

Dynamic Switching with Heterogeneous Channels in Multichannel 802.11 WLANs

Juan J. Galvez

Coordinated Science Lab
University of Illinois at Urbana-Champaign
Email: jjgalvez@illinois.edu

Nitin Vaidya

Department of ECE and Coordinated Science Lab
University of Illinois at Urbana-Champaign
Email: nhv@illinois.edu

Abstract—In a multichannel wireless system, the performance of a link depends on many factors. The link quality is subject to temporal, spatial and spectral diversity, i.e. the SNR is time varying, link-dependent and channel-dependent. In addition, the performance also depends on MAC dynamics and the degree of congestion present in the channel. As a result, different links can experience different performance on the same channel, and the performance of a link varies across channels and time. An effective way to exploit and cope with the diversity in the wireless system is to use opportunistic channel switching. This technique allows a link to dynamically search for a channel/spectrum where it can maximize its performance at a given point of time. In addition, we observe that the link diversity can make it beneficial to have channels configured with different PHY/MAC parameters (e.g. different transmit power, data rates, or carrier sensing threshold). We refer to these as *heterogeneous channels*. A group of links may perform better under a set of parameters while a different group may perform better under a different set. In this paper we combine an opportunistic channel switching scheme with heterogeneous channels in multichannel Wireless LANs (WLANs) and show that the combined approach is effective in increasing link throughput and fairness.

I. INTRODUCTION

The recent explosion of wireless devices and traffic is putting increasing demands on existing WLANs. Today's enterprise networks use a dense deployment of Access Points (APs) in an attempt to increase capacity [1]. In many cases, however, interference between devices, limited spatial reuse and inefficiencies of the Distributed Coordination Function (DCF) of 802.11 make it difficult to scale throughput. There is an ongoing need to improve the performance of WLANs to satisfy increasing traffic demands [2], [3].

The use of multiple frequency channels is a strategy which has been effectively used in the past to increase capacity. In many real-world scenarios, this is achieved by having APs or cells transmit in different channels and separating them in space [1], [4]. However, this strategy proves to be limited, because nodes in the same cell cannot transmit concurrently, and the load between APs need not be the same in general, which will result in load imbalance between channels.

In addition, it is generally assumed that system parameters (e.g. transmission power, backoff window size, carrier sense threshold) are the same in different channels. In practice, this need not be the case. A benefit can be achieved by having channels with different PHY/MAC parameters, referred to

in this paper as *heterogeneous channels*. Moreover, a link can perform differently in different channels even when the channels have the same parameters and load, due to temporal and channel diversity, the dynamic nature of the topology and the operation of DCF. For example: back-off, retransmissions, hidden and exposed nodes, and differences in transmission time between links are factors that can affect performance [5]. This dynamic aspect can be unpredictable and difficult to model.

In this paper we assume that multiple channels are co-located in space (e.g. APs are equipped with multiple radios, or multiple APs are in the same area using different channels). Based on this, we propose an online approach with heterogeneous channels where links dynamically search for the best channel, and which ultimately is capable of increasing capacity and fairness. The solution proposed uses channels with different carrier sensing thresholds. In this way, links with a high signal strength or low interference can transmit in channels with less aggressive carrier sensing and exploit spatial reuse, while low throughput links or those susceptible to hidden nodes can transmit on channels with more aggressive carrier sensing.

The rest of the paper is organized as follows. Section II reviews related work. In section III we describe the system model and opportunistic channel switching problem with heterogeneous channels. Based on this, section IV describes the specific solution we have developed. Section V presents simulation results in *ns-3*. Finally, section VI concludes the paper.

II. RELATED WORK

Opportunistic channel switching has been used previously to exploit channel diversity. In many cases, the diversity is understood as time-varying fluctuations which cause each link to experience a different time-varying SNR in each channel. Examples of these works are [6], [7], [8]. The idea is that the transmitter and/or receiver periodically obtain channel quality information in order to dynamically switch to the channel with better instantaneous conditions. Kanodia et al. [6] proposed a protocol to exploit this kind of diversity. In [8], the authors study if it is possible to exploit this diversity in practical conditions, concluding that to do so requires fast channel sensing and low overhead in the control channel.

In the work we present here we don't restrict ourselves to this kind of diversity involving fast changing channel

conditions. Instead, we consider channels with heterogeneous parameters (which may remain with this configuration for long periods of time), and where varying conditions are induced not only by fading but also by the current traffic and interaction between devices.

In [9], Wang et al. consider opportunistic channel switching as a way for links to dynamically search for channels with better quality. In particular, their scheme focuses on achieving equal-channel occupancy (i.e. having the same number of active links in each channel). It does not consider heterogeneous parameter channels nor does it account for other types of diversity between channels like SNR or MAC unfairness.

Many studies have considered the problem of channel allocation in multichannel systems (e.g. [10], [11]). The problem consists in allocating a channel for each link, generally with the goal of reducing interference and maximizing network capacity. The assignment can be static or dynamic and is intended to be used for a long period of time (minutes or hours). These works assume that channels are homogeneous. In [12], the authors consider a channel assignment scheme which separates high-rate links from low-rate links by assigning them to different channels. In this way, they induce a specific channel diversity with the purpose of avoiding the performance anomaly of multirate 802.11 networks.

There have been many studies concerning the improvement of carrier sensing in 802.11 networks (e.g. [13], [14], [15], [16]). The goal is to maximize spatial reuse while attempting to minimize the hidden node problem. These studies however are for single channel networks or assume that all channels follow the same carrier sensing scheme.

In this paper we propose the use of heterogeneous channels in multichannel systems, specifically through heterogeneous carrier sensing, which has not been proposed earlier to the best of our knowledge. Furthermore, the channel switching scheme we propose requires only minimal information and does not make any assumptions on the characteristics or behavior of each channel.

III. OPPORTUNISTIC CHANNEL SWITCHING PROBLEM

In this section we describe the system model and the channel switching problem with heterogeneous channels.

A WLAN is composed of wireless stations and APs. Each station associates with an AP, typically a nearby one based on signal strength. A station can send packets to its AP (uplink) or receive packets from the AP (downlink). Fig. 1 shows an example of a network with three APs.

