TCP Throughput over Full-Duplex WLANs: Novel Implications of the AP's New Capability

Vinay U. Pai*, Neelesh B. Mehta†, Chandramani Singh‡
Indian Institute of Science, Bangalore, India
Email: *vinaypu@iisc.ac.in, †nbmehta@iisc.ac.in, ‡chandra@iisc.ac.in

ICC 2024 - IEEE International Conference on Communications | 978-1-7281-9054-9/24/$31.00 ©2024 IEEE | DOI: 10.1109/ICC51166.2024.10623126

***Abstract*—Full-duplex (FD) communication promises to double the throughput of wireless local area networks (WLANs) by allowing simultaneous transmission and reception of data. While the benefits of FD on the medium access control (MAC) layer throughput of IEEE 802.11 WLANs are well-studied, the interaction between the transmission control protocol (TCP) and the FD MAC layer is less explored. We consider TCP file uploads and downloads between stations (STAs) and a server via an FD access point (AP). Using a novel FD-specific saturation approximation, Markov renewal theory, and fixed point analysis, we derive novel expressions for the TCP upload and download saturation throughputs. These expressions differ from those derived in the literature for TCP over half-duplex (HD) WLANs, and bring out how the AP is no longer a bottleneck. Despite greater contention between STAs, cross-link interference between transmitting and receiving STAs, and asymmetric payloads due to the different sizes of TCP data packets and TCP acknowledgments, we find that an FD WLAN achieves a significantly higher TCP throughput than a conventional HD WLAN.**

I. Introduction

IEEE 802.11 wireless local area networks (WLANs) have been ubiquitously deployed in commercial and residential settings. Physical (PHY) layer technologies such as multiple-input multiple-output and orthogonal frequency division multiple access, and enhancements to the contention-based medium access control (MAC) protocol such as frame aggregation and spatial reuse have been introduced over multiple generations to improve the WLAN data rates and efficiency. However, the MAC protocol is still designed for a half-duplex (HD) PHY, in which an access point (AP) or station (STA) can transmit or receive data, but cannot do both simultaneously.

Full-duplex (FD) communication is a promising new PHY technology in which signals can be transmitted and received simultaneously over the same frequency band. This promises to double the throughput and spectral efficiency of conventional HD wireless systems. Sophisticated analog and digital signal processing innovations to reduce self-interference have made FD practically feasible. For this reason, a variant of FD is being considered for Release 18 5G-Advanced cellular standards [1]. FD has also attracted considerable attention in the 802.11 WLAN literature [2]–[5], although hardware implementations have been limited to experimental prototypes in laboratory environments. These works focus on handling the cross-link interference (CLI) between the uplink (UL) STAs, which transmit to the AP, and downlink (DL) STAs, which receive from the AP.

However, the interaction between the MAC layer that employs the IEEE 802.11 carrier sense multiple access with collision avoidance (CSMA/CA) protocol to let the STAs and the AP contend for medium access, and the higher layers, such as the transport layer that uses the transmission control protocol (TCP), in 802.11 FD WLANs has not yet been studied. Delays suffered by the TCP acknowledgments (ACKs) at the AP or the STAs due to contention at the MAC layer can slow the TCP congestion window growth and degrade the overall throughput. Moreover, TCP uses short ACK packets to acknowledge the reception of the longer TCP data packets. However, the MAC treats the TCP data and TCP ACK packets equally, which leads to asymmetries in the UL and DL packet durations. This asymmetry affects the FD MAC layer throughput and, thus, the TCP throughput.

A. Related Literature

We summarize the related literature on TCP over WLANs or FD networks.

TCP over HD 802.11 WLANs: Reference [6] studied the interaction between the TCP protocol and the 802.11 contention-based MAC protocol when the STAs download TCP files from a remote server via the AP. It showed that the AP was the bottleneck in transmitting TCP data packets because it had to contend with the STAs to transmit on the channel. This formed the basis for the TCP throughput analyses in [7]–[10], which analyzed the upload and download throughput [7], the effect of transmission errors [8], delayed ACKs [9], and the joint impact of channel errors and a finite buffer at the AP [10].

TCP over FD Networks: To the best of our knowledge, no works in the existing literature analyze TCP throughput in 802.11 FD WLANs. In cellular networks, [11] showed that the TCP congestion window grew faster when the base station and the user equipment were FD-capable. This enabled FD to outperform HD in both throughput and delay, except in scenarios with significant inter-cell interference.

