

Optimal Sequential Frame Synchronization

Venkat Chandar, Aslan Tchamkerten, and Gregory Wornell, *Fellow, IEEE*

Abstract—We consider the “one-shot frame synchronization problem,” where a decoder wants to locate a sync pattern at the output of a memoryless channel on the basis of sequential observations. The sync pattern of length N starts being emitted at a random time within some interval of size A , where A characterizes the asynchronism level. We show that a sequential decoder can optimally locate the sync pattern, i.e., exactly, without delay, and with probability approaching one as $N \rightarrow \infty$, if the asynchronism level grows as $O(e^{N\alpha})$, with α below the *synchronization threshold*, a constant that admits a simple expression depending on the channel. If α exceeds the synchronization threshold, any decoder, sequential or nonsequential, locates the sync pattern with an error that tends to one as $N \rightarrow \infty$. Hence, a sequential decoder can locate a sync pattern as well as the (nonsequential) maximum-likelihood decoder that operates on the basis of output sequences of maximum length $A + N - 1$, but with far fewer observations.

Index Terms—Frame synchronization, pattern recognition, quickest detection, sequential analysis.

I. INTRODUCTION

Frame synchronization refers to the problem of locating a sync pattern embedded into data and received over a channel (see, e.g., [4]–[7]). In [5], Massey considered the situation of binary data transmitted across a white Gaussian noise channel. He showed that given received data of fixed size which the sync pattern is known to belong to, the maximum-likelihood rule consists of selecting the location that maximizes the sum of the correlation and a correction term.

We are interested in the situation where the receiver wants to locate the sync pattern on the basis of sequential observations, which Massey refers to as the “one-shot” frame synchronization problem in [5]. A receiver observes data sequentially, with the foreknowledge that a sync pattern will occur within a certain time interval with probability one. The length of this time interval represents the asynchronism level. The receiver’s goal is to locate the sync pattern exactly and without delay. Our result is an asymptotic characterization of the largest asynchronism level with respect to the size of the sync pattern for which a decoder can correctly perform with arbitrarily high probability.

We note that a similar problem formulation has been studied in [1], where various pattern detection rules have been studied, yet without deriving fundamental limitations in terms of sync pattern size, asynchronism level, and mislocation probability.

II. PROBLEM FORMULATION AND RESULT

We consider discrete-time communication over a discrete memoryless channel characterized by its finite input and output alphabets \mathcal{X} and \mathcal{Y} , respectively, transition probability matrix $Q(y|x)$, for all $y \in \mathcal{Y}$ and

Manuscript received August 17, 2007; revised February 19, 2008. This work was supported in part by the National Science Foundation (NSF) under Grant CCF-0515122, by the Department of Defense (DoD) under MURI Grant N00014-07-1-0738, and by a University IR&D Grant from Draper Laboratory.

The authors are with the Electrical Engineering and Computer Science Department, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: vchandar@mit.edu; tcham@mit.edu; gww@mit.edu).

Communicated by H. Bölcskei, Associate Editor for Detection and Estimation.

Digital Object Identifier 10.1109/TIT.2008.926444




Fig. 1. Time representation of what is sent (upper arrow) and what is received (lower arrow). The “*” represents the “data” symbol. The sync pattern starts being sent at time ν , and is detected at time τ .

$x \in \mathcal{X}$, and “data” symbol $\star \in \mathcal{X}$.¹ We discuss the extension of our main result to the Gaussian channel at the end of the correspondence.

The sync pattern s^N consists of $N \geq 1$ symbols from \mathcal{X} —possibly also the \star symbol. The transmission of the sync pattern starts at a random time ν , uniformly distributed in $[1, 2, \dots, A = e^{\alpha N}]$, where the integer $A \geq 1$ characterizes the asynchronism level and α the asynchronism exponent.

We assume that the receiver knows A but not ν . Before and after the transmission of the sync pattern, i.e., before time ν and after time $\nu + N - 1$, the receiver observes random data. Specifically, conditioned on the value of ν , the receiver observes independent symbols Y_1, Y_2, \dots distributed as follows. If $i \leq \nu - 1$ or $i \geq \nu + N$, the distribution is $Q(\cdot|\star)$. At any time $i \in [\nu, \nu + 1, \dots, \nu + N - 1]$ the distribution is $Q(\cdot|s_{i-\nu+1})$, where s_n denotes the n th symbol of s^N .

