Ce-Fi: Centralized Wi-Fi over Packet-Fronthaul

Simon Schmitz-Heinen, Anatolij Zubow, Christos Laskos, Falko Dressler
School of Electrical Engineering and Computer Science, TU Berlin, Germany
{schmitz-heinen,zubow,laskos,dressler}@ccs-labs.org

***Abstract*—Cloud radio access network (C-RAN) in cellular systems offers improved coverage and efficiency due to statistical multiplexing. Applying centralization to Wi-Fi networks could yield similar benefits, particularly in dense deployments like campus networks. However, this is challenging due to the usage of unlicensed, shared spectrum and strict protocol timing constraints like carrier sensing, making such systems highly sensitive to C-RAN induced delays. Existing approaches typically rely on radio-over-fiber and either demand very low delay or neglect neighboring Wi-Fi networks. In this paper, we present `Ce-Fi`, a centralized Wi-Fi architecture operating over a packet-based fronthaul (FH) connecting the central unit with the remote radio heads. `Ce-Fi` employs a two-level design with two bands: the network discovery and user association is handled standard compliantly on one band while another band is used for the actual data transmission with the centralized Wi-Fi. To tolerate the FH-induced delay, we propose minor changes to the 802.11 medium access control protocol like NAV extension and the usage of a piggybacking mechanism enabling the C-RAN-AP to transmit without prior channel contention. We evaluate `Ce-Fi` through simulations, demonstrating that it supports high throughput even under high FH-delays and in the presence of neighboring legacy networks.**

***Index Terms*—Cloud RAN, IEEE 802.11, Wireless**

I. Introduction

The cloud radio access network (C-RAN) architecture is becoming an increasingly popular technology in cellular networks [1] and is, for example, used to increase radio coverage by densely deploying inexpensive remote radio heads (RRHs) instead of full-fledged radio towers. The RRHs handle the actual radio transmission and forward digitized IQ samples via fiber to the central unit (CU), which is connected to the mobile backhaul network and performs all further signal processing, as illustrated in Figure 1. Among other benefits, this architecture enables statistical multiplexing gain, whereby the total computing power required by the CU is less than or equal to what would be needed if each RRH operated as an independent, full-function radio tower [2].

These benefits could also be of interest to other wireless access technologies, such as Wi-Fi. While it has been shown that dense deployment of Wi-Fi access points (APs) can improve network performance, an excessive number of APs leads to increased interference, ultimately degrading performance [3]. In this context, centralized Wi-Fi could offer significant advantages: instead of deploying full APs, inexpensive, centrally coordinated RRHs communicating at low power could be densely deployed, for example, by placing one or more in each room of a building. As a result, there would be less interference and each station (STA) would contend with fewer neighboring STAs, ultimately increasing the throughput per STA. Thus, centralized Wi-Fi could enable the cost-effective construction of large enterprise or campus networks, offering high coverage and excellent performance. The RRHs could utilize packet-switched high-throughput Ethernet infrastructure to connect to the CU for coordinated processing. Using high-throughput fiber-based Ethernet would lead closer to transitioning from fibre-to-the-home (FTTH) to fibre-to-the-room (FTTR), where users can benefit from fiber Internet connections, without requiring a wired connection.

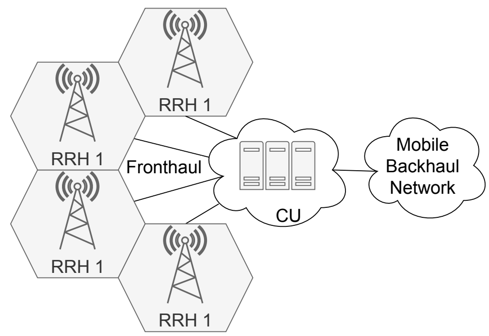

Figure 1. Architecture of C-RAN consisting of RRHs connected over a fronthaul (FH) channel to the CU (Based on Checko *et al.* [2, Fig. 2b])

However, the practical implementation of centralized Wi-Fi is not trivial due to Wi-Fi's usage of unlicensed, shared spectrum and tight protocol timing requirements for example during the channel access using carrier sense multiple access with collision avoidance (CSMA/CA). Here, STAs are required to assess the state of the channel, an assessment that becomes nearly impossible, as the AP perceives the medium's state highly delayed. This makes Wi-Fi very susceptible to delays and jitter. While existing literature offers solutions to this problem, they primarily focus on centralized long-distance Wi-Fi utilizing radio-over-fiber (RoF) technology, where the added delay is limited to propagation delay [4]–[11]. Although RoF would mitigate delay-related issues in indoor environments thanks to its near-negligible delays over short distances, dedicated fibers between the RRHs and the CU become expensive and complex when scaling to a large number of RRHs, making packet-fronthaul (FH) solutions more feasible, albeit at the cost of additional delay [12].

In this paper, we present `Ce-Fi`, an approach to enable centralized Wi-Fi even with FH-delays in the order of more than 100 µs. Our solution employs a two-level architecture for association and communication, where communication with centralized APs is made possible through a combination of network allocation vector (NAV) extension and piggybacking, augmented by a forced traffic (FT) scheme to ensure sufficient piggybacking opportunities.

II. Related Work

The majority of previous work on centralizing Wi-Fi focuses on using RoF links between the RRH and the CU to enable long-distance Wi-Fi networks with CU-RRH distances of several kilometers [5], [8]–[11].

Kalantari-Sabet *et al.* [13] investigated the delay-induced performance impairments of such a system and have shown that the length of the fiber cable and consequently the associated delay, negatively impacts communication, e.g. due to the expiry of acknowledgment (ACK) or clear-to-send (CTS) timeouts. Deronne *et al.* [4] also conducted research on the impact of this propagation delay using network simulations. While they highlight advantages such as cost savings through the use of simple RRHs connected to a more complex CU, the authors observed that the likelihood of collisions increases with every additional half slot-time of propagation delay, which has an immediate negative impact on network performance. To avoid the expiry of the aforementioned timeouts, they suggest to simply increase these parameters proportionally to the additional delay. In a subsequent publication, Deronne *et al.* [5] propose to also increase the slot-time parameter proportionally to the delay to mitigate collisions and improve performance. The downside of this approach is that, as slot-times increase, so do the waiting times, which consequently reduces network efficiency. This issue can be partially mitigated by using frame aggregation [6]. However, it is shown that their approach is not capable of operating in the presence of neighboring overlapping basic service sets (OBSSs) [10].

