

# UnChannelize the Channels in WLAN

Yuan Yuan  
University of Maryland  
College Park, MD 20742

Paramvir Bahl, Ranveer Chandra,  
Thomas Moscibroda and Yunnan Wu  
Microsoft Research, Redmond, WA 98052

## I. INTRODUCTION

IEEE 802.11-based WLANs use a simple, fixed channelization structure, which divides the available spectrum into a set of channels of equal bandwidth. We argue that the fixed channelization has severely constrained the total capacity and fairness of WLANs. Our argument stems from recent measurements of WLANs in real world [4], [2], [3], which show significant spatial and temporal variations in user and traffic load. More specifically, user populations served by different access points (APs) in a WLAN fluctuate considerably over time, and are extremely unbalanced. Certain APs become hotspots while others remain unused. The fixed channel structure offers very limited flexibility in handling such variations.

We propose a radically new channelization structure. Our scheme dynamically creates the appropriate number of channels to accommodate all interfering APs, and adjusts the channel bandwidth (channel-width) depending on the instantaneous traffic loads at the AP and its the neighboring APs. The dynamic channelization is enabled by recent advances in radio technologies, such as cognitive radios, which make it possible to reconfigure the key operating parameters of radios, including center-frequency, bandwidth, and power, with a very low time overhead [5].

The structure of the paper is as follows: The next section presents a motivating example, and Section III describes the system design for creating dynamic channels. We conduct the simulations using real-world traces. Our results, presented in Section IV, show that our dynamic scheme significantly increases network capacity and improves per-client fairness.

## II. MOTIVATING EXAMPLE

We illustrate the advantages of the dynamic channelization through the following example. Fig. 1 shows a scenario with four APs within mutual interference distance of one another. In case 1,  $AP_1$  has 6 clients,  $AP_3$  has 3 clients, while the remaining two APs have one client each. In case 2, client A moves away from  $AP_2$  and associates to  $AP_4$ . We compare the performance of using fixed channels (F) with dynamic channels (D). In the fixed channel case, the spectrum is divided into 4 channels of 20 MHz each. Using the dynamic channels, channel-width may be 10, 20, or 40 MHz. Table I lists the bandwidth received *per client* at each AP. Also included is the total bandwidth used (B), and Jain's fairness index (FI), which is calculated using  $(\sum c_i)^2 / n \sum c_i^2$ .  $c_i$  is the bandwidth obtained by client  $i$ , and  $n$  is the total number of clients.

In case 1, the fixed channels lead to severe un-fairness among different clients. A client in the crowded location ( $AP_1$ ) receives 1/6 of unit bandwidth while the clients served by  $AP_2$  and  $AP_4$  receiving 1. In contrast, with an allocation of 40 MHz to  $AP_1$ , 20 MHz to  $AP_2$  and 10 MHz to the remaining APs, fairness improves significantly to 0.97. The dynamic channelization is not only important for fairness, but also for system capacity. For instance, in case 2 if client A moves from  $AP_2$  to  $AP_4$ , a dynamic approach can reallocate the 10 MHz spectrum formerly used by  $AP_2$  to  $AP_4$ , and thus improve the network capacity by 33%.




Fig. 1. A network with four mutually interfering APs. Using fixed channels, both throughput and fairness are suboptimal.

TABLE I  
BANDWIDTH RECEIVED BY EACH CLIENT (NORMALIZED BY 20MHz)

Scenario	$AP_1$	$AP_2$	$AP_3$	$AP_4$	B	FI
Case 1:F	1/6	1	1/3	1	4	0.58
Case 1:D	2/6	1/2	1/3	1/2	4	<b>0.97</b>
Case 2:F	1/6	X	1/3	1/2	3	0.82
Case 2:D	2/6	X	1/3	1/2	<b>4</b>	<b>0.97</b>

Our study on real-world traces reveals that in a large WLAN, fairness and capacity problems illustrated in Fig. 1 occur frequently. These scenarios motivate the need for dynamic channels in next-generation WLANs. By dynamically allocating variable-width channels, the dynamic approach is able to cope with the variations in user demands. If few APs are in the system, the structure creates channels with larger widths, enabling the clients to communicate at a higher speed. This new channelization provides better fairness than IEEE 802.11, because heavily-loaded APs get a larger spectrum to balance the per-client throughput across the network.

