Cooperation and Cognition in Wireless Networks

Natasha Devroye*, Patrick Mitran*, Oh-Soon Shin†, Hideki Ochiai‡, and Vahid Tarokh*
* Division of Engineering and Applied Sciences
Harvard University, U.S.A.
ndevroye, mitran, vahid @deas.harvard.edu
‡ Department of Electrical and Computer Engineering
Yokohama National University, Japan
hideki@ynu.ac.jp
† Samsung Electronics
Suwon, Korea
osshin@deas.harvard.edu

***Abstract*— The fundamental limits of communication over a wireless network of nodes has recently become an area of intense research. Even in the simplest case, the information theoretic capacity region of multi-terminal networks remains an open problem. The analysis of networks is made even richer through the recent introduction of *cognitive radios,* which could allow for cooperation between network nodes. In this paper we present a decomposition of arbitrary wireless networks with cognitive and non-cognitive nodes, which reduces the network to a set of clusters which behave in *competitive*, *cognitive* and *cooperative* fashions. We survey results on the more classical competitive and cooperative behaviors, and explore two examples of the recently introduced cognitive behavior in more depth. These two examples highlight different aspects of the asymmetric cooperation implied by cognitive behavior, and demonstrate the potential of asymmetric cooperation in wireless networks.**

I. Introduction

In the past 20 years, wireless communication and wireless technology have exploded. Cellular phone base-stations pepper most cities, and are a perfect example of a wireless network with *infrastructure*. Such wireless networks, which allow for a large amount of centralized control, have been extensively studied. On the other hand, groups of mobile devices which spontaneously and opportunistically form *ad-hoc networks* are gaining in popularity and technical feasibility. Ad hoc networks are the subject of extensive and vigorous research. Of the numerous questions which remain to be answered for ad hoc networks, here we limit ourselves to that of cooperative communication at the physical layer. Specifically, we ask the question: how can multiple wireless nodes cooperate in order to communicate?

In this tutorial paper, we will start by showing that any wireless network may be decomposed into components according to the amount of cooperation potentially possible between nodes. In the literature, different models, protocols, and coding schemes suitable to cooperative communications are rapidly being proposed. We hope to provide a framework which is able to classify these various types of proposed cooperative communications. To do so, we present a *wireless decomposition* of networks which indicates the level of cooperation between *transmitters*. We have chosen to group nodes in terms of transmitter cooperation because such cooperation is often more insightful and general than receiver-side cooperation [1], [2]. We show that when given certain information about a network, it may be decomposed into clusters that behave in *competitive, cognitive,* or *cooperative* manners. These three categories may be applied to behavior *within* a cluster, or *between* clusters. We survey results on competitive behavior and cooperative behavior. These two types of behavior have been studied extensively in the literature. We then focus on the less studied *cognitive behavior*.

Cognitive behavior captures the concept of asymmetric cooperation between transmitters. Although this is by no means a new concept in the communications and information theory communities, the asymmetry possibly leads to a more complex problem which has been less studied than its symmetric extremes; competition and cooperation. In this paper, we study two examples of *cognitive behavior*. These two examples highlight different aspects of the cognitive problem: the first considers a case where two users asymmetrically cooperate to send two independent messages to two separate non-cooperating receivers. Here full channel state information at the transmitter (CSIT) is assumed, and cognition is aided, conceptually, by a *genie* who allows for non-causal message knowledge. In the second case we also consider asymmetric cooperation, but here only one message is to be transmitted from a single source to a single destination possibly with the (asymmetric) help of a relay. In this example of cognitive behavior, full CSIT is no longer assumed, and quasi-static, half-duplex systems are considered. Both these examples illustrate some of the issues involved in characterizing what communication is possible over the vast middle-ground between full and no cooperation that we will call cognitive behavior. We note that the term *cognitive behavior* is inspired by the abilities of cognitive radio, but should not be thought of as the behavior of cognitive radios in the spectrum sensing sense. That is, we use the term cognitive behavior to mean asymmetric transmitter cooperation, which would be possible using the cognitive radio technology, rather than a behavior which senses the spectrum looking for available spectral slots.

The paper is structured as follows. In Section II we demonstrate that given three pieces of information about a wireless network at an arbitrary snapshot in time, that we can decompose it into clusters. These clusters may be classified as behaving in *competitive, cooperative,* or *cognitive* manners. In Section III we survey classical results on competitive wireless channels, including the interference and multiple access channels. In Section IV we survey results on cooperative wireless channels, including the broadcast and MIMO channels. In Section V and VI we present two key and different examples of cognitive behavior, as well as outline strategies and the resulting rates achievable in these channels. In Section VII we conclude.

II. Wireless Decomposition

We consider an arbitrary network of wireless devices, which may be cognitive, denoted as (C), or non-cognitive (NC) radios. At any given point in time, certain transmitting nodes (T) have information which they wish to transmit to certain receiving nodes (R). Nodes that do not have any information of their own to transmit are denoted as "extra" nodes (E). We assume that nodes are not able to simultaneously transmit and receive, i.e., they must obey the half-duplex constraint. This is a reasonable assumption given current technology. Thus, a node is classified as either a (T), (R) or (E) node, but never more than one, and as either cognitive (C) or non-cognitive (NC).

If all devices simultaneously transmit, the network may suffer from interference. However, we wish to exploit the nature of cognitive radios to reduce this interference. The key to doing so is transmitter cooperation, which could lead to interference mitigation. At each point in time, depending on the device capabilities, as well as the geometry and channel gains between the various nodes, certain cognitive nodes may be able to hear and/or obtain the messages to be transmitted by other nodes. In reality, these messages would need to be obtained in real time, and could exploit the geometric gains between cooperating transmitters relative to receivers in a, for example, 2 phase protocol [3]. However, as a first step, we idealize the concept of message knowledge: whenever a (T) or (E) node is cognitive and in principle able to hear and decode the message of another transmitting node, we assume it

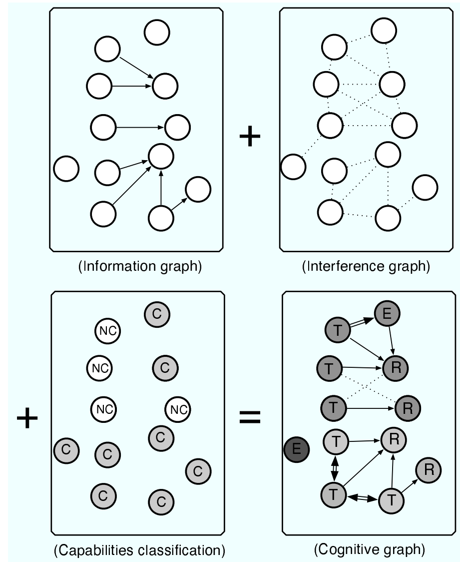

Fig. 1. The information and interference graphs, together with the capabilities classification yield the cognitive decomposition graph.

has full *a-priori* knowledge. We call this the *genie assumption*, as these messages could have been given to the appropriate transmitters by a genie. Notice that we explicitly allow for asymmetric message knowledge, and that this message knowledge is between potentially transmitting nodes only. We ignore cognitive receiving nodes for now. In this decomposition theorem, all transmitter cooperation occurs under the genie assumption. Protocols which remove this assumption are discussed in [3].