Tanaka *et al.* [7] analyze the performance of delay-affected Wi-Fi systems using both a Markov model and OPNET simulations, also demonstrating throughput degradation with increasing delay. Okamoto *et al.* [8] describe the occurrence of collisions between ACK and data frames when the delay exceeds the virtual reservation (NAV) time. Thus, they propose extending the NAV and suggest that the AP transmits its data a short inter-frame space (SIFS) after acknowledging uplink (UL) traffic to increase downlink (DL) throughput. This approach consequently synchronizes UL and DL throughput. Nishio *et al.* [9] build on this with a piggybacking-based medium access control (MAC) protocol, where the duration-field of CTS-frames in response to incoming ready-to-send (RTS)-frames are extended by twice the expected delay to set the NAV and virtually reserve the channel. Their approach also allows to adjust the UL-to-DL ratio by setting the probability with which the NAV will be extended and used to piggyback data. While this mechanism does only requires protocol changes from the AP, it is dependent on sufficient UL traffic and the use of the RTS/CTS mechanism. Furthermore, Funabiki *et al.* [10] fill in some shortcomings of the solution by Nishio *et al.* [9] by proposing a mechanism which enables the fair coexistence of long-range and legacy Wi-Fi. This works by overhearing legacy transmissions of a legacy network and timing DL transmission so that data reaches the channel without causing collisions. However, this approach requires deterministic delay knowledge, stable conditions, and the use of RTS/CTS in the neighboring network, which cannot be guaranteed. Also, the system is limited by the expected delay as it only works if the delay is lower than half the time virtually reserved by the original RTS-frame.

Another approach is presented by Valkanis *et al.* [11] where a hybrid time division multiple access (TDMA) system modifying IEEE 802.11ac is used to support environments impacted by delays. They optimize channel access by adjusting inter-frame space (IFS) timings and limiting contention to stations only, defining separate UL/DL states to avoid collisions. Yet, in comparison to the aforementioned solution, this method may struggle in legacy coexistence scenarios due to altered IFS. Also, it requires protocol changes from both, AP and STAs.

While the aforementioned studies focus on extending Wi-Fi coverage via RRHs, centralization can also simplify AP management. Dely *et al.* [14] propose to implement virtual APs, where the distributed coordination function (DCF) remains in the RRH and further MAC processing is centralized, which supports new services but requires more complex RRH. Kim *et al.* [15] suggest the usage of immediate ACK-transmissions after header decoding to mitigate delay-related performance degradation to allow for dense RRH-deployment. However, they focus on data UL and only consider low delays.

III. IEEE 802.11 Channel Access

In Wi-Fi networks, the distributed coordination function (DCF) is responsible for coordinating channel access to the wireless medium using the CSMA/CA mechanism [16]. In CSMA/CA, STAs contend for the channel by first performing a clear channel assessment (CCA) to sense whether the medium is currently free. If the channel is free, a contending STA waits for a distributed inter-frame space (DIFS) and then an additional random backoff period. If the channel remains idle after the backoff expires, the STA begins its transmission and, in the case of a unicast data transfer, waits for an acknowledgment. However, if the channel becomes busy during the backoff period, the countdown is paused. Once the channel is idle again (after another DIFS), the station resumes waiting for the remaining backoff time. The backoff duration is determined by the contention window (CW), as defined in Equation (1), where t_{slot} depends on the Wi-Fi standard (e.g., 9 µs in 802.11ac):

$$t_{\text{Backoff}} = \text{uniform}(0,\ \text{CW}) \times t_{\text{slot}} \tag{1}$$

If a transmitted packet is not acknowledged after a SIFS, a timeout is triggered. The STA then assumes the packet was lost, possibly due to a collision caused by high network traffic, and initiates a retransmission using an exponential backoff. This means the CW, which initially starts at $\text{CW}_{\min}$, is doubled with each failed attempt until it reaches $\text{CW}_{\max}$. After a successful transmission, the CW is reset to $\text{CW}_{\min}$. The station retries transmission until a retry limit is reached, at which point the packet is dropped. DCF employs two mechanisms to determine whether the channel is idle: The first mechanism is physical carrier sensing, which detects whether energy is being received on the channel. The second is virtual carrier sensing. When a frame is received, STAs inspect the duration field in the MAC header. This field indicates how long the transmitting STA expects the medium to remain busy for the current frame exchange. In a unicast transmission, for example, the duration field includes the time required for a SIFS and the transmission of the ACK frame. This value sets the NAV, a timer that virtually reserves the channel for the specified period, preventing another station from attempting to transmit.

IV. Limits of Carrier Sensing

The primary issue that arises with the FH induced delay is that the channel access mechanism of 802.11 fails to function reliably. Since CSMA/CA relies on accurate and timely assessment of the channel status to determine whether a transmission can begin without causing a collision, a high delay in this assessment proves detrimental. Figure 2 illustrates that a transmission from STA *A* is only received by STA *B* after the FH-delay δ, effectively causing the FH-affected STA to perceive the channel state as it was in the past. Similarly, when the FH-affected station initiates a transmission, it reaches the channel only after a one-way FH-delay δ. When the FH-affected STA *B* starts a transmission and while its data traverses the FH, other STAs can assess the channel to be idle and start their own transmission, leading to a collision. This becomes particularly problematic when the FH-delay exceeds a multiple of half a slot-time, as the CSMA/CA mechanism is based on discrete backoff slots [5]. When two STAs select the same number of backoff slots and begin transmission simultaneously, a collision occurs, thereby reducing throughput. However, if one of the STAs is affected by FH, collisions may occur even more frequently, not only when identical backoff slots are selected, because successful CCA requires both the channel state and the actual transmission to traverse the FH. As a result, the number of overlapping backoff slots increases with each half slot-time increment of the FH-delay. With reference to Figure 2, the collision window is 2δ.

In a non-centralized system, the overall sensing delay is determined by the sum of propagation delay between STAs (which is negligible in conventional Wi-Fi scenarios) and processing delay. In centralized Wi-Fi systems, however, this delay is further increased by the delays introduced by the FH-channel, such as additional propagation or transmission delays.

V. System Design

This section provides an explanation of the proposed system design and its components.

