Performance Anomaly in Download TCP Flows Over IEEE 802.11n Wireless LAN

Yoshiki Hashimoto, Masataka Nomoto, Celimuge Wu, Satoshi Ohzahata, and Toshihiko Kato
Graduate School of Information Systems
University of Electro-Communications
Tokyo, Japan
e-mail: hys3224@net.is.uec.ac.jp, noch@net.is.uec.ac.jp, clmg@is.uec.ac.jp, ohzahata@is.uec.ac.jp, kato@is.uec.ac.jp

***Abstract*— The performance anomaly is one of well-known performance problems in IEEE 802.11 wireless LANs (WLANs). It reduces the throughput of all stations managed by one access point when some of them use low data rates even if the others use high rates. In the recent WLANs, such as 802.11n WLAN, providing higher data rates than the legacy ones, the performance degradation due to the performance anomaly may give larger impacts. In the previous paper, we showed the evaluations of the performance anomaly for UDP and TCP flows over 802.11n WLAN. The results show that, although the performance anomaly occurs for UDP flows, the situation is different for TCP flows. This paper presents more detailed experimental performance study on the performance anomaly for download TCP flows analyzing the throughput, the congestion window size (cwnd) and the round trip time (RTT). It concludes that, although the throughput degrades where there are low data rate stations, the so called performance anomaly does not occur in download TCP flows.**

Keywords- WLAN; IEEE802.11n; Performance Anomaly, TCP flows.

I. INTRODUCTION

Recently, a variety of equipment based on IEEE 802.11 WLAN standard [1], such as smart phones, tablets and notebooks, are widely deployed. When a number of stations access to one WLAN access point, they suffer from several kinds of performance problems. Among them, the *performance anomaly* [2][3] is a well-known problem. When some stations are located far from their access point and others are near it, the performance of the near stations is degraded to that of far located stations. This is caused by the fair assignment of channel access based on the *carrier sense multiple access with collision avoidance* (*CSMA/CA*) principle, and the multiple Media Access Control (MAC) level data rate support. Those allow stations with low bit rate to capture the channel for a long time, and it penalizes other stations with higher data rates.

Nowadays, 802.11n [1] is one of most widely adopted IEEE WLAN standards. It establishes high speed data transfer using the higher data rate support (e.g., 150 Mbps), the frame aggregation in Aggregated MAC Protocol Data Unit (A-MPDU) and the Block Acknowledgment mechanism. In spite of those improvements, 802.11n standard does not resolve the performance anomaly. So, the performance degradation by the performance anomaly will give larger impacts than the legacy standards.

There are a few papers describing the performance anomaly in IEEE 802.11n WLANs. Abu-Sharkh and Abdelhadi [4] reports that the performance anomaly still exists in 802.11n WLANs. It describes the results of the upload TCP data transfer focusing on the A-MPDU function. Our previous paper [5] shows the results of examining performance anomaly in UDP and TCP download data transfer. It describes that the performance anomaly surely occurs in UDP flows, but there may be only performance degradation in download TCP flows even in near and far located stations coexist.

This paper presents more detailed experimental performance analysis on the performance anomaly for download TCP flows. The feature of our analysis is as follows.

- During a TCP flow, the time variation of the throughput, cwnd and RTT are examined in detail, instead checking their average values.
- From the cwnd and RTT values, the traffic load to the access point is estimated, and it is compared with the actual MAC level data rate.
- The experiments with two stations and with four stations are performed.
- The commercially available access point is used with installing the firmware provided by the OpenWRT project [6]. This allows us to obtain various MAC level performance metrics and to adopt different queue management schemes at the access point.

The rest of this paper consists of the following sections. Section 2 shows the experimental settings. Sections 3 and 4 describe the results of the two station and four station experiments, respectively. In the end, Section 5 gives the conclusions of this paper.

II. EXPERIMENTAL SETTINGS

Fig. 1 shows the configuration of our experiment. Up to four stations (STAs) conforming to 802.11n with 5GHz band are associated with one access point, which is connected with a server through 1Gbps Ethernet. Some STAs are located at a near position to the access point, and another STA is located in various positions in the experiment.