To identify the instant when the sync pattern starts being emitted, the receiver uses a sequential decoder in the form of a stopping time τ with respect to the output sequence Y_1, Y_2, \dots .² If $\tau = n$, the receiver declares that the sync pattern started being sent at time $n - N + 1$ (see Fig. 1).

The associated error probability is defined as

$$\mathbb{P}(\tau \neq \nu + N - 1).$$

We now define the *synchronization threshold*.

Definition: An asynchronism exponent α is achievable if there exists a sequence of pairs sync pattern/decoder $\{(s^N, \tau_N)\}_{N \geq 1}$ such that s^N and τ_N operate under asynchronism level $A = e^{\alpha N}$, and so that

$$\mathbb{P}(\tau_N \neq \nu + N - 1) \xrightarrow{N \rightarrow \infty} 0.$$

The *synchronization threshold*, denoted $\alpha(Q)$, is the supremum of the set of achievable asynchronism exponents.

Our main result lies in the following theorem.

Theorem: The synchronization threshold as defined above is given by

$$\alpha(Q) = \max_x D(Q(\cdot|x)||Q(\cdot|\star))$$

where $D(Q(\cdot|x)||Q(\cdot|\star))$ is the Kullback–Leibler distance between $Q(\cdot|x)$ and $Q(\cdot|\star)$. Furthermore, if the asynchronism exponent is above the synchronization threshold, a maximum-likelihood decoder that is revealed the maximum length sequence of size $A + N - 1$ makes an error with a probability that tends to one as $N \rightarrow \infty$.

¹Throughout this note, we assume that for all $y \in \mathcal{Y}$ there is some $x \in \mathcal{X}$ for which $Q(y|x) > 0$.

²Recall that a (deterministic or randomized) stopping time τ with respect to a sequence of random variables $\{Y_i\}_{i=1}^\infty$ is a positive, integer-valued random variable such that the event $\{\tau = n\}$, conditioned on $\{Y_i\}_{i=1}^n$, is independent of $\{Y_i\}_{i=n+1}^\infty$ for all $n \geq 1$.

A direct consequence of the theorem is that a sequential decoder can (asymptotically) locate the sync pattern as well as the optimal maximum-likelihood decoder that operates on a nonsequential basis having access to sequences of maximum size $A + N - 1$.

Note that the synchronization threshold is the same as the one in [8], which is defined as the largest asynchronism level for which reliable communication can be achieved over point-to-point asynchronous channels. This should not come as a surprise since, for asynchronous communication in the zero-rate regime, decoding errors are mainly due to incorrectly identifying the location of the transmitted message.

We now prove the theorem by first presenting the direct part and then its converse. Recall that a type, or empirical distribution, induced by a sequence $z^N \in \mathcal{Z}^N$ is the probability measure \hat{P} on \mathcal{Z} where $\hat{P}(a)$, $a \in \mathcal{Z}$, is equal to the number of occurrences of a in z^N divided by N . Similarly, the empirical distribution induced by two sequences $z^N \in \mathcal{Z}^N$ and $w^N \in \mathcal{W}^N$ is the probability measure \hat{P} on $\mathcal{Z} \times \mathcal{W}$ such that $\hat{P}(a, b)$, $a \in \mathcal{Z}$, $b \in \mathcal{W}$, is equal to the number of occurrences of the pair (a, b) among the pairs (z_i, w_i) , $1 \leq i \leq N$, divided by N . We define $P(x, y) = \hat{P}_s(x)Q(y|x)$, where $\hat{P}_s(x)$ is the empirical distribution of the sync pattern.

Proof of Achievability: We show that a suitable sync pattern together with the sequential typicality decoder³ achieves an asynchronism exponent arbitrarily close to $\max_x D(Q(\cdot|x)||Q(\cdot|\star))$. The intuition is as follows. Let \bar{x} be a “maximally divergent symbol,” i.e., so that $D(Q(\cdot|\bar{x})||Q(\cdot|\star)) = \max_x D(Q(\cdot|x)||Q(\cdot|\star))$. Suppose the sync pattern consists of N repetitions of \bar{x} . If we use the sequential typicality decoder, we have almost all the properties we need. Indeed, if $\alpha < \max_x D(Q(\cdot|x)||Q(\cdot|\star))$, with negligible probability the random data generates a block of N output symbols that is jointly typical with the sync pattern. Similarly, the block of output symbols generated by the sync pattern is jointly typical with the sync pattern with high probability. The only problem occurs when a block of N output symbols is generated partly by the data and partly by the sync pattern. Indeed, consider for instance the block of N output symbols from time $\nu - 1$ up to $\nu + N - 2$. These symbols are all generated according to the sync pattern, except for the first. Hence, whenever the decoder observes this portion of symbols, it makes an error with constant probability. The argument extends to any fixed length shift.

