

Socially Conscious Channel Selection in 802.11 WLANs for Coexistence in a Non-cooperative Environment

Joo Ghee Lim

Chun Tung Chou

Sanjay Jha

School of Computer Science and Engineering
University of New South Wales
NSW 2052, Australia
{jool, ctchou, sanjay}@cse.unsw.edu.au

ABSTRACT

The increasing number of independent IEEE 802.11 WLANs owned and managed by autonomous users has led to increased interference, resulting in performance degradation and unfairness. Performance can be improved by allowing these networks to operate on different channels. Due to the autonomous nature of the networks, a suitable channel selection scheme should be distributed, adaptive and require no explicit coordination. In this paper, we model the channel selection of WLANs as a non-cooperative game in a learning setting. Using a novel method of acquiring a disruption factor value, we propose a class of socially conscious channel selection schemes based on game-theoretic learning. These schemes are distributed, adaptive and are able to improve fairness without explicit inter-network communication. These features allow the WLANs to coexist in an interference-limited but non-cooperative environment. They also have the advantage of not requiring any modification to the existing 802.11 standards. Simulations show improved fairness and aggregate throughput compared with two existing schemes.

Categories and Subject Descriptors: C.2.1 [Computer- Communication Networks]: Network Architecture and Design - *Wireless Communication*

General Terms: Algorithms, Design, Performance

Keywords: IEEE 802.11, channel selection, game theory, learning, coexistence

1. INTRODUCTION

IEEE 802.11 has become the predominant technology to enable Internet access of many wireless devices. As a result, it is common to have multiple Wireless Local Area Networks (WLANs) deployed in a single locality. Apart from locations like university campuses or corporate offices, most WLANs can be characterized as a single Access Point (AP) providing Internet connectivity to one or more clients. They are often set up by individuals (e.g. residential occupants, small businesses) and are therefore owned and managed by separate entities. We term these networks *Independent WLANs*. Henceforth, the terms independent WLAN and network will be used interchangeably in this paper.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MSWiM'09, October 26–29, 2009, Tenerife, Canary Islands, Spain.
Copyright 2009 ACM 978-1-60558-616-8/09/10 ...\$5.00.

These networks generally have the following characteristics:

- *Uncoordinated*: They have variable and uncoordinated operating times. Over time, APs are also set up and removed with similar irregularities. In addition, this non-coordination leads to uneven density of deployment, with more APs located in highly concentrated residential and business areas.
- *Non-cooperative*: Unlike enterprise or campus WLANs, these networks do not have any network management software that can ensure efficient use of the radio resources (e.g. channel usage or power control) in a centralized or cooperative manner.

One way to improve the performance of multiple WLANs in an interference-limited environment is to make use of the different channels available in the standards. As a result, a lot of research has been devoted to developing effective and efficient channel selection schemes [1, 3, 4, 11, 13–15, 18, 21, 22]. With the exponential growth of independent WLANs in recent years [2], there is a need for channel selection schemes that are *distributed* and *adaptive* in operation, without explicit control messages exchanged among them.

Fairness is another issue that arises when 802.11-based WLANs are deployed over an area spanning multiple cells or *collision domains*. In wireless networks, a collision domain is the region where links located within it interfere with one another. In [7], the authors show that, due to the inherent MAC protocol, 802.11 can exhibit unfairness in such situations. Depending on their locations, some links experience much lower throughput performances compared to others. In this paper, we define them as *starved* links or networks.

While using a channel selection scheme, a starved network may not always be able to improve its throughput by unilaterally switching to a different channel. We will show that fairness among independent WLANs can be improved if WLANs that are causing the starvation can detect this unfairness and take steps to alleviate it. We term these networks *socially conscious* networks, since they proactively improve the “welfare” of disadvantaged networks.

In this paper, we propose a class of channel selection schemes that is distributed, adaptive as well as socially conscious with the aim of increasing overall throughput performance and inter-network fairness among independent WLANs.

Following are the primary contributions of this paper:

- We propose a class of channel selection schemes based on game theoretic learning that is practical to be implemented in existing 802.11 networks. Our schemes require no modification to the standards and hence can interoperate with existing networks.
- We present a *disruption factor* value for each independent WLAN that seeks to inform it of the unfairness it is causing to the surrounding networks. We describe a novel approach to acquire this value without explicit exchange of messages. This is incorporated into our channel selection schemes to create socially conscious WLANs.

- Through extensive simulations, we show that our schemes achieve higher overall throughput (as high as 30%) as well as better fairness (as high as 17%) when compared to two existing channel selection schemes.

The paper is organized as follows: In the next section, we briefly discuss the issue of unfairness in 802.11-based WLANs. In Section 3, we review the existing channel selection schemes proposed for WLAN deployments. This is followed by introduction to a class of learning algorithms based on non-cooperative game theory, which we use to develop socially conscious channel selection schemes. In Section 5, our channel selection schemes are described in detail. We present some simulations to evaluate and compare our schemes in Section 6 and conclude with Section 7.

2. UNFAIRNESS IN IEEE 802.11

The most common MAC in 802.11 is the Distributed Coordination Function (DCF), which primarily consists of the CSMA/CA mechanism. A station first senses the channel if it intends to transmit. If the channel is free, the station may still defer for a backoff duration that depends on the previous transmission success or failure. In addition, during this period, the backoff process stops anytime the channel becomes busy and only resumes when the channel becomes free again.¹

While DCF seeks to ensure some level of fairness within a single collision domain, it has been shown to result in significant unfairness when the stations span over multiple collision domains. Through detailed modelling and analysis, Garetto et al. [7] show that one of the major causes of unfairness is the difference in channel state perceived by the transmitters. This is shown in Figure 1. In this example, transmitter T2 is in the sensing range of transmitters T1 and T3 that cannot sense each other. T2 freezes its backoff counter whenever T1 or T3 is transmitting. T1 (T3), on the other hand, can keep decreasing its backoff counter when T3 (T1) is transmitting. In fact, both T1 and T3 can even be transmitting simultaneously. This results in a much limited transmission opportunity for T2, compared to T1 and T3. Subsequently, traffic on Link 2 becomes starved while those on Links 1 and 3 remain high. In [7], this is known as the Flow-In-the-Middle (FIM) effect.