Devices use the 802.11 standard for communication (e.g. 802.11a). Let C be the set of non-overlapping frequency channels available for transmission, and let R be the number of radios in a station. Stations can be equipped with one or two radios. If a station has only one radio it will use it for both the uplink (transmission) and downlink (reception). If it has two radios one will be used for the uplink and the other for the downlink. Two radios allow a station to transmit and receive on different channels simultaneously. We assume that APs can use multiple channels at the same time. This can be achieved in a single device with multiple radios or emulated by having multiple co-located APs each tuned to a different channel. The

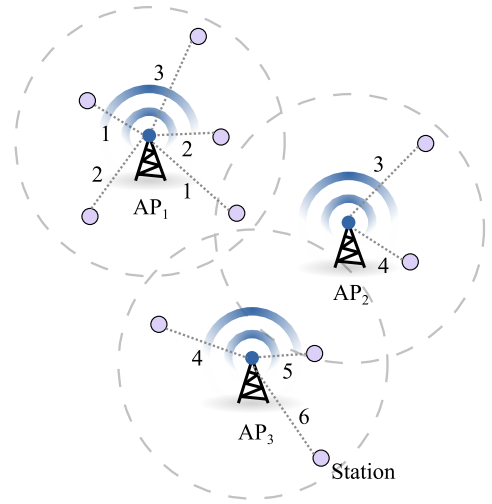


Fig. 1. Multichannel WLAN with three APs. Each station is associated with an AP. There are two links between a station and its AP (uplink and downlink). This example shows the use of six channels, with each AP using a subset of three channels. The labels on the links indicate the channel used by the link.

radio interfaces have multiple modulation and coding schemes at their disposal, which allow them to transmit using one of several link rates. For example, the 802.11a standard supports eight different data rates (from 6 to 54 Mbps).

The throughput a wireless link achieves with 802.11 depends on a variety of factors, which include its data rate, how frequently it can access the channel, and the losses it experiences. The links in a WLAN are heterogeneous: they have different SNR, transmission rate and interference tolerance. Also, the degree of interference suffered, contention and number of retransmissions varies between links. A link can suffer from hidden and exposed nodes and this too varies from link to link. To complicate matters, many of these factors depend on the current topology, traffic and channel, and can change dynamically. Even if channels have the same PHY/MAC parameters, links cannot be expected to perform in the same way in each.

The above phenomena may cause a link to behave badly in a particular channel for a number of reasons, e.g. it might take up too many resources, transmit too slow, too often, retransmit frequently, or unnecessarily prevent other links from transmitting (exposed node problem). It is also possible for a link to starve, due to frequent collisions or hidden nodes.

In addition, based on the difference between links, it makes sense to have *heterogeneous channels* which allow links to perform in a way that maximizes utility. There are many parameters at the PHY and MAC layers that can be varied between channels, including the transmission power, back-off intervals, carrier sensing threshold, or the set of accepted data rates. Links which behave badly in a channel can be switched to a channel with different parameters in order to maximize utility.

An opportunistic channel switching strategy will allow links to dynamically search for optimal channels to cope with channel diversity. Note that what we propose is not simply a case of load balancing; channel diversity and inefficiencies in DCF can produce suboptimal performance even when load

is balanced between channels. In the opportunistic switching problem, links dynamically choose channels in a way that maximizes aggregate utility¹. Ultimately, this means maximizing the average flow throughput while avoiding unfairness and starvation. As we will see, this can be achieved with a protocol that requires only minimal information about the channels. In the next section we describe our solution.

IV. OPPORTUNISTIC CHANNEL SWITCHING PROTOCOL WITH HETEROGENEOUS CARRIER SENSING

In this section we propose the use of channels with heterogeneous carrier sensing threshold, combined with an opportunistic channel switching protocol.

A. Heterogeneous carrier sensing

The 802.11 DCF is based on Carrier Sense Multiple Access (CSMA), which uses carrier sensing in order to avoid collisions between devices, but is well known to result in suboptimal scheduling [15]. If the sensing range is too high, links which can transmit in parallel may be prevented from doing so (exposed node problem). On the other hand, if the sensing range is too low, nodes might not be able to sense conflicting transmissions (hidden nodes) and produce collisions, which means that more sensitive links will starve. Clearly, there are differences between links. Some links will benefit from an increased sensing range and protection from collisions, while others will benefit from reduced sensing range and more transmission opportunities.

The use of channels with heterogeneous carrier sensing can help take advantage of this link diversity. Let the carrier sensing range (CSr) denote an estimate of the distance from which it is possible to detect transmissions of other nodes. Let the multiset \mathcal{S} whose elements are carrier sensing ranges and $|\mathcal{S}| = |C|$ be the set of carrier sensing ranges used by the channels in the system. For example, in a two channel system we might have a configuration of $\mathcal{S} = \{600, 600\}$, which indicates that both channels 1 and 2 have a CSr of 600 m. As we will see in section V, it is possible that a configuration $\mathcal{S} = \{600, 300\}$ gives better performance than one with $\mathcal{S} = \{600, 600\}$.

The variety of links and the dynamic nature of the topology and traffic means that there can be a spectrum of optimal sensing ranges \mathcal{S} at a given time. Determining the best combination of sensing threshold for the channels is outside the scope of this paper. In our experiments we have obtained results under many different combinations of thresholds.

B. Opportunistic channel switching protocol (O-CS)

We develop an online channel switching strategy (which we will refer to as O-CS) capable of exploiting channel diversity in general. O-CS requires only minimal information that is readily available at the APs, i.e. the current throughput of flows. Therefore, O-CS is not restricted to the specific case of heterogeneous carrier sensing. As we will see, it works with any combination of carrier sensing, and also proves beneficial when the channels have the same sensing thresholds.

¹In the experimental results we will measure utility using the proportional fairness metric [17].

