Impact of Medium Access Control Strategies on the Effectiveness of Advanced Cooperative Hybrid ARQ Techniques

Andrea Munari, *Member, IEEE*, Francesco Rossetto, *Member, IEEE*, and Michele Zorzi, *Fellow, IEEE*

Abstract—This paper explores the relationship between cooperation and two very different medium access strategies: CSMA (epitomized by IEEE 802.11 DCF) and one instance of TDMA (represented by IEEE 802.15.3). By cooperation we mean basic decode-and-forward relaying as well as more advanced forms thereof based on network coding. The essential features that make each system more or less suitable to support the potentially high performance gains of hybrid ARQ are analyzed and studied by means of extensive simulations. The discussion carried out in this paper shows that centralized systems are able to reap far higher gains from a cooperative paradigm than their distributed counterparts, and investigates in depth the reasons that lead to such different behaviors, prompting also some interesting insights for proper design of advanced hybrid ARQ protocols.

Index Terms—Cooperation, CSMA, PAN, Network Coding, Cooperative diversity, Hybrid ARQ.

I. INTRODUCTION

THE importance and potential of cooperation in wireless networks has become manifest in the past half decade [1]. An ever growing amount of work in the literature has repeatedly reported the capability of cooperative diversity to yield significant performance gains over conventional protocols [2], [3] and also some first experimental testbeds have shown encouraging results [4], [5]. However, most of these studies focus on Hybrid ARQ (HARQ) from the perspective of signal processing and information theory, deal overwhelmingly with rather simplified topologies like the three-node source-relay-destination network, and assume idealized multiple access schemes that perfectly coordinate transmissions among terminals.

On the other hand, in recent years some works have started to analyze cooperation from a networking viewpoint [6], highlighting the impact that several issues arising in larger and more realistic scenarios have on the effectiveness of such techniques. In this perspective, an interesting contribution is provided by [7], where the authors show that many common

yet simplistic modeling assumptions may lead to the prediction of non-achievable gains, and emphasize how the medium access policy chosen to support a cooperative paradigm can dramatically affect its efficacy. The inferences derived in [7] stem from the nature of medium access control rather than from a specific cooperative strategy, and then they also hold for advanced forms of relaying. From this point of view, for instance, the idea of using Network Coding (NC) to make HARQ more efficient has appeared not long after cooperation itself [8]–[10], and a remarkable example is the usage of coded retransmissions to both relay data and serve additional traffic. This principle has proved to be very powerful on paper [9], [10], yet once more it was noted that the achievable gains heavily depend on the link layer protocols used, e.g., Carrier Sense Multiple Access (CSMA) [11] or Time Division Multiple Access (TDMA) [12].

Starting from these remarks, the main goal of this paper is to analyze the impact of two realistic and widespread MAC schemes, namely IEEE 802.11g [13] and IEEE 802.15.3 [14], on the efficiency of cooperation. These protocols are some of the most widely used implementations of the carrier sense- and centralized scheduling-based multiple access policies, respectively. Therefore, their analysis has first of all practical relevance. Also, the differences between the physical layers employed by the two standards may lead to different data rates, but do not fundamentally alter the network capacity, for instance in terms of simultaneous packet receptions. This ensures a comparison that is not biased by non link layer-related aspects. Even more relevant, they implement two radically dissimilar philosophies, and hence enable the investigation of a wide range of issues for cooperative networking. Finally, since 802.11g and 802.15.3 are highly representative of their own class of wireless MACs, such study provides valuable insight into the essential relationships and problems between CSMA/TDMA and HARQ.

In order to infer broadly applicable conclusions, we consider both a widely studied approach such as decode-and-forward [15], [16], which stands as a benchmark in the field of relaying, and newer ideas such as network-coded cooperation, which promises further performance and efficiency improvements. The key contributions of this paper can be summarized as follows:

- to the best of the authors' knowledge, this is the fist work that presents a comparative study on the effectiveness of relaying considering such different classes of link layers;
- the network level analysis of HARQ developed in this

paper deals also with network coded cooperation, while, with the exception of [9], the research on this topic focuses almost exclusively on physical layer design. Our study provides some insight into the system level design of such relaying strategies;

- we show that the intrinsically distributed nature of CSMA prevents a large fraction of the positive effects of cooperation. Instead, the coordination brought by TDMA avoids most of the problems that occur in distributed scenarios, increases the efficiency of relaying, and can thus boost the benefits of sophisticated versions of HARQ, like network-coded cooperation;

- not only does this paper quantify the impact of different channel access policies on cooperative solutions, showing twice as large a gain in terms of network performance offered by centralized over distributed link layers in a variety of environments, but also it investigates in depth the issues behind these differences. A key conclusion is that such a diverse behavior is rooted in the intrinsic nature of the considered medium access policies and does not stem from the specific protocol implementations under analysis;

- the thorough discussion carried out in this work provides some hints for the efficient design of cooperative protocols in both CSMA- and TDMA-based networks.

The rest of the paper is structured as follows. Section II studies how apt the IEEE 802.11 and 802.15.3 MAC protocols are to yield the potential gains of cooperation, and in particular it highlights whether the fundamental properties of these two schemes boost or reduce the efficiency of cooperative retransmissions. The specific multiple access schemes that will be studied are reported in Section III and are compared by means of simulation in Section IV. Finally, the conclusions are drawn in Section V.

II. COOPERATION IN CSMA AND TDMA ENVIRONMENTS

Cooperative techniques have recently attracted a great deal of attention in the wireless research community, thanks also to their capability of improving ARQ schemes. In the event of a communication failure between a source S and a destination D, basic retransmission policies require the former to retry after a certain time interval, whose duration depends on the underlying medium access control. However, if the decoding has not succeeded due to a bad state of the propagation channel between S and D (e.g., deep fading), the recovery mechanism may involve several attempts, in light of the time-correlation property of the channel, with detrimental effects on energy and bandwidth consumption. On the other hand, thanks to the broadcast nature of the wireless medium, other nodes may have decoded the packet sent by the source even if it was not intended for them. The decode-and-forward cooperative paradigm proposes that one such terminal immediately performs a retransmission on behalf of S. The destination, then, has at its disposal two copies of the same packet received over statistically independent channels, and can take advantage of spatial diversity, e.g., by performing Chase combining.

Such a simple technique has the potential of both shortening the recovery phases and improving their reliability. However, nodes may be discouraged from cooperating, as relays are
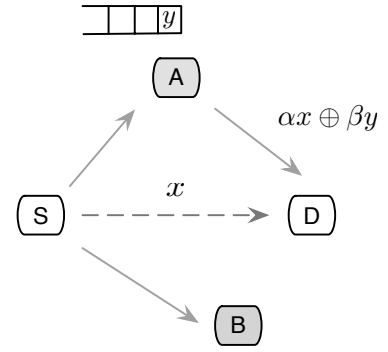

Fig. 1. Reference topology for a HARQ phase for a communication between S and D. x is the original packet sent by S, while y is a packet taken from A's queue. $\alpha x \oplus \beta y$ represents a linear combination of the two payloads, according to the rules of MIMO_NC [17].

asked to behave in a selfless way, spending their own resources to help other terminals in the hope that similar support will be offered to them when needed. This intrinsic limitation has been recently mitigated by the introduction of hybrid cooperative-Network Coded (Coop-NC) ARQ [10], [11]. Let us refer to the topology depicted in Fig. 1, and assume that the transmission of packet x over the S-D link has failed. If another terminal, say A, has decoded x and has in its own queue a data unit y addressed to the same destination, Coop-NC ARQ exploits a cooperative phase to transmit a linear combination of x and y instead of just a copy of the former. The advantage offered by this approach is twofold. In the first place, involvement in cooperation is boosted, since relays are offered the opportunity to pursue their own interest. On the other hand, encoded retransmissions allow more traffic to be served at no additional cost in terms of energy and bandwidth and with only a slight increase of the error probability, inducing beneficial effects on the overall network performance. This advanced form of cooperation is based on a novel type of physical layer, called MIMO_NC [17], which makes it possible to decode both x and y at the destination starting from a corrupted version of the former while keeping the diversity order provided by Chase combining and introducing low additional cost in terms of complexity.

