RICE UNIVERSITY

Dual Wi-Fi: Dual Channel Wi-Fi for Congested WLANs with Asymmetric Traffic Loads

by

Adriana B. Flores Miranda

A THESIS SUBMITTED
IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE

Master of Science

APPROVED, THESIS COMMITTEE:

Dr. Edward W. Knightly, *Chair*
Professor of Electrical Engineering

Dr. Ashutosh Sabharwal
Professor of Electrical Engineering

Dr. Lin Zhong
Associate Professor of Electrical Engineering

HOUSTON, TEXAS
DECEMBER 2013

Abstract

Dual Wi-Fi: Dual Channel Wi-Fi for Congested WLANs with Asymmetric Traffic Loads

by

Adriana B. Flores Miranda

In many WLANs scenarios, the load transmitted from the AP to the clients (Downlink), far outweighs traffic demand from the clients to the AP (Uplink), thereby yielding *traffic asymmetry.* Moreover, when many clients associate with a single AP, the clients can cause a disproportional amount of medium contention compared to the AP, producing *contention asymmetry.* We present Dual Wi-Fi, a MAC architecture and protocol that minimizes MAC overhead by matching spectrum resources to traffic asymmetry. Dual Wi-Fi separates uplink and downlink data traffic into two variable-width independent channels, each allocated in accordance to the network's traffic demand. Our experimental and simulation results demonstrate that Dual Wi-Fi matches downlink vs. uplink throughput ratio to demand ratio within 1% under any client density and traffic load. Through this matching capability, Dual Wi-Fi offers unbounded downlink gains as congestion increases, minimizing and in some cases eliminating retransmissions and contention time.

Acknowledgements

I express my gratitude to my committee members, Dr. Knightly, Dr. Sabharwal and Dr. Zhong, for their careful and valuable examination and their advice that has fine-tuned this work. Foremost, I would like to thank my advisor, Dr. Edward Knightly, for always being an inspiring figure that has always shown me the correct path, through encouragement and infinite support that has thought me that through careful analysis and hard work, goals can be reached. Additionally, I thank Dr. Knightly for believing in me and pushing me to create the best work I can.

I extend my gratitude to my research group Rice Network Group: Ryan Guerra, Naren Anand, Oscar Bejarano and Eugenio Magistretti for their support, guidance and advice. Specially, Oscar and Eugenio for helping me shape my work by reviewing the early steps of my work and providing technical support with NS2 hacking. Furthermore, I would like to thank Chris Hunter for his infinite help, always providing me with perspective, technical support, and for being a person I can discuss and review concepts always obtaining significant feedback. This work would have never been possible without the assistance and advice of all these colleagues.

Finally, I would like to thank the most important people in my life, my family. I thank my parents, Pedro and America, and my brother and sister, Pedro and Margarita, for their unconditional encouragement and faith in me and my goals. For inspiring me and supporting me to reach any milestone I set for myself. Lastly, and most importantly, I would like to thank Gareth Middleton, for being my rock, for supporting me throughout this work, not only technically by reading every draft and reviewing ideas and concepts with me, even after a long workday. But also supporting me emotionally, helping me get through any difficult moment. I thank him, for always believing in me and making me a stronger person, without his support this work would not have been achievable. Their unconditional encouragement and advice has meant more to me than I could ever express. I dedicate this work to my family: Gareth, Margarita, Pedro, Mama y Papa.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Server-based multimedia streaming applications (e.g., YouTube, Netflix, Pandora) and many client-server applications such as web browsing yield highly asymmetric traffic demand with the vast majority of traffic being transmitted from the server to the client as opposed to vice-versa. Unfortunately, WLAN technologies such as IEEE 802.11 are oblivious to *traffic asymmetry*: First, uplink and downlink traffic share a common channel over time. Second, CSMA targets to provide all nodes (AP or client) with equal medium access probability such that with N clients, CSMA targets the AP transmissions to receive $1/(N+1)^{th}$ of spectrum resources. Finally, since each AP can serve a large number of clients, the clients create a disproportionate amount of contention yielding *contention asymmetry* between uplink and downlink traffic.

In this paper, we present the following two contributions. First, we introduce Dual Wi-Fi, a novel architecture and MAC protocol that provides spectrum independence between uplink and downlink MAC data traffic. In contrast to FDD, Dual Wi-Fi's channels are physically bi-directional and separate traffic between two channels according to the direction of *data*. Namely, we create two configurable-width physical channels: the "downlink data channel" carries MAC-layer data from the AP to clients as well as the corresponding ACKs flowing in the opposite direction; the second "uplink data channel" carries data originating from clients as well as the ACKs associated with that client data (sent by the AP).

Our channel architecture separates uplink from downlink MAC data traffic, removing competition for the same spectrum resources by allocating each traffic direction into one of two independent and asynchronous channels, and thereby removing *contention asymmetry*. Moreover, Dual Wi-Fi matches spectrum resources to the network's underlying *traffic asymmetry*, enabling support of a network-operator targeted division of spectrum resources to client vs. server generated traffic. Isolating uplink and downlink spectrum resources and matching spectrum resources to the network's traffic asymmetry allow Dual Wi-Fi to increase performance and spectrum efficiency by largely reducing contention overhead and retransmissions.

Second, we implement a prototype of Dual Wi-Fi on commodity Wi-Fi hardware as well as on the NS-2 simulator. We also implement a simple two channel Wi-Fi scheme modeling WiFi-NC [1] and use IEEE 802.11a [2] on both platforms as baselines for comparison. We modify commodity hardware drivers to support 5/10/20 MHz, and target to maintain equivalent spectral efficiency across the different bandwidths by adapting transmit power accordingly. Furthermore, we calibrate our simulations based on hardware measurements and tune bandwidth-dependent parameters.

Dual Wi-Fi's objective is to provide throughput performance to APs and clients in proportion to their traffic demand. Thus, ideally, the ratio of downlink to uplink *throughput* is the same as the ratio of downlink to uplink *offered load*. Our evaluation demonstrates that 802.11 and WiFi-NC fail to equalize throughput ratio and load ratio under any traffic load in a WLAN with 20 clients and 1 AP. For example, in the case of *symmetric traffic*, where downlink load is equal to uplink load, 802.11 and WiFi-NC provide 78% of throughput to the uplink traffic and 22% to downlink traffic. To the contrary, Dual Wi-Fi provides 50.97% to downlink traffic and 49.03%

to uplink traffic, <1% away from the ideal division of 50/50. Furthermore, in the case of a downlink vs. uplink asymmetry load of 4, Dual Wi-Fi nearly achieves the ideal throughput ratio of 4 (80.54/19.46), whereas WiFi-NC and 802.11 have a throughput ratio of approximately 2.4 (70.82/29.18). Nevertheless, our small-scale evaluation showed all three systems provide throughput performance to APs and clients that is proportional to their traffic demand. This reveals that a key limitation for 802.11 systems to match downlink vs. uplink performance to offered load is the disproportionate contention between uplink and downlink traffic caused by client density, i.e. contention asymmetry. Our evaluation demonstrates that downlink performance severely degrades with increasing client density for both 802.11 and WiFi-NC. For example, in a network with 20 clients with symmetric traffic, downlink traffic is significantly underserved due to uplink traffic repeatedly winning contention and utilizing approximately 25% more bandwidth than desired. This situation aggravates as the number of clients increases. For example, in the case of 100 clients, 802.11 and WiFi-NC allow downlink traffic to only utilize 10% of the band. In contrast, Dual Wi-Fi approximates the ideal throughput ratio (50/50 in the case of symmetric traffic) within 1% in all network densities.

High contention between uplink and downlink MAC level traffic is the main cause for baseline systems to deviate from ideal throughput to load ratio. WLANs with symmetric traffic or asymmetric traffic under high client densities have highly disproportional uplink vs. downlink contention and therefore present the most deviation form ideal throughput to load ratio. Consequently, performance independence between uplink and downlink traffic is the key element for Dual Wi-Fi's significant gains.

Furthermore, we explore Dual Wi-Fi's division of bandwidth to achieve the desired downlink vs. uplink performance proportion for a given traffic asymmetry and client density. Through exhaustive exploration we determine the value of the bandwidth division parameter that provides Dual Wi-Fi with a performance metric closest to 1, i.e. the optimal bandwidth division. Our evaluation shows that optimal bandwidth division of downlink to uplink spectrum is slightly sensitive to client density. Thus, increasing bandwidth is required for uplink traffic to counter increasing contention in dense networks. We observed that for symmetric traffic the optimal bandwidth division, in the case of 5 clients is 10/10 MHz (DL/UL); however in the case of 40 clients is 8/12 MHz. With high contention, the uplink channel requires a greater bandwidth allocation to counter the negative effect of high uplink contention and provide a downlink vs. uplink throughput ratio of 1.

Dual Wi-Fi's ability to provide proportional performance under any load asymmetry or network density translates into significant gains. We observe that Dual Wi-Fi achieves the highest gain in downlink throughput, up to 626% for 100 clients with symmetric traffic. These significant gains are obtained in high congestion scenarios and thus result from the capability of Dual Wi-Fi to maintain downlink performance across all client densities. In contrast, 802.11 and WiFi-NC downlink performance rapidly decays with network density, yielding Dual Wi-Fi increased relative performance. Because Dual Wi-Fi provides a downlink-only channel, we achieve up to 90% reduction of downlink's average backoff time. Furthermore, Dual Wi-Fi reduces backoff time of all transmissions by 41% to 52% under any traffic asymmetry, meaning Dual Wi-Fi's medium can be utilized approximately 50% more for data transmissions. Moreover, Dual Wi-Fi completely removes downlink retransmissions in the case of small number of APs and by 80% in other cases. Overall, uplink transmissions are reduced by 5 to 25% and on aggregate retransmissions are reduced by approximately 20 to 30%, values varying based on client density and traffic asymmetry. Dual Wi-Fi retransmission reductions are result of isolating uplink and downlink contention.

Chapter 2

Background

2.1 Medium Access Control

Medium access control protocols can be abstracted to be responsible for two tasks: first, to allocate spectrum to each traffic direction and second, to coordinate the access to the allocated spectrum resources. In other words, medium access control is responsible for determining what time and which spectrum resources a node will access.

2.2 Allocation of Spectrum Resources in Practice

Spectrum today is allocated in fixed sized frequency blocks that serve either licensed or unlicensed services [3]. This spectrum is mostly utilized in contiguous blocks, however due to the increasing spectrum demand non-contiguous spectrum allocation has been introduced.

Contiguous Spectrum. Commonly, the complete contiguous spectrum block is allocated to nodes' transmissions. For example, 802.11 uses spectrum blocks of 10, 20, 40, 80 and recently 160 MHz for transmissions. One node uses the complete block of spectrum at a given instant, thus nodes share the complete blocks of spectrum in time.

Moreover, to provide a more efficient use of large spectrum blocks, such as the 160 MHz channel bandwidth of 802.11ac, different approaches have been introduced to divide the contiguous spectrum blocks into sub-channels to allow multi-channel access [4, 1]. The key issue in operating multiple channels in a contiguous frequency block is the feasibility of implementing asynchronous operation of the sub-channels. The common solution to this issue is the implementation of *guard bands* to eliminate inter-channel interference. Determining guard band size has been widely studied [5, 6, 1]. Some studies have shown that the dimension of a guard band is determined by multiple factors such as radio hardware, network topologies, and location of wireless links [5]. Other studies have demonstrated sub-channel isolation with narrow guard bands of 100 KHz, through the usage of cognitive radios and digital filters. Overall, contiguous spectrum is utilized either in its unitary form or fragmented with guard bands in between the sub-channels.

Non-contiguous Spectrum. New approaches for spectrum allocation try to remove the fixed block approach and instead attempt to use spectrum as a shared medium without fixed licensed or unlicensed allocations [7]. The implementation of spectrum sharing requires the usage of non-contiguous spectrum allocation. In order to achieve high rate transmissions in fragmented spectrum, we require to use channel aggregation to form wide channels that can support high data rates. An example of fragmented spectrum exists in white spaces, where scattered 6 to 8 MHz channels are available in a large range of spectrum, and channel aggregation is required to form wider channels and higher transmission rates.

Dual Wi-Fi can be implemented in contiguous or non-contiguous spectrum. In order to utilize Dual Wi-Fi in contiguous spectrum, the downlink data channel must be isolated from the uplink data channel through the use of a guard band. Moreover, Dual Wi-Fi can be implemented on non-contiguous spectrum, by allocating the downlink data channel to a set of sub-channels in one part of the frequency band and the uplink data channel on another set of sub-channels in a different part of the frequency band.

2.3 Medium Access Coordination in Practice

Two main approaches exist to grant access to the medium: random access and scheduled access. Next, we provide a brief overview of the key characteristics of each approach.

Random Access. A decentralized form to grant usage to the medium is through random access. The key in this medium access is the ease of implementation. Because a random access approach is decentralized, it is a scalable and easy to deploy solution. However, a decentralized medium access has downsides. For example, in order to coordinate the nodes in a decentralized manner, each node is required to independently countdown to attempt to access the medium while it remains idle, and must stop the count down in the case that the medium becomes busy. This uncoordinated medium access may lead to collisions and retransmissions that cause inefficient use of the spectrum.

