Backoff Reset

Finding the limits of IEEE 802.11 modelling

MSc. Embedded Systems, Software & Networking

Thom van der Steenhoven

EEMCS, Delft University of Technology

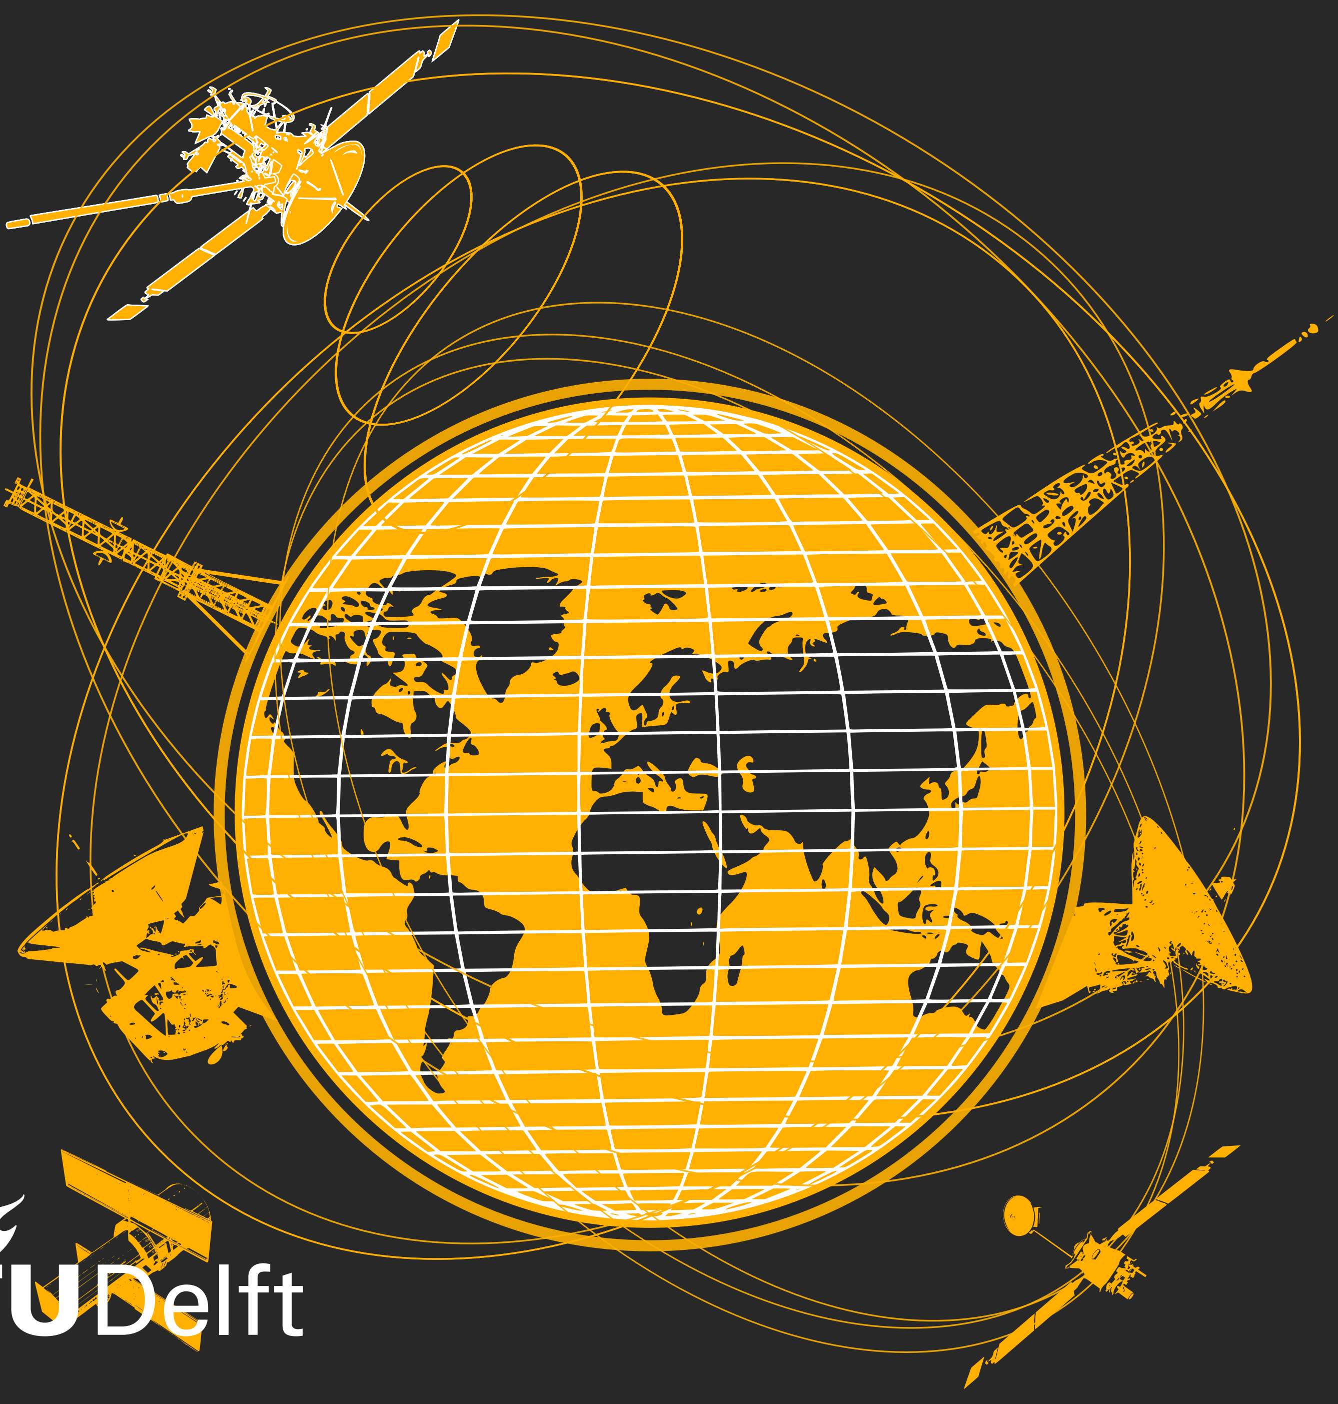

Backoff Reset

Finding the limits of IEEE 802.11 modelling

by

Thom van der Steenhoven

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Monday June 5th, 2023 at 16:00.

Student number:	4392566	
Project duration:	November 1st, 2021 – June 5th, 2023	
Thesis committee:	Prof. Dr. Ir. Koen Langendoen,	TU Delft, Daily supervisor
	Prof. Dr. Ir. RangaRao Venkatesha Prasad,	TU Delft, Academic supervisor
	Prof. Dr. Ir. Soham Chakraborty	TU Delft

Cover: Satellites Space Orbit by OpenClipart-Vectors from Pixabay (Modified)
Style: TU Delft Report Style, with modifications by Daan Zwaneveld

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

Preface

The thesis that lies before you was created as part of the Master of Science programme Embedded Systems at the TU Delft in the Netherlands.

The author would like to thank many people, first and foremost his loving girlfriend, Ruth Alkema, without whose support this thesis would never have been completed. Secondly, the author owes a great debt of gratitude to Prof. Koen Langendoen, who took up the torch when others let theirs fall. Thirdly, a heart-felt thank you goes out to all the friends and family who have listened, helped, and provided support and advice during the many months that this thesis has been in the works. You are all truly wonderful, thank you.

Thom van der Steenhoven
Rotterdam, May 2023

Summary

This thesis has studied an oversight in the scientific literature on Wi-Fi 6 modelling and attempts to quantify the prevalence and effect of this omission. This effect was dubbed the ”backoff reset” and roughly means that a device connected to a Wi-Fi 6 network that recently received the green-light to send from the access point and, as such, has successfully sent a message, should reset the remaining time it will wait to send its next message. The problem statement was then based on this effect and read as follows:

”What is the effect on performance metrics of resetting a station’s backoff stage and counter after an UL-OFDMA transmission?”

To allow readers new to the Wi-Fi space to be able to follow along, a brief introduction was given to some of the core concepts of Wi-Fi 6 and a preliminary literature survey was performed to find out how often the backoff reset is omitted. From this census, it was found that only five out of fifteen papers likely included the backoff reset in their results.

To analyze the effects of the backoff reset on network performance, a series of experiments were performed in simulator. In these experiments, the system with a backoff reset, named the ”reset” system, was pitched against a system without the reset, called the ”continue” system, while varying various network parameters.

From these results, the following conclusion is drawn: ”The ”reset” system outperforms the ”continue” system in certain situations on (a combination of) latency and collision probability, but is generally strictly worse on all five key metrics in a typical, flexible Wi-Fi deployment, where data rates and the amount of connected devices can change rapidly. We thus accept the research hypothesis.” Finally, an overview of the situations in which the ”reset” system can outperform the ”continue” system was given and the limitations of this study were discussed, after which future work was proposed.

Acronyms

Contents

1 Introduction

Look at your smartphone, then look at the device you are reading this on. Chances are, both of these computers are connected to a Wi-Fi network right now. Wi-Fi has become one of the most dominant and ubiquitous connective technologies in the western world today, with just the number of public Wi-Fi hotspots already forecast to number over 600 million globally in 2023 [1]. Not only are more people using Wi-Fi each year, but the speed of these wireless links is improving year over year as well [1].

This improvement in data rates is driven partly by wireless researchers coming up with more efficient and faster ways of transmitting data over the air to meet the ever-growing demand for bandwidth of all these Wi-Fi devices. Such work is thus essential in our always-connected world, where going off-line for even a day can feel like being excommunicated by one's friends and family. However, wireless researchers are often working on the cutting edge of technology, meaning that hardware to test their theories on can be months or even years away. There are efforts to research and create open-source Wi-Fi hardware, such as the openwifi project [2], however these efforts are, by necessity, lagging behind the development of the newest Wi-Fi standards and are thus not always useful for the development of new Wi-Fi protocols. Closed-source hardware solutions, meanwhile, offer their own challenges, as the author of this report can attest to. Hardware vendors tend to keep the workings of their hardware and firmware a secret for as long as possible to maintain a competitive edge, while providing the open-source community with the bare essentials to get their hardware working with, for instance, the Linux ecosystem.

As such, a lot of Wi-Fi research is done by extending existing wireless networking simulators like ns-3 [3] and OMNeT++ [4] with new features or writing proprietary ones from scratch. These simulators are then validated against theoretical frameworks to make sure that the data they produce is at least somewhat accurate. While this obviates the need for advanced hardware to perform research, this does leave some uncertainty as to whether the simulation results are actually representative of real-world performance and oversights can easily be introduced without anyone noticing. This report aims to address one of those oversights and to investigate the effects it has on some key performance metrics. However, before addressing this oversight, some explanation of the working of Wi-Fi is in order.

1.1. Essential Wi-Fi concepts

This subsection explains some core concepts for the working of the newest Wi-Fi standard, known as Wi-Fi 6 or IEEE 802.11ax, at a high level. A full explanation of the Wi-Fi protocol is outside the scope of this report, but an interested reader might refer to (in order of ascending complexity): The free IEEE 802.11 crash course on Udemy [5]; the excellent paper by Khorov et al. on the working of IEEE 802.11ax [6]; and the most recent IEEE 802.11 standard [7] and IEEE 802.11ax amendment [8].

Instead, this report will limit itself to the bare essentials required for a thorough understanding of the research conducted. Those essentials are, in order: the legacy distributed coordination function (DCF) medium access control (MAC) protocol; the current enhanced distributed channel access (EDCA) MAC protocol; and finally orthogonal frequency-division multiple access (OFDMA) statistical multiplexing.

1.1.1. Distributed coordination function (DCF)

The wireless signals that make up a Wi-Fi transmission share a transmission medium: the frequency band, or channel, in which they are transmitted. While this allows wireless devices, also knows as stations (STAs), within the same channel to communicate with each other, it also means that simultaneous data transmissions on the same channel will interfere with each other. Such interference will generally result in at least a partial loss of the transmitted data, which is called a collision.

Collisions are not unique to wireless media, as two systems on a wired network will cause a similar kind of interference if they transmit simultaneously. However, the problem is more pronounced in wireless systems, as they can not sense the wireless medium while transmitting to detect whether a collision has occurred, as their own transmission effectively deafens them to those of others. As such, a protocol is required that manages the when and how of accessing the channel. For IEEE 802.11 networks before 2007, this medium access control (MAC) protocol was called the distributed coordination function (DCF). The relevance of this legacy protocol is that the current MAC protocol, EDCA, is heavily based on DCF. To perform its function, DCF made use of carrier-sense multiple access with collision avoidance (CSMA/CA) in combination with a binary exponential backoff algorithm.

In CSMA/CA, whenever a STA has data it wants to transmit across the wireless network, it enters a state called contention. In this state, it first senses the channel for a predefined time period of 34 μs, called the DCF inter-frame spacing (DIFS). If it finds the channel idle for the DIFS period, it is then allowed to immediately proceed with sending its data. However, if the STA finds the channel to be occupied during the DIFS duration, it defers transmitting its payload for a random duration of time. The duration of this backoff period is determined by the previously mentioned binary exponential backoff algorithm, which is the next step in understanding DCF.

This algorithm works as follows: Each STA maintains a variable called the contention window (CW)[1], of the form

$$CW(s) = min(2^s * (CW_{min} + 1) - 1, CW_{max}), \tag{1.1}$$

where s is called the backoff stage, CW_{min} is a predefined minimum value and CW_{max} is a predefined maximum value, both of the form $2^x - 1$, with x some positive integer. Whenever a STA finds the channel busy during a DIFS period, a random integer b, for backoff counter, in the range of $0 \leq b \leq CW$ is chosen. The STA then waits for b backoff slots, each with a duration of 9 μs. After each backoff slot, the value of b is decremented, unless another STA transmits during one of those backoff slots. Then, the backoff counter is not decremented until a full idle backoff slot passes. Once the backoff counter hits zero, the STA attempts to transmit again. As a wireless device can not detect a collision on the wireless medium, it expects an acknowledgement (ACK) message from its recipient[2]. If this ACK is not received, it assumes a collision has occurred, increments s, and chooses a new random backoff counter value b, but with the new value of CW. Whenever a message is successfully transmitted and the corresponding ACK message is received, the stations resets its backoff stage, s, to 0, ending the contention state.

[1] The term contention window (CW) in literature is used to denote both the variable as defined in Equation (1.1), and the range $[0, CW(s)]$. While this double usage might cause confusion, the conscious decision was made to keep in line with convention in this thesis.

[2] The ACK message does not have to contend for the channel, as instead of waiting for a DIFS duration of channel idle time, it is allowed to wait for the much shorter short inter-frame spacing (SIFS) of 16 μs, thus attaining the channel before any other STA would be able to.

Access category	CW_{min}	CW_{max}	AIFS (μs)	TXOP_{max} (ms)
Background (BK)	aCW_{min}	aCW_{max}	79	2.528
Best effort (BE)	aCW_{min}	aCW_{max}	43	2.528
Video (VI)	$(aCW_{min}+1)/2-1$	aCW_{min}	34	4.096
Voice (VO)	$(aCW_{min}+1)/2-1$	$(aCW_{min}+1)/2-1$	34	2.080
DCF	aCW_{min}	aCW_{max}	34	0[3]

Table 1.1: Default EDCA parameters for a Wi-Fi 6 STA. Adapted from [7] and [8], table 9-155.

1.1.2. Enhanced distributed channel access (EDCA)

The DCF protocol was succeeded in 2007 by the enhanced distributed channel access (EDCA) protocol. EDCA's uses the same CSMA/CA protocol, but its main improvement over DCF is the inclusion of priority levels, or access categories (ACs), for differing classes of data. Consider, for instance, that emails should likely have a lower priority than a real-time video call.

EDCA defines four of these ACs, which are, in ascending priority: background (BK), best effort (BE), video (VI) and voice (VO). Each AC has its own set of contention parameters, as shown in Table 1.1. As can be seen from the table, the CW minima and maxima change depending on the AC and the static DIFS from DCF is replaced by a variable arbitration inter-frame spacing (AIFS), which is used in the same way. Each AC now contends independently for the channel, with ties between ACs of the same STA being won by the highest priority AC. Effectively, this means that data belonging to a higher priority AC will wait for an idle channel for a shorter duration and will select a smaller value of b on average.

Another change from DCF is that each AC now defines a transmit opportunity (TXOP) duration. This TXOP allows a STA's AC that wins contention to transmit more than one data packet at a time. In fact, any AC that wins contention is allowed to transmit as much data as can fit and be acknowledged in the allotted TXOP.

Finally, note that EDCA, like the rest of the Wi-Fi specifications, is backwards compatible with DCF. This means that whenever a STA that only supports DCF is part of the network, it can contend fairly with the EDCA-enabled STAs, but it cannot make use of the prioritization EDCA's ACs provide.