We now demonstrate that given a snapshot of a network and three pieces of information: an *information graph*, an *interference graph* and a *capabilities classification* as in Fig.1, transmission scenarios in which there is some form of transmitter cooperation are captured in a *cognitive graph*: a set of disjoint non-interfering *groups* of nodes, each of which consists of a set of *clusters* behaving in an inter/intra cluster competitive, cognitive, or cooperative manner.

The information graph: This directed graph captures which nodes have independent information to be sent to which receivers at a given moment in time.

The interference graph: This undirected graph captures the interference in a network. If two nodes can hear each other, and thus potentially interfere with each other, then an edge exists between them. Notice that for a (T) node to be able to transmit to an (R) node, an edge in the interference graph should appear between them.

The capabilities classification: This partition of the nodes then labels them as cognitive (C) or non-cognitive (NC). A node is (C) when it is able and willing to sense and adapt to its environment. Note that an (NC) node could model either a wireless device that does not have cognitive capabilities, or could alternately model devices that do not require cognition to communicate. For example, in vertical spectrum sharing, the (possibly paying) primary users are guaranteed spectrum access; secondary users must avoid interfering with these primary users, so primary user cognition may not be necessary for transmission. While receivers can be (C) or (NC) in our formulation, this has no impact, as we do not allow for receiver cooperation in our current model.

Cognitive graph: From the information graph, interference graph, and capabilities classification, we can form a cognitive graph as done in [4].

A. Cognitive Graph Decomposition

The cognitive graph gives us information on the interference seen, and the transmitter cooperation that is possible. We assume all (T), (R) and (E) nodes have full channel knowledge. This assumption is used to simplify and idealize the problem, and will provide an upper bound to any real world scenario.

In order to fully describe all transmitter cooperation strategies in a wireless network employing cognitive radios as described by the cognitive graph, the following notions are needed. A *group* is a set of connected nodes (ignoring the direction of arcs). It is easy to see that a cognitive graph may be partitioned into groups, and that, by construction, these groups will not interfere with each other. They may independently encode their messages and simultaneously transmit with no interference. Thus, it is of interest to calculate the capacity region of each group. Within a group, we may further divide the nodes into *clusters*. A *cluster* is defined as a set of nodes connected only through solid arcs to a single receiver. We assume all receivers are independent and unable to cooperate. Thus, there exists one cluster per receiver.

<u>Intra-Cluster behavior:</u> within a single cluster, we may partition transmitter cooperation into three classes:

- *Competitive:* all (T) within a cluster encode their messages independently. They compete for the channel. If there are no arcs between any of the (T) and (E) nodes within a cluster, that cluster behaves competitively.
- *Cooperative:* all the (T) / (E) in a cluster know the messages of all the other (T) in that cluster a priori. These require bi-directional cognitive (double) arcs between all (T) nodes of that cluster. A cluster consisting of a single transmitter is said to be cooperative.
- *Cognitive:* all clusters that are not competitive or cooperative, i.e., some but not all of the (T) / (E) in a cluster know the messages to be transmitted by other (T) in the cluster *a-priori* (solid double arcs). This is an asymmetric form of cooperation, which may allow the user with the message knowledge to mitigate interference, or aid in the transmission of the *a-priori* known messages.

<u>Inter-cluster behavior:</u> when two (or more) clusters within one group are connected through undesired interference (dotted) edges or share (T) / (E) nodes, we can speak of inter-cluster behavior.

- *Competitive:* when all (T) / (E) nodes of one cluster are independent of all (T) / (E) nodes of another cluster, the *clusters* compete for the channel during simultaneous transmission. Note that competitive inter-cluster behavior does not imply anything about the competitive, cooperative, or cognitive behavior of nodes within one cluster. The clusters will be linked through interference (dotted) edges.
- *Cooperative:* all the (T) / (E) nodes in one cluster know the messages of a second cluster and vice-versa. Clusters under consideration know each others' messages and so the clusters can cooperate, at the cluster level, to transmit their messages, potentially reducing interference.
- *Cognitive:* encompasses all clusters that do not behave competitively or cooperatively, that is, when a subset of the (T) nodes in one cluster knows the messages to be transmitted by a subset of the (T) nodes of the other clusters, we call this inter-cluster cognitive behavior. The cluster with the message knowledge may be able to at least partially mitigate some interference from the other cluster(s).

Note that if nodes $(X) \Leftrightarrow (Y)$ and $(Y) \Leftrightarrow (Z)$ (where $\Leftrightarrow$ indicates two-way cognition, or cooperation) then one may suppose $(X) \Leftrightarrow (Z)$. This only makes sense if there is no overhead to cognition and all message knowledge is assumed to be non-causal and instantaneous. This transitivity property may break down once messages must be causally obtained, and our model does not enforce such transitivity of cognition. We have the following theorem, which follows directly from the construction and definitions above.

Theorem 1: At a point in time, if given information and interference graphs as well as a capabilities classification, we may construct a cognitive graph which identifies the non-interfering *groups*, and the interfering *clusters* within each group. All forms of user cooperation within a cluster is described as *competitive, cognitive,* or *cooperative* behaviors. Furthermore, between clusters in the same group, we may have *competitive, cognitive,* or *cooperative* behavior.

We demonstrate this decomposition by example and construct the cognitive graph from the given information, interference and capabilities graphs, and indicate the groups, clusters, and their inter and intra-cluster behaviors in Fig.2.

III. Competitive Behavior

Competitive behavior in wireless networks has been considered in the communications literature, although not formally under this name. A cluster behaves in an *intra-cluster competitive manner*, as defined in Section II when all the transmitters in that cluster wish to transmit to a particular receiver (possibly simultaneously), and must do so independently.

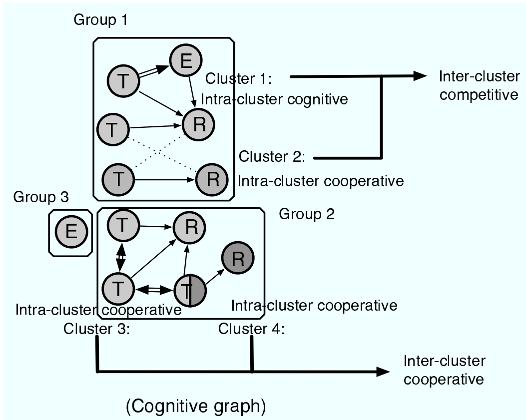

Fig. 2. The resulting groups, clusters, and their behaviors.