Analytical studies have extensively proven that both decode-and-forward and hybrid Coop-NC ARQ elicit important gains over basic schemes in simple scenarios [1], [9], [11]. However, when such techniques are implemented in larger and more articulated environments, several networking issues which are typically neglected from a theoretical perspective arise. In order to investigate the impact of these aspects, we start by identifying the requirements that have to be met for cooperation to successfully take place. From this viewpoint, and with reference to the topology in Fig. 1, three key conditions can be identified:[1]

R1) D is able to identify the corrupted packet x sent by S.
    In order to request a cooperative retransmission, the

---

[1]While the necessity of the presented conditions stems from the definition given earlier in this section of a decode-and-forward based cooperative phase, we stress that they are not sufficient, i.e., they do not guarantee successful decoding at the destination. In fact, the outcome of a relaying process also depends on other factors such as the SINRs of the involved communications.

destination must have identified the packet that has been lost. Not only does this mean that D has to be available for reception, i.e., neither transmitting nor synchronized to other ongoing communications, but also that it has to be aware that an attempt addressed to it for a specific data unit is being performed. The latter condition can be met either by decoding the header of x, which contains information on nodes involved in the transmission as well as the ID of the packet,[2] or by having been informed in advance of the incoming transmission, as in the case of a scheduled system. Moreover, a MIMO_NC based retransmission is possible only if D has also been able to cache a copy of the corrupted signal.

R2) *One ore more relay candidates are available.*
The possibility for a terminal to act as cooperator is subject to specific constraints. In the first place, a potential relay must not be involved in other ongoing communications. Secondly, according to the decode-and-forward paradigm considered throughout this paper, the node must have successfully decoded the packet originally sent by the source. Finally, a terminal must be granted access to the medium to perform a HARQ transmission. Such a requirement is not trivial, and may have a severe impact on the availability of relays. Indeed, the communication between S and D does not reserve in general bandwidth for a cooperative stage, and relayed transmissions have to be carefully handled, so as not to harm other ongoing links in the network.

R3) *A policy to trigger and define the HARQ phase is specified.*
When several nodes are involved in a cooperative process, additional coordination at the link layer is required. In particular, HARQ should be triggered by the destination only when it is useful, i.e., D has not decoded the packet sent by S, and if it can be successfully supported, e.g., for a MIMO_NC based retransmission if D has cached a corrupted version of the data unit. Transmissions performed under different conditions would only result in undesired interference and wasted bandwidth. Moreover, a strategy to identify who among multiple relay candidates may act as cooperator has to be provided, so as to avoid collisions and inefficiencies. Incidentally, the policy that defines HARQ phases should be designed in order to maximize the success probability. While this is not a strictly necessary requirement, it is certainly desirable from an engineering point of view.

The capability of a network to cope with these general requirements is tightly related to the control policy used for medium access. Link layer strategies, thus, play a role of paramount importance for the effectiveness of cooperative solutions. Starting from this observation, in this paper we investigate the capability of two radically different networking environments, namely *completely distributed* and *centralized* systems, of effectively supporting relaying schemes. We regard as reference for the former class of solutions non-centralized ad hoc networks based on carrier sensing, epitomized by

the IEEE 802.11 DCF working in ad hoc mode. As far as centralized systems are concerned, we focus on TDMA-based schemes as per the rules of the IEEEE 802.15.3 standard for high rate personal area networks. In the remainder of this section we deal separately with such medium access strategies, analyzing their impact on HARQ implementations.

A. Distributed systems: Carrier Sense Multiple Access

The distributed medium access policy implemented by the IEEE 802.11 DCF[3] [13] requires a node that has data to send to choose a random backoff interval and sense the power level on the wireless channel (*carrier sensing*). If the perceived value exceeds a given threshold, the terminal infers the presence of other ongoing links and freezes its backoff until the medium is sensed idle again. Conversely, when the contention window expires with no other communication being detected, the terminal transmits its packet. If the destination successfully decodes the payload, an acknowledgment (ACK) is sent in response, and the data exchange comes to an end. Otherwise, the source starts another backoff interval (following the Binary Exponential Backoff, BEB, mechanism [13]) and iterates the procedure until either a positive feedback is received or a maximum number of attempts (Short Retry Limit, SRL) is reached.

While simplicity and lack of a coordinating unit make this link layer particularly suitable for ad hoc networks, its completely distributed nature poses some intrinsic challenges that hamper the effective implementation of cooperative solutions.

In the first place, a terminal becomes aware of a communication in which it should be involved only when the transmission is actually performed. This implies that, in case of a data loss, cooperation can solely be enabled when the addressee of the payload successfully receives the header of the incoming packet (*requirement R1*). Such a condition is not easily met in carrier sense-based environments. Indeed, while payload and header may be separately encoded, they experience very similar channels, e.g., in terms of fading and interference. Therefore, the decoding probabilities for the two of them are correlated. In addition, the destination may lose data if it is already synchronized to another ongoing transmission not overheard by the source. Also in this case the incoming packet cannot be cached for later combining, as it is regarded as interference at the receiver. This issue, often referred to as the hidden terminal problem, stems from the asynchronous and decentralized behavior of 802.11, according to which each node in the network is asked to decode (and thus can be locked by) any communication taking place in the surroundings, even if not addressed to it.

Carrier sense-based access control also influences the availability of relays (*requirement R2*), for two primary reasons. First of all, the hidden terminal problem may affect potential cooperators as well, as nodes other than the intended destination may fail to decode a packet from S, e.g., if they are synchronized to the reception of other data units. Secondly, terminals that could perform a HARQ phase may not be allowed to transmit. In fact, although the active direct link

---

[2]We point out that the header may be successfully received even if the payload is corrupted, since it can be independently channel encoded and have a separate CRC.

[3]Throughout this paper, we consider carrier sense multiple access without channel negotiation, whereas CSMA with collision avoidance is left as part of future work.
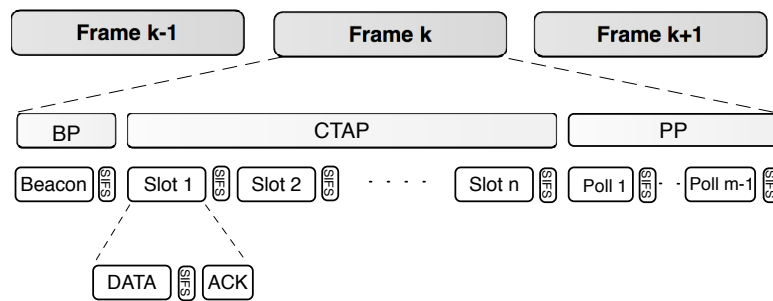
Fig. 2. Frame structure for the time division access-based systems under analysis. In this example, the cluster is composed by m nodes, including the clusterhead.

protects an area centered at S, relay candidates, possibly located at the borders of such a region, are likely to sense a significantly higher aggregate interference and may thus be forced to refrain from accessing the channel.

Finally, we remark that distributed systems lead to non-trivial policies to identify who has to act as cooperator (*requirement R3*). Static approaches that associate a fixed relay to each source-destination pair may turn out to be highly inefficient, as the chosen terminal may be unavailable or experiencing bad channel conditions towards the destination when its action is needed. Conversely, decentralized algorithms to opportunistically select a cooperator among a set of candidates (i.e., nodes that have decoded the original packet sent by the source and are allowed to transmit) may lead to collisions or suboptimal decisions, since contenders are reciprocally unaware of each other's condition.

B. Centralized systems: Time Division Multiple Access

As a reference for time division-based multiple access we focus on high-rate Wireless Personal Area Networks (WPANs), as described in the IEEE 802.15.3 standard [18]. These systems, indeed, represent a good scenario for the implementation of HARQ due to their demand for high throughput even under harsh interference conditions. In WPANs, the network is divided in clusters, each comprising a clusterhead and other wireless devices. All the terminals in a cluster are located within each other's communication range, and single hop peer-to-peer traffic can take place among any pair of nodes. Time is organized in successive frames, each of them subdivided in three periods, as shown in Fig. 2: a *Beaconing Phase* (BP), a *Channel Time Allocation Period* (CTAP), and a *Polling Phase* (PP). During the polling phase of frame k, every node in a cluster is allocated a slot for sending a POLL packet to its clusterhead, requesting channel access for data traffic in frame $k + 1$. Once all the polls have been collected and considering the status of its own queue, the cell master runs an algorithm to assign the n data slots of the subsequent CTAP. The computed schedule is then broadcast to the terminals in the cell during the beaconing phase of frame $k + 1$. Upon receiving the beacon, all the members of the cluster become aware of how to coordinate medium access for the following data exchanges. In particular, when a slot allocated to a given communication starts, the source node transmits the corresponding data unit and waits for a feedback. If the packet is successfully received, the destination responds

with an ACK. Otherwise, no reply is provided, and the source keeps the packet in its queue for later retransmission until the SRL is reached.

