In Proc. of 10th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (IEEE WoWMoM 2009), Kos, Greece, June 2009

A Framework for Interference Mitigation in Multi-BSS 802.11 Wireless LANs

Murad Abusubaih, Berthold Rathke, and Adam Wolisz*
Technical University Berlin
Telecommunication Networks Group (TKN)
Einsteinufer 25, 10587 Berlin, Germany
abusubaih@tkn.tu-berlin.de

Abstract

This paper proposes a framework for interference mitigation in multi-BSS infrastructure 802.11 WLANs. Our interference mitigation approach is based on Access Point (AP) Coordination. With this approach, interfering BSSs negotiate and switch from the 802.11 CSMA/CA to a time slotted mechanism if users' QoS is observed to be degraded, diagnoses conclude that the cause is high interference, and the switch to the time slotted modus is expected to be useful and feasible. The proposed algorithms within the framework are driven by measurements. We utilize the wireless bandwidth and improve the fairness among WLAN users. We present results of detailed simulation experiments as well as real implementation.

1. Introduction

Due to the diminishing costs of wireless devices, Wireless Local Area Networks (WLANs) have been massively deployed in public places such as: university campuses, offices, apartments, airports and hotels. In alignment with the growth of WLANs, users demands are also increasing and their satisfaction becomes a challenging task for both network designers and administrators.

In multi-BSS infrastructure WLANs, each Access Point (AP) is usually assigned a fixed channel. As in all communication systems, the 802.11 spectrum is a scarce resource. The number of supported channels by any IEEE 802.11 standard is limited and among all channels, only few of them do not overlap. WLAN administrators try to improve the coverage of their premises by deploying a high density of APs. However, the dense deployment of APs can introduce additional mutual interference unless the network is carefully planned and tuned.

In current 802.11 WLANs, channel access is governed by the CSMA/CA mechanism. Although this mechanism is robust within a single BSS, it fails to provide acceptable service for many WLAN users in an Extended Service Set (ESS) when the traffic load gets high. As the traffic load increases, interference among neighboring BSSs increases, leading to collisions and retransmissions, which in turn add to the load and consequently to more collisions.

In this paper, we propose a framework to combat interference in infrastructure 802.11 WLANs. With the help of the stations (STAs) they accommodate, WLAN APs that operate in a CSMA/CA modus over the same channel monitor the QoS in their BSSs. They negotiate and switch from the 802.11 CSMA/CA to a time slotted mechanism if users' QoS is observed to be degraded, diagnoses conclude that the cause is high interference, and the switch to the time slotted modus is expected to be useful and feasible. When operation conditions improve, BSSs negotiate and switch back to the CSMA/CA modus.

The paper is organized as follows: Section 2 discusses background and relevant work. In section 3 we present our framework and discuss its components. Our proposed methods for interference estimation are presented in section 4. Section 5 discusses potential rules for switching the access mode. Section 6 details the slot allocation criterion we used within the proposed framework. In section 7 we evaluate the performance of the framework before we conclude the paper in section 8.

2. Background and Relevant Work

2.1. A brief overview of 802.11 MAC

The 802.11 MAC DCF protocol is based on CSMA/CA. The CSMA/CA works as follows: A node wishing to transmit a data packet first has to sense the medium, and, if no activity is detected, the node waits a randomly selected additional period of time before it transmits if the medium is still free. If the receiving node receives the packet intact, it issues an ACK frame to confirm the reception of a data packet. The ACK frame completes the process if successfully received by the sender. The sender assumes a collision to have occurred

1. ***Professor Adam Wolisz is also with the University of California, Berkeley, BWRC.**

978-1-4244-4439-7/09/$25.00

if the ACK frame is not successfully received. In this case, the data packet is transmitted again after deferring another random amount of time.

2.2. Interference Mitigation

In this work we constrain ourselves to the case of non-overlapping channels. Interference will denote hereafter a phenomenon where signals transmitted from one BSS spread to a neighboring BSS that operate over the same channel. While a node in a BSS is receiving a frame, a coincident in time signal from a neighboring BSS may corrupt the frame under reception if the interfering signal has comparable strength relative to the signal strength of the frame being received. This is known as the *Hidden Node Problem*. It leads to collisions and errors which will cause discards and retransmissions. Similarly, signals from neighboring BSSs on the same channel can prevent local nodes from transmitting their frames, even if intended receivers might not be within an interference region. This is known as the *Exposed Node Problem*.

The 802.11 standard provides the Request to Send/Clear to Send (RTS/CTS) mechanism to reduce interference. However, this mechanism does not alway help due to the following shortcomings: First, RTS/CTS is not efficient enough in case of overlapping BSSs. The main design assumption with RTS/CTS is that all nodes within sender and receiver vicinity will hear the RTS or CTS packets and set their Network Allocation Vector (NAV) accordingly. This assumption does not necessarily hold in multiple BSS deployments, whereby a node(s) may be busy receiving a frame generated within its BSS and therefore will not get the RTS or CTS sent in a neighboring BSS. Second, RTS/CTS introduces considerable overhead and may unnecessarily decrease the communication efficiency (see [1]). Third, the RTS/CTS does not solve the exposed node problem, under which possibly successful transmissions are inhibited.

Channel assignment policies have been usually proposed in the literature to combat interference in WLANs. However, solutions addressing exclusively channel selection have a limited improvement potential especially under high load. Therefore, we expect to achieve additional performance improvement if multiple interfering BSSs coordinate their transmissions.

2.3. Coordinated Channel Access

Time slotted access is another option for channel access in 802.11 WLANs. The 802.11e standard (enhanced to support QoS with multimedia) coordinates channel access within a BSS. Nevertheless, the standard does not address the problem of overlapping BSSs/cells that use the same channel. There is no mechanism beyond CSMA/CA to coordinate the channel access across BSSs, thereby there is no guarantee that during the transmission of a frame by some STA in a time slot other STAs belong to neighboring BSSs will remain silent. This is due to the fact that BSSs operate asynchronously and independently.