TABLE I. CHANNEL SWITCHING PROTOCOL SYMBOLS

C	Set of channels available for transmission
R	Number of radios in stations
α_l^c	Number of transmission attempts of link l on channel c
$\alpha_{l,r}^c$	Number of transmission attempts of link l on channel c at rate r
Φ_l^c	Failure rate of link l on channel c
$\Phi_{l,r}^c$	Failure rate of link l on channel c at rate r
x_l	Current channel assigned to link l
r_l^*	Target rate for link l
$T_{l,c}$	Time link l has been assigned to channel c

Before describing the protocol we define two metrics. The *packet throughput* of link l is the throughput when the link has gained access to a channel and transmits a single packet. More specifically:

$$PKT_THR = \frac{pSucc \times packetSize}{T} \quad (1)$$

where $pSucc$ is the probability of successfully transmitting one packet and T is the time required to transmit a packet (depends on the modulation scheme used).

The *link throughput* is the average throughput achieved by a link over a larger period of time (multiple packets) in a channel, and accounts for time sharing between links. Table I lists the symbols used to describe the protocol.

The protocol aims to improve the throughput and fairness of active links in the WLAN. It has two main components:

- *Distributed* component, whereby every node dynamically searches for a channel on which it can achieve high PKT_THR (per-packet throughput) for its outgoing links.
- *Centralized* component, whereby a central controller dynamically balances load between channels based on the observed performance in each channel, and prevents unfairness and starvation.

Devices are assumed to collect local transmission statistics for each of their outgoing links. These are shown in Table I. Note that this is typically done by existing rate control algorithms (e.g. minstrel [18]). APs keep track of the set of active links and their throughput, and pass this information to the controller. All traffic (downlink and uplink) passes through the APs, so this information is readily known.

In the following subsections we explain O-CS in detail.

1) *Distributed channel switching*: Every link l has a target rate r_l^* which is the rate that it should use for transmission. The goal is for the target rate to be as high as possible, thus allowing a high PKT_THR (per-packet throughput). The distributed phase is detailed in Algorithm 1.

As we can see, if a channel is not valid for the current target rate r_l^* , the transmitter must look for an alternative channel (lines 2-7). If no valid channel is found, the target rate is decreased one step (line 7). The idea is to try other channels before using lower data rates.

Algorithm 1 Protocol periodically executed by every node n

```

1: for  $l \in outgoingLinks(n)$  do
2:   if not  $validChannel(l, x_l, r_l^*)$  then
3:      $vc \leftarrow \{c \in C \mid validChannel(l, c, r_l^*)\}$ 
4:     if  $vc \neq \emptyset$  then
5:        $x_l \leftarrow$  Choose a channel from  $vc$ 
6:     else // no valid channels at  $r_l^*$ 
7:        $r_l^* \leftarrow r_l^* - 1$ 

```

Algorithm 2 Decide if channel c is valid for link l at rate r

```

1: function  $validChannel(l, c, r)$  do
2:   if  $\alpha_l^c > A$  then
3:     if  $\Phi_l^c > F$  then
4:       return False
5:     else if  $\alpha_{l,r}^c > A$  and  $\Phi_{l,r}^c > F_t$  then
6:       return False
7:     else if  $\alpha_{l,0}^c > A$  and  $\Phi_{l,0}^c > F_0$  then
8:       return False
9:     else
10:      return True
11:   else
12:     return True

```

APs implicitly determine when a station's uplink radio has changed channel when they receive packets via the new channel². If an AP decides to switch the channel of a downlink and this requires switching the channel of an interface of the affected station, the AP communicates this to the station via a channel switch message before switching. The APs will communicate a link's current channel to the controller.

Algorithm 2 shows the rules followed locally to determine whether a channel is valid or not. They are based on the number of transmission attempts and failure rate on the channel. The parameters A , F , F_t , F_0 serve as threshold values. Table III shows the values used for these parameters in the simulation tests.

It is important to note that the distributed channel switching algorithm works in conjunction with existing rate control algorithms like minstrel. The rate control algorithm is responsible for adapting the actual link rate, while r_l^* indicates the current target rate, which need not be the same as the current rate being used by the rate control algorithm. Initially r_l^* can be set to the highest achievable data rate of the link assuming no interference from other sources. If this is unknown, the highest data rate offered by the 802.11 standard can be used.

2) *Centralized channel switching*: The distributed channel switching component does not balance load between channels nor does it guarantee that nodes get fair access to the channel. There can be links with high per-packet throughput but low average throughput if they don't get sufficient access to the channel. DCF can even unnecessarily starve some links. The job of the controller is to improve the median throughput and fairness of the links currently active in the network. To do this, the controller periodically moves a low-throughput link from

the channel with worst median throughput to the channel with best median throughput. This is shown in Algorithm 3.

Algorithm 3 Centralized control loop

```

1:  $wC \leftarrow$  channel with lowest median throughput
2:  $bC \leftarrow$  channel with highest median throughput
3: for  $l \in activeLinks(wC)$  do // in ascending order of link throughput
4:   if  $validChannel(l, bC, r_l^*)$  and  $T_{l,wC} > \Delta$  then
5:      $x_l \leftarrow bC$ 
6:   return

```

To know if a channel is valid or not for a given link, the controller obtains this information from the source periodically. It can also deduce that a channel is not valid when it observes a link abandon its current channel. A change in channel assignment is communicated via control messages (see [19] for an example). To prevent oscillations and frequent channel switching, only links which have been in their current channel for a period of time Δ can be switched.

It is important to note that the centralized component, in conjunction with the distributed component of O-CS, has the capability of dynamically exploiting the channel diversity and heterogeneity of the channels. For example, channels with lower carrier sensing range have higher potential spatial reuse and therefore capacity, but not all links will perform well in them due to hidden nodes. The distributed component will only allow links which perform well to stay on such a channel, while the centralized component will detect the higher capacity of the channel (higher median throughput of flows) and move more links to it.

As we can see, O-CS does not need to know the particular PHY/MAC parameters of each channel. It adaptively finds the best channel for a link, i.e. a channel where it can perform efficiently and access sufficient channel resources.

V. SIMULATION EXPERIMENTS

In this section we evaluate the O-CS protocol in 802.11 WLANs through simulation in *ns-3* [20].