1.1.3. Orthogonal frequency-division multiple access (OFDMA)

Contrary to DCF and EDCA, orthogonal frequency-division multiple access (OFDMA) is not a MAC protocol. Instead, it is a form of statistical multiplexing, which means that it allows multiple devices to use the same Wi-Fi channel simultaneously by subdividing the channel in some way. In OFDMA, this is achieved by using the fact that a single Wi-Fi 6 channel consists of many smaller carrier waves, called subcarriers or tones, each 78.125 kHz wide. By splitting the channel up into smaller groups of these subcarriers, called resource units (RUs), OFDMA effectively creates dynamically assignable mini channels for data transmission. These RUs can consist of 26, 52, 106, 242, 484, 996 or 2x996 tones[4], and RUs across a transmission may vary in size, allocating more bandwidth to STAs that require it. In this scheme, the access point (AP) decides how these RUs are allocated, by selecting the optimal supported RU distribution[5] for the current set of queued data packets.

In practice, this means that whenever the AP wins contention, it will first schedule OFDMA in the downlink (DL) direction (from AP to STAs), which allows it to send data to many STAs at once. When DL-OFDMA is completed, it will then schedule uplink (UL)-OFDMA allowing multiple STAs to transfer data to the AP at once.

[3] A TXOP duration of 0 means that only a single data packet may be sent at once.

[4] The 2x996-tone RU is designated as such because the RU can be split along two non-contiguous sets of subcarriers of equal size.

[5] For the full list of supported RU allocations, see [8], Table 27-26.

1.2. Problem statement and hypothesis

The flaw that was alluded to earlier lies in the combination of UL-OFDMA and the backoff mechanism of CSMA/CA. In theory, a station that was able to send data through UL-OFDMA, after which it still has data to send, should reset its backoff stage and counter. However, in practice, a lot of simulation studies do not actually include this ”backoff reset”[6] into their design, instead, they simply continue with the previous values of b and s after each UL-OFDMA transmission. While it is more likely that this omission is a simple oversight rather than a conscious design choice, the effects of this discrepancy with respect to the workings of a real-world Wi-Fi network are unknown. As such, this thesis aims to answer the question:

> **”What is the effect on performance metrics of resetting a station’s backoff stage and counter after an UL-OFDMA transmission?”**

To formulate a hypothesis for this problem statement, it is instructive to look at the effects of the value of b on the performance metrics and to look at the mathematical effects at play on the expected value of b ($E[b]$) during a backoff reset.

To examine the mathematical effects on $E[b]$, it is required to first define $E[b]$. As b is chosen uniformly from the range $[0, CW(s)]$, $E[b(s)]$ can be defined as

$$E[b(s)] = \frac{CW(s)}{2}. \tag{1.2}$$

CW was defined in Equation (1.1), which has been reproduced below for convenience.

$$CW(s) = min(2^s * (CW_{min} + 1) - 1, CW_{max}). \tag{1.1}$$

Next, remember that a backoff reset causes s to become 0, after which a new value for b is chosen. As such, the expected value of b after a backoff reset, $E[b_p]$, can be defined as

$$E[b_p] = E[b(0)] = \frac{CW(0)}{2} = \frac{CW_{min}}{2}, \tag{1.3}$$

where b_p is the value of b after (post) a backoff reset. To aid in analysis, define the difference between the values of b before and after a backoff reset (Δ_b) as follows:

$$\Delta_b(b_p, b_a) = b_p - b_a, \tag{1.4}$$

where b_a is the value of b before (ante) a backoff reset. The expected value of Δ_b is then

$$E[\Delta_b] = E[b_p] - E[b_a] = \frac{CW_{min}}{2} - E[b_a]. \tag{1.5}$$

$E[b_a]$ depends on the chance that the AP successfully transmits in a given time-slot. Bellalta and Kosek-Szott [9] have shown that this probability, dubbed a_4 in their paper, can be calculated as

$$a_4 = (1 - \alpha)(1 - \beta)\tau_{ap}(1 - \tau_{sta})^N, \tag{1.6}$$

where $1 - \alpha$ is the chance that the AP initiates a multi-user (MU) transmission, $1 - \beta$ is the chance that the AP initiates a UL transmission, N is the number of non-AP STAs in the network and τ_{ap} and τ_{sta} are the chances that the AP and a STA transmit in a given time-slot, respectively. This equation can be simplified by assuming a system where the AP always initiates MU transmissions and always schedules an UL transmission after each DL transmission, as is assumed in this thesis. This simplifies Equation (1.6) to the chance that the AP successfully contends for the channel in a given time-slot:

$$a = \tau_{ap}(1 - \tau_{sta})^N. \tag{1.7}$$

Unfortunately, both τ_{ap} and τ_{sta} do not have closed-form solutions in the general case. Thus, Bellalta and Kozek-Szott spend the rest of their paper calculating their values through fixed-point analysis.

[6] In this study, the term ”backoff reset” only ever relates to the reset of both backoff counter and stage that should occur after UL-OFDMA.

As such, a full mathematical analysis is out of scope for this preliminary analysis. Instead, one can turn to intuition. In a situation where at least some contention occurs, s is likely greater than 0. As each increment to s results in a doubling of $E[b(s)]$, up to $CW_{max}/2$, $E[b_a]$ is likely greater than $E[b_p]$. Therefore, the distribution of Δ_b is expected to skew negative for all but the smallest values of s. In other words, b is expected to decrease as a result of a backoff reset for all but the lowest contention systems.

Continuing on to the effect of b on the performance metrics of throughput, collision probability and latency, one can first see that a lower b as a result of a backoff reset after UL-OFDMA would likely result in more collisions in the short term, as the chance of multiple STAs picking the same value for b increases with the smaller CW. Secondly, from the foundational paper by Bianchi [10], it is known that collision probability is inversely related to throughput. Finally, it is relatively easy to see that in a system in which all other things remain equal, a reduction in throughput should result in an increase in latency, as the same size data packets are sent at a lower successful transmission rate.

Combining these two findings, the author formulates the following hypothesis:

”Resetting a station’s backoff stage and counter after an UL-OFDMA transmission results in an increase in collision probability, a decrease in throughput and an increase in latency for all but the lowest contention networks.”[7]

1.3. Contributions

The main contributions that this thesis makes are as follows. First, it studies and explains an effect, coined the ”backoff reset”, that is often overlooked in the scientific literature on Wi-Fi 6. Second, it makes a rough estimate that about two-thirds of the literature is affected by this omission. Thirdly, it investigates through simulation the effect on the five performance metrics of throughput, UL and DL latency and UL and DL collision probability. Through these simulations, it is found that not including the backoff reset in Wi-Fi 6 models will generally result in significant overestimation of the performance of a given Wi-Fi 6 network. Fourth and finally, it gives recommendations for future work and urges future Wi-Fi researchers to include the backoff reset in their future simulations of Wi-Fi 6.

[7] Note that if this hypothesis turns out to be true, the IEEE 802.11ax protocol could be improved by removing the backoff reset.

2
Prior art

This chapter has the goal of providing a scientific context to the rest of this thesis in two ways. First, it will list related work to draw a background to the intended research into the effects of the backoff reset. Secondly, a literature survey is used to estimate the prevalence of the backoff reset in the current scientific corpus.

2.1. Related work

The area of CSMA/CA has been thoroughly explored after its introduction by Kleinrock and Tobagi [11], due to its inclusion in many different network protocols. Research has been done, among others, on its performance in the general sense [12, 13], its stability under normal conditions [14], as well as on the total network collapse the protocol will experience under extreme conditions [15].

From this line of inquiry, more specialised papers were created specific to the Wi-Fi standard, such as Bianchi's foundational paper on DCF performance [10], followed by papers analyzing the performance of networks based solely on EDCA [16] or OFDMA [17].

Indeed, OFDMA has proven to be a popular topic of investigation for Wi-Fi researchers. Besides its performance having been holistically analyzed [18, 19], research has focused on many different aspects of its operation, such as UL-OFDMA [20, 21], RU allocation [22], its random access mode [17, 23] and its efficiency [24].

It can be thus be seen that both the CSMA/CA-based protocols and OFDMA have been extensively analyzed and documented. Some attention has also gone to their interplay in the Wi-Fi standard [9]. However, at least one oversight remained in the form of the backoff reset. As far as the author of this report is aware, no papers explore its effects or indeed have even identified its existence. Thus, the topic of this thesis was born.

2.2. Literature survey

This section aims to make a cursory survey of the literature available on IEEE 802.11ax research. To reiterate: no papers were found that specifically focus on the effects of the backoff reset. Instead, the goal of this survey is to give an estimate on the amount of literature that incorrectly disregards it. To accomplish this, a list of fifteen prominent papers analyzing 802.11ax performance was compiled and scoured for any mentions that might indicate that the backoff reset was taken into account. If no such mentions were found, it was assumed that the paper did not include the backoff reset. Now, it is good to mention that this approach is not fool-proof, as this is effectively an attempt to prove a negative. The backoff reset is a small implementation detail that is easy to overlook or under-report. However, a more accurate assessment was often impossible, as many years had passed since the authors of the paper had written it and in many circumstances the original code was lost.

This section will first list the papers that are believed to omit the backoff reset, after which the papers that did include the reset are discussed. Finally, a rough estimate is proposed of the amount of literature affected by this oversight.

2.2.1. Papers lacking the backoff reset

This section lists the papers that likely did not include the backoff reset mechanism.

To start off with, the paper ”Optimal Resource Allocation in IEEE 802.11ax Uplink OFDMA with Scheduled Access” by Dovelos and Bellalta [22] takes both a mathematical and simulated look at resource allocation schemes proposed for 802.11ax and proposes its own. Unfortunately, the paper does not mention the backoff reset or any of its related terms. The simulator used seems to be a proprietary one that was not made public and the mathematical model does not concern itself with backoff procedures. As such, it can be reasonably concluded that this paper does not take the backoff reset into account.

Next, the paper ”AP-initiated Multi-User Transmissions in IEEE 802.11ax WLANs” by Bellalta and Kosek-Szott [9] proposes a mathematical model for modelling Wi-Fi 6 MU transmissions that is then again validated in a proprietary simulator. From communications with Prof. Bellalta, it has become clear that the code for this simulator is unfortunately lost, however the mathematical model is quite extensively documented and seems to not include any backoff reset mechanism.

Continuing on, the paper ”Several EDCA Parameter Sets for Improving Channel Access in IEEE 802.11ax Networks” by Khorov et al. [25] aims to solve some of the unfairness that can occur when Wi-Fi 6 STAs co-exist with older Wi-Fi systems in the same network. For this, it proposes a solution based on differing EDCA parameter sets and analyzes the results through an analytical model. The paper or the model make no mention of the backoff reset, leading to the conclusion that it was not part of the considerations for this publication.

Next, the paper ”Performance evaluation of OFDMA and MU-MIMO in 802.11ax networks” by Daldoul et al. [19] analyzes the performance of OFDMA in pure Wi-Fi 6 networks. To accomplish this, an open-source simulator [26] was developed. From the source code, it becomes clear that the simulator does not correctly reset the backoff window and counter after a successful UL-OFDMA transmission, only resetting after a successful EDCA transmission.

Going further, the paper ”Performance Analysis of the IEEE 802.11ax MAC Protocol for Heterogeneous Wi-Fi Networks in Non-Saturated Conditions” by Lee [27] analyzes the 802.11ax MAC protocol for mixtures of Wi-Fi 6 and legacy devices in situations where the channel was not saturated[1]. This analysis is performed through a combination of numerical analysis and a proprietary simulator. Unfortunately, this simulator is not open-sourced, but the numerical analysis does not include a backoff reset-like mechanism, so it seems prudent to assume that this paper also does not take it into account.

Next, the paper ”OFDMA Uplink Scheduling in IEEE 802.11ax Networks” by Bankov [21] specifically only looks at scheduling schemes for UL-OFDMA, disregarding the EDCA system completely. It thus does not include any kind of backoff mechanism.

Then, the paper ”Enabling Massive Real-Time Applications in IEEE 802.11be Networks” by Avdotin et al. [28] puts forward a modification of the 802.11ax access rules that should allow future Wi-Fi 7 networks to be more suited for real-time applications. These rules were then implemented in a proprietary simulator and no mention is made of EDCA or its backoff systems. As such, this paper likely does not contain the backoff reset.

Going further, the paper ”An Efficient Backoff Procedure for IEEE 802.11ax Uplink OFDMA-Based Random Access” by Kozek-Szott and Domino [23] is concerned with finding a more efficient backoff algorithm for the random access mechanism that can be enabled for UL-OFDMA, called uplink orthogonal frequency-division multiple access random access (UORA). The paper considers only pure Wi-Fi 6 networks that have replaced EDCA completely with UORA and thus does not include the EDCA backoff mechanisms.

Similarly, the paper ”Performance analysis of the 802.11ax UL OFDMA random access protocol in dense networks” by Lanante et al. [17] only concerns itself with UORA-capable devices and does not include EDCA functionality at all. As such, the interplay between UORA and EDCA is not explored and the backoff reset is not included.

Finally, the paper ”Resource Management for OFDMA based Next Generation 802.11ax WLANs” by Karaca et al. [24] only concerns itself with algorithms to allocate RUs and transmission durations in a energy-constrained network. It thus similarly does not include any EDCA mechanisms.

[1]Saturated conditions mean that a STA always has more data to ready to send and so is never idle.

2.2.2. Papers that include the backoff reset

This section lists the papers that were known, or at least strongly suggest, to have included the backoff reset in their modelling.

Starting off, ”Will OFDMA Improve the Performance of 802.11 WiFi Networks?” by Avallone et al. [18] aims to forecast the expected performance improvements that the inclusion of UL and DL OFDMA will create, compared to single-user (SU) transmissions. This is done through simulations of both saturated and unsaturated network conditions. While the paper itself does not mention the backoff reset, the authors were the ones who implemented the correctly functioning 802.11ax model in the ns-3 [3] simulator. As such, it is known that their simulation included the backoff reset.

Continuing on, the paper ”Impact of MU EDCA channel access on IEEE 802.11ax WLANs” by Naribole et al. [29] investigates the impact of having STAs that recently transmitted data through UL-OFDMA defer further contention for a small amount of time. For this research, the authors used a modified version of ns-3 [3], which is known to contain the backoff reset. While one cannot be sure whether the modification disabled the reset, it is highly likely that care was taken to model it accurately, due to the subject matter being closely related to the backoff reset.

Next, the paper ”Uplink Resource Allocation in IEEE 802.11ax” by Bhattarai et al. [30] aims to define a scheme that maximizes network throughput by assigning the RUs in an UL-OFDMA transmission to be either pre-allocated or open for contention. To analyze the effects, it uses the ns-3 simulator [3], with its known-good implementation of the backoff reset.

Further, the paper ”Performance Analysis of Uplink Multi-User OFDMA in IEEE 802.11ax” by Naik et al. [20] investigates the performance of the 802.11ax MAC layer through an analytical model and simulations. It also concerns itself with finding a balance between high network throughput and providing quick access to the network for newly joining STAs. As the simulations are once again implemented in ns-3 [3], it is extremely likely that the backoff reset was included.

Finally, the paper ”Validation of the ns-3 802.11ax OFDMA Implementation” by Magrin et al. [31] verifies the correct workings of the 802.11ax OFDMA implementation of ns-3 [3]. As it is known from previous papers that the ns-3 [3] implementation correctly implements the backoff reset and that this paper specifically sought not to modify the system under test, it can be concluded that the backoff reset was part of the simulations.

Authors and citation	Backoff Reset	Unsaturated/ Saturated	Simulation/ Numerical	Open-source simulator
Dovelos and Bellalta [22]	No	Saturated	Both	No
Bellalta and Kosek-Szott [9]	No	Saturated	Both	No
Khorov et al. [25]	No	Saturated	Numerical	N/a
Daldoul et al. [19]	No	Both	Simulation	Yes, custom
Lee [27]	No	Unsaturated	Both	No
Bankov [21]	No	Unsaturated	Simulation	No
Avdotin et al. [28]	No	Both	Simulation	No
Kozek-Szott and Domino [23]	No	Both	Simulation	No
Lanante et al. [17]	No	Saturated	Both	No
Karaca et al. [24]	No	Saturated	Both	No
Avallone et al. [18]	Yes	Both	Simulation	Yes, ns-3
Naribole et al. [29]	Yes	Saturated	Simulation	Yes, ns-3
Bhattarai et al. [30]	Yes	Saturated	Both	Yes, ns-3
Naik et al. [20]	Yes	Saturated	Both	Yes, ns-3
Magrin et al. [31]	Yes	Both	Both	Yes, ns-3

Table 2.1: Overview of papers considered for backoff reset inclusion.

