Experimental Analysis on Performance Anomaly for Download Data Transfer at 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—**It is reported that, even in IEEE 802.11n wireless local area network (WLAN), the performance anomaly occurs which reduces the throughput of all stations when there are some stations communicating with low data rate. But, the previous paper uses a manually configured environment in the Transmission Control Protocol (TCP) uplink performance evaluation. In this paper, we show the results of experiments indicating the performance anomaly at downlink User Datagram Protocol (UDP) and TCP data transfer over 802.11n WLAN. In the experiment, we adopted the actual parameter settings in the stations and the access point, and carefully examined the relationship of the frame aggregation, the queue management at the access point, UDP traffic load, and TCP congestion window size with the performance anomaly. We show that a phenomenon like the performance definitely occurs in both the UDP and TCP data transfer, but the reasons seem to be different for each of them.**

Keywords- WLAN; IEEE802.11n; Performance Anomaly, Queue Management.

I. Introduction

Resulting from the wide deployment of WLANs based on IEEE 802.11 standard [1], a varieties of WLAN environments including home, office and public hot spots are used by a lot of terminals such as smart phones, tablets and notebooks. When a number of stations access to one WLAN access point, they suffer from several kinds of performance problems. The *performance anomaly* [2][3] is a typical one among those problems.

It is a problem such that: 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 following two reasons. First, the IEEE 802.11 WLAN is based on the carrier sense multiple access with collision avoidance (CSMA/CA) principle, which tries to assign fair chances to send data frames among all the stations. The other reason is that 802.11 WLAN provides multiple Media Access Control (MAC) level data rates. So, a station with low bit rate captures the channel for a long time, and it penalizes other hosts with higher data rates.

The IEEE 802.11n [1] WLAN introduces several enhancements to the legacy standards such as 802.11a, b and g. Among them, supporting higher data rates (e.g., 150 Mbps), the Aggregated MAC Protocol Data Unit (A-MPDU) and the Block Acknowledgment mechanism provide high performance data transfer.

In spite of the drastic improvement of the data transfer performance, 802.11n standard does not resolve the performance anomaly problem. Abu-Sharkh and Abdelhadi [4] reported that the performance anomaly still exists in 802.11n WLAN. The report [4] describes the results of experiment where four wireless stations perform uplink TCP data transfer. Among the stations, three are located near the access point and one is located far from it. As a result, the performance of near three stations is affected by the far station. The result also shows that the aggregation is closely related with the performance. If the aggregation is used for both near and far stations, the throughput of near stations becomes the same as that of far station. If the aggregation (four MPDUs in an A-MPDU) is used only for near stations, the throughput is four times higher for near stations than far stations.

Although the report [4] mentions the performance anomaly in 802.11n, it uses a manually configured experimental environment, such as a controlled aggregation scheme. In this paper, we show the experimental results of performance anomaly for downlink data transfer in an actual WLAN environment. The feature of our experiment is as follows.

- Both TCP and UDP data transmissions are examined.
- The aggregation scheme in the access point is used as it is.
- The other parameters such as the queue management scheme at the access point are considered.
- For the UDP data transmission, various traffic loads are examined. For the TCP data transmission, the congestion window size is examined in detail together with MAC level data rate.

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 UDP and TCP experiments, respectively. In the end, Section 5 gives the conclusions of this paper.

II. Experimental Settings

Figure 1 shows the configuration of our experiment. Two stations (STAs) conforming to 802.11n with 5GHz band are associated with one access point, which is connected a server through 1Gbps Ethernet. One STA (STA1) is located at a near position to the access point, and the other STA (STA2) is located in various positions in the experiment.

