers," in *Proceedings IEEE International Symposium on Information Theory*, (Lausanne, Switzerland), p. 496, June/July 2002.

[7] C. Windpassinger, R. F. H. Fischer, T. Vencel, and J. B. Huber, "Precoding in multi-antenna and multi-user communications," *Accepted for publication in IEEE Transactions on Wireless Communications*, 2003.

[8] U. Fincke and M. Pohst, "Improved methods for calculating vectors of short lengths in a lattice, including a complexity analysis," *Mathematics of Computation*, vol. 44, pp. 463–471, April 1985.

[9] E. Agrell, T. Eriksson, A. Vardy, and K. Zeger, "Closest point searches in lattices," *IEEE Trans. Information Theory*, vol. 48, pp. 2201–2214, August 2002.

[10] B. M. Hochwald and S. Vishwanath, "Space-time multiple access: Linear growth in the sum rate," in *Proceedings 40th Allerton Conference on Computers, Communications and Control*, (Monticello, Illinois), October 2002.

[11] R. Kannan, "Improved algorithms for integer programming and related lattice problems," in *Proceedings of the ACM Symposium on the Theory of Computing*, (Boston), pp. 193–206, April 1983.

[12] B. Hassibi and H. Vikalo, "On the expected complexity of integer least-squares problems," in *Proceedings ICASSP*, (Orlando), 2002.

[13] M. O. Damen, H. El Gamal, and G. Caire, "On maximum likelihood detection and the search for the closest lattice point," *IEEE Transactions on Information Theory*, pp. 2389–2402, October 2003.

[14] W. H. Mow, "Universal lattice decoding: Principle and recent advances," *Wireless Communications and Mobile Computing*, vol. 3, no. 5, pp. 553–569, 2003.

[15] M. O. Damen, A. Chkeif, and J.-C. Belfiore, "Lattice code decoder for space-time codes," *IEEE Communications Letters*, vol. 4, pp. 161–163, May 2000.

[16] B. Hochwald and S. ten Brink, "Achieving near-capacity on a multi-antenna channel," *IEEE Trans. Comm.*, pp. 389–399, March 2003.

[17] A. Edelman, *Eigenvalues and Condition Numbers of Random Matrices.* PhD Thesis, MIT, Dept. of Math., 1989.

[18] G. Caire and S. Shamai, "On the achievable throughput of a multi-antenna Gaussian broadcast channel," *IEEE Transactions on Information Theory*, vol. 43, pp. 1691–1706, July 2003.

[19] G. Ginis and J. Cioffi, "A multi-user precoding scheme achieving crosstalk cancellation with application to DSL systems," in *Conference Record of the Thirty-Fourth Asilomar Conference on Signals, Systems and Computers*, vol. 2, pp. 1627–1631,

2000.

[20] Z. Tu and R. S. Blum, "Multiuser Diversity for a Dirty Paper Approach," *IEEE Communications Letters*, vol. 7, pp. 370–372, August 2003.

[21] G. J. Foschini, G. D. Golden, R. A. Valenzuela, and M. J. Gans, "Simplified processing for high spectral efficiency wireless communication employing multi-element arrays," *IEEE JSAC*, vol. 17, pp. 1841–1852, Nov 1999.

[22] C. Windpassinger, T. Vencel, and R. F. H. Fischer, "Precoding and loading for BLAST-like systems," in *Proceedings of the IEEE International Conference on Communications*, vol. 5, (Anchorage, Alaska), pp. 3061–3065, May 2003.

[23] M. Costa, "Writing on dirty paper," *IEEE Trans. Information Theory*, vol. 29, pp. 439–441, May 1983.

[24] U. Erez, S. Shamai, and R. Zamir, "Capacity and lattice strategies for cancelling known interference," in *Proceedings International Symposium on Information Theory and its Applications*, (Honolulu, Hawaii), pp. 681–684, Nov. 2000.

[25] S. ten Brink, "Private communication," 2003.

[26] U. Erez and R. Zamir, "Lattice decoding can achieve $\frac{1}{2}\log(1+SNR)$ on the AWGN channel using nested codes," in *Proceedings IEEE International Symposium on Information Theory*, (Washington D.C.), p. 125, June 2001.

[27] H. El Gamal, G. Caire, and M. O. Damen, "On the Role of MMSE in Lattice Decoding: Achieving the Optimal Diversity-vs-Multiplexing Tradeoff," in *Proc. Allerton Conference on Communication, Control, and Computing*, October 2003.