B. Contributions

We consider a system in which a fraction of the STAs upload TCP data while the rest download TCP data from a remote server that is connected to an AP. The STAs communicate with the AP over a wireless channel. We consider an asymmetric FD WLAN, in which only the AP is FD-capable since FD is likely to be first implemented in the more expensive APs than the cheaper STAs. We make the following contributions:

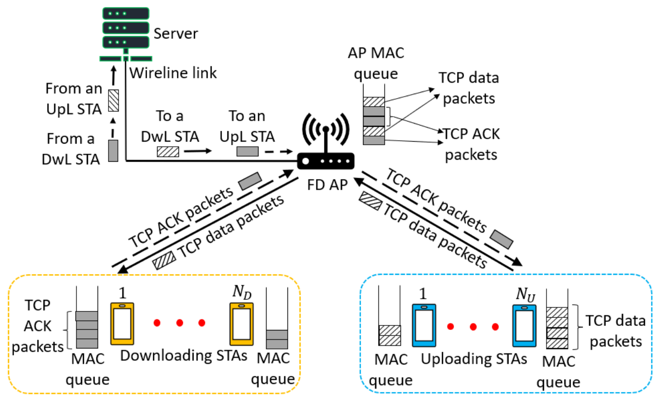

Fig. 1. System model showing a server connected to an FD AP via a wireline link, and uploading and downloading STAs. Also shown is the flow of TCP data packets and TCP ACK packets between the server and the STAs.

1) We derive expressions for the total TCP upload and download throughput in an 802.11 FD WLAN where the AP is FD-capable and can simultaneously transmit to an STA while receiving from another STA. We use mean-field analysis and renewal theory [12] and introduce a novel, FD-specific saturation approximation that is based on the insight that an FD AP is no longer a bottleneck because of the additional transmit opportunities that it can exploit. This leads to novel expressions for the AP's and STAs' attempt and collision probabilities, which differ from those derived for an HD WLAN in [6]–[8], [10].
2) We consider a novel two-rate model that incorporates the impact of CLI in FD WLANs. During FD, the AP uses a CLI-limited PHY layer data rate r_2 for the DL transmission, which is lower than the PHY layer data rate r_1 used by the UL STA. The different rates combined with the different payloads of the TCP data and TCP ACK packets lead to different UL and DL packet transmission durations. Our model captures the MAC layer inefficiencies that arise due to this asymmetry.
3) Our simulation results show an FD WLAN has a much higher TCP throughput than an HD WLAN. However, the gains are less than 100% due to increased contention and CLI between STAs.

C. Outline and Notation

In Section II, we present the FD MAC protocol and the TCP model. We analyze the TCP throughput in Section III. Numerical results are presented in Section IV, followed by our conclusions in Section V.

Notation: The expectation of a random variable (RV) X is denoted by $\mathbb{E}[X]$. The probability of an event E is denoted by $\mathrm{P}(E)$, and the conditional probability of an event E given an event F is denoted by $\mathrm{P}(E|F)$.

II. System Model

We consider an infrastructure 802.11 WLAN with N_D downloading (DwL) STAs, each of which is downloading a large TCP file from the server, and N_U uploading (UpL) STAs, each of which is uploading a large TCP file to the server. Let $N = N_U + N_D$. Each STA has a single long-lived TCP connection with the server. The server is connected to the AP via a high-speed wireline link and communicates with the STAs via the AP. The system model is illustrated in Figure 1.

The server breaks the TCP file intended for a DwL STA into data packets. Each data packet is of length L_d bits. It sends them over the wireline link to the AP's MAC queue according to the TCP congestion control protocol described in Section II-A. The AP then sends these packets to the DwL STA over the 802.11 FD WLAN. Each data packet received by a DwL STA generates a TCP ACK packet of length L_a bits that resides in the STA's MAC queue. The STA transmits the ACK packet to the AP over the WLAN MAC, which sends it to the server over the wireline link.

Similarly, an UpL STA breaks the TCP file to be uploaded into data packets, each of length L_d bits. These are injected into the UpL STA's MAC queue as per the TCP congestion control protocol. The STA sends these packets over the WLAN to the AP, which forwards them to the server. The server generates a TCP ACK packet of length L_a bits in response, which its sends to the UpL STA via the AP over the WLAN.

Thus, a DwL STA receives TCP data packets from the AP and transmits TCP ACK packets to the AP. On the other hand, an UpL STA transmits TCP data packets to the AP and receives TCP ACK packets from the AP.

A. TCP Congestion Control Model

We make the following assumptions, which are typical of most analyses of TCP over 802.11 WLANs [8]–[10]:

1) There are no TCP packet losses due to duplicate ACKs or retransmission timeouts [9]. The AP's MAC queue is large enough to ensure there is no buffer overflow [8], [9]. Thus, the TCP congestion window of each connection grows to the maximum value $w_{\max}$ and stays there.
2) Each TCP data packet generates one TCP ACK packet at the receiver [8]. Conversely, each TCP ACK packet acknowledges the successful reception of one TCP packet.[1]
3) There is a negligible delay between the server and the AP [10], [13].

This model applies to TCP NewReno [14] and TCP CUBIC [15] because the TCP connections in them are always in the congestion avoidance phase in the absence of congestion events.