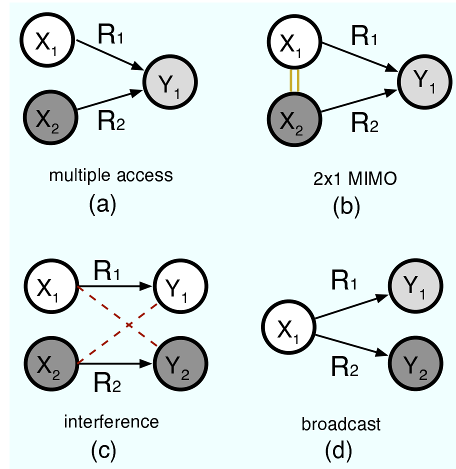

Fig. 3. Four classical multi-user information theoretic channels. The multiple access (a) is an example of intra-cluster competitive behavior, (b) is the 2×1 MIMO channel, which is an example of intra-cluster cooperative behavior, (c) is an example of inter-cluster competitive behavior, and the broadcast channel (d) is an example of inter-cluster cooperative behavior.

The *intra-cluster competitive behavior* corresponds to the classical information theoretic multiple access channel (MAC) [5], [6] , shown in Fig.3(a). This channel is a two transmitter, single receiver channel where the two transmitters have independent messages and cannot cooperate. For definitions of a code, an achievable rate, an achievable region, and the capacity region we refer the interested reader to Chapters 8 and 14 of [6]. The capacity region of this channel is a well-known result in information theory, and is given by the convex hull of the union over all distributions of the input random variables (X_1, X_2) of the form $p(x_1, x_2) = p(x_1)p(x_2)$ of all rate pairs (R_1, R_2) satisfying

$$R_1 \leq I(X_1; Y_1|X_2) \quad (1)$$

$$R_2 \leq I(X_2; Y_1|X_1) \quad (2)$$

$$R_1 + R_2 \leq I(X_1, X_2; Y_1) \quad (3)$$

In the Gaussian noise channel often considered, this rate region becomes a convex pentagon in the positive quadrant with two sides along the axes.

Consider now an example of *inter-cluster* competitive behavior: the two sender, two receiver interference channel, as shown in Fig.3(c). Relating it to our decomposition theorem, our two clusters consist of single-sender single receiver point-to-point channels. They are connected however through the interference graph, and so the two receivers experience interference from the other channel's transmitter. The two transmitters are independent, and so they must compete for the channel. The interference channel, introduced in [5], [7] was consequently studied by, among many others [8], [9], [10], [11], [12]. Although the capacity region of this channel is known for a few cases, its capacity in its most general setting remains an open problem.

IV. Cooperative Behavior

We now turn to two examples of cooperative behavior. For the case of *intra-cluster cooperative behavior* the simplest example is a two sender, single receiver channel in which the two senders may cooperate. This can be viewed as a 2×1 Multiple Input Multiple Output (MIMO) channel [13], [14], as shown in Fig.3(b), where the double line between transmitters indicates full transmitter cooperation. This is an example of a more general $M \times N$ MIMO channel, which consists of a single sender, single receiver wireless channel in which the transmitter and receiver have M and N antennas respectively.

The use of multiple antennas equips a system with spatial diversity which helps to combat fading and can thus improve the rate (bits/second) or quality (probability of error) of the link. If, conceptually, each antenna is viewed as a transmitting node, a MIMO system behaves in an intra-cluster cooperative manner. That is, all receive antennas, which may fully cooperate, may be viewed as a single receiver, and all the transmitters may fully cooperate, thus satisfying our definition of *intra-cluster* cooperative behavior. The transmitters make use of this perfect cooperation by taking advantage of the spatial dimension in wireless channels, that is, spatial separation in fading channels often provides nearly independent fading on all links. Intuitively then, it is less likely to see bad fades on all links. Furthermore, one can judiciously combine the signals sent and received on the various transmit/receive antennas so as to maximize the average combined signal level, and minimize the received interference.

We now look at a classical result on MIMO channels in additive white Gaussian noise. Assume the channel inputs and output are related by the vector equation

$$\mathbf{Y} = \mathbf{H}\mathbf{X} + \mathbf{N}$$

where $\mathbf{Y} \in \mathbb{C}^{N \times 1}$, $\mathbf{X} \in \mathbb{C}^{M \times 1}$, $\mathbf{N} \in \mathbb{C}^{N \times 1}$ is symmetric complex additive white Gaussian noise with iid components, and $\mathbf{H} \in \mathbb{C}^{N \times M}$ is the channel matrix with independent identically distributed elements h_{ij} which describes the channel between the i-th transmit and the j-th receive antenna. Without loss of generality, assume that all transmit antennas have unit average power. Likewise, the receive antennas have unit power Gaussian noise. If this is not the case, the $\mathbf{H}$ matrix may be appropriately scaled row-wise and column-wise. Under the above unit transmit power per transmit antenna and unit noise power per receive antenna constraint, it is well known that a MIMO system with Gaussian codebook and with rate R bits/channel use can reliably communicate over any channel with transfer matrix $\mathbf{H}$ such that $R < \log_2 \det(I + \mathbf{H}\mathbf{H}^\dagger) \triangleq C(\mathbf{H})$ [14], [13], where I denotes the identity matrix and $\mathbf{H}^\dagger$ is the conjugate transpose of $\mathbf{H}$.

The simplest example of *inter-cluster* cooperation is the information theoretic broadcast channel, shown in Fig.3(d). In its simplest form, the broadcast channel [15] consists of one transmitter sending two independent messages to two non-cooperating receivers. This can be viewed as inter-cluster cooperation if we think of the two links $(X_1 \rightarrow Y_1)$ and $(X_1 \rightarrow Y_2)$ as the two clusters, where full cooperation between transmitters is achieved by the fact that transmitter node X_1 may cooperate with itself. The capacity region of the general broadcast channel is still unknown, save for certain cases [16], [6], [17], [18]. Of course, many achievable regions have been developed; the largest to date was computed in [19]. In Gaussian noise, much more can be said about the broadcast channel. The capacity region of a broadcast channel with single antennas coincides with the region of a degraded broadcast channel and is a classical result, outlined in Section 14.6.3 of [6]. In contrast, the capacity region of the Gaussian MIMO broadcast channel was recently found in the most non-trivial result of [20]. The gain of dirty-paper coding over time division multiple access (TDMA, the users time share the channel) when broadcasting information from a single base station to multiple users is explored in [21]. There, the authors find the sum-rate of broadcasting using the optimal dirty-paper coding strategy is at most min(number transmit antennas, number receivers) that of TDMA. We thus see that cooperation is a relatively well understood and beneficial strategy.