The authors of [2] and [3] attempted to develop a time slotted access scheme for 802.11 networks. Nonetheless, their main concern has been to solve implementation challenges of a time slotted approach with 802.11 adapters in small testbeds. Hence, the interference mitigation problem was not directly addressed.

The work of Bejerano et. al. [4] presents a managed WiFi system to support QoS in 802.11 WLANs with multiple BSSs. It uses Inter-AP coordination to allow overlapping BSSs coordinate their operation during up-link transmissions of the Point Coordination Function (PCF) modus so as to improve fairness among STAs. The presented solution proposes to assign disjoint time slots to BSSs that interfere with each other, whereby during a time slot assigned to one BSS other interfering BSSs remain silent (i.e Blocked). The length of the time slot that each BSS gets depends on the number of users the BSS accommodates. Although the solution has shown improvement, still it has some drawbacks: First, the authors assume a purely attenuation-based propagation model which is not the case in practice due to fading. Second, the PCF modus is not supported by most IEEE 802.11-compliant products. Third, the BSS-based scheduling does not efficiently utilize the wireless bandwidth since it does not exploit exposed nodes within interfering BSSs which can simultaneously send their packets.

Recently, there has been a significant amount of research activities in the area of wireless mesh and sensor networking, aiming for network performance enhancement through channel access coordination [5], [6], [7]. While we are following the same general ideas of scheduling transmissions, our work differs from the foregoing efforts in that we are aiming at development of a holistic framework for interference mitigation, covering interference estimation and switching between a CSMA/CA and a time slotted access schemes depending on interference conditions. We also consider a different approach for solving the scheduling problem, the features of which are discussed latter in this paper.

3. A Framework for Interference Mitigation

3.1. System Model

We consider an ESS 802.11 WLAN (see figure 1) composed of N APs and M stationary STAs, ($M \geq N$). APs are

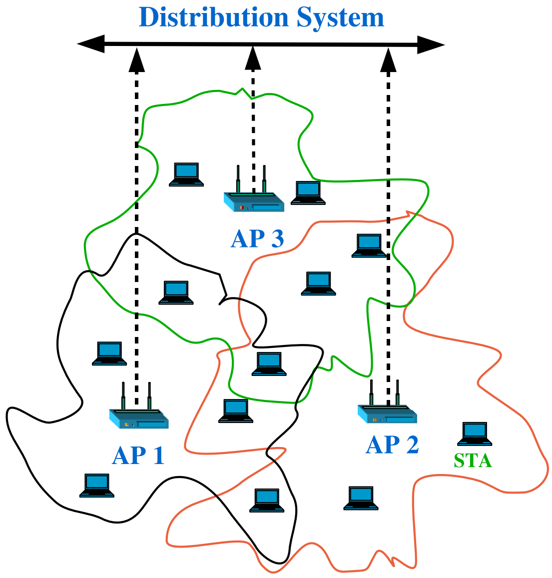

Figure 1. Assumed Network Model

assumed to operate on non-overlapping channels. Nonetheless, extensions to the case of partially overlapping channels is easily possible in the framework. Some APs are assigned the same channel. APs are connected to a single distribution system (DS). They provide communication services to the M STAs that reside within their unknown a priori coverage area. The coverage areas of APs are assumed to overlap. Neither the location of an AP nor its operational channel is known to the other APs. At any time instant, a STA is associated exactly to one AP.

3.2. Solution Idea

We exploit the efficiency of a temporal separation approach to mitigate interference in multi-BSS 802.11 WLANs. The 802.11 CSMA/CA channel access scheme provides best effort service. It is easy to implement, does not need synchronization among contending nodes, and works well at low traffic load. At increased traffic levels, frequent collisions, contention, and retransmissions due to interference occur, degrading the QoS the wireless users experience. On the other hand, a collision-free channel access scheme, such as a time slotted access scheme, is known to perform better than the CSMA/CA at high traffic loads despite the signaling overhead it adds [13], [4]. We combine the CSMA/CA and a time slotted channel access schemes, suggesting their alternative usage depending on operation conditions. BSSs switch from the CSMA/CA access mechanism to the time slotted mechanism only if high interference is detected and the operation is a time slotted modus is expected to be useful and feasible. A switch-back to the CSMA/CA modus takes place when operation conditions are observed to improve.

3.3. Framework Description

An architectural block diagram for each AP is shown in Figure 2. The basic system components are: **Interference Conditions Estimator, Channel Access Scheme Selector, a Slot Scheduler, and a Coordination Protocol.** In this section and the following ones, we elaborate our design principles of the various system components.

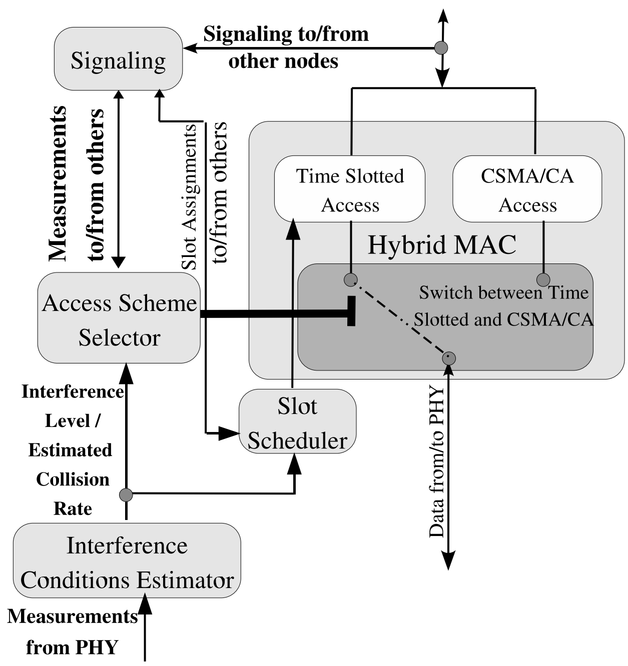

Figure 2. Architectural Block Diagram

3.3.1. Interference Conditions Estimator.

The Interference Conditions Estimator resides at each AP. It processes AP's local ("own") observations and interference measurement information reported by STAs and produces an estimate of the interference in the BSS. The measurement information will then be used as input to the access scheme selector as well as the slot scheduler as will be described in the following subsections. To improve the accuracy of estimating interference conditions and determining interfering links, we propose to use two complementary approaches:

- Packet Loss Diagnosis.
- Passive interference estimation using packet decoding.