The detailed specifications of stations and the access point are shown in Table 1. We use commercially available notebooks and access point in the experiment. As described above, we use the access point firmware provided by the OpenWRT project. By using this firmware, we can obtain several performance metrics in the access point described below.

By use of the OpenWRT firmware, we can configure the queue management schemes used in the access point. It supports the following schemes.

- *FIFO*: A scheme to use one queue to store all frames being sent by the access point.

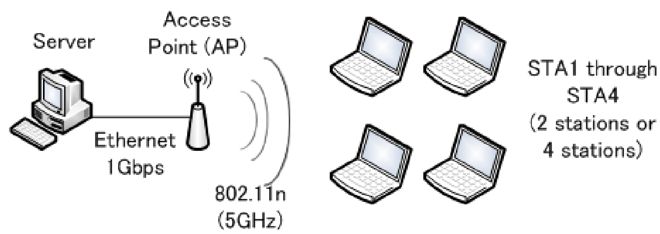

Figure 1. Configuration of experiment.

TABLE I. SPECIFICATIONS OF STAS AND ACCESS POINT (AP)

STA	Manufacturer/Model	DELL Insilon 14
	Operating system	Ubuntu 14.04LTS (kernel 3.13)
AP	Manufacturer/Model	BUFFALO AirStation WZR-HP-AG300H
	Firmware	OpenWRT (BarrierBraker, r444, selfbuild)
	WLAN chip	Atheros AR7161
	WLAN driver	ath9k

- *CoDel* [7]: An active queue management scheme designed to resolve the Bufferbloat problem [8]. It uses packet-sojourn time in a queue as a control parameter, and drops one packet among those staying in the queue too long.
- *Stochastic Fare Queueing* (*SFQ*): A scheme to provide a separate queue for packets of an individual flow.
- *FQ_CoDel*: A scheme which combines SFQ and CoDel. In OpenWRT, FQ_CoDel is the default queue management scheme.

In the experiment, we used all those queueing management schemes for the performance evaluation.

The access point uses two streams in the spatial division multiplexing with each channel using 60 MHz bandwidth, and, as a result, the data rate ranges from 6.5 Mbps to 300 Mbps.

In the experiment, data is transferred from the server to two or four stations through the access point. The server uses *iperf* tool [9] to generate TCP data segments. As for TCP parameter settings, we used the native ones in the Linux operating system. Specifically, the TCP version is CUBIC TCP.

During the data transfer, the following performance metrics data are collected for the detailed analysis for the communication;

- MAC level data rate (an average during one second, collected at the access point for individual A-MPDUs from the ath9k device driver [10]),
- the number of MPDUs per A-MPDU (an average during one second, collected at the access point for individual A-MPDUs from the device driver),
- TCP throughput (an average during one second, measured at the server by use of *tcpdump*),
- cwnd (an average during one second, measured at the server for every data segment sent by use of *tcpprobe* [11]), and
- TCP level RTT (an average during one second, measured at the server for every ACK segment received).

From measured cwnd and RTT values, we calculate the estimated TCP load by the following equation.

$$Estimated\ load\ (Mbps) = \frac{cwnd * 1{,}500 * 8}{RTT * 1{,}000{,}000} \qquad (1)$$

The experiment is conducted in a building constructed with reinforced concrete. Fig. 2 shows the layout inside the building and the positions of network equipment. The thick black line represents the exterior wall of the building and the thin black line represents the interior wall, which is made from wood.

In the case of two station experiment, the black circles named “AP” and “STA1” correspond to the positions of the access point and the near station, STA1, respectively. These are fixed throughout the experiment. The black circles named “Position0” through “Position7” represent the positions of the far station, STA2. It is located in one of these eight positons in the experiment.

In four station experiment, three stations, STA1 through STA3, are located at position “STA1” and the other station, STA4, is located at one of positions from “Position0” to “Position7.”

III. RESULTS FOR TWO STATION EXPERIMENT

A. Experimental Scheme

In the two station experiment, we measured the performance of TCP data transfer from the server to the stations, by changing the position of STA2 and the queue management scheme in the access point. For each STA2 position and queue scheme, we executed three experiment runs, each of which is 120 second TCP data transfer.