Scheduled Access. A centralized approach to grant usage of the medium is through scheduled access. A central unit determines a time slot in which a node is given access to the medium. Because this approach is centralized, there is no opportunity for collisions and thus once the schedule is formed, the medium usage is highly efficient. The significant downside to this approach is that the creation of the schedule access requires a lot of overhead and the implementation is complex.

Dual Wi-Fi grants medium access in a decentralized manner, following the existing random access techniques. The key Dual Wi-Fi brings to the random medium access approach is that we reduce the amount of collisions, retransmissions and backoff through the isolation of uplink and downlink data spectrum resources. In this way, we achieve efficient spectrum usage.

Chapter 3

Dual Wi-Fi Channel Architecture

Dual Wi-Fi's channel architecture addresses traffic asymmetry and contention asymmetry by separating uplink and downlink medium access contention and allowing physically bi-directional data-ACK handshakes within the same channel. In this chapter, we describe how these key features treat asymmetries and reduce medium contention leading to increased spectral efficiency. Moreover, Dual Wi-Fi enables network managers to set a predefined bandwidth division parameter to control the throughput of uplink and downlink traffic.

3.1 Separating Uplink and Downlink Medium Access Contention

3.1.1 Channel Design

Dual Wi-Fi decouples uplink and downlink medium access by allocating a *"downlink data channel"* for MAC-layer data originating from the AP and transmitted to the clients (downlink), with the corresponding ACKs flowing in the opposite direction. Likewise, the *"uplink data channel"* transmits data from the clients to the AP (up-

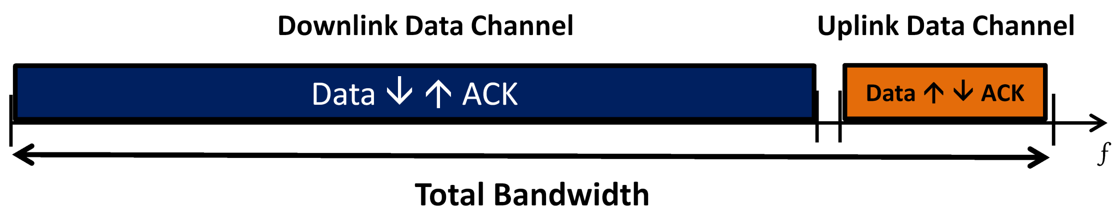

Figure 3.1: Dual Wi-Fi channel architecture: two bi-directional physical channels each carrying a single uni-directional MAC traffic direction, with configurable asymmetric spectrum resource allocation.

link) and the reciprocal ACKs from the AP to clients. Through this allocation, the downlink and uplink data traffic never compete for the same spectrum resources, allowing simultaneous, asynchronous and independent uplink and downlink MAC layer data-handshakes. Namely, the spectrum isolation per traffic direction is achieved through two configurable-width physical bi-directional channels that each carry a single uni-directional MAC data traffic.

Figure 3.1 illustrates the Dual Wi-Fi channel design. Dual Wi-Fi divides the total allocated bandwidth (spectrum) into two independent configurable-width channels: the "downlink data channel" and the "uplink data channel". Our channel architecture can be implemented on contiguous or non-contiguous frequencies. In the case of contiguous frequency, downlink and uplink channels require a guard band between them to eliminate inter-channel interference. The selection of guard band size has been widely studied [5, 6, 1]; prior work has shown a minimum guard band requirement of 100 KHz to isolate the channels [1]. Further background on guard bands is presented in Chapter 2. Moreover, Dual Wi-Fi channel design permits configurable bandwidth allocations. The configuration of the downlink vs. uplink bandwidth ratio is dependent on the network's traffic and contention asymmetry, as explored in Section 5.3.

3.1.2 Configurable Spectrum Isolation

The uplink and downlink spectrum isolation allows Dual Wi-Fi to address contention asymmetry by eliminating the uneven number of uplink and downlink devices competing for a single channel. Since downlink and uplink data traffic only contend with its own traffic direction, medium access contention is directly weighted on the traffic load of that direction. Furthermore, the spectrum isolation allows Dual Wi-Fi to flexibly allocate spectrum resources to either traffic direction, providing a simple way to provide more spectrum where needed.

Dual Wi-Fi provides a wide and configurable downlink channel. Based on the premise that downlink data traffic load differs from (and is typically greater than) uplink data traffic load, our design can allocate the downlink channel with more bandwidth, i.e. wide downlink channel and narrow uplink channel. This allows the downlink traffic to independently transmit more data than the uplink traffic at any instant. Overall, the amount of bandwidth assigned to the downlink traffic is flexible and can be configured to fit the traffic characteristics of any system. Asymmetric resource allocation can be used to benefit a dominant load and symmetric bandwidth allocation for systems with equal loads.

3.2 Bi-directional traffic within channels

The second feature of Dual Wi-Fi is physical bi-directional communication within each channel to support the complete MAC-layer data-ACK handshake. Specifically, MAC data traffic in an individual channel consists of the data flowing in one direction and acknowledgments (ACKs) only associated with that data flowing in the opposite direction. Unlike FDD systems, in Dual Wi-Fi no generic control messages use the channel, only the ACKs for the data flow traveling in that channel. Thus, this feature

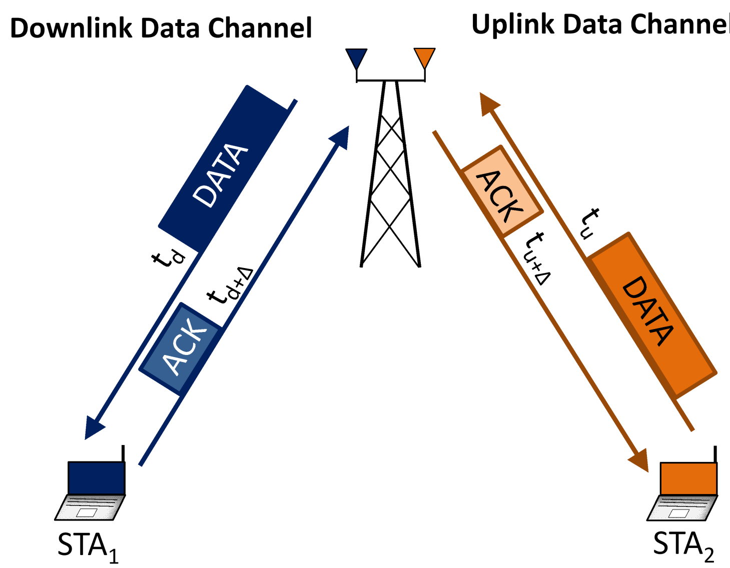

Figure 3.2: Bi-directional Data-ACK communication performed individually and asynchronously in each downlink and uplink data channels. Isolated downlink and uplink data channels, allow simultaneous transmissions even in clique networks.

provides in-channel control feedback that is paired with transmitted data allowing asynchronous and independent uplink and downlink MAC traffic operation.

Figure 3.2 depicts an example of a network operating our channel design in a clique of one AP and two stations (STAs), where each transmit in one of the two uplink or downlink data channels. In this example, STA_1 receives a downlink data packet from the AP in the downlink data channel and responds with the associated ACK in the same channel. Independently and asynchronously to AP-STA_1 downlink data transmission, STA_2 transmits an uplink data packet to the AP and receives the related ACK in the uplink data channel. Overall, this figure demonstrates the MAC level bi-directional communication performed in the frequency independent uplink and downlink data channels.

The fundamental benefit of physical bi-directionality is the capability for performing *independent and asynchronous* uplink and downlink MAC level data-ACK transmissions. Having the complete MAC data-ACK handshake in a single channel allows Dual Wi-Fi to adjust MAC and PHY level mechanisms to independently address uplink and downlink resource allocation. For example, if a STA receiving a downlink packet has the capability to achieve a high data rate transmission, while another STA has low quality channel to transmit an uplink packet, in our system, both transmissions can occur independently and asynchronously without affecting each other's performance. Moreover, uplink medium access factors such as hidden terminals will only affect uplink performance, while downlink transmission operate normally. Overall, the bi-directional capability of our channel design allows Dual Wi-Fi to provide independent performance between uplink and downlink MAC level traffic.

3.2.1 Contrast of Dual Wi-Fi

Pre-allocating spectrum resources and achieving isolation between the traffic directions, differentiates Dual Wi-Fi's medium access from 802.11 medium access techniques, where uplink and downlink traffic share a common channel over time.

Moreover, the physical bi-directional communication within each sub-channel contrasts Dual Wi-Fi from existing scheduling systems. An example is Frequency Division Duplex (FDD) utilized in cellular systems, where these implement two channels each allocated based on the physical direction of the transmission. The FDD channels perform uni-directional physical downlink (AP to STAs) or uplink (STAs to AP) transmissions, which means the division is performed at a physical level. On the contrary, Dual Wi-Fi divides the traffic at the MAC level, our downlink and uplink data channels are physically bi-directional and thus support the bi-directionality of MAC Data-ACK handshakes. The key difference of Dual Wi-Fi to FDD is the support of the complete MAC-layer data-ACK handshake within each channel. In FDD systems a data-ACK handshakes is performed in two channels, e.g. a downlink data transmission occurs in the physical downlink channel while the response ACK is transmitted on the physical uplink channel. This key difference brings numerous benefits to our system, such as the in-channel control feedback that is paired with transmitted data that allows asynchronous and independent uplink and downlink MAC traffic operation.

The channel bi-directionality allows our system to perform non-slotted, asynchronous transmissions in each channel without requiring synchronization and scheduling overhead of FDD systems. Moreover, FDD systems have static and equal bandwidth allocation between uplink and downlink traffic, whereas Dual Wi-Fi provides flexible and asymmetric spectrum allocation. Furthermore, the bi-directionality capability of Dual Wi-Fi allows in-channel feedback that can allow the implementation of MU-MIMO, key limitation of FDD systems.

As a distinction to multi-channel random access techniques such as FICA [4] and WiFi-NC [1], Dual Wi-Fi operates specifically with two variable-width channels, instead of multiple equal width narrow channels and more importantly our design allocates each of the channels to uplink and downlink MAC layer data traffic, removing contention between the two. Because these systems do not differentiate between uplink and downlink traffic, increase in spectral efficiency is limited by uplink and downlink contention within the narrow channels. The core distinction between these approaches and Dual Wi-Fi is that our system pre-allocates the sub-channels to downlink and uplink traffic and thus presents the capability to adapt spectrum resources to the traffic load, instead of presenting limited fixed width sub-channels.

Furthermore, our work differentiates from patent 20040213180 [8] through the specific allocation of downlink Data-ACK and uplink Data-ACK handshakes to the ”downlink data channel” and ”uplink data channel“ accordingly. Therefore, both of our channels, always operate bi-directional transmissions. More importantly, Dual Wi-Fi supports the configuration of bandwidth resources of each channel to match the system's traffic asymmetry or a desired downlink and uplink service.

3.3 Bandwidth Parameter Selection

Dual Wi-Fi presents a flexible boundary that is the channel bandwidth allocated to uplink and downlink data traffic. This boundary is selected based on the network's targeted uplink to downlink service or the measured traffic demand. The channel bandwidth allocated to uplink and downlink data traffic, referenced in this thesis as the bandwidth parameter or bandwidth ratio (DL/UL). To select the bandwidth allocated to either sub-channel, a network-operator estimates the realistic traffic demand of the network and targets the downlink to uplink bandwidth ratio to match the estimated downlink to uplink traffic asymmetry or a desired downlink and uplink service.

The results of a study we performed over 5 months on Rice University campus network that has excess capacity with 867 APs in an area of 1 km^2, suggest that the network-operator would select a bandwidth allocation giving the downlink channel 6.5x more spectrum resources than the uplink channel. This allocation is based on the measured *average* traffic demand of this network, which was observed to remain constant across multiple time frames (minutes, hours, days, weeks and months).

This study leads us to assume a static or slowly varying bandwidth division based on a network's *average* traffic demand will be appropriate and only be required to be varied at slow time scales. However, at different moments in time a particular bandwidth allocation will be mismatched to the instantaneous traffic demand of the network. In Section 5.3, we evaluate the robustness of the selected bandwidth division parameter setting to be mismatched to the instantaneous traffic asymmetry. This evaluation demonstrates that even a fixed bandwidth parameter setting can achieve significant gains over 802.11 uplink and downlink division of spectrum resources, solely based on the concept of statically allocating spectrum resources to the average traffic demand. Although not studied in this thesis, an adaptive measurement-based bandwidth selection scheme could be implemented in Dual Wi-Fi. We consider the design and implementation of an adaptive bandwidth parameter as part of our future work.

Chapter 4

Dual Wi-Fi Medium Access Control Design

The two key features of Dual Wi-Fi's channel architecture: *data*-based downlink and uplink spectrum isolation; and physically bi-directional channels for data-ACK handshakes within channels, allow the separation of the data-downlink MAC from the data-uplink MAC. This key feature provides more medium access opportunities for both uplink and downlink traffic.