For the sake of organization, details on interference recognition and estimation are separately provided in section 4.

3.3.2. Access Scheme Selector.
The access scheme selector is mainly responsible for selecting the proper access scheme to be employed within the BSSs. Also, it has to determine the scope of BSSs which shall employ a channel access mode. Currently, we are developing a method to achieve this objective. The key idea is to first construct a conflict graph using interference estimations among the BSSs. Then, from the conflict graph, the set of loosely coupled or independent clusters of BSSs are identified using clustering techniques, particularly the density-based class which are originally developed to recognize dense areas within an object space.

The decision on the mode to be used within BSSs is based on:

- Access Mode Switch Rules.
- Observations and diagnosis of the QoS degradation reported by local interference conditions estimator and measurements signaled from other APs.

Potential rules for access mode selection are discussed in section 5.

3.3.3. Slot Scheduler.
Basically, slot scheduler is an algorithm for assigning disjoint time slots to all links within a group of BSSs which have been selected for potentially simultaneous switching to the time slotted modus. The input to the algorithm is interference matrix among the links within the group. The scheduling algorithm should find out the set of transmissions/links that can go in parallel without collision. This becomes extremely important as the number of STAs and cooperating APs increases. In this case, the sequential assignment of time slots (i.e. the assignment of one long time slot to each participating BSS as done in [4]) becomes not possible since other BSSs cannot be blocked (wait) for long time. Section 6 elaborates on slot assignment algorithms.

3.3.4. Coordination Protocol.
The interference mitigation approach we propose in this paper involves two types of signaling. The complete specification of signaling mechanisms is a future work. In this paper, we assume that information exchange works perfectly, i.e. we do not consider errors in data exchange for the sake of coordination. Here, we just generally outline the signaling requirements and challenges for the described system operation.

The first signaling will be needed for ***the passing of interference measurements from STAs to their respective APs***. STAs report to their respective APs: the measured interference level, the collision rate estimated as described in [11], and the identity of nodes from which interference in coming.

The second type signalling will be needed for:

- The sharing of interference measurements among coordinating APs.
- The distribution of access scheme selector decisions.
- The distribution of slots allocation results.
- The decision on the scope of nodes within which an operation mode (CSMA/CA or slotted time) shall be used since the usage of a mode can not happen for an arbitrary subset of BSSs.
- Achieving reliable mode switching (i.e. assuring that all nodes switch operation mode at the same time.
- Synchronization.

Primarily, the information includes: Access scheme change messages, Interference measurement information, and Slot assignment results once a change to the time slotted access mode is decided. Interference measurement information includes the amount of estimated interference in the BSS and the set of interferers for each node in the BSS. Slot assignment results include the identity of nodes that can access the channel at the beginning of each time slot.

3.3.5. Synchronization.
The coordinated channel access operation modus basically defines an orderly access to the jointly shared wireless channel. In such a time slotted access schemes clock synchronization is needed. It is well known that STAs in a WLAN BSS are synchronized with their AP via the beacon frames transmitted periodically by APs. This is completely specified in the IEEE 802.11 standard. A STA that receives a beacon frame from its AP adjusts its local clock considering the potential propagation and processing delays. However, APs operate independently from each other and time offsets may appear in different BSSs.

Within the proposed framework, we require to synchronize each group of coordinating APs to a certain degree of accuracy to avoid slot overlap due to timing inaccuracies and the different propagation delays of individual nodes. Practically, several approaches can be used to achieve the required synchronization among coordinating APs. For instance, a technique similar to the one already used in Independent Basic Service Set (IBSS) may be used. With IBSS, STAs transmit a special time-stamped frame like the beacon frame. Any STA receives this frame has to update its clock if the time stamp in the received frame is latter. Neighboring APs may use either the DS or the STAs that are located in the overlapping areas to send synchronization frames. An AP that receives a frame with a time stamp latter to its time clock has to perform the appropriate adjustments to its local clock. One could also use the well known Network Time Protocol (NTP) to achieve clock synchronization. Another option would be the usage of control channels, similar to those proposed for HiperLAN.

4. Interference Measurement

4.1. Methods for Determining Interference Relations

In fact, there is no way for measuring the amount of interference while a node is receiving a signal. Therefore, studies in the literature follow two different approaches to infer the effect of interference while assessing the performance of wireless communication systems:

- Using simplified models for interference approximation.
- Performing active interference measurements that are shifted in time.

Two models proposed by Gupta and Kumar [8] are being widely used:

1) **The Simple Interference Model:**
With this model, the euclidean distance between wireless nodes is used to infer wether a transmission can be correctly received or not. Particularly, a receiver N_j is assumed to successfully receive a frame from a transmitter N_i , if and only if there is no other simultaneous sender within a guard zone, determined by a factor d, from the receiver N_j. In equation form, a transmission from node N_i to node N_j is successful if, for every other node N_k,

$$|N_k - N_j| > (1 + d)|N_i - N_j| \quad (1)$$

The factor d models the radius of the guard zone and specified by a protocol. Hence, the model is referred to as the protocol interference model. The main characteristic of the model is that it only accounts for the path loss as the source of signal attenuation and applies under same transmit power levels.