The reason that the decoder is unable to locate the sync pattern exactly is that the sync pattern used above has the undesirable property that when it is shifted to the right, it still looks almost the same. Therefore, to prove the direct part of the theorem, we consider a sync pattern mainly composed of \bar{x} 's, but with a few \star 's mixed in⁴ so that shifts of the sync pattern look sufficiently different from the original sync pattern. This allows the decoder to identify the sync pattern exactly, with no delay, and with probability tending to one as N goes to infinity, for any asynchronism exponent less than $\max_x D(Q(\cdot|x)||Q(\cdot|\star))$. We formalize this below.

Suppose that, for any arbitrarily large K , we can construct a sequence of patterns $\{s^N\}$ of increasing lengths such that each $s^N = s_1, s_2, \dots, s_N$ satisfies the following two properties:

³The sequential typicality decoder operates as follows. At time n , it computes the empirical distribution \hat{P} induced by the sync pattern and the previous N output symbols $y_{n-N+1}, y_{n-N+2}, \dots, y_n$. If this distribution is close enough to P , i.e., if $|\hat{P}(x, y) - P(x, y)| \leq \mu$ for all x, y , the decoder stops, and declares $n - N + 1$ as the time the sync pattern started being emitted. Otherwise, it moves one step ahead and repeats the procedure. Throughout the argument we assume that μ is a negligible strictly positive quantity.

⁴Indeed, any symbol different than \bar{x} can be used.

- I. all s_i 's are equal to \bar{x} , except for a fraction at most equal to $1/K$ that are equal to \star ;
- II. the Hamming distance between the pattern and any of its shifts of the form

$$\underbrace{\star, \star, \dots, \star}_{i \text{ times}}, s_1, s_2, \dots, s_{N-i} \quad i \in [1, 2, \dots, N]$$

is linear in N .

Now let $A = e^{N(\max_x D(Q(\cdot|x)||Q(\cdot|\star)) - \epsilon)}$, for some $\epsilon > 0$, and consider using patterns with properties I and II in conjunction with the sequential typicality decoder, i.e., we take τ_N to be the sequential typicality decoder operating on blocks of size N .

By [2, Lemma 2.6, p. 32] and property I, the probability that N output symbols entirely generated by the random data are typical with the sync pattern is upper-bounded by

$$\exp(-N(1 - 1/K)(\max_x D(Q(\cdot|x)||Q(\cdot|\star)) - \delta))$$

where $\delta > 0$ goes to zero as the typicality constant μ goes to zero.⁵ Hence, by the union bound

$$\mathbb{P}(\{\tau_N < \nu\} \cup \{\tau_N \geq \nu + 2N - 1\}) \leq e^{-N(\epsilon - \delta - (\max_x D(Q(\cdot|x)||Q(\cdot|\star)) - \delta)/K)}$$

which tends to zero for μ small enough and K sufficiently large.⁶ If the N observed symbols are generated partly by the data and partly by the sync pattern, by property II, the Chernoff bound, and the union bound we obtain

$$\mathbb{P}(\tau_N \in [\nu, \nu + 1, \dots, \nu + N - 2]) \leq (N - 1)e^{-\Omega(N)}$$

which vanishes as N tends to infinity. We then deduce that

$$\mathbb{P}(\tau_N = \nu + N - 1) \rightarrow 1$$

as $N \rightarrow \infty$.

To conclude, we give an explicit construction of a sequence of sync patterns satisfying the properties I and II above. To that aim, we use maximal-length shift register sequences (see, e.g., [3]). Actually, for our purposes, the only property we use from such binary sequences of length $l = 2^m - 1$, $m \in [1, 2, \dots]$, is that they are of Hamming distance $(l + 1)/2$ from any of their circular shifts.