A. Packet-FH

In opposition to existing approaches, which use RoF [4], [9]–[11], `Ce-Fi` uses a packet-FH. Packet-FH offers a cost-effective, resilient and flexible solution for FH-systems with the possibility of reusing existing infrastructure [12], [17]. However, in contrast to RoF, where the FH-delay primarily consists of the propagation delay, the main latency for packet-FH-based systems in indoor environments come from packetization, processing, and transmission [18]. Furthermore, since packet-FH operates over packet-switched networks, the resulting delay is non-deterministic and subject to jitter, unlike the deterministic delay in RoF systems. While advanced scheduling algorithms can help reduce the jitter, it is hard to eliminate it entirely without further increasing the overall delay [19]. The delay experienced in packet-FH varies significantly, even in zero-hop network configurations. Measurements by Laskos *et al.* [20] indicate that, even under best-case conditions using high-end hardware and 100 Gbit/s Ethernet links, the FH-delay ranges between 15–30 µs. Under less favorable hardware conditions, this delay can reach up to 200 µs. To implement centralized Wi-Fi, for the packet-FH enhanced common public radio interface (eCPRI) could be used, which is employed in cellular networks with packet-FH.

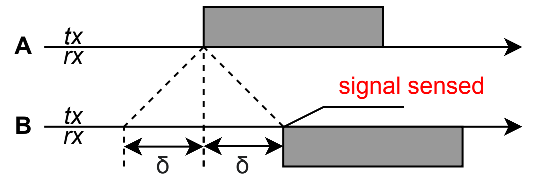

Figure 2. Transmission by node B (with FH-delay) during the interval $2 \times \delta$ cannot be prevented by carrier sensing, resulting in simultaneous transmissions with node A (without FH-delay)

B. Two-Level Architecture

`Ce-Fi` follows a two-level architecture, depicted in Figure 3, where each level is characterized by the use of a specific frequency band and a task to perform:

The first level, using the wider range 2.4 GHz band, is used for discovery and association with the network. This association level is necessary here, as the FH-delay impairs the centralized AP's ability to transmit the required messages for association. The association is achieved by introducing the association unit (AU), which is essentially a regular Wi-Fi-AP operating in the 2.4 GHz band and connected to the rest of the network via an Ethernet backhaul. The AUs are strategically placed to provide radio coverage over the entire application area. The main difference between AUs and APs is that the AUs check whether associating STAs support `Ce-Fi` and then perform channel steering [21], to shift the associating STAs from the 2.4 GHz band to the 5 GHz band for the actual data exchange. In cases where a STA does not support `Ce-Fi`, the STA can remain connected to the AU and use it for data communication in the 2.4 GHz band.

In the second level, which operates in the 5 GHz band, STAs perform their actual data transmissions. Instead of communicating with a regular AP, they communicate with RRHs, which are low-cost antennas connected to a CU via a delay-inducing FH channel, such as a high-throughput Ethernet link. The RRH transmits its radio signals in packetized form to the CU, which

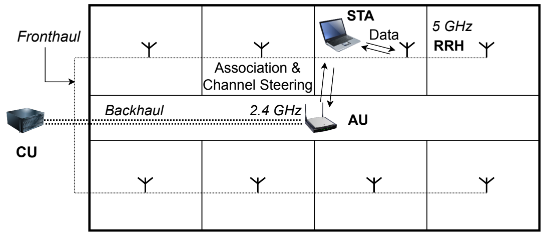

Figure 3. Example Ce-Fi deployment with RRHs in each room, connected to the CU via a packet FH, and with a centrally located AU connected to the CU via a backhaul.

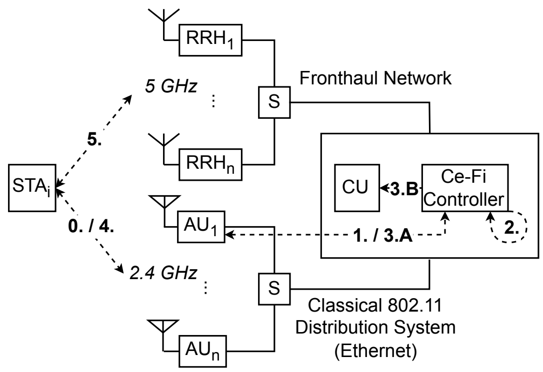

Figure 4. General Ce-Fi architecture, with a STA communicating to an AU for association and to RRHs for data transmission

can consist of general-purpose computing hardware and takes over the processing and all MAC functionality, analogous to the C-RAN architecture in mobile networks. In contrast to the AU, the RRH covers only a very small area with its radio, for example, a single room, and communicates within that area using low power. However, RRHs are densely deployed, with one or more per room, to provide excellent overall coverage.

C. Detailed Specification

In the following, details of Ce-Fi will be described. The general architecture and communication flow is shown in Figure 4.

1) Channel Steering: Once a STA has been associated to the network via the AU (**0.**), the AU communicates the STA's capabilities to the Ce-Fi-controller (**1.**), which then decides whether to accept the STA (**2.**). If the STA is Ce-Fi-capable, the Ce-Fi-controller will tell the AU to channel-steer the STA (**3.A**) by sending a unicast beacon frame with a channel-switch announcement and FH-delay information to the STA (**4.**) [21]. This will cause the STA to change its channel from the 2.4 GHz band to the 5 GHz band used by the RRHs. At the same time, the new STA is initiated at the CU (**3.B**). Once the channel has been changed, the STA will use all the mechanisms required for Ce-Fi and start its communication with the RRH (**5.**).

2) Dealing with High FH-Delay: Based on the information provided by the AU, the STA increases its ACK timeout by twice the FH-delay to account for the additional time required for the reception of the ACK. In addition, the STA increases the duration field of each packet sent via unicast, thereby extending the virtual channel reservation using the NAV to

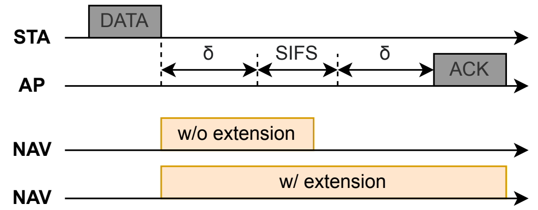

Figure 5. NAV with and without expansion in case of a centralized Wi-Fi network with FH-delay

ensure that no other STA will send before the packet has been acknowledged, as shown in Figure 5. This NAV extension was proposed by Okamoto *et al.* [8] and has the advantage of ensuring that neighboring networks within range will refrain from transmitting if they hear the data transmission. While the STAs can correctly perform CCA for CSMA/CA, the RRH and, by extension, the centralized AP cannot. The AP therefore piggybacks its data onto STA transmissions, sending its own data after a SIFS interval following the ACK. Since the STAs use the NAV extension and the AP transmits its data SIFS after transmitting the ACK, no other STA can access the channel before it, thereby allowing for an undisturbed DL transmission. In this sense, the UL acts as a reverse trigger for the DL.

