Toward Enabling High-Five Over WiFi A Tactile Internet Paradigm

Gokhale, Vineet; Eid, Mohamad; Kroep, Kees; Prasad, R. Venkatesha; Rao, Vijay S.

DOI

[10.1109/MCOM.007.2100271](https://doi.org/10.1109/MCOM.007.2100271)

Publication date

2021

Document Version

Accepted author manuscript

Published in

IEEE Communications Magazine

Citation (APA)

Gokhale, V., Eid, M., Kroep, K., Prasad, R. V., & Rao, V. S. (2021). Toward Enabling High-Five Over WiFi: A Tactile Internet Paradigm. *IEEE Communications Magazine*, 59(12), 90-96.
<https://doi.org/10.1109/MCOM.007.2100271>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Toward Enabling High-Five Over WiFi: A Tactile Internet Paradigm

Vineet Gokhale[†], Mohamad Eid[‡], Kees Kroep[†], R. Venkatesha Prasad[†], and Vijay Rao^{*}

[†]Delft University of Technology, The Netherlands

[‡]New York University Abu Dhabi, UAE

^{*}Cognizant Technology Solutions, The Netherlands

Abstract—The next frontier for immersive applications is enabling sentence over the Internet. *Tactile Internet (TI)* envisages transporting skills by providing Ultra-Low Latency (ULL) communications for transporting *touch senses*. In this work, we focus our study on the first/last mile communication, where the future generation WiFi-7 is pitched as the front-runner for ULL applications. We discuss a few candidate features of WiFi-7 and highlight its major pitfalls with respect to ULL communication. Further, through a specific implementation of WiFi-7 (*vanilla WiFi-7*) in our custom simulator, we demonstrate the impact of one of the pitfalls – standard practice of using jitter buffer in conjunction with frame aggregation – on TI communication. To circumvent this, we propose *Non-Buffered Scheme (NoBuS)* – a simple MAC layer enhancement for enabling TI applications on WiFi-7. NoBuS trades off packet loss for latency enabling swift synchronization between the master and controlled domains. Our findings reveal that employing NoBuS yields a significant improvement in RMSE of TI signals. Further, we show that the worst-case WiFi latency with NoBuS is 3.72 ms – an order of magnitude lower than vanilla WiFi-7 even under highly congested network conditions.

Index Terms—Tactile Internet, WiFi-7, ultra-low latency

I. INTRODUCTION

The Covid-19 pandemic has dramatically increased our reliance on interactive audio-video teleconferencing. The disruptive vision of *Tactile Internet (TI)* takes this interactivity a step further by enabling transportation of physical actions coupled with *haptic* (touch) feedback over long distances [1]. This facilitates exchange of touch-based gestures, such as ‘*High-Fives*’, hugs, and handshakes, revolutionizing our ability to be sentient over the Internet.

Such an immersive experience enables seamless *skill transfer*, manifesting the potential to revolutionize many fronts of human lives. TI will lay a strong foundation for accelerating the progress of the *Industry 4.0* revolution that enables industrial processes to be remotely executed through controlling telerobots. Critical healthcare services can be democratized through telesurgery; surgeons can conduct medical procedures over long distances with the same precision and speed as a conventional surgery. TI will also significantly redefine several other sectors such as education, entertainment, e-commerce, and disaster management.

Fig. 1 depicts a high-level representation of various domains and use-cases of TI. The human operator and the teleoperator in the *master* and *controlled domains*, respectively,

are equipped with sensor and actuator devices. Such devices include robotic arms, exoskeletons, and haptic suits. The motion commands (position, velocity, and torque) generated by the operator’s actions are communicated over the *first mile* to the network domain that facilitates their fast communication. These commands are then delivered over the *last mile* to the controlled domain where they are utilized by the teleoperator to replicate operator’s actions. The dynamics of the controlled domain, captured as force, audio, and video, are fed back to the operator.