Figure 1: FIM example: T2 is within the sensing range of T1 and T3 but T1 and T3 are out of each other’s sensing range.

It is clear that the unfairness in 802.11-based networks that results from FIM is related to the relative locations of the links. In independent WLANs, the location of each network is often constrained by the location of the user. For example, a resident could only set up an AP within the confines of her home, or a cafe owner could only install an AP within his business premises. If a network happens to be located between two other WLANs that cannot sense each other, it is unduly penalized for no fault of its own.

In this paper, to ensure interoperability with existing WLANs, we assume that the fundamental 802.11 MAC protocol is unlikely to be changed in the short term. Instead, we propose a channel selection scheme that is able to improve fairness in independent

¹IEEE 802.11 DCF consists of many other intricate components that are not described here. Interested reader is directed to the standards and the vast literature available for a full description.

WLAN deployments. The 802.11 DCF MAC remains unchanged, since channel selection schemes are not part of the standards.

3. RELATED WORK

Akella et al. [2] are one of the first to highlight the challenges of independent WLANs coexistence. Using data from existing WLAN deployments, they study the impact of interference on client performance. In [5], Ergin et al. provide experimental and simulation results of similar unplanned WLAN deployments. However, they assume a single collision domain for all the WLANs and do not capture the unfairness that arises over multiple collision domains.

Least Congested Channel Search (LCCS) [1] is a simple channel selection scheme where an AP scans its channels to use the one with the fewest number of neighboring APs. This scheme is currently implemented in some APs. In [13], Mishra et al. propose a Hminmax distributed algorithm that formulates the problem as a weighted graph coloring one where an edge is defined for every WLAN that is within the communication range of a particular WLAN, say i . For every edge, they define as W the number of i ’s clients that will be affected by the interference of the WLAN sharing that edge (if they share the same channel). Along with the channel separation I between the WLANs sharing the edge, they compute the weight of each edge as $I \times W$. Hminmax is performed periodically and ensures that i selects the channel that minimizes the $I \times W$ value of its maximum weighted edge. Hminmax has been shown to outperform LCCS, although it requires the clients to also scan the available channels and provide feedback to the AP, thereby increasing complexity and communication overheads.

In both LCCS and Hminmax, a WLAN scans a channel by listening for frames from neighboring WLANs in order to compute the channel’s utility. However, they are unable to detect the interference that comes from WLANs lying outside the communication range, as their frames cannot be correctly received. LCCS and Hminmax also do not take into account the traffic load of the networks in each channel. A number of works have highlighted these points and have proposed solutions that compute channel utilities using more detailed metrics. In [11], Leith et al. propose a simple learning algorithm, making use of frame error rate as the metric.

Chen et al. [3] introduce a few measurement-based frequency allocation algorithms. One of the algorithms uses the clients’ measurement of the channel interference I_c , the received signal power from the AP to the client R_c and the traffic volume Y_c between the AP and client. The utility of each channel is computed by summing up $(\frac{Y_c}{S_c} \cdot I_c)$ of all clients belonging to the AP. The No-Coord User (No-U) algorithm uses these utilities, where each AP performs periodic scanning and independently chooses the channel with the lowest value. Because it takes into account the clients’ view of the channel, No-U is shown to perform better than that of [11].

In [4], the authors make use of IEEE 802.11k parameters – the channel load info and noise histogram – as the metrics to compute the channel utility. When the channel load falls below a threshold, it triggers the AP to choose a new channel with the lowest noise histogram value. Yoo and Kim [22] use a similar threshold-based approach in their solution.

All the schemes described so far are suitable for independent WLANs as they require no inter-WLAN coordination. However, all of them are selfish in nature and unlike our schemes, do not specifically address the fairness issue. Since *Hminmax* [13] and *No-U* [3] have been shown to perform better than the majority of schemes described here, we will compare our schemes against them.

There are also schemes that have been proposed that assign the channels in a centralized manner, e.g. [14, 18]. They assume all the WLANs belong to a single managing entity and are therefore

not applicable to independent networks. Their discussion is beyond the scope of this paper. In [21], even though the authors do not assume a single managing entity, they assume explicit communication among the independent WLANs. Thus it is also not a suitable solution to our problem.

4. GAME THEORY AND LEARNING

Game theory [6] has been applied to network interactions [8] and in particular, wireless networks [20]. In our prior work [12], it is asserted that non-cooperative game theory is suitable to model the interactions of independent wireless networks, where each network constitutes a *player* and the actions available to the network, e.g. channels, transmit power, represent the *strategies* available.

While classical non-cooperative game theory is able to provide insights into the interactions of independent wireless networks, it may not be suitable for practical schemes. This has been pointed out by Greenwald et al. [8], as they explore the application of game theory in the networking environment. Classical game theory assumes common knowledge of the set of players and strategies. This is highly unlikely given the distributed nature of independent networks. An independent WLAN would not be aware of all the WLANs that could affect its performance (e.g. those deployed beyond its communication range but within the carrier-sensing range). In addition, the utility function and the players in the game change as new WLANs are deployed and old ones are taken down over time. In the next section, we will introduce the concept of learning in game theory that can provide practical channel selection solutions to the independent networks.