## III. DYNAMIC CHANNELIZATION STRUCTURE

We consider a WLAN in which all APs are connected to a central controller in a backbone network. Each AP reports the interfering APs and its traffic load, measured by the client population, to the central controller. Periodically, the central controller runs an algorithm that creates an optimal or near-optimal channel allocation. Then, the APs and clients reconfigure their radio parameters to reflect the new allocation. In order to make the dynamic channels compatible to the new hardware [5], we restrict that the spectrum assigned to one channel to be contiguous. The channel-width options are discrete and limited. Furthermore, we currently make the conservative assumption that overlapping channels always interfere, and we seek to assign non-overlapping frequency interval to any two potentially interfering APs.

Putting aside other practical concerns, we focus on the allocation algorithm running in the central controller. We present a heuristic algorithm that decides the center-frequency and the channel-width for APs as a function of their loads. The algorithm is based on the greed-pack subroutine shown in Alg. 1. This greedy-pack routine takes as its input an ordering of the APs and a channel-width requirement for each AP. It then proceeds in the given order and, when considering  $AP_i$ , greedily attempts to pack a non-overlapping frequency interval of channel-width  $B_i$  into the spectrum. The intervals are packed at the lowest possible frequency at which the

---

**Algorithm 1** *greedy-pack*( $B_1, \dots, B_N, \mathcal{O}$ ) Routine

---

Input: Channel-width options  $B_1, \dots, B_N$  and an ordering  $\mathcal{O}$  of APsOutput: If possible, a non-overlapping packing of various channel-widths into the available spectrum.  
Return false if no packing is found.

- 1: In the order of  $\mathcal{O}$ : **for each**  $AP_i \in V$  **do**
  - 2:   pack an interval of channel-width  $B_i$  in the lowest possible non-overlapping frequency.
  - 3: **end for**
  - 4: **if** the interval of all APs was successfully packed within the total bandwidth  $[F_{bottom}, F_{top}]$  **then**
  - 5:   **return** for each  $AP_i \in V$  its starting frequency  $S_i$  in the successful packing.
  - 6: **else return** false
  - 7: **end if**
- 

---

**Algorithm 2** GreedyRaising Algorithm

---

Input: An ordering  $\mathcal{O}$  of APs, the load  $D_i$  at  $AP_i$ 

Output: A non-overlapping packing of channels of variable width in the available spectrum.

- 1: Set parameter  $\theta := 1$  and let *successful* := *FALSE*;
  - 2: **while not successful** **do**
  - 3:   Let  $\phi'_i := \theta \cdot D_i / (D_i + \sum_{j \in N(i)} D_j)$  for each  $AP_i \in V$ .
  - 4:   Let  $B_i$  be the largest channel-width option s.t.  $B_i \leq \phi'_i \cdot B$
  - 5:   *successful* := *GreedyPack*( $B_1, \dots, B_N, \mathcal{O}$ ).
  - 6:    $\theta := \theta/2$ ;
  - 7: **end for**
  - 8: In the order of  $\mathcal{O}$ : **for each**  $AP_i \in V$  **do**
  - 9:   Let  $\hat{B}_i$  be the next higher channel-width option of  $B_i$ .
  - 10:   *successful* := *GreedyPack*( $B_1, \dots, \hat{B}_i, \dots, B_N, \mathcal{O}$ ).
  - 11:   **if** *successful* = *true* **then**  $B_i := \hat{B}_i$ .
  - 12: **end for**
- 

interval is non-overlapping with any interval previously assigned to the neighbors. The subroutine returns false if the desired channel-widths are too wide to be packed in the total bandwidth.