3) Forced Traffic in Low UL Scenarios: The aforementioned mechanism allows for communication when there is a significant amount of UL traffic to piggyback on. However, this is not always guaranteed, as there may be situations with mostly DL traffic, such as live TV streaming. While Nishio *et al.* [9] suggest switching from piggybacking to regular channel access when there is low UL, this approach would be prone to failure in the presence of neighboring OBSS. To mitigate this, we propose a forced traffic (FT) mechanism, in which the associated STAs are requested to regularly send null data frames (NDFs) whenever they neither have data to transmit nor overhear UL traffic from other STAs in the same basic service set (BSS). This ensures that there are sufficient transmission opportunities for the AP to send its DL. As with the NAV extension, the AU informs the STAs to participate in the FT mechanism.

The FT mechanism works as follows: Whenever a STA associated with a Ce-Fi network has no data to transmit, it will enqueue a NDF and begins contending for the channel to provide the AP with a piggybacking opportunity. To prevent STAs with enqueued NDF from interfering with other STAs in the same BSS that have actual data to transmit, the backoff calculation is offset according to Equation (2). This prioritizes actual UL traffic and avoids negative impact on the system throughput.

$$t_{\text{Backoff}} = (\text{uniform}(0,\ \text{CW}) + \text{CW}_{\text{offset}}) \times t_{\text{slot}} \quad (2)$$

$\text{CW}_{\text{offset}}$, however, is not static and depends on the traffic situation. Generally, the $\text{CW}_{\text{offset}}$ is related to the $\text{CW}_{\min}$ and $\text{CW}_{\max}$ parameters. The default value for $\text{CW}_{\text{offset}}$ is $\text{CW}_{\min}$, while the maximum is equal to $\text{CW}_{\max}$. Setting the default $\text{CW}_{\text{offset}}$ to $\text{CW}_{\min}$ is done to prioritize STAs with actual UL.

In the following, we assume that the considered STA has no UL traffic and has a NDF enqueued. Whenever the STA transmits a NDF, it doubles its current $\mathrm{CW_{offset}}$, analogous to the exponential backoff in CSMA/CA. If the STA does not overhear a subsequent transmission from the AP, this suggests that there is no DL and no need for further NDF transmission, which is why the $\mathrm{CW_{offset}}$ is doubled. Since there is no immediate feedback mechanism on whether the RRH will not use the NDF for piggybacking, this doubling needs to be done preemptively.

If, on the other hand, the STA successfully overhears DL as the next transmission in its BSS, it will reset the $\mathrm{CW_{offset}}$ to zero, recalculate the backoff, and restart the contention. The successful usage of the NDF for DL indicates low UL in the BSS, resulting in fewer piggybacking opportunities, which is why it is necessary to increase the chance of transmitting a NDF by setting the $\mathrm{CW_{offset}}$ to zero.

If a STA receives DL without having sent a NDF beforehand, and its $\mathrm{CW_{offset}}$ is greater than $\mathrm{CW_{min}}$, it will reset the $\mathrm{CW_{offset}}$ to $\mathrm{CW_{min}}$, as this indicates the general presence of DL. Furthermore, the contention will restart with the new $\mathrm{CW_{offset}}$, provided the remaining backoff slots are larger than $\mathrm{CW_{offset}} + \mathrm{CW}$.

Additionally, the $\mathrm{CW_{offset}}$ is changed from zero to $\mathrm{CW_{min}}$ whenever the STA overhears UL data within its BSS, as this UL contains valuable user data and provides piggybacking opportunities, making NDFs unnecessary.

Finally, whenever a STA with an enqueued NDF receives UL within its BSS, it recalculates the random backoff and restarts its contention process. This is done to prevent the STA's backoff from declining too quickly, which could lead to it accessing the channel even when other stations still have data to send.

VI. Evaluation

A. Methodology

To evaluate `Ce-Fi`, various scenarios were simulated using the OMNeT++ simulator [22]. In OMNeT++, messages are exchanged over channels between different modules and submodules. These channels can be configured to delay a transmission using a parameter. To model the FH-delay, a fixed delay was added to the channel between the MAC and radio submodules of the WLAN module. While this setup generally models the FH quite well, additional modifications were required in other modules to ensure that the rest of the stack functions properly. The default simulation parameters are listed in Table I. Any deviations from these parameters will be explicitly stated for each scenario. In every scenario with an increased FH-delay, timeouts will be adjusted accordingly. The actual timeouts are vendor-specific and can be easily modified in the device firmware.

An adjusted INET4.5 framework together with examples as used for the evaluation in this paper is provided to the community as open-source under a GPL on github: `https://github.com/tkn-tub/OMNeT_Ce-Fi`.

Table I
Simulation Parameters

Parameter	Value
# Simulation runs	20
Simulation run time	30 s
# STAs per BSS	6 (Greenfield) / 3 (OBSS)
Mobility	Random static positioning in 15 m×15 m area
Traffic Intensity	Full Buffer
Transport Protocol	User datagram protocol (UDP)
Wi-Fi Standard	IEEE 802.11ac
Bit Rate	693.3 MByte/s
MAC payload	8000 Byte
Slot-Time	9 µs
SIFS	16 µs
$\mathrm{CW_{min}}$	15
$\mathrm{CW_{max}}$	1023

B. Impact of FH-Delay in Greenfield Deployment

The objective of the first scenario is to evaluate whether, and to what extent, `Ce-Fi` can maintain performance under the influence of FH-delay. This is achieved by measuring the throughput for varying FH-delays and comparing the results to both the 802.11 DCF and the approach proposed by Okamoto *et al.* [8], within a greenfield deployment.

The corresponding results are presented in Figure 6. The first notable observation is that the throughput of `Ce-Fi` and the approach by Okamoto *et al.* are virtually identical. This is expected, as both rely on the concept of NAV extension and piggybacking data. Both the total throughput and the AP throughput exhibit a graceful degradation with increasing FH-delay. This behavior can be approximated as seen in Equation (3) and Equation (4) for the total and the AP's throughput, respectively:

$$\mathrm{Throughput_{Total}}(FH) \approx 219 \times e^{-0.00253 \times FH} \qquad (3)$$

$$\mathrm{Throughput_{AP}}(FH) \approx 110 \times e^{-0.00253 \times FH} \qquad (4)$$

The decline in performance with increasing FH-delay is primarily due to the reduced network efficiency, as the extended NAVs cause longer waiting periods, growing with the delay.