Pick some large K that satisfies $\lfloor \frac{N}{K} \rfloor = 2^m - 1$ for some $m \in [1, 2, \dots]$.⁷ We start by setting $s_i = \bar{x}$ for all $\lfloor \frac{N}{K} \rfloor < i \leq N$. With this choice, property I is already satisfied regardless of the first $\lfloor \frac{N}{K} \rfloor$ symbols of the pattern. To specify the latter, pick a maximal length shift register sequence $m_1, m_2, \dots, m_{\lfloor \frac{N}{K} \rfloor}$, and set $s_j = \bar{x}$ if $m_j = 0$ and $s_j = \star$ if $m_j = 1$, for any $j \in [1, 2, \dots, \lfloor \frac{N}{K} \rfloor]$. It can be readily verified, using the circular shift property of maximal length shift register sequences, that this construction yields patterns satisfying property II. \square

Proof of the Converse: We assume that $A = e^{N\alpha}$ with

$$\alpha > \max_x D(Q(\cdot|x)||Q(\cdot|\star))$$

and show that the (optimal) maximum-likelihood decoder that operates on the basis of sequences of maximum length $A + N - 1$ yields a probability of error going to one as N tends to infinity.

⁵See footnote 3.

⁶If $\max_x D(Q(\cdot|x)||Q(\cdot|\star)) = \infty$, the upper bound is zero if μ is small enough.

⁷We use $\lfloor x \rfloor$ to denote the largest integer smaller than x .




Fig. 2. Parsing of the entire received sequence of size $A + N - 1$ into r blocks $y^{(t_1)}, y^{(t_2)}, \dots, y^{(t_r)}$ of length N , where the i th block starts at time t_i .

We assume that the sync pattern s^N is composed of N identical symbols $s \in \mathcal{X}$. The case with multiple symbols is obtained by a straightforward extension. Suppose the maximum-likelihood decoder not only is revealed the complete sequence

$$y_1, y_2, \dots, y_{A+N-1},$$

but also knows the value of $\nu \bmod n$, i.e., the decoder knows that the sync pattern was sent in one of

$$r = \left\lfloor \frac{A + N - (\nu \bmod N)}{N} \right\rfloor$$

disjoint blocks of duration N , as shown in Fig. 2.

Assuming $Q(y|\star) > 0$ for all $y \in \mathcal{Y}$,⁸ straightforward algebra shows that the decoder outputs the time $t_i, i \in [1, 2, \dots, r]$, that maximizes⁹

$$f(y^{(t_i)}) = \frac{Q(y^{(t_i)}|s^N)}{Q(y^{(t_i)}|\star)}. \quad (1)$$

Note that $f(y^{(t)})$ depends only on the type of the sequence $y^{(t)}$ since s^N is the repetition of a single symbol. For conciseness, from now on we adopt the notation $Q_s(y)$ instead of $Q(y|s)$ and $Q_\star(y)$ instead of $Q(y|\star)$.

Let $Q_s \pm \epsilon_0$ denote the set of types (induced by sequences y^N) that are $\epsilon_0 > 0$ close to Q_s with respect to the L_1 norm, and let E_1 denote the event that the type of the ν th block (corresponding to the sync transmission period) is not in $Q_s \pm \epsilon_0$. It follows that

$$\mathbb{P}(E_1) \leq e^{-N\epsilon} \quad (2)$$

for some $\epsilon = \epsilon(\epsilon_0) > 0$.¹⁰ Let $\bar{Q}_s = \arg \max_{P \in Q_s \pm \epsilon_0} f(P)$,¹¹ where with a slight abuse of notation $f(P)$ is used to denote $f(y^N)$ for any sequence y^N having type P . Now consider the event E_2 where the number of blocks generated by Q_\star that have type \bar{Q}_s is smaller than

$$\frac{1}{2(N+1)^{1+|\mathcal{X}|}} e^{-N(D(\bar{Q}_s||Q_\star) - \alpha)}.$$

Using [2, Lemma 2.6, p. 32], the expected number of blocks generated by Q_\star that have type \bar{Q}_s is lower-bounded as

$$\begin{aligned} & \mathbb{E}(\text{number of type } \bar{Q}_s \text{ blocks generated from } Q_\star) \\ & \geq \frac{1}{(N+1)^{|\mathcal{X}|}} e^{-ND(\bar{Q}_s||Q_\star)} (r-1) \\ & \geq \frac{1}{(N+1)^{1+|\mathcal{X}|}} e^{-N(D(\bar{Q}_s||Q_\star) - \alpha)} \end{aligned}$$

⁸If $Q(y|\star) = 0$ for some $y \in \mathcal{Y}$ we have $\alpha(Q) = \infty$, and there is nothing to prove.

⁹In the case of a tie, we assume that the decoder chooses one of the maximizing t_i 's uniformly at random.

¹⁰Here we implicitly assume that N is large enough so that the set of types $Q_s \pm \epsilon_0$ is nonempty.