2) **The Physical Interference Model:**
This model predicts that a transmission can be successful if the signal to interference ratio SINR exceeds some threshold. Specifically, a sender i transmits a frame successfully to receiver j, if and only if:

$$\frac{P_i}{P_t + P_n} > SINR_{TH} \quad (2)$$

where P_i is the power level of the signal received from i at node j, P_t is the total power received by receiver j from other potential simultaneous senders, P_n is the noise power level. $SINR_{TH}$ is the threshold value necessary for a successful decoding of i's transmission at receiver j.

Obviously, the physical interference model is less restrictive than the simple model. With this model, it may happen that a packet is successfully received by a receiver, even if there is another node located within the interference range of this receiver is simultaneously transmitting. Additionally, the model is more related to physical layer effects as it consider attenuation sources like fading other than the path loss.

While the above interference models simplify the calculation of interference, their use in realistic networks has been shown to be erroneous [9]. A second approach for determining interference relations among links is through measurements that are shifted in time. Interference measurement can be performed actively during the deployment phase or passively while the network is operating. The core of active interference measurement approaches is the measurement of throughput or signal strength [10]. With the throughput-based approach, two links i and j are assumed to interfere iff the throughput of one degrades when the other is active. The determination of interfering links takes place in the dedicated configuration phase and the start of network operation. With the signal strength based approach, each node sends in turn a series of broadcast packets. All other nodes measure the signal level of the received packets. The signal strength is used to indicate the potential interference level from the transmitting node to each other node. Such active measurements delivers, however only an estimate of the real interference. This is due to the following: ***First,*** *the signal strength varies in time, dependent on environmental changes. Hence, initial measurements are not valid all the time.* ***Second,*** *the estimated interference is usually valid in the scenario used to estimate it. Due to the variable nature of traffic, the potential interference is not observed all the time. Additionally, protocol based dependencies on the node state (transmitting, receiving) change the dynamic pattern of the real interference.*

4.2. Suggested Interference Estimation Approach

In this section, we develop a passive measurement-based approach for interference relations determination as well as interference level estimation. We determine interference relations among links and estimate interference level at a node through measurements conducted while the network is operating. This trend is advocated by standardization bodies which develop mechanisms to facilitate measurements during network operation(e.g the 802.11k standard [14])

In fact, the real impact of interference depends both on the interference signal level and the frequency of the interference event. The latter is strongly dependent on the traffic profile. While we confine our attention to the Received Channel Power Indicator (RCPI), recently standardized in 802.11k [14], other signal level indicators such as RSSI can be used if the RCPI measure is not supported. As an IEEE 802.11 standard feature, the RSSI is defined in the standard as a measure by the Physical Layer (PHY) of the power level observed at the antenna used to receive the current Physical Layer Protocol Data Unit (PPDU) at the receiver antenna during packet reception, measured during the PLCP (Physical Layer Convergence Protocol) of an arriving packet. In contrast, the RCPI value is measured over the entire frame at the antenna connector used to receive that frame. Hence, the RCPI value seems to be a better metric to represent the signal power level of a received packet.

The following method for determining interference relations and estimating interference conditions will be considered. We first use a method for determining interference conditions using ***Packet Loss Discrimination.*** For the sake of space limitation, we just give an overview of this method. More details can be found in [11]. The idea is to compare the power level of a corrupted packet (measured over the entire packet length, e.g using RCPI) with a quantile value acquired from power levels of correctly received packets. Packet corruption is attributed to collision if the power level of this packet is higher than the quantile value, otherwise the loss is attributed to fading. ***In [11], we also use the packet loss discrimination method to produce an estimate of total interference at the node side. Nonetheless, the method does not tell us from which interferers this estimated interference is coming ?***
To achieve this latter goal, we use ***passive observation of interference.*** This approach works as follows:

- An AP requests the STAs it accommodates to monitor the wireless medium for a period of time T.
- During the measurement period, a measuring STA monitors all transmitted frames over the medium and records the following information elements: The number of transmitted frames from each source address, the length of each frame, the rate at which each frame was transmitted, and the power level at which each frame is received.
- Since frames have different lengths and can be transmitted using different physical rates, an interference metric has to account for these facts. A STA k captures the interference level from a source address as follows:

$$InterferenceLevel_k = \frac{1}{T}\sum_{i=1}^{N}\frac{L_i P_i}{R_i} \qquad (3)$$

  where L_i and P_i denote the length in bits and received power level in dBm of frame i, respectively. P_i is captured from RCPI or RSSI. R_i denotes the physical rate in bits/second at which frame i is received, and T denotes the length of the measurement period.
- Each measuring STA k sends the measurement information to its AP. ***From this report, the set of potential interferers for each STA as well as an estimate of interference level that each STA experiences can be identified.***
- The duration of the measurement is fundamental. This period should be as small as possible to reduce the time a STA spends listening to the channel but large enough to assure that transmissions from interferers fall within the measurement time and consequently improve the accuracy of estimation. A value of 50ms has been used in the evaluations. We think that this is a reasonable value in order not to harm real time applications (e.g. VoIP).
- A STA periodically conducts this measurement and post the measurement report to its AP. In order to reduce the measurement overhead, the time between successive measurements shall be adaptively set based on differences between measurement results which obviously depends on traffic patterns.
- Similarly, an AP measures interference level coming from nodes that belong to neighboring BSSs.

