

Potential Pitfalls of the Message in Message Mechanism in Modern 802.11 Networks

Wei Wang, Wai Kay Leong, and Ben Leong
School of Computing
National University of Singapore
{weiwang, waikay, benleong}@comp.nus.edu.sg

ABSTRACT

We study the performance impact of the Message in Message (MIM) mechanism in modern 802.11 networks. The MIM mechanism refers to the capability of receiver to abandon an ongoing reception of an 802.11 MAC frame and shift to decode another frame with a higher signal strength. MIM is a common feature in modern 802.11 adapters and it has been shown to improve spatial concurrency. However, our measurement study in a campus WLAN shows that under certain conditions, MIM could cause a throughput degradation of more than 30% when enabled, instead of improving it as expected. With comprehensive experiments using commercial 802.11n adapters, we characterize the impact of MIM for a range of parameters and for different scenarios. We show that a simple adaptive MIM scheme can potentially achieve throughput that is close to the optimal. Our method is practical because it can be easily implemented in existing commodity 802.11 adapters.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: [Wireless communication]; C.4 [Performance of Systems]: Measurement techniques

General Terms

Experimentation; Measurement; Performance

Keywords

802.11n; frame aggregation; message in message mechanism

1. INTRODUCTION

Message in Message (MIM) is a physical layer mechanism in modern 802.11 adapters [12, 18]. When MIM is enabled, the adapter will abandon the ongoing reception of a frame to start decoding another frame with a stronger signal. In other words, a frame with a higher signal strength can “knock out”

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
WiNTECH'14, September 7, 2014, Maui, Hawaii, USA.
Copyright 2014 ACM 978-1-4503-3072-5/14/09 ...\$15.00.
<http://dx.doi.org/10.1145/2643230.2643231>.

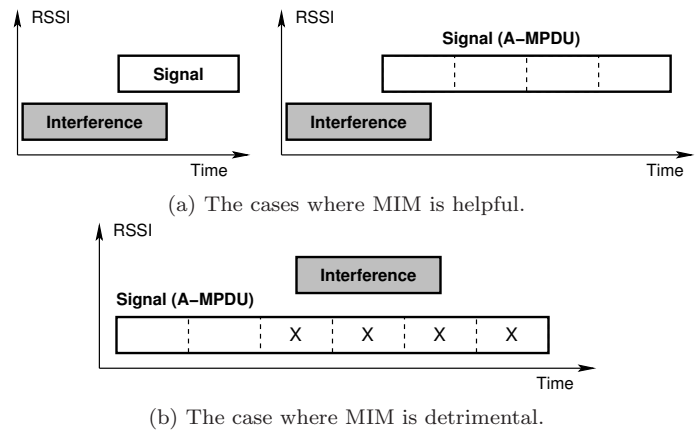


Figure 1: MIM may not always be helpful.

a frame with lower signal strength. This mechanism is also present in the adapter hardware for sensor networks [23].

It has been shown that MIM can help to improve spatial reuse in 802.11g WLANs [15] and also reduce packet loss in sensor networks [5]. Figure 1(a) serves to illustrate why MIM can often be helpful. In this example, the target data frame is stronger, but arrives in the midst of a weaker interfering frame. In this situation, MIM will cause the interfering frame to be “knocked out” and the receiver to switch to receive the target frame. Without MIM, both frames would have been corrupted and the sender of the target frame will have to retransmit the lost frame.

The 802.11n standard now includes a new feature called aggregate MAC Protocol Data Unit (A-MPDU) to reduce transmission overhead [19, 17, 11], by aggregating several frames into one physical data frame. What we have found is that MIM can be detrimental when a short but stronger interfering frame collides midway with a long, but weaker A-MPDU. We illustrate the problem in Figure 1(b). Without MIM, the receiver would successfully decode the first two and also the last frame of the A-MPDU and reply with a Block ACK (BA) indicating that the middle three frames were corrupted. The sender would then only have to retransmit the corrupted frames. However, with MIM enabled, the A-MPDU would be knocked out by the interfering frame and thus the last frame in the A-MPDU cannot be decoded. In addition, the receiver would not send a BA. This will cause the sender to either time out and retransmit the entire A-MPDU, or to send a Block ACK Request (BAR) frame. The 802.11 standard [1] does not mandate that the BAR be sent

and we found that the Cisco 1140 AP in our campus WLAN does not do so. This means that the entire A-MPDU will be retransmitted, incurring unnecessary overheads.

To better understand how MIM could potentially degrade performance under unfavorable scenarios, we conducted a comprehensive set of experiments on the MIM mechanism using commercial 802.11n adapters. In particular, we investigated various scenarios by varying A-MPDU size, interfering frame air time, and received signal strength difference. We also studied the effect of channel bonding and adjacent-channel interference. Interestingly, we found that strong interference at an adjacent channel is still able to knock out an A-MPDU that is being received. To the best of our knowledge, this paper is the first comprehensive study of how enabling and disabling the MIM mechanism can affect the reception of A-MPDU under interference.