B. 802.11 FD WLAN: Pertinent Aspects and Notation

Time is discretized into slots of duration δ. Each node that has a packet to transmit uniformly samples a back-off timer value in the range $\{0, \ldots, C-1\}$, where C is the node's MAC contention window (CW) size. Each node decrements its back-off timer value by one in a slot if it senses the channel to be idle. Instead, the node freezes its back-off timer if it senses a

[1]In practice, the receiver can send one TCP ACK for several TCP data packets to reduce the protocol overhead. Analyzing this delayed ACK mechanism is beyond the scope of this work.

transmission. It resumes decrementing its back-off timer after a distributed coordinate function inter-frame space (DIFS) of duration T_{DIFS} from the time it senses the channel to be idle.

The request-to-send/clear-to-send (RTS/CTS) access mechanism is used. Thus, the node transmits an RTS packet of duration T_{RTS} when its back-off timer value reaches 0. We assume that rate r_1 is used for all RTS packets and that an STA will not be able to decode the AP's RTS packet under CLI. This is consistent with the collision model used in the WLAN literature [12]. In the FD MAC, the following scenarios arise depending on which nodes transmit RTS packets.

a) STA-initiated FD Transmission: When an STA, say A, transmits an RTS packet to the AP, the AP responds by sending a CTS packet of duration T_{CTS} after a short inter-frame space (SIFS) of duration T_{SIFS}. STA A starts its packet transmission to the AP with rate r_1 after T_{SIFS} seconds from the end of the CTS packet. The FD AP simultaneously transmits its packet to another STA, say B, that it has a packet for, using rate r_2.[2] With high probability, the head-of-line (HoL) packet in the AP's MAC queue will not be for STA A. The FD AP transmits a packet to an STA in the DL without needing a CTS packet while receiving a packet in the UL from another STA. This can be implemented, for example, using the MAC layer enhancement proposed in [16].

The packet transmission durations of STA A and the AP depend on whether they transmit TCP data or TCP ACK packets. Since the TCP ACKs have a much smaller payload than TCP data packets, we assume that $\frac{L_d}{r_1} > \frac{L_a}{r_2}$. The AP and STA B send a MAC layer ACK packet of duration T_{ACK}, a T_{SIFS} duration after the channel becomes idle, i.e., from the end of the longer transmission. We call STA A and the AP the primary transmitter and the secondary transmitter, respectively. They both reset their CW size to $C_{\min}$, which is the minimum CW size. This scenario is illustrated in Figure 2. A success occurs for both AP and STA A.

However, if two or more STAs transmit RTS packets with rate r_1 at the same time, the AP cannot decode any of them. The colliding STAs double their respective CW values, up to a maximum of $C_{\max}$. An STA or the AP that has collided k times since its last successful transmission is said to be in back-off stage $\min\{k, m\}$, where $m = \log_2 (C_{\max}/C_{\min})$.

b) AP and STA-initiated Scenario: Here, an STA A transmits an RTS packet at the same time when the AP transmits an RTS packet to STA B. While STA B cannot decode the RTS from the AP due to CLI from STA A, the AP decodes the RTS from STA A successfully and sends a CTS in response. It then follows the mechanism mentioned above and uses its FD MAC enhancement to transmit a packet to STA B.

If STA A transmits an RTS packet at the same time that the AP transmits its RTS to STA A, the AP's RTS packet cannot be decoded due to self-interference at STA A. However, the AP decodes STA A's RTS and sends a CTS in response. As above, the FD AP sends a packet to another STA even though its HoL packet is for STA A. We call STA A and the AP the primary transmitter and the secondary transmitter, respectively. In this scenario, a success occurs for both AP and STA A, and they reset their CW size to $C_{\min}$.[3] If the AP and at least two STAs transmit an RTS packet at the same time, a collision occurs at the AP and at the STA receiving the RTS from the AP. The nodes involved in the collision double their CW values.

c) AP-initiated Scenario: If the AP wins the contention process, it transmits its packet to an STA with rate r_1. Since it is the sole transmitter, its FD capability is not exploited.

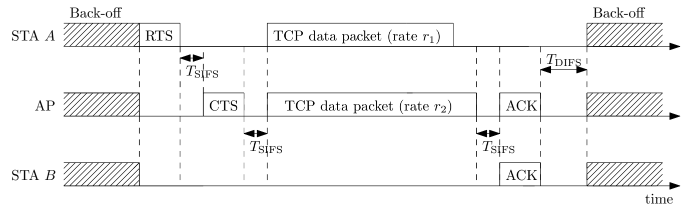

Fig. 2. Illustration of the STA-initiated FD transmission where STA A transmits a TCP data packet to the AP with rate r_1 and the AP transmits a TCP data packet to STA B with rate r_2.

III. TCP Throughput Analysis

An exact Markov chain model of the joint back-off process of the nodes in an 802.11 WLAN is analytically intractable because the number of states is $O\left(N^{m+1}\right)$ [13]. We use the following classical decoupling approximations to make the problem tractable [13]:

1) The back-off process of a given node is independent of the aggregate attempt process of all the other nodes.
2) Each attempt by the AP or an STA collides with a probability that is independent of its back-off stage.