We make the following discussion on the above interference estimation approach:

1) It does not only consider the power levels of transmissions from interfering nodes, but also the duration of these transmissions. This is important since the probability of collision due to interfering transmissions and the collision cost depend on the time period collided packets occupy the medium.
2) By considering the time of each frame and dividing over the whole measurement duration, we capture the activity level of an interferer.
3) **Improving Interferers Identification:**
   The passive observation approach can lead to identification of interfering transmitters only if the interfering packet is captured and decoded, providing the source address. In order to assess whether it was able to recognize potential interferers or not, a measuring node uses the two estimates of interference conditions from the two methods, packet loss discrimination and the passive observation of interference. Note that the packet loss discrimination method provides an estimate of interference from all interferers (i.e. total interference) while the interference estimate computed with the passive observation method is just part of the total interference, because it is just computed over the set of interferers whom the measuring node was able to decode their packets. The two estimates are compared. If the interference estimate computed with

the passive observation method is small compared to the total interference estimated with the packet loss discrimination approach, a measuring node concludes that it was not possible to identify all interferers. In this case, the node shall either repeat the passive measurement and again try to identify interferers or consider a different method (e.g. actively sending broadcast probe packets at increased power level and reduced physical rate to assure higher coverage, or utilize the estimations of other nodes in the BSS for inferring unrecognized interferers).

5. Access Scheme Selection Rules

The general reasons for changing the operation modus are:

- *a) Potential improvement of users' satisfaction level.*
- *b) Improvement of the channel usage efficiency in terms of the ratio of number of packets successfully ACKed from the first transmission to the total number of transmissions.*

In this work, we only focus on the users' satisfaction level in terms of goodput and try to improve that while designing a mode switching rule by considering the user load and the MAC ability to deliver this load.

Intuitively, the crucial question here is whether retransmissions, due to interference, block incoming frames from upper layers from being transmitted or not. When a frame arrives from upper layers at the sender MAC, the MAC can either pass this frame to the physical layer for transmission or send a signal to the upper layer if it is still busy transmitting or retransmitting a previous frame. Let us consider an observation period of length T. Denote the number of times the MAC layer sends a busy signal after receiving a frame from upper layers during T as N_r. Also, denote the total number of frames that arrived from upper layers at the MAC layer during T as N_t. Note that N_r depends on the number of retransmissions, i.e. the MAC would have attempted to send some of the N_r frames during the time period spent in retransmissions. If the ratio N_r / N_t is relatively small, then retransmissions do not introduce much disturbance (i.e. they do not block data frames coming from upper layers from being transmitted). In this case, there will be no need to switch to the time slotted modus. However, if the ratio N_r / N_t is high and receivers indicate an increase in collision rate, then the MAC is not able to pass frames due to retransmissions and a switch to the time slotted modus might be useful, depending on load volume and scheduling feasibility. This seems to be a candidate rule to infer if a switch **to the time slotted modus** will be useful for users in a BSS. For the switch-back to the CSMA/CA, the traffic load is a good candidate.

6. A Heuristic Slot Assignment Algorithm

In [12], we have developed an optimal algorithm for time slots allocation. Despite that the optimal algorithm provides optimal time slots assignment, it is rather computationally expensive, especially for large number of STAs and APs. Hence, the question about a heuristic algorithm is relevant. This section proposes such algorithm. The main design objectives are:

- Maximizing the number of active links in each slot.
- Minimizing number of required slots.
- Achieving some fairness in terms of number of slots during which a link is active.

The slot assignment problem has to be solved for each channel. Denote $L(i,j)$ as the communication link from transmitter i to receiver j. Basically, two links $L(i,j)$ and $L(m,n)$ can simultaneously be active (i.e. within the same time slot) iff: i does not interfere n, m does not interfere j, j does not interfere m, and n does not interfere i. The first two constraints are for data packets protection and the latter two ones are for ACKs protection. Therefore, what do we really need to know is the set of interferers for each node.

The scheduler starts with the link that experiences the highest interference and finds out all links that can run parallel to it. This set of links are marked as **done** and should be assigned a time slot. Then, it proceeds with the next link and again finds out all links that can receive in parallel with it starting with those that are not marked as **done** yet. The algorithm proceeds until all links are marked as **done**. The algorithm is shown in Algorithm 1. It has the following features which differentiate it from other algorithms proposed in the literature:

1) Due to delay constraints, the algorithm sets an upper bound on the number of time slots to be used in scheduling. Note that a node can not be blocked from accessing the channel for a long time. This may happen when the number of interfering nodes gets large. So, we extend other scheduling algorithms by allowing some (probably) small interference between links whenever it is impossible to schedule all links without exceeding a pre-defined maximum SlotCount threshold.
2) In order to minimize the number of needed slots and the search time, the algorithm first sorts the set of links in descending order according the interference each measures.
3) Note that step (10) achieves the objective of maximizing the number of links that use a slot, while in the meanwhile it tries also to improve fairness by considering the ones that already got minimal slots as first candidates.

Algorithm 1 Heuristic Slot Assignment

INPUT: ***S***= {Set of all links};
OUTPUT: The scheduled links in each slot;
Initialization: SlotCount=0; **Done**={};
Sort ***S*** (descending) according the interference level;
MAX : a maximum upper bound on the number of slots that can be allocated;
Repeat {
Select the Next link l from ***S*** **AND** $\notin$ **Done**.
Find the set of links ***K*** $\subset$ ***S*** that can be active parallel to each other and to link l **AND** $\notin$ **Done**.
Done= Done U l U ***K***.
Find the subset ***T*** $\subset$ **Done** that can be active parallel to each other and to l and every link $n \in$ ***K***, starting with those that occupy less slots.
Assign SlotCount to links l U ***K*** U ***T***
SlotCount = SlotCount + 1
if (SlotCount $>$ MAX) distribute all remaining links among the slots in a way that keeps interference among scheduled links in each slot minimal.
} **Until** all links $\in$ **Done**

7. Performance Evaluation of a Framework Instance

In this section we assess the performance of the interference mitigation framework developed in this paper. We have conducted a number of simulation experiments using the NCTUns simulation package [15]. The MAC layer goodput is used as a first metric to be observed. Every STA and AP measures it during a time interval of one second. Additionally, we use Jain's fairness index [16] to capture the fairness level among WLAN users. The slot assignment algorithm, interference estimation algorithms are fully implemented in the simulation, while the signaling protocol for the exchange of information among STAs and their respective APs and among the APs themselves is not. We simply make the measurement information accessible to APs. On the other hand, we also implemented the heuristic slot assignment (Algorithm 1) of section 6 on top of the 802.11 MAC. Additionally, we have implemented the coordinated channel access in a small infrastructure WLAN of two APs and five STAs.