4.1 Downlink Medium Access

Dual Wi-Fi's architecture ensures that AP's do not contend with clients. Thus, a single AP on the channel operates collision-free and performs minimal medium contention. Nonetheless, with the potential for multiple APs per channel (studied in Section 5.4), Dual Wi-Fi uses 802.11 CSMA in the downlink channel. This enables Dual Wi-Fi to obtain significant gains in downlink performance via a small number of contending nodes. That is, because the Dual Wi-Fi channel architecture isolates downlink and uplink medium access, the number of contending nodes in the downlink data channel is reduced to only the number of same-channel APs. Consequently, smaller number of nodes competing for the medium reduces coordination time, the number of collisions and retransmissions, and thus increases spectral efficiency. Moreover, in the case when only 1 AP is operating in the downlink channel, obtained by frequency planning, contention is close to eliminated, allowing the downlink data channel to approximate back-to-back data transmission achieving close to full channel utilization.

4.2 Uplink Medium Access

Similarly, because Dual Wi-Fi separates traffic directionally, medium access for client-to-AP data follows 802.11 CSMA basic access. Since Dual Wi-Fi's channel architecture removes uplink and downlink contention for a common channel, only STAs contend for the uplink data channel. This feature benefits uplink contention because uplink traffic does not share the medium with the heavy and continuous downlink traffic, providing it with more medium access opportunities. Thus, the gains Dual Wi-Fi achieves (shown in Chapter 5.4) for uplink traffic are reflection of the downlink to uplink contention.

4.3 Flexibility for Alternative Medium Access Mechanisms

We utilize IEEE 802.11 basic access as the base medium access to Dual Wi-Fi, however our system presents the flexibility to separately address uplink and downlink resource allocation. This feature permits to independently address any medium access issue found in either channel. Thus, in the presence of an uplink or downlink medium access issue, our system permits the implementation of an in-channel solution that will only be applied where required.

As an example, in the presence of uplink hidden terminals, Dual Wi-Fi allows the implementation of any existing in-channel hidden terminal MAC for the uplink data channel. An RTS-CTS mechanism such as [9] or an efficient control mechanism such as correlated symbol sequences [10] can be implemented in the uplink channel, but not the downlink channel as they are not required there. The benefit our system brings to the applied solutions is that the solution's overhead will only affect the performance of the channel it is applied to. In this case, only uplink packets (not downlink packets) will be required to include RTS-CTS overhead in each transmission.

It is important to note that Dual Wi-Fi provides the base channel architecture to which we apply 802.11 CSMA basic access. However, in this thesis we do not propose solutions to issues that arise in either uplink and downlink traffic, such as hidden terminals. Instead, we introduce a system that isolates uplink and downlink factors, allowing the system designer to separately address the issues found in either uplink or downlink transmissions.

Chapter 5

Dual Wi-Fi Implementation and Evaluation

In this section, we first present the implementation of Dual Wi-Fi. Moreover, we explore the performance impact of traffic and contention asymmetry through over-the-air experiments and simulations. Lastly, we study Dual Wi-Fi's bandwidth division parameter and explore the gains of our system.

5.1 Dual Wi-Fi Implementation

We implement a Dual Wi-Fi prototype on a commodity Wi-Fi experimental platform as well as on the NS-2 simulator. We also implement WiFi-NC [1] and use IEEE 802.11a [2] on both platforms as baselines for comparison. In particular, 802.11a serves as a baseline in which a full 20 MHz channel is shared in time by uplink and downlink traffic, as is common practice in commercial systems. Moreover, we implement WiFi-NC as a two-channel baseline in which the nodes split traffic evenly among two equal 10 MHz channels. Unlike Dual Wi-Fi, WiFi-NC's two sub-channels are fully bi-directional instantiations of 802.11, i.e., they split the channel but do not allocate the sub-channels to uplink and downlink data traffic, as done by Dual Wi-Fi.

5.1.1 Experimental Platform

We implement Dual Wi-Fi and the baseline protocols on a hardware platform comprising custom hardware running Debian Linux with Atheros wireless chipset version AR5413 supporting 802.11a/b/g and utilizing an ath5k wireless driver. In order to support 5/10/20 MHz channels, we implement a modified version of the ath5k driver. For the wireless interface, we utilize the dbii F-50 pro off-the shelf miniPICI Wi-Fi cards which operate in the 5 GHz band.[1] Due to standard specifications, commodity hardware presents a limitation on the granularity of bandwidth selection to only 5/10/20 MHz. We target to maintain an equivalent spectral efficiency across the channels by reducing the transmit power in relation to the bandwidth size and noise level. We set the transmit power to 22 dBm for a 20 MHz channel and reduce transmit power by 3 dBm for every bandwidth division.

5.1.1.1 Testbed Setup

We conduct our experiments with a five node indoor testbed set in a BSS topology in which 4 clients perform single-hop over the air transmissions to and from a single AP. All nodes are in carrier sensing range of each other with no client mobility. The 802.11a protocol operates under a 20 MHz channel with bi-directional traffic, such that in aggregate, uplink and downlink traffic saturate the channel. Similarly, WiFi-NC operates under two 10 MHz channels without co-channel interference, and each channel operates under saturation conditions with bi-directional traffic. Dual Wi-Fi operates in 5/10/20 MHz channels, and combines these bandwidths to match load asymmetries of 0.5, 1, 2 and 4. Co-channel interference and cross-talk across wireless interfaces is eliminated by separating transmissions in time. Refer to Table 5.1 for testbed parameters.

[1] http://www.dbii.com/f50-PRO.html

Parameter	20 MHz	10 MHz	5 MHz
Carrier Frequency:	5805 MHz	5785 MHz	5745 MHz
Transmit Power:	22 dBm	19 dBm	16 dBm
Contention Window:	15 to 1023		
RTS/CTS:	Inactive		
Capture:	Inactive		
Max Payload Rate:	54 Mbps		
Packet Size:	1500 bytes		
Traffic Pattern:	Fully Backlogged Flows		
Application:	UDP		

Table 5.1: Commodity hardware parameters.

5.1.2 Simulation Platform

We implement Dual Wi-Fi, 802.11a and WiFi-NC on the network simulator NS-2. For WiFi-NC and Dual Wi-Fi, we modified NS-2 to allow multiple interfaces and simultaneous multi-channel transmissions. Dual Wi-Fi employs two channels, each allocated to either downlink or uplink traffic, where each channel operates asynchronously at an independent data rate based on system's load asymmetry. WiFi-NC devices have two interfaces where each operates in one of two equal data rates channels.

5.1.2.1 Calibration

We calibrate the simulation parameters with over-the-air experiments as follows. We perform these experiments overnight with no fading and utilize 5 GHz interference-free channels, with non-line-of-sight transmissions. The detailed hardware parameters can be found on Table 5.1. We select the log-normal shadowing model in NS-2 to represent this environment. Figure 5.1 depicts the results in which both simulation and hardware platforms achieve the same Packet Delivery Ratio (PDR) at different received signal strengths. The parameters modified for the physical channel calibration were path loss exponent, shadowing deviation, receiver sensitivity, data rate and

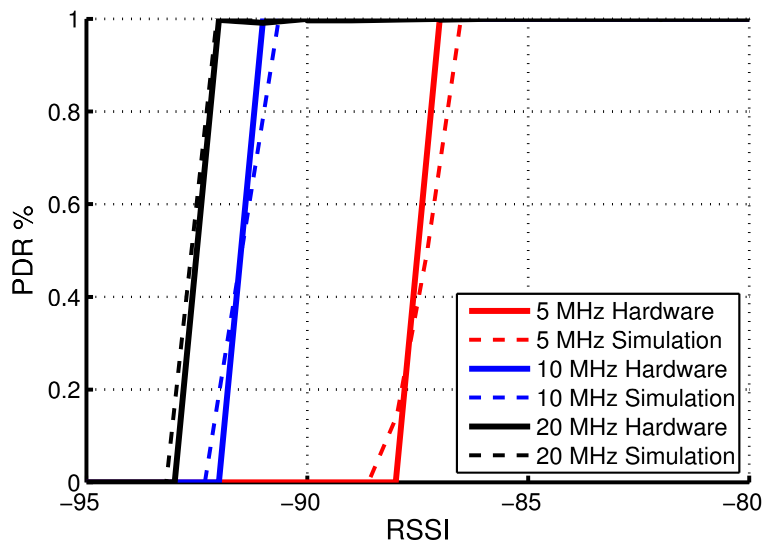

Figure 5.1: Packet Delivery Ratio (PDR) comparison between NS-2 simulation and commodity hardware.

transmit power. Moreover, we determined bandwidth-dependent parameters including receiver sensitivity, data rate and transmit power. Parameters such as slot time, SIFS time, CCA time and preamble rate are modified based on the specifications of the IEEE 802.11 Standard on Table 18-17 [2]. For further simulation parameters refer to Table 5.2.

5.1.2.2 Simulation Setup

The evaluation scenarios use BSS topologies where all nodes are in carrier sensing range of each other. A high density of clients communicates with 1 or several APs. Simulations follow 802.11a parameters [2], carefully adapted for any bandwidth selection. Traffic asymmetry is varied across evaluations, however full channel congestion is maintained. Uplink traffic is equally spread among the number of stations in the network. Uplink and downlink perform UDP transmissions with constant bit rates and packet size of 1500 bytes.

Parameter	20 MHz	10 MHz	5 MHz
Carrier Frequency:	5805 MHz		
Transmit Power:	22 dBm	19 dBm	16 dBm
Max Payload Rate:	20 Mbps	10 Mbps	5 Mbps
Receiver Threshold:	-93 dBm	-92 dBm	-88 dBm
SIFS:	16 μs	32 μs	64 μs
Slot Time:	9 μs	13 μs	21 μs
CCA:	4 μs	8 μs	16 μs
Contention Window:	15 to 1023		
RTS/CTS:	Inactive		
Capture:	Inactive		
Packet Size:	1500 bytes		
Traffic Pattern:	Fully Backlogged Flows, CBR		
Application:	UDP		

Table 5.2: NS-2 simulation parameters for 5/10/20 MHz. Parameters are adapted for in-between granular bandwidths.

5.2 Throughput Ratio and Load Ratio

Dual Wi-Fi's objective is to provide throughput performance to APs and clients in proportion to traffic demand. Thus, ideally, the ratio of downlink to uplink *throughput* is the same as the ratio of downlink to uplink *offered load*. We utilize the ratio of these two values (ideally 1) as a key metric for performance evaluation. Namely, a throughput ratio to load ratio of 1 corresponds to providing throughput that is directly proportional to demand. In this section we evaluate the impact of traffic and contention asymmetry on the relationship between throughput ratio and load ratio.

5.2.1 Impact of Traffic Load Asymmetry

In this experiment, we analyze the impact of traffic asymmetry (difference between downlink and uplink offered load or traffic demand) on network performance, as measured by the throughput to load ratio.

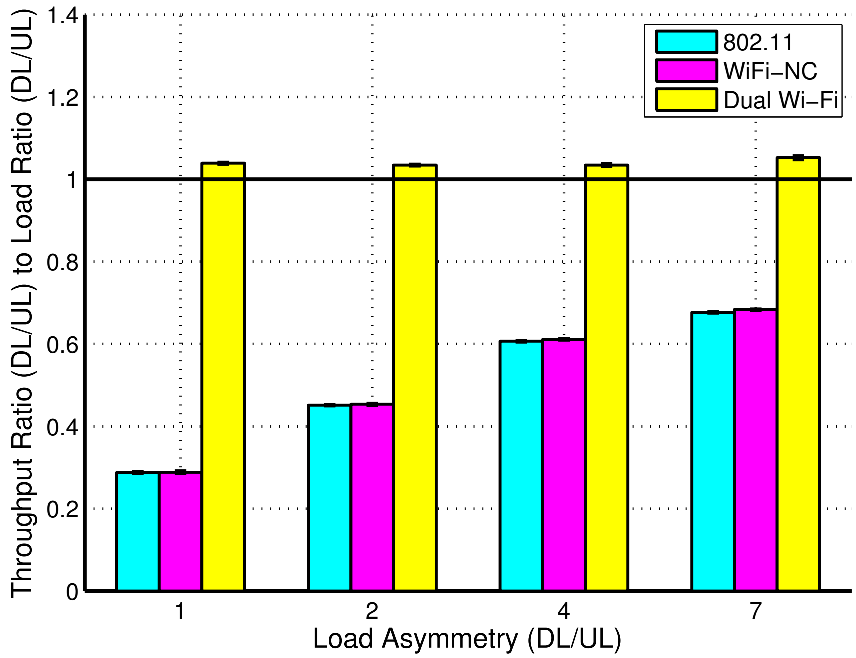

Figure 5.2: Proportion of throughput ratio (DL/UL) to load ratio (DL/UL) for 20 clients. Note that Dual Wi-Fi approximates ideal ratio (value of 1) for all load asymmetries.

Setup. For this experiment we fix the number of clients to 20 to represent a WLAN scenario. All clients communicate with a single AP and are within range of each other. For all experiments, the aggregate load is fixed to saturation, and the ratio of the downlink load to the uplink load is varied and depicted on the x-axis of Figure 5.2. For each of the 3 systems (802.11, Dual Wi-Fi and WiFi-NC), we measure the corresponding throughput for both downlink and uplink traffic. The resulting throughput ratio to load ratio is depicted on the y-axis of the figure. Any value above 1 on the y-axis indicates that more throughput is given to the downlink traffic than demanded, whereas downlink traffic is underserved when the value is below 1.

