Cooperative Diversity in Wireless Networks with Stochastic and Bursty Traffic

Petros Elia, Somsak Kittipiyakul and Tara Javidi
ECE Department
University of California at San Diego
San Diego, CA 92037
{elia,skittipi,tara}@ucsd.edu

Abstract— **This work investigates the asymptotic error performance of outage-limited communications in cooperative wireless relay networks, with fading that is quasi-static, with an information-arrival process that is stochastic and bursty, and with bits that have a strictly limited lifespan. Employing large-deviation techniques, we analyze the probability of bit error where such errors are due to both erroneous decoding as well as due to delay violation. We derive a tradeoff between, on one hand, the optimal negative SNR exponent of the total probability of error, and on the other hand, the ratio of the average bit-arrival rate to the ergodic capacity of the channel. This is for the case of the orthogonal amplify-and-forward protocol, constant system loading, many flows and asymptotically high values of SNR. The tradeoff holds for any delay limitation. As a practical consequence, the tradeoff tells us how to better balance the effects of channel atypicality (outage) and burstiness atypicality, by proper choice of transmission rate and by optimizing the cooperative cluster size, i.e., limiting the cooperation to a specific subset of the cooperative users.**

I. Introduction

Cooperation among users in a wireless network can substantially improve the reliability of communication [1], [2]. These improvements relate to encoding across space and they fully appear only after some minimum finite amount of time-averaging. As a result, when the bit-arrival process is stochastic and bursty, and when the bits are limited by a strict delay requirement, this required temporal averaging and the corresponding reduction in queue responsiveness, results in an increase of the probability of error due to delay violation. In this paper, we explore this above tradeoff, i.e., we explore the relation between mitigating for channel errors and mitigating for delay-violations of delay-sensitive and bursty information.

A. Time averaging: queue responsiveness v.s. diversity gain

In a cooperative setting and in the presence of fading, network users exchange and transmit functions of each others' signals in a manner that emulates multiple-input, multiple-output (MIMO) communications. In the setting where errors are solely due to channel (deterministic and non-bursty bit-arrival process and/or no delay limitations), cooperative diversity techniques reduce the probability of decoding error by diversifying resources and exploiting fading independence across nodes. In this same setting, error performance relates to, among other things, the transmission's rate-to-power ratio and to the average number of fading paths associated with each bit of information. The above relations were nicely captured, using large deviation techniques, by the concept of the diversity-multiplexing gain tradeoff (DMT) [3], which will be recalled later in the context of cooperative networks. DMT analysis offers the important result that in the outage-limited quasi-static setting of fixed fading, the asymptotically optimal error performance can be achieved with coding over some finite time duration [3], as long as this time duration is larger than some minimum encoding time which is usually a function of the number of transmitting nodes [7].

It is the case though that when cooperation aims to assist the communication of stochastic and bursty information with delay limitations, then this spatiotemporal averaging, which mitigates the stochastic nature of the channel, gets in the way of queue responsiveness, resulting in slower emptying of the queue, and in more bits violating the delay limitation placed by the receiving application.

In this paper, we explore this tradeoff in the relation between mitigating for channel errors and mitigating for delay-violations of delay-sensitive and bursty information. More specifically, we analyze the asymptotic probability of overall bit loss due to both transmission error as well as delay-violation. Under the assumption of asymptotically high scaling of the signal-to-noise ratio (SNR), fixed average packet size, and fixed system loading, we provide an analytical expression for the overall probability of error due to both the channel and delay. This expression is then used to quantify the tradeoff between performance enhancement due to diversity gain and performance degradation due to the reduced responsiveness of the queue. We present this *responsiveness-diversity-multiplexing tradeoff (RDMT)* for the orthogonal symmetric amplify-and-forward (OAF) [4] channel, with minimum encoding delay and for the compound-Poisson bit-arrival process. This tradeoff is in form of a simple expression between channel and application related parameters such as SNR, Rayleigh fading statistics, the rate-to-power ratio, the compound Poisson parameters, and the delay limitation. As a practical consequence, the tradeoff tells us how to pick the information rate and cluster size of cooperative users, in order to balance the effects of channel atypicality (outage) and burstiness atypicality.