The promising benefits of TI accompany several challenges in the context of sensing, actuation, and communication. The foremost challenge with respect to communication is the ultra-low latency (ULL) requirement of round-trip latency of sub-10 ms [2]. High quality of teleoperation can only be maintained if the haptic feedback corresponding to an action is provided to the operator within the ULL deadline. While innovations in 5G are making giant strides towards meeting ULL demands, designing specialized solutions for TI would further accelerate the progress towards realizing seamless teleoperation. Additionally, the TI community also recommends an ultra-reliability requirement of up to 99.9999 % [2]. However, unlike the ULL requirement, which is experimentally quantified, the ultra-reliability requirement has no reasonable substantiation and is only a speculation of the necessary performance guarantee. In fact, a recent work [3] that involved rigorous subjective experiments reveals that even a significant amount (nearly 50 %) of packet losses barely introduces any disturbance to the users. This indicates the possibility of ultra-reliability being an overkill for at least some TI applications that have the same characteristics as the TI task investigated in work [3]. Extending this observation to a broader range of TI applications necessitates further systematic studies. Until the TI reliability requirement is well-established, treating ultra-reliability as a key requirement may lead to over-designing TI solutions. Hence, in this work, we focus on the ULL requirement of TI. Note that in this work, we provide objective measurements that align with the claims on reliability requirement made in [3]. This further justifies our decision to focus on ULL requirement only.

Satisfying TI’s ULL constraint demands significant advancements at every hop in the network including the network domain and first/last mile. In this work, we focus on the first/last mile communication that presents several interesting

Fig. 1: Schematic representation of TI depicting the master, network, and controlled domains as well as a few popular use-cases.

research opportunities. The ubiquity of WiFi makes it one of the front-runners in the first/last mile communication. Further, extensive efforts are underway to enable ULL applications through the future generation *WiFi-7* (IEEE 802.11be). A few works, such as [4], [5], have conducted simulation-based studies to provide a preliminary understanding of the latency performance of some candidate features of *WiFi-7*. However, these works are generic and do not consider the specifics of TI applications.

In this work, we take the first step towards answering the following questions – (i) How well are the candidate features of the future *WiFi-7* suited for TI applications? (ii) How to overcome their limitations?

Our contributions in this work are the following. First, we provide an overview of a few candidate features of *WiFi-7*. We highlight major pitfalls of *WiFi-7* that might constrain its usage for TI communications. Secondly, we investigate one of the pitfalls by considering a specific implementation of *WiFi-7* involving the 320 MHz bandwidth candidate feature, Orthogonal Frequency Division Multiple Access (OFDMA), and frame aggregation in our open-sourced custom MAC simulator. We call this implementation as *vanilla WiFi-7* and take it as the baseline for our performance evaluations. By using realistic TI data traces, we objectively demonstrate the severe limitations of *vanilla WiFi-7* on the quality of teleoperation. To address this issue, we propose *Non-Buffered Scheme (NoBuS)* – a MAC layer enhancement tailored for TI applications to facilitate swift synchronization between the master and controlled domains. Our experiments demonstrate that NoBuS can enable TI communication with a worst-case WiFi latency of approximately 3.72 ms – an order of magnitude lower than that of *vanilla WiFi-7*. Further, we demonstrate that employing NoBuS yields a significant reduction in RMSE of the reconstructed signal compared to *vanilla WiFi-7*.

II. A WiFi PRIMER

For providing a concrete idea of WiFi developments, we discuss some existing features of WiFi systems crucial for understanding our work (Sec. II-A) and also discuss candidate features [6] that are expected to play key roles in driving WiFi's ULL objectives (Sec. II-B). Further, we present some of the potential pitfalls of *WiFi-7* (Sec. II-C).

A. Existing features of WiFi

1) *Frame aggregation*: This is a key bandwidth-saving mechanism of WiFi introduced in 802.11n. If a device's (access point (AP) or station (STA)) data generation rate exceeds its transmission rate, the MAC layer aggregates the MAC Protocol Data Units (MPDUs) into an aggregated MPDU (AMPSDU). AMPDUs are transmitted once the wireless medium becomes accessible. Transmitting an AMPDU instead of separate MPDUs leads to a substantial reduction in the protocol header overhead as well as the medium access latency.

2) *OFDMA*: This is a multiple access feature of WiFi introduced in 802.11ax (WiFi-6) [7]. Unlike in legacy WiFi systems (pre WiFi-6), where a STA utilizes the entire channel bandwidth during the medium access (single-user mode), OFDMA enables simultaneous medium access by multiple STAs (multi-user mode) [4]. When a STA wins the contention, it employs the single-user mode for AMPDU transmissions, akin to legacy WiFi. OFDMA is triggered only when AP wins the contention by dividing the channel bandwidth into several resource units (RUs). OFDMA is employed on both uplink (UL) and downlink (DL).