4.1 Game Theoretic Learning

Learning in game [17, Ch. 4] allows initially uninformed players to acquire information about the state of the world they are in, as the game is repeatedly played. Learning has been applied to networking research in [8], where the authors study what strategies players will play in the long run as they learn about their environment. The attribute of game theoretic learning is that information like the number and identity of players in the game and their utility functions is not required by a player in order to play efficient strategies in the long run.

4.2 Model and Notations

Let $\Gamma = \langle \mathcal{N}, \mathcal{S}, \{U_i\}_{i \in \mathcal{N}} \rangle$ be a finite N -player game in normal form, where \mathcal{N} is a finite set of players, $|\mathcal{N}| = N$, and \mathcal{S} is the Cartesian product of the set of strategies available to each player in \mathcal{N} ; i.e., $\mathcal{S} = \times_{i \in \mathcal{N}} \mathcal{S}_i$ where \mathcal{S}_i is the set of strategies available to player i . $S = (s_1, s_2, \dots, s_N) \in \mathcal{S}$ is a strategy profile consisting of a strategy from every player in \mathcal{N} . $U_i: \mathcal{S} \rightarrow \mathbb{R}$ is defined as a utility function of player i representing the value of the outcome resulting from a strategy profile S . For a particular strategy profile S , if the strategy used by player i is $s_i \in \mathcal{S}_i$, we collectively term the strategies of the other players as s_{-i} .

In game theoretic learning, let U_i^t denote the utility of player i at time t . $S^t = (s_1^t, s_2^t, \dots, s_N^t), \forall s_i^t \in \mathcal{S}_i$ denotes the strategy profile of the players at time t . The s_i^t that is played by player i arises from a probability density q_i^t which denotes the probability of playing each strategy $s_i \in \mathcal{S}_i$. Over time, the probability density q_i^t will evolve with more favorable strategies taking higher values, as determined by the learning algorithm. In each period, player $i \in \mathcal{N}$, chooses a strategy $s_i^t \in \mathcal{S}_i$ in accordance to q_i^t . q_i^t can also be interpreted as a function $q_i^t(s_i)$ returning the probability of playing strategy $s_i \in \mathcal{S}_i$, at time t , which is done in this paper.

We will present two different learning algorithms in the subsequent sections.

4.3 Best Response Learning

In best response (BR) learning, q_i^t is updated by each player in the following manner (assuming no tie among utilities):

$$q_i^{t+1}(j) = \begin{cases} 1, & \text{if } j = \arg \max_{s_i \in \mathcal{S}_i} U_i^t(s_i); \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

Essentially, BR learning (as the name suggests) *always* uses the strategy that yields the highest utility during the last period when the game is played.

4.4 Internal Regret Minimization Learning

While BR learning uses the immediate past period to determine its strategy choice, internal regret minimization (IRM) learning can be viewed as using a history of periods to make the decision. In IRM learning, the notion of *internal regret* must first be defined.

At time t , we denote the internal regret R_i^t that player i feels for playing strategy s_i^t rather than $s_i \neq s_i^t$ as

$$R_i^t(s_i^t, s_i) = [D_i^t(s_i^t, s_i)]^+ \quad (2)$$

where $[\cdot]^+ = \max\{\cdot, 0\}$ and

$$D_i^t(s_i^t, s_i) = \frac{1}{t} \sum_{\tau \leq t; s_i^\tau = s_i^t} [U_i(s_i, s_{-i}^\tau) - U_i(s_i^\tau, s_{-i}^\tau)] \quad (3)$$

The value $D_i^t(s_i^t, s_i)$ can be interpreted as the average difference in utilities a player would have obtained if for every time he had played s_i^t in the past, he had instead played $s_i \neq s_i^t$. In [9], Hart and Mas-Colell introduce an IRM learning algorithm using the following q_i^t updating scheme:

$$q_i^{t+1}(j) = \begin{cases} \frac{1}{\mu} R_i^t(s_i^t, j), & \text{for all } j \neq s_i^t, \\ 1 - \sum_{j \in \mathcal{S}_i; k \neq j} q_i^{t+1}(k), & \text{otherwise.} \end{cases} \quad (4)$$

where $\mu > 0$ is a sufficiently large value².

Briefly, the IRM learning algorithm of Hart and Mas-Colell updates the probability that a player would switch strategy as a linear function of the average regret. The IRM learning algorithm ensures that as $t \rightarrow \infty$, the expected internal regret over the probability density q_i^t almost surely becomes zero [9].

One can see that unlike BR learning, a better utility of another strategy in the previous period does not trigger an immediate strategy change in IRM learning. This is because it uses a probability density and a regret value that are computed over the history of play.

5. SOCIALLY CONSCIOUS CHANNEL SELECTION SCHEMES

In this section, we describe how we incorporate the learning algorithms into practical channel selection schemes. We also introduce a novel way of detecting unfairness in the network environment that requires no explicit message exchange among the independent WLANs. By adding this capability into our channel selection schemes, we are able to develop socially conscious schemes.

5.1 WLANs Channel Selection Game

We first define the WLANs Channel Selection Game. The game is played by a set of players \mathcal{N} , where each player $i \in \mathcal{N}$ is an independent WLAN deployed within a predefined area. We assume

²In most cases, it suffices for μ to be $|\mathcal{S}_i| - 1$, which is the value used in our simulations.

each WLAN consists of an AP connected to the Internet via a wired connection, and a collection of one or more wireless clients. Henceforth, the terms player and WLAN will be used interchangeably.