The detailed specifications of STAs and the access point are shown in Table 1. We use commercially available notebooks and access point in the experiment. As for the

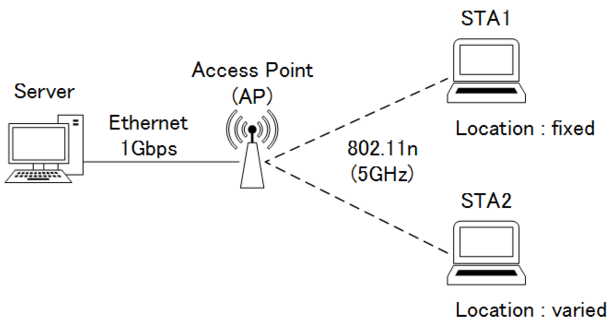

Figure 1. Configuration of experiment

TABLE 1. 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

access point, we do not use the software originally installed by the vendor but that provided by the OpenWRT project [5]. By using this software, we can obtain the performance metrics in the access point such as the MAC level data rate, the number of A-MPDUs sent, and the number of MPDUs aggregated in an A-MPDU.

We can also configure the queue management schemes used in the access point. The OpenWRT firmware supports the following schemes.

- *FIFO*: A scheme to use one queue to store all packets being sent by the access point, independently of flows the packets belong to.
- *CoDel* [6]: An active queue management scheme designed to resolve the Bufferbloat problem [7]. It uses packet-sojourn time in a queue as a control parameter, and drops a packet in the situation that packets stay in the queue too long.
- *Stochastic Fare Queueing* (*SFQ*): A scheme to provide a separate queue for packets of an individual flow. When sending packets, each queue is examined in the round robin scheduling, which avoids a large delay of packets against an aggressive flow.
- *FQ_CoDel*: A scheme which combines SFQ and CoDel. A queue is prepared for an individual flow and the delay within one queue is controlled by CoDel scheme. 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 transmitted from the server to two STAs through the access point (downlink data transfer in the WLAN). The server uses *iperf* tool [8] to generate UDP and TCP data flows and to measure their performance such as throughput. As for parameter settings for UDP and TCP, we used the native ones in the Linux operating system. Specifically, the TCP version is CUBIC TCP and the TCP small queues mechanism is used.

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

- packet trace at the server and the STAs, by use of *tcpdump*,
- WLAN related metrics, such as the MAC level data rate, the number of A-MPDUs sent and the number of MPDUs per A-MPDU, from the device driver at the access point and the stations,
- the throughput and packet loss rate in UDP data transmission, by iperf, and
- TCP connection information such as the congestion window size at the server, by use of *tcpprobe* [9].

The experiment was conducted in a building constructed with reinforced concrete. Figure 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. The black circles named "AP" and "STA1" correspond to the positions of the access point and STA1, respectively. These are fixed throughout the experiment. The circles named "Position0" through "Position7" represent the positions of STA2. STA2 is located in one of these eight positons in the experiment.

III. RESULTS FOR UDP DATA TRANSMISSION

A. Experimental Scheme

In the experiment evaluating the performance anomaly using UDP data transmission, we changed the UDP traffic load. For STA1, the server generates UDP datagrams at 100 Mbps, and for STA2, UDP datagrams whose traffic load is 10, 40, 70 and 100 Mbps are generated by the server. In the traffic,

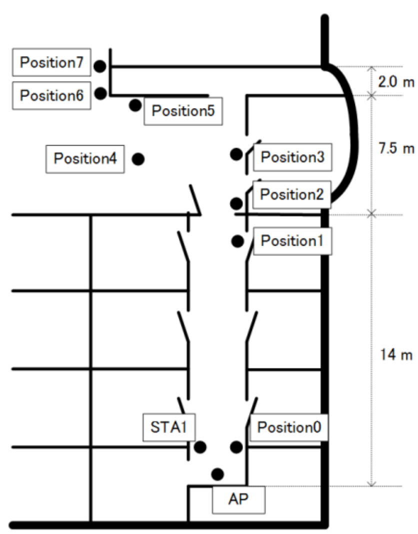

Figure 2. Layout inside building and position of equipment

 ISBN: 978-1-61208-450-3

the size of UDP datagram user data is set to 1472 bytes, which makes the size of IP datagrams 1500 byte.

Eight positions for STA2 and four queue management schemes at the access point are examined. The UDP traffic is generated for 60 seconds in one experiment run, and three runs are examined for one parameter setting.