2.2.3. Prevalence of backoff reset in literature

Table 2.1 shows an overview of the fifteen papers included in this census. Out of these papers, only five are known to have very likely included the backoff reset in some capacity. All of these papers utilized the ns-3 [3] simulator, whose source-code is open-sourced and could thus be inspected for its implementation of the backoff reset. Out of the papers who did not include the backoff reset, only ”Performance evaluation of OFDMA and MU-MIMO in 802.11ax networks” by Daldoul et al. [19] used a simulator that was open to inspection. Most other papers used a closed-source simulator, which meant that assumptions had to be made as to whether the backoff reset was included, based on the contents of the papers. The final category were articles that focused on numerical analysis alone, which could be easily inspected for inclusion of the backoff reset.

From both the related work and from this analysis, it is clear that the backoff reset is not only not studied directly in literature, it is generally not even included by accident in models and simulations. Thus, answering the problem statement of this thesis can provide valuable insight in the severity of the effects of this oversight.

3 Method

This chapter lays out the method by which this report answers the problem statement.

3.1. Simulator

The simulator used in this research is a more advanced version of the Wi-Fi 6 MAC Simulator by Gokhale and Kroep [32] first introduced in the paper by Gokhale et al. on "Toward Enabling High-Five Over WiFi: A Tactile Internet Paradigm" [33]. This simulator is flexible enough to be used in a more general setting, and this report is not likewise specific to the tactile internet. The choice for this simulator was a practical one: Dr. Gokhale is one of the supervisors of the author and as such was available to provide technical support and modifications where necessary. For this research, it was necessary to add support for bounded random traffic generation rates and RU allocation proportional to the amount of data queued at each STA.

3.2. Network conditions

Parameter	Value	Unit
Channel width	80	MHz
Spatial streams	1	
Modulation coding scheme (MCS) index	9	
Guard interval (GI)	0.8	µs
Access category (AC)	Voice (VO)	
Channel capacity	480.4	Mbit/s

Table 3.1: Wi-Fi channel parameters

This report considers a situation with perfect channel conditions. That means that the channel experiences no noise, is completely isolated from other wireless signals and propagation delays are uniform across all devices. Additionally, data is generated directly at the MAC layer of the AP and STAs, so no further network or computational delays are modelled. Table 3.1 shows the Wi-Fi channel parameters used. Note that all data is sent over the VO AC to guard against effects resulting from inter-AC contention.

Variable	Default value	Unit	Non-default experiments	Inapplicable experiments
Number of STAs	8		1,8	
AP CW$_{min}$	32		2,3,5-7,9-12	
STA CW$_{min}$	32		2,3,5-7,9-12	
CW$_{max}$	1024			
Fixed data generation rate	5	kHz	4,7	8-12
Minimum random data generation rate	4	kHz		1-7
Maximum random data generation rate	6	kHz		1-7
Proportional RU allocation	False		5-7,11,12	
Packet size	300	B		
Simulation time	10	s		
Observations per data set	10			

Table 3.2: Controlled variables for all experiments. Experiment numbers where values other than the defaults are used or where variables do not apply are listed.

3.3. Intention behind experiments

The overall goal of the experiments in this thesis is to investigate the differences in performance on five metrics between a system implementing the backoff reset correctly, called the ”reset” system, and a system that does not and thus continues with its previous value of b and s, the ”continue” system. As will be explained in Section 3.4, each experiment modifies some set of input variables to find out how the different systems behave under different circumstances. The combination of all these experiments will then act as a design-space exploration on the effects of the backoff reset.

3.4. Inputs and outputs

Table 3.2 lists the controlled input variables and defaults for all experiments, as well as the experiment numbers in which non-default values are used or variables do not apply. Refer to the relevant experiment for the modified values.

Out of these controlled variables, four are constant. First, CW$_{max}$ is not modified as it only serves as a bound on another variable, CW$_{min}$. Secondly, packet size is not expected to provide new insights within the bounds of the simulator’s capabilities. Finally, simulation time and observations per data set are simply constraints on the amount of data generated for each experiment. Furthermore, both proportional RU allocation and the minimum and maximum data generation rates serve as toggles for realism. They are not expected to provide new insights beforehand, but are taken into account to make sure that more abstract systems without them still provide accurate data in a realistic scenario. Finally, that leaves the controlled variables that are expected to generate actually interesting results. First, the number of non-AP STAs in the network aims to find whether the ”reset” system performs worse or better for certain network sizes. Second, the AP and STA CW$_{min}$ variables can be varied independently to find out whether the ”reset” and ”continue” systems require different values for CW$_{min}$ to perform optimally. Finally, the data generation rate allows one to evaluate the two systems under differing levels of network load. Note that the DL data rate is per STA. For example, for an AP in a network with 8 other STAs, generating data at a fixed data generation rate of 5 kHz, the total rate of packet generation at the AP is 5 kHz $* 8 = 40$ kHz.

For each experiment, the collision probability and latency in both DL and UL directions are recorded, as well as the the overall network throughput.

3.5. Method of performance evaluation

In Chapter 4, the results from each experiment are listed and evaluated. To prevent unnecessary repetitions, this section will explain the analysis for both two-dimensional and three-dimensional dataset collections[1].

3.5.1. Two-dimensional dataset collections

Experiments with only a single controlled variable will result in a two-dimensional dataset collection for each performance metric. To analyze the significance of the outcomes, a paired, two-tailed t-test with $\alpha = 0.05$ was performed on the means of each pair of data sets. Subsequently, the effect size was determined through Glass' Δ, with the "continue" system data as the control group. For each two-dimensional experiment, a table shows these effect sizes for each of the output variables, given that the effect is significant. These effect sizes are interpreted based on the rules of thumb defined by Sawilowsky in [34].

3.5.2. Three-dimensional dataset collections

The dataset collections from experiments with two controlled variables are three-dimensional in nature. This makes plotting them in a comprehensible way difficult. As such, it was decided not to present the raw results for these tests, but instead to solely present the significant effect sizes in heatmaps.

As in the two-dimensional experiments, a paired, two-tailed t-test with $\alpha = 0.05$ was performed on the means of each pair of data sets to analyze the significance of the outcomes. Subsequently, the effect size was determined through Glass' Δ, with the "continue" system data as the control group. These effect sizes are then interpreted based on the rules of thumb defined by Sawilowsky in [34].

[1] The term dataset denotes a single set of observations on a single performance metric for a single input. In contrast, a dataset collection is the set of datasets for a single performance metric across all inputs for an experiment.

4
Results

This chapter details the experiments that gave significant and/or surprising results. Additional experiments that did not provide new insights can be found in Appendix A. These experiments are Experiments 9 (Section A.1), 10 (Section A.2) and 12 (Section A.3). Each experiment details first its intended purpose and non-default inputs, then displays the outputs and the results from statistical analysis.

Number of STAs	Throughput	DL latency	UL latency	DL coll. prob.	UL coll. prob.
1	-	-5.85	24.12	2.81	-2.00
2	-	-	9.45	-1.33	-1.39
3	-	-	-	-2.33	-
4	-1.03	1.33	-	-1.67	1.57
5	-0.79	0.70	-	-	1.83
6	-	-	-	2.50	4.03
7	-	0.69	-	4.43	6.49
8	-0.82	1.53	1.35	7.33	6.57

Table 4.1: Significant Glass' Δ for experiment 1. Bold values are controlled variable values. Positive values indicate improvement for the "reset" system over the "continue" system for throughput, while the same is indicated by negative values for the other measurements. Dashes indicate insignificant results.

4.1. Experiment 1: Fixed rate traffic with varying number of stations

This first experiment aimed to give a rough insight into the way that performance differs between networks of differing sizes with and without backoff resets. To accomplish this, all parameters were kept at their default values, except the number of non-AP STAs, which varied from 1 through 8. Note that further increases were not possible, as the simulator did not support larger networks.

4.1.1. Results

The sub-figures of Figure 4.1 show the outputs of this experiment. Figure 4.1a and Figure 4.1b show the UL and DL collision probability, Figure 4.1c and Figure 4.1d show the UL and DL latency, and Figure 4.1e shows the overall throughput of the network.

4.1.2. Performance evaluation

Table 4.1 shows these effect sizes for each of the output variables, given that the effect is significant, as explained in Section 3.5.1. From Table 4.1, it can be seen that the system with a correct backoff reset mechanism will occasionally experience a large reduction in throughput and a medium to huge increase in latency, except on the DL for a single STA. However, on average, the two systems perform mostly identically on throughput and UL latency. More interesting are the results for the collision probability: they show very large to huge effects across the board, but the effect is a pure improvement over "continue" system only as long as the number of connected STAs is between one and four. For systems with one or four STAs, conflicting effects appear on the DL and UL. After this point, the "reset" system will consistently perform worse than "continue" system.

Taking a more holistic view, it can be seen that the "reset" system will perform worse on all metrics compared to "continue" system once the number of STAs in the system increases beyond four. Before this point, the "reset" system may be seen as a specialized tool: in very niche situations, where one metric is much more important than the others, it might be useful to consciously enable or disable the backoff reset. These results are mostly in line with expectations, in that as more STAs join the network, contention for the channel increases and the "reset" system will start performing worse due to a greater number of collisions. The lack of differentiation between the two systems on throughput can be explained by the fact that the channel is not saturated and thus delayed packets due to collisions will eventually still be transmitted.

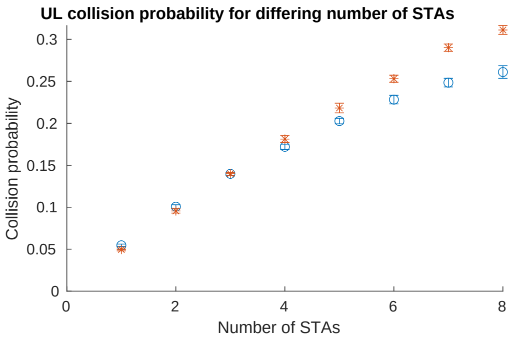

(a) UL collision probability for experiment 1

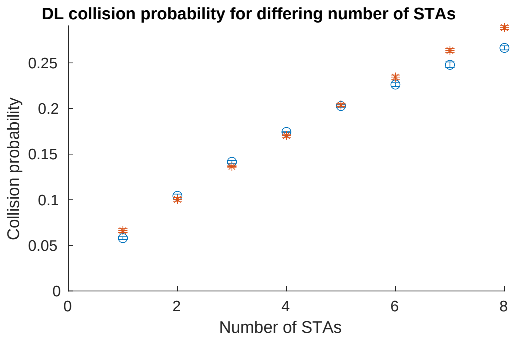

(b) DL collision probability for experiment 1

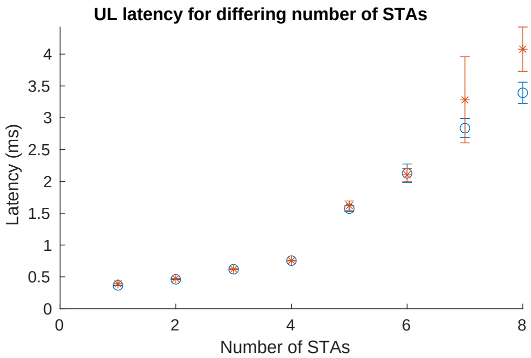

(c) UL latency for experiment 1

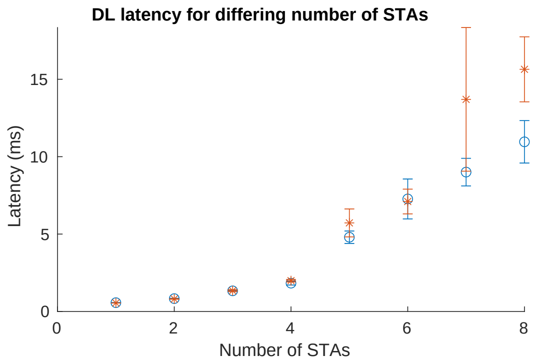

(d) DL latency for experiment 1

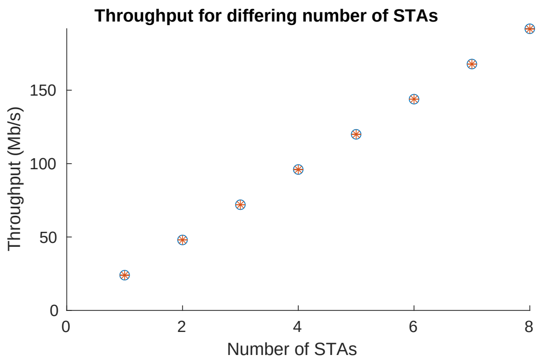

(e) Throughput for experiment 1

Figure 4.1: Results for experiment 1. "Continue" system data is plotted in blue circles, "reset" system data is plotted in orange asterisks.

CW_{min}	Throughput	DL latency	UL latency	DL coll. prob.	UL coll. prob.
1	-4.11	1.71	2.18	-4.23	17.04
2	-1.23	0.97	2.44	1.94	20.56
4	-1.29	2.20	2.09	1.74	7.88
8	-1.18	0.66	0.72	2.29	4.99
16	-	0.97	0.93	6.09	9.11
32	-	-	-	4.90	5.56
64	-	-	-	-	-
128	-	-	-	-2.56	-
256	-	-1.09	-	-2.88	-
512	-	-2.82	2.96	-2.92	-1.46
1024	-	-2.04	6.85	-	-

Table 4.2: Significant Glass' Δ for experiment 2. Bold values are controlled variable values. Positive values indicate improvement for the "reset" system over the "continue" system for throughput, while the same is indicated by negative values for the other measurements. Dashes indicate insignificant results.

4.2. Experiment 2: Fixed rate traffic with varying symmetric CW_{min}

After getting a feel for the effects of the backoff reset in different sizes of network, it was thought that varying the size of CW_{min} symmetrically[1] would give a different perspective on the effects of the backoff reset. After all, increasing CW_{min} increases the effect of the linear increase from b to $E[b]$ and decreases the exponential reduction as a result of setting s to 0. Therefore, for this experiment CW_{min} varied symmetrically from 1 to 1024 (CW_{max}).

4.2.1. Results

The sub-figures of Figure 4.2 show the outputs of this experiment. Figure 4.2a and Figure 4.2b show the UL and DL collision probability, Figure 4.2c and Figure 4.2d show the UL and DL latency, and Figure 4.2e shows the overall throughput of the network.

4.2.2. Performance evaluation

Table 4.2 shows these effect sizes for each of the output variables, given that the effect is significant, as explained in Section 3.5.1. Table 4.2 shows a consistent very large to huge decrease in throughput compared to "continue" system for small CW_{min}, after which the two systems appear to perform identically. Latency sees a medium to huge increase in both the UL and DL directions. Notable exceptions are found in the DL direction for bigger CW_{min}, instead showing large to huge decreases in latency. Collision probability shows very large to huge increases for CW_{min} less than or equal to 32, with the notable exception of a huge decrease for $CW_{min} = 1$. Beyond that point, the collision probability reduces by a very large to huge effect.

Looking at the bigger picture, it can be seen that the "reset" system has, in general, much worse performance than the "continue" system for $CW_{min} \leq 32$, while performance increases for higher CW_{min}, so long as a greater UL latency is acceptable for the intended use-case. While it was expected that the "reset" system would perform worse when CW_{min} is small, as the STAs will have a smaller number of values to choose from for b after a backoff reset, the increased performance for greater values of CW_{min} is not so easily explained. From theory it can be concluded that as CW_{min} increases, the negative effects of a backoff reset lessen, as the probability of two or more STAs choosing the same value of b decreases. As such, the author theorizes that other, smaller effects become more pronounced, such as the fact that in the "reset" system, all STAs that choose an initial value of b larger than the AP effectively do not contend for the channel as long as the AP wins contention, as they will reset their b during the backoff reset. Thus, the number of UL and DL collisions will decrease, as the network effectively becomes smaller from a contention standpoint. However, as these STAs with a b larger than the AP will have to wait for at least the AP DL transmission to complete, the UL latency will increase. In

[1] In this context, symmetric means that AP and STAs have the same CW_{min}

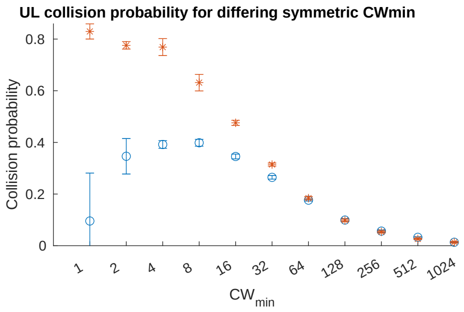

(a) UL collision probability for experiment 2

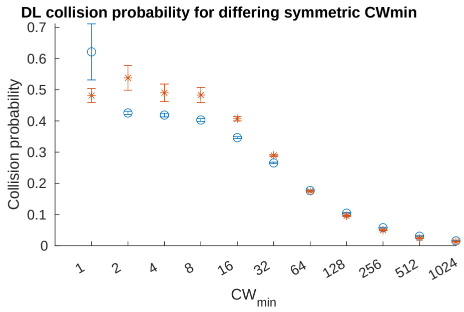

(b) DL collision probability for experiment 2

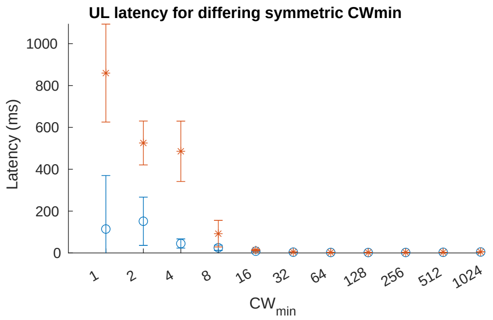

(c) UL latency for experiment 2

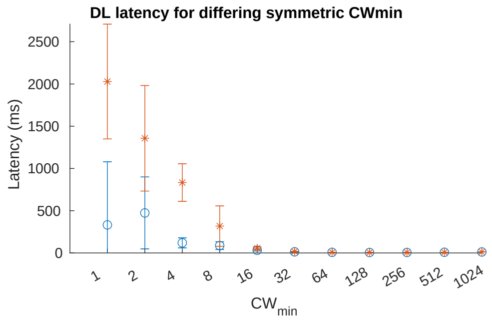

(d) DL latency for experiment 2

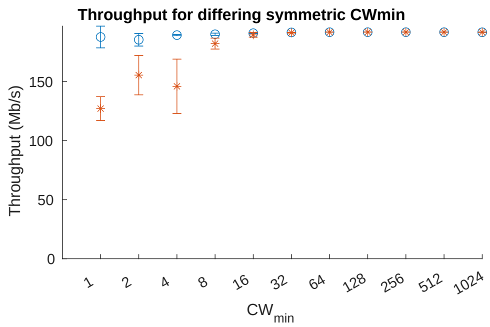

(e) Throughput for experiment 2

Figure 4.2: Results for experiment 2. "Continue" system data is plotted in blue circles, "reset" system data is plotted in orange asterisks.

contrast, the AP will have to wait for less STA on average to complete their UL transmissions before it can transmit, thus reducing its DL latency.