The performance of the 802.11 DCF is more significantly impacted due to an increase in collisions. This effect becomes particularly apparent at FH-delays between 40–50 µs, where (as described by Okamoto *et al.* [8]) ACK frames begin to collide with data frames, severely degrading network performance. It is also worth noting that the AP's throughput in 802.11 DCF is significantly lower when compared to `Ce-Fi` or the approach by Okamoto *et al.* This is because `Ce-Fi` inherently prioritizes the DL relative to the UL, resulting in more balanced UL and DL traffic. Moreover, overall throughput improves because the DL does not require its own contention phase; instead, it piggybacks on the existing UL contention. This improves channel efficiency and overall network performance.

C. Impact of FH-Delay with OBSS

In the second scenario, the impact of the FH-delay on network performance is evaluated, when a FH affected network coexists with a neighboring OBSS. As `Ce-Fi` and the proposal

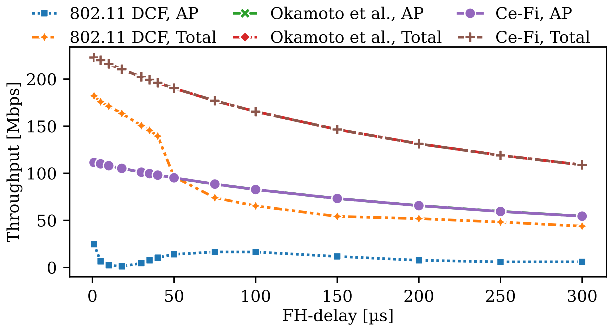

Figure 6. Throughput comparison for different channel access mechanisms in a greenfield deployment for varying FH-delays

by Okamoto *et al.* behave identically in full buffer scenarios, the results for the latter are omitted.

Figure 7a shows that both BSSs are able to communicate, regardless of the FH-delay when the FH-affected network uses `Ce-Fi`. This is possible because the STAs can still contend for the channel on behalf of the AP. By doing so and correctly setting the NAV, communication can proceed without collisions after a successful contention, even under high FH-delays. It also shows that a `Ce-Fi`-based BSS has a competitive advantage over neighboring networks, as the `Ce-Fi`-AP does not need to contend for DL traffic separately. Consequently, the `Ce-Fi`-BSS can dominate the channel and occupy it for longer periods than neighboring BSSs, which could be considered unfair. Additionally, as the FH-delay increases, the resulting longer NAV-based waiting times cause the overall channel occupancy to rise, penalizing not only the `Ce-Fi`-BSS but also any adjacent OBSSs. This issue could be mitigated by adopting a probabilistic access control scheme, as suggested by Funabiki *et al.* [10], where the `Ce-Fi`-network is restricted to transmit DL e.g. with a probability of $\frac{1}{\#\text{STAs}}$.

Figure 7b illustrates that the usage of 802.11 DCF for the FH-affected network causes a near complete loss of throughput for FH-delays larger than around 50 µs. Here, the OBSS monopolizes the channel, as in the FH affected network, due to faulty CSMA/CA, neither DL nor the transmission of ACK to the STAs is possible. This increases the CW, eventually causing the FH-affected network to cease participating in the channel access altogether.

D. *Impact of UL-Intensity in Greenfield Deployment*

The goal of the third scenario is to analyze the impact of the UL-intensity on the DL throughput in a greenfield deployment. The UL-intensity is defined as the number of application-layer packets per second passed to the MAC-layer for transmission. The inter-arrival time between two application-layer packets is modeled using an exponential distribution.

Figure 8a makes clear that without FH-delay and for a high UL-intensity (left part of the figure), the approach by Okamoto *et al.* and `Ce-Fi` yield identical DL-performance. In contrast, when using the 802.11 DCF the DL-performance is significantly lower due to the earlier discussed, increased channel access probability for DL in `Ce-Fi`, which scales

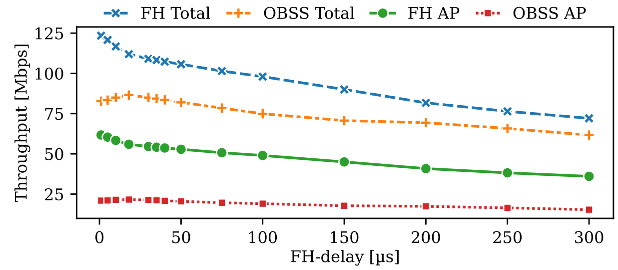

(a) `Ce-Fi`

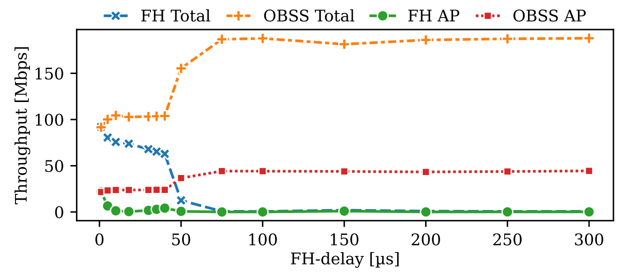

(b) IEEE 802.11 DCF

Figure 7. Throughput comparison of OBSS and centralized Wi-Fi BSSs using different channel access mechanisms under varying FH-delay

with the number of STAs. However, as the UL transitions from high to low intensity (between 100–200 frames/s), the approach by Okamoto *et al.* experiences immense throughput losses due to the lack of piggybacking opportunities, so that the DL approaches zero for very low UL-intensities. As seen in the results, `Ce-Fi`'s FT mechanism mitigates the lack of UL so that the overall throughput is comparable to 802.11 DCF. The cost of transmitting an additional NDF using `Ce-Fi` compared to 802.11 DCF is negligible, as it carries no payload and thus requires little airtime due to its small size.