B. *Results*

Figure 3 shows the relationship between the position of STA2 and the MAC level data rate (the average during one experiment run) in STA1 and STA2. STA1 located near the access point keeps high data rate such as 200 Mbps through 300 Mbps. On the contrary, the data rate of STA2 decreases along with its position being far away from the access point. More specifically, the data rate of STA2 remains a similar value for Position3 through Position6, and decreases at Position7. The data rate at Position7 has a larger variance than that of the other positions. Figure 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 experiment setup described above, we measured the throughput and the number of MPDUs per A-MPDU for STA1 and STA2, by changing the UDP traffic load of STA2, the position of STA2, and the queue management schemes in the access point. The measured value is the average through three experiment run.

As for the queue management schemes, the results for FIFO and CoDel, and that for FQ_CoDel and SFQ showed similar trends, respectively. It is considered that the packet losses invoked by CoDel do not give any influence to the throughput and the MAC level aggregation behavior in the UDP data transmission. In this section, the results with FIFO and FQ_CoDel schemes are given.

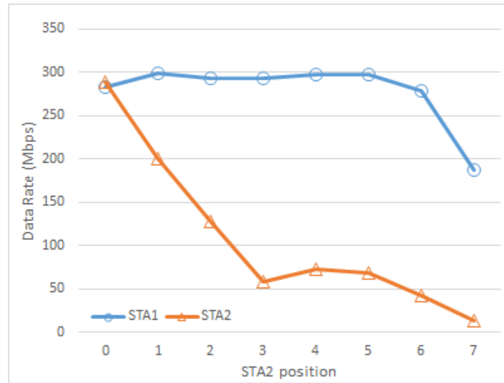

Figure 3. Average MAC level data rate vs. STA2 position (FIFO)

Figure 4 (a) shows the UDP throughput of STA1 and STA2 when the UDP traffic load to STA2 is changed from 10 Mbps to 100 Mbps. (Remember that the UDP traffic load to STA1 is 100 Mbps.) The queue management scheme at the access point is FIFO. The graph shows the results with STA2 located at Position0 through Position7. The solid line is the result of STA1 and the dashed line is for STA2. The color of line discriminates the STA2 position.

When STA2 is located at Position0 where two station can communicate with high MAC level data rate, the UDP

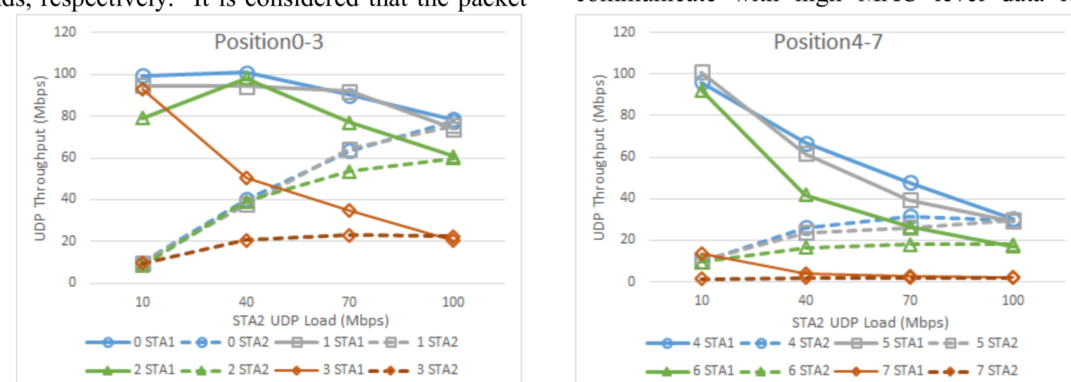

(a) UDP throughput of STA1 and STA2 vs. UDP traffic load to STA2 (FIFO)