4.3. Experiment 3: Fixed rate traffic with varying asymmetric CW$_{min}$

The author theorized that varying CW$_{min}$ asymmetrically, meaning the CW$_{min}$ of AP and STAs could vary independently, could prove beneficial for Wi-Fi 6 networks, as a situation where the AP CW$_{min}$ is smaller than the CW$_{min}$ of the STAs would allow the AP to win contention more often. This would then translate in UL-OFDMA being scheduled more often, possibly reducing collision probability and latency for all STAs in the process. To study this, CW$_{min}$ was allowed to vary from 1 to 1024 (CW$_{max}$) independently for both AP and STAs. For completeness sake, situations where the AP CW$_{min}$ was greater than or equal to the STA CW$_{min}$ were included in the experiment, although no beneficial effect of this situation was expected a priori.

4.3.1. Results and performance evaluation

The significant effect sizes for this experiment are plotted in the heatmaps of Figure 4.3, as explained in Section 3.5.2. Figure 4.3a and Figure 4.3b show the UL and DL collision probability, Figure 4.3c and Figure 4.3d show the UL and DL latency, and Figure 4.3e shows the overall throughput of the network.

From Figures 4.3a and 4.3b, it can be gleaned that the effects on collision probability of modifying the AP and STA CW$_{min}$ values are almost uniformly huge, when compared to the ”continue” system. For UL collision probability, the ”reset” system performs better than ”continue” system roughly when STA CW$_{min}$$\geq 64$ and AP CW$_{min}$ ≥ 8, while on the DL the AP CW$_{min}$ does not seem to matter much for the same STA CW$_{min}$ range, except for the extreme case of AP CW$_{min}$$= 1$. Meanwhile, Figures 4.3c and 4.3d show a much more cohesive picture for the UL and DL latency: values of AP CW$_{min}$$\leq 16$ result in medium to huge increases in both latencies, but this effect can be offset once STA CW$_{min}$$\geq 64$ for higher values of AP CW$_{min}$. Outside of these two ranges, the two systems experience similar latencies. Finally, the effects on throughput are almost universally strongly negative for small AP and STA CW$_{min}$ values. For other values, there is no evidence to suggest that the two systems differ in throughput. Two positive effects are observed, but their placement in a sea of insignificant results does not allow for generalizations as to the interplay between higher AP and STA CW$_{min}$ values on throughput.

Overall then, the effect of small AP and STA CW$_{min}$ values results in significantly lower performance on all key metrics for the ”reset” system. The relation flips once STA CW$_{min}$$\geq 64$, with the ”reset” system strongly outperforming the ”continue” system on all measures except throughput. This fits in the hypothesis set forth in Experiment 2, Section 4.2.2. As the STA CW$_{min}$ increases, the performance reduction of the ”reset” system lessens as the negative effects of the backoff reset reduce and the positive effects of having the AP contend for the channel for all systems that chose a larger initial value of b increase.

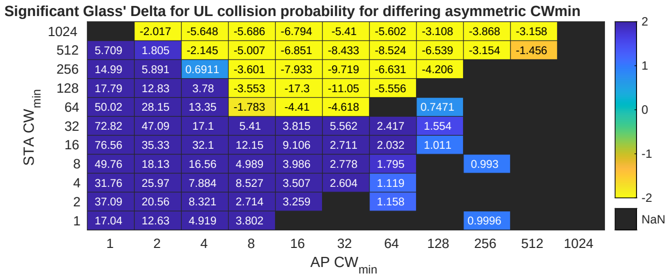

(a) Significant effect sizes on UL collision probability for experiment 3

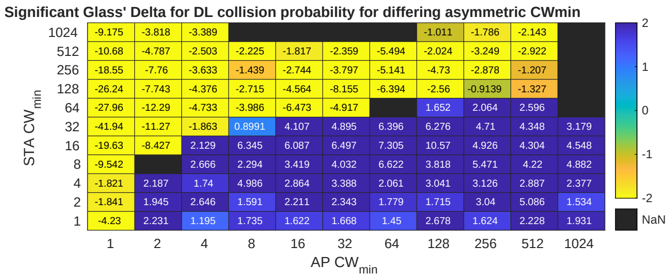

(b) Significant effect sizes on DL collision probability for experiment 3

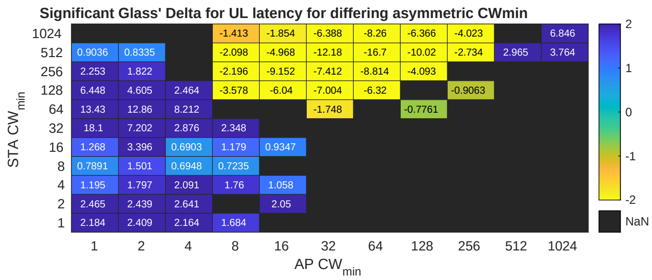

(c) Significant effect sizes on UL latency for experiment 3

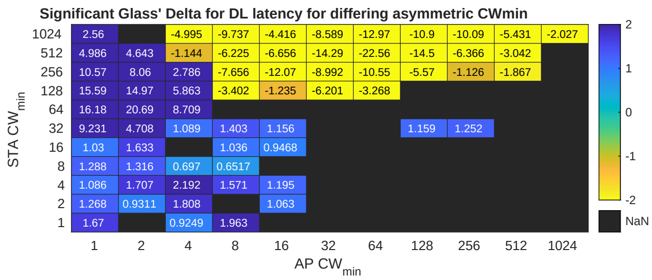

(d) Significant effect sizes on DL latency for experiment 3

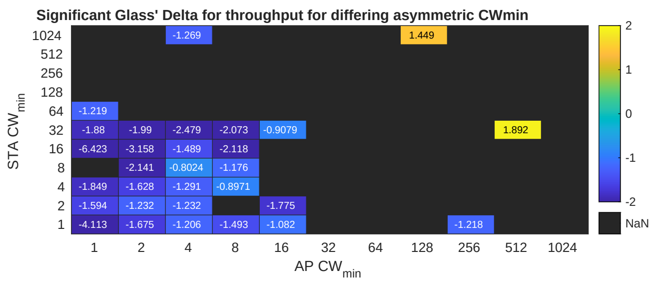

(e) Significant effect sizes on throughput for experiment 3

Figure 4.3: Significant Glass' Δ for experiment 3. Yellow shaded values indicate improvement for the ”reset” system over the ”continue” system, while blue shaded values indicate regression. NaN values, colored black, indicate insignificant results.

4.4. Experiment 4: Fixed rate traffic with varying traffic generation rate

The goal of this experiment was to see how the effect of a backoff reset changes as the load on the channel increases. To study this, the rate at which data was generated at both AP and STAs was varied from 100 Hz to 7000 Hz, the limit of what the simulator could support.

4.4.1. Results

The sub-figures of Figure 4.4 show the outputs of this experiment. Figure 4.4a and Figure 4.4b show the UL and DL collision probability, Figure 4.4c and Figure 4.4d show the UL and DL latency, and Figure 4.4e shows the overall throughput of the network.

4.4.2. Performance evaluation

Table 4.3 shows these effect sizes for each of the output variables, given that the effect is significant, as explained in Section 3.5.1. From Table 4.3, we see sporadic, medium to very large reductions in throughput for the ”reset” system. The effects on the latencies are slightly less sporadic, indicating an overall medium to very large increase in latency across the board. In general, however, the two systems perform similarly on throughput and latency. A much more concrete effect can be seen when it comes to the collision probabilities. Both UL and DL see consistent, large to huge increases in collision probability when comparing the ”reset” system to the ”continue” system for sample rates greater than 1100 Hz.

Overall then, it can be seen that the ”reset” system strongly under-performs on collision probability, regardless of network load, with the other metrics on average appearing unaffected. This increase in collision probability for the ”reset” system can be explained by the fact that the AP often schedules UL-OFDMA and thus the negative effects of the backoff reset on collision probability for low CW_{min} values is more often realized. The lack of a similar effect on throughput and the latencies can be explained by the fact that the channel is not saturated and thus data will not be dropped.

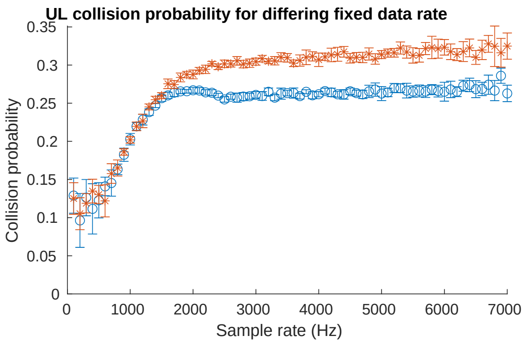

(a) UL collision probability for experiment 4

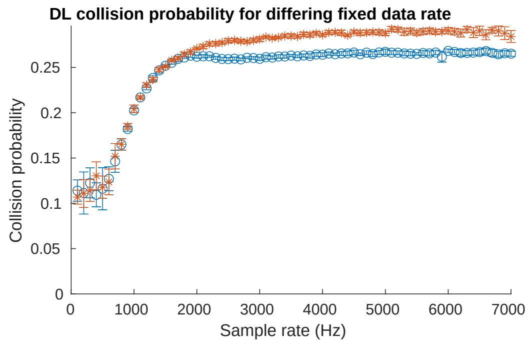

(b) DL collision probability for experiment 4

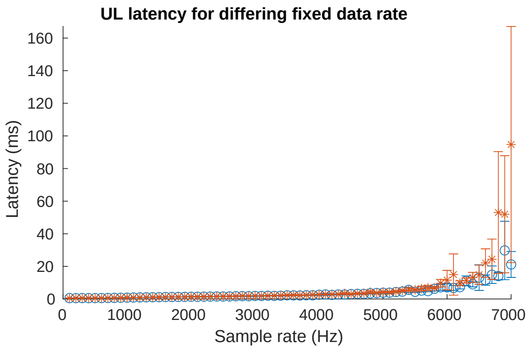

(c) UL latency for experiment 4

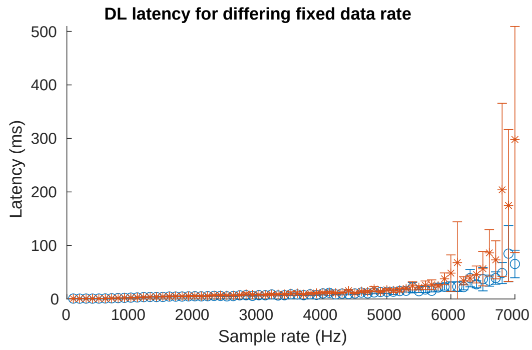

(d) DL latency for experiment 4

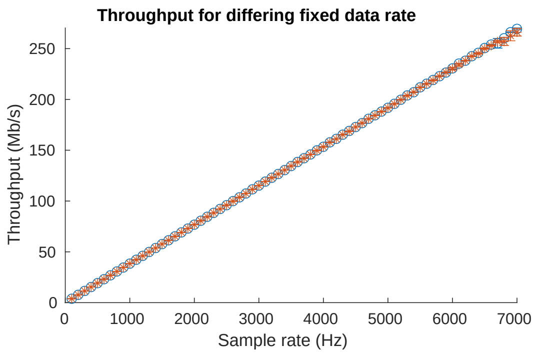

(e) Throughput for experiment 4

Figure 4.4: Results for experiment 4. ”Continue” system data is plotted in blue circles, ”reset” system data is plotted in orange asterisks.

Sample rate (Hz)	Throughput	DL latency	UL latency	DL coll. prob.	UL coll. prob.
100	-	-	-	-0.64	-
200	-	-	-	-	-
300	-	-	-	-	-
400	-	-	-	-	-
500	-	-	-	-	-
600	-	-	-	-	-
700	-	-	-	-	-
800	-	-	-	-	-
900	-	-	-	-	-
1000	-	-	-	-	-
1100	-	-	-	-	-
1200	-	-	-	1.00	-
1300	-	-	-	-	1.19
1400	-	0.69	0.64	-	1.20
1500	-	-	-	-	-
1600	-	-	-	1.63	1.65
1700	-0.92	-	-	-	1.38
1800	-	-	-	1.07	2.15
1900	-	-	-	1.24	3.51
2000	-	-	-	4.74	2.47
2100	-	-	-	3.02	4.25
2200	-	-	-	4.24	3.84
2300	-	-	-	4.51	8.41
2400	-	-	-	7.45	7.06
2500	-	1.09	0.88	5.79	6.35
2600	-	1.16	1.25	7.22	7.19
2700	-	-	-	6.49	5.91
2800	-	-	0.89	6.85	5.88
2900	-	3.13	2.46	5.94	5.41
3000	-	-	-	5.64	6.28
3100	-	-	-	11.92	7.72
3200	-	-	-	9.59	8.48
3300	-	1.34	1.66	10.65	6.06
3400	-	1.24	1.20	8.30	5.49
3500	-	-	-	5.90	5.67
3600	-	-	-	8.56	6.11
3700	-0.64	1.38	1.25	6.55	4.29
3800	-	-	-	6.48	3.21
3900	-	1.80	2.17	7.98	5.81
4000	-	-	-	8.94	3.52
4100	-	-	-	7.04	6.49
4200	-0.68	1.46	1.61	8.87	3.79
4300	-	-	-	5.19	3.92
4400	-1.43	1.59	1.79	9.51	5.46
4500	-	-	-	7.97	4.73
4600	-	-	-	5.50	5.16
4700	-	-	-	6.03	4.85
4800	-	0.97	-	7.67	4.22
4900	-	-	-	5.20	3.99
5000	-	-	-	5.70	6.74
5100	-	-	-	6.26	6.60
5200	-	-	-	6.32	5.58
5300	-	1.08	-	4.11	4.57
5400	-	-	-	5.21	4.33
5500	-0.79	1.49	1.46	5.02	3.22
5600	-	-	-	4.69	3.59
5700	-	0.83	0.91	5.61	4.86
5800	-	-	-	5.38	2.57
5900	-	0.94	0.74	7.89	3.06
6000	-	-	-	4.34	3.92
6100	-	-	-	4.24	3.30
6200	-	1.03	1.15	3.47	4.22
6300	-	-	-	5.13	2.37
6400	-	0.79	0.92	2.76	2.86
6500	-	-	-	3.13	3.12
6600	-1.34	0.82	0.82	2.22	2.77
6700	-	0.66	0.53	5.31	3.31
6800	-	-	0.71	3.23	1.50
6900	-	-	-	2.16	1.09
7000	-	0.75	0.70	1.99	2.46

Table 4.3: Significant Glass' Δ for experiment 4. Bold values are controlled variable values. Positive values indicate improvement for the ”reset” system over the ”continue” system for throughput, while the same is indicated by negative values for the other measurements. Dashes indicate insignificant results.