The outline of the rest of the paper is as follows: In Section II, we discuss the system model in more detail and recall some existing results. In Section III, we derive an analytical expression for the overall probability of error due to decoding and delay violation. Using this expression, we derive the optimal choice of rate-to-power ratio, the optimal choice of cluster size, and eventually the optimal negative SNR

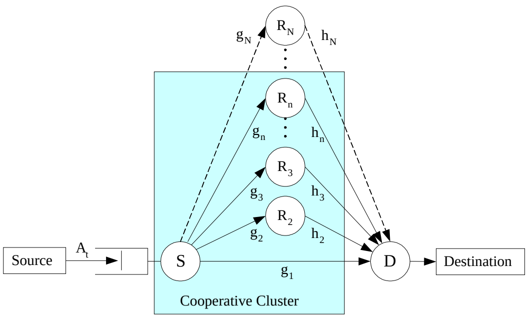

Fig. 1. Cooperative Relay Network with a single bursty source

exponent (RDMT tradeoff) of the overall probability of error. We conclude the paper in Section IV.

II. SYSTEM MODEL

We consider a snapshot of a cooperative network (see Figure 1), where the source-node, denoted by $\mathcal{S}$, receives information according to a stochastic process and stores the bits that cannot be sent immediately in a buffer which is assumed to be infinite. Information bits stochastically arrive at discrete *time-slots*. The source-node transmits at the beginning of *symbol-slots* with the assistance of some subset $\{R_i\}_{i=2}^{n}$ of all possible nodes $\{R_i\}_{i=2}^{N}$. The duration of a symbol-slot is assumed to match with that of a time-slot[1]. The delay requirement asks that each bit of information be decoded at the destination-node $\mathcal{D}$ within a *maximum allowable delay* of D time-slots from the time it arrives at the source-node. Otherwise, the bit will be obsolete, discarded and counted as erroneous. The statistics of the source and channel, as well as the value of D, are assumed to be known and fixed for all time. The performance of the system is a direct function of the number n of transmitting nodes (source-node and $n-1$ relaying participants).

We now provide details and notation for the channel and the arrival process.

A. Cooperative Relay Channel Model Under the OAF protocol

We consider the case where cooperation utilizes the symmetric, minimum-delay OAF protocol. In the symmetric OAF protocol [4], the source transmits for a specific time T' (broadcast phase of the round, duration T'), and after that the relays start forwarding the vectors they have received (cooperation phase, duration T'). The source-node remains silent during the second phase. We focus on the minimum-delay version of this protocol, where $T' = n$. Consequently, we consider the overall duration of communication to be $T = 2n$. This means that the source-node commences transmission of a new message at times multiple of T.

Communication takes place in the presence of additive receiver noise, and in the presence of spatially independent and statistically symmetric quasi-static fading. We assume complete knowledge of the fading channel at the receiver of the final destination, and no knowledge of the fading at the receivers of the assisting relays. Each node has a single receive and transmit antenna, operating in half-duplex.

[1]The general case of more than one symbol-slots per time-slot can be handled easily.

a) Transmission Rate: We limit our analysis to the case where the transmission rate is fixed at some R bits per channel use (bpcu), and does not adapt to the occupancy of the queue. This specifically means that for some time duration of T channel uses, from the beginning to the end of the communication between the source-node and the destination, the maximum number of information bits received by the destination (correctly or incorrectly), is RT.

The choice of R and n will be discussed later, and will be the outcome of an optimization consider performance measures for delay and channel errors.

b) Asymptotic Performance Measure for Decoding Errors: Given a choice of R and n, the performance of the decoder at the receiver of the destination is given by the probability $P_{r,e}$ of codeword decoding error and will be described in terms of the DMT. In a network where each user operates at rate R bpcu and at SNR ρ, performance is described in terms of the *diversity gain*

$$d(r) = -\lim_{\rho\to\infty} \log(P_{r,e})/\log(\rho)$$

as a function of the *multiplexing gain*

$$r = R/\log(\rho).$$

For sake of presentation, we use the notation $\doteq$ corresponding to the exponential equality, i.e. $y \doteq \rho^x$ is equivalent to $\lim_{\rho\to\infty} \frac{\log(y)}{\log(\rho)} = x$.

c) Scale of Interest in Transmission: In this asymptotic performance analysis, we allow SNR (hence R) to scale to infinity. Furthermore, in this scale of interest we have that $\rho^\epsilon \log\rho \doteq \rho^\epsilon,\ \epsilon > 0$, allowing us to consider $P_{r,e}$ to be both the probability of codeword error as well as the probability of bit error, both due to decoding (channel). Henceforth, the two probabilities will be considered to be the same.