DL-OFDMA: The AP employs DL-OFDMA to simultaneously transmit to multiple STAs by assigning a dedicated RU per STA. This assignment is communicated with the STAs so they can restrict listening to their own RUs.

UL-OFDMA: The AP also provisions UL transmissions during its medium access through UL-OFDMA. In a simplistic implementation, the AP could trigger UL-OFDMA once it completes DL-OFDMA. The AP allows for UL-OFDMA using two modes – scheduled access (SA) and random access (RA). In the SA mode, the AP assigns RUs based on the transmit queue of STAs to conduct collision-free UL transmission. In the RA mode, the RU access by the STA is random based on a backoff counter. The AP can employ the two modes in conjunction, in which case a fraction of the total RUs is dedicated for SA. The RA STAs contend for the remaining RUs.

B. Candidate features of WiFi-7

1) *PHY layer enhancements*: To enhance multi-user, multiple-input, multiple-output (MU-MIMO) capacity, WiFi-

Fig. 2: Schematic representation of a generic, WiFi-7 enabled TI setup used in this work. The MAC transmit queue of AP is depicted to contrast the frame aggregation of vanilla WiFi-7 and *non-buffering* principle of NoBuS.

7 proposes to increase the number of spatial streams from 8 to 16. Additionally, the maximum bandwidth will be doubled from 160 to 320 MHz. Furthermore, a higher modulation and coding scheme, 4096-QAM, will likely be included in the standards.

2) *Multi-Link Operation (MLO)*: As opposed to previous versions of WiFi, MLO enables WiFi-7 devices to simultaneously contend across multiple frequency bands, namely 2.4, 5, and 6 GHz. This allows them to access the band that becomes available at the earliest. Preliminary investigations show MLO's potential to reduce the worst-case medium access latency by an order of magnitude compared to single-link approach [4].

3) *Separate AC for ULL streams*: Currently, WiFi provides a different access categories (AC) for audio and video streams, with voice having the highest priority. Some WiFi-7 submissions are proposing to introduce a new AC for ULL streams with a priority higher than voice's [8]. This is expected to significantly reduce the medium access latency.

4) *Preemptive scheduling*: Longer MAC frames, such as those carrying video samples, can hold up subsequent frame transmissions from other devices leading to higher latency. While this is acceptable for conventional voice and video applications, it could be catastrophic for TI applications due to their ULL constraint. To circumvent this issue, WiFi-7 candidate solutions propose to pause an ongoing transmission of a lower priority frame in order to pave way for transmission of a higher priority frame when generated.

C. WiFi-7 Pitfalls

Despite the above proposed enhancements, we identify three major inadequacies of WiFi-7 in supporting TI applications.

1) *Standard practice causes sluggishness*: Typically, humans are much more sensitive to audio-video loss than haptic loss. A loss of up to 1 % for conventional audio-video communication is acceptable [9], whereas even up to 90 % haptic loss can go unnoticed [10]. On the contrary, if the haptic feedback is not provided within ULL limits, human operators cannot

continue to teleoperate effectively [2]. Hence, haptic feedback is *latency-sensitive*.

As a standard practice, WiFi utilizes frame aggregation (Sec. II-A1) to serve loss-sensitive applications as it aims to maximize the number of samples delivered. A jitter buffer is employed at the receiver's application layer to cope with this bursty arrival of samples. This increases the waiting time of the samples (prior to display). This higher latency could result in WiFi-7 not supporting turnkey use-cases such as safety-critical TI applications. Although this approach was designed to boost the performance of conventional audio-video applications, it is a *shot-in-the-foot* for TI applications.

2) *Haptics AC*: While using a separate AC for haptic stream seems to be an addition of a highly-desired feature, it creates an undesired consequence: TI applications are characterized by a heavy temporal correlation between voice, video, and haptic streams. Hence, maintaining inter-stream synchronization is crucial during media display. Therefore, the medium access should also be driven by the constraints for inter-stream synchronization, the lack of which may heavily obstruct TI applications over WiFi-7.