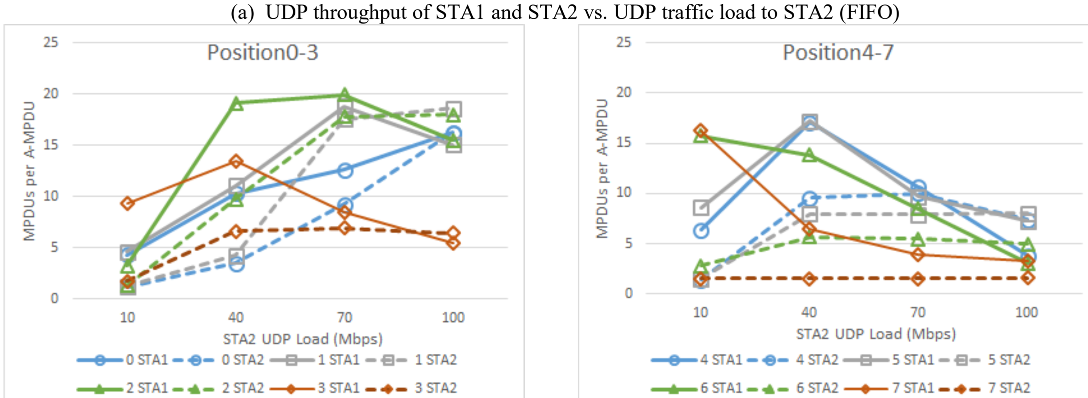

(b) Number of MPDUs per A-MPDU vs. UDP traffic load to STA2 (FIFO)

Figure 4. Results for UDP data transmission in FIFO queue management scheme

throughput of STA1 is as high as 80 Mbps. STA2 also provides 80 Mbps throughput at this point. According to locating STA2 far from the access point, the UDP throughput drops. This is because the MAC level data rate used by STA2 becomes low. The drop is larger for higher MAC level data rate. It should be noted that the UDP throughput drop indicates that there several losses of UDP datagrams.

Although the location of STA1 is fixed near the access point and the MAC level data rate STA1 uses is high, the UDP throughput of STA1 becomes lower as STA2 is located far from the access point. When the UDP traffic load to STA2 is 100 Mbps, which is the same as STA1, the UDP throughput is the same for STA1 and STA2. The reason is considered that the low MAC level data rate of STA2 occupies the WLAN channel and the chance of STA1 to transmit data frames will be the same as STA2. This is the performance anomaly.

Figure 4 (b) shows the number of MPDUs per A-MPDU of STA1 and STA2 with the UDP traffic load to STA2 changed from 10 Mbps to 100 Mbps. The queue management scheme in the access point is FIFO. When STA2 is located at Position0, the number of MPDUs per A-MPDU becomes large in both STA1 and STA2, as the UDP traffic load to STA2 increases. This is closely related with the packets queued in the buffer of the WLAN device. At Position0, the MAC level data rate is high for STA1 and STA2. So, while the UDP traffic load to STA2 is low, the buffer of the WLAN device does not contain many data frames to send since the MAC level data rate is high compared with the UDP traffic load. This situation decreases the number of MPDUs aggregated in an A-MPDU. As the UDP traffic load to STA2 increases, the number of MPDUs per A-MPDU increases, and at the STA2 UDP load of 100 Mbps, the numbers for STA1 and STA2 become the same.

In the case that STA2 is located at Position1 and Position2, the number of MPDUs per A-MPDU of STA1 and STA2 increases along with the UDP traffic load to STA2, until the load is less than and equal to 70 Mbps. When the UDP traffic load to STA2 is more than 70 Mbps, it decreases. As the location of STA2 becomes far from the access point, the number of MPDUs per A-MPDU come to decrease at a lower UDP traffic load to STA2. We infer that this is caused by the aggregating behavior in the WLAN device driver Ath9k [10]. The device driver aggregates MPDUs under the requirement that the transmission time of A-MPDU is below 4 msec. When STA2 is located far from the access point and the MAC level data rate of STA2 decreases, this requirement limits the number of MPDUs aggregated in an A-MPDU in STA2. For STA1, the MAC level transmission rate is equivalently decreased because of the performance anomaly, the number of MPDUs per A-MPDU becomes low.

Figure 5 (a) shows the relationship between the UDP traffic load to STA2, and the UDP throughput of STA1 and STA2, when the access point uses the FQ_CoDel queue management scheme. In this case, similarly to FIFO, the UDP throughput of STA1 becomes lower as the UDP traffic load to STA2 increases. We can say that the performance anomaly also occurs in this experiment.