¹¹Note that \bar{Q}_s may not be equal to Q_s .

where the last inequality holds for sufficiently large N . Hence, using Chebyshev's inequality, we get

$$\mathbb{P}(E_2) \leq \text{poly}(N) e^{-N(\alpha - D(\bar{Q}_s||Q_\star))} \quad (3)$$

where $\text{poly}(N)$ denotes a term that increases no faster than polynomially in N .

Finally, consider the event E_3 defined as the complement of $E_1 \cup E_2$. Given that E_3 happens, the decoder sees at least

$$\frac{1}{2(N+1)^{1+|\mathcal{X}|}} e^{-N(D(\bar{Q}_s||Q_\star) - \alpha)}$$

time slots whose corresponding ratios (1) are at least as large as the correct ν th. Hence, the probability of correct decoding given that the event E_3 happens is upper-bounded as

$$\mathbb{P}(\text{corr.dec}|E_3) \leq \text{poly}(N) e^{-N(\alpha - D(\bar{Q}_s||Q_\star))}. \quad (4)$$

We deduce from (2), (3), and (4) that the probability of correct decoding is upper-bounded as

$$\begin{aligned} \mathbb{P}(\text{corr. dec.}) &= \sum_{i=1}^3 \mathbb{P}(\text{corr.dec}|E_i) \mathbb{P}(E_i) \\ &\leq \mathbb{P}(E_1) + \mathbb{P}(E_2) + \mathbb{P}(\text{corr.dec}|E_3) \\ &\leq (e^{-N\epsilon} + e^{-N(\alpha - D(\bar{Q}_s||Q_\star))}) \text{poly}(N). \end{aligned}$$

Therefore, if

$$\alpha > D(\bar{Q}_s||Q_\star)$$

the probability of successful decoding goes to zero as N tends to infinity. Since $D(\bar{Q}_s||Q_\star)$ tends to $D(Q_s||Q_\star)$ as $\epsilon_0 \downarrow 0$ by continuity of $D(\cdot||Q_\star)$,¹² the result follows by maximizing $D(Q_s||Q(\cdot|\star))$ over $s \in \mathcal{X}$. \square

III. THE GAUSSIAN CHANNEL

The result presented in the previous section can be extended to the Gaussian channel with a peak power constraint on the input. Specifically, consider a Gaussian channel described by the conditional probability distribution function

$$Q(y|x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(y-x)^2}{2\sigma^2}}$$

where $x, y \in \mathbb{R}$. Suppose the random data has a Gaussian distribution with mean 0 and variance σ^2 , i.e., $Q(y|\star) = Q(y|0)$. The input constraint is that the sync pattern s^N must satisfy $s_i^2 \leq P$ for $1 \leq i \leq n$, where P represents the available power. For this channel, a straightforward extension of the arguments from the previous section shows that the synchronization threshold is given by

$$\alpha(Q) = \frac{1}{2} \text{SNR}$$

where SNR denotes the *signal to noise ratio*, and is defined as P/σ^2 .

¹²We may assume $D(\cdot||Q_\star)$ is continuous because otherwise $\alpha(Q) = \infty$, so there is nothing to prove.

IV. CONCLUDING REMARKS

Similarly to Massey's setting [5], ours assumes that the decoder observes either a noisy version of the sync pattern or a sequence of independent and identically distributed, random variables. In this setting, we derived an optimal tradeoff between pattern size and asynchronism level in the form of the synchronization threshold. In certain applications, however, the output random variables outside the sync pattern period are not i.i.d., and optimal pattern detection procedures remain to be found.

Another issue concerns the definition of optimality. Ours requires that "the decoder must isolate the sync pattern with arbitrarily high probability." However, we have imposed no constraints on how quickly the error probability approaches 0 as $N \rightarrow \infty$. For the problem where we require that the error probability decays with a certain error exponent, the maximum achievable asynchronism exponent remains to be identified.