d) Asymptotic Probability of Channel Error: DMT analysis of the 3-node OAF protocol with large *and equal* durations of the first and second phases, was first provided in [4]. For the $(n+1)$-node ($n \geq 2$) symmetric minimum-delay OAF protocol, the negative SNR exponent of the probability of codeword error, which coincides with the negative SNR exponent of the *probability of bit error*, is given ([5]) by

$$d(n,r) = \begin{cases} n(1-2r) & 0 \leq r \leq \frac{1}{2}, \\ 0 & r > \frac{1}{2}. \end{cases} \quad (1)$$

In addition, the DMT tradeoff for direct transmission (i.e. $n = 1$, no cooperation) is given by

$$d(1,r) = 1 - r, \quad 0 \leq r \leq 1 \quad (2)$$

where the minimum coding time is $T = 1$ ([3]).

B. Source Model

We consider a stationary information source which generates at time-slot t, a random number of packets N_t, where each packet $i,\ i \in [1, N_t]$, has some random size $Y_{i,t}$ (bits).

e) Scaling of Average Arrival Rate: The adopted asymptotic scaling of the transmission rate R asks for a similar scaling of the *average bit arrival rate* R_{in} at the source-node. In accordance with these scales of interest, we scale the average packet arrival rate, $E[N_t]$, with $\log\rho$ while keeping the average packet size, $E[Y_{i,t}]$, fixed.

f) Compound Poisson Arrival Process: We consider the compound Poisson arrival process with exponential distribution of the packet size. More specifically, treating ρ as an index, we consider a compound Poisson bit-arrival process A^ρ that generates (during time-slot t) A_t^ρ bits

$$A_t^\rho = \sum_{i=1}^{N_t^\rho} Y_{i,t},$$

where N_t^ρ is the number of packets that have arrived over the t^{th} time-slot, and $Y_{i,t}$ is the number of bits in packet i. N_t^ρ is drawn from an independent Poisson distribution with mean $\nu = \mu\lambda \log \rho$, and $Y_{i,t}$ is drawn independently from an exponential distribution with mean of $1/\mu$ (independent of ρ) bits per packet. Consequently, the average bit arrival rate R_{in} (bits per time-slot) is

$$R_{\text{in}} = \mathbb{E}\{A_t^\rho\} = \frac{\nu}{\mu} = \lambda \log \rho.$$

g) Performance Measure for Delay Violation: Jointly considering source and channel, we arrive at a first-come-first-serve (FCFS) queuing model. The delay violation probability P_{dv}, is defined as

$$P_{\text{dv}} := P[\text{steady-state end-to-end delay} > D], \tag{3}$$

where the end-to-end delay is the time interval between a bit's arrival to the queue and its departure from the decoder. In summary, we have a single server queuing system with constant service of RT bits at every T time-slots and compound Poisson arrival with rate R_{in} bits per time-slot. Note that we require $\lambda < r$ to ensure queue stability.

III. Problem Analysis

In this section, we derive an analytical expression for the probability of bit loss, due to delay violation and erroneous decoding, as a function of $r, \lambda, \mu, T,$ and D. We distill the tradeoff to describe the negative SNR exponent that is optimized over all coding schemes, constant service rates, and cluster sizes. The tradeoff then describes the probability of having error-producing atypicalities in the arrivals (bursts) and in channel fading (outage).

A. Asymptotic Probability of Delay Violation

For the batch queue-service taking place every T time-slots, we find the asymptotic approximation of the probability of delay violation P_{dv} as follows:

Theorem 1: Assume $D - 2T + 1 > 0$, then the probability of delay violation P_{dv} is given by

$$P_{\text{dv}} \doteq \rho^{-I(r,T,D)} \tag{4}$$

where

$$I(r,T,D) = \min_{t\in\{0,1,2,\ldots\}} (Tt+T_k)\Lambda^*\left(\frac{(n+t-1)rT}{Tt+T_k}\right) \tag{5}$$

for $n = \left\lfloor \frac{D}{T} \right\rfloor$, $k = D - nT$, $T_k = T - 1 - k$ and, for $x > 0$,