In contrast to FIFO, where the UDP throughput of STA1 decreases slowly as STA2 is located at Position3 through

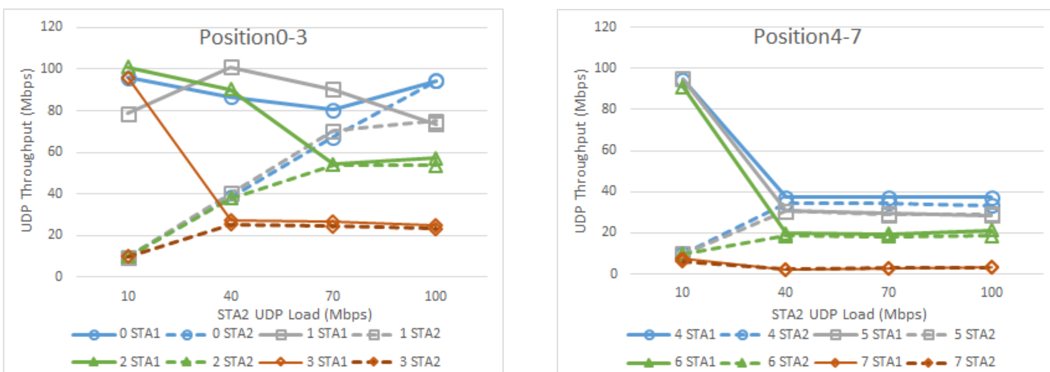

(a) UDP throughput of STA1 and STA2 vs. UDP traffic load to STA2 (FQ_CoDel)

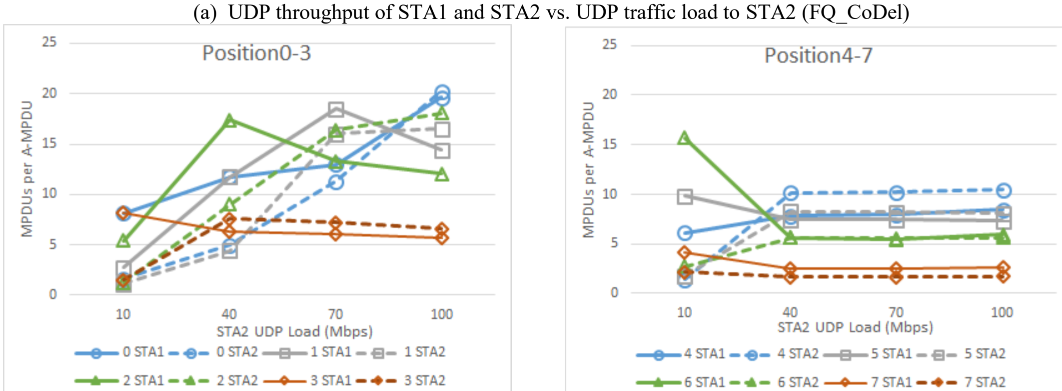

(b) Number of MPDUs per A-MPDU vs. UDP traffic load to STA2 (FQ_CoDel)

Figure 5. Results for UDP data transmission in FQ_CoDel queue management scheme

 ISBN: 978-1-61208-450-3

Position7, FQ_CoDel introduces a sharp drop in the UDP throughput of STA1 with UDP traffic load to STA2 larger than 40 Mbps. Figure 5 (b) shows that the number of MPDUs per A-MPDU at STA1 becomes the same as STA2 in this range of UDP traffic load to STA2.

C. Discussions

(1) In the download UDP data transmission over 802.11n WLAN, the performance anomaly surely happens. When the traffic load to a far located station becomes larger than the MAC level data rate which the station uses, it affects the throughput of a near located station. In the experiment, the UDP traffic load to STA2 of 10 Mbps does not invoke the performance anomaly except that STA2 is located at Position7. This is because the MAC level data rate of STA2 is more than 10 Mbps at Position 0 through Position6. The data rate at Position7 is less than 10 Mbps, and this case caused the performance degradation invoked by the performance anomaly.