The described link layer requires synchronization among terminals within a cluster and introduces additional complexity with respect to the procedures that characterize CSMA. However, the presence of a coordinating unit significantly favors the implementation of HARQ-based relaying. First of all, nodes become aware of all the transmissions that take place in their cell during the beaconing phase, and therefore successful decoding of the packet headers is no longer a requirement for cooperation to take place. From this viewpoint, a destination can cache corrupted versions of data units addressed to it as long as the beacon for the current frame has been received (*requirement R1*). Indeed, the receivers can easily avoid being locked to undesired communications, since only one transmission per cluster can be active at a time and packets coming from other cells can be discarded by taking advantage of cluster-based synchronization. Moreover, in contrast to CSMA systems, non-decodable data can be stored even in the presence of extremely unfavorable fading conditions between the source and the destination, as beacon and payload are in general received over different and statistically independent channels. Incidentally, we remark that the absence of the hidden terminal problem can have a beneficial effect on the availability of relay candidates too, since nodes can properly synchronize to, and thus potentially cache, the only transmission taking place in their cell at any given time.

The centralized policy of TDMA also eases the identification and selection process of relays (*requirements R2 and R3*), as the organization of cooperative phases can be completely entrusted to the clusterhead. Polling processes can easily be extended to include notifications of both lost data units at intended destinations, i.e., HARQ chances, and successfully decoded packets addressed to other nodes, i.e., availability to act as relay. Upon receiving this information, the cell master can decide whether and how to suitably allot CTAP slots to cooperative communications and notify all the involved terminals exploiting the beaconing phase. Following this approach, inefficient contentions among candidates are avoided, and optimal retransmissions may be possible.

Two concluding observations are in order. In the first place, we notice that the issues discussed in this section for distributed and centralized systems are related to the capability of effectively triggering any HARQ procedure, and thus

affect both plain, e.g., decode-and-forward, and advanced, e.g., network coded, cooperative techniques, possibly undermining their beneficial contributions. Secondly, the problems that have been highlighted stem from the interaction of several nodes, and therefore become more pronounced in larger networks or at high traffic loads, which are exactly the conditions in which cooperation could be exploited at its utmost.

III. PROTOCOL DESCRIPTIONS

To evaluate the impact of medium access control policies on cooperation, we focus on the implementation of both decode-and-forward and cooperative-Network Coded ARQ in CSMA- and TDMA-based systems. The protocols that we propose and study are extensions of the standard solutions described in Section II, and their key features are highlighted in the remainder of this section. A complete description of the schemes can be found in [11] and [12], respectively.

A. Carrier Sense Multiple Access

1) Decode-and-forward cooperation: Let us refer to the topology depicted in Fig. 1, and assume that the transmission of x between S and D has not succeeded. If the destination has cached a corrupted version of the payload, a Negative ACKnowledgement (NACK) packet is broadcast, asking for HARQ. Nodes that receive this request and have successfully decoded the data unit from the source enter a distributed contention based on carrier sense to select the relay. Each candidate picks a random backoff, whose duration, in slots, is uniformly drawn in $\{0, 1, \ldots, CW_{rel}\}$, and listens to the medium. If the power level on the channel exceeds a threshold $\Gamma_{CS}^{(rel)}$, the terminal infers that either somebody else is cooperating or other communications are going on in the surroundings and refrains from transmitting. Otherwise, the node wins the contention and, at the end of the backoff, sends a copy of x. Should the decoding at D fail again, no feedback is provided and the source will perform another attempt according to the BEB mechanism. Conversely, an ACK addressed to S is transmitted. In this case, the source transmits a new packet (if available), while the cooperator resumes the activity in which it was involved prior to the relay contention, e.g., if it was in a backoff stage, it continues with the countdown.

2) Cooperative-Network Coded ARQ: With reference to Fig. 1, suppose D has cached a corrupted version of x, and assume that A has in its own traffic queue a packet y addressed to the same node. As discussed in Section II, in such a situation the Coop-NC paradigm proposes to exploit a HARQ phase to convey both messages to the destination. However, if the quality of the stored version of x at D is extremely poor, a network-coded retransmission would have too low a success probability with respect to decode-and-forward cooperation.[4] In order to cope with this aspect, the NACK packet can be modified to include an *NC_flag* field. Before asking for HARQ, the destination checks the average Signal to Interference and Noise Ratio (SINR) that

characterized the cached corrupted payload. If this value is below a given threshold Λ_{Th}, the *NC_flag* is set to 0, and only basic cooperation is permitted. Otherwise, the flag is set to 1, signaling that a coded retransmission can be successfully supported. Terminals that receive the request for cooperation and that have decoded the packet sent by the source enter the distributed contention procedure discussed earlier. If the *NC_flag* is not active or the chosen relay does not have traffic for D, decode-and-forward is carried out as described in Section III-A1. In the opposite case, i.e., the *NC_flag* is active and the winner of the contention has traffic for D, the latter, say A, generates a random linear combination of x and y following the principles of network coding and transmits the obtained packet. In this condition, as a result of the working principles of MIMO_NC [17], D can either retrieve both the encoded data units or neither of them. If decoding fails, no feedback is provided. Instead, when Coop-NC ARQ succeeds, a single ACK is sent. In any case, both the original source and the relay can infer the outcome of the cooperative phase and act consequently, i.e., discarding their packet or keeping it in the queue for later retransmission.[5]

B. Time Division Multiple Access

1) Decode-and-forward cooperation: The basic TDMA scheme described in Section II-B can easily be extended to take advantage of decode-and-forward techniques. In our solution, each node keeps a so called cooperative queue to cache packets for which it may act as relay. In particular, during CTAPs, terminals store successfully decoded data units that are not addressed to them and whose acknowledgment they do not overhear. Moreover, in order to increase the success probability of retransmissions, a node, say A, proposes itself as relay only for destinations it shares a good link with (i.e., only if the SINR of the last packet decoded over the D-A channel was higher than a threshold Λ_{coop}).

The polling phase is used also to notify availability to perform HARQ. This can be achieved by reserving p_{coop} out of the p entries of a POLL packet[6] to describe payloads cached in the cooperative queue during the current frame. Once all the polls have been received, the clusterhead can proceed with the schedule computation, allocating up to n_{coop} of the n slots in the successive CTAP to cooperative communications. Notice that nodes may erroneously propose to act as relays for data units that have been successfully delivered. Should such a decode-and-forward phase be allocated, it would result in a waste of energy and bandwidth. In order to prevent this issue, the cell master[7] allots cooperative slots only if S has not received an ACK and one or more terminals are available as relays.

The designed schedule is then distributed as usual by means of a beacon. When a node is reserved a slot for HARQ, it

---

[4]With MIMO_NC, the success probability of a coded transmission, i.e., decoding of both x and y, is slightly lower than that of basic Chase combining, i.e., retrieval of just x, even though the diversity order is preserved [17].

[5]In contrast to [11], we do not implement the possibility of using network coded retransmissions to serve traffic addressed to secondary destinations, i.e., the case in which y is for a node other than D, nor do we give priority to NC phases by letting relay candidates with packets for D pick shorter contention windows. These choices have been made in order to simplify the protocols and take into account the results of some studies we performed, which showed how the aforementioned features do not lead to significant improvements.

[6]For a detailed description of the packet structures please refer to [12].

[7]The words *cell* and *cluster* will be used interchangeably.

simply transmits the corresponding packet, without waiting for any feedback. Conversely, if Chase combining succeeds, the destination sends an ACK to S.

2) Cooperative-Network Coded ARQ: As discussed in Section III-A2, network coded retransmissions should be triggered only when the SINR on the corrupted payload available at the destination is sufficiently high. However, differently from the immediate feedback implemented in CSMA, we propose to exploit polling to identify potential Coop-NC phases. To this aim, the POLL packet is split in three sections: an *NC-request* part, a *cooperative* part, and a *basic traffic* part. The first part is composed by p_{NC} fields, describing the non-decoded packets addressed to the node that have been cached during the current frame with average SINR greater than Λ_{Th}. The successive p_{coop} entries are used by the source of the POLL to propose itself as cooperator as in the decode-and-forward case. Finally, the remaining $p - p_{NC} - p_{coop}$ fields are devoted as usual to requests for channel access to serve own traffic. When computing the schedule for a frame, the clusterhead scans through the *NC-request* sections of all the received polls, and tries to allocate hybrid coop-NC phases for each of the notified packets. This procedure is iterated until either n_c slots have been assigned or no other Coop-NC transmissions can be supported. In the latter case, the remaining part of the n_c slots reserved to HARQ are distributed, if possible, to decode-and-forward cooperation as in Section III-B1.