$$\Lambda^*(x) = \mu(\sqrt{x} - \sqrt{\lambda})^2. \tag{6}$$

Proof: (sketch) We use the result from [8] which shows that

$$P_{\text{dv}} \doteq \Pr[Q_{T-1-k}^\rho > (n-1)RT]$$

where Q_{T-1-k}^ρ is the steady-state queue length at timeslot $T - k - 1$ relative to the start of service time (commencement of transmission). Given the scaling of packet arrival rate $\nu = \mu\lambda \log \rho$, we choose to consider the compound Poisson arrival A_t^ρ as an aggregation of $M = \log \rho$ independent compound Poisson streams $\left\{A_t^{(j)}, j = 1, \ldots, M\right\}$; each with average packet rate $\mu\lambda$ and average packet size $1/\mu$, i.e.,

$$A_t^\rho = A_t^M = A_t^{(1)} + \cdots + A_t^{(M)}.$$

At this point the *many-flow large-deviations* techniques relating to Cramer's theorem, apply as in [10] (see also [16, Theorem 1.8]), and give that

$$\lim_{M\to\infty} \frac{1}{M} \log P(Q_{T-1-k}^\rho > (n-1)RT) = -I(r,T,D) \tag{7}$$

with $I(r,T,D)$ as in (5). The complete proof can be found in the journal version [9] of this work. □

Consequently, the optimal tradeoff is given by

$$\frac{T^*}{2}(1 - 2r^*)$$

where $r^* = r^*(T^*)$, with $r^*(T)$ (optimal r for any given T) is the solution to

$$\frac{T}{2}(1 - 2r) = I(r,T,D), \tag{8}$$

and with T^* (the optimal T given the optimal r) being the solution to the differential equation

$$\frac{d}{dT}\big(T(1 - 2r^*(T))/2\big) = 0.$$

Remark 1: We emphasize that the result in Theorem 1 holds for any positive D. This is due to scaling of the average number of packets in comparison to the variance of the packet size, and hence due to the subsequent reduction in the overall burstiness. Consequently, for a given system loading, the event of having any queue buildup becomes rare with increase in ρ even for small time D and can thus be approximated using large deviations techniques. This is an improvement over the approximation used in [8], which was based on effective bandwidth techniques (see for example [15]), where we required an asymptotically large value of D.

We note that the closed form expression in (5) and (8) can be calculated, but provides for reduced intuition due to its complexity. We instead use a simple approximation for large D, which is exact for some specific values of D and becomes more accurate as D becomes larger: Following the techniques used to prove Lemma 1.7 in [16], we have that

Lemma 1: For $D \gg \frac{2rT}{\lambda} - 1$,

$$P_{\text{dv}} \approx \rho^{-\mu(r-\lambda)(D-2T+1)}. \tag{9}$$

Proof: See the journal version [9] of this work. □

1) Effective Delay: We have seen that the coding delay reduces the delay-tolerance of the system to a new *effective delay bound*, call it D'. Given (9) and given the minimum-coding duration $T = 2n$, we consider a reduction of the effective delay tolerance, from D to

$$D' := D - 2T + 1 = D - 4n + 1. \tag{10}$$

B. Overall Probability of Bit Error

The total probability of bit error P_{tot} is given by

$$\begin{aligned} P_{\text{tot}} &= P_{r,e} + P_{\text{dv}} - P_{r,e}P_{\text{dv}} \\ &\doteq P_{r,e}(n,r) + P_{\text{dv}}(\lambda, n, r) \\ &\doteq \rho^{-d(n,r)} + \rho^{-\mu(r-\lambda)D'} \end{aligned} \quad (11)$$

where the last equality is the result of (1), (2) and (9). The exponential order of P_{tot} is minimized when the two summands in (11) are of the same exponential order. Following this, we can optimize the cluster size n as well as the multiplexing gain r. We now consider the optimal performance described by

$$\xi(\lambda, \mu, D) := -\lim_{\rho\to\infty} \frac{\log P^*_{\text{tot}}}{\log \rho}, \quad (12)$$

the *optimal SNR exponent* of P^*_{tot}, corresponding to the optimal cluster size and multiplexing gain.