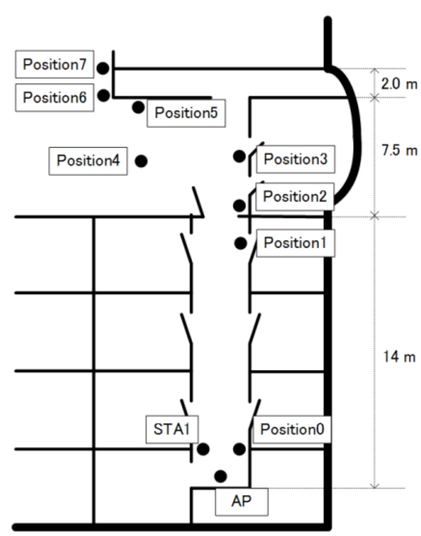

Figure 2. Layout inside building and position of equipment.

B. Overall Results

Fig. 3 shows the relationship between the position of STA2 ("PS" in the figure means "Position") and the average of MAC level data rate of STA1 and STA2. STA1 located near the access point keeps high data rate around 250 Mbps. On the contrary, the average data rate of STA2 decreases along with its location being far away from the access point. More specifically, the data rate of STA2 takes a similar value for Position3 and Position4, and Position6 and Position7. Fig. 3 is the result when the FIFO queue management scheme is used at the access point. The cases when the other schemes are used showed similar results. In the rest of this paper, the position of STA2 (far located station) is represented its average data rate.

Fig. 4 shows the relationship between the STA2 average data rate and the average TCP throughput (average values throughout three experimental runs). In this graph, solid lines indicate the results of STA1 and dashed lines indicate those of STA2, and the color of lines corresponds to the individual queue management scheme. The TCP throughput of not only STA2 but also STA1 decreases as the position of STA2 becomes far from the access point. There was no difference among the queue management schemes in the access point.

Fig. 5 (a) shows the average cwnd versus the STA2 average data rate. In this case, the results largely depend on the queue management schemes. In FIFO, the average cwnd is large and varied from 1 to 900 packets. In SFQ, the queue length for an individual flow is limited to 127 packets, and this in turn limits the cwnd. CoDel and FQ_CoDel drop packets which stay in the queue for a long time, and so the cwnd is suppressed. In all queue schemes, the average cwnd is small when STA2 is located at Position6 and Position7. In this situation, there seems to be a lot of packet losses in both STA1 and STA2 TCP flows.

Fig. 5 (b) shows the average TCP level RTT versus the STA2 average data rate. Here, the results also depend on the queue management scheme. Especially, FIFO has a large RTT compared to the other schemes. In FIFO, both cwnd and RTT are larger than the others, but, since both of them are larger in the similar magnitude, the throughput is also similar with the others.

We consider that the overall results are not enough to explain the results for TCP flows. So, we explain the detailed analysis below.

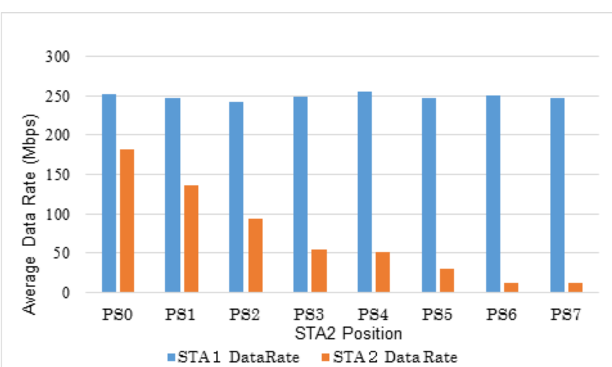

Figure 3. Average MAC level data rate vs. STA2 position (FIFO).

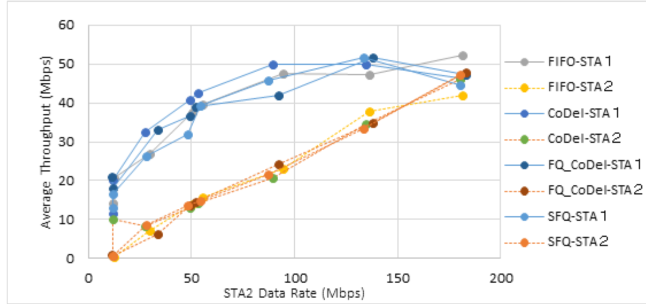

Figure 4. Average TCP throughput vs.STA2 data rate.