Each WLAN i is able to switch between $|\mathcal{S}_i|$ numbers of *non-overlapping* channels. We assume that each WLAN can only be on one channel at any given time. The channel $s_i \in \mathcal{S}_i$ that WLAN i decides on thus constitutes the strategy chosen by player i out of the available strategy set of \mathcal{S}_i . Henceforth, the terms strategy and operating channel will be used interchangeably. As mentioned above, the classical way of computing U_i as a known function of the strategy profile of all the players in the game is not possible here. Instead, U_i is computed by player i by some measurement process.

This game is played repeatedly through time: $t = 1, 2, \dots$, where after every $T_A \in \mathbb{R}^+$ period of normal operation, each player i will perform some process that will determine U_i^t and choose a channel s_i^{t+1} for the next T_A period (shown in Figure 2). The operation performed during the T_P period differs for the different learning schemes. Note that we do not assume that the times when the players perform the channel switching operation are synchronized.




Figure 2: Timing diagram of the channel selection game, where each iteration contains an active period of T_A duration and (passive) scanning period of T_P duration.

5.2 Channel Selection using BR Learning (CSBRL)

In the CSBRL scheme, during each T_P period, a player i performs a passive scanning operation of all channels in \mathcal{S}_i where each channel is scanned for t_s time units. In each t_s scanning duration, player i measures $t_{b_i}^t(s_i)$, which is the total time the channel s_i is sensed busy at time period t . Practically, this is the time the clear channel assessment (CCA) function, as defined in the standards, is set to busy within the t_s period.

For each channel scanned, we compute the utility of the channel,

$$U_i^t(s_i) = 1 - \frac{t_{b_i}^t(s_i)}{t_s}, \forall s_i \in \mathcal{S}_i \quad (5)$$

The utility can be seen as an estimation of the fraction of the channel non-busy time. A higher U_i^t suggests that player i could have more opportunity to transmit data on that channel.

With the utilities acquired for each channel, the player updates q_i^{t+1} using (1), which is essentially choosing the channel with the lowest channel utilization.

5.3 Channel Selection using IRM Learning (CSIRML)

In the CSIRML scheme, we apply IRM learning to the channel selection process. The scanning process and utility remain as described in (5) for the CSBRL scheme. The difference is in the updating of the probability density, using (2), (3) and (4). At the end of the updating process, the new channel will be chosen over the probability density q_i^{t+1} .

5.4 Disruption Factor

While most channel selection schemes, including the ones proposed in this paper so far, allow a starved player to switch to a channel with a higher utility (i.e., lower utilization). The player is unlikely to see any improvement in its situation if no such channel

exists. Figure 3 provides an example of this case, consisting of 5 players using 2 channels.

One can see that if both channels were occupied by exactly 2 outer players, Player P1 would not be able to get a better utility whichever channel it chooses, since P1 will be in a FIM situation in either channel. However, if any of the outer players can detect that P1 is unfairly starved, and switches to the other channel, then P1 could potentially share this channel with the remaining player. In fact, one can do even better. If both outer players switch channel such that only P1 remains in that channel, all players will get maximum performance, since there is now no interference.




Figure 3: Channel Selection Game with $\mathcal{N} = \{P1, P2, \dots, P5\}$ and $\mathcal{S}_i = \{C1, C2\}$. Dotted lines denote interference if players are on the same channel.

The key challenge is to allow a player to detect that it is causing unfairness to some players without the exchange of explicit control messages. That is because in independent WLANs, the players are not likely to cooperate with explicit feedbacks. In addition, interference often extends beyond the communication range of a network.

We will now describe a novel way for a player to make this detection, using a disruption factor value δ . To compute δ , we make use of the fact that during the operation of our channel selection schemes, there is both an active phase, and a passive phase. The active phase occurs when a player is sending data over the channel it has chosen, with duration of T_A . The passive phase occurs when the player is scanning the channel set, with duration of T_P , which is $(|\mathcal{S}_i| \times t_s)$ plus the processing and channel switching times.

During the active phase, a player i can compute the utility when it is actively transmitting on the channel s_i^t using

$$\bar{U}_i^t(s_i^t) = 1 - \frac{T_{b_i}^t(s_i^t)}{T_A - T_{d_i}^t(s_i^t)} \quad (6)$$

where $T_{b_i}^t$ is the total time in the duration T_A that the channel is sensed busy by player i and $T_{d_i}^t$ is the time that player i spent in transmission mode.

For the period t , let

$$\delta^t(i) = [U_i^t(s_i^t) - \bar{U}_i^t(s_i^t)]^+ \quad (7)$$

The value δ thus gives a sense of the difference between the state of the channel activity when player i is participating actively in the medium to when it is not. A high δ value would mean that there is more channel activity when player i is passive compared to when it is active. This gives an indication that player i may be unfairly causing starvation to one or more other players due to different perceptions of the channel conditions (e.g. in the FIM case). In Section 6.1, we show that δ is able to detect starvation in a FIM setting.

5.5 Incorporating Social Consciousness

As defined earlier, a channel selection scheme is *socially conscious* if it enables a player to detect unfairness and takes actions to improve it. Using the disruption factor acquired as described in the previous section, we show how social consciousness can be incorporated into our channel selection schemes.

To enable social consciousness in our schemes, we define a new utility function, $V_i^t(s_i^t)$ which is computed according to algorithm 1.