Our study suggests that instead of necessarily always being helpful, the MIM mechanism is a double-edged sword, and it should be enabled judiciously. To this end, we proposed and evaluated a simple yet effective method to adaptively decide when to enable MIM by monitoring the received data frames. We show that our proposed adaptive MIM scheme is able to achieve near-optimal throughput with commodity 802.11 adapters.

2. IMPACT OF MIM: A QUALITATIVE STUDY

In this section, we describe a measurements study where we investigate the impact of MIM with modern 802.11n adapters. The experiments were conducted with our campus enterprise WLAN. First, we identified two adjacent access points (AP) that were hidden from each other, and placed a client at three different positions as illustrated in Figure 2. Position-1 had a stronger received signal strength (RSS) from AP-1 than that from AP-2, while Position-3 had a stronger RSS from AP-2 than that from AP-1. Position-2 was a point where the RSS from both AP-1 and AP-2 were similar. The corresponding RSS values at each position are shown in the following table:

Position	1	2	3
RSS from AP-1 (dBm)	-59	-66	-73
RSS from AP-2 (dBm)	-73	-65	-57

The client is an Advantech board with a Compex WLE350NX adapter in which we were able to enable and disable the MIM mechanism (more details in Section 3.1). The experiment was conducted during the vacation period to minimize interference from user traffic. The first client is associated with AP-1 and a saturating UDP flow with 1,500-byte payload was sent to the client from a machine in our laboratory. To generate interference from AP-2, we connected another client (not shown in Figure 2) to AP-2 and sent a 1,500-byte non-saturating UDP flow to this interfering client from another machine in our laboratory. To mitigate MAC ACK interferences [21] from the interfering client, the interfering client was positioned further away from AP-1’s client (i.e., on the “other” side of AP-2) and also operated at reduced transmission power. For each test, the UDP flows ran for 30 s and we measured the throughput obtained by AP-1’s client. We repeated the test three times at each position to mitigate the impact of external factors.




Figure 2: Campus WLAN experiment setup.

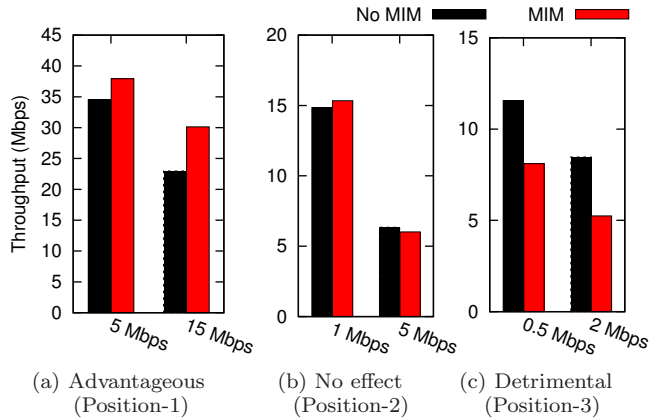


Figure 3: Impact of MIM mechanism for three different scenarios. The x-axis refers to the throughput of the interfering frames from AP-2. We did not impose the same level of interference as the one for Position-1, as it would cause the throughput from AP-1 to drop to nearly zero for the other two positions regardless of MIM.

Figure 3 shows the average UDP throughput of the client that was connected to AP-1, for both with MIM enabled and with MIM disabled. At Position-1, the average throughput was higher with MIM enabled than with MIM disabled, which shows that MIM is indeed helpful as expected when the RSS of the interfering signal is lower than the target frame. At Position-2, enabling or disabling MIM had little effect since both the target and interfering frames had similar RSS. At Position-3, the average throughput with MIM enabled was much lower (more than 30%) than that when MIM was disabled. When examining the first trace at Position-3, we found that about 28% of the A-MPDU that the client had started to decode from AP-1 were “knocked out” by the interfering frames from AP-2 as these frames had a higher RSS than the target A-MPDUs. While the scenario with the client at Position-3 will be uncommon in enterprise WLANs (as the client will likely associate with the stronger AP-2), this scenario is likely to occur in unplanned 802.11 networks where a client does not have the choice of associating with a stronger AP [2]. In addition, the client may still be subject to strong interference from frames sent by the clients of neighboring APs [21].

Thus, we conclude from this simple experiment that enabling MIM is not always helpful, and can in fact, potentially be detrimental when the interfering signal is stronger than the target A-MPDU.

3. EFFECT OF MIM ON A-MPDU RECEPTION

In this section, we investigate how the MIM mechanism affects the reception of A-MPDUs under interference in com-

Table 1: Data rate and the corresponding size of the A-MPDU.

MCS index	0	1	2	3	4	5	6	7
Data rate (Mbps)	6.5	13	19.5	26	39	52	58.5	65
A-MPDU size	2	4	6	8	12	16	18	20




Figure 4: Experiment setup for MIM characterization. $d_s > d_i$ ensures that the interfering node has a stronger signal than the sender at the receiver.

mercial 802.11n adapters. The Frame Delivery Ratio (FDR) of A-MPDU is used as the performance metric.