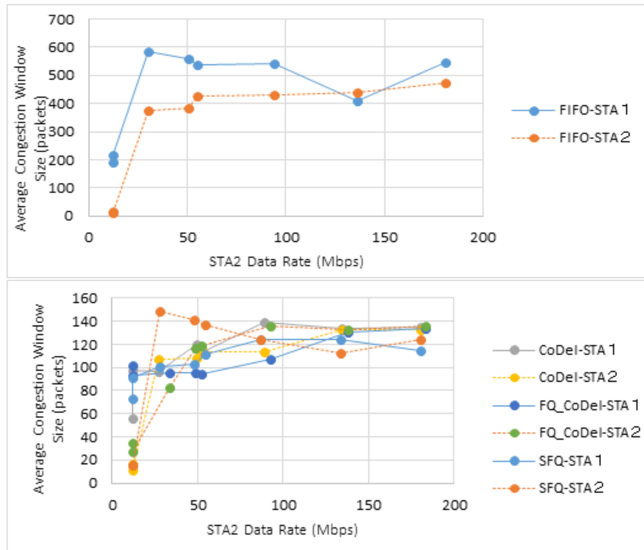

(a) Average cwnd vs. STA2 data rate

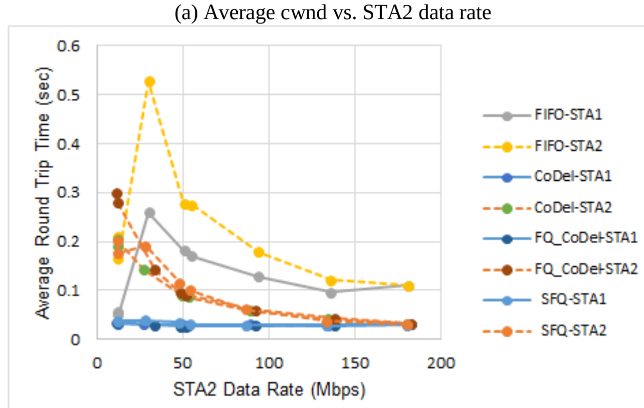

(b) Average RTT vs. STA2 data rate

Figure 5. Average cwnd and RTT.

C. Detailed Results

Figs. 6, 7 and 8 show the time variation of TCP throughput of STA1 and STA2 when STA2 is located at Position0, Position 4 and Position7, respectively. Those results are obtained by use of FIFO at the access point.

When STA2 is located at Position0, the TCP throughput of STA1 and STA2 varies around 50 Mbps. Throughout 120 sec. data transfer, two stations provide high TCP throughput. In the time frames from 50 sec. to 70 sec. and after 90 sec., the

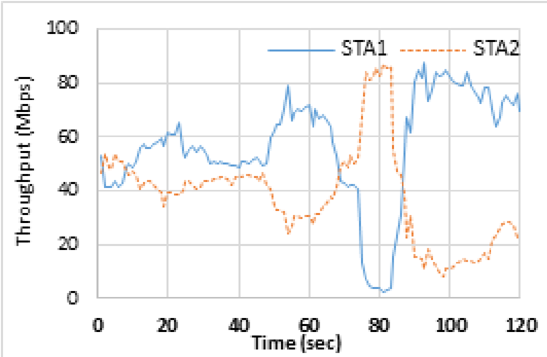

Figure 6. Time variation of TCP throughput when STA2 is located at Position0.

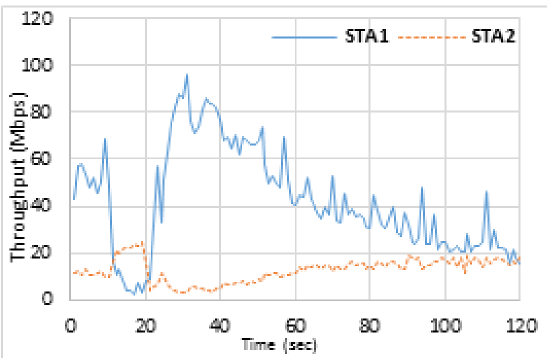

Figure 7. Time variation of TCP throughput when STA2 is located at Position4.