CW$_{min}$	Throughput	DL latency	UL latency	DL coll. prob.	UL coll. prob.
1	-2.44	2.32	2.89	-	27.06
2	-1.14	1.12	2.03	1.22	9.38
4	-1.00	1.01	1.30	2.36	9.53
8	-	-	0.81	7.33	23.27
16	-	0.94	1.12	7.59	6.10
32	-	1.51	1.29	8.21	4.33
64	-	-	-	-	1.94
128	-	-1.16	-1.23	-3.16	-1.15
256	-	-3.18	-	-6.09	-2.49
512	-	-3.43	2.59	-1.55	-1.32
1024	-	-2.17	5.47	-1.18	-

Table 4.4: Significant Glass' Δ for experiment 5. Bold values are controlled variable values. Positive values indicate improvement for the ”reset” system over the ”continue” system for throughput, while the same is indicated by negative values for the other measurements. Dashes indicate insignificant results.

4.5. Experiment 5: Fixed rate traffic with proportional RU allocation and varying symmetric CW$_{min}$

A parameter that was so far left untouched was the way in which RUs were allocated for UL-OFDMA. Previously, all STAs that had data to send would be scheduled simultaneously and with equal RU widths. However, it can be far more efficient to allocate wider RUs to STAs with more queued data and to allocate smaller RUs to STAs with less data to send. Indeed, this is exactly what the IEEE 802.11ax standard proposes and this experiment therefore tried to come slightly closer to the real-world behaviour of the protocol. During this experiment, CW$_{min}$ was allowed to vary symmetrically again from 1 to 1024 (CW$_{max}$) to allow for comparisons with Experiment 2.

4.5.1. Results

The sub-figures of Figure 4.5 show the outputs of this experiment. Figure 4.5a and Figure 4.5b show the UL and DL collision probability, Figure 4.5c and Figure 4.5d show the UL and DL latency, and Figure 4.5e shows the overall throughput of the network.

4.5.2. Performance evaluation

Table 4.4 shows these effect sizes for each of the output variables, given that the effect is significant, as explained in Section 3.5.1. Table 4.4 shows a large to huge reduction in throughput for low values of CW$_{min}$. Both types of latency are increased for CW$_{min} \leq 32$ by large to huge effects. Above this point, the relation flips on the DL and the results show large to huge reductions in latency for the ”reset” system. A similar situation occurs with the collision probabilities: for CW$_{min} \leq 64$ the collision probability is increased by large to huge magnitudes, while above this point the probability drops by similar effect sizes.

Zooming out, one can see that the ”reset” system strongly outperforms the ”continue” system on all metrics, except UL latency, once the symmetric CW$_{min}$ exceeds 64. Below this point, the ”continue” system should be strongly preferred. Comparing these effects to the effects of Experiment 2 in Table 4.2, one can see that the inclusion of proportional RU allocation did not influence the relationship between the CW$_{min}$ and the key metrics. This is to be expected, as all STAs have the same amount of queued data at any point in time due to having the same packet size and data rate, so a proportional RU allocation scheme would choose the same even distribution of RUs as the previous fixed scheme. Experiment 11 in Section 4.9 investigates whether a bounded random data rate will significantly change these results.

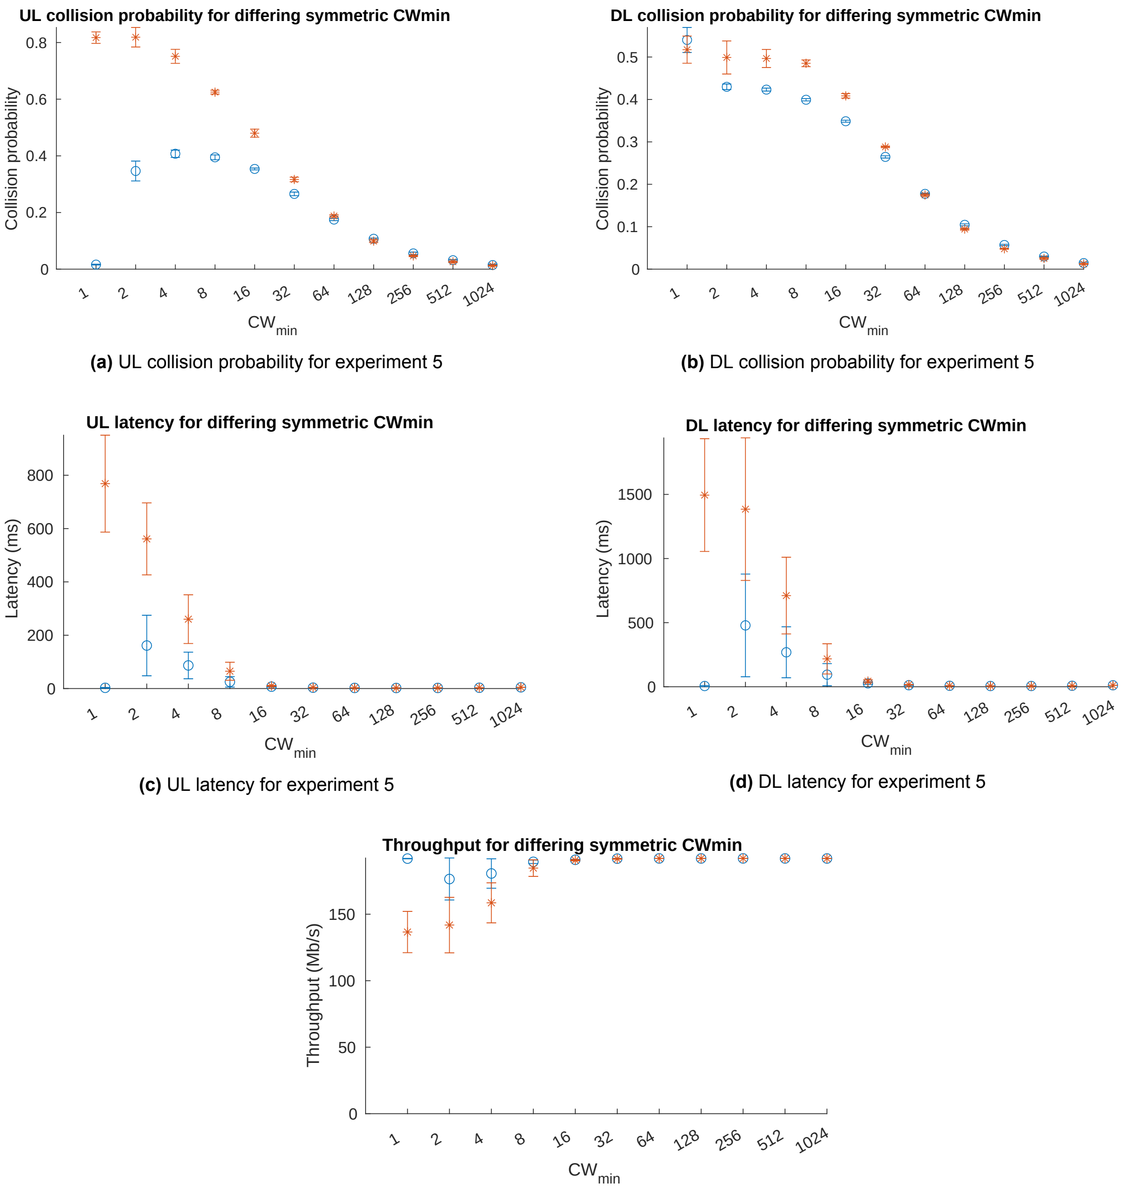

(a) UL collision probability for experiment 5

(b) DL collision probability for experiment 5

(c) UL latency for experiment 5

(d) DL latency for experiment 5

(e) Throughput for experiment 5

Figure 4.5: Results for experiment 5. ”Continue” system data is plotted in blue circles, ”reset” system data is plotted in orange asterisks.

4.6. Experiment 6: Fixed rate traffic with proportional RU allocation and varying asymmetric CW$_{min}$

This experiment aimed to see if the hypothesised latency and collision probability improvements of Experiment 3 would continue to present themselves when RU allocation was proportional. To study this, CW$_{min}$ was allowed to vary from 0 to 1024 (CW$_{max}$) independently for both AP and STAs. Again, situations where the AP CW$_{min}$ was greater than or equal to the STA CW$_{min}$ were included in the experiment, although no beneficial effect of this situation was expected a priori.

4.6.1. Results and performance evaluation

The significant effect sizes for this experiment are plotted in the heatmaps of Figure 4.6, as explained in Section 3.5.2. Figure 4.6a and Figure 4.6b show the UL and DL collision probability, Figure 4.6c and Figure 4.6d show the UL and DL latency, and Figure 4.6e shows the overall throughput of the network.

Figure 4.6 shows a very similar trend across all metrics as Figure 4.3 did for Experiment 3. Once again, using small AP and STA CW$_{min}$ values results in significantly worse performance for the ”reset” system over ”continue” system. However, once STA CW$_{min}$ ≥ 64, overall performance improves for all metrics. Analogously to Experiment 5, the effect of proportional RU allocation on the key metrics when varying CW$_{min}$ asymmetrically is negligible. The reasoning is also similar: due to the fixed-rate data generation and fixed packet size, the proportional scheme allocates the same RUs to each STA as the fixed scheme. The random-rate counterpart of this experiment is Experiment 12, found in Section A.3.

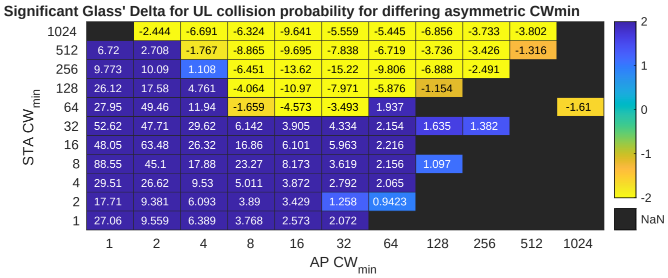

(a) Significant effect sizes on UL collision probability for experiment 6

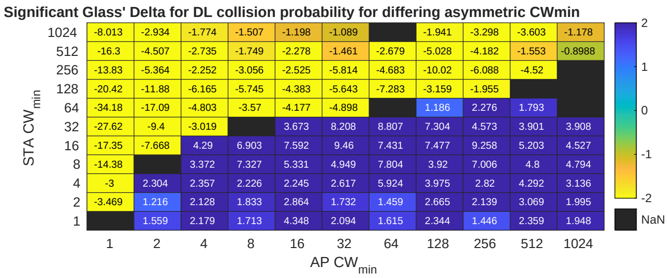

(b) Significant effect sizes on DL collision probability for experiment 6

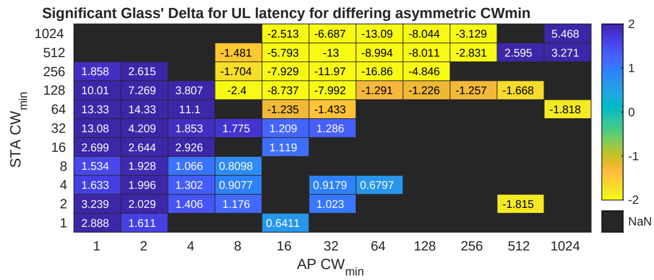

(c) Significant effect sizes on UL latency for experiment 6

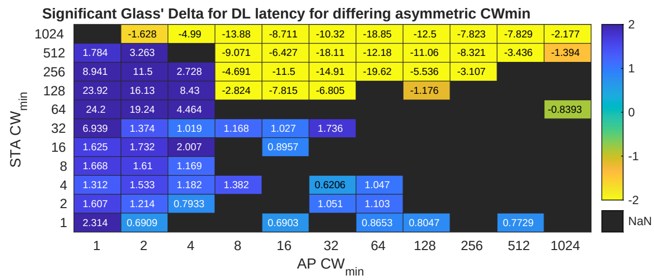

(d) Significant effect sizes on DL latency for experiment 6

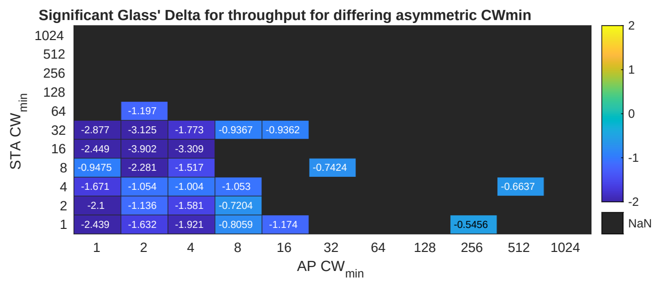

(e) Significant effect sizes on throughput for experiment 6

Figure 4.6: Significant Glass' Δ for experiment 6. Yellow shaded values indicate improvement for the "reset" system over the "continue" system, while blue shaded values indicate regression. NaN values, colored black, indicate insignificant results.

4.7. Experiment 7: Fixed rate traffic with proportional RU allocation, optimal asymmetric CW$_{min}$ and varying traffic generation rate

This experiment was performed as a sort of miniature design space exploration. The intention was to see what would become of the effect of a backoff reset in an ”ideal situation” given the constraints of the parameter set that could be manipulated. To create this situation, the best performing asymmetric CW$_{min}$ pair from Experiment 6 was taken in terms of average performance improvement across all three output variables. This optimal parameter set was determined to be AP CW$_{min}$$= 32$ and STA CW$_{min}$$= 256$, due to its consistently huge performance increases and lack of drawbacks. Then, proportional RU allocation was enabled and the data generation rate was varied from 100 Hz to 7000 Hz.

4.7.1. Results

The sub-figures of Figure 4.7 show the outputs of this experiment. Figure 4.7a and Figure 4.7b show the UL and DL collision probability, Figure 4.7c and Figure 4.7d show the UL and DL latency, and Figure 4.7e shows the overall throughput of the network.

4.7.2. Performance evaluation

Table 4.5 shows these effect sizes for each of the output variables, given that the effect is significant, as explained in Section 3.5.1. From Table 4.5, we see very sporadic, large to very large improvements in throughput, with a single large reduction at 3300 Hz. Otherwise, throughput is similar between the two systems. More interestingly, both the latencies and collision probabilities show only a single large increase for 200 Hz, with all further effects being mostly huge decreases, with a couple of large to very large ones for the lower sample rates.

Taking a look at the full picture, then, it becomes clear that the ”reset” system can outperform the ”continue” system quite spectacularly, if given the right parameter set. While this will, in general, not result in throughput improvements as no saturation is achieved and thus all available data will be transmitted by both systems, the system will experience much lower latencies and collision probabilities. Comparing these results to those found in Experiment 4, it can be seen that where the ”reset” system under-performed there, it excels against ”continue” system once the right parameter set is introduced.

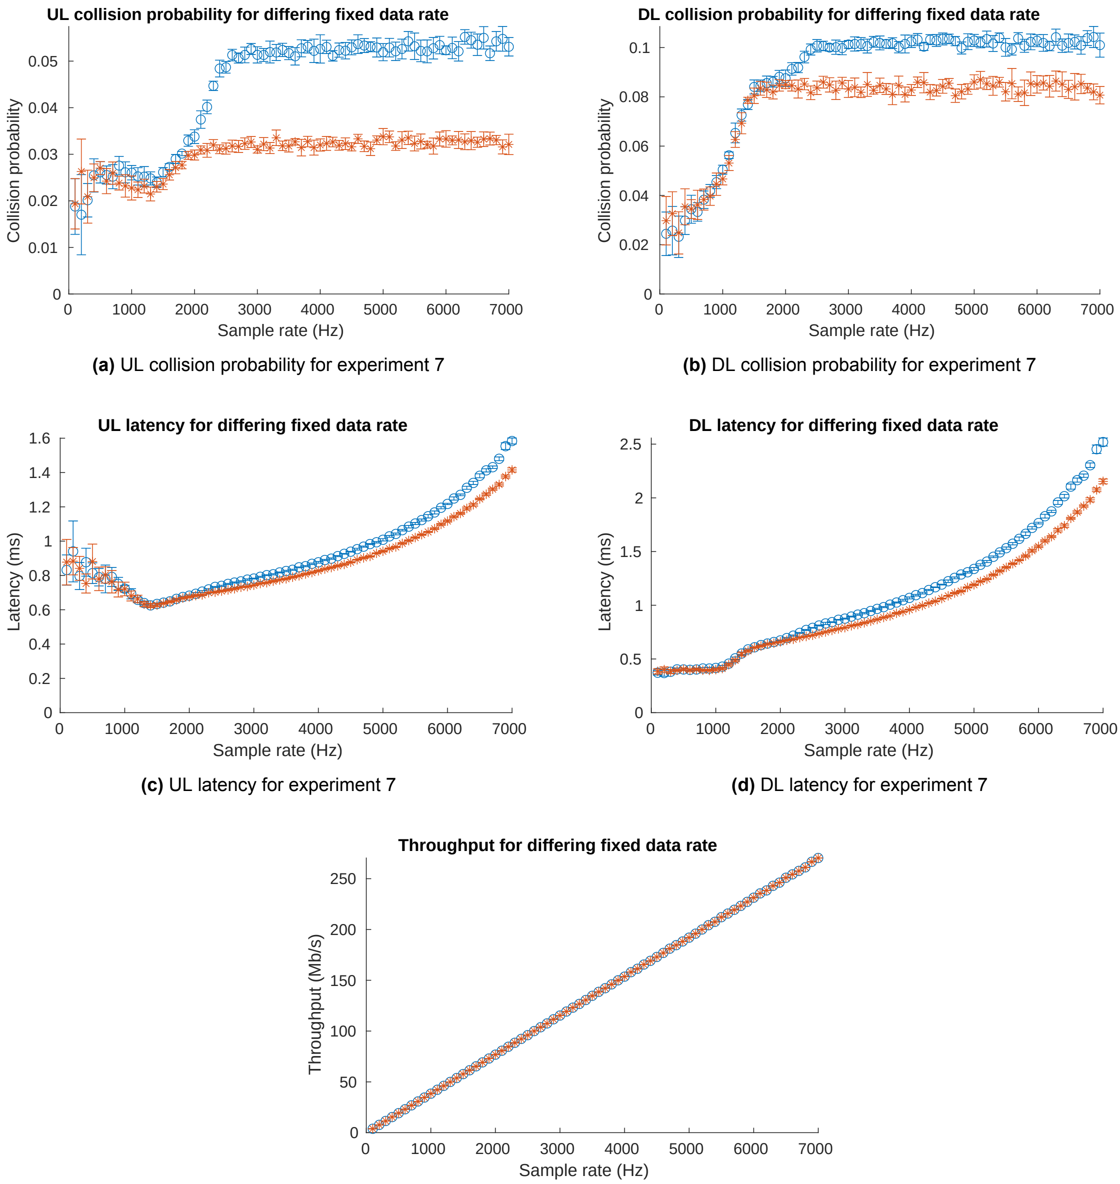

(a) UL collision probability for experiment 7

(b) DL collision probability for experiment 7

(c) UL latency for experiment 7

(d) DL latency for experiment 7

(e) Throughput for experiment 7

Figure 4.7: Results for experiment 7. ”Continue” system data is plotted in blue circles, ”reset” system data is plotted in orange asterisks.