3.1 Experimental Methodology & Setup

Three physical WiFi adapters were used to conduct our experiments. They were placed a few meters apart in our laboratory according to the topology shown in Figure 4, with the interfering node placed closer to the receiver so the receiver will receive a stronger signal from the interfering node than that from the sender. The receiver is an Advantech board with a Compex WLE350NX adapter (Atheros AR9580 chipset). The sender and interfering node are both ALIX boards with a Wistron DNMA92 adapter (Atheros AR9220 chipset). All the three boards run OpenWRT (Attitude Adjustment release) with the ath9k WiFi driver. The MIM mechanism is enabled by default at the receiver and can be disabled by clearing a bit in the `AR_PHY_RESTART` hardware register. We set the adapters to operate in *monitor* mode as other modes like *managed* mode and *ibss* mode will implicitly transmit some management frames, which may affect the FDR results. All the experiments (except the ones in Section 3.5 and Section 3.6) were conducted over channel 56, which we had verified to have little interference from our campus WLAN.

Figure 5 illustrates a simple scheme we used to ensure that the interfering frame arrives later than the target A-MPDU at the receiver. The receiver first broadcasts a small poll message at the lowest data rate which will trigger the sender to immediately transmit an A-MPDU (with no retry) to the receiver. The interfering node with carrier sensing disabled, will wait for a random time t before sending the interfering frame with a broadcast MAC destination address. t is uniformly distributed between zero and the air time duration of the A-MPDU. This ensures that the interfering frame will collide with the A-MPDU that is being received at the receiver with a high probability. However, this method of generating interference is not foolproof and a small number of interfering frames will not collide with the A-MPDU. We omit such instances as we are only interested in investigating the cases where collisions occur.

In our experiments, the poll messages were sent every 50 ms, i.e., we collected 20 samples every second, over 100 s. To mitigate the impact of any external factors, we toggled MIM between enabled and disabled every 10 s. We also conducted a separate experiment to verify that the frame de-




Figure 5: Polling scheme ensures that the interfering frame arrives t time later than the A-MPDU. t is uniformly distributed between zero and the air time duration of the A-MPDU.

livery ratio of A-MPDU was close to 1 when there was no interference.

We modified the ath9k driver to allow the sender to set the data rate and the number of frames aggregated for each A-MPDU. In our experiments, each frame in an A-MPDU has a 1,500-byte payload, and the number of frames aggregated into an A-MPDU was chosen to achieve a total air time duration of about 3.8 ms, which is slightly smaller than the maximum 4 ms allowed by the ath9k driver. Table 1 summarizes the number of frames in an A-MPDU (A-MPDU size) according to the data rate. For the rest of the paper, we will refer to data rate using MCS index since different MCS indices may use the same data rate (in Mbps), and there might be ambiguity if we specified the data rate in Mbps.

3.2 A-MPDU Size

In this section, we investigate the impact of MIM by varying the A-MPDU size. The payload of the interfering frames was set at 50 bytes, which is the size of a typical TCP-ACK packet observed in our traces. The interfering frames were transmitted at data rate MCS-2 (19.5 Mbps) and had a stronger signal strength of about 14 dB higher than the A-MPDU. The air time duration of an interfering frame under these settings was about $60 \mu\text{s}$.

Figure 6 shows the normalized distribution of the number of frames delivered per A-MPDU for A-MPDUs of size 2, 4, 12 and 20. With MIM enabled, the probability density of the number of frames delivered per A-MPDU is evenly distributed for all the graphs. This is because a collision in one frame of an A-MPDU causes all subsequent frames to be lost, and the interfering frame collides uniformly across the A-MPDU (see Figure 5). With MIM disabled, there is a small non-zero probability that no frames were received due to the interfering frame colliding with the PLCP header of the A-MPDU, causing a failure to decode the entire A-MPDU. The remaining probability density is skewed towards losing only one frame for A-MPDUs of size 2 and 4 (Figures 6(a) and 6(b)), and is approximately two frames as the A-MPDUs size increases (Figures 6(c) and 6(d)).

Intuitively, it seems likely that the two-frame loss is caused by the interfering frame “straddling” two consecutive frames when colliding with an A-MPDU. For A-MPDU size of 20, each frame in an A-MPDU takes about $190 \mu\text{s}$ air time, as compared with the $60 \mu\text{s}$ of interfering frame. Theoretically, the interfering frame would straddle two frames in an A-MPDU at a probability of 0.3. However, in Figure 6(d), we see that the majority of A-MPDU lose two frames. We suspect that the reason for this higher than expected loss ratio is because the impact of an interfering frame is more than its total airtime and it takes receiver some additional amount of time to re-adjust itself to receive the original A-




Figure 6: Distribution of the number of frames delivered per A-MPDU.

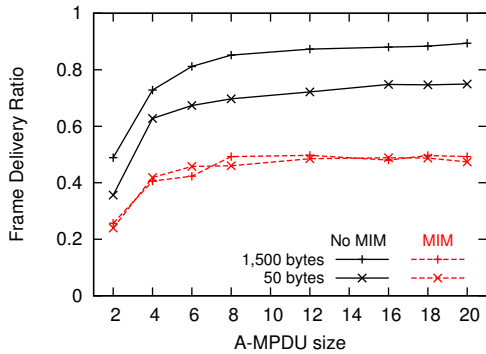


Figure 7: Impact of A-MPDU size with interfering frame payload of 50 and 1,500 bytes.