[28] B. M. Hochwald and C. B. Peel, "Vector Precoding for the Multi-Antenna, Multi-User Channel," in *Proc. Allerton Conference on Communication, Control, and Computing*, October 2003.

[29] S. Shi and M. Schubert, "Precoding and power loading for multi-antenna broadcast channels," in *Proceedings Conf. on Information Sciences and Systems*, (Princeton, NJ), March 2004.

[30] C. Windpassinger, R. F. H. Fischer, and J. B. Huber, "Lattice-reduction-aided broadcast precoding," in *Proceedings 5th ITG Conference on Source and Channel Coding*, (Erlangen, Germany), Jan. 2004.

[31] A. K. Lenstra, H. W. Lenstra, and L. Lovász, "Factoring polynomials with rational coefficients," *Mathematische Annalen*, pp. 515–534, 1982.

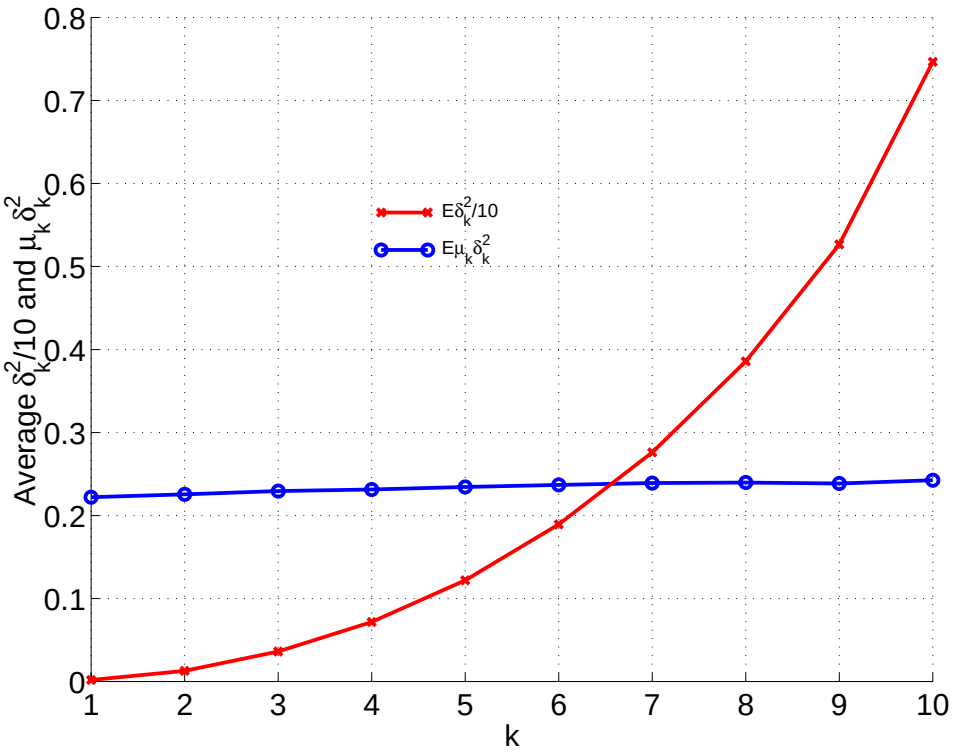


Fig. 1. Numerical example of average δ_k^2 and $\mu_k \delta_k^2$ for $K = 10$. A constellation of 16-QAM symbols was simulated and the cost function (8) was minimized over $\boldsymbol{l}$. We have $\mathrm{E}\,\tilde{\mathbf{u}} = 23.7$, $\mathrm{E}\,\gamma = 2.34$, and $\nu = 0.25$. Thus $\mathrm{E}\,\gamma$ nearly achieves its lower bound $\mathrm{E}\,\gamma \geq 1.55$ given in (17). We also see that $\mathrm{E}\,\mu_k \delta_k^2$ approximately obeys (18).