A. Simulation environment

The experiments are conducted in WLANs with two APs. We have generated 10 random topologies. In each topology, there are two APs separated by 350 m, and 20 stations located randomly around each AP, within a radius of 200 m of their AP³. For the purpose of establishing Internet traffic flows there is a server on the Internet connected to the APs by a point-to-point 1 Gbps link (in this way the performance observed is only limited by the WLAN). Traffic can be both uplink and downlink. The parameters used to configure wireless interfaces and propagation in *ns-3* are shown in Table II.

To set the carrier sensing threshold of radio interfaces, we will use combinations from a set of CSr values (expressed in meters). Note that although carrier sensing is set as an energy threshold (in dBm), we will express values in meters for easier visualization. This range is translated to the equivalent received

²In this system APs use multiple channels in parallel (see section III) and don't switch the channels of their interfaces. Stations choose among the channels used by their AP.

³There are a total of 80 links in each network, counting uplinks and downlinks.

TABLE II. SIMULATION PARAMETERS

Wifi standard	802.11a
Frequency band	5 GHz
Set of data rates	{6, 9, 12, 18, 24, 36, 48, 54} Mbps
Path loss model	Log-distance, exponent = 2.7
Tx power	18.0 dBm
RTS/CTS	Disabled
Rate control	minstrel
CSr values	{600, 500, 400, 300, 200, 100} m

TABLE III. O-CS PARAMETERS USED IN SIMULATIONS

A	10
F	0.4
F_t	0.33
F_0	0.2
Δ	500 ms

energy level according to the path loss model used in the simulations. The set of possible CSr values is shown in Table II. To generate heterogeneous channels, we use all combinations of k -multisets \mathcal{S} from the possible CSr values, where $k = |C|$. For example, with $|C| = 2$ a valid configuration (2-multiset) is $\mathcal{S} = \{400, 400\}$ and both channels would be assigned the same CSr of 400 m. Another valid configuration is $\mathcal{S} = \{600, 300\}$, where one channel is assigned 600 m and the other 300 m.

In the experiments we generate random traffic flows and test them with UDP and TCP connections. Given a topology, we generate 10 different traffic patterns. In each pattern, 40 flows are chosen randomly (a flow can be either downlink or uplink). Each traffic pattern is tested using UDP and TCP. Flows start at a random instant in 30 ± 0.5 seconds and transmit continually for 10 seconds. For the case of UDP flows, the sending rate is chosen such that it can saturate the network.

The data rate of each link is dynamically controlled during the simulation by the minstrel algorithm (implemented in *ns-3*). Every device executes the minstrel rate control algorithm on each of its radio interfaces.

B. Channel switching strategies

In these experiments we evaluate two different strategies: equal occupancy (EO) and O-CS. As the name suggests, EO implements an equal occupancy policy, where the goal is that each channel have the same number of active links. Channel switching only occurs when the number of links in each channel is different. In other words, once the load is balanced, links will remain in a channel. EO is the policy used in [9]. We have implemented both strategies in *ns-3*. The balancing component of both EO and O-CS is executed in a centralized controller connected to the APs; this is executed every 500 ms. Table III shows the values used to configure the O-CS protocol. These are not necessarily optimal, but rather values which have given good experimental results.

C. Performance metrics

We use the following metrics to study protocol performance: (i) *Flow throughput* - the number of bytes successfully delivered divided by the duration of the flow; (ii) *Aggregate*

flow utility - measure of the proportional fair throughput of the network, given by:

$$\sum_{f \in \mathcal{F}} \log(\text{throughput}(f)) \quad (2)$$

where \mathcal{F} is the set of flows in the network. Proportional fairness [17] is a useful metric because it allows to compare in one metric the overall network performance in terms of minimum flow rate, average rate and aggregate flow rate. It takes into account the total flow throughput, fairness and starvation. In this way, we determine the best strategy to be the one that maximizes the utility.

Regarding flow throughput, in each scenario we measure the minimum, median and mean throughput of flows. Both the aggregate flow utility and minimum throughput are important metrics to detect unfairness and flow starvation.

Points shown in graphs represent the average (mean) result of the scenarios that fall in that class. Confidence intervals are shown at the 95% level.

In the following C refers to the number of channels used.

D. Analysis of O-CS with heterogeneous carrier sensing

In this section we evaluate O-CS and EO under a variety of heterogeneous channel configurations \mathcal{S} , as described in section V-A. In the following figures, CSr combinations are shown in a concise form. For example, “4/3” means $\mathcal{S} = \{400, 300\}$. The total number of CSr combinations with $C = 2$ is 21, and the total number with $C = 3$ is 56.

First we show in detail the effect of CSr on various performance metrics. For this we will focus on the case with $C = 2$ channels and UDP flows. The results are shown in Fig. 2. In this test $R = 2$, i.e. stations have two radios. This option allows the most flexibility in channel allocation (if needed, a station can transmit and receive simultaneously in different channels). Later we compare with $R = 1$. The bottom axis shows the 21 different CSr combinations. Note that it extends from a low enough value $\{100, 100\}$ where nodes won't be able to sense most transmissions in the network, to a high enough value $\{600, 600\}$ that allows a node to hear most transmissions.

As we can see, the CSr chosen for the channels has an important effect on the performance metrics in both strategies. As a general trend with both protocols, starting with the lowest CSr values (left of x axis), performance starts to increase with higher CSr up to a point where it then starts to decrease. This is expected, as too low carrier sensing range suffers from many collisions while too high a sensing range will suffer from low spatial reuse. Note that stations can also cope with the hidden node problem through rate adaptation and opportunistic channel switching, making it undesirable to use to use a very high sensing range on any channel.