V. Cognitive Behavior: The Cognitive Radio Channel

We now turn to the behavior which spans the gap, or in some sense interpolates between cooperative and competitive behavior, that will undoubtedly form an important part of future wireless and cognitive radio channels: that of asymmetric channel knowledge and co-operation. In this Section we will consider one example of cognitive behavior: a two sender, two receiver (thus with two independent messages) interference channel with asymmetric and a priori message knowledge at one of the transmitters. Certain asymmetric (in transmitter cooperation) channels have been considered in the literature: for example in [22], among other results, the capacity region of a multiple access channel with asymmetric cooperation between the two transmitters is computed. The authors in [23] consider a channel which could involve asymmetric transmitter cooperation, through the use of a *conference*. In [24] and [25] we introduced the *cognitive radio channel*, which captures the most basic form of asymmetric transmitter cooperation for the interference channel. We now study the information theoretic limits of interference channels with asymmetric transmitter cooperation, or *cognitive radio channels*. To this end, we compare the achievable region for the cognitive radio channel of [25] to inner and outer bounds on the region.

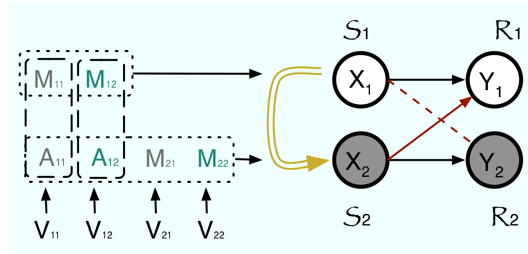

Fig. 4. The modified cognitive radio channel with auxiliary random variables M_{11}, M_{12} and M_{21}, M_{22}, inputs X_1 and X_2, and outputs Y_1 and Y_2. The auxiliary random variable A_{11}, A_{12} associated with $\mathcal{S}_2$, aids in the transmission of M_{11} and M_{12} respectively. The vectors V_{11}, V_{12}, V_{21} and V_{22} denote the effective random variables encoding the transmission of the private and public messages.

We define a 2×2 *genie-aided cognitive radio channel* C_{COG}, as in Fig.4, to be two point-to-point channels $\mathcal{S}_1 \rightarrow \mathcal{R}_1$ and $\mathcal{S}_2 \rightarrow \mathcal{R}_2$ in which the sender $\mathcal{S}_2$ is given, in a non-causal manner (i.e., by a genie), the message X_1 which the sender $\mathcal{S}_1$ will transmit. Let X_1 and X_2 be the inputs to the channel, and let Y_1 and Y_2 be the outputs of the channel. The conditional probabilities of the discrete memoryless C_{COG} are fully described by $P(y_1|x_1,x_2)$ and $P(y_2|x_1,x_2)$.

In [12], an achievable region for the interference channel is found by first considering a modified problem and then establishing a correspondence between the achievable rates of the modified and the original channel models. We follow a similar approach.

The channel C^m_{COG}, defined as in Fig.4 introduces many new auxiliary random variables, whose purposes can be made intuitively clear by relating them to auxiliary random variables in previously studied channels. They are defined and described in Table I. Standard definitions of achievable rates and regions are employed [6], [24]. Then an achievable region for the 2×2 cognitive radio channel is given by:

Theorem 2: Let $Z \triangleq (Y_1, Y_2, X_1, X_2, V_{11}, V_{12}, V_{21}, V_{22}, W)$, be as shown in Fig.4. Let $\mathcal{P}$ be the set of distributions on Z that can be decomposed into the form

TABLE I

DESCRIPTION OF RANDOM VARIABLES IN THEOREM 2.

(Random) variable names	(Random) variable descriptions
M_{11}, M_{22}	Private info from $\mathcal{S}_1 \to \mathcal{R}_1$ and $\mathcal{S}_2 \to \mathcal{R}_2$ resp.
M_{12}, M_{21}	Public info from $\mathcal{S}_1 \to (\mathcal{R}_1, \mathcal{R}_2)$ and $\mathcal{S}_2 \to (\mathcal{R}_1, \mathcal{R}_2)$ resp.
A_{11}, A_{12}	Variables at $\mathcal{S}_2$ that aid in transmitting M_{11}, M_{12} resp.
$V_{11} = (M_{11}, A_{11}), V_{12} = (M_{12}, A_{12})$	Vector helping transmit the private/public (resp.) info of $\mathcal{S}_1$
$V_{21} = M_{21}, V_{22} = M_{22}$	Public, private message of $\mathcal{S}_2$.
	Also the auxiliary random variables for Gel'fand-Pinsker coding
W	Time-sharing random variable, independent of messages

$$
\begin{aligned}
&P(w) \times [P(m_{11}|w)P(m_{12}|w)P(x_1|m_{11}, m_{12}, w)] \\
&\times [P(a_{11}|m_{11}, w)P(a_{12}|m_{12}, w)] \\
&\times [P(m_{21}|v_{11}, v_{12}, w)P(m_{22}|v_{11}, v_{12}, w)] \\
&\times [P(x_2|m_{21}, m_{22}, a_{11}, a_{12}, w)]\, P(y_1|x_1, x_2)P(y_2|x_1, x_2),
\end{aligned} \tag{4}
$$

where $P(y_1|x_1, x_2)$ and $P(y_2|x_1, x_2)$ are fixed by the channel. Let $T_1 \stackrel{\triangle}{=} \{11, 12, 21\}$ and $T_2 \stackrel{\triangle}{=} \{12, 21, 22\}$. For any $Z \in \mathcal{P}$, let $S(Z)$ be the set of all tuples $(R_{11}, R_{12}, R_{21}, R_{22})$ of non-negative real numbers such that there exist non-negative reals $L_{11}, L_{12}, L_{21}, L_{22}$ satisfying:

$$
\begin{aligned}
\bigcap_{T \subset \{11,12\}} \left(\sum_{t \in T} R_t \right) &\leq I(X_1; \mathbf{M}_T | \mathbf{M}_{\overline{T}}) && (5)\\
R_{11} &= L_{11} && (6)\\
R_{12} &= L_{12} && (7)\\
R_{21} &\leq L_{21} - I(V_{21}; V_{11}, V_{12}) && (8)\\
R_{22} &\leq L_{22} - I(V_{22}; V_{11}, V_{12}) && (9)\\
\bigcap_{T \subset T_1} \left(\sum_{t_1 \in T} L_{t_1} \right) &\leq I(Y_1, \mathbf{V}_{\overline{T}}; \mathbf{V}_T | W) && (10)\\
\bigcap_{T \subset T_2} \left(\sum_{t_2 \in T} L_{t_2} \right) &\leq I(Y_2, \mathbf{V}_{\overline{T}}; \mathbf{V}_T | W), && (11)
\end{aligned}
$$

$\overline{T}$ denotes the complement of the subset T with respect to T_1 in (10), with respect to T_2 in (11), and $\mathbf{V}_T$ denotes the vector of V_i such that $i \in T$. Let S be the closure of $\cup_{Z \in \mathcal{P}} S(Z)$. Then any pair $(R_{11} + R_{12}, R_{21} + R_{22})$ for which $(R_{11}, R_{12}, R_{21}, R_{22}) \in S$ is achievable for C_{COG}. $\square$

Proof outline: The main intuition is as follows: the equations in (5) ensure that when $\mathcal{S}_2$ is presented with X_1 by the genie, the auxiliary variables M_{11} and M_{12} can be recovered. Eqs. (10) and (11) correspond to the equations for two overlapping MAC channels seen between the effective random variables $\mathbf{V}_{T_1} \to \mathcal{R}_1$, and $\mathbf{V}_{T_2} \to \mathcal{R}_2$. Eqs. (8) and (9) are necessary for the Gel'fand-Pinsker [26] coding scheme to work ($I(V_{21}; V_{11}, V_{12})$ and $I(V_{22}; V_{11}, V_{12})$ are the penalties for using non-causal side information). Intuitively, the sender $\mathcal{S}_2$ could aid in transmitting the message of $\mathcal{S}_1$ (captured by the A_* random variables) or it could dirty paper code against the interference it will see (captured by the M_{2*} variables). We smoothly interpolate between these two options.