The heuristic algorithm called GreedyRaising is shown in Alg. 2. GreedyRaising starts from a feasible initial assignment, and “probes” APs one-by-one to check whether greedy-pack remains successful if the AP’s channel-width is raised. More specifically, GreedyRaising considers all APs in a given sequence  $\mathcal{O}$ , which starts from AP with the heaviest to the one with the lightest load. When considering an AP, its channel-width is increased to the next higher channel-width option, and the greedy-pack subroutine is called in order to see whether it still succeeds. If it does, the higher channel-width is adopted; if not, its channel-width is reset to its original value. In addition, we propose 3 algorithms, called ILP, LP, and LP2 that are based on integer linear programs and linear program formulations of the problem. The ILP program determines a theoretically optimal solution, but it is computationally inefficient. LP and LP2 require a linear program solver to deliver a near-optimal solution.

#### IV. SIMULATION STUDY

We study the performance of dynamic channels in a large WLAN. The available spectrum is 86 MHz, i.e., the size of 2.4 GHz ISM band, and the set of bandwidth options is 5, 10, 20, 40 MHz. We reconstruct the number of clients associated to each AP from the real traces in [4]. In these traces, 50% of APs serve less than 5 users, while 10% of APs serve over 15 users. The average client size served by each AP is 8. Since the traces do not record locations of APs and clients, we randomly place the 50 APs in a flat area of 1000m x 1000m. For each AP, we randomly place the client within its transmission range. The clients are static. We study two different scenarios: a 20 AP WLAN and a 50 AP WLAN. We vary the interference among APs by changing the transmission power




Fig. 2. Throughput and fairness in a WLAN of 20 APs




Fig. 3. System throughput and per-client collisions in a WLAN of 50 APs

from -1.6 dbm to 4.2 dbm. All results are averaged over 20 runs. We compare the GreedyRaising (GR) algorithm with RaC [1], the well-known channel assignment scheme. RaC uses 4 20-MHz channels.

Fig. 2 shows the throughput and fairness index of all clients in sparse and dense deployments of 20 APs. We emulate a sparse case by reducing the power of each AP to -1.6 dbm, such that each AP has 2 to 3 neighboring APs. We use ILP to compute the optimal solution as reference. In this scenario, GR achieves 47% more throughput than RaC. This can be explained by GR’s attempt to allocate all the available spectrum to the APs. RaC is unable to utilize all the channels, as each AP might not have sufficient interfering neighbors. As the AP density increases, GR allocates bandwidth to APs proportional to the number of clients associated to it, which achieves significantly better per-client fairness.

Fig. 3 shows the system throughput in a WLAN of 50 APs with the increase of AP density. GreedyRaising achieves up to 150% throughput improvement over RaC. The throughput obtained by RaC decreases substantially with the increased interferences among APs. This can be explained by the second graph, which plots the number of collisions per client with the increase of interfering APs.

#### V. CONCLUSION

We argue that by using dynamic channels of varying channel-width, a significant increase of both system capacity and per-client fairness can be achieved. We suggest that among other parameters, WLAN designers should consider the channel-width as a configurable parameter in the design of future WLANs.

#### REFERENCES

- [1] A. Mishra, V. Brik, S. Banerjee, A. Srinivasan and W. Arbaugh. Client-driven Channel Management for Wireless LANs. *INFOCOM*, 2006.
- [2] D. Kotz and K. Essien. Analysis of a Campus-wide Wireless Network. *MobiCom*, 2002.
- [3] D. Tang and M. Baker. Analysis of a Local-Area Wireless Network. *MobiCom*, 2000.
- [4] M. Balazinska and P. Castro. Characterizing Mobility and Network Usage in a Corporate Wireless Local-Area Network. *MobiSys*, 2003.
- [5] Y. Yuan, P. Bahl, R. Chandra, P.A. Chou, J.I. Ferrell, T. Moscibroda, S. Narlanka and Y. Wu. KNOWS: Kognitive Networking Over White Spaces. *Dyspan*, 2007.