In addition, we introduce the following FD-specific saturation approximation. It is based on the intuition that an FD AP can transmit a packet even when an STA succeeds in the MAC contention process. We shall verify the approximation's accuracy in Section IV. This approximation is a departure from the analyses of TCP over HD WLANs [6]–[10], which assume that any STA has at most one packet in its MAC queue, and only two or three STAs contend for the channel at any time.

3) The AP and all STAs always have packets to transmit in their MAC queues.

Using these approximations, we define the following steady-state probabilities:

1) *Attempt Probability*: The AP and each STA transmit an RTS packet in an idle slot with attempt probabilities β_{AP} and β_{STA}, respectively.
2) *Conditional Collision Probability*: Conditioned on the event that an RTS packet is transmitted by the AP or an STA, a collision occurs with the conditional collision probabilities γ_{AP} and γ_{STA}, respectively.

[2]In practice, the rate r_2 will depend on the locations of the STAs and the AP, and their transmit powers. It is lower than r_1 because the DL STA has to decode the AP's packet in the presence of CLI, which is typically larger than the residual self-interference at the AP.

[3]The AP resets its CW size even after an RTS collision because it can still successfully transmit a packet using its FD capability. This is an FD-specific deviation from the conventional 802.11 HD MAC protocol.

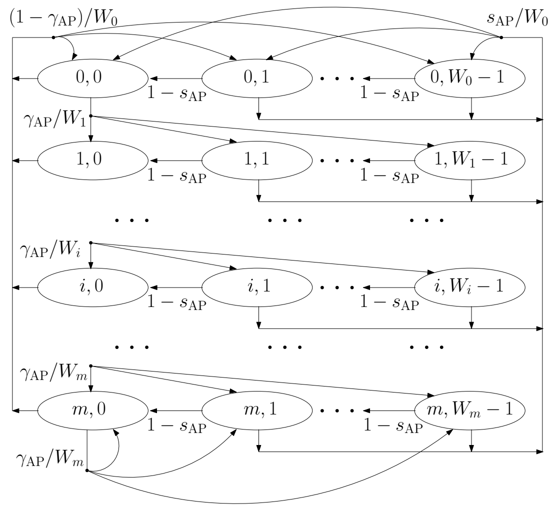

Fig. 3. Markov chain model of the back-off process for the AP, where γ_{AP} and s_{AP} denote the AP's conditional collision probability and the probability of being selected as a secondary transmitter, respectively. The model for an STA is similar except that γ_{AP} is replaced with γ_{STA}, and s_{AP} is replaced with 0 since an STA is never selected as a secondary transmitter.

We now derive expressions for β_{AP}, β_{STA}, γ_{AP} and γ_{STA}.

1) Conditional Collision Probabilities: From Section II-B, an STA's RTS packet collides if at least one other STA transmits an RTS packet in the same slot. Hence, using the decoupling approximations, we have $\gamma_{\text{STA}} = 1-(1-\beta_{\text{STA}})^{N-1}$. On the other hand, an RTS packet sent by the AP collides only if at least two STAs transmit RTS packets in the same slot. Hence, $\gamma_{\text{AP}} = 1-(1-\beta_{\text{STA}})^N - N\beta_{\text{STA}}(1-\gamma_{\text{STA}})$.

2) Attempt Probabilities: Let s_{AP} denote the probability that the AP is chosen as the secondary transmitter. Let t and $t+1$ denote the beginnings of two consecutive intervals in which the AP's back-off timer is decremented. Since the back-off timers are frozen during transmissions, these intervals may have different durations. Let $b(t)$ and $g(t)$ be the AP's back-off timer value and back-off stage, respectively, in the t^{th} interval.

The AP's back-off process state is given by the tuple $(g(t), b(t))$. From the decoupling and saturation approximations, it follows that the back-off process is a two-dimensional discrete-time Markov chain (DTMC) embedded at the instants at which the AP decrements its back-off timer. This DTMC's transition probability diagram is shown in Figure 3. We note that the Markov chain in [12] is a special case and corresponds to $s_{\text{AP}} = 0$.[4] Let $\mathrm{P}(i,j|k,l)$ denote the transition probability from state (k,l) at time t to state (i,j) at time $t+1$. We specify it below. Let $W_i = 2^i C_{\min}$.

The AP decrements its back-off timer by one unless it transmits a packet. If $b(t) \neq 0$, this happens only if the AP is selected as a secondary transmitter. Thus,

$$\mathrm{P}(i,k|i,k+1) = 1-s_{\text{AP}},\ \forall\ 0 \le k \le W_i-2, 0 \le i \le m. \quad (1)$$

[4]A similar DTMC is used in [5]. However, it is for a bidirectional FD MAC, and it considers the MAC retransmission threshold to be exactly m.