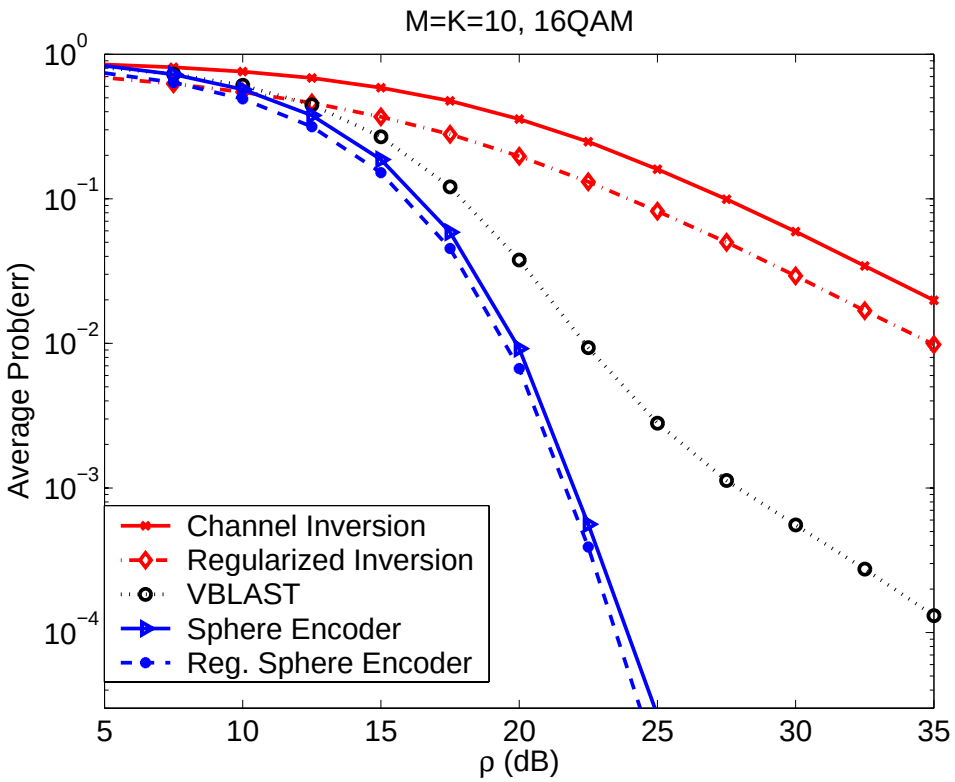


Fig. 2. Uncoded probability of symbol error for channel inversion [1] (x's, solid line), regularized inversion [1] (diamonds, dash-dotted line), the VBLAST algorithm (19) (o's, dotted line), the vector perturbation algorithm (8) (triangles), and the regularized perturbation technique (22) (dashed line, using $\alpha = 1/\rho$—see Section IV).

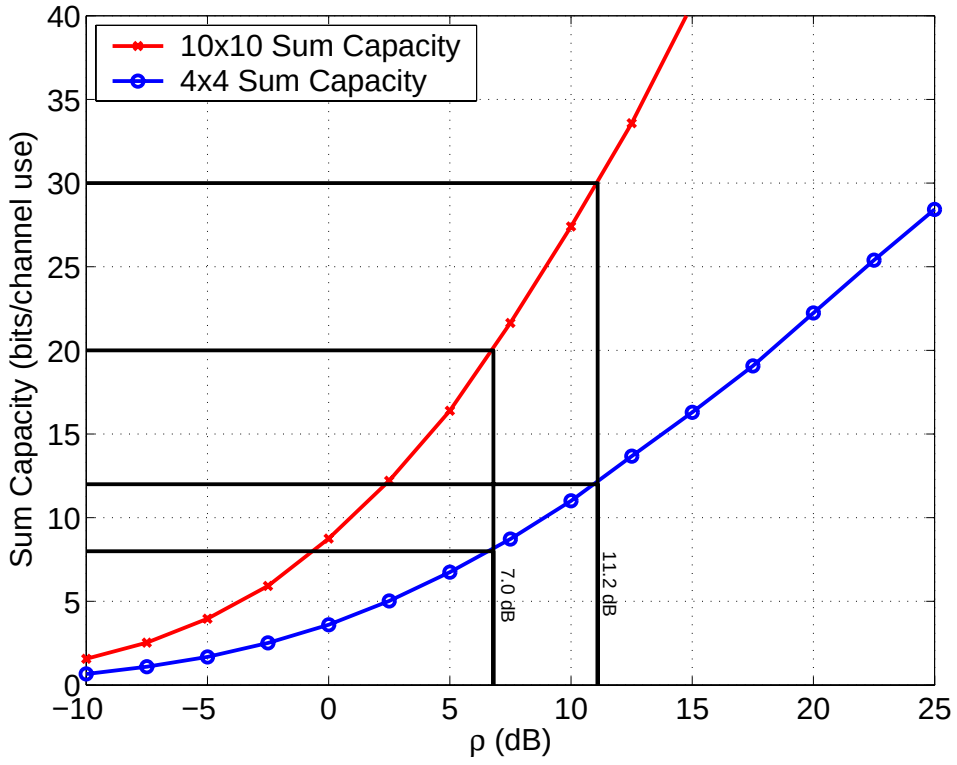


Fig. 3. Sum-capacity for $M = K = 4$ (lower curve) and $M = K = 10$ (upper curve) as a function of the receiver additive noise variance. The marker lines show that to achieve $C = 8$ ($K = 4$) or $C = 20$ ($K = 10$) the (reciprocal) noise variance must be $\rho = 1/\sigma^2 = 7$ dB. For $C = 12$ ($K = 4$) or $C = 30$ ($K = 10$) the noise variance must be $\rho = 11.2$ dB.