3) *Preemptive scheduling*: This feature, although beneficial for ULL applications, needs a thorough investigation, especially when employed in conjunction with a new AC. Video frames may be continually preempted for haptic frames causing non-compliance to video latency deadlines. Therefore, such context-unaware preemptive scheduling may impede TI sessions.

The standard practice discussed in Sec. II-C1 will undoubtedly be a part of the WiFi-7 standards as it is based on existing WiFi features. As we show later in Sec. III-C, it introduces a few tens of milliseconds of additional latency. Hence, in this work, we take an important step towards addressing this issue.

III. WiFi-7 PERFORMANCE ASSESSMENT

A. Tactile Internet Setup

The TI setup considered in this work is generic across several use cases, such as telesurgery and a connected factory.

Parameter	Value
Frequency band	5GHz
Channel bandwidth	320MHz
Subcarrier spacing	78.125kHz
Access category	Voice (AC_VO)
Min. contention window	3
Max. contention window	7
Retransmission limit	7
MCS index	9
Slot size	9 μ s
AIFS	34 μ s
SIFS	16 μ s
Max. PPDU duration	5.4ms
Guard interval	0.8 μ s
Aggregation type	AMPDU

TABLE I: WiFi-7 configuration parameters used in our simulations.

Let us take the example of a connected factory. It consists of manufacturing processes remotely executed by operators via controlling teleoperators located inside the factory as shown in Fig. 2. For simplicity, we consider a scenario where the controlled domain is WiFi-7 enabled and each operator is controlling a separate teleoperator (the pairs are color-coded). Each WiFi-7 Basic Service Set (BSS) consists of an AP and multiple STAs (teleoperators). Since we focus on studying the intra-BSS dynamics, we depict a single BSS in Fig. 2. Since our focus is only on the WiFi network, we assume that other network links are uncongested and the WiFi network serves only TI streams. Background traffic, like loss-based TCP, is known to be detrimental for TI applications [11]. Therefore, it is important to carry out TI communication in tightly-controlled WiFi networks.

B. Network Simulations

1) *The Simulator*: Since this is a first-of-its-kind work investigating TI communication over WiFi-7, we resort to a simulation-based study. At the time of writing this paper, NS-3 has no complete OFDMA implementation in its official release. Therefore, we leverage a recently open-sourced 802.11ax MAC simulator [12] as a starting point of our implementation and add significant extensions to arrive at our *vanilla WiFi-7* implementation. It mainly consists of the candidate 320 MHz bandwidth of WiFi-7, fully functional UL- and DL-OFDMA, and frame aggregation features. In order to enable TI research community to conduct further experiments, we have open-sourced our simulator [13]. In developing our custom simulator, we make a reasonable assumption that the channel conditions and received power for the chosen modulation and coding scheme (MCS-9) result in a negligible path loss. The fixed MCS setup helps to minimize the impact of link adaptation, enabling us to isolate the effects of frame aggregation and jitter buffer on the signal reconstruction quality. This means that the PHY layer packet losses occur purely due to collisions. This is a common approach adopted by many existing works, such as [4]. The configuration details of the WiFi-7 network are given in Table I. In this work, we transmit TI traffic using voice AC (AC_VO) since it has the highest priority as per existing WiFi standards.

Fig. 3: Networked VE game using a haptic device. The displays on the left and right belong to the master and controlled domains, respectively.

Accounting for several sensors and HD-quality video, we consider a scenario where the UL (DL) traffic from (to) each STA is 20 Mbps. To reiterate, UL traffic is served each time a STA wins the contention as well as when AP schedules UL-OFDMA. On the other hand, the DL traffic is served only when AP schedules DL-OFDMA. While the UL and DL traffic are equal, the UL and DL access is asymmetric. Hence, in our implementation the AP schedules DL-OFDMA transmissions first. Only after the DL transmit queue is exhausted, UL-OFDMA is triggered. We use the standard 1 kHz sampling of haptic updates. Each haptic sample elicits an MPDU. To leverage the periodic nature of MPDU generation in reducing channel collisions, we employ only SA mode for UL-OFDMA.