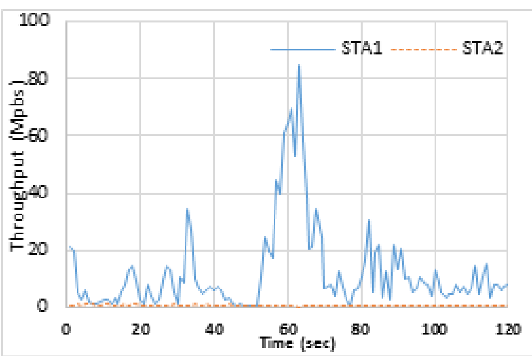

Figure 8. Time variation of TCP throughput when STA2 is located at Position7.

TCP throughput of STA1 is higher, and between 70 sec. and 90 sec., STA2 provides higher throughput. It can be said that two stations conflict for the WLAN channel rather fairly.

When STA2 is located at Position4, the TCP throughput of STA1 is 40 Mbps through 60 Mbps in the beginning, but it falls down in 10 sec. through 20 sec. It grows to 80 Mbps again, but after that, it decreases gradually to the value similar to STA2. In the case of STA2, the TCP throughput is less than 20 Mbps.

When STA2 is located at Position7, the TCP throughput of STA2 is very low and around 1 Mbps. The throughput of STA1 is lower than 20 Mbps in most of time, but is as high as 80 Mbps around 60 sec.

Throughout those experimental runs, the MAC level data rate is almost constant both for STA1 and STA2. On the other hand, the TCP throughput, especially that of the near station, fluctuates largely. When STA2 is located at Position7, the throughput of STA1 is lower than the other cases, but it has a time frame when the throughput is high.

In order to examine the reason of this fluctuation, Figs. 9 and 10 show the time variation of cwnd and RTT, when STA2 is located at Position4 and Position7, respectively. In Fig. 8, the cwnd of STA1 decreases sharply at 10 sec., goes to 1000 packets and then decreases gradually. The cwnd of STA2 is around 200 packets until 20 sec., but it decreases sharply and then it is going up gradually. These results show that the variation of cwnd is closely related with the variation of TCP throughput. On the other hand, the time variation of RTT has a frequent fluctuation but the value is smaller than 300 m sec. in most of time. Fig. 10 gives similar results. The graph of the time variation of cwnd is similar with that of the TCP throughput. The RTT values are also less than 300 m sec. in most of time. From those results, it can be decided the TCP throughput of STA1 is determined by the time variation of cwnd, which is decreased by packet losses.

In the previous paper [5], we showed that the performance anomaly occurs in UDP flows when the UDP traffic load to far located station is larger than its MAC level data rate. So, we calculated the estimated TCP load from equation (1) using the results given in Figs. 9 and 10. Figs. 11 and 12 show the time variation of the estimated TCP load and the data rate in STA2 when it is located at Position4 and Position7, respectively. Those results indicate that the estimated load is much smaller than the MAC level data rate. That is, the traffic

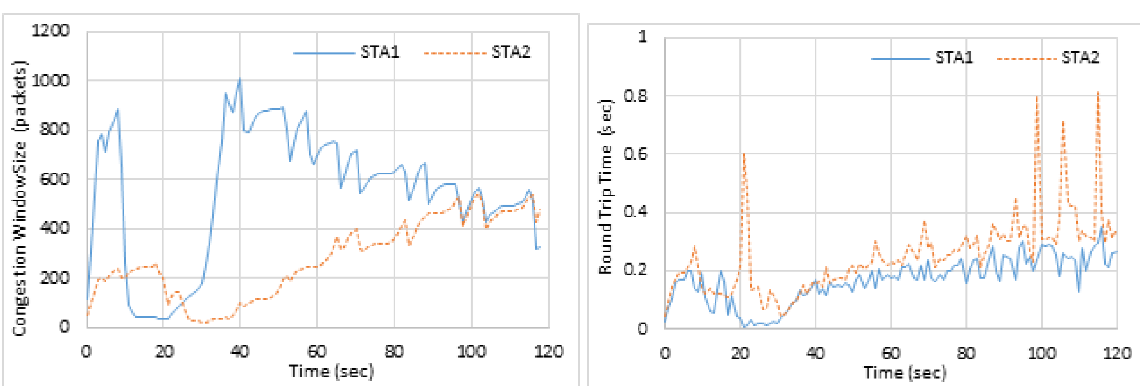

Figure 9. Time variation of cwnd and RTT when STA2 is located at Position4.