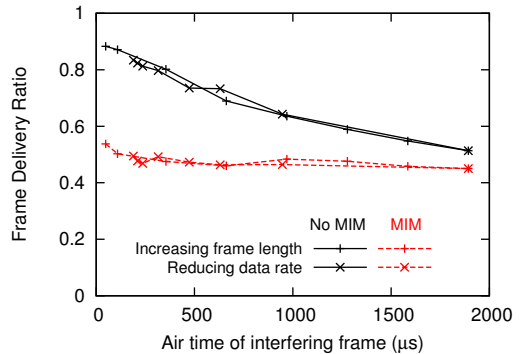


Figure 8: Impact of the air time of interfering frames. The A-MPDU size is 20.

MPDU after being interrupted by a strong interfering frame. Unfortunately, we do not possess the equipment that would allow us to make the measurements required to verify this hypothesis.

We plot the average frame delivery ratio (FDR) for each A-MPDU size in Figure 7, and confirmed that with MIM disabled, the average FDR for interference with 50-byte payloads converges to a value of about 0.9, which is the approximate ratio of the air time duration of the interfering frame to the total A-MPDU. With MIM enabled, the FDR converges to 0.5, which is expected as the interfering frame would collide uniformly across the A-MPDU. When the payload of the interfering frame is increased to 1,500 bytes, the results are similar to that of 50-byte payload when MIM is enabled. But with MIM disabled, the average FDR now converges to a much lower value of about 0.75. This shows that with MIM disabled, the air time duration of the interfering frame will affect the average FDR of the A-MPDU.

3.3 Interfering Frame Air Time Matters

It turns out that the achieved frame delivery ratio for an A-MPDU depends on the air time duration of the interfering frame. To demonstrate this, we performed two experiments: in one, we increased the air time of the interfering frame by increasing its payload, while in the other, we reduced the data rate. In the first experiment, we varied the payload of the interfering frame from 1 to 1,500 bytes, keeping the data rate at MCS-0 (6.5 Mbps). In the second experiment, we fixed the payload of the interfering frame at 1,500 bytes and varied the data rate of the interfering frame from MCS-7 (65 Mbps) to MCS-0 (6.5 Mbps).

Figure 8 shows the achieved average FDR against the air time duration of the interfering frame from both experiments. The results show that the air time duration has a direct effect on the FDR of the A-MPDU and that there is no difference between increasing the payload or reducing the data rate of the interfering frame. With MIM enabled, the FDR remains at a stable value just below 0.5 even as the interfering frame air time increases. With MIM disabled, the FDR decreases as the interfering frame air time increases.

Our results suggest that as the air time duration of the interfering frame decreases, the negative impact of enabling MIM unnecessarily increases. Therefore, to understand the impact of enabling MIM in real networks, we investigated the distribution of the air time duration of frames “in the wild”. Figure 9 shows the distribution of the air time duration of frames captured at different locations in a university library over a 2-hour period during office hours. The results show that most of the frames have a very short air time duration. In particular, the median air time is only about $30 \mu\text{s}$, and 90% of the frames have air time durations shorter than $300 \mu\text{s}$. We can also see two distinct bandings with about 25% of the frames having an air time of $20 \mu\text{s}$ and 20% at $190 \mu\text{s}$. Upon closer inspection, we found that the former $20 \mu\text{s}$ frames were MAC ACK frames [21] and the latter frames were multicast frames arising from the Neighbor Discovery Protocol of IPv6.

3.4 Impact of Received Signal Strength Differences

In the previous experiments, the received signal strength (RSS) of interference was set to about 14 dB higher than that

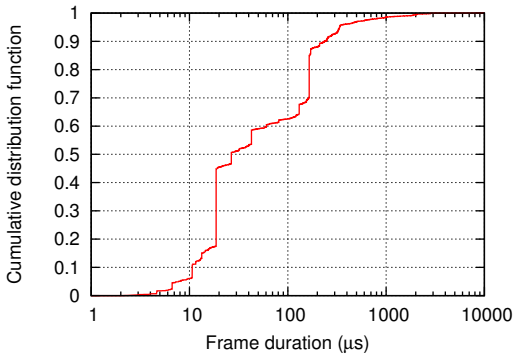


Figure 9: Distribution of frame air time duration in a university library.

of the A-MPDU signal. In this section, we investigate and characterize the impact that the difference of RSS has on the MIM mechanism. The transmission power of A-MPDU sender was held constant while the `iwconfig` command was used to change the transmission power of the interfering node to set the RSS difference.

Figure 10 shows the average FDR against the RSS difference with A-MPDUs of size 6 and 20, and interfering frames with 1,500-byte payload sent at MCS-7 (65 Mbps) and MCS-15 (130 Mbps). With MIM disabled, the FDR decreases slightly as the RSS difference increases. This effect is more pronounced when the air time duration of the interfering frame is longer (i.e., using MCS-7). Like the results for interference highlighted in Section 3.2, we suspect that it takes time for the receiver to re-adjust itself to receive the original A-MPDU after being interrupted by an interfering frame. Again, we do not currently have the means to verify this and leave this as future work.