More importantly, both protocols react differently to the use of heterogeneous channels. Fig. 2 (a) shows that the best overall performance (aggregate utility) of O-CS is achieved with $\{400, 300\}$, proving the advantage of using heterogeneous channels. In particular, compared to its closest performing homogeneous configurations, there is an average improvement

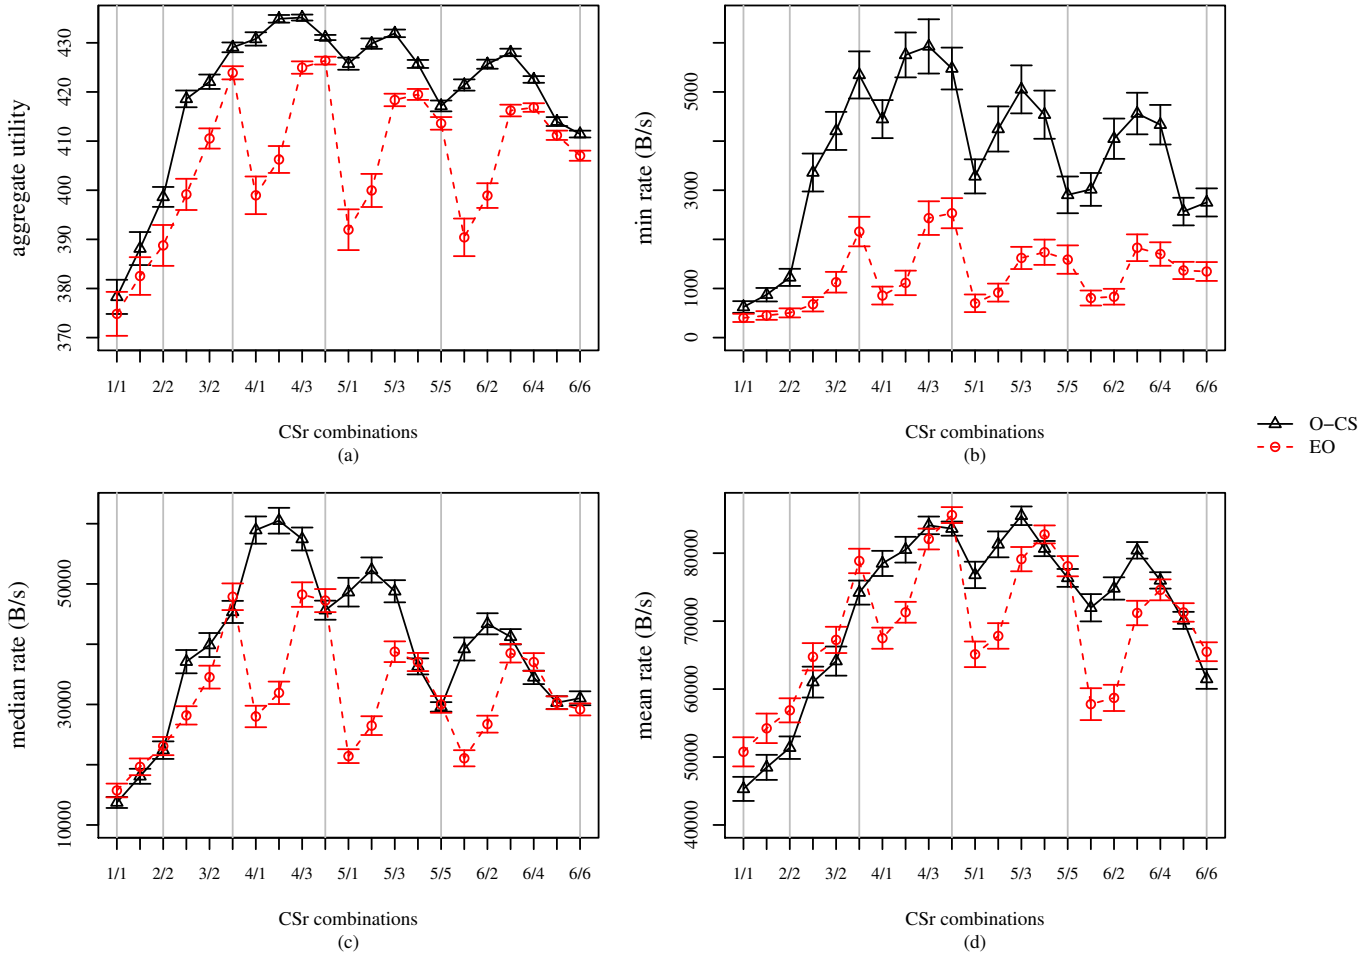


Fig. 2. Performance comparison of O-CS and EO with $C = 2$ and UDP flows and $R = 2$. Shown are the average values of aggregate utility, minimum, median and mean flow rates under different CSr combinations, which include both heterogeneous and homogeneous configurations. A configuration of “1/1” means CSr combination $\mathcal{S} = \{100, 100\}$, “3/1” is $\mathcal{S} = \{300, 100\}$ and so on. There is a vertical line on each homogeneous configuration, i.e. 1/1, 2/2, 3/3, 4/4, 5/5 and 6/6.

of 11% in minimum throughput over $\{300, 300\}$ and 8% over $\{400, 400\}$; and an improvement of 27% in median throughput over $\{300, 300\}$ and 26% over $\{400, 400\}$.

Regarding EO, it achieves peak utility at $\{400, 400\}$, which means that it cannot exploit heterogeneous channels. Indeed, its performance with $\{400, 300\}$ is lower than $\{400, 400\}$. As another example, while O-CS is capable of improving performance in $\{500, 300\}$ with respect to $\{500, 400\}$, EO experiences the opposite. And there are heterogeneous configurations where EO experiences substantial drops in performance (e.g. $\{400, 200\}$ and $\{400, 100\}$) where this does not happen with O-CS. This is a result of EO being unaware of any differences between the channels and treating them in the same way.

Comparing O-CS to EO, the peak utility of EO is lower than that of O-CS. One of the main problems with EO is that its minimum throughput is much lower than O-CS. This is due to inefficiencies in the 802.11 MAC which causes some flows to starve. Comparing the best case of O-CS with the best of EO, there is an average improvement of 134% in minimum rate of O-CS to EO, and 22% in median rate. One category where

EO outperforms O-CS is in mean throughput, but this comes at the expense of decreased fairness. Even in homogeneous channels, O-CS outperforms EO: in the homogeneous configuration of $\{400, 400\}$, it achieves 116% better minimum rate than EO. Opportunistic channel switching is thus effective in coping with MAC inefficiencies and exploiting channel diversity.