Results. First, we observe in Figure 5.2 Dual Wi-Fi achieves the closest to the ideal ratio (value of 1) under any load asymmetry. Dual Wi-Fi matches the throughput ratio within 1% of the demand ratio. In contrast, both 802.11 and WiFi-NC systems fail to equalize throughput ratio and load ratio. For example, when the demand for downlink and uplink traffic is identical, (corresponding to a value of 1 on the x-axis) the y-axis depicts a value of approximately 0.28 for both 802.11 and WiFi-NC. This indicates that both baseline systems have a throughput division of 0.28 corresponding to 21.87/78.13 (DL/UL) vs. 50.97/49.03 (DL/UL) from Dual Wi-Fi, <1% away from the ideal division of 50/50.

Second, the role of uplink vs. downlink contention is revealed as asymmetry varies. 802.11 and WiFi-NC systems fail to provide ideal throughput to load ratio under low traffic asymmetries but improve under high load asymmetries. For example, with load asymmetry (x-axis value) of 4, Dual Wi-Fi nearly achieves the ideal throughput ratio of 4 (80.54/19.46). In contrast, WiFi-NC and 802.11 have a throughput ratio of 2.4 (70.82/29.18) yielding a value of approximately 0.6 on the y-axis, slightly closer to the desired value 1, in comparison to low asymmetry. Thus, because the majority of the demand is from the AP, uplink vs. downlink contention is lower compared to load asymmetry of 1. For a particular demand ratio, the AP alone generates less contention than the clients in aggregate. Therefore, baseline systems improve under high asymmetries due to the reduction of uplink demand that allows more spectrum resources and medium access for downlink demand. This finding reveals that uplink vs. downlink medium contention has a critical effect on throughput to load proportion. This conclusion drives the key reason for 802.11 and WiFi-NC to attain nearly identical performance in these cases. Both systems allow uplink and downlink traffic to contend with each other and neither provides the performance isolation of Dual Wi-Fi, limiting these from achieving ideal throughput to load ratio.

Over-the-Air Validation. In order to validate the above simulations, we perform small-scale over-the-air experiments. We implement asymmetries of 0.5, 1 and

Load Asymmetry	0.5	1	2
Hardware	1.146 ($\pm$ 0.14)	1.051 ($\pm$ 0.09)	0.938 ($\pm$ 0.15)
Simulation	1.030 ($\pm$ 0.00)	0.972 ($\pm$ 0.00)	0.918 ($\pm$ 0.00)

Table 5.3: Mean and standard deviation of the throughput ratio (DL/UL) to load ratio (DL/UL) with 2 clients for experimental and simulation platforms.

2, under realistic channel conditions with a two client - one AP topology. Due to the coarse-grain bandwidth selection of commodity hardware, we are not permitted to compare baseline systems to Dual Wi-Fi. This is because we cannot provide the same amount of total bandwidth to all systems while adapting Dual Wi-Fi's bandwidth to a given traffic asymmetry (e.g. 10 MHz/5 MHz vs. 15 MHz).

Table 5.3 presents Dual Wi-Fi's relationship between throughput ratio and load ratio for both simulation and experimental platforms. Dual Wi-Fi approximates the throughput ratio to the demand ratio across all load asymmetries in both simulation and over-the-air experiments. Furthermore, we observe high proximity between all simulation and hardware results. For example, in the case of load asymmetry 1, the implementation achieved throughput to load ratios from 0.961 to 1.141 whereas simulation obtains 0.972 with negligible variance. Since all simulation results lie within one standard deviation of the hardware results, we consider our simulation results to be a comprehensive representation of over-the-air network performance.

5.2.2 Impact of Contention Asymmetry

Medium congestion (attempts to access the medium, collisions, etc.) increases as the number of backlogged clients increase. This can yield contention asymmetry in which the many clients provide disproportionate contention compared to the single AP. That is, with all nodes (AP or STAs) counting down from the same uniform distribution, many clients will have a greater chance to choose the lowest value and win the medium than will a single AP. This in turn can yield a greater probability for a client to gain access to the medium rather than the AP. Thus, here we evaluate the impact of contention asymmetry caused by client density.

Setup. We select an uplink to downlink symmetric load, in which downlink and uplink traffic demand is identical, because it presents the highest medium access congestion. Additionally, we select a load asymmetry of 4 as a representation of the impact contention asymmetry presents under asymmetric traffic. For this experiment we vary the network density to observe the effect of increasing contention asymmetry on the downlink to uplink throughput to load ratio. When multiple APs share the same spectrum, there is a need for downlink contention. Thus, we define topologies with 1 to 5 APs, all transmitting on the same channel, to introduce a higher downlink data contention to counteract the increasing uplink data contention. We vary the total number of clients between 20 to 100, and divide these equally across the number of APs. The throughput ratio for a given number of clients is obtained by averaging across iterations and the 1-5 in-channel APs evaluated for that given set of clients. All three protocols utilize the same total bandwidth. Dual Wi-Fi adapts the downlink to uplink bandwidth ratio based on the network's traffic load and client density.

Results for symmetric traffic. Figure 5.3 illustrates the impact of client density on throughput ratio to load ratio under symmetric traffic. Figure 5.3 indicates that downlink performance severely degrades with increasing client density for both 802.11 and WiFi-NC. Even in a network with 20 clients, downlink traffic in baseline systems is significantly underserved due to uplink traffic repeatedly winning contention. In the case of 20 clients, 802.11 and WiFi-NC uplink traffic utilizes approximately 25% more bandwidth than desired 25.57/74.43 (DL/UL) vs. Dual Wi-Fi's throughput division of 50.25/49.75 (DL/UL). This situation aggravates as the number of clients increase, for example in the case of 100 clients baseline systems provide a throughput

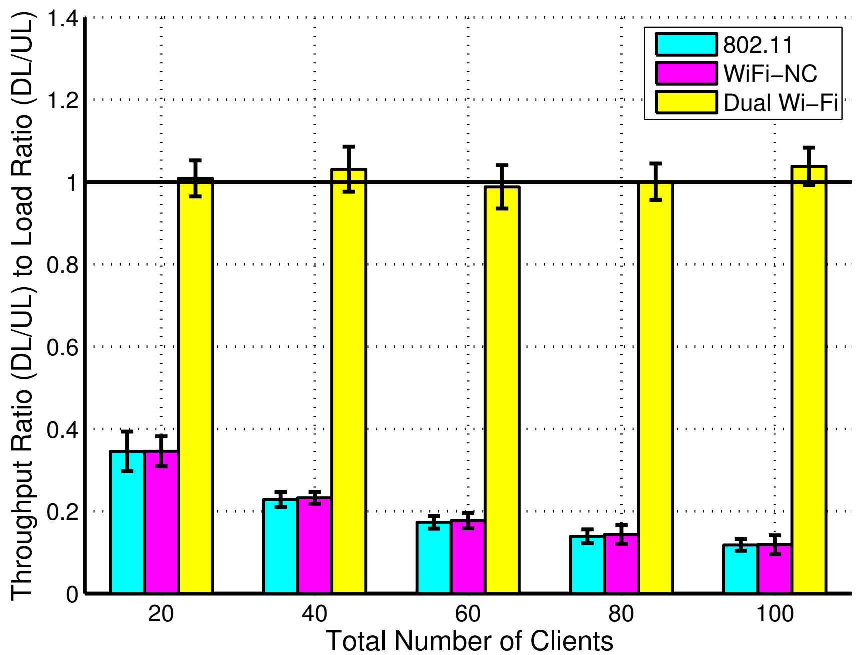

Figure 5.3: Mean throughput ratio (DL/UL) to load ratio (DL/UL) with symmetric load for different network densities.

ratio (DL/UL) of 0.118 corresponding to a 10.52/89.48 (DL/UL) spectrum division, meaning uplink traffic is using close to 40% of downlink's intended throughput and thus allowing downlink traffic to only utilize 10% of the band. To the contrary, Dual Wi-Fi in the case of 100 clients has a throughput division of 50.98/49.02 (DL/UL), less than 1% from ideal division of 50/50.

Dual Wi-Fi is able to approximate the ideal throughput ratio within 1% in all network densities. The reason is the ability of Dual Wi-Fi to isolate uplink and downlink performance. Even though WiFi-NC is able to provide a second contention domain by splitting the band, it does not remove uplink to downlink contention for the same medium, and consequently has a dependency between uplink and downlink performance.

Results for asymmetric traffic. Figure 5.4 illustrates the impact of client density on throughput ratio to load ratio under asymmetric traffic where downlink

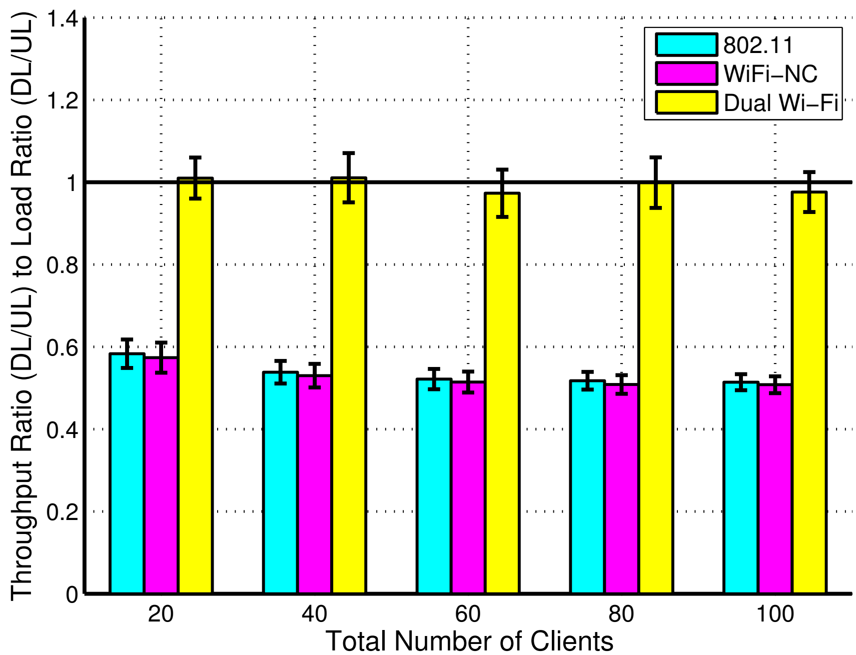

Figure 5.4: Mean throughput ratio (DL/UL) to load ratio (DL/UL) with load asymmetry (DL/UL) of 4 for different network densities.

load is 4x greater than uplink load. We observe in Figure 5.4 that client density has a smaller impact on throughput ratio to load ratio under asymmetric traffic. However, the uplink vs. downlink contention for the medium still causes baseline systems to deviate from ideal through to load proportion. For example in the case of 20 clients, 802.11 and WiFi-NC, obtain a throughput ratio to load ratio of 0.57 corresponding to a spectrum division of 69.5/30.5 (DL/UL), 10.5% away from ideal spectrum division of 80/20 (DL/UL) for a load asymmetry of 4. To the contrary, Dual Wi-Fi approximates ideal throughput ratio to load ratio with a value of 1.01, corresponding to 80.2/19.8 (DL/UL) spectrum division.

Distinct from the case of symmetric traffic, client density has a smaller effect on throughput ratio to load ratio on baseline systems. We observe, in the case of 100 clients, baseline systems achieve a throughput ratio to load ratio of 0.5, corresponding to a spectrum division of 66.67/33.33 (DL/UL), 13.3% away from ideal

Clients:	2	3	4
Dual Wi-Fi	0.968 (± 0.036)	1.101 (± 0.065)	1.054 (± 0.013)
802.11	0.991 (± 0.003)	0.991 (± 0.003)	1.027 (± 0.017)
WiFi-NC	0.992 (± 0.004)	1.015 (± 0.012)	1.039 (± 0.018)

Table 5.4: Small-scale experimental evaluation of throughput ratio (DL/UL) to load ratio (DL/UL) with symmetric load for different number of clients. Table reports mean and standard deviation.

division 80/20 (DL/UL). Meanwhile, Dual Wi-Fi approximates ideal spectrum division 79.67/20.32 (DL/UL) for a load asymmetry of 4 with 100 clients. We observe the divergence in baseline systems between 20 and 100 clients is only of 3% for asymmetric traffic, meanwhile for symmetric traffic is 15%. These results show that medium congestion is the key problem of baseline system to provide proportional throughput ratio to load ratio.

Over-the-Air Evaluation. We performed a small-scale analysis on our experimental platform to confirm contention asymmetry of large-scale scenarios is the reason for unmatched throughput ratio. In the case of small-scale scenarios, uplink and downlink medium contention is approximately symmetric, therefore we expect a near ideal throughput ratio for all systems. We evaluate all three systems under a symmetric load with various topologies formed by 1 AP with 2 to 4 clients.

As expected from the simulation results, Table 5.4 indicates that all three systems achieve ideal throughput to load ratio within 3%. This result supports that contention between uplink and downlink traffic caused by the increasing number of clients is the bottleneck of shared band systems.

5.3 Dual Wi-Fi Bandwidth Division Parameter

Here, we study Dual Wi-Fi's bandwidth division parameter which determines which fraction of the available spectrum is used for downlink or uplink traffic. We observed in previous evaluation how Dual Wi-Fi is able to provide a proportional performance ratio to a given offered load and client density. In order to achieve the desired performance proportion, Dual Wi-Fi requires setting the bandwidth division accordingly. We next examine the impact of the selection of the bandwidth division parameter and consider Dual Wi-Fi's robustness to the case that the bandwidth division parameter is set in a sub-optimal way.