Network coded retransmissions are handled as usual, with the relay node sending a linear combination of the involved payloads and the destination either responding with a single ACK addressed to both the cooperator and the original source, or providing no feedback at all.

IV. PERFORMANCE EVALUATION

In this section, an extensive investigation is carried out by means of Omnet v3.0 [19] simulations so as to quantify the impact of the considered access control strategies on the effectiveness of cooperative techniques. We start our discussion by presenting the reference system model used in our studies in Section IV-A. Then, the focus is moved to the analysis of the collected performance results. From this viewpoint, we distinguish two main types of metrics: end-user and cooperative. The former are those quantities that directly affect the end-user satisfaction, i.e., aggregate throughput, delay, transmit energy consumption. The latter are those metrics that are not immediately relevant for the application layer, but shed light on the relationship between the access scheme and cooperation. Examples are the reasons why a cooperative retransmission is not carried out or the percentage of HARQ phases that employ NC. Since the discussion of these two families of performance criteria is quite different, the former are presented in Section IV-B, while Section IV-C is devoted to the latter.

A. System Model

In order to have a fair and insightful comparison of the considered cooperative solutions in centralized and distributed systems, we have focused on networks composed by several clusters arranged in a circular pattern so that each of them

has two adjacent cells.[8] Such a configuration, obtained by positioning the centers of the clusters at the vertices of a regular hexagon, is symmetric and leads to a common inter-cell interference statistics for all the subnetworks, avoiding border effects that may mask some key results. The reference scenario for our studies is composed by six cells, each populated by six static nodes within communication range, and represents a harsh test for cooperative protocols in view of the high level of co-channel interference it entails. As far as CSMA is concerned, every cluster operates in ad hoc mode, so as to stress the completely distributed nature of such medium access policy. Conversely, TDMA-based protocols rely on the presence of a coordinator, which is chosen as the node closest to the center of each cell. Single hop all-to-all communications take place within a cluster, with every terminal generating traffic according to a Poisson model unless otherwise specified.

A transmission between two nodes δ meters apart is affected by inter- and intra-cell interference as well as by correlated frequency flat Rayleigh fading, so that the average useful power at the receiver is proportional to $P\delta^{-\alpha}$, where P is the transmitter power and α is the path loss exponent. The temporal coherence of the wireless channel is modeled according to Jakes' approach for land mobile fading [20], with correlation between two instances of the same channel spaced in time by τ seconds given by $J_0(2\pi f_d \tau)$, where $J_0(\cdot)$ is the zero-order Bessel function of the first kind and f_d is the maximum Doppler frequency for the environment under consideration. The reference channel parameters for our studies are reported in Table I.

Three schemes are compared, and each of them is instantiated for CSMA and TDMA: on the one hand "Basic", representing the standard solutions discussed in Section II; on the other hand "Coop" and "NC", implementing HARQ policies as explained in Section III. Moreover, we have designed and studied protocol bounds that idealize some aspects of medium access contention in the two networking environments under consideration. The rationale behind this approach is twofold. Indeed, not only can the performance analysis of the bounds provide an insight on the share of potential gains attained by the implementations that we propose, but also it sheds light on the intrinsic limitations of different link layers in exploiting cooperation. As for CSMA, the bound that we consider always selects the best relay among the available candidates, i.e., the node that experiences the most favorable channel conditions to the destination. Moreover, the selection process is idealized so that all potential cooperators are immediately informed of how the HARQ phase will be organized, i.e., no contention is required and collisions among relays are avoided. In TDMA, conversely, the protocols have been bounded by letting terminals have unlimited cache for packets involved in cooperative processes and include in their polls the complete description of their queues, at no additional cost in terms of overhead. With this approach, the clusterheads can maximize the number of allocated decode-and-forward and network coded retransmissions. The SRL values have been chosen so as to achieve comparable packet delivery

---

[8]The system works under universal frequency reuse, i.e., all nodes transmit and receive on the same frequency.

TABLE I
PARAMETERS USED IN OUR SIMULATIONS

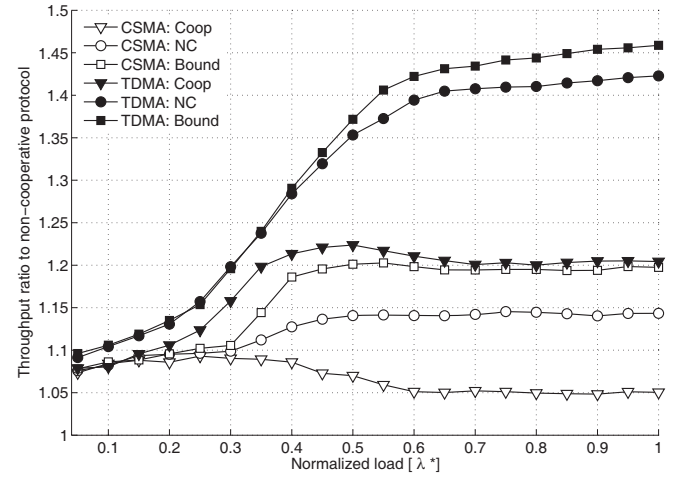|	CSMA	TDMA
Transmission power (dBm)	10	10
Noise Floor (dBm)	-91.6	-91.6
CS threshold, Γ_{CS} (dBm)	-89	//
CS threshold for relay contention, $\Gamma_{CS}^{(rel)}$ (dBm)	-89	//
Detection threshold (dBm)	-89	-89
Path loss exponent	3.5	3.5
Maximum Doppler shift (Hz)	11.1	11.1
Saturation load (pk/s/node)	128	960
Slot, DIFS, SIFS duration (μs)	20, 128, 28	//, //, //
Transmission/reception switching time (μs)	//	20
Carrier Frequency (GHz)	2.4	2.4
Data Rate (Mbit/s)	6	22
Initial maximum contention window (slots)	128	//
Short Retry Limit - Coop and NC	3	3
Short Retry Limit - Basic	4	4
Number of slots used for relay contention, CW_{rel}	32	//
Minimum SINR to trigger a Coop-NC phase, Λ_{Th} (dB)	3	3
Minimum SINR with a destination to act as relay, Λ_{coop} (dB)	//	5
Simulation Time after transient (s)	9	9
DATA header (bit)	272	//
Payload (bit)	4000	4000
ACK/NACK (bit)	112	224
Polling pkt Basic/Coop/NC (bit)	//	200/224/272
Beacon Basic/Coop/NC (bit)	//	192/208/220
Number of data slots per frame, n	//	8
Number of HARQ slots per frame, n_{coop}	//	3
Tx req. p per polling pkt Basic/Coop/NC	//	8/10/10
NC req. per polling packet, p_{NC}	//	2
Coop. proposals per polling packet, p_{coop}	//	3

Fig. 3. Ratio of aggregate throughput for cooperative schemes over that for the basic protocol vs. normalized load.

have been obtained by averaging the outcome of 25 different topologies to achieve the desired statistical confidence.

The normalized aggregate throughput is depicted in Fig. 3. Several observations are in order. First of all, at very low loads the gain enabled by HARQ solutions is about 7-9% and is almost independent of the specific protocol and of the medium access policy. The rationale is that for very light traffic the system is just noise limited and no network coding opportunities are possible. Thus, the different protocols are simple cooperative systems with no medium access contention and exhibit very similar performance because they work in very comparable environments. As the traffic increases, CSMA-Coop and TDMA-Coop gains first rise and then fall. The reason is that the throughput of the Basic protocols grows more for heavier loads and hence the relative gain of the Coop systems decreases. Such a behavior is due to the fairer resource distribution induced by cooperation: when relaying is in place, the nodes with good channels (and hence potentially high throughput) use part of their resources to support terminals in worse conditions in their retransmissions. The outcome is a more balanced distribution of the bandwidth, at the expense of a slightly reduced throughput. It is important to notice that this issue does not affect Coop-NC ARQ, as coded retransmissions boost the throughput also of the high performing nodes when they help other terminals.

It is also remarkable that TDMA cooperative systems attain about twice as large an improvement as their CSMA counterparts, a trend that will characterize the other aggregate metrics as well. Accordingly, TDMA-NC is able to achieve over 40% gain on TDMA-Basic and 20% on TDMA-Coop, far more than the 15% and 5% improvement of CSMA-NC and CSMA-Coop over CSMA-Basic, respectively. Such an advantage also holds for the Bounds, proving that the TDMA environment is inherently more suitable for cooperation, in accordance with the discussion of Section II. Incidentally, we remark that both NC-enhanced protocols are close to the respective Bounds, which verifies the efficiency of the schemes proposed in [11] and [12].