For scheduling the STAs within UL- and DL-OFDMA, we use a simple approach for maximizing the bandwidth utilization, similar to the one recently proposed in [14]. We employ a *highest-first* approach where STAs with higher transmit queue occupancy are scheduled earlier than others. The AP sorts the STAs based on the decreasing order of their queue occupancy. It then divides them into disjoint groups of two if the number of STAs with data is lower than four, and groups of four otherwise. The channel bandwidth is then divided such that the number of RUs equals the group size. The group with the highest queue occupancy is served first. In case of DL-OFDMA, the AP simply looks at its transmit queue to determine the scheduling order of STAs. In case of UL-OFDMA, we leverage the periodic nature of MPDU generation to minimize the control overhead. The STAs inform about their queue occupancy to the AP by piggybacking the information in QoS control field of the packets – this is known as *unsolicited buffer status report* [7].

2) *Realistic TI data generation*: Since we intend to conduct performance assessment under realistic TI settings, we leverage a TI game that we developed in-house. In this, the human subject employs a Novint Falcon haptic device to move a rigid cube to a pre-determined target location inside a virtual environment (VE), as shown in Fig. 3. The VE runs on a remote server; further details on the VE setup are available in [15]. Through human-in-the-loop experiments, we record the operator's motion commands, force feedback, and graphics data generated during this interaction. We then utilize the

Fig. 4: Plot showing UL and DL latency and corresponding mean AMPDU size over a range of teleoperators (WiFi STAs) for vanilla WiFi-7 implementation.

recorded data traces in our network simulator, where each operator/teleoperator uses samples from these traces as the payload. This helps us to conduct realistic network simulations while using data generated through an actual TI setup.

C. Vanilla WiFi-7

1) *Packet generation*: In practice, the operator/teleoperator could have multiple sensors, each operating at a different sampling rate than the other. To minimize packetization delay under heterogeneous sampling rates, the application layer generates messages at the highest sampling rate of all sensors. A message carries the latest sample from each sensor and each message elicits an MPDU. Whenever a WiFi device gets medium access, it composes AMPDU by aggregating all MPDUs in its transmit queue corresponding to a destination. A packet is then transmitted consisting of an AMPDU. Essentially, in vanilla WiFi-7, each packet may consist of several samples from each sensor due to aggregation. The MAC transmit queue of vanilla WiFi-7 is depicted in Fig. 2.

2) *Performance Evaluation*: We begin by presenting UL and DL latency with increasing STA density (number of STAs) in Fig. 4. Note that since we are proposing an enhancement to WiFi-7, we measure only the latencies arising from WiFi layers – MAC and PHY. With twelve STAs, the worst-case DL and UL latencies are approximately 26.10 ms and 23.53 ms, respectively, resulting in a worst-case round-trip latency of 49.63 ms. To get an idea of the extent of frame aggregation, Fig. 4 also presents the mean AMPDU size which is the mean number of MPDUs per AMPDU. For low STA density, the AMPDU size is close to 1, indicating that the transmission rate can cope with the data generation rate without aggregation. With increasing STA density, the mean AMPDU size also starts to rise quickly and exceeds 12 and 8 on DL and UL, respectively, with twelve STAs. For 1 kHz haptic sampling rate, this signifies that each AMPDU carries samples worth more than 12 ms and 8 ms, respectively, corresponding to each sensor.

Note that the latency presented in Fig. 4 correspond to only channel access and does not take into account the latency due jitter buffer at the receiver. While the earliest sample in an AMPDU is displayed first, the later samples are queued up in

Fig. 5: Performance improvement of NoBuS in terms of position RMSE at teleoperator despite its lower fraction of transmitted samples in comparison with vanilla WiFi-7.

the jitter buffer for their turn to be displayed serially at 1 kHz. Accounting for this, the round-trip latency is even higher than above specified values. Note that for TI applications, it is crucial to contain the worst-case round-trip latency within the ULL bound (10 ms). It can be seen that this deadline gets violated with as few as five STAs. Therefore, with vanilla WiFi-7, TI applications are extremely vulnerable to the detrimental effects of high latency.