Algorithm 1 Compute SC Utility V_i^t

```
1: for  $t = 1, 2, 3, \dots$  do
2:   Compute  $U_i^t(s_i)$  using (5)
3:   Compute  $\bar{U}_i^t(s_i^t)$  using (6)
4:   Compute  $\delta^t(i)$  using (7)
5:   if  $t = 1$  or  $s_i^{t-1} \neq s_i^t$  then
6:     cumDel  $\leftarrow$  0
7:   end if
8:   cumDel  $\leftarrow$  cumDel +  $\delta^t(i)$ 
9:   for every  $s_i \in \mathcal{S}_i$  do
10:    if  $s_i = s_i^t$  then
11:       $V_i^t(s_i) \leftarrow U_i^t(s_i) - \alpha$  cumDel, where  $\alpha \in \mathbb{R}^+$ 
12:    else
13:       $V_i^t(s_i) \leftarrow U_i^t(s_i)$ 
14:    end if
15:  end for
16: end for
```

Algorithm 1 can be understood as follows. For every time period, the values U_i^t , \bar{U}_i^t and the disruption factor δ are computed. As long as a player continues using a channel consecutively, each δ is added to a cumDel value (line 8). The cumDel value can be viewed as the cumulative effect of a player’s disruption factor and it gets larger the longer this player stays on a particular channel. This counter is reset to 0 when a player chooses to switch channel (line 6).

The utility of the current channel is discounted by a factor α of this cumDel value, while those of the other channels remain unchanged. The effect of this is to penalize a player for continuing to use a channel that it is causing disruption to (i.e., consistently having a high δ). U_i^t will be substituted with V_i^t in either (1) or (3) to compute the probability density for BR Learning or IRM Learning respectively. We will term the socially conscious schemes CSBRL with Social Consciousness (CSBRL-SC) and CSIRML with Social Consciousness (CSIRML-SC) respectively.

Note that the value α determines how much a player is conscious about its disruption to other networks. When $\alpha = 0$, CSBRL-SC is essentially CSBRL and CSIRML-SC is CSIRML. We investigate the effect of this SC factor α in Section 6.2.

6. PERFORMANCE EVALUATION

In this section, we present results of simulations conducted to evaluate the performance of our channel selection schemes. All the simulations have been conducted on the Qualnet simulator [19], allowing us to evaluate the performance using realistic channel conditions. In addition, the 802.11 DCF MAC and PHY layers have been realistically implemented in the simulator. We build our channel selection schemes on top of these layers to illustrate their backward compatibility with existing WLANs. Since as highlighted in Section 2, unfairness results from the MAC protocol, we use 802.11b PHY without loss of generality.

In evaluating the channel selection schemes, we run a total of 20 random topologies for each simulation set. Unless stated otherwise, the following parameters apply to all the simulations. In each topology, 10 WLANs are deployed with each WLAN consisting of an AP and 4 clients. There are 3 channels available for selection. Each WLAN appears randomly in time at the beginning of the simulation and begins the channel selection process. Since most clients currently attached to APs are mobile devices, the predominant traffic are downlink flows originating from the Internet [16]. Therefore, application data packets of size 1460B flow from every AP to each of its clients, with each flow lasting 2000s.

To evaluate the performance, we look at 3 different metrics:

1. **Fairness** To investigate overall system fairness, we compute

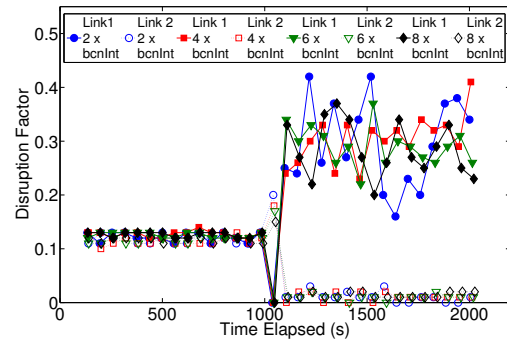


Figure 4: The change in the disruption factor over time (Links 1 and 2), for different t_s .

Jain’s fairness index [10], given as $(\sum x_i)^2 / (n \sum x_i^2)$, where x_i is the application throughput of each flow i of the n flows in the system. A number that is closer to 1 signifies that the WLANs are able to achieve a better fairness.

2. **Aggregate Network Throughput** The total application throughput of all the networks in each simulation run tells us how well the various schemes utilize the channel resources.
3. **Minimum Flow Throughput** As we are interested in the performance of the starved networks, the minimum flow throughput captures the performance of the worst-performing link in the simulation.

6.1 Evaluation of Disruption Factor δ

We first evaluate the effectiveness of the disruption factor δ to detect the unfairness in the network region through passive scanning. We deploy 3 links in the configuration of Figure 1. At the beginning of the simulation, only links 1 and 2 are active, transmitting saturated traffic. During this time, both links should experience similar throughput as they share the channel equally. After about 1000s, link 3 starts transmitting saturated traffic, resulting in link 2 being starved. At the end of each interval ($T_A = 60$ s) of actively sending traffic, the links will passively scan the channel for t_s , after which δ will be computed. We vary t_s to investigate the effect of passive scan time on δ .

Figure 4 shows the change in the disruption factor of links 1 and 2 over time, for different t_s , varying from 2 to 8 beacon intervals of around 200ms in duration. We can see from the figure that for link 1 (the outer link), there is a marked increase in the disruption factor when link 2 becomes starved. At the same time, link 2’s disruption factor decreases. This is a desired outcome, as it means that a starved link will not try to be socially conscious.

We can also see that the value of t_s has minimal effect on the disruption factor. A smaller t_s only results in a marginally larger variance in the disruption factor. This is also a favorable outcome, as the higher t_s is, the more time a network would have to spend doing passive scanning, leading to lower throughput. For the rest of our simulations, we set t_s to be $(2 \times \text{beacon interval})$.