A. Achievable rates for Cognition in Gaussian noise

Consider the 2×2 genie-aided cognitive radio channel described by the input, noise and output relations:

$$
\begin{aligned}
Y_1 &= X_1 + a_{21} X_2 + Z_1 \\
Y_2 &= a_{12} X_1 + X_2 + Z_2
\end{aligned}
$$

where a_{12}, a_{21} are the crossover (channel) coefficients, $Z_1 \sim \mathcal{N}(0, Q_1)$ and $Z_2 \sim \mathcal{N}(0, Q_2)$ are independent AWGN terms, X_1 and X_2 are constrained to average powers P_1 and P_2 respectively, and $\mathcal{S}_2$ is given X_1 non-causally. In order to determine an achievable region for the modified Gaussian genie-aided cognitive radio channel, specific forms of the random variables described in Theorem 2 are assumed, and are analogous to the assumptions found in [25].

The resulting achievable region, in the presence of additive white Gaussian noise for the case of identical transmitter powers ($P_1 = P_2$) and identical receiver noise powers ($Q_1 = Q_2$), is presented in Fig.5. The ratio of transmit power to receiver noise power is 7.78 dB. The cross-over parameters in the interference channel are $a_{12} = a_{21} = 0.55$.

In the figure, we see 4 regions. The time-sharing region (1) displays the result of pure time sharing of the wireless channel between users X_1 and X_2. Points in this region are obtained by letting X_1 transmit for a fraction of the time, during which X_2 refrains, and vice versa. The interference channel region (2) corresponds to the best known achievable region [12] of the classical information theoretic interference channel. In this region, both senders encode independently, and there is no message *a-priori* knowledge by either transmitter of the other's message. The cognitive channel region (3) is the achievable region described here and in [25]. In this case X_2 received the message of X_1 non-causally from a genie, and X_2 uses a coding scheme which combines interference mitigation with relaying the message of X_1. We see that both users – not only the incumbent X_2 which has the extra message knowledge – benefit from using this scheme. This is as expected, as the selfish strategy boosts R_2 rates, while the selfless one boosts R_1 rates, and so gracefully combining the two will yield benefits to both users. Thus, the presence of the incumbent cognitive radio X_2 can be beneficial to X_1, a point which is of practical significance. This could provide yet another incentive for the introduction of such schemes.

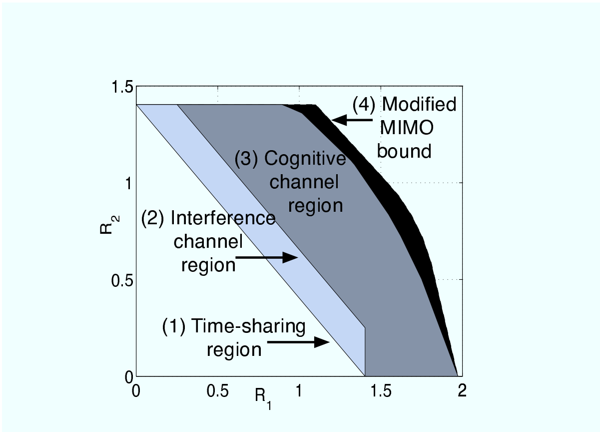

Fig. 5. Rate regions (R_1, R_2) for 2×2 wireless channels.

The modified MIMO bound region (4) is an outer bound on the capacity of this channel. It consists of the $2{\times}2$ MIMO Gaussian Broadcast Channel capacity region [20], where we have restricted the form of the transmit covariance matrix to be of the form $\begin{pmatrix} P_1 & c \\ c & P_2 \end{pmatrix}$, to more closely resemble our constraints. That region is then intersected with the capacity bound on $R_2 \leq I(Y_2; X_2|X_1)$ for the channel for $X_2 \rightarrow Y_2$ in the absence of interference from X_1 to yield the outer bound (4) of Fig.5.

After the *cognitive radio channel* was introduced, [27] presented a slightly different formulation of the problem for which not only an achievable but a capacity region could be computed. They first consider the set of rate pairs that the primary user and secondary user can achieve assuming that no change in the encoding and decoding of the primary user is made, that is to say, the primary transmitter and receiver are unaware of the cognitive user. This however has no effect on the achievable regions as compared to the formulation presented here. The capacity region thus derived is given in Theorem, 3.1 of [27], repeated below in our notation.

Theorem 3: The capacity of the *Gaussian* cognitive radio channel with direct coefficients 1 and cross-over parameters a_{12}, a_{21} is given by

$$\begin{aligned} R_1 &= \tfrac{1}{2}\log_2\left(1 + \tfrac{P_1}{N}\right) \\ R_2 &= \tfrac{1}{2}\log_2\left(1 + (1-\alpha^*)P_2\right) \end{aligned}$$

as long as $a_{21} < 1$,and α^* is

$$\alpha^* = \left(\frac{\sqrt{P_1}\left(\sqrt{1 + a_{21}^2 P_2 (1+P_1)} - 1\right)}{a_{21}\sqrt{P_2}(1+P_1)} \right)^{\frac{1}{2}}.$$

VI. Cognitive behavior: Collaborative Communications

We now consider a second example of cognitive behavior where rather than having two independent messages to be transmitted, there is only one message to be sent from given source to a given destination, possibly with the help of a relay. This relay help can be considered as asymmetric transmitter cooperation, and falls into the category of intra-cluster cognitive behavior. We first survey some relay channel results before moving onto the case considered in [28], which has removed many of the classical, and somewhat unrealistic constraints,

A. The relay channel

The relay channel, which in its simplest and most classical form is a three-terminal channel with one source, one relay (without its own information to transmit) and one destination, is another example of *cognitive behavior*. Relay channels were introduced by van der Meulen [22], and various variations of the problem were later studied by others [29], [30]. The current state of the art is well summarized in the recent [31]. The classical relay channel, like the cognitive radio channel, assumes full channel-state information at all terminals. The paper [30] introduced two fundamental coding schemes for the relay channel often called Decode-and-Forward (DF, Theorem 1) and Compress-and-Forward (CF, Theorem 6). This formulation may be extended to multiple relays, as done in [29], [32] and improved in [33], [34]. We defer to the very informative and insightful [31] for further information on relay channels.