(2) The UDP throughput also depends on the number of MPDUs aggregated in an A-MPDU. This number depends on the MAC level data rate and the UDP traffic load to the far station. It should be noted that the variation of the number of MPDUs per A-MPDU was large. For example, there was a case where the number of MPDUs changes from 1 to 32. But, the average UDP throughput is determined by the average number of MPDUs per A-MPDU.

(3) The results were affected by the queue management scheme used by the access point. In the case that the access point uses FIFO and CoDel, the data to be sent to both the far and near stations are stored in one queue. So, while the UDP traffic load to the far station is low, more data to the near station are stored in the queue. This makes the number of MPDUs in the near station larger than that of far station. As the UDP traffic load to the far station increases, the number of aggregated MPDUs and the throughput of the far station increases. On the contrary, FQ_CoDel and SFQ maintain separate queues for individual traffic flows. The aggregation is performed queue by queue basis, and so, even in the situation where the performance anomaly occurs, e.g., when the UDP traffic load to STA2 is more than 40 Mbps, the aggregation, and the UDP throughput were similar for the far and near stations.

IV. RESULTS FOR TCP DATA TRANSMISSION

A. Experimental Scheme

Similarly with the experiment for UDP data transmission, we measured the performance of TCP data transmission form 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, and obtained the averages. We measured the MAC level data rate at the access point, the TCP throughput and the number of MPDUs aggregated in an A-MPDU at the receivers (stations), and the TCP congestion window size (cwnd) and the TCP level round trip time (RTT) at the server.

B. Results

As for the MAC level data rate, the similar results were obtained as those in the UDP experiment depicted in Figure 3. STA1 keeps high data rate such as 250 Mbps. The data rate of STA2 decreases as its position is far away from the access point.

Figure 6 (a) shows relationship between the STA2 position and the average TCP throughput. In this graph, solid lines indicate the results of STA1 and dashed lines indicate those of STA2, and the color of lines correspond to the 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. The results seem to be the performance anomaly. In the case of TCP, there was no difference among the queue management schemes in the access point.

Figure 6 (b) shows the relationship between the position of STA2 and the number of MPDUs per A-MPDU. Here, the number of aggregated MPDUs of STA1 and STA2 goes down as the STA2 position becomes far from the access point. Similarly with the TCP throughput, there was no difference among the queue schemes.

Figure 7 (a) shows the average cwnd versus the STA2 position. In this case, the results largely depend on the queue management scheme. In FIFO, the average cwnd is large and it 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