6.2 Evaluation of SC Factor α

We next investigate the effect of the SC factor α on the channel selection schemes we have proposed. As discussed in Section 5.5, α is directly linked to how fast a player reacts to the disruption it detects in its environment. An α value of 0 means that the player is not socially conscious at all. We compute Jain’s fairness index and channel change frequency for the networks deployed over a 1000m by 1000m area, shown in Figure 5.

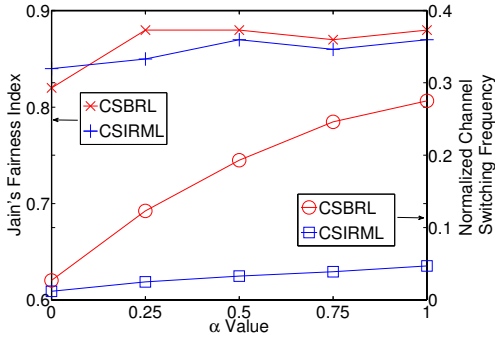


Figure 5: Throughput fairness (top 2 lines, left axis) and channel changing frequency (bottom 2 lines, right axis) for different α .

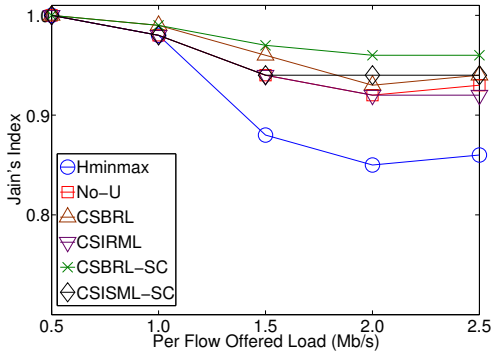


Figure 6: Throughput fairness for different offered load.

From Figure 5, we can see that increasing α has the effect of improving the fairness among the networks. With social consciousness, the BR learning scheme achieves higher throughput fairness compared to IRM learning. The tradeoff is in the channel change frequency, which represents the number of times the respective schemes trigger a change in the channel, normalized over the total number of iterations. BR learning with social consciousness results in about 5 times more changes in channel. This is to be expected, since BR learning immediately triggers a change in strategy whenever it acquires a higher utility for another strategy. This high rate of channel switching may not be desirable, as there are always costs associated with a WLAN changing its operating channel.

From the simulation result, we set the value of $\alpha = 0.5$ as the SC factor for both the CSBRL-SC and CSIRML-SC schemes in all subsequent experiments.

6.3 Comparison with Existing Schemes

We now evaluate the performance of our schemes against two existing channel selection schemes described in Section 3 – Hminmax [13] and No-U [3]. These 2 schemes are chosen for comparison because they do not require any explicit communication and coordination among the WLANs, and thus are suitable for use in independent WLAN deployment. They also show superior performance when compared with other existing schemes.

6.3.1 Offered Load

We deployed the WLANs in a 1500m by 1500m area and varied the offered load for each AP-client link from 0.5 Mb/s to 2.5 Mb/s. Figures 6, 7 and 8 show the throughput fairness, aggregate throughput and minimum link throughputs of the different channel selection schemes for varying offered load.

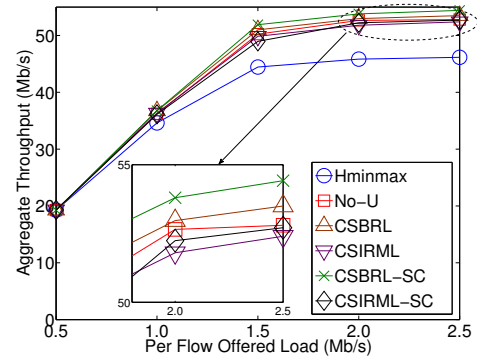


Figure 7: Aggregate throughput for different offered load.




Figure 8: Minimum flow throughput for different offered load.

At low traffic load, no starvation is taking place, as the channels are under-utilized. Consequently, all schemes perform similarly. As the networks become more congested, unfairness becomes noticeable. Hminmax performs less well compared to the other schemes in terms of fairness, aggregate throughput as well as minimum flow throughput. This shows that information from neighboring networks that are within the communication range is clearly not sufficient. Across the varying offered loads, we find little difference among the fairness and aggregate throughput results of CSBRL, CSIRML and No-U (within 3% difference), even though No-U requires the additional complexity of client feedback.

When we incorporate social consciousness into our schemes, we see that CSBRL-SC and CSIRML-SC increase the system fairness (Figure 6) compared to their non-SC counterparts. In fact, the SC schemes result in a slightly higher aggregate throughput than their non-SC counterparts, as shown in Figure 7. We believe this can be explained by cases similar to the example of Figure 3. Finally, the SC schemes prevent starvation by providing a much higher minimum throughput compared to the other schemes (Figure 8), as high as 180% when comparing with Hminmax and 50% when compared with No-U.

Since the SC schemes outperform their non-SC counterparts, only the results comparing the SC schemes with Hminmax and No-U will be shown in subsequent sections.

6.3.2 Network Area Size

Figures 9, 10 and 11 show the same triplet of performance metrics as the area where the WLANs are deployed is varied from 500m by 500m to 1500m by 1500m, with saturated traffic in all links. This gives an indication of how the different schemes perform with respect to how close the WLANs are located. In addition, the chance of uneven distribution across the area increases with the




Figure 9: Throughput fairness for different network area size.