Three major issues are ignored in the classical relay channel framework. The first is the *duplex* constraint. Most of the results on relay channel assume full-duplex relays, that is, relays which may receive and transmit simultaneously. In realistic wireless channel, this assumption does not hold, since the intensity of the near-field of the transmitted signal is much higher than that of the far field of the received signal. In essence, a full-duplex relay would start to interfere with itself. Cognitive relay schemes which operate under a *half-duplex* constraint, that is, where a node cannot simultaneously transmit and receive data, must be considered.

The second assumption often made in the context of wireless communications is the quasi-static fading model. That is, traditionally many authors assume that the fading coefficients remain fixed for the entire duration of the transmission frame. In order to drive the probability of error to zero, in the information-theoretic sense, frame lengths must often tend to infinity. Thus, all realizations of a channel are not experienced, and ergodic capacity results seem limited in their applicability. This, in addition to the fact that the channel state is often not known to the transmitters but only to receivers motivates the study of more realistic *compound* channels [35], [36].

Finally, while the degraded relay channel has been completely solved [30], [37], in wireless systems most noise is due to thermal noise in the receiver frontend. While it may be reasonable to assume that the relay has a better signal to noise ratio (SNR) than the ultimate receiver, it is unrealistic to assume that the receiver is a degraded version of the relay.

These three drawbacks of traditional approaches to the relay channel motivate the study of non-degraded compound relay channels which satisfy the half-duplex constraint. In [28], the authors investigate a bandwidth efficient decode and forward approach that does not employ predetermined phase durations or orthogonal sub-channels to resolve the half-duplex constraint: each relay determines based on its own receive channel when to listen and when to transmit. Furthermore, the transmitters are not aware of the channel and no assumption of degradedness are made: the noise at the relays is independent of that at the destination. Also, as opposed to previous relay and collaborative literature, the results still hold under a bounded asynchronous model. Finally, in the case of multiple relays assisting the source, the approach permits one relay to assist another in receiving the message, a feature not present in much of the early work on communications over compound channels. However, more recent work along this line may be found in [38], [39].

B. *Collaborative Communications*

We now present a brief summary of this important and alternate view of the compound relay channel [28], which is a prime example of cognitive behavior in a network where the cognitive nodes do not have information of their own to send, and where the channel is compound (or quasi-static, non-ergodic). The authors of [28] use the term *collaborative communications* to describe their category of work. This falls into the category of cognitive behavior in the setting considered here.

The use of the spatial dimension is known to greatly increase the reliability of such channels. Spatial diversity is the term often used to capture the potential gain (reliability in this case) of independent paths between sources and destinations which result from spatial separation of nodes or antennas. Of primary interest then is to determine if one can achieve the genie bound on diversity: the diversity gain that would be achieved if all the transmit antennas of the source and relay nodes were in fact connected to a single node (in [40], [41] this is referred to as the transmit diversity bound).

For example, suppose we consider the three transmit collaborators and one receiver node scenario (each equipped with a single antenna) as illustrated in Fig.6. If all the collaborators were aware of the message *a priori*, we could in principle achieve the ideal performance of a 3×1 space-time system between the transmit cluster and the receiver node. However, only the source node in the transmit cluster is aware of the message *a priori*. The other two nodes in the cluster must serve as relays and are not aware of the message *a priori*. There will be a loss in performance (as measured by the probability of outage) compared to the idealized 3×1 space-time system. In particular, we shall be interested in determining sufficient conditions on the geometry and

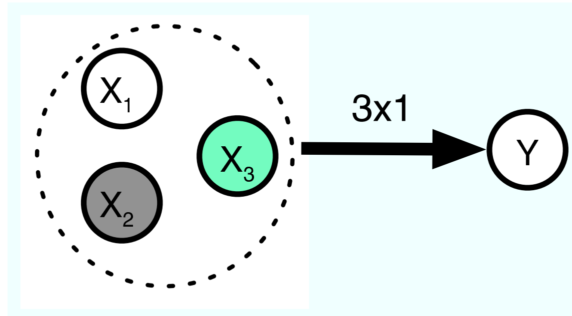

Fig. 6. Is an ideal 3×1 space-time gain achievable with three separate transmit nodes and one receive node?

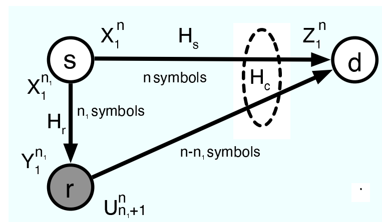

Fig. 7. The collaborative communications problem for two transmit collaborators and one receiver.

signal path loss of the transmitting cluster for which performance close to the genie bound can be guaranteed.

To determine an upperbound on this loss, the authors [28] derive a novel approach to the compound relay channel. This approach is best summarized as follows. In a traditional compound channel, a set of possible channel realizations are given and one seeks to prove the existence of a code (with maximal rate) which is simultaneously good on all channel realizations. In [28], the problem is framed in the opposite direction. They fix a rate and ask how large the set of compound channels can be made while guaranteeing that the code is still good.

Consider three nodes denoted as source (s), relay (r) and destination (d) as illustrated in Fig.7 and each equipped with N_s, N_r and N_d antennas respectively (the results readily generalize to multiple relay nodes).

We assume that while listening to the channel, the relay may not transmit. Hence, the communications protocol we propose is as follows. The source node wishes to transmit one of 2^{nR} messages to the destination employing n channel uses. While not transmitting, the relay node listens. Due to the relay node's proximity to the source, after n_1 samples from the channel (a number which the relay determines on its own and for which the source has no knowledge), it may correctly decode the message. After decoding the message, it then proceeds to transmit for the remaining $n - n_1$ transmissions in an effort to improve the reception of the message at the destination. The destination is assumed to be made aware of n_1 before attempting to decode the message. This may be achieved by an explicit low-rate transmission from the relay to the destination. Alternatively, if the value of n_1

is constrained to some integer multiple of a fundamental period n_0 (say $n_0 \sim \sqrt{n}$), then the destination may estimate n_1 accurately using power detection methods. We denote the first phase of the n_1 transmissions as the *listening* phase while the last $n - n_1$ transmissions as the *collaboration* phase.

We assume that all channels are modeled as additive white Gaussian noise (AWGN) with quasi-static fading. In particular, X and U are column vectors representing the transmission from the source and relay nodes respectively and we denote by Y and Z the received messages at the relay and destination respectively. Then during the listening phase we have that

$$Z = H_s X + N_Z \tag{12}$$
$$Y = H_r X + N_Y, \tag{13}$$

where the N_Z and N_Y are column vectors of statistically independent complex AWGN with variance 1/2 per row per dimension, H_s is the the fading matrix between the source and destination nodes and likewise, H_r is the fading matrix between the source and relay nodes. During the collaboration phase, we have that

$$Z = H_c [X^T, U^T]^T + N_Z, \tag{14}$$

where H_c is a channel matrix that contains H_s as a submatrix (see Fig.7).