The AP transmits an RTS packet as a primary transmitter if $b(t) = 0$. If it is successful, it resets its CW size to $C_{\min}$, and $g(t+1)$ is reset to 0. Hence,

$$\mathrm{P}(0,k|i,0) = \frac{1-\gamma_{\text{AP}}}{W_0},\ \forall\ 0 \le k \le W_0-1, 0 \le i \le m. \quad (2)$$

The AP increases its back-off stage $g(t)$ by 1 if its RTS packet collides. It then samples $b(t+1)$ with uniform probability from the updated CW. Thus,

$$\mathrm{P}(i,k|i-1,0) = \frac{\gamma_{\text{AP}}}{W_i},\ \forall\ 0 \le k \le W_i-1, 1 \le i \le m. \quad (3)$$

However, the CW size remains the same if the AP collides when $g(t) = m$. Hence,

$$\mathrm{P}(m,k|m,0) = \frac{\gamma_{\text{AP}}}{W_m},\ \forall\ 0 \le k \le W_m-1. \quad (4)$$

Finally, if the AP is selected as the secondary transmitter, it successfully transmits a packet. This happens only when $b(t) \neq 0$. The AP resets its CW size to $C_{\min}$, and $g(t+1)$ is reset to 0. Thus,

$$\mathrm{P}(0,k|i,j) = \frac{s_{\text{AP}}}{W_0},\ \forall\ 0 \le k \le W_0-1,\ 0 \le i \le m, \quad 1 \le j \le W_i-1. \quad (5)$$

Result *1:* The attempt probability β_{AP} of the AP equals

$$\beta_{\text{AP}} = \frac{1+\sum_{i=1}^{m} z_i'(s_{\text{AP}},\gamma_{\text{AP}})}{z_0'(s_{\text{AP}},\gamma_{\text{AP}}) + \sum_{i=1}^{m} z_i(s_{\text{AP}}) z_i'(s_{\text{AP}},\gamma_{\text{AP}})}, \quad (6)$$

where $s_{\text{AP}} = N\beta_{\text{STA}}(1-\gamma_{\text{STA}})$, $z_i(s) = \frac{sW_i}{s[1-(1-s)^{W_i}]} - \frac{1-s}{s}$, and

$$z_i'(s,\gamma) = \begin{cases} \left(\frac{\gamma}{s}\right)^i \frac{\prod_{j=1}^{i}[1-(1-s)^{W_j}]}{\prod_{j=1}^{i} W_j}, & 0 \le i \le m-1, \\ \left(\frac{\gamma}{s}\right)^m \frac{\prod_{j=1}^{m}[1-(1-s)^{W_j}]}{\prod_{j=1}^{m} W_j \left(1-\frac{\gamma[1-(1-s)^{W_m}]}{W_m s}\right)}, & i = m. \end{cases} \quad (7)$$

The attempt probability β_{STA} of an STA is given by

$$\beta_{\text{STA}} = \frac{2(1-2\gamma_{\text{STA}})}{(1-2\gamma_{\text{STA}})(W_0+1) + \gamma_{\text{STA}} W_0(1-(2\gamma_{\text{STA}})^m)}. \quad (8)$$

Proof: The proof follows from global balance equations. We skip it to conserve space. ■

A. TCP Throughput Analysis

From the decoupling and saturation approximations, the time instants when the back-off timers are decremented are the renewal instants. The renewal duration T is an RV that depends on whether a success, collision, or idle event has occurred.

Let Θ_{DwL} be the DwL STAs throughput and Θ_{UpL} be the UpL STAs throughput. They are given as follows.

Result *2:* The DwL STAs throughput is given by

$$\Theta_{\text{DwL}} = hL_d \frac{N\beta_{\text{STA}}(1-\gamma_{\text{STA}}) + \beta_{\text{AP}}(1-\gamma_{\text{AP}})}{\mathbb{E}[T]}, \quad (9)$$

where $h = \frac{N_D}{N}$ is the probability that the AP's HoL packet is a TCP data packet. The UpL STAs throughput is given by

$$\Theta_{\text{UpL}} = N_U L_d \frac{\beta_{\text{STA}}(1-\gamma_{\text{STA}})}{\mathbb{E}[T]}. \quad (10)$$

TABLE I
SIMULATION PARAMETERS

Parameter	Value	Parameter	Value	Parameter	Value
T_{RTS}	36 μsec.	T_{CTS}	44 μsec.	T_{PHY}	24 μsec.
T_{DIFS}	34 μsec.	T_{SIFS}	16 μsec.	T_{ACK}	44 μsec.
r_1	54 Mbps	r_2	36 Mbps	δ	9 μsec.
$C_{\max}$	1024	$C_{\min}$	32	$w_{\max}$	64 KB
L_d	16 Kb	L_a	320 bits	L_h	432 bits

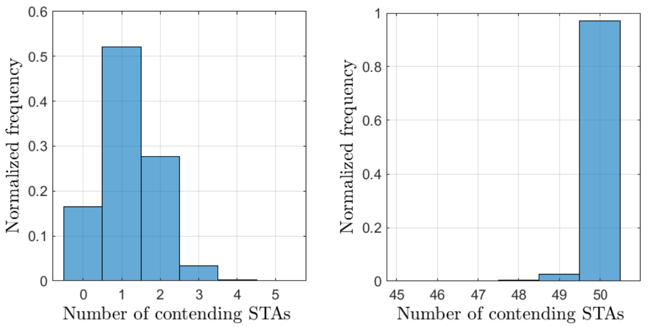

(a) TCP over HD (b) TCP over FD

Fig. 4. Histogram of the number of contending STAs for TCP over HD and FD WLANs ($N = 50$ and $f_{\text{UpL}} = \frac{1}{4}$).