REFERENCES

- [1] M. Chiani and M. G. Martini, "On sequential frame synchronization in additive white gaussian noise channels," *IEEE Trans. Commun.*, vol. 54, no. 2, pp. 339–348, Feb. 2006.
- [2] I. Csiszàr and J. Körner, *Information Theory: Coding Theorems for Discrete Memoryless Channels*. New York: Academic, 1981.
- [3] S. Golomb, *Shift Register Sequences*. San Francisco, CA: Holden-Day, 1967.
- [4] G. L. Lui and H. H. Tan, "Frame synchronization for direct-detection optical communication systems," *IEEE Trans. Commun.*, vol. COM-34, no. 3, pp. 227–237, Mar. 1986.
- [5] J. L. Massey, "Optimum frame synchronization," *IEEE Trans. Comm.*, vol. 20, no. 2, pp. 115–119, Apr. 1972.
- [6] P. T. Nielsen, "On the expected duration of a search for a fixed pattern in random data," *IEEE Trans. Inf. Theory*, vol. IT-19, no. 5, pp. 702–704, Sep. 1973.
- [7] R. A. Scholtz, "Frame synchronization techniques," *IEEE Trans. Commun.*, vol. COM-28, no. 8, pp. 1204–1212, Aug. 1980.
- [8] A. Tchamkerten, V. Chandar, and G. W. Wornell, "Communication under strong asynchronism," *IEEE Trans. Inf. Theory*. [Online]. Available: <http://arxiv.org/abs/0707.4656>, submitted for publication

A Case for Amplify-Forward Relaying in the Block-Fading Multiple-Access Channel

Deqiang Chen, *Member, IEEE*, Kambiz Azarian, and
J. Nicholas Laneman, *Senior Member, IEEE*

Abstract—This correspondence demonstrates the significant gains that multiple-access users can achieve from *sharing* a single amplify-forward relay in slow-fading environments. The proposed protocol, namely, multiple-access amplify-forward (MAF), allows for a low-complexity relay and achieves the optimal diversity-multiplexing tradeoff (DMT) at high multiplexing gains. Analysis of the protocol further reveals that it outperforms both the compress-forward strategy at low multiplexing gains and the dynamic decode-forward protocol at high multiplexing gains. An interesting feature of the proposed protocol is that, at high multiplexing gains, it resembles a multiple-input single-output (MISO) system, and at low multiplexing gains, it provides each user with the same DMT as if there were no contention for the relay from the other users.

Index Terms—Amplify-forward, block-fading channel, cooperative diversity, diversity-multiplexing tradeoff (DMT), multiple-access relay channel (MARK), wireless networks.

I. INTRODUCTION

A. Motivation

In recent years, cooperative communications has received significant interest (e.g., [1]–[7]) as a means of providing spatial diversity for applications in which temporal, spectral, and antenna diversity are limited by delay, bandwidth, and terminal size constraints, respectively. Cooperative techniques offer diversity by enabling users to utilize one another's resources such as antennas, power, and bandwidth. As a consequence, most cooperative protocols share the characteristic that they require substantial coordination among the users. In a wireless setting, establishing this level of user cooperation may be impractical due to cost and complexity considerations. Inspired by this observation, this correspondence focuses on an alternative architecture, namely, the multiple-access relay channel (MARC) [4], [8] and proposes a strategy called the multiple-access amplify-forward (MAF) that allows the users to operate as if in a normal (noncooperative) multiple-access channel (MAC). In this system, the users need not be aware of the existence of the relay, i.e., all cost and complexity of exploiting cooperative diversity is placed in the relay and destination. Such an architecture may be suitable for infrastructure networks, in which the relay and destination correspond, respectively, to a relay station and a base station deployed and managed by the service provider. It is worth noting that since a single relay is *shared* by multiple users in the MARC, the extra cost of adding the relay is amortized across many users and may thus be more acceptable, especially as the number of

Manuscript received January 5, 2007; revised September 13, 2007. This work was supported in part by the National Science Foundation (NSF) under Grants CCF05-15012 and CNS06-26595. The material in this correspondence was presented in part at the Conference on Information Sciences and Systems (CISS), Princeton, NJ, March 2006 and the IEEE Communication Theory Workshop (CTW), Dorado, Puerto Rico, May 2006.

D. Chen was with the Department of Electrical Engineering, University of Notre Dame, Notre Dame, IN 46556 USA (e-mail: dchen2@alumni.nd.edu).

K. Azarian was with the Department of Electrical Engineering, University of Notre Dame, Notre Dame, IN 46556 USA. He is now with Qualcomm Inc., San Diego, CA 92121 USA (e-mail: kazarian@nd.edu).

J. N. Laneman is with the Department of Electrical Engineering, University of Notre Dame, Notre Dame, IN 46556 USA (e-mail: jnl@nd.edu).

Communicated by E. Modiano, Associate Editor for Communication Networks.

Color versions of Figures 2–4 in this correspondence are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIT.2008.926336