We further assume that the source has no knowledge of the H_r and H_c matrices (and hence the H_s matrix too). Similarly, the relay has no knowledge of H_c but is assumed to know H_r. Finally, the destination knows H_c.

Without loss of generality, we will assume that all transmit antennas have unit average power during their respective transmission phases. Likewise, the receive antennas have unit power Gaussian noise. If this is not the case, the respective H matrices may be appropriately scaled row-wise and column-wise.

Intuition for the above problem then suggests the following. During the listening phase, the relay knowing H_r listens for an amount of time n_1 such that $nR < n_1 C(H_r)$. During this time, the relay receives at least nR bits of information and may reliably decode the message. The destination, on the other hand, receives information at the rate of $C(H_s)$ bits/channel use during the listening phase and at the rate of $C(H_c)$ bits/channel use during the collaborative phase. It may reliably decode the message provided that $nR < n_1 C(H_s) + (n - n_1) C(H_c)$. In the limit as $n \to \infty$, the ratio n_1/n approaches a fraction f and we may conjecture that there exists a "good" code of rate R for the set of channels (H_r, H_c) which satisfy

$$R \le f C(H_s) + (1-f) C(H_c) \tag{15}$$
$$R \le f C(H_r), \tag{16}$$

for some $f \in [0, 1]$. We note that if the channel between the source and the relay is particularly poor, we may fall back on the traditional point-to-point communications paradigm and add the following region to that given in (15) – (16)

$$R \le C(H_s). \tag{17}$$

The above intuition is not a proof of achievability but it does provide an upper bound on the performance of the protocol. The essential difficulty in proving that there exists a code which is "good" for any such pair of channels (H_r, H_c) is two-fold. The problem we are dealing with is a *relay channel* which is also a *compound channel*: we seek to prove the existence of a code which performs well over an entire set of channels (unknown to the transmitters). The key will be to show the existence of a code that may essentially be refined. Regardless of the actual value of n_1, there exists a codebook for the source which, starting at time $n_1 + 1$, may be layered with the transmission of the relay and perform just as well as if the value of n_1 had been known to the source. For a formal statement and proof of these results, we defer to [28].

The authors simulated the outage probability of their scheme under a quasi-static Rayleigh fading assumption. These numerical and simulation results showed that if the intra-cluster communication has a 10 dB path loss advantage over the receiver at the destination node, in most cases there is essentially no penalty for the intra-cluster communication. Physically, in a two collaborator scenario with square path loss, this corresponds to a transmit cluster whose radius is 1/3 the distance between the source and destination nodes. By comparison, for a time-division scheme (first the source sends to the relay for a *half* of the time rather than the adjustable fraction f allowed by the authors, then the relay and source send to the destination for the remaining half) with a 5 dB geometric penalty, the allowable cluster size is at most 0.178 times the distance between the source and the destination. This work demonstrates the power of this flexible technique with more realistic assumptions on the wireless channel.

VII. Conclusion

In this paper we presented a decomposition theorem which reduces a network to a set of *competitively, cognitively,* and *cooperatively* behaving clusters. These form fundamental blocks for the study of the fundamental limits of communication in wireless networks with cognitive and non-cognitive nodes. We explored what is currently and classically known about competitive and cooperative clusters, and presented two examples of cognitive behavior. In the first, the *cognitive radio channel*

[0]Here, $C(H)$ does not, in general, designate the capacity of each link as is witnessed by the fact that only for a special subset of matrices is capacity achieved by placing an equal transmit power on each antenna.

is introduced as a two sender two receiver channel where one sender knows the message to be transmitted by the other and thus may cooperate in an *asymmetric* manner. Our second example of cognitive behavior considered *collaborative communications*, where a single sender may be aided by one or more cognitive users, or relays to transmit to a single receiver over a compound channel. In both examples, fundamental limits where obtained in the form of achievable rate (regions) which demonstrate the potential gains of such schemes.

REFERENCES

[1] S.A. Jafar. Capacity with causal and non-causal side information – a unified view. Submitted to *IEEE Trans. Inf. Theory*, October 2005.

[2] C.T.K. Ng and A. Goldsmith. Capacity gain from transmitter and receiver cooperation. In *ProcISIT*, September 2005.

[3] N. Devroye, P. Mitran, and V. Tarokh. Achievable rates in cognitive radio channels. *IEEE Trans. Inf. Theory*, 52(5):1813–1827, May 2006.

[4] N. Devroye, P. Mitran, and V. Tarokh. Cognitive decomposition of wireless networks. In *Proceedings of CROWNCOM*, March 2006.

[5] R. Ahlswede. Multi-way communcation channels. In *ProcISIT*, September 1973.

[6] T. Cover and J.A. Thomas. *Elements of Information Theory*. New York: John Wiley & Sons, 1991.

[7] E.C. van der Meulen. Transmission of information in a T-terminal discrete memoryless channel. Technical report, University of California, Berkeley, 1968.

[8] H. Sato. Two user communication channels. *IEEE Trans. Inf. Theory*, IT-23, November 1985.

[9] G. Kramer. Outer bounds on the capacity of Gaussian interference channels. *IEEE Trans. Inf. Theory*, 50(3), March 2004.

[10] H. Sato. The capacity of Gaussian interference channel under strong interference. *IEEE Trans. Inf. Theory*, IT-27(6), November 1981.

[11] M. Costa. On the gaussian interference channel. *IEEE Trans. Inf. Theory*, 31(5):607–615, September 1985.

[12] T.S. Han and K. Kobayashi. A new achievable rate region for the interference channel. *IEEE Trans. Inf. Theory*, IT-27(1):49–60, 1981.

[13] G. J. Foschini and M. J. Gans. On limits of wireless communications in a fading environment when using multiple antennas. *Wireless Personal Communications*, 6:311–335, 1998.

[14] I. E. Telatar. Capacity of multi-antenna Gaussian channels. *European Trans. on Telecomm.*, 10(6):585–595, 1999.

[15] T. Cover. Broadcast channels. *IEEE Trans. Inf. Theory*, IT-18(1):2–14, January 1972.

[16] T. Cover. Comments on broadcast channels. *IEEE Trans. Inf. Theory*, 44(6):2524–2530, September 1998.

[17] A. E. Gamal. A capacity of a class of broadcast channels. *IEEE Trans. Inf. Theory*, 25(2):166–169, March 1979.

[18] J. Korner and K. Marton. General broadcast channels with degraded message sets. *IEEE Trans. Inf. Theory*, 23(1):60–64, January 1979.

[19] K. Marton. A coding theorem for the discrete memoryless broadcast channel. *IEEE Trans. Inf. Theory*, 25:306–311, May 1979.

[20] H. Weingarten, Y. Steinberg, and S. Shamai. The capacity region of the Gaussian MIMO broadcast channel. *IEEE Trans. Inf. Theory*, 9(52):3936–3964, September 2006.