Another key end-user metric is the MAC delay for successfully delivered packets (Fig. 4). At low loads, TDMA-Coop
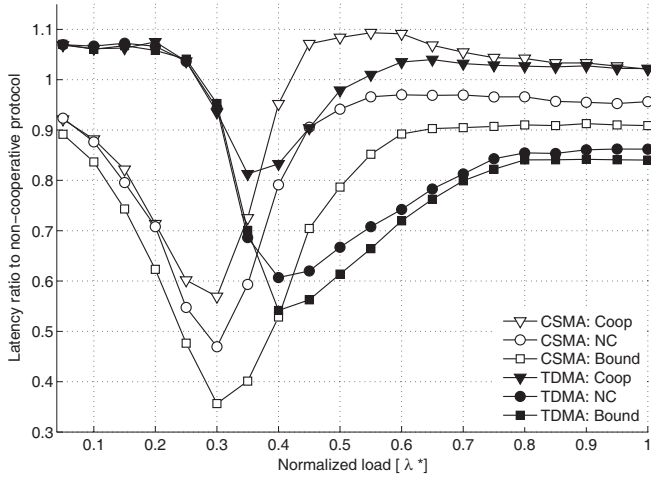
ratios at saturation (around 98%). This led to 4 attempts for the basic protocols and 3 for the cooperative approaches in both centralized and distributed systems. All other relevant simulation parameters are reported in Table I.

B. Simulation Results: end-user metrics

Before delving into a detailed discussion of the results, a few important remarks have to be made. In the first place, in view of the differences between the two transmission policies under analysis (802.11 and 802.15.3), a comparison of the absolute values of end-user performance would be misleading. Therefore, the plots that we report always show the gains attainable by advanced cooperative solutions over plain ARQ, i.e., the results achieved with Coop, NC, and Bound for a given metric are normalized to the performance of the Basic protocol under the same medium access policy. Moreover, due to the different data rates employed, the injected traffic that can be supported is significantly higher in the TDMA setting than in CSMA. Hence, it is more insightful to refer the metrics to the normalized load λ^*, defined as the absolute load in terms of generated packets per node per second divided by the saturation load, i.e., the per node load at which 98% of the saturation aggregate throughput is achieved. Finally, we point out that our simulations take into account all necessary signaling and overhead required for cooperative phases to be triggered and performed, thus being accurate from both a PHY and a MAC point of view, and that the presented results

Fig. 4. Ratio of latency for cooperative schemes over that for the basic protocol vs. normalized load.

Fig. 5. Ratio of transmission energy consumption for cooperative schemes over that for the basic protocol vs. normalized load.

and TDMA-NC actually underperform with respect to TDMA-Basic. Such a behavior stems from channel correlation: with low traffic, the SINR is almost equal to the SNR and hence its fluctuations are due almost exclusively to the channel time variability. In TDMA, the interval between two retransmissions can be quite short (one frame, at a rate three times larger than the one used in CSMA), while in CSMA it may be much longer (one backoff interval). Hence, with CSMA successive retransmissions experience less correlated SNR and also nodes in unfavorable positions may in the end enjoy some time diversity (albeit after a long time). Conversely, such an advantage does not emerge with the considered TDMA implementations, so that at light loads the packets are often either immediately delivered at the first attempt (and thus are affected by a small latency) or dropped.

At medium loads, cooperation starts being efficient and yields very significant improvements in both scenarios. Finally, when saturation is approached, many retransmissions become network coded and thus the NC protocols have truly an edge over their decode-and-forward counterparts, reducing the delay by up to 10% and 18%, respectively. Such result is possible because of the shorter queueing time experienced by all packets and the absence of delay due to channel contention for the retransmitted frames. It is also remarkable that in this condition both TDMA-Coop and CSMA-Coop have slightly worse delay than the reference protocol. This is again due to their more fair behavior towards terminals with severe loss rates: high-throughput nodes sacrifice part of their bandwidth (and of their low-delay packets) to help other terminals which intrinsically suffer a worse delay. The impact of this drawback, instead, does not appear for NC protocols, as it is counterbalanced by the beneficial effects of coded retransmissions.

The final end-user metric is transmit energy consumption, depicted in Fig. 5 and defined as the ratio of the total energy spent in transmission to the number of delivered data bits. The observations follow what has been discussed so far: the HARQ solutions improve energy efficiency at medium load, while for high loads the unfairness of CSMA-Basic or TDMA-Basic may curb the gains. The NC protocols are however
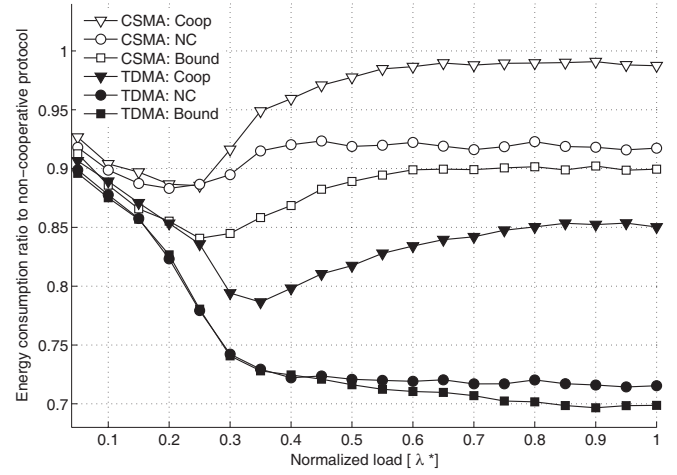
far less affected by (CSMA) or even immune to (TDMA) this effect because of the aforementioned mechanisms. Notice the significant improvement achieved by TDMA-Bound (and TDMA-NC) over CSMA-Bound, which testifies once again the inherent superiority of cooperation in a TDMA environment.

In our work, we have also carried out several simulation studies over relevant parameters in order to assess the generality of the conclusions drawn so far. Specifically, we have investigated the impact of the wireless channel time correlation, of the number of nodes per cell, of the carrier sense threshold to drop the relay contention $\Gamma_{CS}^{(rel)}$ (for CSMA), and of the maximum number of cooperative slots per frame n_{coop} (for TDMA). The throughput gain for cooperative protocols over their basic counterparts as a function of the maximum Doppler frequency f_d of the environment is reported in Fig. 6. The plot shows that the performance gap between all relaying schemes and their correspondent benchmarks diminishes as the channel correlation time decreases. This general trend stems from the fact that the gain brought to HARQ approaches by spatial diversity tends to be compensated for high Doppler frequencies by time diversity, which is enjoyed by all the protocols regardless of the type of retransmission mechanism they implement. It may also be noticed that this effect is less pronounced in CSMA. Not only is the data rate in this case lower, but also the BEB mechanism induces a longer delay between two successive attempts, enabling some form of time diversity already in the presence of slowly-varying channels. All in all, even though the gains of the cooperative protocols reduce as the Doppler frequency increases, they remain rather large especially for the NC based approaches in both environments.