Setup. We evaluate multiple load asymmetries under different network densities and step through all possible downlink to uplink bandwidth allocations. We exhaustively determine the value of the bandwidth division parameter that provides a performance metric closest to 1.

Results. Figure 5.5 depicts the experimentally optimal setting of the bandwidth division parameter as a function of the number of clients for symmetric traffic. Namely, the x-axis depicts the number of clients. The y-axis depicts the optimal ratio of spectrum allocated to downlink and uplink traffic for Dual Wi-Fi with symmetric offered load. The figure indicates that the optimal bandwidth division of downlink to uplink spectrum *decreases* with client density. Thus, increasing uplink bandwidth is required for uplink traffic to counter increasing contention in dense networks. For example, in the case of 5 clients the downlink and uplink bandwidth ratio should be the same as the load ratio, in the case of symmetric traffic i.e. 10/10 MHz bandwidth ratio. However in the case of 40 clients with symmetric traffic, the y-axis depicts 0.667 corresponding to 8/12 MHz. Thus, with high contention, the uplink requires a greater bandwidth allocation of 12 MHz (vs. the downlink's 8 MHz) to counter the negative effect of high uplink contention and provide a throughput to load ratio of 1. Namely, the selection of Dual Wi-Fi bandwidth parameter depends on the network density. This dependency arises from increased medium access contention which degrades uplink performance resulting in a shift of the bandwidth selection towards the

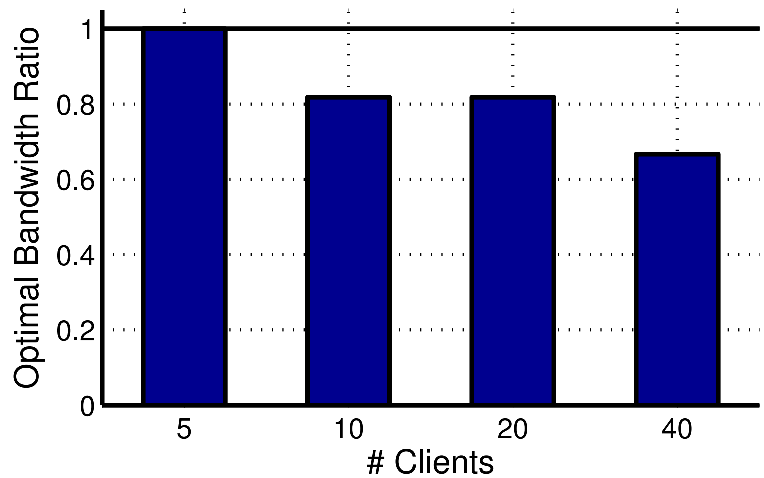

Figure 5.5: Dual Wi-Fi's optimal bandwidth ratio (DL/UL) based on network density for symmetric traffic.

uplink channel.

5.3.1 Robustness to Bandwidth Parameter Mismatch

Next, we evaluate Dual Wi-Fi's robustness to the bandwidth parameter setting and we explore the case in which the bandwidth parameter setting is sub-optimal compared to best setting available for the given traffic demand and network density.

5.3.1.1 Dual Wi-Fi Robustness

The goal of this evaluation is to examine the change on Dual Wi-Fi's throughput ratio (DL/UL) to load ratio (DL/UL) as we deviate from the optimal bandwidth setting.

Setup. We examine the *throughput (DL/UL) to load (DL/UL) ratio* obtained by all possible bandwidth parameter allocations for a given traffic load and client density. We select a network with 1 AP and 20 clients to represent a WLAN. We first select symmetric traffic because it has the most congestion, thus a mismatch from the optimal bandwidth setting will present a higher impact than with asymmetric loads. Second, we select traffic asymmetry of 4, to exemplify the ratio change under asymmetric traffic. In Figures 5.6 and 5.7, the y-axis presents the *throughput to load ratio* obtained by each bandwidth parameter setting in the x-axis. The optimal bandwidth setting (value 0 in the x-axis) approaches the ideal throughput to load ratio, y-axis value of 1. Moreover, any value above 1 on the y-axis indicates that more throughput is given to the downlink traffic than demanded, whereas downlink traffic is underserved when the value is below 1. The x-axis depicts the mismatch from the ideal bandwidth setting as a percentage of the total bandwidth. In both figures, the values on the x-axis to the right of the optimal bandwidth (value 0 in the x-axis) represent parameter settings with more bandwidth given to the downlink channel and the values to the left of the optimal bandwidth give more bandwidth to the uplink channel.

Results. In Figure 5.6, we explore how the throughput to load ratio changes with different bandwidth settings for symmetric traffic. We observe the optimal bandwidth (x-axis value of 0) has a throughput to load ratio of 1.03, corresponding to spectrum division 50.7/49.3, less than 1% away from the ideal 50/50 division. We observe the ideal throughput to load ratio 1.00 is crossed under a 2% mismatch towards the uplink channel, thus Dual Wi-Fi can obtain the ideal ratio of 1 under non-integer bandwidths, however we limit our simulation results to integer bandwidth values for realistic representation.

Figure 5.6 shows as the bandwidth setting deviates from the optimal bandwidth setting, Dual Wi-Fi diverges from ideal throughput to load ratio of 1. For example, if 10% more bandwidth is given to the downlink channel (DL/UL = 11/9 vs. optimal 9/11), the throughput ratio (DL/UL) is *1.591*, where downlink data obtains approximately *1.6x* more throughput that uplink data traffic, with a performance division of 61.41/38.59 (DL/UL). Meanwhile if 10% more bandwidth is given to the uplink channel (DL/UL = 7/13 vs. optimal 9/11) we obtain a throughput ratio of *0.673*, where uplink data has approximately *1.5x* more throughput than downlink data traffic with a 40.23/59.77 (DL/UL) utilization division. We notice in this case downlink benefits more from added bandwidth. However, anywhere above 20% mismatch uplink traffic benefits more than downlink traffic from the added bandwidth. For example, in the case of 35% mismatch, downlink obtains *4.136x* more throughput than uplink traffic, meanwhile uplink with 35% more bandwidth obtains *6.31x* more throughput than downlink traffic, with a throughput ratio of 0.1584.

The driver for this discrepancy is the medium contention within each channel. Because we operate a downlink-only channel, the minimized downlink contention allows downlink traffic to directly benefit from the added bandwidth. The downlink channel reaches the maximum bandwidth required when the bandwidth added is 20% more, any extra bandwidth will minimally affect the throughput of this channel and thus will only reduce uplink throughput. On the other hand, due to moderate uplink contention, the uplink data traffic will benefit from any added bandwidth because the time-length of medium contention parameters are reduced as bandwidth increases.

In Figure 5.7, we explore the change in throughput to load ratio under asymmetric traffic where downlink is 4x greater than uplink traffic. In this case, since the optimal bandwidth follows the traffic asymmetry, downlink can be only provided with up to 15% more bandwidth. In the case where downlink is provided with 10% more bandwidth, downlink throughput is 1.8x higher than uplink throughput. Meanwhile when the uplink data channel is provided with 10% more bandwidth, it achieves 1.5x more throughput than uplink traffic. As more bandwidth is given to the uplink channel the ratio deviates from the ideal ratio of 1, this deviation accelerates past a 10% mismatch (10% = UL **1.5x** > DL, 20% = UL **1.84x** > DL, 30% = UL **2.27x** > DL, 40% = UL **2.9x** > DL, 50% = UL **4x** > DL) due to the severe impact of reducing downlink bandwidth with asymmetric load where downlink load is greater

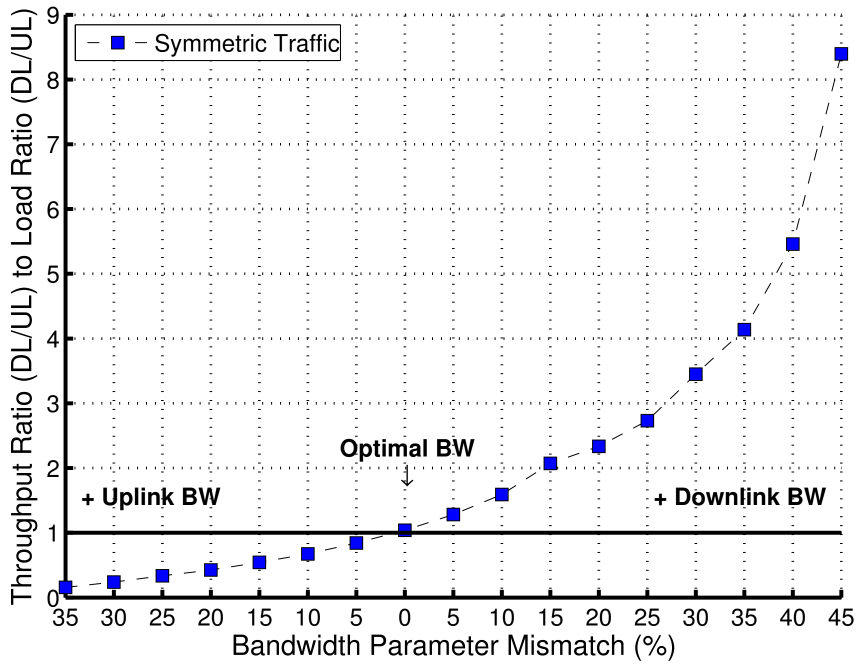

Figure 5.6: Impact of mismatch in optimal bandwidth setting on throughput to load ratio for symmetric traffic.

than uplink load.

5.3.1.2 Robustness Dual Wi-Fi vs. 802.11

Next, we examine the set of bandwidth allocations in which Dual Wi-Fi provides a closer to ideal *throughput to load ratio* than 802.11 under the same network and load conditions.

Setup. In Figure 5.8, the x-axis depicts load asymmetry. The y-axis shows the range of bandwidth settings in which Dual Wi-Fi provides a closer ideal throughput to load ratio than 802.11, as a fraction of the total bandwidth. For example, if the range of estimation error is 100%, this means that under any bandwidth allocation Dual Wi-Fi is closer than 802.11 to a throughput ratio to load ratio of 1. However, if the range of estimation error is 0%, Dual Wi-Fi requires directly selecting the optimal bandwidth setting in order to outperform 802.11. A large robustness for estimation

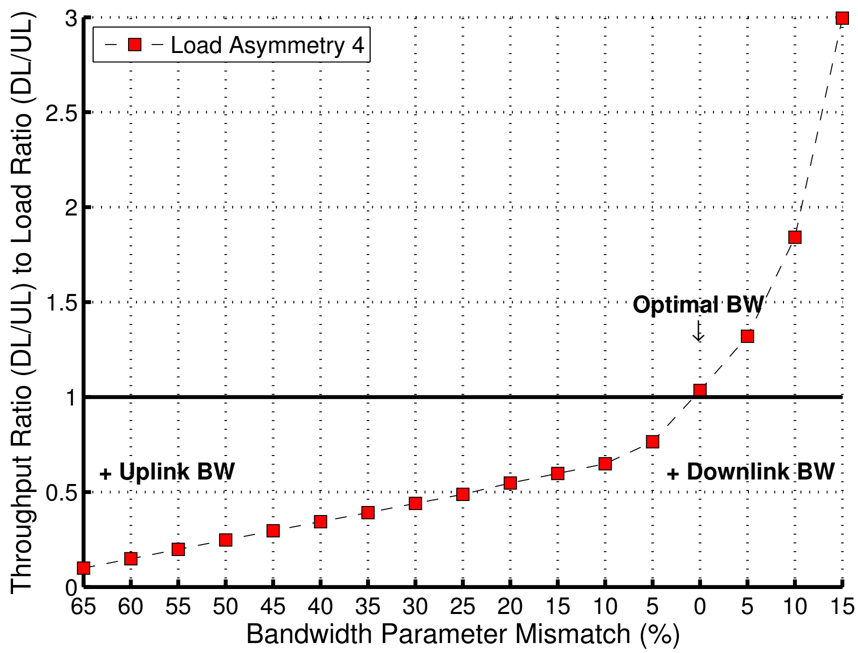

Figure 5.7: Impact of mismatch in optimal bandwidth setting on throughput to load ratio for asymmetric traffic.

error allows an open-loop bandwidth adaptation scheme with long response time to demand-to-allocation mismatch, however if the range of robustness to error is small a closed-loop approach with continuous demand-to-allocation feedback is suited.

Results. In Figure 5.8 we observe that robustness decreases as load asymmetry increases. In the case of low network density, the highest value is 30% with symmetric loads and the lowest is 10% with load asymmetry of 7. The reason for robustness decrease is that under high load asymmetries, 802.11 approaches ideal performance to load ratio, as shown in Figure 5.2. Moreover, in the case of high network density, Dual Wi-Fi achieves higher robustness values, up to 45%, because as the number of clients increase, 802.11 diverges rapidly from the ideal performance to load ratio, as shown in previous evaluations.