Figure 8b shows the results for the same scenario, yet this time with FH-delay of 100 µs. In a low UL scenario, 802.11 DCF outperforms `Ce-Fi`, as the overhead of `Ce-Fi` increases with the FH-delay. Since all transmitted packets must traverse the FH, the actual channel utilization decreases, leading to reduced throughput. This also increases the cost of transmitting a NDF because while it is small, the FH delays the reception and hence decreases channel utilization. In general, introducing a FH-delay shifts the throughput curve downwards due to longer waiting times and reduced network efficiency. The FH-delay also shifts the curve to the right which can be attributed to the slower queue depletion, as it takes longer for packets to be acknowledged. With 802.11 DCF, the FH-delay additionally increases the likelihood of collisions, which causes packets to remain in the queue even longer, further decreasing the UL queue reduction rate.

E. *Impact of UL-Intensity on Throughput with OBSS*

In the fourth scenario, the effect of the UL-intensity on DL with a neighboring OBSS is investigated.

Figure 9a shows the effect of the UL-intensity on the DL for different channel access mechanisms without a FH-delay. The results show that the total throughput of the OBSS increases with decreasing UL, when 802.11 DCF or the approach by

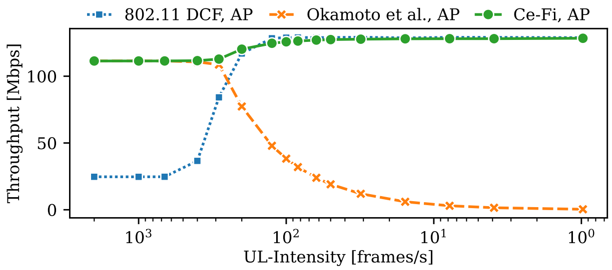

(a) Without FH-delay

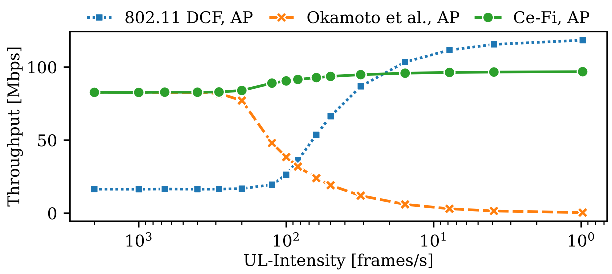

(b) 100 µs FH-delay

Figure 8. Throughput comparison of a Wi-Fi system using different channel access mechanisms in a greenfield scenario with varying UL-intensities

Okamoto *et al.* is used. This is because with decreasing UL, there will be fewer piggybacking opportunities when using the approach by Okamoto *et al.* while with 802.11 DCF, there are overall fewer STAs contending, resulting in a better throughput for both, the OBSS and the DL of the centralized network.

When using `Ce-Fi`, the DL remains relatively stable, since the STAs transition from sending UL to NDFs, thus creating piggybacking opportunities for DL traffic. The total throughput of the OBSS reaches a local maximum at a UL-intensity of around 100 frames/s, where UL traffic in the centralized BSS is so regular, that the STAs often have an extended backoff before transmitting a NDF, so that the OBSS has a slight advantage in channel access. The DL of the centralized network has again an advantage over the OBSS, since all STAs content for the channel to enable DL. With very low UL-intensity, the DL is even superior to the DL at a high UL-intensity, since the NDF occupies less air time than does UL data. As mentioned before, this could be mitigated by decreasing the probability of accessing actually using a piggybacking opportunity for DL to $\frac{1}{\#\text{STAs}}$.

The results with a FH-delay of 100 µs are illustrated in Figure 9b. With a OBSS and a FH-delay, the 802.11 DCF mechanism fails to access the channel, therefore resulting in complete loss of DL. The overall trajectories with `Ce-Fi` and the approach by Okamoto *et al.* [8] remain the same, however, the curves moved down and right again since the increased waiting times influence both BSS negatively.

F. Impact of DL-Intensity on STA Overhead

In the fifth scenario, we investigate the overhead imposed on the STAs associated with the `Ce-Fi` BSS by the FT mechanism. For that, we examine how the DL-intensity, defined analogously to the UL-intensity, affects the number of NDFs sent per STA across different BSS sizes. The number of

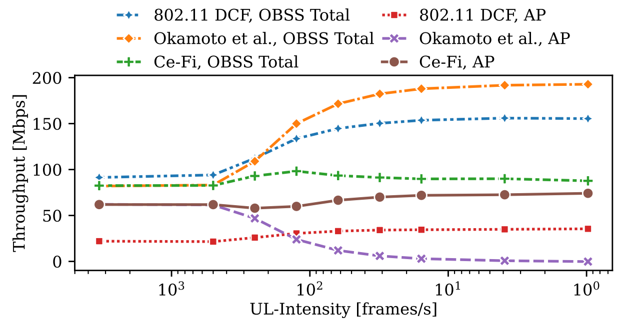

(a) Without FH-delay

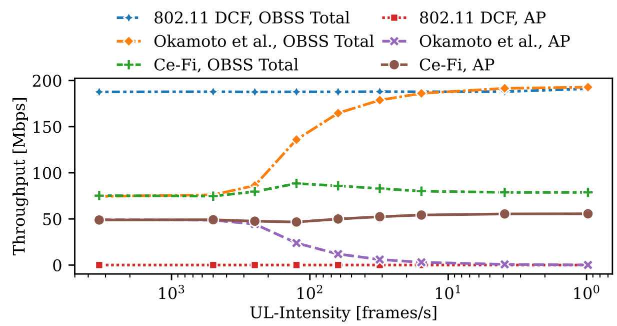

(b) With a FH-delay of 100 µs

Figure 9. Throughput comparison of a Wi-Fi system using different channel access mechanisms in concurrence with OBSS and varying UL-intensity

transmitted NDF frames serves as a proxy to evaluate the additional overhead. The STA have no UL to transmit.