The mean renewal duration $\mathbb{E}[T]$ is given by

$$\begin{aligned}\mathbb{E}[T] &= (1-\beta_{\text{STA}})^N(1-\beta_{\text{AP}})\delta \\ &\quad + \beta_{\text{AP}}(1-\gamma_{\text{AP}})\left[hT_p(r_1)+(1-h)T_a(r_1)\right] \\ &\quad + N_U\beta_{\text{STA}}(1-\gamma_{\text{STA}})\left[hT_p(r_2)+(1-h)T_p(r_1)\right] \\ &\quad + N_D\beta_{\text{STA}}(1-\gamma_{\text{STA}})\left[hT_p(r_2)+(1-h)T_a(r_2)\right] \\ &\quad + (1-(1-\beta_{\text{STA}})^N(1-\beta_{\text{AP}})-N\beta_{\text{STA}}(1-\gamma_{\text{STA}}) \\ &\quad -\beta_{\text{AP}}(1-\gamma_{\text{AP}}))[T_{\text{RTS}}+T_{\text{DIFS}}], \qquad (11)\end{aligned}$$

where $T_p(r)$ and $T_a(r)$ are given by

$$T_p(r) = T_{\text{RTS}} + 3T_{\text{SIFS}} + T_{\text{CTS}} + T_{\text{PHY}} + \frac{L_h+L_d}{r} + T_{\text{ACK}} + T_{\text{DIFS}},$$

$$T_a(r) = T_{\text{RTS}} + 3T_{\text{SIFS}} + T_{\text{CTS}} + T_{\text{PHY}} + \frac{L_h+L_a}{r} + T_{\text{ACK}} + T_{\text{DIFS}}.$$

T_{PHY} denotes the PHY preamble and header duration, and L_h denotes the total TCP and MAC header length in bits.

Proof: The proof is relegated to Appendix A. ■

The total TCP throughput is equal to $\Theta_{\text{DwL}} + \Theta_{\text{UpL}}$.

IV. SIMULATION RESULTS

The simulation parameters are listed in Table I. The source code for simulating FD in 802.11 WLANs is not available in existing network simulators such as ns-3 or Opnet. We, thus, implement a custom simulator in MATLAB that integrates the FD WLAN MAC protocol with the slow start and congestion avoidance phases of TCP NewReno [14]. The simulations are run for a duration of 45 sec. It takes 21.11 sec. and 28.35 sec. for FD and HD, respectively, for the congestion windows of all TCP connections to reach $w_{\max}$. Let $f_{\text{UpL}} = \frac{N_U}{N}$ denote the fraction of UpL STAs.

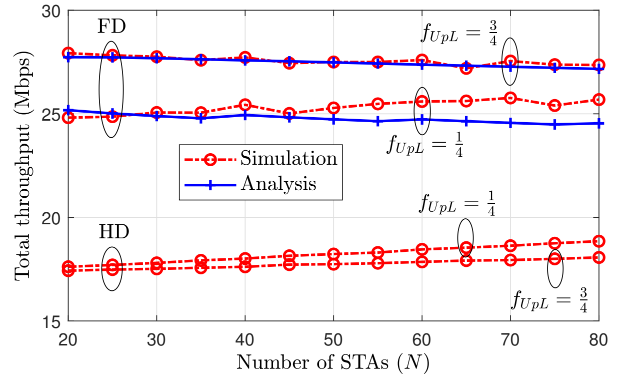

Fig. 5. Total TCP throughput as a function of N for different values of f_{UpL}.

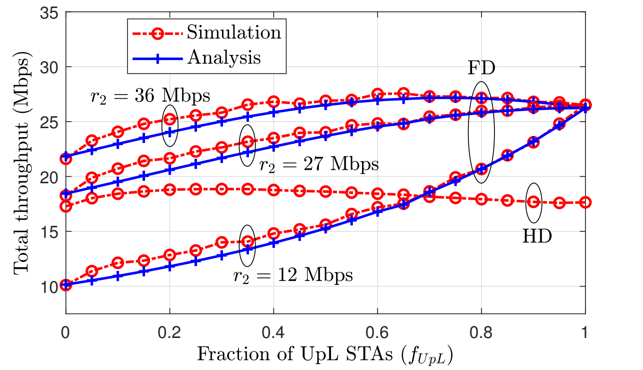

Fig. 6. Total TCP throughput as a function of f_{UpL} for different values of r_2 ($N = 80$).

Figures 4a and 4b plot the histogram of the number of STAs contending for the channel in TCP over HD and FD WLANs, respectively, for $N = 50$. In an HD WLAN, no more than three STAs contend for the channel at any given instant. This is consistent with the behavior reported in [6]–[10]. However, in an FD WLAN, the number of contending STAs is N with a probability of 0.97, which validates our saturation approximation.