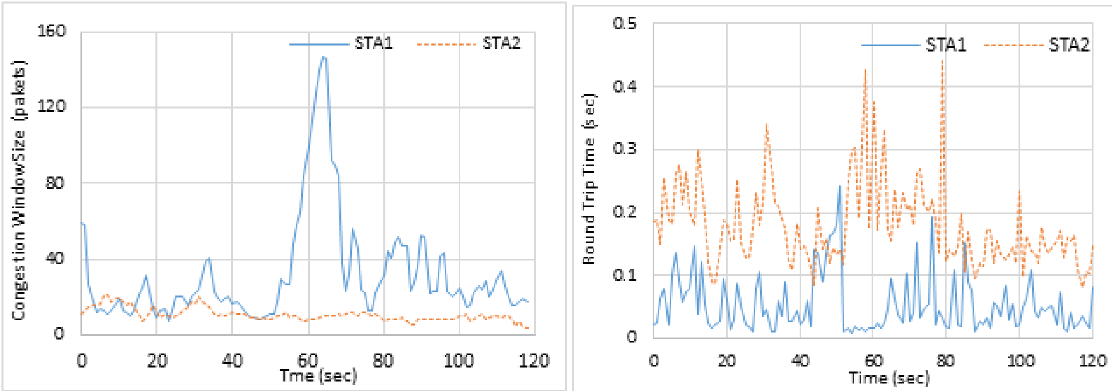

Figure 10. Time variation of cwnd and RTT when STA2 is located at Position7.

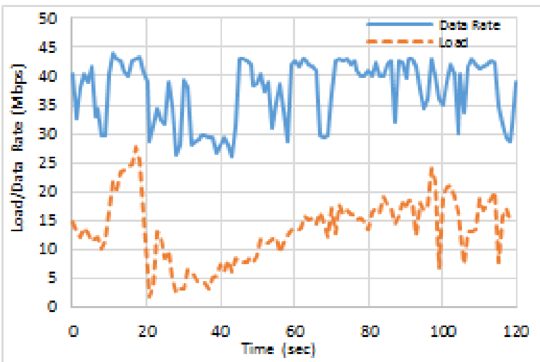

Figure 11. Time variation of estimated TCP load and data rate in STA2 when it is located at Position4.

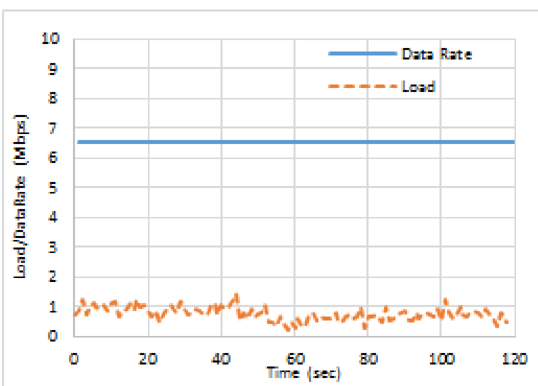

Figure 12. Time variation of estimated TCP load and data rate in STA2 when it is located at Position7.

load to STA2 is smaller than the transmission speed for STA2. Therefore, it can be concluded that, in the case of TCP download data transfer, the performance anomaly does not occur in a strict sense. That is, the throughput of a near station is not reduced to the similar value of a far located station. The reason the throughput is degraded is that the cwnd value in a near station is decreased due to packet losses.

So far, we used the FIFO queue management scheme at the access point. We also tried the other queue management schemes; CoDel, SFQ and FQ_CoDel. Although the cwnd values in STA1 and STA2 are smaller than the case in FIFO, we had a similar trend that the estimated TCP load to STA2 is much smaller than its MAC level data rate.