The results are shown in Figure 10. With only one STA associated to the `Ce-Fi` BSS, the overhead steadily declines starting at a DL-intensity of approximately 1000 frames/s, decreasing from around 1460 frames/s to approximately 100 frames/s at very low DL-intensity. In scenarios with 5 or 10 STAs, the initial number of transmitted frames is lower, 360 frames/s and 190 frame/s, respectively, because the STAs compete for channel access, distributing the total number of NDF among them. The observed increase in transmitted NDFs at a DL-intensity of around 500 frames/s can be attributed to the fact that not every NDF is followed by a DL transmission. As a result, the channel remains idle more often, which increases the amount of time in which STAs content for channel access. This leads to a higher number of NDF transmissions while the DL-intensity remains high enough such that, in accordance with the mechanism described in Section V-C, the CW backoff is not yet fully reset across all stations. As the DL-intensity continues to decrease, the number of transmitted NDF also drops, eventually approaching around 100 frames/s. The reason why the number of NDF does not drop further is the coupling of the maximum CW offset to the CW_{max} value, which is typically set to 1023. Even in the absence of any DL, in a 1 STA scenario with the current $CW = CW_{min}$, the STA will transmit at the latest after $CW_{max} + CW_{min} = 1023 + 15 = 1038$ slot-times. Given a slot-time of $t_{slot} = 9\,\mu s$, this results in a maximum waiting time of 9376 µs (including a DIFS of 34 µs) before transmitting the next NDF, which in the worst case (always waiting the full backoff period) equates to approximately 100 frames/s. This behavior

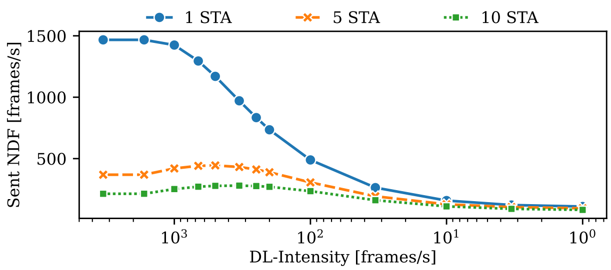

Figure 10. Number of transmitted NDF per DL-intensity in a greenfield scenario without UL and a FH-delay of 100 µs

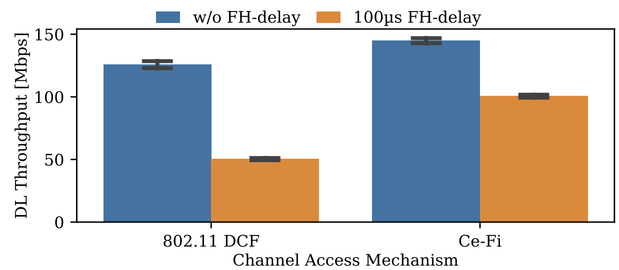

Figure 11. DL throughput for different channel access mechanisms when using TCP in a greenfield deployment

could be altered by adjusting CW_{max} or by decoupling the maximum CW_{offset} from CW_{max} and defining an independent maximum value. However, this would increase the time it takes for the AP to resume communication following a low DL period.

G. Impact of Using TCP on DL Throughput

In the aforementioned scenarios, UDP was used as the transport protocol. Regarding the issue of insufficient piggybacking opportunities for DL when using `Ce-Fi` under low UL traffic conditions, UDP relies on the presence of NDFs. In contrast, with the transmission control protocol (TCP), every frame sent for DL is acknowledged by a TCP-ACK, which can be used to piggyback additional DL.

In this scenario, each STA establishes a TCP connection with the AP and expects continuous DL traffic. Figure 11 shows that the achieved DL throughput using TCP with `Ce-Fi` and without FH-delay is approximately 14 % higher than with 802.11 DCF, owing to the improved network efficiency resulting from reduced waiting times through piggybacking data onto TCP-ACKs. When the system is affected by FH-delay, the advantage becomes even more apparent. While the throughput drops by 60 % with 802.11 DCF, mainly due to collisions between UL and DL, it only degrades by 31 % when using `Ce-Fi`, primarily due to increased waiting times caused by the FH-delay.

Thus, when using TCP and `Ce-Fi`, the system becomes less dependent on NDFs, thanks to the transmission opportunities created by TCP-ACKs. However, NDFs remain crucial in mixed traffic scenarios, e.g., when there is no UL traffic but both UDP and TCP DL traffic are present, as would be the case during simultaneous live-TV streaming and movie downloading. In such cases, the UDP DL traffic does not generate new piggybacking opportunities, making NDFs necessary.

VII. Discussion and Conclusion

In this paper, we have presented `Ce-Fi`, an architecture designed to enable centralized Wi-Fi with a packet-fronthaul. We have demonstrated that `Ce-Fi` facilitates high throughput communication, even in systems affected by large FH-delays or when coexisting with OBSS for both, UDP and TCP traffic. This is achieved through a combination of NAV extension and a piggybacking scheme, which, thanks to an forced traffic scheme, operates effectively even in scenarios with low UL traffic.

The results show that the proposed `Ce-Fi` architecture distinguishes itself from the existing literature by maintaining performance even under high FH-delays, in the presence of neighboring networks, and with low UL traffic available for piggybacking DL traffic. However, the presented `Ce-Fi` architecture is not without its drawbacks. The most prominent of these is that it requires modifications to the protocol stack, not only to the AP side but also to the STA, which is in contrast to the solution presented by Funabiki *et al.* [10] where only timeouts needed to be adjusted. Moreover, `Ce-Fi` imposes additional strain on the STAs, as many of them need to expend energy transmitting NDFs, even if they have no actual data to transmit. This issue could potentially be addressed in the future by considering device classes or battery status, allowing stations to decide whether to participate in transmitting NDF. Additionally, even with the current implementation, STAs could announce their intention to enter a sleep state, thus saving energy.

Please note that null data frames are not a novel concept in this protocol, as they were originally designed for other purposes, such as signaling a sleeping state or sending keep-alive messages, which could still interfere with the setting of the CW_{offset}. Yet, it is worth noting that the expected number of NDFs to be transmitted is relatively low.

Also, our approach relies on virtual channel reservation, which results in the channel being physically unoccupied for extended periods (depending on the actual FH-delay).

Regarding the simulations conducted, we focused on a static, deterministic FH-delay, whereas the actual delay in a packet-FH can vary and exhibit significant jitter. To address this, when determining the FH-delay of a system, it should be set to the worst-case FH-delay measured, in order to mitigate the negative effects of jitter. This approach may lead to unnecessarily high channel reservation, resulting in worse channel occupancy, but it will also contribute to fewer collisions and, consequently, fewer retransmissions.

VIII. Acknowledgments

This work was supported by the Federal Ministry of Education and Research (BMBF, Germany) within the 6G Research and Innovation Cluster 6G-RIC under Grant 16KISK020K as well as by the German Research Foundation (DFG) within the project ML4WiFi under grant DR 639/28-1.

References

[1] "Cloud RAN to Take Share from Purpose-Built RAN." Accessed: 2025-04-28, Dell'Oro Group. (2025), [Online]. Available: https://www.delloro.com/news/cloud-ran-to-take-share-from-purpose-built-ran/.