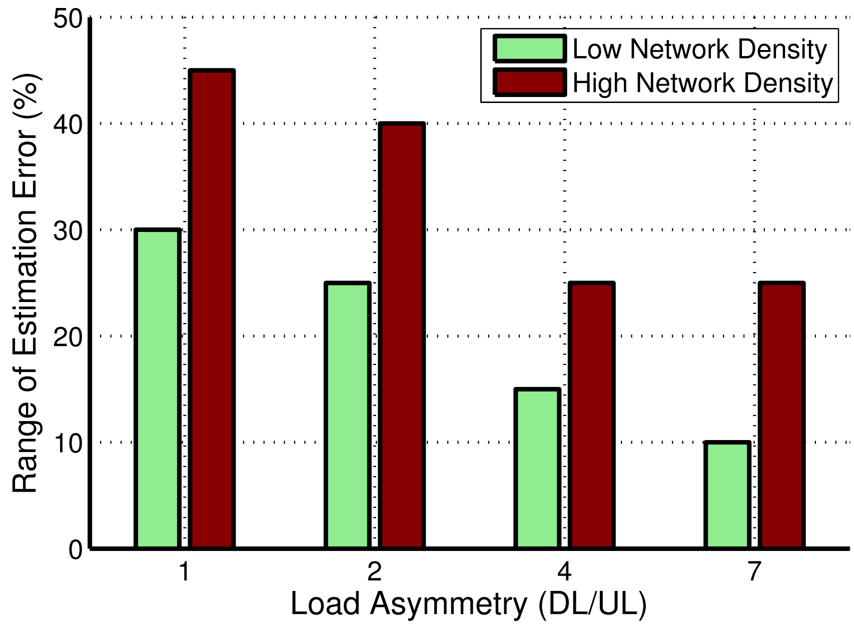

Figure 5.8: Range of bandwidth settings in which Dual Wi-Fi provides a closer ideal throughput ratio (DL/UL) to load ratio (DL/UL) than 802.11, as a fraction of the total bandwidth.

5.4 Gains of Dual Wi-Fi

Dual Wi-Fi's ability to provide proportional performance under any load asymmetry or network density translates into significant throughput gains. In this evaluation, we explore these gains and their sources.

The configuration for the following experiments mirrors that of the "Impact of Contention Asymmetry" evaluation. We explore multi-AP topologies, to introduce downlink contention. We vary the total number of clients between 20 to 100, equally divided across the number of APs. We present the gains for a set number of clients in multiple forms. First, the gains from a 1, 2 or 5 in-channel APs topology, and the gains obtained by averaging across the 1-5 in-channel APs topologies evaluated.

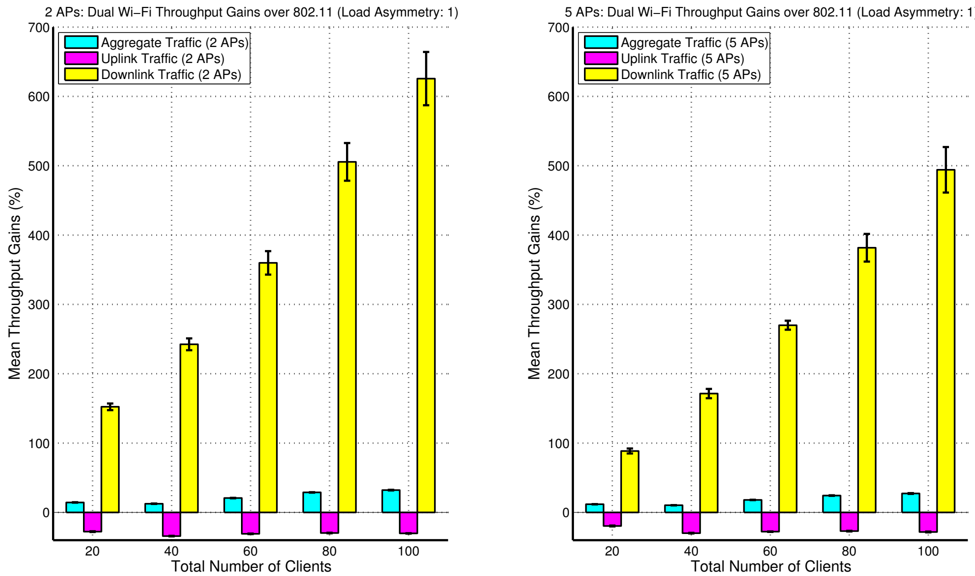

Figure 5.9: Dual Wi-Fi throughput gains over 802.11 with symmetric traffic. Left: 2 in-channel AP topology, right: 5 in-channel AP topology.

5.4.1 Throughput Gains

In this experiment, we explore the throughput of Dual Wi-Fi in comparison to 802.11 and WiFi-NC. We vary the number of clients and measure the aggregate, uplink, and downlink throughput gains of Dual Wi-Fi as compared to baseline systems.

Results for symmetric traffic. Figures 5.10 and 5.9 depict the results for *symmetric traffic* and omit the case of WiFi-NC because these results are within 2 to 3 % of the gains over 802.11. Figure 5.9 depicts Dual Wi-Fi throughput gains over 802.11 for topologies with 20 to 100 clients under 2 in-channel APs (left) and 5 in-channel APs (right). As expected we observe the highest gains in downlink traffic, due to performance isolation between downlink and uplink traffic of Dual Wi-Fi. The highest downlink gains are obtained in the case of 2 APs (left plot), in this case downlink traffic with 100 clients achieves average throughput gains of 626%. Meanwhile downlink traffic under 100 clients with 5 in-channel APs (right plot) obtains average gains of 494%. We observe, the larger downlink traffic gains are obtained in the case of higher client density with smaller number of in-channel APs. This is because 802.11 in the case of smaller number of in-channel APs has a higher disproportion in contention between uplink and downlink traffic making uplink have a higher probability to win the medium, thus Dual Wi-Fi performance isolation allows our system to achieve higher gains. Consequently, as the number of clients increase in the network, Dual Wi-Fi achieves unbounded gains over 802.11 systems. Furthermore, an additional reason for Dual Wi-Fi to achieve smaller gains in the case of 5 in-channel APs is that as more APs are added to the downlink channel, these cause a small increase in contention for the downlink channel, thus slightly degrading Dual Wi-Fi's downlink performance.

Moreover, we observe the aggregate gains between the 2 and 5 AP networks are only <4% apart from each other. The 2 AP topology provides aggregate gains from 14 to 32 % meanwhile the 5 AP network achieves aggregate gains from 12 to 27%. The main cause for the aggregate gains is the removal of contention between uplink and downlink traffic.

On the other hand, Dual Wi-Fi decreases uplink throughput in both networks by 19 to 34%. However, this decrease is targeted by Dual Wi-Fi's re-allocation of resources since uplink traffic in 802.11 systems utilizes more spectrum than required by ideal performance to load proportion. Dual Wi-Fi decreases 2% more the uplink throughput in a network of 2 APs than in a 5 AP topology. Again, the reason for this difference is the higher contention asymmetry between uplink and downlink traffic in 802.11, under smaller number of in-channel APs.

Figure 5.10 depicts the average throughput gains for the evaluated 1-5 APs topologies. As shown before, Dual Wi-Fi achieves the highest gain in downlink throughput, up to 530% on average for 100 clients. Such gains are obtained in scenarios with high contention caused by high client density. Our previous evaluation demonstrated

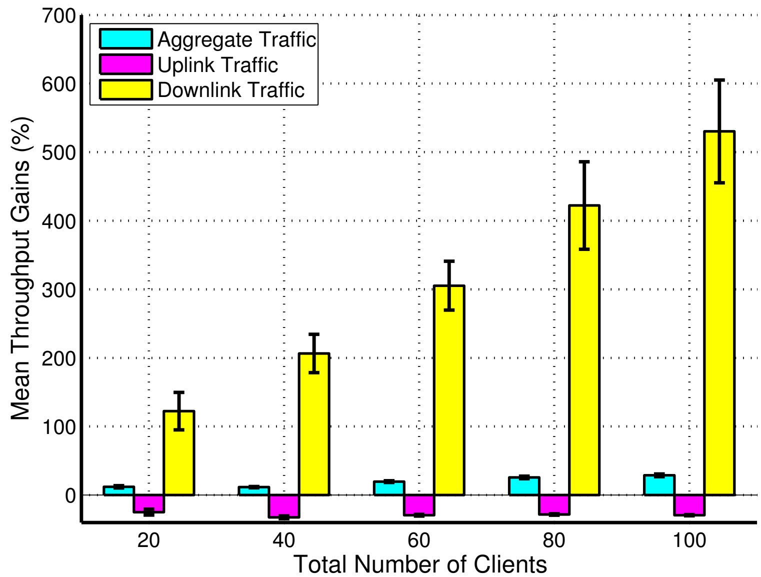

Figure 5.10: Dual Wi-Fi throughput gains over 802.11 with symmetric traffic for topologies with 1-5 in-channel APs.

that contention asymmetry severely affects downlink performance in 802.11 but not Dual Wi-Fi. Our system maintains downlink performance across all client densities, whereas 802.11 and WiFi-NC downlink performance rapidly decays yielding increasing relative performance. Downlink gains yield a large standard deviation because in cases where the number of APs is large, downlink performance in 802.11 systems achieves a better medium access probability than in the case of a single AP, as shown in Figure 5.9. Furthermore, as observed in Figure 5.9 gains over uplink and aggregate traffic are similar between topologies, thus in Figure 5.10 a small deviation is observed for uplink and aggregate gains across the 1-5 in-channel APs topologies evaluated.

Results for asymmetric traffic. Figure 5.11 depicts the average throughput gains over an 802.11 system with traffic asymmetry of 4 (DL 4x > UL). The results are evaluated with topologies composed by 1 to 5 APs with 20 to 100 clients equally spread across the number of APs.

As seen in the case of symmetric traffic, downlink traffic also obtains the highest

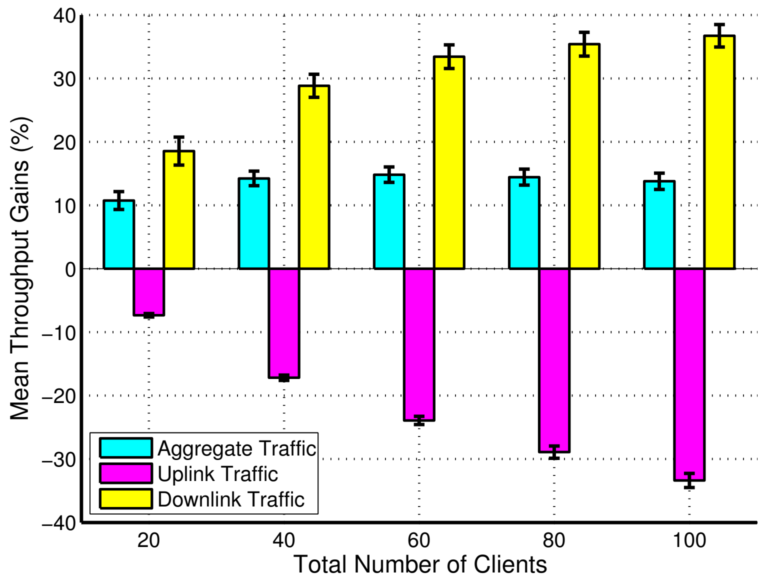

Figure 5.11: Dual Wi-Fi throughput gains over 802.11 with load asymmetry (DL/UL) of 4 for different network densities.

gains with asymmetric traffic. In Figure 5.11 we observe downlink traffic obtains gains from 19 to 37 %. Because 802.11 systems provide more throughput to uplink traffic than downlink traffic, Dual Wi-Fi intends to reduce uplink performance to provide downlink traffic with a better performance that is proportional to its demand. Thus in Figure 5.11 we observe the volume of the downlink gains closely follows the intended reduction of uplink traffic especially under high client density. We reduce uplink traffic from 7 to 33 % while downlink traffic presents throughput gains from 19 to 37 %. However, even though we reduce uplink traffic, overall we achieve aggregate traffic gains ranging from 10 to 15 %. The gains from aggregate traffic reflect the effect uplink vs. downlink contention has in system's performance.

5.4.2 Reduction of Contention Time

Backoff time before every data transmission leads to spectrum underutilization because the medium must remain idle. Here we evaluate the reduction Dual Wi-Fi

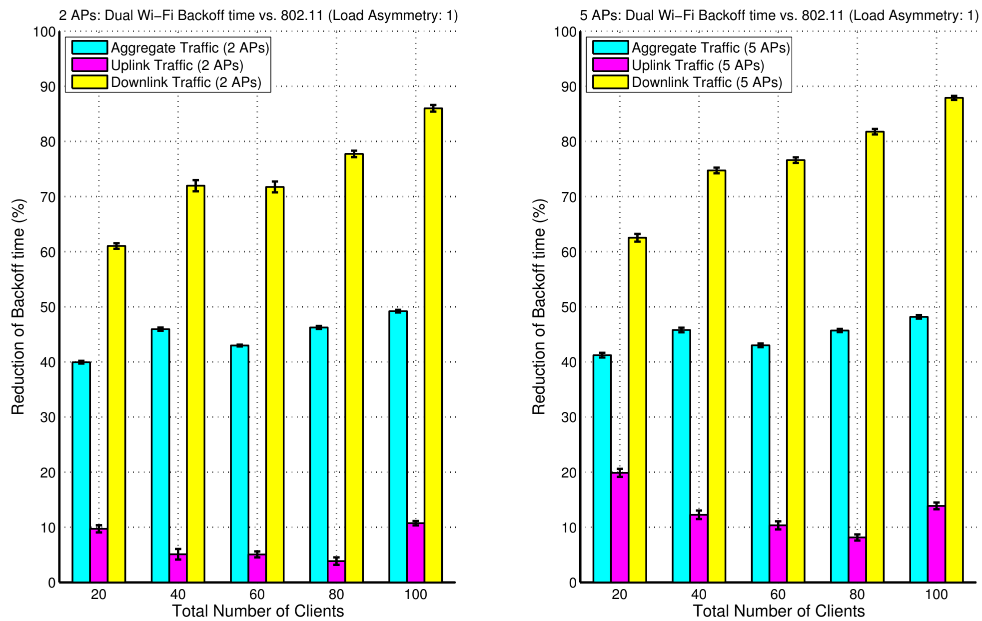

Figure 5.12: Mean reduction of backoff time per successful data transmission over 802.11 with symmetric traffic. Left: 2 in-channel AP topology, right: 5 in-channel AP topology.

presents over 802.11 systems for the average backoff time per successful transmissions.