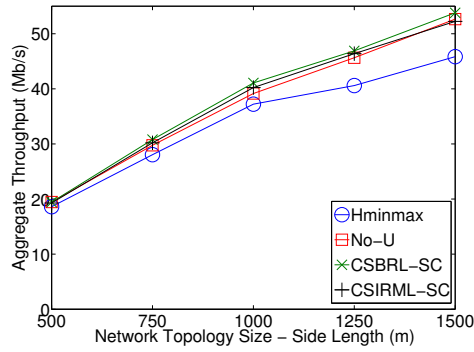


Figure 10: Aggregate throughput for different network area size.

size of the deployment area. This situation is similar to actual deployment, as some areas (e.g. residential) will see a higher density of WLANs compared to others (e.g. a nearby park).

From the figures, it can be seen that the SC schemes outperform the existing schemes by as much as 12% in terms of fairness and 10% in terms of the aggregate throughput. The minimum flow throughput also increases by as much as 2.6 times.

6.3.3 Number of Channels

As the total number of channels provided for WLANs may vary depending on the standards (IEEE 802.11b/g or IEEE 802.11a), we evaluate our schemes with respect to the number of channels available. In our simulations, we deployed 24 WLANs consisting of an AP-client connection in a 1000m by 1000m area. The simulation time is extended to 4000s. As the number of available channels increases, we would expect an effective channel selection scheme to have better fairness and overall throughput. This is because the increased number of channels reduces the chance of networks interfering with each other.

Figures 12, 13 and 14 show the network fairness, aggregate and the minimum per-link throughput for different numbers of available channels. The figures show again that the SC schemes result in a higher fairness among the networks, as high as 17% compared to Hminmax and 13% compared to No-U. In terms of aggregate throughput, CSBRL-SC performs as much as 30% and 10% better than Hminmax and No-U respectively. Both SC schemes are also able to achieve higher minimum flow throughputs compared to the existing schemes.

6.4 Channel Switching Frequency

Figure 15 shows the channel switching frequencies of the different schemes for the 1000m by 1000m network area size. The chan-

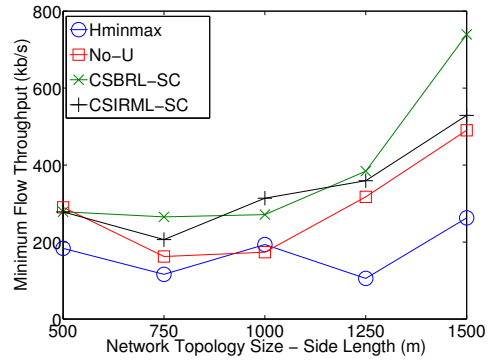


Figure 11: Minimum flow throughput for different network size.

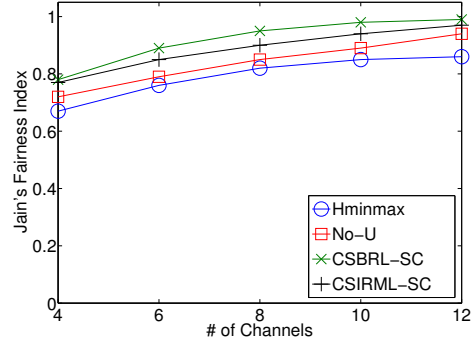


Figure 12: Throughput fairness for different number of channels.

nel switching frequencies indicate how often the players switch channels during the simulation run, normalized over the number of scanning periods. It shows that even though CSBRL-SC performs better compared to the other schemes, it results in more frequent channel switches. The reason for this, as discussed previously, is that the best response algorithms trigger immediate strategy change when the previous strategy results in a lower utility. This immediate change accounts for the frequent channel switch. On the other hand, the IRM learning algorithms are able to provide a much lower channel switching frequency at a slightly lower performance cost (at times).

This observation presents a tradeoff for the choice of channel selection schemes – if the channel switching cost is high (e.g. when the wireless devices have high channel switching time), the IRM learning schemes can be seen as better solutions. Finally, even though Hminmax has the lowest channel switching frequency, it also consistently produces the worst performance. This shows that the information it acquires is not sufficient for it to make effective channel switching decisions.

7. CONCLUSION

Due to the inherent nature of the CSMA/CA mechanism in IEEE 802.11 DCF MAC, it has been shown that unfairness can occur among WLANs set up in a region spanning multiple collision domains. As more and more WLANs are being deployed, there is a need to ensure some level of fairness concerning the amount of traffic each WLAN can support. In this paper, we look at the use of channel selection to achieve fairness. We have shown that using information gathered from networks within the communication range of a WLAN is not sufficient. As a result, we have described a number of channel selection schemes that make use of a more

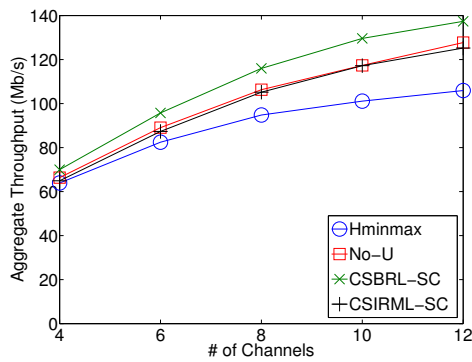


Figure 13: Aggregate throughput for different number of channels.

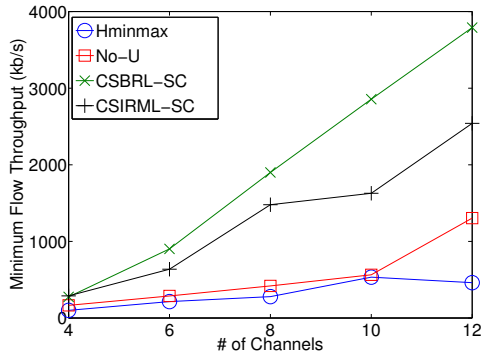


Figure 14: Minimum flow throughput for different number of channels.

accurate assessment of the channel condition by using a game theoretic learning approach. We also introduce an innovative method for a WLAN to detect that it is unfairly causing starvation to a neighboring network and have incorporate this capability into the learning schemes. This has resulted in socially conscious channel selection schemes, which we have shown through simulations to perform better than existing schemes in providing a higher system fairness, aggregate throughput and minimum flow throughput.