7.1. Coordinated Channel Access

7.1.1. Simulation Setup.

The scenario is composed of 4 BSSs and 75 stationary STAs deployed as shown in figure 3, where the distance between two adjacent APs was 200 meters APs operate over the same channel. All nodes implement the 802.11b technology. STAs are randomly uniformly distributed in the coverage

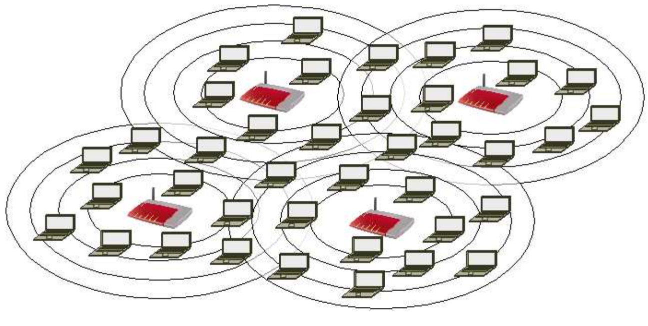

Figure 3. Network Topology as used in the experiments

area of the APs. At the physical layer, we have used a two ray ground reflection path loss model. The received power is further influenced by Rayleigh fading. A Rayleigh fading model provided by the NCTUns simulator is used. It takes as parameters the received power P_{rx} and a fading variance set to its default value of 10dB. The received power level of a packet (with respect to both path loss and fading attenuations) is passed to an error module provided by the simulator along with packet length and modulation type. This module determines whether a received packet is correct or corrupted due to fading and path loss attenuation. A sender selects a physical transmission rate based on the distance d to the receiver and the rate remains fixed during the simulation time (i.e no rate adaptation is used). Table 1 lists values of other parameters as used in simulations. Over

Parameter	Value	Parameter	Value
PLCP header T_H	48 μs	T_{SIFS}	10 μs
PLCP preamble T_P	144 μs	T_{DIFS}	50 μs
Tx Power	100 mW	T_{Slot}	20 μs
$W_{\max}$	1023	$W_{\min}$	31
d $\leq$ 40	11Mbps	40 $<$ d $\leq$ 80	5.5M
80 $<$ d $\leq$ 120	2Mbps	d $>$ 120	1M

Table 1. Constant Parameters

a measurement period of 50ms, a STA monitors the wireless channel, it computes the interference level as described in section 4.2. An AP is identified as interferer to a STA if the measured interference level from that AP is greater than a cutoff value of -83dBm. Throughout this study, the length of a slot in the time slotted modus was selected to be 15ms and the maximum number of slots was set to 15. These constant values were selected after conducting intensive simulation experiments.

7.1.2. Traffic Model.

In a first experiment, each user downloads infinite number of UDP packets from a server via its AP. The interval between two successive packets is drawn from an exponential distribution with 10ms mean, while all packets are of same size chosen to be 1500 Bytes. In a second experiment, each user downloads UDP packets for 300 seconds using the traffic profile provided in table 2. It starts with a low load phase, followed by a high low phase and then back to low load. Since the interference depends on users' workload, we also tested our solution using realistic WLAN traffic traces provided in [17], wherein it has been shown that these traffic measurements are consistent with an analysis of SIGCOMM conference traces. We used the realistic WLAN traces in the following way: Using CoralReef Software [18], we extracted users flows from the dump file. We selected the flows of 75 different users during 10 minutes. We used the total number of bytes, number of packets of a flow to characterize a user load and compute an average packet length. Then, the **stg** tool which comes with the NCTUns simulation package was used to emulate users' flows.

Simulation Time	Offered Load (Pkt/s)	Pkt Size (B)
0 - 100	10	1500
101 - 200	200	1500
201 - 300	10	1500

Table 2. Traffic Profile

7.1.3. Simulation Results with Synthetic Traffic.

a) Effect of Coordinated Channel Access on MAC Goodput:

For different load levels, figure 4 shows the aggregate MAC goodput experienced by users when the network just employs CSMA/CA and when it employs coordinated channel access. From these results, we draw the following

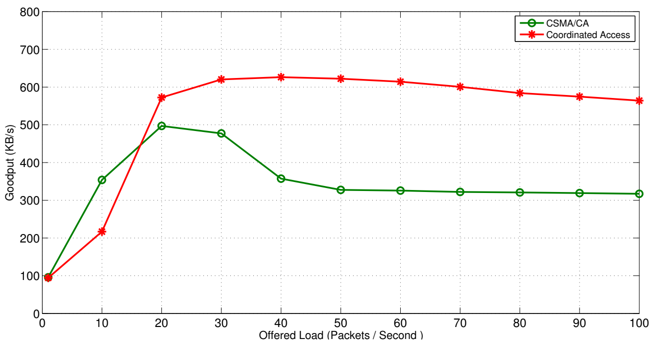

Figure 4. Aggregate Goodput experienced by all users with CSMA/CA and with coordinated access for different load levels.

observations: **(i)** At high load, the aggregate goodput has been improved if APs coordinate channel access. Note that the goodput starts to degrade again at extremely high load conditions. We attribute this to the allocation of same time slot to some interfering STAs, where the impact of this interference starts to be harmful at very high loading. **(ii)** However, coordination degrades the goodput when the load becomes low. This is because we employed fixed slot assignment during our experiments, meaning that a slot is wasted if slot owner(s) has no data to send at the beginning of this time slot. Additionally, the probability of collisions with low load is lower and the CSMA/CA MAC can handle corrupted frames through retransmissions between successive arriving frames.

b) Tracking high interference conditions:

Now we run the simulation with the traffic profile of table 2 (subsection 7.1.2). In this experiment, APs use the rules described in section (5) for deciding on the operation modus. The observation period was set to 5 seconds. After operating in the time slotted modus for 20 seconds, APs change back to CSMA/CA and again decide on the operation modus to be employed. Figure 5(a) plots the aggregate goodput for two cases, where in the first case APs just use the CSMA/CA for channel access while in the second case they employ the time slotted channel access during high interference periods. The figure shows that, the aggregate goodput has been improved when APs coordinate channel access during the high load period. Further, figure 5(b) plots Jain's fairness level [16] among the 75 users, which also indicates a gain in fairness level among users as a result of coordinated channel access during the high load period.