With MIM enabled, the FDR drops steeply to below 0.5 when the RSS difference is larger than 10 dB, regardless of the size of the A-MPDU. This shows enabling MIM is effective only when the RSS difference is larger than 10 dB and corroborates an earlier study on Atheros 802.11a adapters [12].

In addition, the 10 dB threshold of RSS difference remains the same when the interfering frame was sent at MCS-7 and MCS-15. Note that the interfering frame with MCS-15 uses two spatial streams for its payload, as compared with one for MCS-7. In our experiments, the interfering node was in direct line-of-sight of the receiver, and the receiver could hardly decode the payload of the interfering frames sent at MCS-15 [4]. However, the preamble of the interfering frame is still sent in a single spatial stream even when the payload is sent across two. Thus, in terms of the capability to knock out A-MPDU, there is no difference when the interfering node uses multiple spatial streams for its frame payload.

3.5 Channel Bonding

The 802.11n standard specifies an extended channel bandwidth of 40 MHz (which is twice that of the conventional 20 MHz bandwidth) to increase network capacity by doubling the data rate. This use of an extended channel bandwidth is known as *channel bonding* [4]. In this section, we investigate the impact of the MIM mechanism when channel bonding is used.

In this experiment, the 40 MHz-wide channel used comprises channel 52 and channel 56 (see Figure 11). We consider all the possible permutations of channel widths for the

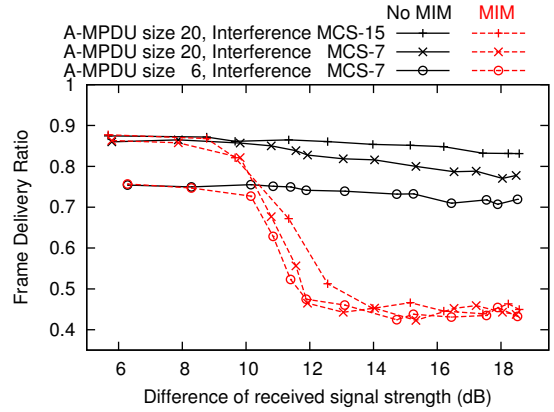


Figure 10: Impact of received signal strength. The payload of interfering frame is 1500-byte.




Figure 11: The two channels used in the channel bonding experiments.

three nodes as shown in Table 2. Note that for the receiver to receive the sender's signal, the receiver channel cannot be narrower than the sender. The data rate MCS-7 (65 Mbps) was used for both the sender and interfering node, and the payload of the interfering frame was 50 bytes.

In Figure 12, we plot the resulting frame delivery ratio against RSS difference for the cases listed in Table 2. Our results show that there is no apparent difference among the different cases. For example, case-1 shows that the interfering frames sent at bonded channel can still knock out A-MPDU when the sender and receiver use 20 MHz-wide channel. Similarly, case-4 shows that when the interfering node uses a 20 MHz-wide channel, it can disrupt the sender and receiver on a bonded channel.

3.6 Adjacent-channel Interference

In the experiments described in the previous sections, all the interference occurred in the same (or overlapped) channel as the sender. In this section, we investigate the impact of adjacent-channel interference when the interfering frame is sent on a channel adjacent to the A-MPDU. Adjacent-channel interference is common as the spectral mask in the radio hardware is not perfect. In this experiment, the sender is set at channel 56 and the interfering node is set at channel 52 (see Figure 11). We investigated both cases where the receiver listens only on channel 56 and where it uses channel bonding and listens to both channels 52 and 56.

One minor complication in our experiment is that the original method of sending a poll message to clock the interfering frames no longer works as the sender and interfering node are now listening on different channels. Thus, we employed a different approach to generate the interference, by having the interfering node continuously broadcast an 50-byte interfering frame at MCS-7 (65 Mbps), at 4 ms intervals. At the same time, the sender sends the A-MPDU at an interval

Table 2: Combinations of channel width (MHz) used in the channel bonding experiments. The 20 MHz channel refers to channel 56.

	Sender	Receiver	Interfering Node
Case-1	20	20	40
Case-2	20	40	20
Case-3	20	40	40
Case-4	40	40	20
Case-5	40	40	40




Figure 12: Effect of MIM with channel bonding. The payload of interfering frame is 50-byte.

that is uniformly distributed from zero to 100 ms. The carrier sensing mechanism at both the sender and interfering node was disabled. In this way, we ensure that the A-MPDU (with air time duration of about 3.8 ms) will collide with a single interfering frame for most of time. This technique cannot guarantee that the A-MPDU will always arrive before the interfering frame. However, since the air time duration of A-MPDU is much longer than that of interfering frame (several tens of μ s), the probability that the A-MPDU will be received later than interfering frame is very small.

Figure 13(a) shows the FDR when the channel width of receiver is set to 20 MHz. For comparison, we also included the scenario where the sender and interfering node are on the same channel using the same experimental setup. In order to get a larger RSS difference, we also reduced the transmission power of the sender, while ensuring that the FDR of the A-MPDU remains close to 1 when there is no interference.