Sample rate (Hz)	Throughput	DL latency	UL latency	DL coll. prob.	UL coll. prob.
100	-	-	-	-	-
200	-	0.87	-	-	0.91
300	-	-	-	-	-
400	-	-	-	-	-
500	-	-	-	-	-
600	-	-	-	-	-
700	-	-	-	-	-
800	-	-1.41	-	-	-1.75
900	-	-1.59	-	-	-0.83
1000	-	-0.97	-	-	-
1100	-	-1.15	-	-0.72	-1.18
1200	-	-	-	-	-
1300	-	-	-	-	-1.56
1400	-	-	-	-	-
1500	-	-1.06	-	-	-1.45
1600	-	-	-0.89	-	-1.03
1700	-	-1.35	-	-	-0.96
1800	-	-	-	-1.03	-2.12
1900	-	-4.26	-	-	-2.56
2000	-	-3.42	-1.32	-	-2.53
2100	-	-3.65	-1.66	-2.11	-5.13
2200	-	-7.80	-3.27	-1.63	-4.28
2300	-	-14.59	-7.53	-3.26	-7.50
2400	-	-16.16	-6.31	-4.31	-12.44
2500	-	-21.16	-6.51	-4.29	-9.86
2600	-	-14.45	-7.82	-4.17	-10.66
2700	-	-16.80	-8.01	-3.57	-10.09
2800	1.30	-16.25	-11.67	-3.22	-10.55
2900	-	-18.12	-7.92	-4.17	-11.54
3000	-	-23.17	-13.15	-3.63	-17.76
3100	-	-16.18	-10.34	-3.18	-9.34
3200	-	-16.18	-10.06	-6.22	-11.84
3300	-0.95	-13.11	-6.70	-2.21	-7.49
3400	-	-33.77	-13.45	-8.56	-13.90
3500	-	-18.28	-10.37	-4.06	-10.74
3600	-	-17.09	-8.32	-5.67	-18.54
3700	-	-27.31	-21.85	-3.36	-8.22
3800	-	-9.78	-9.94	-2.81	-12.44
3900	0.80	-26.21	-12.43	-4.64	-8.88
4000	-	-22.93	-11.38	-6.15	-10.72
4100	-	-22.77	-15.72	-5.55	-9.54
4200	-	-18.82	-10.90	-2.43	-12.32
4300	-	-22.98	-11.04	-4.76	-13.28
4400	0.79	-18.57	-9.97	-6.24	-10.86
4500	-	-24.28	-12.79	-4.68	-16.78
4600	-	-26.60	-20.31	-6.30	-14.26
4700	-	-14.50	-11.49	-5.59	-12.88
4800	-	-15.90	-13.85	-4.17	-10.32
4900	-	-18.97	-16.00	-6.14	-14.57
5000	-	-15.46	-10.47	-3.92	-7.12
5100	-	-17.48	-16.15	-3.06	-10.54
5200	-	-21.63	-18.69	-5.33	-15.06
5300	-	-15.48	-11.82	-3.06	-7.44
5400	-	-14.96	-13.51	-6.25	-9.92
5500	-	-16.73	-14.06	-3.29	-9.02
5600	-	-20.42	-15.34	-1.55	-8.27
5700	-	-27.30	-21.50	-5.24	-18.99
5800	-	-14.08	-12.27	-2.53	-8.32
5900	-	-13.48	-11.82	-2.43	-9.17
6000	-	-13.04	-10.89	-3.42	-7.35
6100	1.25	-15.61	-13.31	-3.87	-14.18
6200	-	-17.04	-13.85	-2.54	-7.14
6300	-	-11.49	-10.28	-3.00	-9.07
6400	-	-12.24	-10.44	-2.13	-7.59
6500	-	-21.63	-19.10	-3.51	-13.21
6600	-	-15.68	-13.71	-2.32	-9.44
6700	-	-17.16	-13.89	-2.73	-9.10
6800	-	-11.72	-11.68	-7.29	-14.50
6900	-	-15.41	-13.65	-4.75	-10.51
7000	-	-11.07	-9.35	-3.94	-6.56

Table 4.5: Significant Glass' Δ for experiment 7. Bold values are controlled variable values. Positive values indicate improvement for the "reset" system over the "continue" system for throughput, while the same is indicated by negative values for the other measurements. Dashes indicate insignificant results.

Number of STAs	Throughput	DL latency	UL latency	DL coll. prob.	UL coll. prob.
1	-	-5.91	56.25	4.67	-3.38
2	-	-	17.80	-1.29	-2.21
3	-	-	-	-2.13	-
4	-	-	-	-1.41	0.89
5	-	-	-	-	2.18
6	-	1.09	-	1.79	3.35
7	-	-	-	4.65	3.96
8	-	-	-	4.33	6.01

Table 4.6: Significant Glass' Δ for experiment 8. Bold values are controlled variable values. Positive values indicate improvement for the "reset" system over the "continue" system for throughput, while the same is indicated by negative values for the other measurements. Dashes indicate insignificant results.

4.8. Experiment 8: Bounded random rate traffic with varying number of stations

This experiment attempted to capture the same effects as Experiment 1, but now under slightly more realistic data generation constraints, in that the data generation rate was no longer fixed, but was uniformly random for each generated packet within the range of $[4000, 6000]$ Hz. The number of non-AP STAs was once again controlled from 1 to 8.

4.8.1. Results

The sub-figures of Figure 4.8 show the outputs of this experiment. Figure 4.8a and Figure 4.8b show the UL and DL collision probability, Figure 4.8c and Figure 4.8d show the UL and DL latency, and Figure 4.8e shows the overall throughput of the network.

4.8.2. Performance evaluation

Table 4.6 shows these effect sizes for each of the output variables, given that the effect is significant, as explained in Section 3.5.1. In Table 4.6, no significant effects were noted for throughput. For the DL latency, a huge reduction in latency is seen for a single STA, while for 6 STAs a large increase is noted. Meanwhile, the UL latency shows huge increases in latency for 1 to 2 STAs. The DL collision probabilities experience very large to huge improvements only for networks with 2 to 4 STAs. In all other cases, a very large to huge increase is observed. Finally, the UL collision probability improves by a huge amount for networks with 1 to 2 STAs, but otherwise sees large to huge increases across the board.

In comparison to the effects of Experiment 1, tabulated in Table 4.1, we see an almost identical trend in effect sizes. Thus, it seems that the random traffic generation rate does not influence the difference between the "continue" system and the "reset" system when varying the number of STAs in the network.

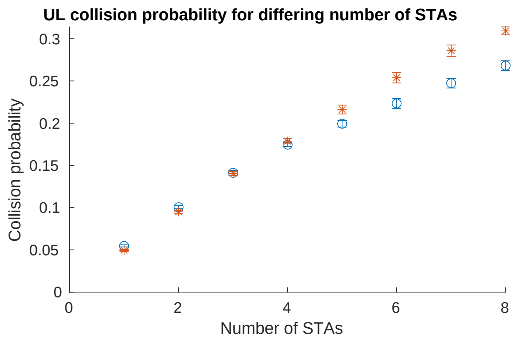

(a) UL collision probability for experiment 8

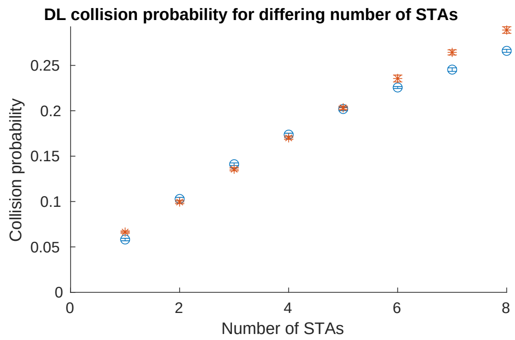

(b) DL collision probability for experiment 8

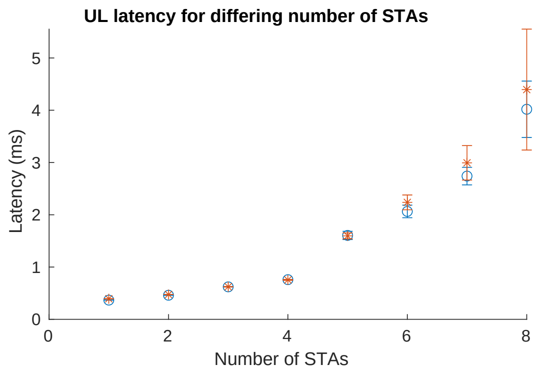

(c) UL latency for experiment 8

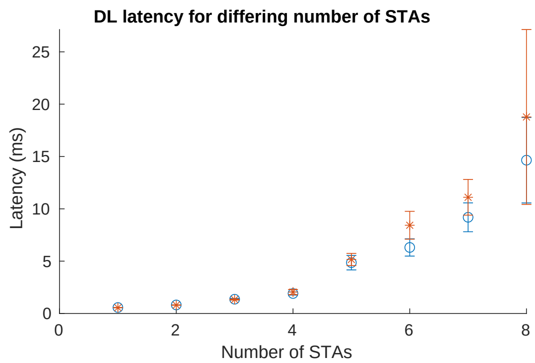

(d) DL latency for experiment 8

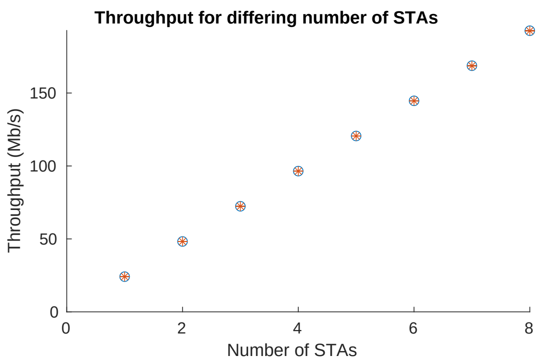

(e) Throughput for experiment 8

Figure 4.8: Results for experiment 8. ”Continue” system data is plotted in blue circles, ”reset” system data is plotted in orange asterisks.

CW$_{min}$	Throughput	DL latency	UL latency	DL coll. prob.	UL coll. prob.
1	-2.30	1.58	3.98	-1.11	21.46
2	-1.04	1.22	2.32	1.47	17.17
4	-1.02	-	1.17	3.26	9.85
8	-	-	-	3.34	11.67
16	-0.79	0.77	0.73	11.59	11.24
32	-	1.45	-	5.49	5.74
64	-	-	-	-	-
128	-	-	-	-1.94	-1.55
256	-	-2.79	-	-2.02	-1.45
512	-0.95	-2.90	3.67	-1.64	-1.25
1024	-	-3.52	7.71	-1.29	-

Table 4.7: Significant Glass' Δ for experiment 11. Bold values are controlled variable values. Positive values indicate improvement for the ”reset” system over the ”continue” system for throughput, while the same is indicated by negative values for the other measurements. Dashes indicate insignificant results.

4.9. Experiment 11: Bounded random rate traffic with proportional RU allocation and varying symmetric CW$_{min}$

This experiment aimed to find out whether the combination of a bounded random data rate and proportional RU allocation would result in significant deviations from the trends observed in Experiment 5. As such, CW$_{min}$ was varied from 1 to 1024 symmetrically, under a uniform random data generation rate of $[4000, 6000]$ Hz.

4.9.1. Results

The sub-figures of Figure 4.9 show the outputs of this experiment. Figure 4.9a and Figure 4.9b show the UL and DL collision probability, Figure 4.9c and Figure 4.9d show the UL and DL latency, and Figure 4.9e shows the overall throughput of the network.

4.9.2. Performance evaluation

Table 4.7 shows these effect sizes for each of the output variables, given that the effect is significant, as explained in Section 3.5.1. From Table 4.7, one can see that the throughput experiences medium to huge decreases when compared to the ”continue” system. Latency shows similarly reduced performance, with a notable exception for DL latency when CW$_{min}$$\geq 256$, where a huge improvement is measured. For CW$_{min}$$\leq 32$, the collision probability experiences a very large to huge increase. Much like in previous experiments, the one notable exception is for CW$_{min}$ is 1, where a large decrease in collision probability is noted. These decreases repeat once CW$_{min}$$\geq 128$, with very large to huge effects.

Comparing these results to Experiment 5 in Table 4.4, one can see that the effect trends are again nearly identical. This leads to the conclusion that the combination of proportional RU allocation and a bounded random data generation rate does not meaningfully impact the overall effects that the ”reset” system experiences in comparison to the ”continue” system.

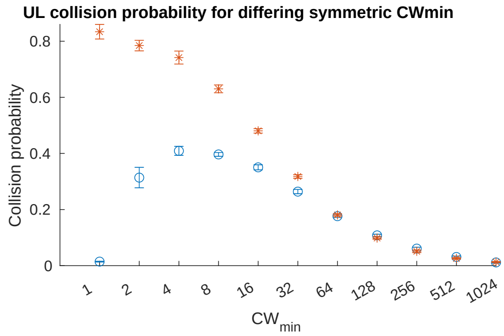

(a) UL collision probability for experiment 11

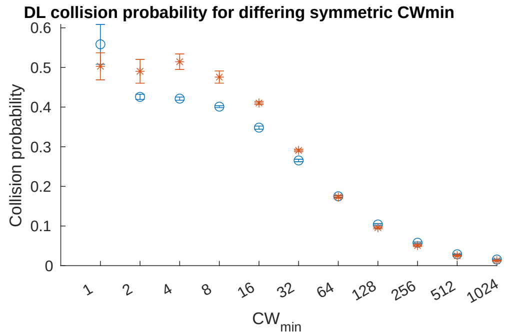

(b) DL collision probability for experiment 11

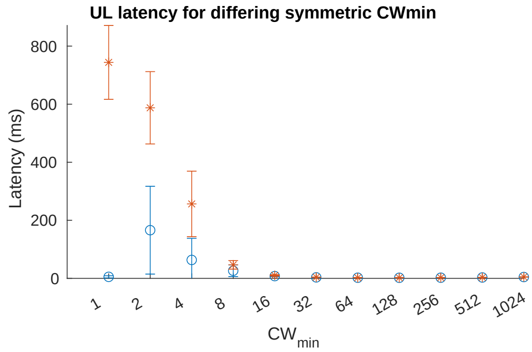

(c) UL latency for experiment 11

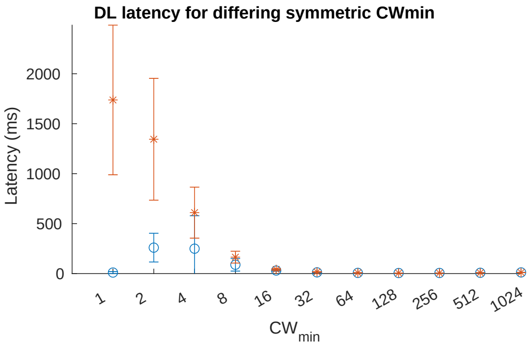

(d) DL latency for experiment 11

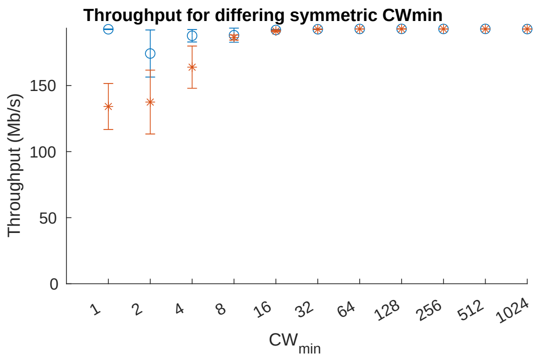

(e) Throughput for experiment 11

Figure 4.9: Results for experiment 11. ”Continue” system data is plotted in blue circles, ”reset” system data is plotted in orange asterisks.