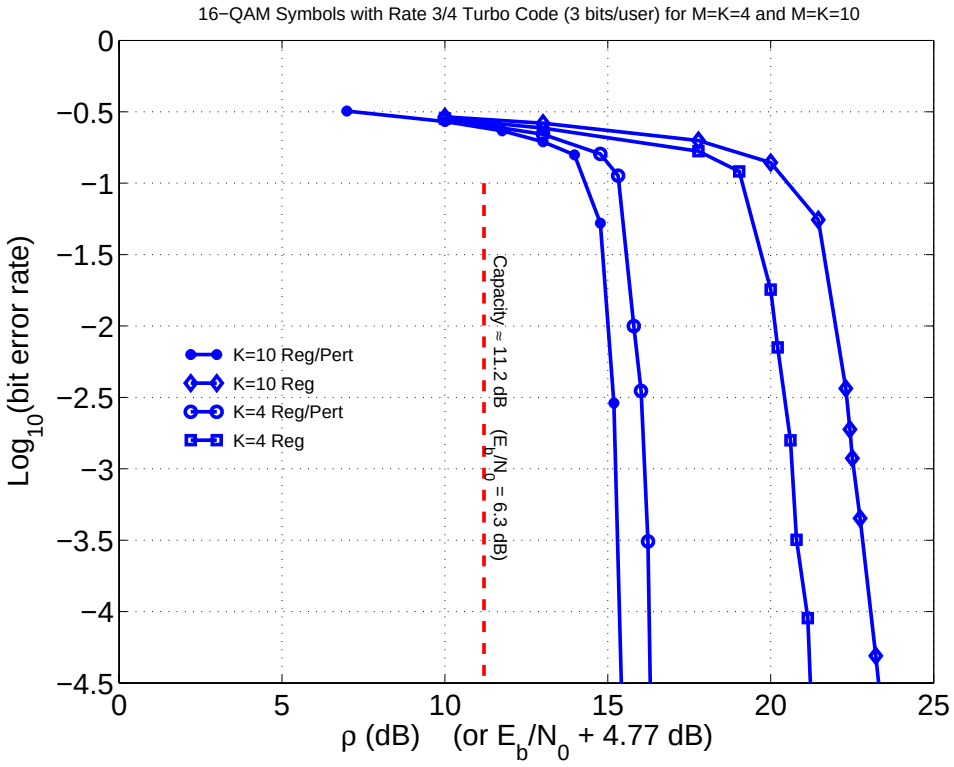


Fig. 4. Bit probability of error for rate $r = 3/4$ turbo-encoded data using 16-QAM symbols, for $M = K = 4$ and $M = K = 10$. We are transmitting three information bits per user for a sum-rate of $3K$ bits/channel-use. The performance of our vector perturbation method for $K = 10$ is the line closest to capacity (approximately 4 dB away). The equivalent line for $K = 4$ is a little further away. The combination of regularization and vector perturbation performs much better than regularized inversion alone.