7.1.4. Simulation Results with Realistic Traffic.

We repeated the experiment described in 7.1.3(b), but with realistic WLAN traffic. Figure 6 shows the results with this experiment. As with synthetic traffic, the results show that coordination has a significant positive impact on the aggregate goodput during high interference periods. Spikes seen in the aggregate throughput are due to bursty traffic and unsaturated channel conditions.

7.2. Coordinated Channel Sharing - Real Experiments

In this experiment, we would like to observe the total system throughput and how this throughput is distributed among the STAs with and without coordinated channel access in a realistic network.

7.2.1. Experiment Setup.

The experiment set-up is shown in figure 7. Two APs and five stationary STAs were deployed in two different LABs. The APs are WLAN adapters from Atheros configured in the master mode (AP mode) through the MADWIFI driver. The APs are connected via an Ethernet Switch. Over the ethernet connection, a master program runs on one AP synchronizes both APs. APs are assigned the same channel. Through

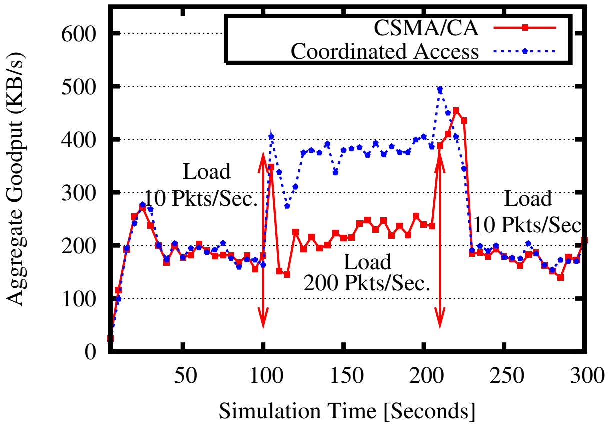

(a) Aggregate Goodput

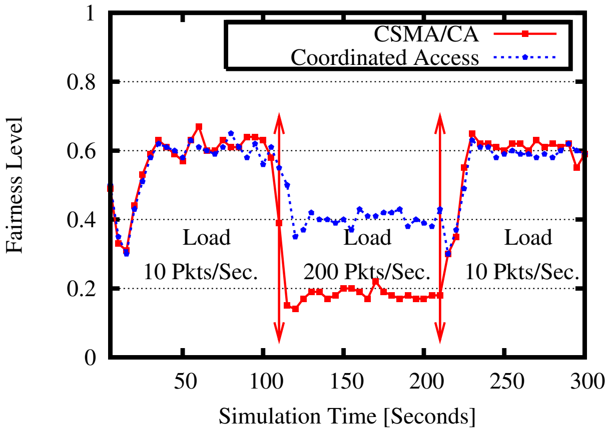

(b) Jain's Fairness Level

Figure 5. Aggregate Goodput and Fairness with CSMA/CA and Coordinated Access during Different Load Conditions

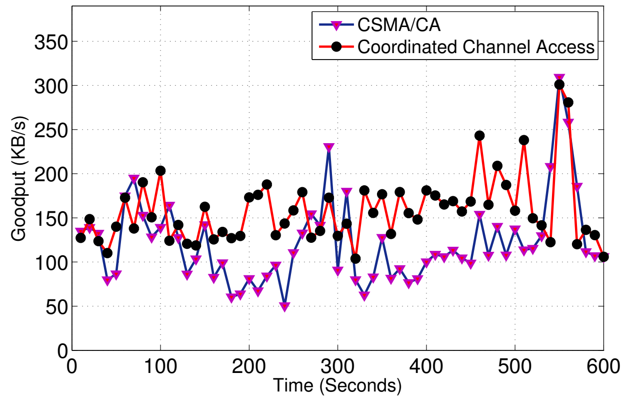

Figure 6. Aggregate Goodput of a 4 APs WLAN of 75 Users with Realistic Traffic

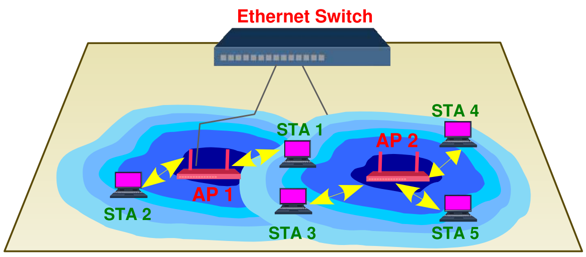

Figure 7. Topology used in Real Implementation

transmit power control, APs are hidden from each other. Two STAs are deployed in overlapping area of the two BSSs. APs transmit UDP traffic to the five STAs.

7.2.2. Experiment Results.

In this experiment, the five STAs are scheduled as shown in table 3. Figure 8 plots the results of the real experiments.