Figures 5 plots the total TCP throughput over FD and HD WLANs as a function of N. Both analysis and simulation results are shown. They are within 5% of each other. The TCP throughput is insensitive to N for both FD and HD. While the TCP throughput is sensitive to f_{UpL} for FD, this is not so for HD. The total TCP throughput of FD is 41% to 54% more than that of HD when f_{UpL} increases from $\frac{1}{4}$ to $\frac{3}{4}$.

Figure 6 plots the total TCP throughput of the FD WLAN as a function of f_{UpL} for different values of r_2. Both analysis and simulation results are plotted and are within 6% of each other. The total TCP throughput is sensitive to f_{UpL}, which affects h and the transmission durations. For $r_2 = 36$ Mbps, FD achieves throughputs that are 26.2% to 50.3% greater than HD as f_{UpL} increases from 0 to 0.65. For $r_2 = 27$ Mbps, FD achieves throughputs that are 5.7% to 48.7% greater than HD as f_{UpL} increases from 0 to 1. However, for $r_2 = 12$ Mbps, FD achieves a lower throughput than HD for $0 \leq f_{\text{UpL}} \leq 0.65$ since the total transmission durations can be longer than for HD for small r_2, which negates the gains from using FD. However, as f_{UpL} increases, FD achieves a higher throughput because transmissions are dominated by UpL STAs sending TCP data in the UL and the AP sending TCP ACKs in the DL. These are of duration $T_p(r_1)$, which is not a function of r_2.

V. Conclusions

We analyzed the TCP throughput over an 802.11 FD WLAN using Markov renewal theory and fixed point analysis. Our analysis captured the interaction between TCP and the IEEE 802.11 CSMA/CA-based FD MAC. We saw that the additional transmission opportunities provided by the FD capability led to the AP and the STAs having a packet to transmit in their queues with high probability. This differed from TCP over HD WLANs, in which the AP was a bottleneck and led to few STAs contending for the channel. Despite the increased contention between STAs, asymmetric uplink and downlink payloads due to the different sizes of TCP data packets and TCP ACKs, and different uplink and downlink rates due to CLI, an FD WLAN achieved significantly higher TCP throughput than an HD WLAN. Interesting avenues for future work include modeling the effect of packet losses and delays, and the delayed ACK mechanism.

Appendix

A. Proof of Result 2

a) Derivation of $\mathbb{E}[R]$*:* The reward R_{DwL} is L_d when the AP transmits a TCP data packet during a success interval and is 0 otherwise. The AP's HoL packet is a TCP data packet with probability h. The AP can be successful as a primary transmitter with probability $\beta_{\text{AP}}(1-\gamma_{\text{AP}})$ or as a secondary transmitter with probability $N\beta_{\text{STA}}(1-\gamma_{\text{STA}})$. Thus, $\mathbb{E}[R_{\text{DwL}}] = hL_d\left(N\beta_{\text{STA}}(1-\gamma_{\text{STA}}) + \beta_{\text{AP}}(1-\gamma_{\text{AP}})\right)$. From the renewal reward theorem [13], $\Theta_{\text{DwL}} = \mathbb{E}[R_{\text{DwL}}]/\mathbb{E}[T]$.

The reward R_{UpL} is L_d when an UpL STA transmits a TCP data packet during a success interval, which occurs with probability $\beta_{\text{STA}}(1-\gamma_{\text{STA}})$, and is 0 otherwise. Thus, $\mathbb{E}[R_{\text{UpL}}] = L_d N_U \beta_{\text{STA}}(1-\gamma_{\text{STA}})$. From the renewal reward theorem, $\Theta_{\text{UpL}} = \mathbb{E}[R_{\text{UpL}}]/\mathbb{E}[T]$.

Derivation of h: The arrival rate of TCP data packets into the AP's MAC queue is equal to the rate at which it receives TCP ACKs from all the DwL STAs. Using the renewal-reward theorem, the rate at which the DwL and UpL STAs send TCP ACK and TCP data packets to the AP is $\frac{N_D\beta_{\text{STA}}(1-\gamma_{\text{STA}})}{\mathbb{E}[T]}$ and $\frac{N_U\beta_{\text{STA}}(1-\gamma_{\text{STA}})}{\mathbb{E}[T]}$, respectively. Thus, $h = \frac{N_D}{N_D+N_U} = \frac{N_D}{N}$.

b) Derivation of $\mathbb{E}[T]$*:* The renewal interval duration T depends on whether an idle, success, or collision event occurred.