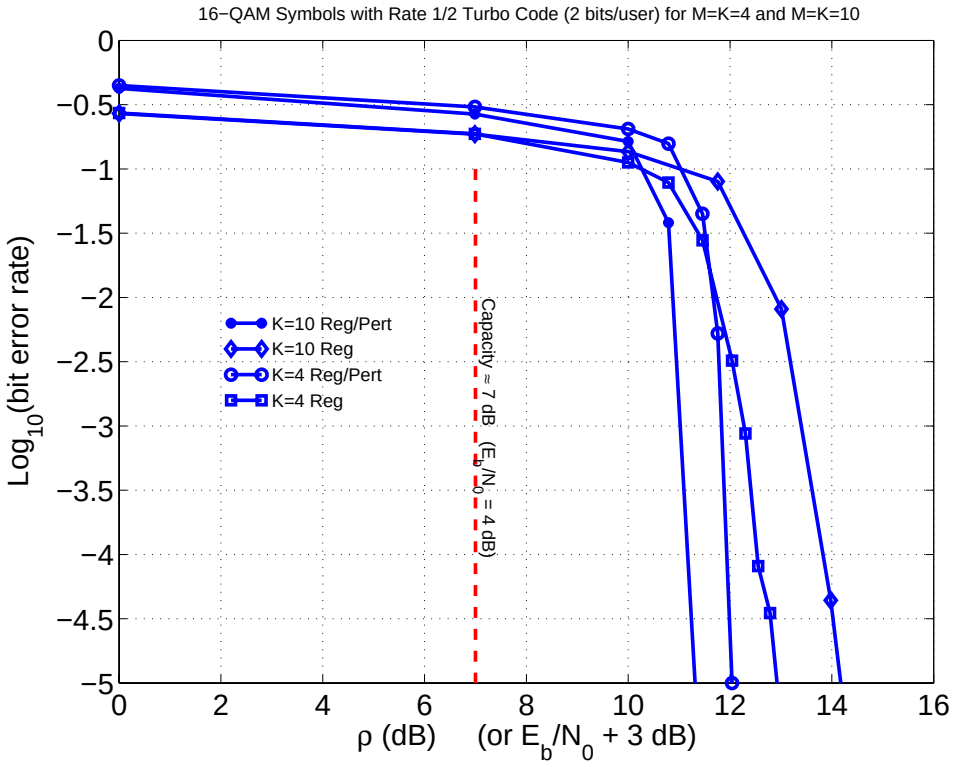


Fig. 5. Bit probability of error for rate $r = 1/2$ turbo-encoded data using 16-QAM symbols, for $M = K = 4$ and $M = K = 10$. The performance of combining regularization and perturbation is much better than regularization alone; the best performance is obtained for $M = K = 10$.

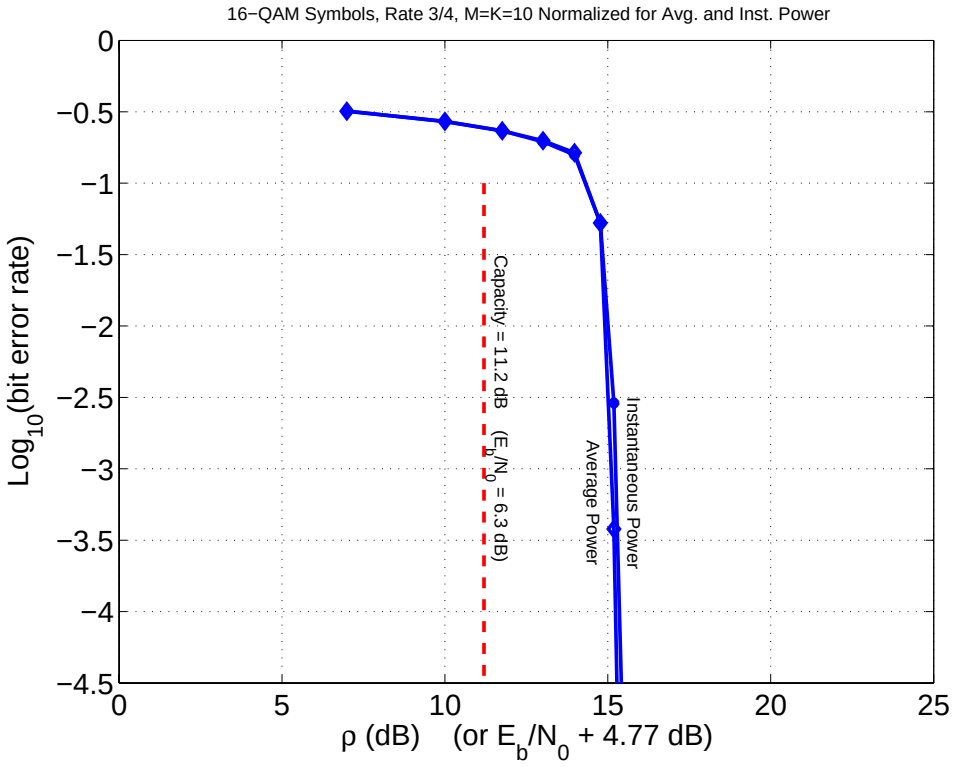


Fig. 6. Bit probability of error for rate $r = 3/4$ turbo-encoded data using 16-QAM symbols, for $M = K = 10$. The curve on the right (Instantaneous power constraint) uses normalization at the transmitter by $\sqrt{\gamma}$; the curve on the left (Average power constraint) uses normalization by $\sqrt{\mathrm{E}\,\gamma}$. The difference in performance is small.