When MIM is disabled and the sender and interfering node are on adjacent channels, the FDR is close to 1 when the RSS difference is small but gradually decreases when the RSS difference increases. This is expected as adjacent-channel interference becomes more pronounced at larger RSS differences. As the RSS difference increases further, the FDR appears to converge to the case where the sender and interfering node are on the same channel.

Next, we examine the case where MIM is enabled and the sender and interfering node are on adjacent channels. At RSS differences smaller than 20 dB, the FDR is similar to when MIM is disabled as the leaked signal from the adjacent channel is weak. However, from an RSS difference of about 20 dB onwards, the receiver begins to pick up and decode the preamble of the interfering frame and kick out the A-MPDU. Thus the FDR begins to decrease sharply until it becomes equivalent to the case where the sender and interfering node are on the same channel when the RSS difference is greater than 30 dB. This shows that when the interfering node has a sufficiently higher signal strength than the sender, MIM can

still take effect and be detrimental even though the sender and interfering node are on adjacent channels. The exact threshold at which this happens is likely to be dependent on the spectral mask of the adapter hardware.

Figure 13(b) shows the FDR when the receiver uses channel bonding and listens to both channels. With MIM disabled, the FDR is similar to the case in Figure 13(a) where the receiver only listens on one channel, as the interfering frame appears to be noise to the receiver. On the other hand, with MIM enabled and RSS difference above 10 dB, the receiver can clearly pick up the interfering frame and abandon the reception of A-MPDU, resulting in a sharp decline of the FDR. In other words, the receiver behaves similarly to when the sender and interfering node are on the same channel. Therefore, we can see that listening on the 40 MHz-wide channel could cause the receiver to be more susceptible to MIM effect due to adjacent-channel interference, as compared with the one listening on a 20 MHz channel.

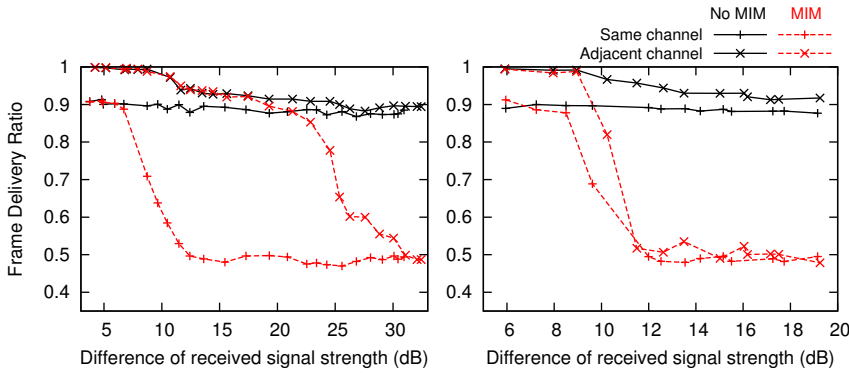
4. ADAPTIVE MIM

Our results clearly suggest that the MIM mechanism can be helpful or detrimental under different conditions. The natural question is whether it is possible to adaptively toggle MIM on and off in a manner such that MIM is turned on only when it is helpful and not when it is harmful. Our key insight is that it is possible to determine whether MIM is helpful by observing the interference patterns on the received A-MPDU frames as illustrated previously in Figure 1. Thus, we can simply count and compare the number of occurrences where MIM was helpful (see Figure 1(a)) versus where MIM was detrimental (see Figure 1(b)) to decide whether and when to enable or disable MIM.

In order to count the occurrences of each case, we modified the ath9k driver to retain the partial frame that is knocked out when the MIM mechanism kicks in. Specifically, the modified driver no longer ignores frames that have the physical error flags OFDM-RESTART or CCK-RESTART set as these error flags indicate that the frame was knocked out by a stronger frame. If the knocked-out frame is not addressed to the receiver and is immediately followed by one addressed to the receiver, then we consider this an occurrence where MIM is helpful. On the other hand, if a frame addressed to the receiver is knocked out by a frame that is not addressed to the receiver, it is an occurrence where MIM is detrimental.

Figure 14 shows the state diagram of our proposed method to adaptively enable and disable MIM. We count the number of helpful (n_{good}) and detrimental (n_{bad}) occurrences in each epoch of 1 s. MIM will be disabled when $n_{good} < n_{bad}$ for a given epoch. A minor complication is that the adapter can only count and update n_{good} and n_{bad} when MIM is enabled. Thus, we need to periodically enable MIM to resume counting after MIM is disabled. We adopt a simple exponential backoff approach by enabling MIM after k epochs, and doubling k every time n_{good} turns out to remain smaller than n_{bad} when we turn MIM on to do the counting. The initial value of k is set to 1 and the maximum value is 10. k will also be reset to 1 whenever $n_{good} \geq n_{bad}$. In this way, we try to strike a balance between effectiveness and responsiveness for the proposed method.