Due to space constraints, we only report the qualitative results for the other parametric studies we performed:

- by varying the cell density it has been observed that the relative gains of CSMA-NC increase as the number of nodes per cluster grows, although the improvements over basic ARQ remain substantially lower than those achievable with TDMA (e.g., the throughput gain of CSMA-NC reaches a maximum of 20% for 12 nodes/cell).
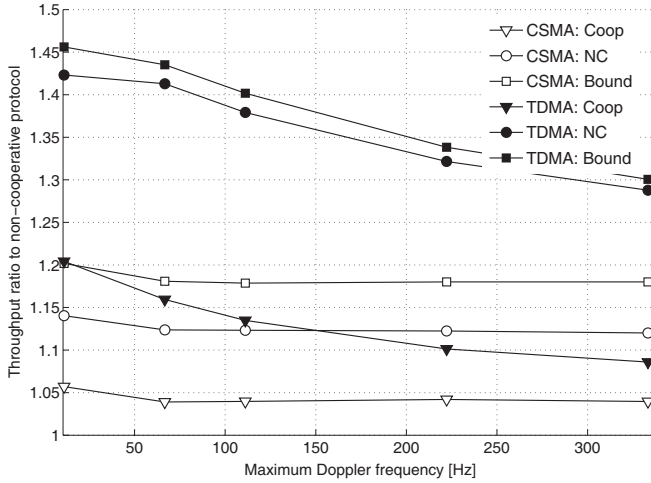
Fig. 6. Ratio of aggregate throughput at saturation for cooperative schemes over that for the basic protocol as a function of the Doppler frequency.
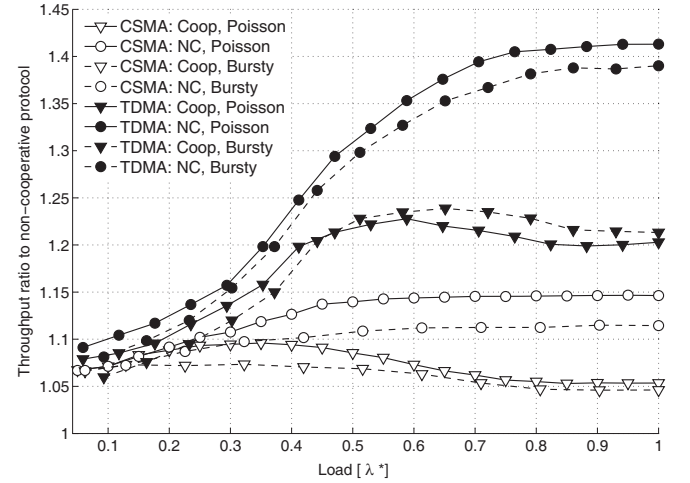

Fig. 7. Ratio of aggregate throughput for cooperative schemes over that for the basic protocol vs. normalized load. Comparison for Poisson traffic (solid lines) and bursty traffic (dashed lines).

Distributed systems incur higher interference levels with the cell size, leading to more retransmissions and thus triggering a larger share of HARQ and NC opportunities. In contrast, schedule based policies are capable of handling larger node populations without triggering additional interference, so that the relative gains increase less noticeably with the cell density;

- as to MAC dependent parameters, in CSMA it is arguable that there exists an optimal value for $\Gamma_{CS}^{(rel)}$. Too low a value will force too many relays to drop the contention, while an exceedingly large threshold will generate excessive interference, again thwarting cooperation. The optimal performance is achieved at -86 dBm (3 dB above the value used in the simulations), but the gain is relatively limited (additional 2% in throughput) and the performance does not critically depend on such choice;

- as far as centralized systems are concerned, the relative gain of TDMA-NC over the other protocols reaches a maximum when at most 3 cooperative slots per frame are allocated. Setting too few slots (2) will reduce the effectiveness of cooperation and NC, whereas at the other edge (4 slots) the performance gain slightly degrades since too much time is devoted to performing retransmissions rather than injecting new traffic.

Finally, it shall be noticed that all the results presented so far have been obtained considering a Poisson traffic pattern, as discussed in Section IV-A. This choice has been made so as to better isolate and highlight the impact of link layer-related effects on the effectiveness of cooperation. Nonetheless, it is reasonable to wonder whether the discussed trends also hold for more realistic traffic models. In order to answer this question, we have performed some simulations that consider bursty transmissions. In particular, we have adopted the traffic model from the UMTS standard [21]. According to this, a traffic source generates bursts of bits. The interarrival times of successive bursts are Poisson distributed, while the number of bits per burst follows a Pareto distribution. Each burst is then fragmented into 4000 bit long packets (the same size as in the Poisson case), that are finally injected into the network. The parameters of this model have been set so that on average 3.75 packets are generated per burst. The outcome

of this simulation study is reported in Fig. 7, that plots against the normalized load λ^* the aggregate throughput gain for Coop and NC schemes over Basic under Poisson and bursty traffic conditions in TDMA- and CSMA-based environments, respectively. As can be noticed, the overall difference between the plotted curves is limited in both scenarios. Indeed, while on the one hand the absolute values of the throughput at low traffic are about 50% larger with bursty than with Poisson arrivals, on the other hand the curves for all systems are shifted upwards by roughly the same amount. Hence, the gain of the studied protocols does not strongly depend on the traffic model, which further supports the generality of the conclusions inferred from our study. There are also two second order effects that deserve to be discussed. In the first place, we notice that at low load the gains of the cooperative schemes under bursty traffic are always slightly smaller than those with Poisson traffic. This effect can be explained observing how, for a given value of aggregate injected traffic, under bursty generation conditions fewer nodes are transmitting on average. Therefore, the interference level is reduced, entailing fewer packet losses and consequently draining cooperative opportunities. Secondly, note that the gains of the NC protocols at high loads with bursty traffic are a few percentage points below those for the Poisson model. The reasons that lead to such a behavior will be discussed in detail in Section IV-C.

C. Simulation Results: metrics on the efficiency of cooperation

The discussion carried out in Section IV-B has shown that TDMA-based cooperative protocols are marked by fundamentally better performance than comparable CSMA-based systems. Such finding has further been supported by the outcome of several parametric studies, which proved how the detected trends do not appreciably vary as the system configuration or the traffic model is modified. In order to get a thorough understanding of the interactions between HARQ and different link layers, however, it is important to explore why cooperation would be more efficient in one environment than in another, and what are the elements that may prevent these techniques from achieving their full potential. These issues are discussed in detail in this section, which reports the
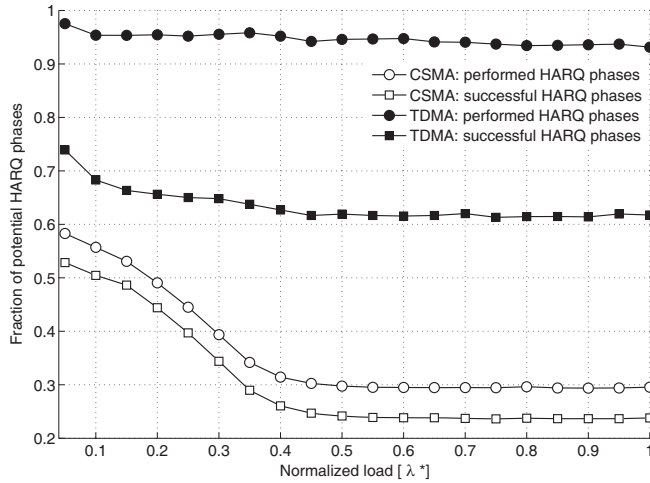
Fig. 8.   Efficiency of HARQ vs. normalized load.

Fig. 9.   Reasons that prevent HARQ when it is needed in a CSMA environment vs. normalized load.

results obtained for both CSMA-NC and TDMA-NC, plotted against the normalized load λ^* as in Section IV-B.

The first question that comes to mind is how efficient cooperation is, i.e., how often a cooperative retransmission (simple or network coded) is carried out given that both source and destination need it, and how often such a retransmission is successful. The answers are reported in Fig. 8. The first apparent observation is that in TDMA a HARQ phase is performed far more often than in CSMA. This is due to the centralized structure of the cell, where one node collects all information and can schedule the necessary relays. In CSMA this knowledge is not available, so that even if nodes that may act as relays are present, they may not be able to do so because of the carrier sense mechanism, as they erroneously believe that other cooperators are acting, or are simply not allowed to access the medium. This is confirmed by the fall of the CSMA curves at high loads: the harsher interference may often fool carrier sense based cooperative mechanisms. Instead, such a behavior is virtually absent in TDMA, as polling packets and allocated cooperative data will always be sent and will not be postponed or canceled because of the medium access policy. It may also be noticed that the success rate in CSMA is higher than in TDMA, and the reason lies in the different way a node may become a cooperator. In our implementation of CSMA-Coop and CSMA-NC, a terminal joins the relay election procedure not only if it correctly received the packet sent by the source, but also if it decodes the NACK sent by the intended receiver. This implies that the $\text{SINR}_{rel-dest}$ to the destination cannot be too weak. On the other hand, TDMA applies the milder requirement $\text{SINR}_{rel-dest} > \Lambda_{coop}$. Thus, on average the relay-destination channel used for HARQ will be worse than in CSMA and cooperation will be somewhat less successful. However, note that the success rate in TDMA cooperative retransmissions is about 2/3, while in CSMA it is about 3/4, hence not dramatically higher.