5
Discussion

This chapter first summarizes the findings from Chapter 4 into one cohesive whole and then reflects on the limitations that are attached to this study.

5.1. Overall results

The aim of this paper was to answer the problem statement:

> **”What is the effect on performance metrics of resetting a station’s backoff stage and counter after an UL-OFDMA transmission?”**

For this, we formulated the hypothesis that **”Resetting a station’s backoff stage and counter after an UL-OFDMA transmission results in an increase in collision probability, a decrease in throughput and an increase in latency for all but the lowest contention networks.”** Using the results gathered from the experiments, this problem statement can now be answered and the hypothesis be accepted or rejected.

To not draw the waiting out any longer, first the main conclusion of this report will be formulated based on the results from the experiments: **the ”reset” system outperforms the ”continue” system in certain situations on (a combination of) latency and collision probability, but is generally strictly worse on all five key metrics in a typical, flexible Wi-Fi deployment, where data rates and the amount of connected devices can change rapidly. We thus accept the research hypothesis.** Noting that, it is now time to discuss in what specific situations the ”reset” system improves over the ”continue” system.

Firstly, from Experiments 1 and 8 it can be concluded that the ”reset” system will outperform the ”continue” system in at least one form of collision probability as long as the number of connected STAs is less than or equal to four. Secondly, from Experiments 2, 3, 5, 6, and 9 through 12, it can be seen that modifying the CW$_{min}$, either symmetrically or asymmetrically, will almost never result in the ”reset” system outperforming the ”continue” system on throughput. However, once the STA CW$_{min}$ meets or exceeds 64, the ”reset” system will perform better than the ”continue” system on both UL and DL latency and collision probability, as long as the AP CW$_{min}$ is at least equal to 8. Thirdly, Experiment 4 shows that the under-performance or over-performance of the ”reset” system versus the ”continue” system is invariant under changes in data rate. Fourthly, Experiments 5, 6, 11 and 12 demonstrate that the same is true for proportional RU allocation. Fifthly, Experiments 8 through 12 reveal that these results are robust to bounded random data generation rates. Finally, Experiment 7 shows that the ”reset” system can outperform the ”continue” system if the right values are chosen for STA and AP CW$_{min}$. It should be noted, however, that no physical layer specification in either the 802.11 standard [7] or the 802.11ax amendment [8] specifies a CW$_{min}$ greater than 32. Thus, this increase in performance will only manifest in a standard-conforming Wi-Fi deployment once the AP detects a high-contention scenario and modifies the CW$_{min}$ of the system.

In summary, it was found that the ”realism” controlled variables of proportional RU allocation and bounded-random data generation rate did indeed not affect the performance of the system, while the other controlled variables all did have significant effects, except for the data generation rate.

5.2. Limitations

While the results of the performed experiments are significant, it is important to note that the research detailed in this report is preliminary and, thus, limited. This subsection aims to explore the main limitations of this research.

5.2.1. Absence of real-world validation

Unfortunately, the simulator used was not validated against a real-world network under similar conditions. This is not a problem unique to this simulator, as many other simulators in the literature are similarly unverified, as Wi-Fi hardware only becomes available once a new standard has already been ratified and generally is not open to modification. Instead, the accepted way of working is to validate a new simulator either against a mathematical model or against another simulator. This leaves some uncertainty on whether the results that the simulator produces hold up in a physical system and thus whether the conclusions drawn in this report generalize well to practical situations. The author of this report has seen preliminary work from the creators of the Wi-Fi 6 MAC Simulator [32] on a mathematical model that validates the working of the simulator. While, these early results seem favorable, they necessarily only provide validation in a theoretical sense.

An obvious solution to this problem is to create a comparable real-world Wi-Fi 6 test-bed with matching capabilities to the simulator. Indeed, the original intention of the author with his thesis was to create such a test-bed for that exact purpose. However, it was found that the scope of such a project was much larger than can be expected of a Master's thesis and as such the project was abandoned.

Another solution, that should take much less time and resources, is to validate the outcomes of the simulator against another simulator that has itself been validated against the real-world. While the author has not been able to find any studies performing a holistic validation, a likely candidate for such a simulator would be ns-3, as both piece-wise mathematical [31, 35] and experimental [36, 37] validation studies are known to exist for the IEEE 802.11 family of protocols.

5.2.2. Lack of mathematical modelling

As mentioned previously, work is underway on a mathematical framework validating the simulator, however providing a mathematical model on the effects the backoff reset was deemed out of scope for this project. As such, the understanding that can be gained from this simulator study on such effects is limited to effectively a ”black box” approach: both the inputs and outputs are known, but the mechanism driving this translation is still opaque.

To alleviate this limitation, a thorough mathematical analysis, possibly based on the mathematical model by Bellalta and Kosek-Szott [9], may be performed.

5.2.3. Limited design space exploration

As a consequence of the desire to keep the scope of research manageable, not all factors could be taken into account or indeed implemented in the time allotted for the project. For instance, only the effects of a backoff reset for the basic access mechanism were explored, instead of including a combination of RTS-CTS and basic access mechanisms. A second parameter that was not explored is the effect of backoff resets when using the full breadth of EDCA ACs. Thirdly, this study had to limit itself to the study of small ($\leq$8 STAs) networks, as the simulator did not support bigger networks. Similarly, the simulator was limited to non-saturated network conditions. Finally, a Wi-Fi network in the real world has to contend with quite possibly many overlapping neighbours whose data traffic could interfere with the effects of a backoff reset.

Further research probing these untapped parameters could prove interesting and would in any case further the scientific understanding of the effects of backoff resets on Wi-Fi 6 networks.

6 Conclusion

The main goal of this thesis was to identify an oversight in the scientific literature on IEEE 802.11ax and attempt to quantify the prevalence and effect of this omission. To allow readers new to the Wi-Fi space to be able to follow along, a brief introduction was given to the concepts of DCF, EDCA and OFDMA. Afterwards, the oversight was defined as a missing return to default for both CW and backoff counter after a successful UL-OFDMA transmission and named the ”backoff reset”. The problem statement was then based on this effect and read as follows:

> **”What is the effect on performance metrics of resetting a station’s backoff stage and counter after an UL-OFDMA transmission?”**

A short symbolic analysis of the effects of the backoff reset lead to the hypothesis that ”Resetting a station’s backoff stage and counter after an UL-OFDMA transmission results in an increase in collision probability, a decrease in throughput and an increase in latency for all but the lowest contention networks.”

While the main investigation of this paper was into the numerical effects of including the backoff reset, a preliminary literature survey was performed to find the prevalence of this defect in the scientific Wi-Fi corpus. This census analyzed fifteen prominent Wi-Fi 6 papers to find mentions or allusions to the inclusion of the backoff reset in their results. Out of these fifteen papers, only five were found to likely include this phenomenon.

Then, to analyze the effects of the backoff reset on the key metrics of throughput, UL and DL latency and UL and DL collision probability, a series of experiments were performed in a slightly modified version of the Wi-Fi 6 MAC Simulator by Gokhale and Kroep [32]. These modifications allowed for bounded random traffic generation rates and proportional RU allocation. In these experiments, the system with a backoff reset, nicknamed ”reset”, was pitched against a system without the reset, nicknamed ”continue”, while varying the input variables of (bounded random) data rate, CW for both AP and STAs, network size and whether proportional RU allocation was enabled. The outcomes of these experiments were then tested for significance with a two-tailed, paired t-test with $\alpha = 0.05$ and, if significant, the Glass’ Δ was calculated to determine the effect size, with the ”continue” system as the control group.

From these effect sizes, the following conclusion was drawn: ”the ”reset” system outperforms the ”continue” system in certain situations on (a combination of) latency and collision probability, but is generally strictly worse on all five key metrics in a typical, flexible Wi-Fi deployment, where data rates and the amount of connected devices can change rapidly. We thus accept the research hypothesis.” Finally, an overview of the situations in which the ”reset” system can outperform the ”continue” system was given and the limitations of this study were discussed.

6.1. Future work

Regardless of the outcomes of this study, the main point of future work that should be engaged in is an inclusion of the backoff reset in all current and future Wi-Fi 6 simulators, as including the backoff reset is a more accurate depiction of the real-world behaviour of Wi-Fi 6 networks. After all, one can not be certain that parameters outside of the scope of this study will not interact strongly in some way with the absence of a backoff reset, resulting in incorrect simulations.

Apart from this inclusion, the suggestions from the previous discussion of the limitations of this research in Section 5.2 provide a good basis for future endeavours. To reiterate and summarize, those were firstly to perform real-world validation of the Wi-Fi 6 MAC Simulator by Gokhale and Kroep [32], either through experimental replication on a real-world Wi-Fi 6 test-bed, or through validation against a simulator that has undergone such validation. Secondly, mathematical modelling could facilitate a much more in-depth understanding of the effects of the backoff reset than the current ”black box” approach can provide. Finally, the design space exploration of this study could be vastly expanded.

References

[1] Cisco. *Cisco Annual Internet Report (2018–2023) White Paper*. 2020. URL: https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html (visited on 12/30/2022).

[2] Xianjun Jiao et al. "openwifi: a free and open-source IEEE802. 11 SDR implementation on SoC". In: *2020 IEEE 91st Vehicular Technology Conference (VTC2020-Spring)*. IEEE. 2020, pp. 1–2.

[3] ns-3. *ns-3 | a discrete-event network simulator for internet systems*. 2022. URL: https://www.nsnam.org/ (visited on 12/30/2022).

[4] OMNeT++. *OMNeT++ Discrete Event Simulator*. 2022. URL: https://omnetpp.org/ (visited on 12/30/2022).

[5] Chris Avants. *802.11 Wireless (Wi-Fi) 30 Minute Crash Course*. 2020. URL: https://www.udemy.com/course/learn-wifi-in-30-minutes/ (visited on 12/20/2022).

[6] Evgeny Khorov et al. "A Tutorial on IEEE 802.11ax High Efficiency WLANs". In: *IEEE Communications Surveys & Tutorials* PP (Sept. 2018), pp. 197–216. DOI: 10.1109/COMST.2018.2871099.

[7] IEEE Std 802.11-2020. "IEEE Standard for Information Technology–Telecommunications and Information Exchange between Systems - Local and Metropolitan Area Networks–Specific Requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications". In: *IEEE Std 802.11-2020 (Revision of IEEE Std 802.11-2016)* (2021), pp. 1–4379. DOI: 10.1109/IEEESTD.2021.9363693.

[8] IEEE Std 802.11ax. "IEEE Standard for Information Technology–Telecommunications and Information Exchange between Systems Local and Metropolitan Area Networks–Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 1: Enhancements for High-Efficiency WLAN". In: *IEEE Std 802.11ax-2021 (Amendment to IEEE Std 802.11-2020)* (2021), pp. 1–767. DOI: 10.1109/IEEESTD.2021.9442429.

[9] Boris Bellalta and Katarzyna Kosek-Szott. "AP-initiated multi-user transmissions in IEEE 802.11ax WLANs". In: *Ad Hoc Networks* 85 (2019), pp. 145–159. ISSN: 1570-8705. DOI: https://doi.org/10.1016/j.adhoc.2018.10.021. URL: https://www.sciencedirect.com/science/article/pii/S1570870518307613.

[10] G. Bianchi. "Performance analysis of the IEEE 802.11 distributed coordination function". In: *IEEE Journal on Selected Areas in Communications* 18.3 (2000), pp. 535–547. DOI: 10.1109/49.840210.

[11] L. Kleinrock and F. Tobagi. "Packet Switching in Radio Channels: Part I - Carrier Sense Multiple-Access Modes and Their Throughput-Delay Characteristics". In: *IEEE Transactions on Communications* 23.12 (1975), pp. 1400–1416. DOI: 10.1109/TCOM.1975.1092768.

[12] Pui King Wong, Dongjie Yin, and Tony T Lee. "Analysis of non-persistent CSMA protocols with exponential backoff scheduling". In: *IEEE Transactions on Communications* 59.8 (2011), pp. 2206–2214.

[13] Jerome Galtier. "Analysis of the Slotted Non-persistent CSMA Protocol with Poissonian Packet Size Using a Semi-Markov Graph Representation". In: *2006 International Conference on Transparent Optical Networks*. Vol. 3. 2006, pp. 258–262. DOI: 10.1109/ICTON.2006.248447.

[14] Jonathan Goodman et al. "Stability of Binary Exponential Backoff". In: *J. ACM* 35.3 (June 1988), pp. 579–602. ISSN: 0004-5411. DOI: 10.1145/44483.44488. URL: https://doi-org.tudelft.idm.oclc.org/10.1145/44483.44488.

[15] D. Aldous. "Ultimate instability of exponential back-off protocol for acknowledgment-based transmission control of random access communication channels". In: *IEEE Transactions on Information Theory* 33.2 (1987), pp. 219–223. DOI: 10.1109/TIT.1987.1057295.

[16] Jie Hui and M. Devetsikiotis. “A unified model for the performance analysis of IEEE 802.11e EDCA”. In: *IEEE Transactions on Communications* 53.9 (2005), pp. 1498–1510. DOI: 10.1109/TCOMM.2005.855013.

[17] Leonardo Lanante et al. “Performance analysis of the 802.11 ax UL OFDMA random access protocol in dense networks”. In: *2017 IEEE international conference on communications (ICC)*. IEEE. 2017, pp. 1–6.

[18] Stefano Avallone et al. “Will OFDMA improve the performance of 802.11 WiFi networks?” In: *IEEE Wireless Communications* 28.3 (2021), pp. 100–107.

[19] Yousri Daldoul, Djamal-Eddine Meddour, and Adlen Ksentini. “Performance evaluation of OFDMA and MU-MIMO in 802.11 ax networks”. In: *Computer Networks* 182 (2020), p. 107477.

[20] Gaurang Naik, Sudeep Bhattarai, and Jung-Min Park. “Performance analysis of uplink multi-user OFDMA in IEEE 802.11 ax”. In: *2018 IEEE international conference on communications (ICC)*. IEEE. 2018, pp. 1–6.

[21] Dmitry Bankov et al. “OFDMA uplink scheduling in IEEE 802.11 ax networks”. In: *2018 IEEE international conference on communications (ICC)*. IEEE. 2018, pp. 1–6.

[22] Konstantinos Dovelos and Boris Bellalta. “Optimal resource allocation in IEEE 802.11 ax uplink OFDMA with scheduled access”. In: *arXiv preprint arXiv:1811.00957* (2018).

[23] Katarzyna Kosek-Szott and Krzysztof Domino. “An efficient backoff procedure for IEEE 802.11 ax uplink OFDMA-based random access”. In: *IEEE Access* 10 (2022), pp. 8855–8863.

[24] Mehmet Karaca et al. “Resource management for OFDMA based next generation 802.11 WLANs”. In: *2016 9th IFIP Wireless and Mobile Networking Conference (WMNC)*. IEEE. 2016, pp. 57–64.

[25] Evgeny Khorov, Vyacheslav Loginov, and Andrey Lyakhov. “Several EDCA parameter sets for improving channel access in IEEE 802.11 ax networks”. In: *2016 International Symposium on Wireless Communication Systems (ISWCS)*. IEEE. 2016, pp. 419–423.

[26] Yousri Daldoul, Djamal-Eddine Meddour, and Adlen Ksentini. *802.11ax lightsim*. 2020. URL: https://github.com/yousri-daldoul/802.11ax-lightsim (visited on 04/14/2023).

[27] Kyu-haeng Lee. “Performance analysis of the IEEE 802.11 ax MAC protocol for heterogeneous Wi-Fi networks in non-saturated conditions”. In: *Sensors* 19.7 (2019), p. 1540.

[28] Evgeny Avdotin et al. “Enabling massive real-time applications in IEEE 802.11 be networks”. In: *2019 IEEE 30th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*. IEEE. 2019, pp. 1–6.

[29] Sharan Naribole, Wook Bong Lee, and Ashok Ranganath. “Impact of MU EDCA channel access on IEEE 802.11 ax WLANs”. In: *2019 IEEE 90th Vehicular Technology Conference (VTC2019-Fall)*. IEEE. 2019, pp. 1–5.

[30] Sudeep Bhattarai, Gaurang Naik, and Jung-Min Jerry Park. “Uplink resource allocation in IEEE 802.11 ax”. In: *ICC 2019-2019 IEEE international conference on communications (ICC)*. IEEE. 2019, pp. 1–6.