IV. Results for Four Station Experiment

Next, we conducted a similar experiment using four stations. As described in Section 2, we used the configuration where three stations, STA1 through STA3, are located at STA1's position so far and the other station, STA4, is located at one of positions “Position0” through “Position7.” Similarly with the two station experiment, we use the average data rate of STA4 to represent its position.

Fig. 13 shows the relationship between the position of STA4 and the average of MAC level data rate of STA1 through STA4. The average data rate of STA4 decreases along with its location being far away from the access point. On the other hand, the average data rate of STA1, STA2 and STA3 located near to the access point is rather high, such as 250 Mbps or higher, except that STA2 sometimes takes 150 Mbps or 200 Mbps average data rate. This is the result when the FIFO queue management scheme is used. The cases when the other schemes are used showed similar results. Similarly with the two station experiment, we use the average data rate of STA4 to represent its position.

Fig. 14 (a) shows the relationship between the STA4 average data rate and the average TCP throughput of all stations. In this experiment, the average TCP throughput of near stations did not decrease so much as the STA4 average data rate goes down. This is in contrast with the result of two station experiment. The average throughput of far station (STA4) decreases much more than the others along with it located far from the access point. This result clearly shows that the performance anomaly does not occur in the case of three near stations and one far station.

Fig. 14 (b) shows the average cwnd of all stations versus the STA4 average data rate. In the four station experiment, the cwnd values are similar among all the stations including STA4, although cwnd of STA4 is slightly reduced when its average data rate is smaller than 50 Mbps. The values are between 200 and 350 packets, which are smaller than the two station experiment, and this means that there are more packet losses and more cwnd drops than the two station case.

Fig. 14 (c) shows the average TCP level RTT of all stations versus the STA4 average data rate. The average RTT for STA1 through STA3 does not change so much even if the STA4 average data rate becomes small. On the other hand,

 ISBN: 978-1-61208-503-6

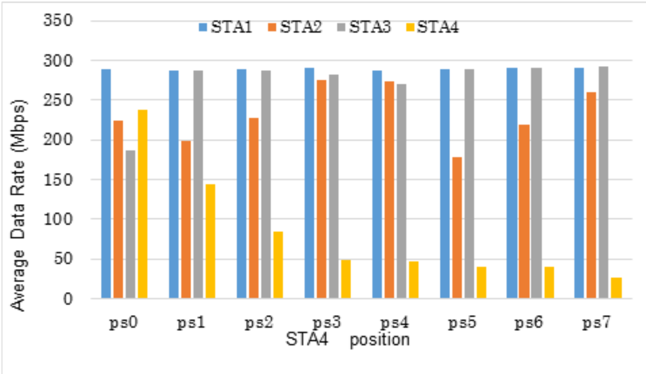

Figure 13. Average MAC level data rate vs. STA4 position (FIFO).

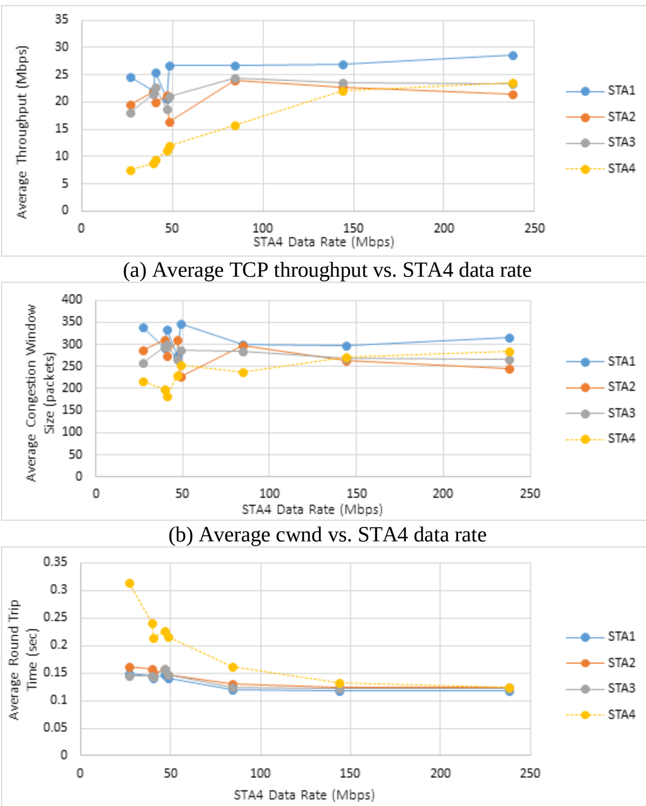

(a) Average TCP throughput vs. STA4 data rate

(b) Average cwnd vs. STA4 data rate

(c) Average TCP level RTT vs. STA4 data rate

Figure 14. Average TCP throughput, cwnd and RTT (FIFO).