Fig. 3 shows results when stations have only one radio ($R = 1$). We also show the performance of O-CS with $R = 2$ for comparison. As we can see, the results have similar behavior although throughput is worse due to less flexibility in channel allocation. It is worth noting that with $R = 1$ O-CS takes less advantage of heterogeneous channels. Even so, the performance of O-CS also peaks at $\{400, 300\}$ and EO at $\{400, 400\}$. There is still improvement of O-CS in heterogeneous configurations, although slightly lower, e.g. an improvement of 23% in median throughput over $\{300, 300\}$ and 21% over $\{400, 400\}$. Comparing the best case of O-CS with the best of EO, there is an average improvement of 79% in minimum rate of O-CS to EO, and 16% in median rate.

From the above results we can observe that opportunistic channel switching is necessary to exploit heterogeneous chan-

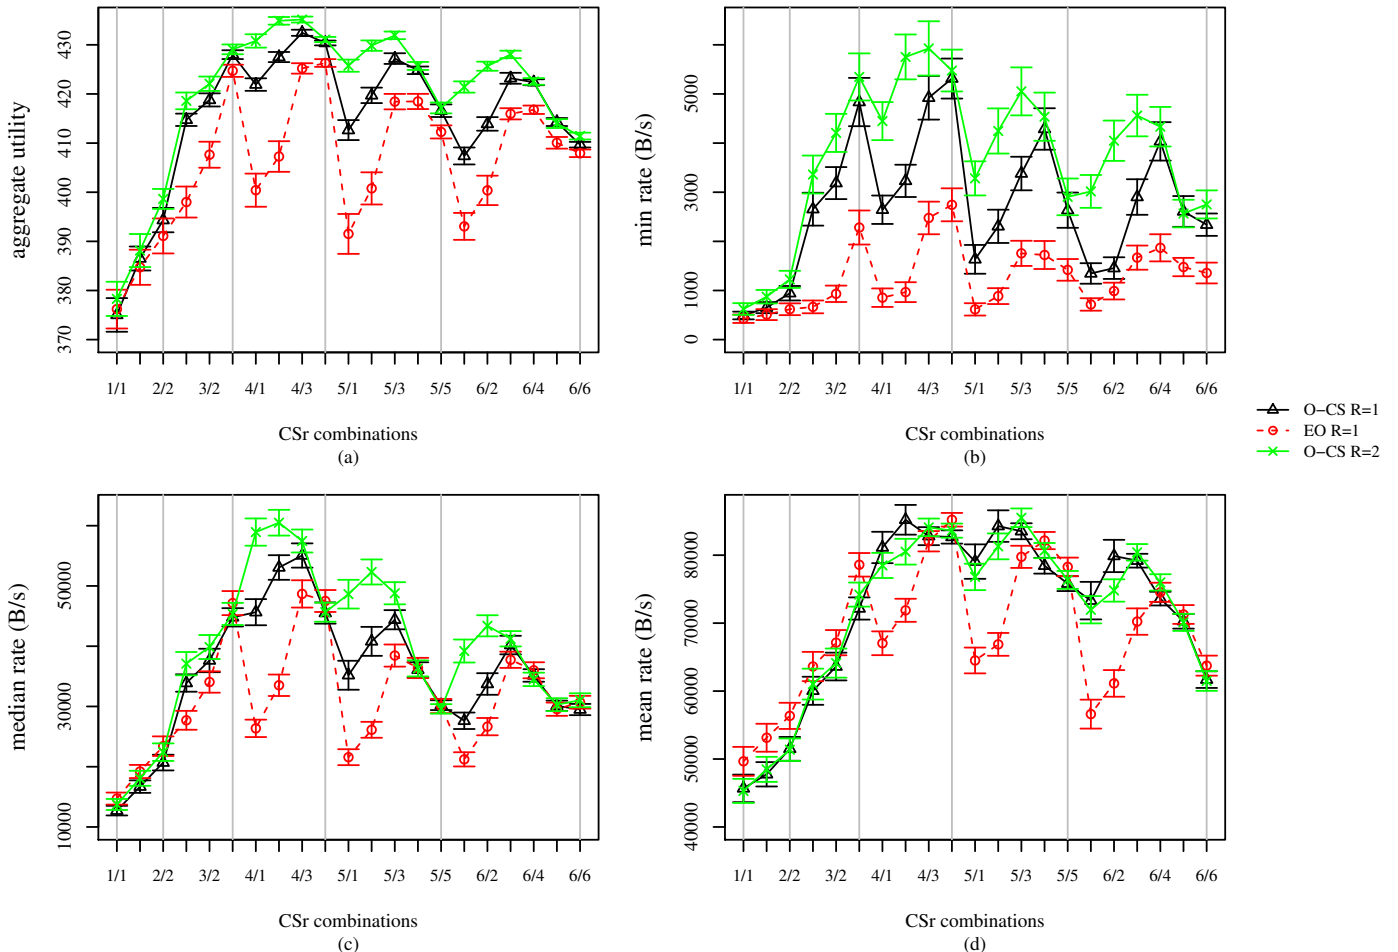


Fig. 3. Performance comparison of O-CS and EO with $C = 2$, UDP and $R = 1$.

nels, and that heterogeneous channels improve performance.

Fig. 4 compares the performance with $C \in \{2, 3\}$, $R = 2$ and using TCP and UDP flows. Here we only show utility for brevity. As we can see, the results show a similar behavior in all cases. The performance depends on the chosen CSr combination, and O-CS finds peak performance in an heterogeneous configuration while EO finds peak performance in homogeneous configurations. With $C = 2$, O-CS obtains peak utility at $\{400, 300\}$ and EO at $\{400, 400\}$. With $C = 3$, O-CS obtains peak utility at $\{400, 300, 300\}$ and EO at $\{400, 400, 400\}$.

In these tests we have seen that spatial reuse notably improves capacity and average link throughput. For example, with O-CS in $C = 2$ and UDP there is a 85% increase in median throughput of $\{400, 300\}$ with respect to $\{600, 600\}$ (see Fig. 2). That's why, in general, it is better to have a carrier sensing range as low as possible, even if this means that links will have to transmit at a lower rate. In this case, more transmission opportunities at lower rate result in more throughput than less opportunities at higher rate and increased contention. Furthermore, the use of heterogeneous channels allows having a subset of channels with even less aggressive carrier sensing than would be possible with homogeneous channels, thus allowing a higher spatial reuse.