C. Main Result

We define n^* and λ_a as follows:

$$n^* = \max\left\{2, \min\left\{\frac{D+1}{4+2\sqrt{\frac{2}{\mu}}}, \left\lfloor \frac{D}{4} \right\rfloor, N\right\}\right\} \quad (13)$$

and

$$\lambda_a = \begin{cases} 0 & \text{if } n^*(1+\frac{1}{\mu(D-1)}) \le 1 + \frac{2n^*}{\mu(D+1-4n^*)}, \\ \frac{\psi-1}{2\psi-1} & \text{if otherwise,} \end{cases} \quad (14)$$

where

$$\psi = \frac{n^*(1+\frac{1}{\mu(D-1)})}{1+\frac{2n^*}{\mu(D+1-4n^*)}}. \quad (15)$$

Theorem 2: The optimal cluster size $n_{\text{opt}}(\lambda)$ is given by

$$n_{\text{opt}}(\lambda) = \begin{cases} n^* & \text{if } 0 \le \lambda \le \lambda_a, \\ 1 & \text{if } \lambda_a \le \lambda \le 1, \end{cases} \quad (16)$$

and the optimal error exponent in (12) is given by

$$\xi(\lambda, \mu, D) = \begin{cases} n^*(1-2r^*_{n^*}(\lambda)) & \text{if } 0 \le \lambda \le \lambda_a, \\ 1 - r^*_1(\lambda) & \text{if } \lambda_a \le \lambda \le 1, \end{cases} \quad (17)$$

where the optimal multiplexing gain when $0 \le \lambda \le \lambda_a$ is

$$r^*_{n^*}(\lambda) = \frac{n^* + \lambda\mu(D-4n^*+1)}{2n^* + \mu(D-4n^*+1)} \quad (18)$$

and when $\lambda_a \le \lambda \le 1$ is

$$r^*_1(\lambda) = \frac{1+\lambda\mu(D-1)}{1+\mu(D-1)}. \quad (19)$$

Theorem 2 states that for arrival rate $\lambda < \lambda_a$, cooperation applies, and the optimal cluster size is n^*. For $\lambda > \lambda_a$, cooperation is not of interest and it is best that the transmitter transmit directly to the receiver. This is mainly due to the fact that the half-duplex relay protocol cannot achieve full multiplexing gain.

To prove Theorem 2, we observe that the DMT expression from (1)-(2) distinguishes between two regions. The first region corresponds to the case where cooperation takes place ($n \ge 2$), and the second region corresponds to the no-cooperation case ($n = 1$). We begin the proof with the first region.

Case 1: Cooperation Region ($n \ge 2$).

For any arbitrary value of $n \ge 2$, we have from (10) that, $D' = D - 4n + 1$. This leaves us with r and n as the two optimizing parameters. Next, we proceed to find the optimal multiplexing gain given a cluster size and then use this result to find the optimal cluster size (maximizing the error exponent in this cooperation region), as follows:

a) Finding the proper r for a given cluster size: For a given value of $\lambda \in [0, 1/2]$ and cluster size n, we find the transmission multiplexing gain $r^*_n(\lambda)$ that equates the exponents of the two summands in (11). Toward this, we rewrite $d(n,r)$ in (1) as $n(1-2r)$ to get

$$\mu(r-\lambda)D' = d(n,r) = n(1-2r)$$

and the optimal multiplexing gain for the given n as

$$r^*_n(\lambda) := \frac{n + \lambda\mu(D-4n+1)}{2n + \mu(D-4n+1)}, \qquad \lambda \in [0, 1/2]. \quad (20)$$

Note that $\lim_{D\to\infty} r^*_n(\lambda) = \lambda$ as expected.

b) Finding the optimal performance for a given cluster size: We apply the above result for all $\lambda \in [0, 1/2]$, and find the negative SNR exponent,

$$\zeta(n, \lambda) := d(n, r^*_n(\lambda)),$$

of the overall probability of error, for any *fixed* number n of cooperative users. Using (20),

$$\begin{aligned} \zeta(n,\lambda) = d(n, r^*_n(\lambda)) &= n(1-2r^*_n(\lambda)) && (21) \\ &= \underbrace{\frac{1}{1+\frac{2n}{\mu D'}}}_{\text{reduction factor}} n(1-2\lambda). && (22) \end{aligned}$$

Comparing $\zeta(n,\lambda)$ with $d(n,\lambda) = n(1-2\lambda)$ (i.e. the error exponent in the case of no burstiness and/or no delay limitation), we see that the term

$$\frac{d(n, r^*_n(\lambda))}{d(n,\lambda)} = 1/(1+\frac{2n}{\mu D'}) \quad (23)$$

describes the performance degradation factor due to burstiness $1/\mu$ and delay limitation D, specifically for a cluster size equal to n. As expected, $\lim_{D\to\infty} d(n, r^*_n(\lambda)) = d(n, \lambda)$.

c) Finding the Optimal Cluster Size: We want to find the optimal cluster size that maximizes $\zeta(n,\lambda)$ in (22). Replacing $D' = D - 4n + 1$ and solving for n^* which is the solution to $\frac{d}{dn}\zeta(n,\lambda) = 0$, we get the optimal cluster size n^* as given in (13), where the bound $\lfloor \frac{D}{4} \rfloor$ comes from requiring $D-4n+1 \ge 1$. This completes the proof for the cooperation region. Now we move to the no-cooperation region.