that for STA4 increases when its average data rate becomes small. This increase of RTT is considered as the reason of the average TCP throughput of STA4 being decreased in the small STA4 data rate.

In summary, in the situation where three near stations and one far station share an access point, the TCP download data transfer in the near stations is not influenced by that of the far station. The impact of far station is smaller than the situation where one near station and one far station exist. This is different from the results stated in [4].

V. CONCLUSIONS

This paper discussed the performance anomaly of TCP flows providing download data transfer over IEEE 802.11n WLAN. We conducted two station and four station experiments, where one of the stations is located far from the access point. We used a commercially available access point using four queuing management schemes in it. We measured several performance metrics including the MAC level data rate, the number of MPDUs aggregated in an A-MPDU, the TCP throughput, the congestion window size, and the TCP level round trip time.

In the two station experiment, the TCP throughput of the near station is degraded along with the distance between the far station and the access point increasing. However, we conclude that this is not the performance anomaly in a strict sense. The reason is that the estimated TCP load to the far station is much lower than its MAC level data rate. Instead, we conclude that the throughput degradation comes from the decrease of the congestion window size caused by packet losses.

In the case of four station experiment, the degradation of the TCP throughput of near stations was smaller than the two station experiment. This means that the data transfer to the far station gives only a small impact to the near stations. This will be also the evidence that the performance anomaly does not occur in the TCP flows.

As for the queue management scheme, there was only a little influence in terms of TCP throughput. We also conducted an experiment using the native firmware of the access point and obtained the similar results.

REFERENCES

[1] IEEE Standard for Information technology: Local and metropolitan area networks Part 11: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications, 2012.

[2] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, “Performance Anomaly of 802.11b,” Proc. INFOCOM 2003, vol.2, pp.836-843, Mar. 2003.

[3] M. Abusubaih, “On Performance Anomaly in 802.11 Wireless LANs: Problem and Solution Approaches,” Proc. Next Generation Mobile Applications, Services and Technologies (NGMSAT) 2010, pp.208-212, Jul. 2010.

[4] O. Abu-Sharkh and M. Abdelhadi, “The impact of multi-rate operation on A-MSDU, A-MPDU and block acknowledgment in greenfield IEEE802.11n wireless LANs,” Proc. Wireless Advanced (WiAd) 2011, pp. 116-121, Jun. 2011.

[5] Yoshiki Hashimoto, Masataka Nomoto, Celimuge Wu, Satoshi Ohzahata, and Toshihiko Kato, “Experimental Analysis on Performance Anomaly for Download Data Transfer at IEEE 802.11n Wireless LAN,” Proc. International Conference on Networks (ICN) 2016, pp.22-27, Feb. 2016.

[6] “Open Wrt Wireless Freedom,” https://www.openwrt.org/, retrieved: Feb. 2016.

[7] K. Nichols and V. Jacobson, “Controlling Queue Delay,” ACM Queue, Networks, vol.10, no.5, pp. 1-15, May 2012.

[8] J. Gettys and K. Nichols, “Bufferbloat: Dark Buffers in the Internet,” ACM Queue, Virtualization, vol. 9, no.11, pp. 1-15, Nov. 2011.

[9] iperf, http://iperf.sourceforge.net/, retrieved: Feb. 2016.

[10] ath9k Linux Wireless, http://wireless.kernel.org/en/users/Drivers/ath9k, retrieved: Feb. 2016.

[11] Linux foundation: tcpprobe, http://www.linuxfoundation.org /collaborate/workgroups/networking/ tcpprobe, retrieved: Feb. 2016.