3) *Contemplating frame aggregation*: Since vanilla WiFi-7 maximizes the amount of samples delivered, the teleoperator can accurately replicate the operator's actions. On the other hand, the jitter buffer latency which is the additional time to synchronize with the operator makes the teleoperator sluggish. From a control systems perspective, it is pragmatic to display only the latest sample. If the teleoperator was programmed to display only the latest received sample while discarding earlier ones, then the synchronization would be achieved in a significantly shorter amount of time. This approach, however, trades off the reliability of teleoperation for latency.

IV. ENABLING TACTILE INTERNET OVER WiFi-7

In this section, we explore this latency-reliability trade-off through the design of a simple MAC layer enhancement that we call *Non-Buffered Scheme (NoBuS)*. We use the performance of vanilla WiFi-7 presented in Sec. III-C2 as the baseline for gauging the improvement of NoBuS.

A. Non-Buffered Scheme (NoBuS)

1) *Working Principle*: In existing WiFi systems, when packet loss occurs MAC layer retransmits the samples up to the retransmission limit (Table I). While attempting to retransmit the older samples increases latency (frame aggregation), the other extreme approach of completely disabling retransmissions would heavily deprive the teleoperator of even the latest samples. Hence, we require an in-between approach where the retransmission happens only until retransmission limit is reached or a new sample arrives – whichever is earliest.

Motivated by this, we design NoBuS as a simple MAC layer enhancement for enabling quick synchronization between the master and controlled domains. The working principle of

NoBuS is as follows: When a device wins the channel contention, it transmits only the latest MPDU. This means that if there are multiple MPDUs buffered for the same destination, then all the previous MPDUs are discarded. Essentially, we use sampling duration as the packet lifetime. This approach minimizes jitter buffer latency and transmission latency. Fig. 2 contrasts the MAC transmit queue of NoBuS with that of vanilla WiFi-7. Note that although we depict the MAC layer of AP alone, we implement NoBuS at the STAs also.

2) *Packet generation*: The application layer messages are generated as explained in Sec. III-C1. Since NoBuS replaces the older MPDU in the transmit queue with the latest, each packet consists of a single MPDU carrying the latest sample from each sensor.

B. Performance Evaluation

Due to space constraints, we present the results for DL only while noting that the UL measurements follow similar trend. As can be seen in Fig. 5, NoBuS transmits a significantly lower amount of samples (up to 50 %) than vanilla WiFi-7. This is because NoBuS drops samples in two scenarios: i) exceeding retransmission limit, and ii) proactive dropping of older samples. On the other hand, vanilla WiFi-7 drops samples only in the former scenario.

We now assess the impact of the above higher losses on the quality of signal reconstruction. To this end, we leverage the recorded TI data traces (Sec. III-B2). For the purpose of presentation, we consider only one axis of the 3D position data. The position coordinates in this axis vary predominantly in the range $[-10, 10]$ cm. We use RMSE as the objective metric. Despite the higher losses, it can be seen that the RMSE of NoBuS increases only marginally with the STA density. Additionally, NoBuS yields a significantly lower RMSE than vanilla WiFi-7, with a maximum reduction of up to 65 %. Further, in TI applications involving very fast hand movements, we expect vanilla WiFi-7 to perform even worse and the NoBuS improvement to be higher. The reason behind this high performance of NoBuS is its ability to minimize the overall latency. As per our simulations, the worst-case DL and UL latency of NoBuS are around 1.43 ms and 2.29 ms, respectively, leading to a worst-case round-trip latency of 3.72 ms. This is an order of magnitude lower than that of vanilla WiFi-7, and is also compliant to the necessary ULL limit of 10 ms even in highly dense environments. This is because the worst-case queuing latency of a transmitted sample is only 1 ms. An additional benefit of NoBuS is that the receiver's MAC layer can forward MPDUs to higher layers without waiting for the earlier ones. This avoids the latency arising from packet reordering.

From the standpoint of RMSE, the performance improvement of NoBuS reveals that fast synchronization between operator and teleoperator is way more crucial for high quality TI communication than maximizing the number of transmitted samples. This aligns very well with the claims made in [3] and hence highlights the need to focus on providing ULL service guarantees for TI applications rather than ultra-reliability. This opens up several opportunities and serves as an important insight for designing efficient TI systems.