Case 2: No-Cooperation Region ($n = 1$).

In this region, $T = 1$ and $D' = D - 2T + 1 = D - 1$. Here, only r is the optimizing parameter. The optimal multiplexing gain $r^*_1(\lambda)$ equates the two summands in (11) and is found by solving

$$\mu(r-\lambda)D' = d(1,r) = 1 - r,$$

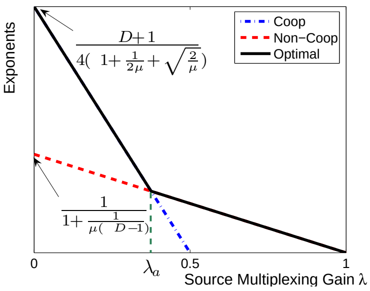

Fig. 2. The exponents for cooperation ($\zeta(n^*, \lambda)$), no-cooperation ($\zeta(1, \lambda)$), and the optimal ($\xi(\lambda, \mu, D)$) cases.

which gives $r_1^*(\lambda)$ as in (19). The negative SNR exponent of the overall probability of error is given by

$$\zeta(1, \lambda) := d(1, r_1^*(\lambda)) = 1 - r_1^*(\lambda) = \frac{1}{1 + \frac{1}{\mu(D-1)}} (1 - \lambda).$$

Now that we have identified the optimal negative SNR exponents under both cooperation and no-cooperation assumptions, we are left to identify the range of the source multiplexing rate λ under which cooperation outperforms no-cooperation. This happens when $\zeta(n^*, \lambda) > \zeta(1, \lambda)$. Since both $\zeta(n^*, \lambda)$ and $\zeta(1, \lambda)$ are linear and decreasing in λ, $\zeta(n^*, 1/2) = 0$, and $\zeta(1, 0) = 0$ (see Figure 2), it is clear that unless $\zeta(n^*, 0) > \zeta(1, 0)$, cooperation is not beneficial. In other words,

$$\zeta(n^*, 0) > \zeta(1, 0) \Rightarrow \exists \lambda_a \in (0, 1/2) : \zeta(n^*, \lambda_a) = \zeta(1, \lambda_a),$$

and

$$n_{\text{opt}}(\lambda) = \begin{cases} n^* & \text{if } 0 \leq \lambda \leq \lambda_a \\ 1 & \text{if } \lambda_a \leq \lambda \leq 1 \end{cases},$$

while

$$\zeta(n^*, 0) \leq \zeta(1, 0) \Rightarrow n_{\text{opt}}(\lambda) = 1, \ 0 \leq \lambda \leq 1.$$

To simplify the presentation in the case of $\zeta(n^*, 0) \leq \zeta(1, 0)$, we fix $\lambda_a = 0$.

What is left is to prove that

$$\zeta(n^*, 0) \leq \zeta(1, 0) \Leftrightarrow n^*(1 + \frac{1}{\mu(D-1)}) \leq 1 + \frac{2n^*}{\mu(D + 1 - 4n^*)},$$

and

$$\lambda_a = \frac{\psi - 1}{2\psi - 1},$$

where ψ is as in (15). The proofs of the above statements simply involve algebraic manipulations and are omitted (see [9] for details). This completes the proof of Theorem 2. □

IV. Conclusion and Future Work

In this paper, we considered transmission of stochastic and delay-limited traffic of a single source, where transmission is assisted by cooperative relay nodes. We derived a tradeoff on the negative SNR exponent of the probability of bit error when errors are due to erroneous decoding as well as due to late arrival of the information. This was achieved for the symmetric and minimum-delay OAF cooperative network with Rayleigh fading, and compound Poisson source. This derivation involved establishing the optimal multiplexing gain and the optimal cooperative cluster size.

This work is a first step toward understanding the benefits of cooperation among users with bursty and delay-sensitive traffic. In this work, we have considered only one source of information along with many relay nodes with no information of their own. In the presence of many delay-sensitive and bursty nodes, the question of optimal selection of source and relays, given queue and channel states, i.e., dynamic cooperation schedule, remains. Extending this work to different bit-arrival processes is also of interest.