Slot	Stations
T1	STA 1
T2	STA 3
T3	STA 4, STA 2
T4	STA 5, STA 2

Table 3. Scheduling of the five STAs

We make the following comments on both figures: **(A)** The total throughput with CSMA/CA and with coordinated channel access is comparable. **(B)** With CSMA/CA, STA 1 (in the overlapping area) experiences degraded performance compared to other STAs due to increased collisions. **(C)** Although STA 2 is outside the interference region of AP 2, it also experiences degraded performance with CSMA/CA due to the time its AP (AP 1) spends retransmitting packets to STA 1. This means, in fact, that the whole BSS of AP 1 suffers communication problems. On the other hand, STA's 3 performance is not degraded with CSMA/CA despite it is located within the interference region of AP 1. By observing the received power at both STAs in the overlapping region, we found that the reason is the capture effect which helps STA 3 to maintain good performance. **(D)** With coordinated access, STA's 2 throughput is higher than other STAs as it is scheduled in two time slots.**(E)** In fact, the experiments have shown two main results: The first is the ability of coordination to improve system performance under high loading. With almost the same aggregate throughput, coordinated channel access was able to relief two users and consequently improve the fairness among WLAN users. Specifically, the Jain's fairness index has increased from 0.72 to 0.98 for the case of coordinated channel access. The second result is the necessity of driving the whole adaptation process by measurements.

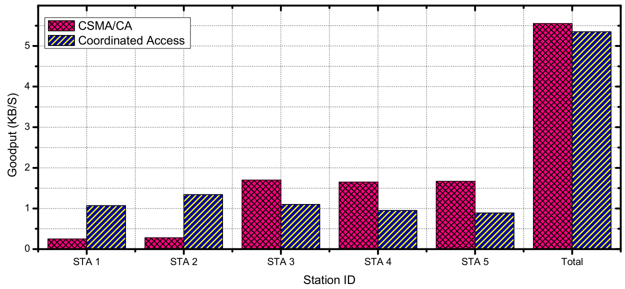

Figure 8. Real Implementation Results

8. Conclusions and Ongoing Work

This paper proposes a framework for interference mitigation in infrastructure WLANs. Neighboring interfering APs negotiate, exchange interference information and agree to switch between a CSMA/CA and a time slotted channel access schemes for delivering packets to their users. Observations and measurements of interference conditions drive the decision on the access scheme to be employed. Detailed simulations and real implementations have shown a good potential of the proposed approach in terms of aggregate goodput and fairness among WLAN users. In our future work, we will focus on the coordination protocol challenges discussed in section 3.3.4.

Acknowledgment

This work was partially supported by the German Academic Exchange Service (DAAD, http://www.daad.de) and the European Commission through project EU-MESH (Enhanced, Ubiquitous, and Dependable Broadband Access Using MESH Networks), FP7 ICT-215320, http://www.eu-mesh.eu.

References

[1] Saikat Ray, and David Starobinski. On False Blocking in RTS/CTS - Based Multihop Wireless Networks. *IEEE Transactions on Vehicular Technology*, March, 2007.

[2] A. Sharma, M. Tiwari and H. Zheng. MadMAC: Building a Reconfiguration Radio Testbed using Commodity 802.11 Hardware, *Proceedings of the IEEE Workshop on Networking Technologies for Software Defined Radio Networks*, September, 2006.

[3] A. Rao and I. Stoica. An overlay MAC layer for 802.11 networks, *Proceedings of the 3rd international conference on Mobile systems, applications, and services*, 2005.

[4] Yigal Bejerano and Randeep S. Bhatia. MiFi: A Framework for Fairness and QoS Assurance for Current IEEE 802.11 Networks With Multiple Access Points, *IEEE/ACM Transactions on Networking, vol. 14, No. 4, pp. 849-862*, August, 2006.

[5] V. Ramamurthi, A. Reaz, S. Dixit, and B. Mukherjee. Link Scheduling and Power Control in Wireless Mesh Networks with Directional Antennas, *Proceedings of IEEE ICC 2008.*

[6] Leonardo Badia and Alessandro Erta,.A General Interference-Aware Framework for Joint Routing and Link Scheduling in Wireless Mesh Networks, *IEEE Network*, Vol. 22, pp. 32-38, 2008.

[7] D. Chafekar, D. Levin, V. S. Anil Kumar, M. V. Marathe, S. Parthasarathy, and A. Srinivasan. Capacity of Asynchronous Random-Access Scheduling in Wireless Networks, *Proceedings of IEEE INFOCOM'08*, April, 2008.

[8] P. Gupta and P. R. Kumar. The Capacity of Wireless Networks *In IEEE Transactions on Information Theory*, 46(2), 2000.

[9] D. Kotz, C. Newport, R. S. Gray, J. Liu, Y. Yuan, and C. Elliott. Experimental evaluation of wireless simulation assumptions *In Proceedings of MSWiM*, October, 2004.

[10] A. Kashyap, S. Ganguly, and S. Das. A measurement-based approach to modeling link capacity in 802.11 based wireless networks, *In Proceedings of ACM MOBICOM*, 2007.

[11] Murad Abusubaih, Berthold Rathke, and Adam Wolisz. Packet Loss Discrimination in Multi-Cell 802.11 Wireless LANs, *Technical Report TKN-08-010, Technical University Berlin*, October, 2008.

[12] Murad Abusubaih, Berthold Rathke, Daniel Hollos, and Adam Wolisz. A Framework for Interference Mitigation in Multi-Cell 802.11 Wireless LANs, *Technical Report TKN-08-011, Technical University Berlin*,November, 2008.

[13] James F. Kurose and Keith W. Ross. Computer Networking: A Top-Down Approach (4th Edition), *Addison Wesley*, April, 2007.

[14] *IEEE 802.11k*, Radio Resource Measurement, IEEE Draft 2.0,2005.

[15] http://nsl.csie.nctu.edu.tw/nctuns.html.

[16] D. Chiu and R. Jain Analysis of the Increase and Decrease Algorithms for Congestion Avoidance in Computer Networks *Journal of Computer Networks and ISDN Systems*, vol. 17,Nr.1, pp. 1-14, June, 1989.

[17] http://www.winlab.rutgers.edu/ ergin/mobicom2007/

[18] http://www.caida.org/tools/measurement/coralreef/