8. REFERENCES

- [1] M. Achanta. Method and Apparatus for Least Congested Channel Scan for Wireless Access Points. International Patent, April 2006. International patent WO/2006/042217.
- [2] A. Akella, G. Judd, S. Seshan, and P. Steenkiste. Self-Management in Chaotic Wireless Deployments. In *ACM Mobicom '05*, August 2005.
- [3] J. K. Chen, G. de Veciana, and T. S. Rappaport. Improved Measurement-Based Frequency Allocation Algorithms for Wireless Networks. In *IEEE Globecom '07*, November 2007.
- [4] M. W. R. da Silva and J. F. de Rezende. A Dynamic Channel Allocation Mechanism for IEEE 802.11 Networks. In *International Telecommunications Symposium (ITS '06)*, September 2006.
- [5] M. A. Ergin, K. Ramachandran, and M. Gruteser. An Experimental Study of Inter-Cell Interference Effects on System Performance in Unplanned Wireless LAN Deployments. *Computer Networks Journal (Elsevier)*, 52(14):2728–2744, October 2008.
- [6] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.
- [7] M. Garetto, T. Salonidis, and E. W. Knightly. Modeling Per-flow Throughput and Capturing Starvation in CSMA Multi-Hop Wireless Networks. *IEEE/ACM Transactions on Networking*, 16(4):864–877, August 2008.
- [8] A. Greenwald, E. J. Friedman, and S. Shenker. Learning in Network Contexts: Experimental Results from Simulations. *Games and Economic Behavior*, 35:80–123, 2001.

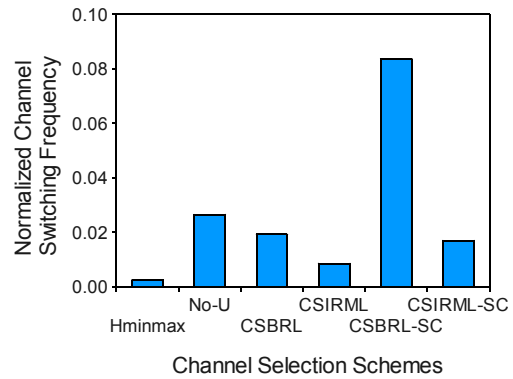


Figure 15: Channel switching frequencies for the different schemes.

- [9] S. Hart and A. Mas-Colell. A Simple Adaptive Procedure Leading to Correlated Equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
- [10] R. K. Jain, D.-M. Chiu, and W. R. Hawe. A Qualitative Measure of Fairness and Discrimination for Resource Allocation in Shared Computer System. Technical Report DEC-TR-301, Eastern Research Lab, September 1984.
- [11] D. J. Leith, P. Clifford, and D. Malone. WLAN Channel Selection Without Communication. Technical report, Hamilton Institute, National University of Ireland, 2006.
- [12] J. G. Lim, C. T. Chou, and S. Jha. Non-Cooperative Coexistence of Co-located Independent Wireless Mesh Networks. In *IEEE Conference on Mobile Adhoc and Sensor Systems (MASS '07)*, October 2007.
- [13] A. Mishra, S. Banerjee, and W. Arbaugh. Weighted Coloring based Channel Assignment for WLANs. *Sigmobile Mobile Computing and Communications Review*, 9(3):19–31, July 2005.
- [14] A. Mishra, V. Brik, S. Banerjee, A. Srinivasan, and W. Arbaugh. A Client-Driven Approach for Channel Management in Wireless LANs. In *IEEE Infocom '06*, April 2006.
- [15] A. Mishra, V. Shrivastava, D. Agrawal, S. Banerjee, and S. Ganguly. Distributed Channel Management in Uncoordinated Wireless Environments. In *ACM Mobicom '06*, September 2006.
- [16] C. Na, J. K. Chen, and T. S. Rappaport. Measured Traffic Statistics and Throughput of IEEE 802.11b Public WLAN Hotspots with Three Different Applications. *IEEE Transactions on Wireless Communications*, 5(11):3296–3305, November 2006.
- [17] N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [18] E. Rozner, Y. Mehta, A. Akella, and L. Qiu. Traffic-Aware Channel Assignment in Enterprise Wireless LANs. In *IEEE International Conference on Network Protocols (ICNP '07)*, October 2007.
- [19] Scalable Networks Technologies, Inc. Qualnet 3.9 User's Guide, 2005.
- [20] V. Srivastava, J. Neel, A. B. MacKenzie, R. Menon, L. A. DaSilva, J. E. Hicks, J. H. Reed, and R. P. Gilles. Using Game Theory to Analyze Wireless Ad Hoc Networks. *IEEE Communications Surveys*, 7(4):46–56, Fourth Quarter 2005.
- [21] C.-F. Wong, S.-H. G. Chan, and J. Chen. PACA: Peer-Assisted Channel Assignment for Home Wireless LANs. In *IEEE Globecom '06*, November 2006.
- [22] K. Yoo and C.-K. Kim. A Channel Management Scheme for Reducing Interference in Ubiquitous Wireless LANs Environment. In *International Conference on Multimedia and Ubiquitous Engineering (MUE '08)*, April 2008.