[31] Davide Magrin et al. “Validation of the Ns-3 802.11ax OFDMA Implementation”. In: *Proceedings of the Workshop on Ns-3*. WNS3 ’21. Virtual Event, USA: Association for Computing Machinery, 2021, pp. 1–8. ISBN: 9781450390347. DOI: 10.1145/3460797.3460798. URL: https://doi.org/10.1145/3460797.3460798.

[32] Vineet Gokhale and Kees Kroep. *Wi-Fi 6 MAC simulator*. 2022. URL: https://github.com/VinGok/Tactile-WiFi (visited on 12/30/2022).

[33] Vineet Gokhale et al. “Toward Enabling High-Five Over WiFi: A Tactile Internet Paradigm”. In: *IEEE Communications Magazine* 59.12 (2021), pp. 90–96. DOI: 10.1109/MCOM.007.2100271.

[34] Shlomo S Sawilowsky. “New effect size rules of thumb”. In: *Journal of modern applied statistical methods* 8.2 (2009), p. 26.

[35] Le Tian et al. “Implementation and Validation of an IEEE 802.11ah Module for Ns-3”. In: *Proceedings of the 2016 Workshop on Ns-3*. WNS3 ’16. Seattle, WA, USA: Association for Computing Machinery, 2016, pp. 49–56. ISBN: 9781450342162. DOI: 10.1145/2915371.2915372. URL: https://doi-org.tudelft.idm.oclc.org/10.1145/2915371.2915372.

[36] Guangyu Pei and Tom Henderson. “Validation of ns-3 802.11 b PHY model”. In: *Online: http://www. nsnam. org/~ pei/80211b. pdf* (2009).

[37] Paul Fuxjaeger and Stefan Ruehrup. “Validation of the NS-3 Interference Model for IEEE802.11 Networks”. In: *2015 8th IFIP Wireless and Mobile Networking Conference (WMNC)*. 2015, pp. 216–222. DOI: 10.1109/WMNC.2015.40.

A

Auxiliary experiments

This appendix contains all experiments whose results did not provide new insights. Each experiment details first its intended purpose and non-default inputs, then displays the outputs and the results from statistical analysis.

CW$_{min}$	Throughput	DL latency	UL latency	DL coll. prob.	UL coll. prob.
1	-2.70	1.38	2.42	-3.87	26.65
2	-2.03	1.42	2.26	1.22	13.45
4	-0.87	-	1.92	2.47	11.24
8	-1.47	0.79	0.72	2.84	9.32
16	-0.64	-	-	9.27	10.26
32	-	-	-	4.16	4.72
64	-	-	-	-1.96	-
128	-	-	-	-3.82	-
256	-1.16	-1.77	-	-2.12	-0.64
512	-	-2.46	4.94	-1.99	-
1024	-	-2.95	7.27	-1.91	-1.14

Table A.1: Significant Glass' Δ for experiment 9. Bold values are controlled variable values. Positive values indicate improvement for the ”reset” system over the ”continue” system for throughput, while the same is indicated by negative values for the other measurements. Dashes indicate insignificant results.

A.1. Experiment 9: Bounded random rate traffic with varying symmetric CW$_{min}$

This experiment continued the exploration of the effect of a bounded random traffic rate on the results of previous experiments. In this case, Experiment 2 was replicated with a uniform random data generation rate in the range of $[4000, 6000]$ Hz, varying the AP and STA CW$_{min}$ symmetrically from 1 to 1024.

A.1.1. Results

The sub-figures of Figure A.1 show the outputs of this experiment. Figure A.1a and Figure A.1b show the UL and DL collision probability, Figure A.1c and Figure A.1d show the UL and DL latency, and Figure A.1e shows the overall throughput of the network.

A.1.2. Performance evaluation

Table A.1 shows these effect sizes for each of the output variables, given that the effect is significant, as explained in Section 3.5.1. Table A.1 shows that the effect of varying the CW$_{min}$ under a bounded random traffic generation rate on throughput is a large to huge reduction. The UL and DL latencies experience large to huge increases, except for large CW$_{min}$ values, where the DL latency improves by a very large to huge effect. In general, the collision probabilities increase by a very large to huge amount for CW$_{min}\leq 32$ and decrease by medium to huge amounts beyond this point. One notable exception is the DL collision probability, which decreases for CW$_{min}$ is 1.

Once again, the conclusion can be drawn that the random generation rate does not influence the effect of the ”reset” system in relation to the ”continue” system, as the results of Experiment 2 show the same trends.

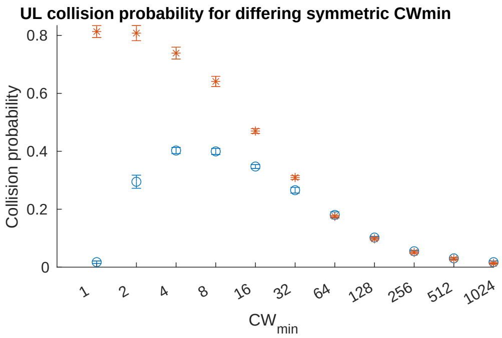

(a) UL collision probability for experiment 9

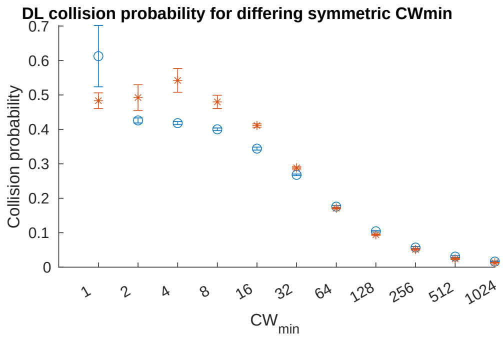

(b) DL collision probability for experiment 9

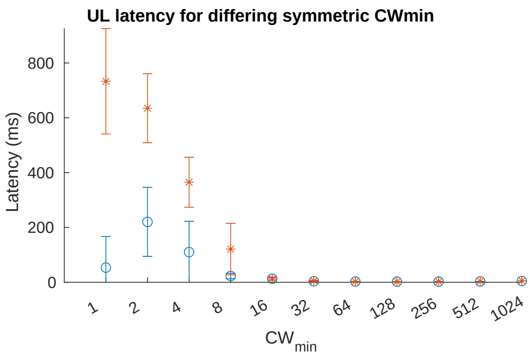

(c) UL latency for experiment 9

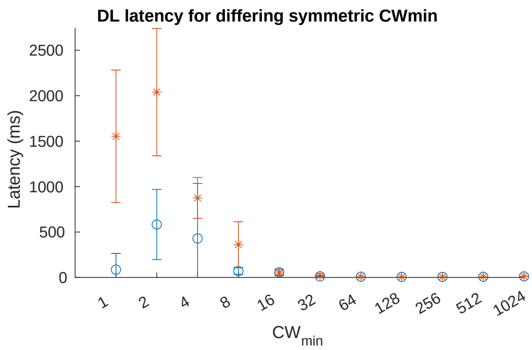

(d) DL latency for experiment 9

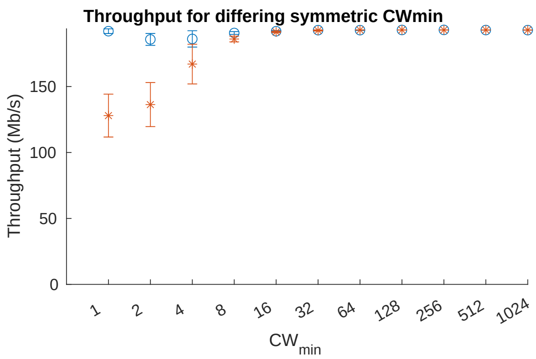

(e) Throughput for experiment 9

Figure A.1: Results for experiment 9. ”Continue” data is plotted in blue circles, ”reset” system data is plotted in orange asterisks.

A.2. Experiment 10: Bounded random rate traffic with varying asymmetric CW$_{min}$

This experiment once again aimed to replicate a previous experiment, this time Experiment 3, under more realistic data generation conditions. As such, CW$_{min}$ was varied from 1 to 1024 independently for both STAs and AP under a uniform random data generation rate of $[4000, 6000]$ Hz.

A.2.1. Results and performance evaluation

The significant effect sizes for this experiment are plotted in the heatmaps of Figure A.2, as explained in Section 3.5.2. Figure A.2a and Figure A.2b show the UL and DL collision probability, Figure A.2c and Figure A.2d show the UL and DL latency, and Figure A.2e shows the overall throughput of the network.

Much like Experiments 8 and 9, this experiment sees the trends occur on the effect sizes for each key metric as its non-random counterpart in Experiment 3. For a full analysis of these trends, refer to Section 4.3.1. This gives further credence to the notion that the effects observed in the non-random experiments generalize to more realistic data generation scenarios.

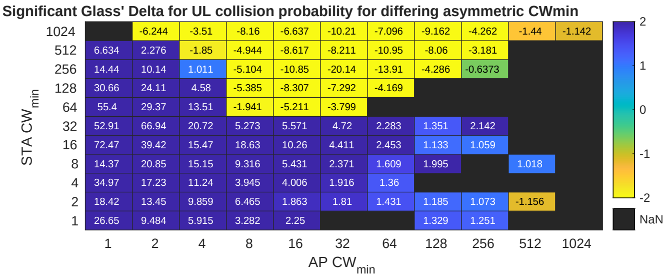

(a) Significant effect sizes on UL collision probability for experiment 10

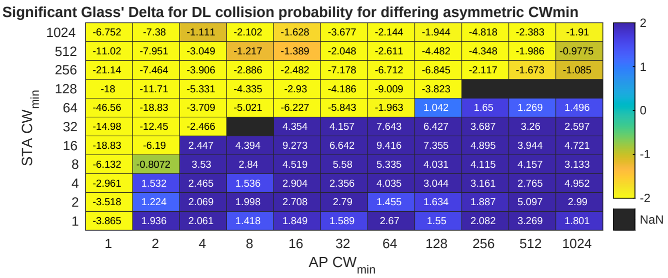

(b) Significant effect sizes on DL collision probability for experiment 10

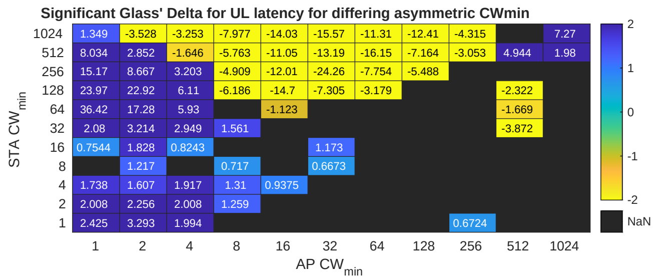

(c) Significant effect sizes on UL latency for experiment 10

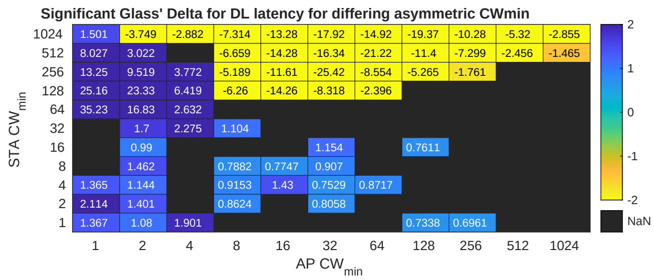

(d) Significant effect sizes on DL latency for experiment 10

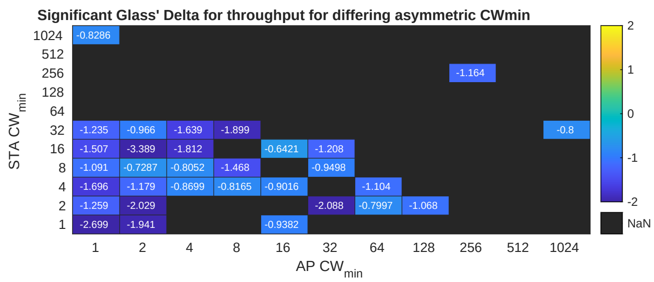

(e) Significant effect sizes on throughput for experiment 10

Figure A.2: Significant Glass' Δ for experiment 10. Yellow shaded values indicate improvement for the ”reset” system over the ”continue” system, while blue shaded values indicate regression. NaN values, colored black, indicate insignificant results.

A.3. Experiment 12: Bounded random rate traffic with proportional RU allocation and varying asymmetric CW$_{min}$

This experiment aimed to expand the search area of Experiment 11 to find any set of asymmetric AP and STA CW$_{min}$ values that would result in a deviation from the trends seen in Experiment 6. As such, proportional RU allocation was turned on, the uniform data generation rate was set to $[4000, 6000]$ Hz, and the AP and STA CW$_{min}$ were controlled to vary independently from 1 to 1024.

A.3.1. Results and performance evaluation

The significant effect sizes for this experiment are plotted in the heatmaps of Figure A.3, as explained in Section 3.5.2. Figure A.3a and Figure A.3b show the UL and DL collision probability, Figure A.3c and Figure A.3d show the UL and DL latency, and Figure A.3e shows the overall throughput of the network.

Once again, this experiment sees the trends occur on the effect sizes for each key metric as its non-random counterpart in Experiment 6. For a full analysis of these trends, refer to Section 4.3.1. These results further confirm that the effect of proportional RU allocation and random data generation rates on the comparison of the ”reset” and ”continue” systems is not strong enough to cause new effects to emerge.

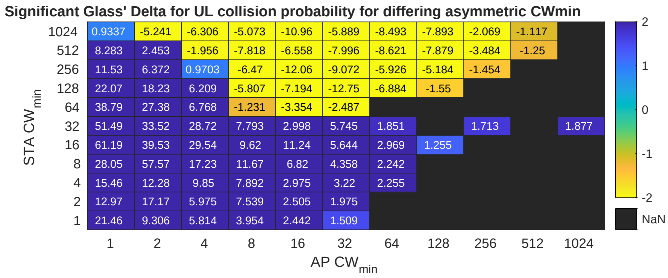

(a) Significant effect sizes on UL collision probability for experiment 12

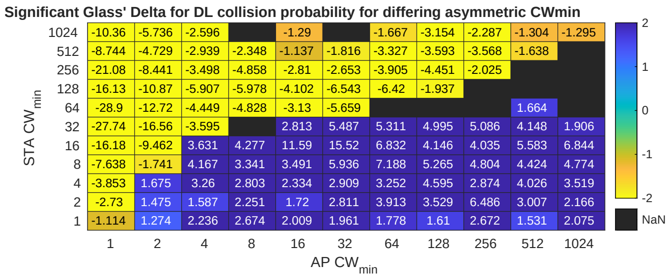

(b) Significant effect sizes on DL collision probability for experiment 12

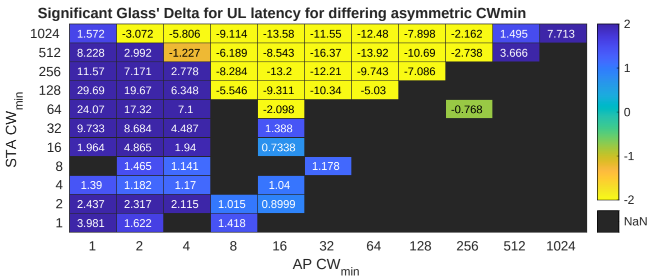

(c) Significant effect sizes on UL latency for experiment 12

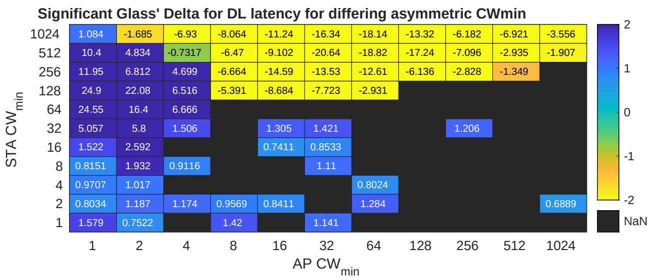

(d) Significant effect sizes on DL latency for experiment 12

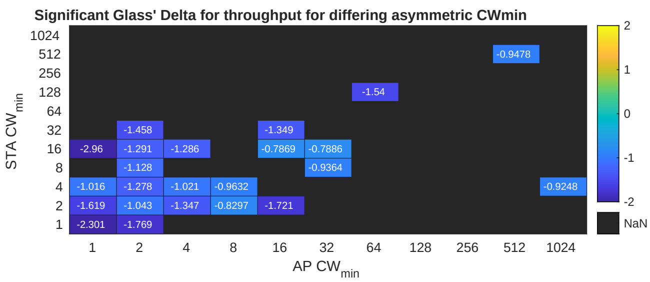

(e) Significant effect sizes on throughput for experiment 12

Figure A.3: Significant Glass' Δ for experiment 12. Yellow shaded values indicate improvement for the ”reset” system over the ”continue” system, while blue shaded values indicate regression. NaN values, colored black, indicate insignificant results.