Acknowledgment

This research is supported in part by the Center for Wireless Communications, UCSD and UC Discovery Grant No. Com04-10176, ARO-MURI Grant No. W911NF-04-1-0224, NSF CAREER Award No. CNS-0533035, and AFOSR Grant No. FA9550-05-01-0430.

References

[1] A. Sendonaris, E. Erkip and B. Aazhang, "Increasing uplink capacity via user cooperation diversity," *In Proceedings of 1998 IEEE ISIT,* pp. 156, August 1998.
[2] A. Sendonaris, E. Erkip and B. Aazhang, "User cooperation diversity-Part I: System description," *IEEE Trans. Communications*, vol. 51, no. 11, pp. 1927-1938, November 2003.
[3] L. Zheng and D. Tse, "Diversity and Multiplexing: A Fundamental Tradeoff in Multiple-Antenna Channels," *IEEE Trans. Info. Theory*, vol. 49, no. 5, pp. 1073-1096, May 2003.
[4] J. N. Laneman, D. N. C. Tse, and G. W. Wornell, "Cooperative Diversity in Wireless Networks: Efficient Protocols and Outage Behavior," *IEEE Trans. Info. Theory*, vol. 50, no. 12, pp. 3062-3080, Dec. 2004.
[5] P. Elia, K. Vinodh, M. Anand and P. V. Kumar, "D-MG Tradeoff and Optimal Codes for a Class of AF and DF Cooperative Communication Protocols," submitted to *IEEE Trans. Info. Theory,* January 2007, available at: http://arxiv.org/pdf/cs.IT/0611156
[6] S. Yang and J.-C. Belfiore, "Optimal Space-Time Codes for the Amplify-and-Forward Cooperative Channel, " To appear in *IEEE Trans. Inform. Theory*.
[7] P. Elia, K. R. Kumar, S. A. Pawar, P. V. Kumar and H. Lu, "Explicit, minimum-delay space-time codes achieving the diversity-multiplexing gain tradeoff," *IEEE Trans. Info. Theory*, vol. 52, no. 9, Sept. 2006.
[8] P. Elia, S. Kittipiyakul, and T. Javidi, "On the Responsiveness-Diversity-Multiplexing Tradeoff," Proceedings of Workshop on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOPT 2007), April 16-20 2007, Limassol, Cyprus.
[9] P. Elia, S. Kittipiyakul, and T. Javidi, "Cooperative diversity in wireless networks with stochastic and bursty traffic," in preparation.
[10] D. D. Botvich and N. G. Duffield, "Large deviations, the shape of the loss curve, and economies of scale in large multiplexers," *Queuing Systems*, v. 20, pp. 293-320, 1995.
[11] J. Abate, G. L. Choudhury, and W. Whitt, "Exponential approximations for tail probabilities in queues I: waiting times," *Operations Research*, vol. 43, no. 5, 1995.
[12] J. Abate, G. L. Choudhury, and W. Whitt, "Exponential approximations for tail probabilities in queues II: sojourn time and workload," *Operations Research*, vol. 44, no. 5, 1996.
[13] C. Courcoubetis, G. Kesidis, A. Ridder, J. Walrand, and R. Weber, "Admission control and routing in ATM networks using inferences from measured buffer occupancy," *IEEE Transactions on Communications*, vol. 43, pp. 17781784, 1995.
[14] George Kesidis, Jean Walrand, and Cheng-Shang Chang, "Effective bandwidths for multiclass Markov fluids and other ATM sources," *IEEE/ACM Transactions on Networking*, 1:424428, 1993.
[15] F. Kelly, "Notes on effective bandwidths," *Stochastic Networks: Theory and Applications*, vol. 4, 1996.
[16] A. Ganesh, N. O'Connel, and D. Wischik, *Big Queues*, Springer-Verlag, Berlin, 2004.
[17] Y. Viniotis, *Probability and Random Processes for Electrical Engineers*, McGraw-Hill, 1998.
[18] G. H. Hardy, *Orders of infinity,* Cambridge Tracts in Mathematics and mathematical physics, no. 12, 1924.
[19] A. Dembo, O. Zeitouni, *Large Deviations Techniques and Applications*, New York: 2nd edition, Springer-Verlag, 1998.