VI. CONCLUSIONS

Opportunistic channel switching is an effective technique to exploit channel diversity. It has been used in the past to adapt to fluctuations in signal quality, allowing links to select the channel which gives them the best quality at a point in time. In addition, there are other phenomena which induce channel diversity, such as differences in congestion and MAC dynamics (e.g. hidden and exposed nodes) between channels.

In this paper we explicitly introduce a new kind of diversity by having *heterogeneous channels*, i.e. channels with different PHY/MAC parameters. The motivation behind this is the high degree of diversity between links in a WLAN. Stations are randomly located around APs, and the signal quality of their link to the AP varies greatly based on factors like distance or presence of obstacles. Therefore, it is expected that each link perform differently under a different set of parameters (e.g. transmission power or carrier sensing threshold).

We have designed an opportunistic channel switching protocol (O-CS) which is capable of exploiting channel diversity in general, including heterogeneous channels. The protocol requires only minimal information which is readily available at the APs, i.e. it adapts channel allocation based on the performance of the flows on each channel. The heterogeneous

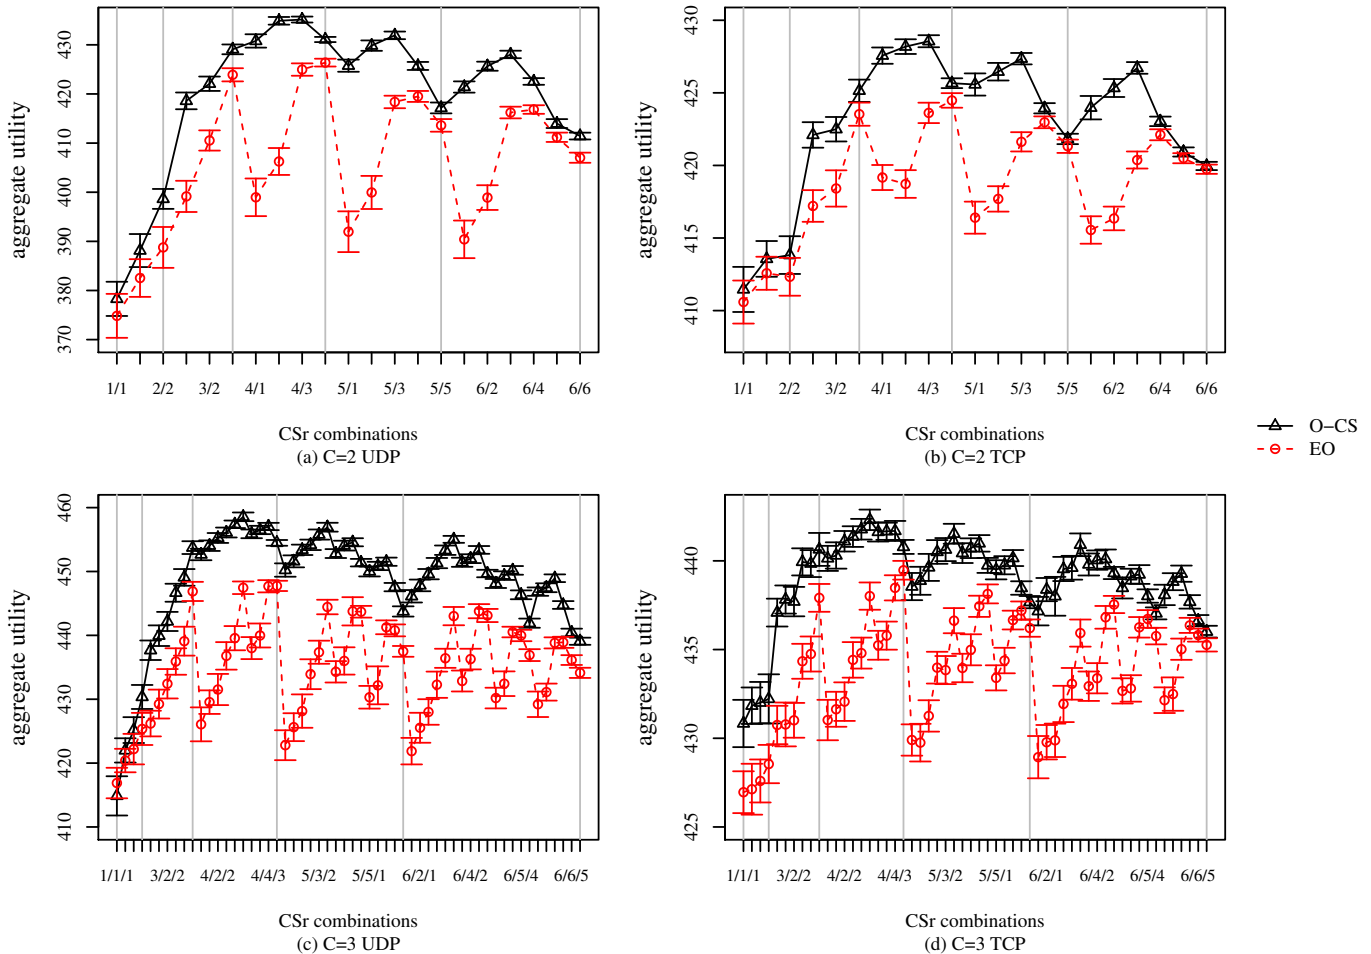


Fig. 4. Comparison of the aggregate utility of O-CS and EO with $C \in \{2, 3\}$ and $R = 2$ using UDP and TCP flows. The bottom axis shows different CSr combinations, which include both heterogeneous and homogeneous channels. There is a vertical line on each homogeneous configuration, e.g. 1/1, 3/3, 2/2/2 and 4/4/4.

channel scheme we evaluate in this paper consists in configuring channels with different carrier sensing threshold.