[2] A. Checko, H. L. Christiansen, Y. Yan, L. Scolari, G. Kardaras, M. S. Berger, and L. Dittmann, "Cloud RAN for Mobile Networks—A Technology Overview," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 1, pp. 405–426, Jan. 2015.

[3] K. Sui, S. Sun, Y. Azzabi, X. Zhang, Y. Zhao, J. Wang, Z. Li, and D. Pei, "Understanding the Impact of AP Density on WiFi Performance Through Real-World Deployment," in *IEEE International Symposium on Local and Metropolitan Area Networks (LANMAN 2016)*, Rome, Italy: IEEE, Jun. 2016, pp. 1–6.

[4] S. Deronne, V. Moeyaert, and S. Bette, "Impact of the Slottime Parameter Value on the MAC Performances in IEEE 802.11 Wireless Systems Using Radio-over-Fiber Technology," in *17th IEEE Symposium on Communications and Vehicular Technology in the Benelux (SCVT 2010)*, vol. 1, Enschede, Netherlands: IEEE, Dec. 2010, pp. 1–6.

[5] S. Deronne, V. Moeyaert, and S. Bette, "Analysis Of The MAC Performances In 802.11g Radio-Over-Fiber Systems," in *18th IEEE Symposium on Communications and Vehicular Technology in the Benelux (SCVT 2011)*, Gent, Belgium: IEEE, Nov. 2011, pp. 1–5.

[6] S. Deronne, V. Moeyaert, and S. Bette, "Simulation of 802.11 radio-over-fiber networks using ns-3," in *6th International ICST Conference on Simulation Tools and Techniques (SimuTools 2013)*, Cannes, France: ICST, Mar. 2013, pp. 190–194.

[7] M. Tanaka, D. Umehara, M. Morikura, N. Otsuki, and T. Sugiyama, "New Throughput Analysis of Long-distance IEEE 802.11 Wireless Communication System For Smart Grid," in *IEEE International Conference on Smart Grid Communications (SmartGridComm 2011)*, Brussels, Belgium: IEEE, Oct. 2011, pp. 90–95.

[8] Y. Okamoto, M. Morikura, T. Nishio, K. Yamamoto, F. Nuno, and T. Sugiyama, "Throughput and QoS Improvement of Wireless LAN System Employing Radio Over Fiber," in *2014 XXXIth URSI General Assembly and Scientific Symposium (URSI GASS)*, Beijing, China: IEEE, Jul. 2014, pp. 1–4.

[9] T. Nishio, K. Funabiki, M. Morikura, K. Yamamoto, D. Murayama, and K. Nakahira, "MAC Protocol for Improving Throughput and Balancing Uplink/Downlink Throughput for Wireless Local Area Networks with Long Propagation Delays," *IEICE Transactions on Communications*, vol. E100.B, Nov. 2016.

[10] K. Funabiki, T. Nishio, M. Morikura, K. Yamamoto, D. Murayama, and K. Nakahira, "ATRAS: adaptive MAC protocol for efficient and fair coexistence between radio over fiber-based and CSMA/CA-based WLANs," *EURASIP Journal on Wireless Communications and Networking*, vol. 2017, pp. 405–426, Dec. 2017.

[11] A. Valkanis, P. Nicopolitidis, and G. Papadimitriou, "A Hybrid Link-TDMA MAC Protocol for Conventional and Radio over Fiber WLANs," *Wireless Communications and Mobile Computing*, vol. 2020, no. 1, 2020.

[12] B. B.S and S. Azeem, "A Survey on Increasing the Capacity of 5G Fronthaul Systems Using RoF," *Optical Fiber Technology*, vol. 74, 2022.

[13] B. Kalantari-Sabet, M. Mjeku, N. J. Gomes, and J. E. Mitchell, "Performance Impairments in Single-Mode Radio-Over-Fiber Systems Due to MAC Constraints," *Journal of Lightwave Technology*, vol. 26, pp. 2540–2548, Aug. 2008.

[14] P. Dely, J. Vestin, A. Kassler, N. Bayer, H. J. Einsiedler, and C. Peylo, "CloudMAC — An OpenFlow Based Architecture for 802.11 MAC Layer Processing in the Cloud," in *IEEE Globecom Workshops*, Anaheim, CA: IEEE, Dec. 2012, pp. 186–191.

[15] Y. Kim, G. Kim, and H. Lim, "Cloud-Based Wi-Fi Network Using Immediate ACK in Uplink Data Transmissions," *IEEE Access*, vol. 6, pp. 37 045–37 054, May 2018.

[16] IEEE, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," IEEE, Std 802.11-2016, Dec. 2016.

[17] M. Gronovius. "Packet fronthaul for Cloud RAN and Open RAN." Accessed: 2025-04-23, Ericsson. (Mar. 2025), [Online]. Available: https://www.ericsson.com/en/blog/2025/3/packet-fronthaul-for-cloud-ran-open-ran.

[18] D. Chitimalla, K. Kondepu, L. Valcarenghi, M. Tornatore, and B. Mukherjee, "5G Fronthaul-latency and Jitter Studies of CPRI Over Ethernet," *Journal of Optical Communications and Networking*, vol. 9, no. 2, pp. 172–182, 2017.

[19] J. Kant Chaudhary, J. Francis, A. N. Barreto, and G. P. Fettweis, "Latency in the Uplink of Massive MIMO CRAN With Packetized Fronthaul: Modeling and Analysis," in *IEEE Wireless Communications and Networking Conference (WCNC 2019)*, Marrakesh, Morocco: IEEE, Apr. 2019, pp. 1–7.

[20] C. Laskos, A. Zubow, and F. Dressler, "Latency Analysis of SDR-based Experimental C-RAN / O-RAN Systems," in *IEEE International Conference on Communications (ICC 2025), Workshop on Tactile Internet with Human-in-the-Loop (TIHL 2025)*, Montréal, Canada: IEEE, Jun. 2025.

[21] A. Zubow, S. Zehl, and A. Wolisz, "BIGAP — Seamless handover in high performance enterprise IEEE 802.11 networks," in *IEEE/IFIP Network Operations and Management Symposium (NOMS 2016)*, Istanbul, Turkey: IEEE, Apr. 2016.

[22] "OMNeT++ Discrete Event Simulator." Accessed: 2025-04-13, OMNeT++. (2025), [Online]. Available: https://omnetpp.org/.