In Figure 15, we compare our adaptive MIM method to cases where MIM is always enabled or disabled for each of the scenarios described in Section 2. In addition, we also



(a) Receiver's channel width is 20 MHz. (b) Receiver's channel width is 40 MHz.

Figure 13: Effect of MIM with adjacent channel interference.

tested with the payload of interfering frame set to 50 bytes. Just like in Section 2, each test was run for 30 s. Our results show that under situations where MIM is helpful (see Figure 15(a)), our adaptive MIM method achieves similar throughput as the case where MIM is always enabled. At Position-2 where MIM has no effect, the throughput results are similar for all cases. In situations where MIM is detrimental, our adaptive MIM method achieves slightly lower throughput than the one with MIM disabled. This is because under such scenarios, our method has to occasionally enable MIM to update the counters n_{good} and n_{bad} , which incurs some overhead. We feel that this is an acceptable overhead.

5. RELATED WORK

There have been a number of studies on characterizing the MIM mechanism as well as the general *physical layer capture effect* for commercial 802.11 adapters. Lee et al. investigated the conditions under which capture effect (including MIM) takes place for 802.11a adapters of Atheros AR5112 chipset [12]. They found that RSSI difference of 10 dB is the threshold for MIM, which is similar to what we found in our experiments. Based on their findings, the authors proposed a capture effect model to enhance the physical layer of QualNet simulator [13]. Kochut et al. studied the capture effect for 802.11b adapters of Prism 2 chipset [10], but they only considered the case where the time difference between two collided frames was within preamble time. Similarly, Judd and Steenkiste focused on the case where the interfering frames are delayed, for 802.11b adapters of Prism 2.5 chipset [9]. The key limitation of these existing characterization works is that they did not consider A-MPDU, which is mandatory in the 802.11n standard (as well as in the latest 802.11ac standard).

Several proposals have utilized capture effect to improve the spatial reuse of wireless networks. Manweiler et al. proposed an MIM-aware link scheduling protocol (called Shuffle) for enterprise WLAN [15]. The basic idea of Shuffle is to order the transmissions of APs in a centralized manner such that the number of concurrent transmissions is maximized. Similarly, Lu and Whitehouse exploited the capture effect to reduce flooding latency in wireless sensor networks by allowing nodes to propagate the flood simultaneously [14]. Capture effect has also been shown to cause serious unfairness problem in 802.11 networks [22]. Ganu et al. evaluated the

degree of unfairness due to capture effect for the 802.11a/b/g adapters of Atheros AR5212 chipset, and studied how various MAC parameters affect the unfairness [6]. More recently, Wang et al. proposed an automatic contention window adjustment algorithm to mitigate the unfairness caused by capture effect in 802.11 mesh networks [20].

There have been a number of studies on the impact of frame aggregation in modern 802.11 networks. Several works showed that frame aggregation is able to greatly reduce transmission overhead and improve throughput for 802.11n links [7, 19, 17, 11]. The advantage of frame aggregation for 802.11ac links was reported by Ong et al. [16]. Ashtaiwi and Hassanein studied the performance of frame aggregation in 802.11 mesh networks and proposed a method to adjust aggregation size based on link quality [3]. Gubner and Lindemann investigated video streaming in 802.11 mesh networks, and showed that frame aggregation has positive impact on both the delay and quality of the streamed video [8].

6. CONCLUSION

In this paper, we made an important observation that the MIM mechanism in modern 802.11 hardware could potentially degrade the performance of A-MPDU reception when there is strong interference. With a comprehensive set of experiments, we characterized how turning MIM on and off would affect the reception of A-MPDU under various parameters and scenarios. We also investigated the impact of MIM in the presence of channel bonding and adjacent-channel interference. We showed that it is possible to avoid the potentially harmful effects of MIM by adaptively turning MIM on and off using a simple mechanism that monitors the interference pattern on received A-MPDUs.

As the latest 802.11ac standard adopts more aggressive frame aggregation (up to 1 MB per A-MPDU vs. 64 KB for 802.11n), we believe that our work would have broad implications on next generation 802.11 networks in the near future. In addition, our detailed measurement study would likely be helpful to researchers who are seeking to improve the existing algorithms for rate adaptation and/or A-MPDU size selection in the presence of strong interference. Furthermore, the results of our measurement studies could also be used to enhance the 802.11n models in existing simulation tools to improve the interference models for MIM implementations.

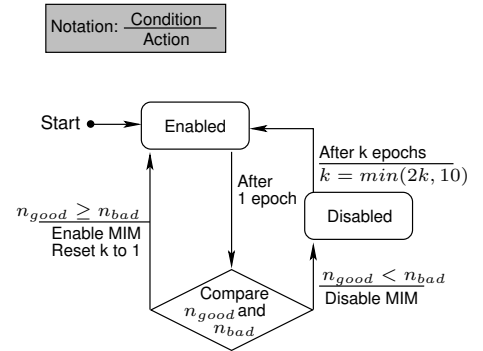


Figure 14: The state diagram of the proposed adaptive MIM method.




Figure 15: Effect of the proposed adaptive MIM method. The x-axis refers to the payload size and throughput of the interfering frames.