Results for symmetric traffic. Figures 5.12 and 5.13 depict the mean reduction of backoff time per successful transmission over 802.11 as a function of client density. Results for WiFi-NC are again not shown since reductions are within 3% of 802.11.

Figure 5.12 depicts Dual Wi-Fi mean reduction of backoff time per successful transmission over 802.11 for topologies with 20 to 100 clients under 2 in-channel APs (left) and 5 in-channel APs (right). As expected, we observe the highest reduction of backoff time is achieved in downlink traffic, even in topologies with 20 clients, Dual Wi-Fi reduces downlink backoff by approximately 60% over 802.11. This means that on average every downlink successful transmission requires 60% less contention time. The reduction of downlink backoff is increased as client density increases. For example in the case of 100 clients with 2 in-channel APs, downlink backoff is reduced

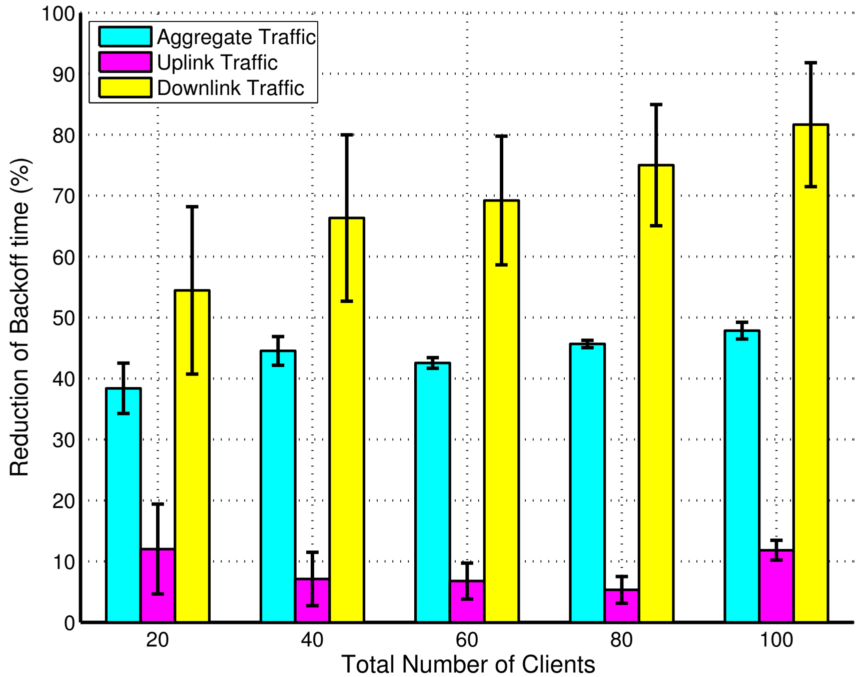

Figure 5.13: Mean reduction of backoff time per successful data transmission over 802.11 with symmetric traffic for topologies with 1-5 in-channel APs.

by 86% and by 88% in the case of 5-channel APs. We observe downlink backoff is reduced on average by 2% more in the case of 5 APs in comparison to 2 in-channel APs.

Moreover, in the case of uplink traffic, Dual Wi-Fi reduces uplink backoff between 4 to 10% in the case of 2 in-channel APs and 8 to 20% in the case of 5 in-channel APs. The cause of the reduction in uplink medium contention is the removal of contention with the downlink traffic. Thus, we achieve higher uplink reduction in the case where more APs operate in the same channel. Overall, Dual Wi-Fi reduces on average contention by approximately 40 to 50 %, for both topologies formed by 2 and 5 in-channel APs. This result demonstrates, Dual Wi-Fi has up to 50% less medium idle time than 802.11, meaning Dual Wi-Fi's medium can be utilized up to 50% more efficiently.

Figure 5.13 depicts the mean reduction of backoff time per successful transmission over 802.11 as a function of client density for the evaluated 1-5 in-channel APs topologies. The figure indicates that Dual Wi-Fi provides significant reduction in backoff for downlink transmissions across all client densities. Because Dual Wi-Fi provides a downlink-only channel, downlink backoff periods are low under any client density and therefore gains increase as the number of clients increase, achieving up to 90% reduction in the case of 1 AP and 100 clients.

Furthermore, Dual Wi-Fi achieves an average of 9% backoff reductions on uplink transmissions across all evaluated topologies, even when utilizing narrower bandwidth than 802.11. The reduction of backoff time in the uplink channel is simply caused by the removal of downlink traffic. Overall, Dual Wi-Fi achieves to reduce backoff time of all transmissions on average by 39% to 48%, meaning Dual Wi-Fi's medium can be utilized on average by 45% more for data transmissions across all the evaluated topologies.

Results for asymmetric traffic. Figure 5.14 depicts Dual Wi-Fi mean reduction of backoff time per successful transmission over an 802.11 system with traffic asymmetry of 4 (DL 4x > UL). The results are evaluated with topologies with 20 to 100 clients under 2 in-channel APs (left) and 5 in-channel APs (right).

Under asymmetric traffic we observe the highest backoff reduction is obtained in aggregate traffic, where backoff is reduced from 28 to 47 % in the case of 2 in-channel APs and from 37 to 50 % in a 5 in-channel AP topology. The backoff reduction of aggregate traffic depends on the reduction of uplink and downlink backoff. Dual Wi-Fi reduces backoff time for downlink traffic by 20 to 27 % in a 2 in-channel AP topology and 28 to 36 % in a 5 in-channel AP topology. Moreover, we observe backoff of uplink traffic is reduced by 20 to 13 % for 2 AP topologies and by 35 to 22 % in the 5 AP case. The reduction of uplink backoff decreases as the number of client

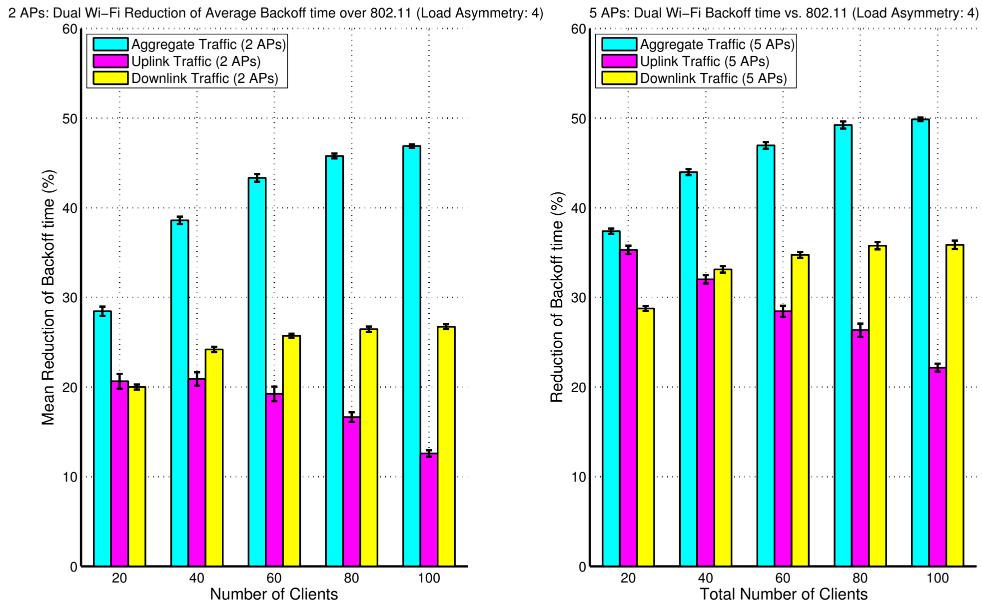

Figure 5.14: Mean reduction of backoff time per successful data transmission over 802.11 with load asymmetry (DL/UL) of 4. Left: 2 in-channel AP topology, right: 5 in-channel AP topology.

increases, mainly because as in 802.11 systems the backoff in Dual Wi-Fi uplink channel is dependent on the number of clients. Reduction of downlink backoff in the case of asymmetric traffic has a smaller magnitude than the case of symmetric traffic, the reason for this discrepancy is the difference in magnitude between downlink and uplink traffic. Because uplink offered load is significantly lower than downlink offered load, downlink traffic has more opportunity to access the channel in 802.11 systems. However, we observe as the number of clients increment the reduction of downlink backoff increases thus portraying the effect of contention asymmetry. Overall, we see even with a large traffic asymmetry, our system is able to significantly reduce overall backoff, thus demonstrating uplink vs. downlink contention affects spectrum efficiency.

5.4.3 Reduction of Retransmissions

Retransmissions affect performance by unnecessarily using spectrum resources through backoff periods and transmission of unsuccessful data, including additional inter-frame times and timeouts. Next, we analyze the reduction of the average number of retransmissions over 802.11 systems.

Results for symmetric traffic. Figures 5.15 and 5.16 depict the mean percentage of retransmission reduction per transmitted packet for Dual Wi-Fi compared to 802.11 with symmetric traffic.

On the left plot of Figure 5.15 we observe the mean reduction of retransmissions for topologies with 1 in-channel AP with 20 to 100 clients. First, as expected, downlink retransmissions are totally eliminated under any client density. This is because Dual Wi-Fi operates a downlink-only channel, thus in the case of having 1 in-channel AP, it operates *collision-free*. Moreover, we observe uplink retransmissions are reduced by 1 to 4%, this percentage represents the amount of uplink collisions caused by downlink traffic, in a shared band system. Overall, Dual Wi-Fi in topologies of 1 in-channel AP reduces retransmissions by 23.4% in the case of 100 clients and 31.5% in the case of 20 clients, thus on average providing a 27% more efficient use of the medium.

The right plot of Figure 5.15 depicts the mean reduction of retransmissions per transmitted packet for Dual Wi-Fi in comparison to 802.11 in topologies with 5 in-channel APs with 20 to 100 clients. First, we observe downlink retransmissions are reduced between 49 to 66 %, the percentage of downlink retransmissions reduction is decreased from the case of 1 in-channel APs, because there are multiple APs contending for the downlink channel. However, even in the case of having multiple APs share the downlink channel, the amount of collisions in the downlink transmission is reduced by more than half than in the case of a shared band system. Moreover, uplink retransmissions in topologies with 5 in-channel APs are reduced between 3 to

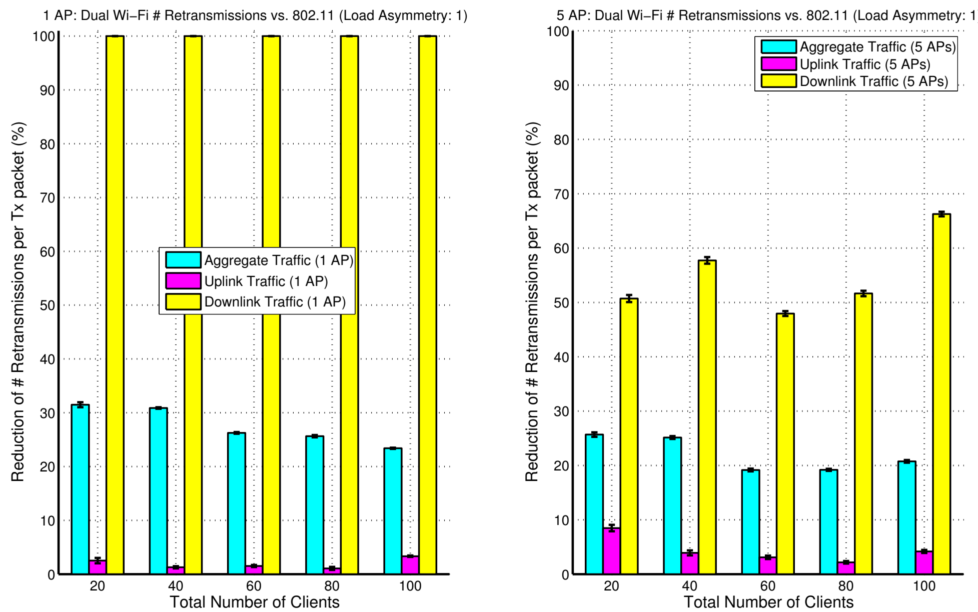

Figure 5.15: Mean reduction of number of retransmissions per transmitted packet over 802.11 with symmetric traffic. Left: 1 in-channel AP topology, right: 5 in-channel AP topology.