If all potential cooperative phases were performed, according to the given definitions, the top curve for CSMA and TDMA in Fig. 8 would be equal to one for all loads. The gap between the actual and ideal behaviors in distributed environments is explained in Fig. 9, which illustrates the causes that prevent a cooperative retransmission from being
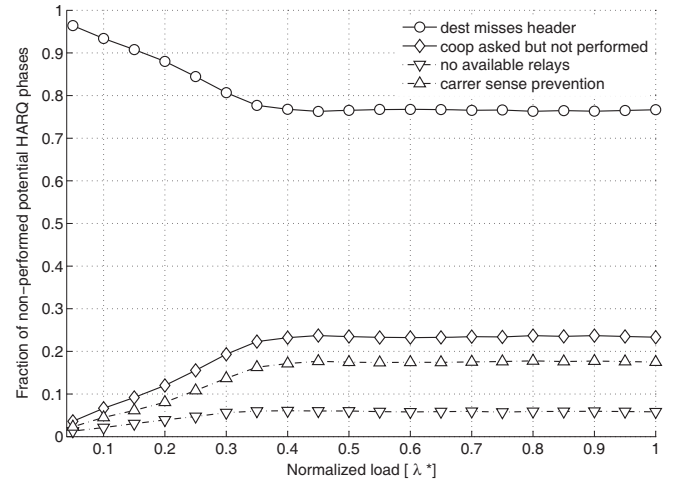
carried out. With CSMA two main reasons may lead to this: i) the destination does not decode the data packet header (and thus does not even request a retransmission), or ii) relaying is requested but not performed. The latter event is in turn due to two problems: ii-a) there may be no available cooperators, or ii-b) there may be such nodes but nobody wins the contention because the carrier sense makes all participants drop. The clear message from Fig. 9 is that the vast majority of the non performed retransmissions is due to the loss of the data header. In this case, the destination is not even aware of the sender's identity, and thus cannot ask for help. This is an intrinsic limitation of CSMA and will be a major difference with respect to scheduled systems. On the other hand, the lack of relay candidates is not particularly relevant, since nodes close to the source are likely to decode the original payload addressed to the destination by virtue of the protection offered by the carrier sense mechanism. However, carrier sense does take its toll during the relay election phase.

While Fig. 9 explored the reasons why cooperation is not performed in CSMA, Fig. 10 analyzes the same phenomenon for TDMA. Three main problems may avert a cooperative phase: i) a signaling packet (i.e., a beacon or a polling packet) is not received, ii) no relays are available, or iii) all the slots reserved for cooperation in the next frame have already been allocated. For high loads, the main cause is the lack of helpers (i.e., no node correctly decoded the lost frame), which is in direct contrast to CSMA, where this problem is negligible in comparison to the header loss or the failure of the election procedure. Instead, at low loads interference is not a major issue and thus most cooperative phases do find a relay. Hence, the main issue may be a lack of sufficient slots reserved to cooperation, but this happens in a negligible fraction of cases (see Fig. 8). In conclusion, the comparison between Fig. 9 and Fig. 10 (for CSMA and TDMA, respectively) suggests that the main reasons for missed cooperation in CSMA and TDMA are the header loss and the lack of relays, respectively.[9]

[9]The parametric studies in the number of nodes per cluster have shown that a larger cell population reduces the probability of not performing a cooperative phase from 7% for 6 nodes/cell (see Fig. 8) to a mere 3% for 12 nodes/cell. This provides additional evidence that the lack of relays is the dominant problem for cooperation in TDMA.
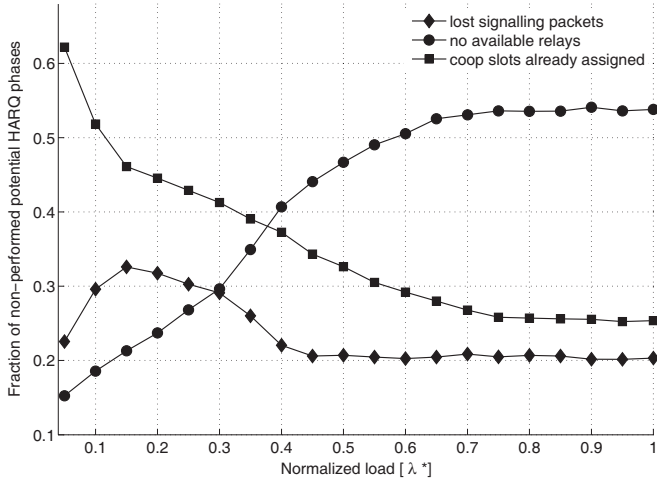
Fig. 10. Reasons that prevent HARQ when it is needed in a TDMA environment vs. normalized load.
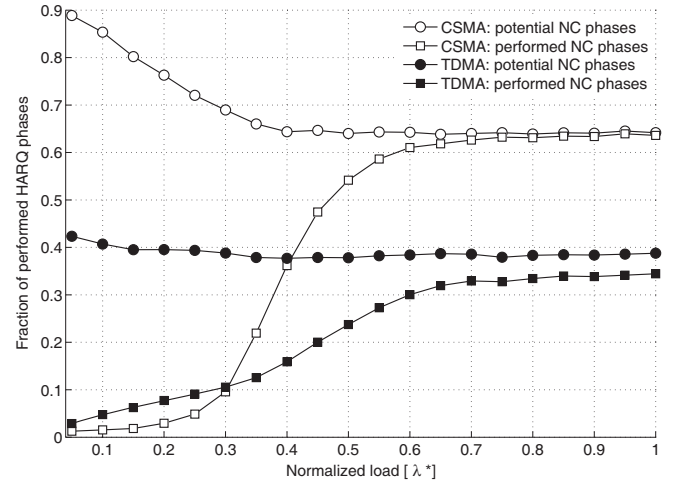

Fig. 11. Efficiency of advanced HARQ techniques vs. normalized load.

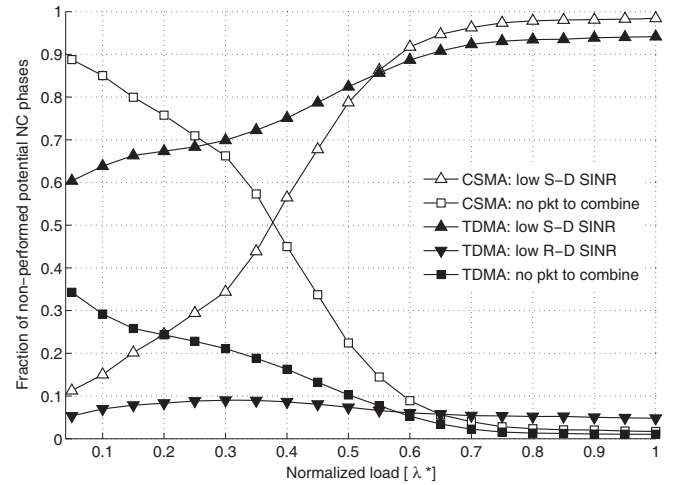Fig. 12. Reasons that prevent advanced HARQ when it would be possible in a CSMA and a TDMA environment vs. normalized load.

Note that each cause is negligible or even impossible in the other scheme, and this highlights the deep differences between CSMA and TDMA towards cooperation.

The previous graphs analyzed the relationship between the two different MAC schemes and cooperative techniques. The last two plots will focus on the efficiency of NC HARQ. The first question is how often NC retransmissions are possible and how many of them are actually carried out. Fig. 11 shows the percentage of potential NC phases and of actually performed NC retransmissions out of all the cooperative stages. It is clear that NC occurs more often in CSMA rather than in TDMA, and this is due to the fact that in CSMA a cooperative phase is requested only if the destination recovered at least the header. Once again, this implies a minimum quality level on the source-destination channel, and increases the probability of meeting the Λ_{Th} requirement. Instead, no such constraint exists in TDMA, and the destination can request cooperation after decoding the beacon but missing the data packet. Such condition makes it possible to ask for cooperation also at lower SINR, where NC is not recommended. Putting Fig. 8 (percentage of actually performed HARQ phases) and Fig. 11 (fraction of NC HARQ) together, it turns out that at high loads $30\% \cdot 65\% = 20\%$ and $90\% \cdot 40\% = 36\%$ of all potential cooperative retransmissions are network coded for CSMA and TDMA respectively. This shows that NC occurs twice more often in TDMA than in CSMA, and this justifies why NC protocols are more effective in TDMA than in CSMA (see Section IV-B).