[21] N. Jindal and A. Goldsmith. Dirty-paper coding versus tdma for mimo broadcast channels. *IEEE Trans. Inf. Theory*, 51(5):1783–1794, May 2005.

[22] E. C. van der Meulen. Three-terminal communication channels. *Adv. Appl. Prob.*, 3:120–154, 1971.

[23] I. Maric, R.D. Yates, and G. Kramer. The discrete memoryless compound multiple access channel with conferencing encoders. In *2005 IEEE International Symposium on Information Theory*, September 2005.

[24] N. Devroye, P. Mitran, and V. Tarokh. Achievable rates in cognitive radio channels. In *39th Annual Conf. on Information Sciences and Systems (CISS)*, March 2005.

[25] N. Devroye, P. Mitran, and V. Tarokh. Achievable rates in cognitive networks. In *2005 IEEE International Symposium on Information Theory*, September 2005.

[26] S.I. Gel'fand and M.S. Pinsker. Coding for channels with random parameters. *Probl. Contr. and Inf. Theory*, 9(1):19–31, 1980.

[27] A. Jovicic and P. Viswanath. Cognitive radio: An information-theoretic perspective. Submitted to *IEEE Trans. Inf. Theory*, 2006.

[28] P. Mitran, H. Ochiai, and V. Tarokh. Space-time diversity enhancements using collaborative communication. *IEEE Trans. Inf. Theory*, 51(6):2041–2057, June 2005.

[29] M.R. Aref. Information flow in relay networks. Technical report, Stanford University, 1980.

[30] T. M. Cover and A. E. Gamal. Capacity theorems for the relay channel. *IEEE Trans. Inf. Theory*, 25(5):572–584, September 1979.

[31] G. Kramer, M. Gastpar, and P. Gupta. Cooperative strategies and capacity theorems for relay networks. *IEEE Trans. Inf. Theory*, 51(9):30373063, September 2005.

[32] P. Gupta and P. R Kumar. Towards an information theory of large networks: an achievable rate region. *IEEE Trans. Inf. Theory*, 49:1877–1894, August 2003.

[33] L.-L. Xie and P. R Kumar. A network information theory for wireless communication: scaling laws and optimal operation. *IEEE Trans. Inf. Theory*, 50:748–767, May 2004.

[34] L.-L. Xie and P. R Kumar. An achievable rate for the multiple level relay channel. Submitted to *IEEE Trans. Inf. Theory*, November 2003.

[35] J. Wolfowitz. *Coding Theorems of Information Theory*. Springer-Verlag, New York, 1978.

[36] I. Csiszár and J. Körner. *Information Theory: Coding Theorems for Discrete Memoryless Systems*. Academic Press, New York, 1981.

[37] A. Reznik, S. Kulkarni, and S. Verdú. Capacity and optimal resource allocation in the degraded Gaussian relay channel with multiple relays. In *Proc. Allerton Conf. Commun., Control and Comp.*, Monticello, IL, October 2002.

[38] M. Katz and S. Shamai. Communicating to co-located ad-hoc receiving nodes in a fading environment. In *Proc. IEEE Int. Symp. Inf. Theory*, page 115, Chicago, IL, July 2004.

[39] K. Azarian, H. El Gamal, and P. Schniter. On the achievable diversity-multiplexing tradeoff in half-duplex cooperative channels. *IEEE Trans. Inf. Theory*, December 2005.

[40] J. N. Laneman and G. W. Wornell. Distributed space-time-coded protocols for exploiting cooperative diversity in wireless networks. *IEEE Trans. Inf. Theory*, 49(10):2415–2425, October 2003.

[41] J. N. Laneman, D. N. C. Tse, and G. W. Wornell. Cooperative diversity in wireless networks: Efficient protocols and outage behavior. *IEEE Trans. Inf. Theory*, 2004.

Natasha Devroye Natasha Devroye is currently a Ph.D candidate in the Division of Engineering and Applied Sciences at Harvard University. She obtained a Honors B. Eng in Electrical Engineering from McGill University in 2001, and a M.Sc from Harvard University in 2003. She has worked as a research intern at Intel Corporation, Santa Clara, as well as Mitsubishi Electric Research Labs, in Cambridge, MA. Her interests include multi-user information theory, wireless communications, and scheduling. She in particularly focused on information theoretic limits of cognitive and cooperative communications.

Patrick Mitran Patrick Mitran(S'01, M'07) received the Bachelor's and Master's degrees in electrical engineering, in 2001 and 2002, respectively, from McGill University, Montreal, PQ, Canada and the Ph.D. degree from the Division of Engineering and Applied Sciences, Harvard University, Cambridge, MA in 2006. In 2005, he interned as a research scientist for Intel Corporation in the Radio Communications Lab. Currently he is a lecturer in the Division of Engineering and Applied Sciences, Harvard University.

Currently he is mostly interested in the study of cooperation and cognition in wireless networks both from an information theorectical viewpoint as well as coding theory and signal processing perspectives. This includes such aspects as capacity analysis, distributed source/channel coding and network coding. He is also interested in channels with timing errors such as seen in magnetic recording channels.

Oh-Soon Shin Oh-Soon Shin received the B.S., M.S., and Ph.D. degrees in electrical engineering from Seoul National University, Seoul, Korea, in 1998, 2000, and 2004, respectively. From 2004 to 2005, he was with the Division of Engineering and Applied Sciences, Harvard University, MA, USA, as a Postdoctoral Fellow. Since April 2006, he has been a Senior Engineer at Samsung Electronics, Suwon, Korea. His research interests include communication theory, wireless communication systems, and signal processing for communication.

Dr. Shin received the Best Paper Award from the CDMA International Conference 2000 (CIC 2000).

Hideki Ochiai Hideki Ochiai received the B.E. degree in communication engineering from Osaka University, Osaka, Japan, in 1996, and the M.E. and Ph.D. degrees in information and communication engineering from the University of Tokyo, Tokyo, Japan, in 1998 and 2001, respectively. Since April 2003, he has been with the Department of Electrical and Computer Engineering, Yokohama National University, Yokohama, Japan, where he is currently an Associate Professor. From 2003 to 2004, he was a Visiting Scientist at the Division of Engineering and Applied Sciences, Harvard University, Cambridge, MA. His research interests include design of modulation and coding for power efficient communications systems.

Vahid Tarokh Vahid Tarokh worked at AT&T Labs-Research and AT&T wireless services until August 2000, where he was the head of the Department of Wireless Communications and Signal Processing. In September 2000, he joined Department of Electrical Engineering and Computer Sciences (EECS) at MIT as an associate professor. In June 2002, he joined Harvard University as a Gordon McKay Professor of Electrical Engineering . Since July 2005, he is a Hammond Vinton Hayes Senior Fellow of Electrical Engineering at Harvard University, and also holds Perkins professorship of applied mathematics. His research is mainly focused in the areas of Signal processing, Communications (wireline and wireless) and Networking . He has received a number of awards and holds 2 honourary degrees.