In simulation tests in $ns-3$ we have observed that the use of heterogeneous channels improves performance and fairness: the aggregate utility (proportional fairness) improves, and we have seen improvement of up to 26% in median flow throughput with respect to the best homogeneous channel configuration. The results also show that O-CS is necessary to exploit heterogeneous channels. For example, a channel equal-occupancy strategy with no opportunistic switching performs worse in a heterogeneous configuration. Finally, O-CS also improves in homogeneous configurations with respect to equal-occupancy, achieving notably better minimum flow rate (up to 116%). This is due to the capability of O-CS of exploiting channel diversity even in homogeneous configurations.

The results observed in the heterogeneous channel scheme are promising. There are many more ways in which heterogeneous channels can be configured and used with opportunistic channel switching (e.g using different transmit power, back-off intervals, or link rates), and which can result in further gains. This will be examined in future work.

ACKNOWLEDGMENT

This research is supported in part by National Science Foundation awards CNS-11-17539 and Futurewei Technologies. Any opinions, findings, and conclusions or recommendations expressed here are those of the authors and do not necessarily reflect the views of the funding agencies or the U.S. government.

REFERENCES

- [1] R. Murty, J. Padhye, R. Chandra, A. Wolman, and B. Zill, "Designing High Performance Enterprise Wi-Fi Networks," in *NSDI*, J. Crowcroft and M. Dahlin, Eds. USENIX Association, 2008, pp. 73–88.
- [2] V. Shrivastava, N. Ahmed, S. K. Rayanchu, S. Banerjee, S. Keshav, K. Papagiannaki, and A. Mishra, "CENTAUR: realizing the full potential of centralized w lans through a hybrid data path," in *MOBICOM*. ACM, 2009, pp. 297–308.
- [3] W. Zhou, T. Bansal, P. Sinha, and K. Srinivasan, "BBN: Throughput Scaling in Dense Enterprise WLANs with Blind Beamforming and Nulling," in *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking*, ser. *MobiCom '14*. New York, NY, USA: ACM, 2014, pp. 165–176. [Online]. Available: <http://doi.acm.org/10.1145/2639108.2639113>
- [4] J. Florwick, J. Whiteaker, A. Cuellar Amrod, and J. Woodhams, "Wireless LAN Design Guide for High Density

- Client Environments in Higher Education.” [Online]. Available: http://www.cisco.com/c/en/us/products/collateral/wireless/aironet-1250-series/design_guide_c07-693245.pdf
- [5] G. Berger-Sabbatel, F. Rousseau, M. Heusse, and A. Duda, “Performance anomaly of 802.11b,” in *INFOCOM*, 2003.
- [6] V. Kanodia, A. Sabharwal, and E. W. Knightly, “MOAR: A Multi-Channel Opportunistic Auto-Rate Media Access Protocol for Ad Hoc Networks.” in *BROADNETS*. IEEE Computer Society, 2004, pp. 600–610.
- [7] N. B. Chang and M. Liu, “Optimal channel probing and transmission scheduling for opportunistic spectrum access.” *IEEE/ACM Trans. Netw.*, vol. 17, no. 6, pp. 1805–1818, 2009.
- [8] Y. Liu, M. Liu, and J. Deng, “Evaluating Opportunistic Multi-Channel MAC: Is Diversity Gain Worth the Pain?” *IEEE Journal on Selected Areas in Communications*, vol. 31, no. 11, pp. 2301–2311, 2013.
- [9] Q. Wang and M. Liu, “Throughput optimal switching in multi-channel w lans,” *IEEE Transactions on Mobile Computing*, vol. 12, 2013.
- [10] A. Raniwala and T.-C. Chiueh, “Architecture and algorithms for an IEEE 802.11-based multi-channel wireless mesh network,” in *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies*, vol. 3, 2005, pp. 2223–2234.
- [11] K. N. Ramachandran, E. M. Belding, K. C. Almeroth, and M. M. Buddhikot, “Interference-Aware Channel Assignment in Multi-Radio Wireless Mesh Networks,” in *Proceedings of the 25th IEEE International Conference on Computer Communications*, ser. INFOCOM ’06, 2006.
- [12] S.-H. Kim, D.-W. Kim, and Y.-J. Suh, “A Cooperative Channel Assignment protocol for multi-channel multi-rate wireless mesh networks,” *Ad Hoc Networks*, vol. 9, no. 5, pp. 893–910, 2011.
- [13] Q. Shen, X. Fang, R. Huang, P. Li, and Y. Fang, “Leveraging spatial reuse with adaptive carrier sensing in 802.11 wireless networks,” in *Proceedings of the 5th International ICST Conference on Heterogeneous Networking for Quality, Reliability, Security and Robustness*, 2008, pp. 38:1–38:7.
- [14] X. Yang and N. H. Vaidya, “A spatial backoff algorithm using the joint control of carrier sense threshold and transmission rate.” in *SECON*. IEEE, 2007, pp. 501–511.
- [15] L. B. Jiang and S. C. Liew, “Improving Throughput and Fairness by Reducing Exposed and Hidden Nodes in 802.11 Networks,” *IEEE Transactions on Mobile Computing*, vol. 7, no. 1, pp. 34–49, Jan. 2008. [Online]. Available: <http://dx.doi.org/10.1109/TMC.2007.1070>
- [16] H. Zhai and Y. Fang, “Physical carrier sensing and spatial reuse in multirate and multihop wireless ad hoc networks.” in *INFOCOM*. IEEE, 2006.
- [17] F. Kelly, “Charging and rate control for elastic traffic,” *European Transactions on Telecommunications*, vol. 8, pp. 33–37, 1997.
- [18] D. Smithies, “minstrel rate control algorithm.” [Online]. Available: <http://wireless.kernel.org/en/developers/Documentation/mac80211/RateControl/minstrel>
- [19] J. J. Galvez and P. M. Ruiz, “Efficient rate allocation, routing and channel assignment in wireless mesh networks supporting dynamic traffic flows,” *Ad Hoc Networks*, vol. 11, no. 6, pp. 1765–1781, 2013.
- [20] “The ns-3 network simulator.” [Online]. Available: <http://www.nsnam.org/>