9%. We observe this reduction is higher than the case of 1 in-channel AP (1 - 4 %), the reason for the increase in percentage reduction of uplink retransmissions in a 5 in-channel topology is due to the increment of uplink collisions caused by the increment of APs. Thus, because in shared band systems the uplink retransmissions increase as more APs operate in the channel, Dual Wi-Fi's reduction of uplink retransmission increases with more in-channel APs. Overall, Dual Wi-Fi reduces retransmissions on average by 22% in the case of 5 in-channel APs with 20 to 100 clients, thus in this case the medium is utilized on average 22% more efficiently.

Figure 5.16 depicts the mean percentage of retransmission reduction per transmitted packet for Dual Wi-Fi compared to 802.11 for the evaluated 1-5 in-channel APs topologies. As expected, the reduction of downlink retransmissions is the most significant in the case of a small number of APs, high end of standard deviation bars, where our system can eliminate (100% reduction) downlink retransmissions.

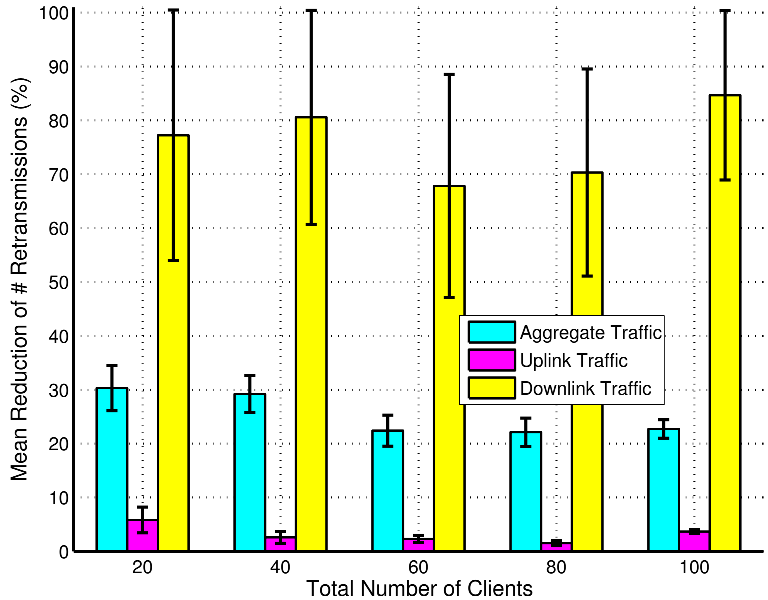

Figure 5.16: Average reduction of number of retransmissions per transmitted packet over 802.11 with symmetric traffic for topologies with 1-5 in-channel APs.

On average, Dual Wi-Fi reduces downlink retransmissions by 75%. This reduction is minimally affected by client density. Uplink retransmissions are reduced on average by a percentage of approximately 3-6%. This reduction is based on the removal of downlink traffic. The percentage reduction of uplink retransmission decreases with client density, as expected, because Dual Wi-Fi's uplink channel implements 802.11, with the only difference of removing downlink traffic. Therefore the gains observed by uplink traffic are the percentage of retransmission caused by downlink transmissions in shared band systems. Overall the system reduces aggregate retransmissions by approximately 30 to 22%. This reduction is the result of removing uplink vs. downlink contention, and allows Dual Wi-Fi to utilize the medium more efficiently.

Results for asymmetric traffic. Figure 5.17 depicts Dual Wi-Fi mean percentage of retransmission reduction per transmitted packet over an 802.11 system with traffic asymmetry of 4 (DL 4x > UL). The results are evaluated with topologies with 20 to 100 clients under 1 to 5 in-channel APs.

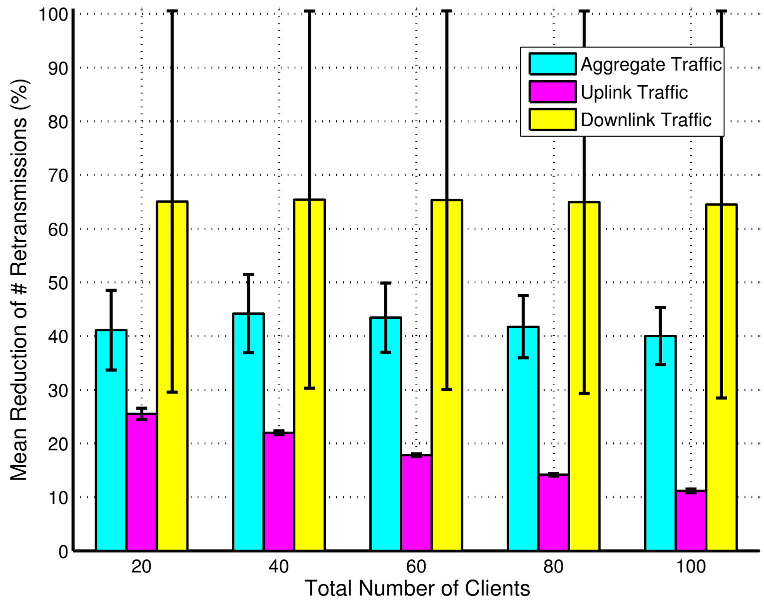

Figure 5.17: Average reduction of number of retransmissions per transmitted packet over 802.11 with load asymmetry (DL/UL) of 4 for topologies with 1-5 in-channel APs.

First, we observe in Figure 5.17 the reduction of downlink retransmissions has a large deviation, the reason for this large variance is the increase in the number of APs evaluated. Retransmissions are eliminated (100% reduction) in the case of small number of in-channel of APs (1 to 2), but as the number of in-channel APs increase the number of retransmissions rises. Across all the evaluated scenarios we maintain an average reduction of downlink retransmissions of 65 %, thus significantly reducing inefficient use of the spectrum. Moreover, uplink retransmissions are reduced by 11 to 26%, this reduction is solely based on isolating uplink medium contention from downlink contention. We observe the reduction of uplink retransmissions is more significant in the case of asymmetric traffic than symmetric traffic, thus demonstrating the impact of downlink transmissions on uplink retransmissions under asymmetric traffic. Overall, we observe Dual Wi-Fi achieves an average reduction of retransmissions of 40%. The reduction of aggregate traffic retransmission under asymmetric traffic (40 %) is larger than the case of symmetric traffic (22%), providing an even more efficient spectrum utilization than in the case of asymmetric traffic.

Chapter 6

Related Work

6.1 Traffic Asymmetry

We observed two common ways in which existing work address heterogeneous traffic loads. First, algorithms and protocols have been propose to dynamically distribute spectrum across APs based on their traffic load [11, 12, 13]. This concept is similar to cell breathing [14], however instead of modifying power to control client association to avoid light and high loaded APs, AP load balancing adapts bandwidth to match traffic. In contrast to these schemes, Dual Wi-Fi combats traffic asymmetry from within the channel, without requiring adjusting the total network bandwidth. Furthermore, these schemes present a limitation on how much bandwidth they can provide to congested WLANs. Instead Dual Wi-Fi combats the issue from the origin by separating uplink and downlink contention to utilize more efficiently the assigned bandwidth without requiring additional resources.

The second approach adapts medium access parameters of a disadvantaged traffic flow to allow a more aggressive medium access. Examples of this approach are 802.11e [15] which differentiates traffic through traffic categories and DCA [16] which targets traffic asymmetry. In these approaches the band is shared between uplink and downlink traffic, and thus even high-priority traffic will be affected by low-priority traffic under congested scenarios, because these are not isolated as long as they coexist in the same channel. Since Dual Wi-Fi isolates uplink and downlink data traffic, traffic asymmetry can be addressed under any environment. Moreover, Dual Wi-Fi creates an easy way to provide more bandwidth where required, however in a shared band system this is more complex because the bandwidth assigned depends on the medium access probability parameters.

6.2 Multi-channel Access

The closest work to Dual Wi-Fi multi-channel access is FICA [4] and WiFi-NC [1]. Both of these systems divide the channel into narrow sub-channels to allow fine-grain access. Even though, these systems do not target traffic asymmetry, they aim to increased spectral efficiency. Because these systems do not differentiate between uplink and downlink traffic, increase in spectral efficiency is limited by uplink and downlink contention within the narrow channels. The core distinction between these approaches and Dual Wi-Fi is that our system pre-allocates the sub-channels to downlink and uplink traffic and thus presents the capability to adapt spectrum resources to the traffic load, instead of presenting limited fixed-width sub-channels.

Chapter 7

Conclusion

In this thesis, we introduce Dual Wi-Fi, a novel architecture and MAC protocol that provides spectrum independence between uplink and downlink MAC data traffic, to support traffic asymmetry under any network density. Our channel architecture consists of two physical bi-directional channels: "downlink data channel" and "uplink data channel", that each carry a single uni-directional MAC traffic direction, i.e. uplink or downlink Data-ACK handshake. The spectrum isolation and bi-directional capability permit Dual Wi-Fi's channel design to configure the downlink and uplink bandwidth resources to match a given traffic asymmetry. Furthermore, our channel design allows independent and asynchronous uplink and downlink MAC level data-ACK transmissions, thus providing independent performance between uplink and downlink MAC level traffic. Through a wide set of evaluation on commodity hardware and simulations, we show that Dual Wi-Fi matches downlink vs. uplink throughput ratio within 1% of the downlink vs. uplink demand ratio, under any traffic asymmetry and network density. Through this matching capability, Dual Wi-Fi achieves significant throughput gains, obtained by large reduction of retransmissions and contention time.

References

[1] K. Chintalapudi, B. Radunovic, V. Balan, M. Buettener, S. Yerramalli, V. Navda, and R. Ramjee, "WiFi-NC: WiFi over narrow channels," in *Proc. of NSDI*, 2012. 1, 2.2, 3.1.1, 3.2.1, 5.1, 6.2

[2] "IEEE Std 802.11-2012 Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," March 2012. 1, 5.1, 5.1.2.1, 5.1.2.2

[3] A. Flores, R. Guerra, E. Knightly, P. Ecclesine, and S. Pandey, "IEEE 802.11af: A Standard for TV White Space Spectrum Sharing," *Communications Magazine, IEEE*, October 2013. 2.2

[4] K. Tan, J. Fang, Y. Zhang, S. Chen, L. Shi, J. Zhang, and Y. Zhang, "Fine-grained channel access in wireless LAN," in *Proc. ACM SIGCOMM*, ser. SIGCOMM '10. ACM, 2010, pp. 147–158. 2.2, 3.2.1, 6.2

[5] L. Yang, B. Y. Zhao, and H. Zheng, "The spaces between us: setting and maintaining boundaries in wireless spectrum access," in *Proc. ACM MobiCom*, ser. MobiCom '10. ACM, 2010, pp. 37–48. 2.2, 3.1.1

[6] W. Hou, L. Yang, L. Zhang, X. Shan, and H. Zheng, "Understanding cross-band interference in unsynchronized spectrum access," in *Proc. ACM workshop on Cognitive radio networks*, ser. CoRoNet '09. ACM, 2009, pp. 19–24. 2.2, 3.1.1

[7] Executive Office of the President's Council of Advisors on Science and Technology, "Realizing the full potential of government-held spectrum to support economic growth," Tech. Rep., July 2012. 2.2

[8] H. Cho and K. Lee, "Wireless terminals supporting communications with an access point over separate wireless channels and related wireless systems," US Patent US20040213180, October 2004. 3.2.1

[9] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang, "MACAW: a media access protocol for wireless LAN's," in *Proc. ACM SIGCOMM*, ser. SIGCOMM '94. ACM, 1994, pp. 212–225. 4.3

[10] E. Magistretti, O. Gurewitz, and E. W. Knightly, “802.11ec: collision avoidance without control messages,” in *Proc. ACM MobiCom*, ser. Mobicom ’12. ACM, 2012, pp. 65–76. 4.3

[11] T. Moscibroda, R. Chandra, Y. Wu, S. Sengupta, P. Bahl, and Y. Yuan, “Load-aware spectrum distribution in Wireless LANs,” in *IEEE International Conference on Network Protocols*, ser. ICNP ’08, 2008, pp. 137–146. 6.1

[12] L. Yang, L. Cao, H. Zheng, and E. Belding, “Traffic-aware dynamic spectrum access,” in *Proc. International Wireless Internet Conference (WICON)*, ser. WICON ’08. ICST, 2008, pp. 1–9. 6.1

[13] E. Rozner, Y. Mehta, A. Akella, and L. Qiu, “Traffic-Aware Channel Assignment in Enterprise Wireless LANs,” in *IEEE International Conference on Network Protocols*, ser. ICNP ’07. 6.1

[14] P. Bahl, M. Hajiaghayi, K. Jain, S. Mirrokni, L. Qiu, and A. Saberi, “Cell Breathing in Wireless LANs: Algorithms and Evaluation,” *IEEE Transactions on Mobile Computing*, vol. 6, no. 2, pp. 164–178, 2007. 6.1

[15] “IEEE standard 802.11e, amendment 8: Medium access control (MAC) Quality of Service Enhancements,” November 2005. 6.1

[16] S. Kim, B. Kim, and Y. Fang, “Downlink and uplink resource allocation in IEEE 802.11 wireless LANs,” *IEEE Transactions on Vehicular Technology*, vol. 54, no. 1, pp. 320–327, 2005. 6.1