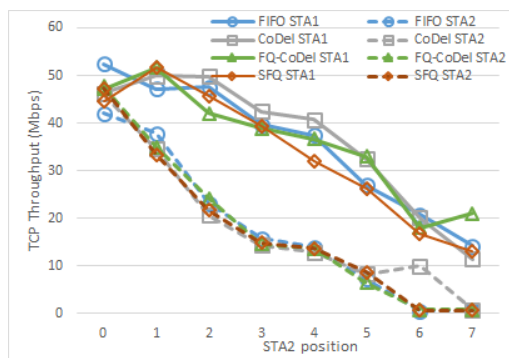

(a) TCP throughput vs. STA2 position

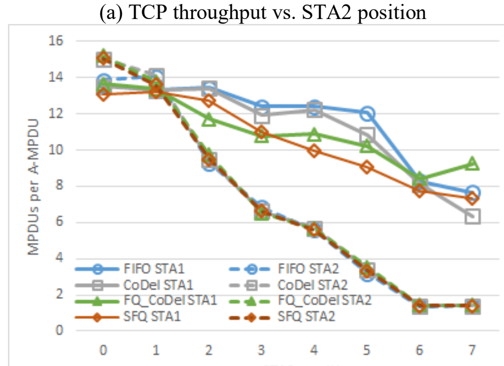

(b) Number of MPDUs per A-MPDU vs. STA2 position

Figure 6. Results of TCP throughput and MPDUs per A-MPDU

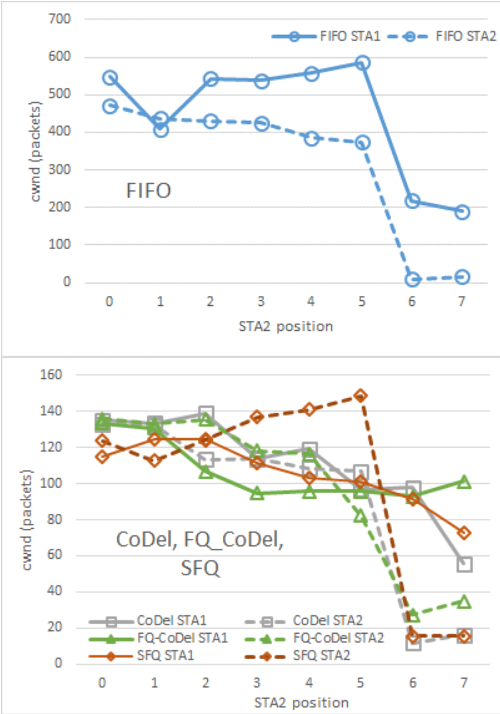

(a) cwnd vs. STA2 position

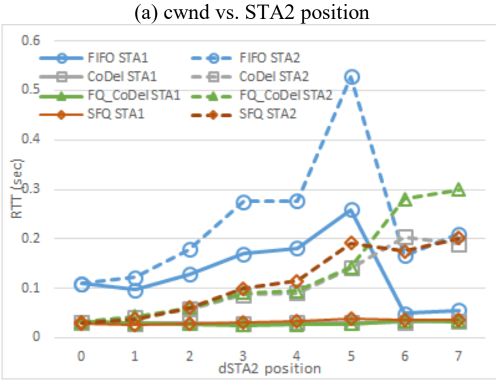

(b) RTT vs. STA2 position

Figure 7. Results of cwnd and TCP level RTT

situation, there are a lot of packet losses in both STA1 and STA2 TCP flows.

Figure 7 (b) shows the average TCP level RTT versus the STA2 position. Here, the results also depend on the queue management scheme. Especially, FIFO has a large RTT compared with 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.

C. *Discussions*

(1) In the download TCP data transmission over 802.11n WLAN, the phenomenon similar to the performance anomaly occurs. The throughput of a station located near the access point becomes lower when a far located station communicates. However, the results in Figure 6 (a) show that the throughput of a far located station, STA2, does not exceed the MAC level data rate of STA2. So, it is considered that the reason of the performance anomaly like phenomenon in the TCP data transmission is different from that of the UDP data transmission.

(2) One possible reason is the decrease of cwnd caused by packet losses. As shown in Figure 7 (a), the average cwnd of STA1 decreases as the position of STA2 moves far from the access point. This means that the performance degradation in STA2 invokes the packet losses in the STA1 communication. This invokes the performance anomaly like phenomenon.

V. CONCLUSIONS

This paper discussed the performance anomaly of UDP and TCP download data transmissions over IEEE 802.11 WLAN. It showed that the performance anomaly problem surely happens for the UDP data transmission. In the situation where the problem occurs, the number of MPDU aggregation of a near located station is also decreased, and this reinforce the performance degradation. It should be noted, however, that this experiment is done in an artificial condition in which an excessive UDP traffic load is applied. As a result, a large number of packets are lost. In an actual communication, such a packet loss is not acceptable. In other words, for UDP, the performance anomaly problem happens only in the situation where excessive data transfer requests are added.

As for TCP data transmission, the phenomenon similar to the performance anomaly occurs. But, it is considered that the reason is different from that in UDP. The throughput (traffic load) is smaller than the MAC level data rate, and instead, the degradation of the congestion window size caused by packet losses decreases the throughput. More investigation is necessary for the performance analysis of TCP data transmission.

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, Mar. 2003, pp.836-843.

[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, Jul. 2010, pp.208-212.

[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, Jun. 2011, pp. 116-121.

[5] "Open Wrt Wireless Freedom," https://www.openwrt.org/, retrieved: Jan. 2016.

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

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

[8] iperf, http://iperf.sourceforge.net/, retrieved: Jan. 2016.

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

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