7. ACKNOWLEDGMENTS

This work was supported by the Singapore Ministry of Education grant T1 251RES1204.

8. REFERENCES

- [1] IEEE 802.11 standard: Wireless LAN medium access control and physical layer specification, Jun. 2012.
- [2] A. Akella, G. Judd, S. Seshan, and P. Steenkiste. Self-management in chaotic wireless deployments. In *Proc. of MobiCom '05*, Aug. 2005.
- [3] A. Ashtaiwi and H. Hassanein. Utilizing IEEE 802.11n to enhance QoS support in wireless mesh networks. In *Proc. of LCN '08*, Oct. 2008.
- [4] L. Deek, E. Garcia-Villegas, E. Belding, S.-J. Lee, and K. Almeroth. The impact of channel bonding on 802.11n network management. In *Proc. of CoNEXT '11*, Dec. 2011.
- [5] B. Firner, C. Xu, R. Howard, and Y. Zhang. Multiple receiver strategies for minimizing packet loss in dense sensor networks. In *Proc. of MobiHoc '10*, Sep. 2010.
- [6] S. Ganu, K. Ramachandran, M. Gruteser, I. Seskar, and J. Deng. Methods for restoring MAC layer fairness in IEEE 802.11 networks with physical layer capture. In *Proc. of REALMAN '06*, May 2006.
- [7] B. Ginzburg and A. Kesselman. Performance analysis of A-MPDU and A-MSDU aggregation in IEEE 802.11n. In *Proc. of Sarnoff '07*, Apr. 2007.
- [8] S. Gubner and C. Lindemann. Evaluating the impact of frame aggregation on video-streaming over IEEE 802.11n multihop networks. In *Proc. of WoWMoM '12*, Jun. 2012.
- [9] G. Judd and P. Steenkiste. Characterizing 802.11 wireless link behavior. *Wireless Networks*, 16(1):167–182, Jan. 2010.
- [10] A. Kochut, A. Vasan, A. U. Shankar, and A. Agrawala. Sniffing out the correct physical layer capture model in 802.11b. In *Proc. of ICNP '04*, Oct. 2004.
- [11] L. Kriara, M. K. Marina, and A. Farshad. Characterization of 802.11n wireless LAN performance via testbed measurements and statistical analysis. In *Proc. of SECON '13*, Jun. 2013.
- [12] J. Lee, W. Kim, S.-J. Lee, D. Jo, J. Ryu, T. Kwon, and Y. Choi. An experimental study on the capture effect in 802.11a networks. In *Proc. of WiNTECH '07*, Sep. 2007.
- [13] J. Lee, J. Ryu, S.-J. Lee, and T. T. Kwon. Improved modeling of IEEE 802.11a PHY through fine-grained measurements. *Computer Networks*, 54(4):641–657, 2010.
- [14] J. Lu and K. Whitehouse. Flash flooding: Exploiting the capture effect for rapid flooding in wireless sensor networks. In *Proc. of INFOCOM '09*, Apr. 2009.
- [15] J. Manweiler, N. Santhapuri, S. Sen, R. Roy Choudhury, S. Nelakuditi, and K. Munagala. Order matters: transmission reordering in wireless networks. In *Proc. of MobiCom '09*, Sep. 2009.
- [16] E. H. Ong, J. Knecht, O. Alanen, Z. Chang, T. Huovinen, and T. Nihtila. IEEE 802.11ac: Enhancements for very high throughput WLANs. In *Proc. of PIMRC '11*, Sep. 2011.
- [17] U. Paul, R. Crepaldi, J. Lee, S.-J. Lee, and R. Etkin. Characterizing WiFi link performance in open outdoor networks. In *Proc. of SECON '11*, Jun. 2011.
- [18] N. Santhapuri, R. R. Choudhury, J. Manweiler, S. Nelakuditi, S. Sen, and K. Munagala. Message in message MIM: A case for reordering transmissions in wireless networks. In *Proc. of HotNets '08*, Oct. 2008.
- [19] D. Skordoulis, Q. Ni, H.-H. Chen, A. P. Stephens, C. Liu, and A. Jamalipour. IEEE 802.11n MAC frame aggregation mechanisms for next-generation high-throughput WLANs. *Wireless Communications, IEEE*, 15(1):40–47, 2008.
- [20] W. Wang, B. Leong, and W. T. Ooi. Mitigating unfairness due to physical layer capture in practical 802.11 mesh networks. *IEEE Transactions on Mobile Computing*, To appear.
- [21] W. Wang, Q. Wang, W. K. Leong, B. Leong, and Y. Li. Uncovering a hidden wireless menace: Interference from 802.11x MAC acknowledgment frames. In *Proc. of SECON '14*, Jun. 2014.
- [22] C. Ware, J. Judge, J. Chicharo, and E. Dutkiewicz. Unfairness and capture behaviour in 802.11 adhoc networks. In *Proc. of ICC '00*, Jun. 2000.
- [23] K. Whitehouse, A. Woo, F. Jiang, J. Polastre, and D. Culler. Exploiting the capture effect for collision detection and recovery. In *Proc. of EmNets '05*, May 2005.