The last plot of interest is Fig. 12, which explains why plain cooperation is performed instead of a NC retransmission. In TDMA, three main problems may arise: the source-destination SINR is too weak; no relay has a packet to transmit for this destination; or the SINR between relay and destination is also too low. In CSMA this last fact does not happen, since a node can become a relay only if it decodes the destination's NACK, thus putting a lower bound on the transmission quality. At low loads, the lack of packets for the destination is the major impairment towards NC in CSMA, while at high loads (as the interference grows) the low source-destination SINR is the foremost problem. The same phenomenon occurs in TDMA, but the weak source-destination SINR becomes a secondary

problem for this MAC only at loads lighter than those shown in Fig. 12. It is interesting to notice how these trends are partly altered when bursty packet generation is considered (see Section IV-B for a specific description of the employed traffic model). In such conditions, even at high loads, the queues tend to be filled with many packets for the same destination. Therefore, when a cooperative NC opportunity shows up, the potential relays may not have a packet for the destination that requested a retransmission. This effect becomes manifest, as up to 15% of the cooperative phases cannot be turned into NC retransmissions because the relay has no packet for the desired destination in both CSMA and TDMA. Such a trend undermines the efficiency of the NC protocols, leading to the performance degradation shown at saturation in Fig. 7 and presented in Section IV-B.

V. CONCLUSIONS AND FINAL REMARKS

This paper has presented a first exploration of the fundamental contrasts between carrier sense- (CSMA) and time division-based (TDMA) medium access control with respect to HARQ. It has been highlighted that structural differences between these two environments have a deep impact on the efficiency of cooperation and especially of its more advanced

versions, such as network coded retransmissions. The completely distributed nature of CSMA has been shown to severely stymie relaying, as it both often prevents nodes from caching packets for later combining and hinders medium access for cooperators. Conversely, TDMA avoids such issues altogether by resorting to a centralized scheduler, and attains far more noticeable performance gains. The thorough and comparative discussion carried out in this work provides interesting insights for the efficient design of cooperative solutions that properly interact with the underlying medium access control policy.

REFERENCES

[1] J. Laneman, *Cooperation in Wireless Networks: Principles and Applications*. Springer, 2006, ch. Cooperative Diversity: Models, Algorithms, and Architectures, pp. 163–188.

[2] G. Kramer, M. Gastpar, and P. Gupta, "Cooperative strategies and capacity theorems for relay networks," *IEEE Trans. Inf. Theory*, vol. 51, no. 9, pp. 3037–3063, Sep. 2005.

[3] A. Özgür, O. Lévêque, and D. N. C. Tse, "Hierarchical cooperation achieves optimal capacity scaling in ad hoc networks," *IEEE Trans. Inf. Theory*, vol. 53, no. 10, pp. 3549–3572, Oct. 2007.

[4] P. Liu, Z. Tao, S. Narayanan, T. Korakis, and S. Panwar, "CoopMAC: a cooperative MAC for wireless LANs," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 2, pp. 340–354, Feb. 2007.

[5] P. Murphy, A. Sabharwal, and B. Aazhang, "On building a cooperative communication system: testbed implementation and first results," *EURASIP J. Wireless Commun. and Netw,*, 2009.

[6] L. Badia, M. Levorato, F. Librino, and M. Zorzi, "Cooperation techniques for wireless systems from networking perspective," *IEEE Wireless Commun. Mag.*, Apr. 2010.

[7] M. Levorato, A. Munari, and M. Zorzi, "On the effectiveness of cooperation in carrier sense based ad hoc networks," in *IEEE SECON*, June 2009.

[8] P. Popovski and H. Yomo, "Bi-directional amplification of throughput in a wireless multi-hop network," in *IEEE VTC Spring*, May 2006.

[9] S. Katti, S. Gollakota, and D. Katabi, "Embracing wireless interference: analog network coding," in *ACM SIGCOMM*, Aug. 2007.

[10] F. Rossetto and M. Zorzi, "Mixing network coding and cooperation for reliable wireless communications," *IEEE Wireless Commun. Mag.*, vol. 18, no. 1, pp. 15–21, 2011.

[11] A. Munari, F. Rossetto, and M. Zorzi, "Phoenix: making cooperation more efficient through network coding in wireless networks," *IEEE Trans. Wireless Commun.*, vol. 8, no. 10, pp. 5248–5258, Oct. 2009.

[12] ——, "Hybrid cooperative-network coding medium access control for high-rate wireless personal area networks," in *IEEE ICC*, May 2010.

[13] "IEEE LAN MAN standards, part 11: wireless LAN medium access control (MAC) and physical layer (PHY) specifications high-speed physical layer in the 5 GHz band"," ANSI/IEEE, Sep. 1999.

[14] J. Karaoguz, "High-rate wireless personal area networks," *IEEE Commun. Mag.*, pp. 96–102, Dec. 2001.

[15] 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*, vol. 50, no. 12, pp. 3518–3539, Dec. 2004.

[16] K. Kumar and G. Caire, "Coding and decoding for the dynamic decode and forward relay protocol," *IEEE Trans. Inf. Theory*, vol. 55, no. 7, pp. 3186–3205, July 2009.

[17] E. Fasolo, F. Rossetto, and M. Zorzi, "Network coding meets MIMO," in *NetCod*, Jan. 2008.

[18] "IEEE LAN MAN standards, part 15.3: wireless LAN medium access control (MAC) and physical layer (PHY) specifications for high-rate wireless personal area networks (WPANs)," ANSI/IEEE, 2003.

[19] A. Varga, "Omnet++," http://www.omnetpp.org, 2001.

[20] W. C. Jakes, *Microwave Mobile Communications*. IEEE Press, 1993.

[21] "Selection procedures for the choice of radio transmission technologies of the UMTS, TR 101 112, v3.1.0," ETSI, 1997.

Andrea Munari (S'06, M'10) was born in Venice, Italy, in 1981. He received the Laurea degree (M.S. equivalent) and the Ph.D. in Telecommunications Engineering from the University of Padova, Italy, in 2006 and 2010, respectively. In 2007 he joined IBM Research in Zurich, Switzerland, where he spent three years working on the design, analysis and implementation of energy-efficient routing protocols for wireless sensor networks. In 2010 he was a research fellow at the University of Padova, doing research on cooperative medium access strategies for wireless networks, and in 2011 he joined the Corp. R&D division of Qualcomm Inc. in San Diego, California, for a post-doctoral internship focused on network coding techniques for LTE cellular scenarios. His main research interests lie in the design and analysis of protocols capable of effectively supporting advanced techniques in wireless networks, such as cooperative relaying, network coding, interference alignment and directional communications.

Francesco Rossetto (S'06, M'09) received the "laurea" (equivalent to MS) and the PhD in Telecommunications Engineering in 2005 and 2009, respectively, from the University of Padova, Padova, Italy. In 2004-2005 he studied electrical engineering at the University of California, San Diego, CA, USA, under a student exchange program. In 2008 he was on leave at the University of California, San Diego, working for the MURI project, a multiuniversity initiative for the development of multihop MIMO networks. Since 2009 he has been with the DLR (German Aerospace Center) in Munich, Germany. His research interests include satellite communication, network coding and cross layer design. His corporate experience includes a summer internship in 2006 at the Ericsson Eurolabs, in Aachen, Germany, working on Hybrid ARQ for 3G/LTE cellular networks.

Michele Zorzi (S'89, M'95, SM'98, F'07) was born in Venice, Italy, on December 6th, 1966. He received the Laurea Degree and the Ph.D. in Electrical Engineering from the University of Padova, Italy, in 1990 and 1994, respectively. During the Academic Year 1992/93, he was on leave at the University of California, San Diego (UCSD), attending graduate courses and doing research on multiple access in mobile radio networks. In 1993, he joined the faculty of the Dipartimento di Elettronica e Informazione, Politecnico di Milano, Italy. After spending three years with the Center for Wireless Communications at UCSD, in 1998 he joined the School of Engineering of the University of Ferrara, Italy, where he became a Professor in 2000. Since November 2003, he has been on the faculty at the Information Engineering Department of the University of Padova. His present research interests include performance evaluation in mobile communications systems, random access in mobile radio networks, ad hoc and sensor networks, energy constrained communications protocols, broadband wireless access and underwater acoustic communications and networking.

Dr. Zorzi was the Editor-In-Chief of the *IEEE Wireless Communications Magazine* from 2003 to 2005, is currently the Editor-In-Chief of the IEEE TRANSACTIONS ON COMMUNICATIONS, and serves on the Editorial Board of the *Wiley Journal of Wireless Communications and Mobile Computing*. He was also guest editor for special issues in the *IEEE Personal Communications Magazine* (Energy Management in Personal Communications Systems) and the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS (Multimedia Network Radios, and Underwater Wireless Communications Networks).

He currently serves as a Member-at-large of the Board of Governors of the IEEE Communications Society.