1) *Idle:* T is δ with probability $(1-\beta_{\text{STA}})^N(1-\beta_{\text{AP}})$. This yields the first summand in (11).
2) *Success:* T depends on who succeeds.
   - *AP Succeeds*: This occurs with probability $\beta_{\text{AP}}(1-\gamma_{\text{AP}})$. Here, $T = T_p(r_1)$ if the AP's HoL packet is a TCP data packet. Else, $T = T_a(r_1)$ if the AP's HoL packet is a TCP ACK packet. The second summand in (11) accounts for this case.
   - *UpL STA Succeeds*: This occurs with probability $N_U\beta_{\text{STA}}(1-\gamma_{\text{STA}})$. Here, $T = \max\{T_p(r_1), T_a(r_2)\} = T_p(r_1)$ if the AP's HoL packet is a TCP ACK packet. Else, $T = \max\{T_p(r_1), T_p(r_2)\} = T_p(r_2)$ if the AP's HoL packet is a TCP data packet. The third summand in (11) accounts for this case.
   - *DwL STA Succeeds*: This occurs with probability $N_D\beta_{\text{STA}}(1-\gamma_{\text{STA}})$. Here, $T = \max\{T_a(r_1), T_a(r_2)\} = T_a(r_2)$ if the AP's HoL packet is a TCP ACK packet. Else, $T = \max\{T_a(r_1), T_p(r_2)\} = T_p(r_2)$ if the AP's HoL packet is a TCP data packet. The fourth summand in (11) accounts for this case
3) *Collision:* T is $T_{\text{RTS}} + T_{\text{DIFS}}$, if it is neither an idle nor a success interval. This yields the last summand in (11).

The above expressions yield the expression for $\mathbb{E}[T]$ in (11).

References

[1] X. Lin, "An overview of 5G advanced evolution in 3GPP release 18," *IEEE Commun. Standards Mag.*, vol. 6, no. 3, pp. 77–83, Sep. 2022.

[2] X. Yang, Y. Mao, Y. Gong, and T. Zhang, "Uplink/downlink initiated based MAC protocol for asymmetric full duplex radio to improve throughput," in *Proc. IWCMC*, Jul. 2021, pp. 1075–1078.

[3] Y. Alkhrijah, J. Camp, and D. Rajan, "HyFDMAC: A hybrid access full-duplex MAC protocol," in *Proc. COMSNETS*, Jan. 2023, pp. 572–578.

[4] L. H. Vu, R. Mareta, and J.-H. Yun, "Full-duplex wireless LAN incorporating successive interference cancellation," *IEEE Trans. Veh. Technol.*, vol. 70, no. 10, pp. 10 293–10 307, Aug. 2021.

[5] M. Murad and A. M. Eltawil, "Performance analysis and enhancements for in-band full-duplex wireless local area networks," *IEEE Access*, vol. 8, pp. 111 636–111 652, Jun. 2020.

[6] R. Bruno, M. Conti, and E. Gregori, "Analytical modeling of TCP clients in Wi-Fi hot spot networks," in *Proc. Networking*, May 2004, pp. 626–637.

[7] ——, "Performance modelling and measurements of TCP transfer throughput in 802.11-based WLAN," in *Proc. MSWiM*, Oct. 2006, pp. 4–11.

[8] S. Krishnasamy and A. Kumar, "Modeling the effect of transmission errors on TCP controlled transfers over infrastructure 802.11 wireless LANs," in *Proc. MSWiM*, Oct. 2011, pp. 275–284.

[9] G. Kuriakose, S. Harsha, A. Kumar, and V. Sharma, "Analytical models for capacity estimation of IEEE 802.11 WLANs using DCF for internet applications," *Wirel. Netw.*, vol. 15, no. 2, pp. 259–277, Feb. 2009.

[10] S. R. Pokhrel, M. Panda, H. L. Vu, and M. Mandjes, "TCP performance over Wi-Fi: Joint impact of buffer and channel losses," *IEEE Trans. Mobile Comput.*, vol. 15, no. 5, pp. 1279–1291, May 2016.

[11] M. G. Sarret, G. Berardinelli, N. H. Mahmood, and P. Mogensen, "Impact of transport control protocol on full duplex performance in 5G networks," in *Proc. IEEE VTC (Spring)*, May 2016, pp. 1–5.

[12] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 3, pp. 535–547, Mar. 2000.

[13] A. Kumar, E. Altman, D. Miorandi, and M. Goyal, "New insights from a fixed point analysis of single cell IEEE 802.11 WLANs," in *Proc. IEEE INFOCOM*, Mar. 2005, pp. 1550–1561.

[14] S. Floyd, T. Henderson, A. Gurtov, and Y. Nishida, "The NewReno modification to TCP's fast recovery algorithm," RFC 6582, Apr. 2012. [Online]. Available: https://www.rfc-editor.org/rfc/rfc6582

[15] I. Rhee, L. Xu, S. Ha, A. Zimmermann, L. Eggert, and R. Scheffenegger, "CUBIC for fast long-distance networks," RFC 8312, Feb. 2018. [Online]. Available: https://www.rfc-editor.org/info/rfc8312

[16] R. Kiran, N. B. Mehta, and J. Thomas, "Design and network topology-specific renewal-theoretic analysis of a MAC protocol for asymmetric full-duplex WLANs," *IEEE Trans. Commun.*, vol. 67, no. 12, pp. 8532–8544, Dec. 2019.