Although we showed the efficacy of NoBuS in a WiFi network, the principle of *non-buffering* can be explored on wired or other wireless networks where schemes akin to frame aggregation are typically employed. While we expect NoBuS' ULL guarantee to facilitate seamless teleoperation, further experimental verification through psychophysical studies is paramount for a systematic substantiation of the efficacy of NoBuS. We intend to take this up as a future extension of this work.

C. Implementation details

Due to its simple working principle, the implementation of NoBuS on WiFi devices is quite straightforward. The MAC layer needs to be informed about the AC used by the TI stream. WiFi's default frame aggregation functionality should be disabled for that AC. At the STAs, the MAC layer can be configured to have a transmit queue of size 1 for the specific AC assigned for TI streams. The MAC layer should be programmed to simply replace the existing MPDU in the queue with the latest one. On the other hand, since the AP holds MPDUs for multiple STAs (for DL transmission), its queue is longer (ideally as long as the number of STAs). On arrival of an MPDU meant for DL transmission, the AP should identify the older MPDU to be replaced based on the destination MAC address. It is worth noting that NoBuS' operation does not entail any cross-layer designs. Therefore, it does not violate any design principles of the WiFi standards and can be easily incorporated with the existing WiFi systems. Further, NoBuS places negligible additional computational and memory requirements on the WiFi devices. In our opinion, at some point, enhancements will be done to support TI traffic on WiFi, and NoBuS can be one of the submissions.

V. VISTAS FOR FUTURE

- **Comprehensive evaluation**: While the RMSE measurements show the potential of NoBuS to facilitate high-fidelity signal reconstruction, subjective experiments are important for further substantiation. Another interesting research direction is to explore the applicability of *age of information* theory to further enhance the performance of NoBuS. Also, we would like to demonstrate the working of NoBuS on a commercial WiFi chipset. This helps us to understand the performance of NoBuS under realistic wireless conditions.
- **Scheduling framework**: Designing a transmission scheduling framework for haptic and video samples taking into consideration their temporal relationship is crucial. Further, doing so in presence of MLO and preemptive scheduling schemes is also a promising research avenue.

VI. CONCLUSIONS

Tactile Internet (TI) will open umpteen possibilities for novel immersive applications involving haptic feedback. To this end, advancements in every hop of the network are a must. In this context, the first/last mile is expected to be invariably served by WiFi-7. In a first-of-its-kind, we considered the performance evaluation of WiFi-7 under realistic TI application conditions. We focused on the standard practice of

using frame aggregation in conjunction with jitter buffer and experimentally demonstrated its negative impact on the quality of teleoperation. Through the design of *NoBuS*, we showed remarkable improvement in latency and RMSE compared to vanilla WiFi-7. The promising findings of our work show potential for enabling TI communications over WiFi-7.

ACKNOWLEDGMENT

This work has been undertaken in the *Internet of Touch* project sponsored by Cognizant Technology Solutions and Rijksdienst voor Ondernemend Nederland under PPS O&I.

REFERENCES

- [1] G. P. Fettweis, "The tactile internet: Applications and challenges," *IEEE Vehicular Technology Magazine*, vol. 9, no. 1, pp. 64–70, 2014.
- [2] O. Holland *et al.*, "The iee 1918.1 "tactile internet" standards working group and its standards," *Proceedings of the IEEE*, vol. 107, no. 2, pp. 256–279, 2019.
- [3] H. Kroep *et al.*, "Etvo: Effectively measuring tactile internet with experimental validation," *arXiv preprint arXiv:2107.05343*, 2021.
- [4] G. Naik *et al.*, "Can wi-fi 7 support real-time applications? on the impact of multi link aggregation on latency," in *IEEE International Conference on Communications (ICC)*, 2021.
- [5] T. Adame *et al.*, "Time-sensitive networking in iee 802.11 be: On the way to low-latency wifi 7," *Sensors*, vol. 21, no. 15, p. 4954, 2021.
- [6] E. Khorov *et al.*, "Current status and directions of iee 802.11 be, the future wi-fi 7," *IEEE Access*, vol. 8, pp. 88 664–88 688, 2020.
- [7] "Wireless lan medium access control (mac) and physical layer (phy) specifications: Enhancements for high efficiency wlan," IEEE standard 802.11ax, October 2020, accessed: 2021-07-12.
- [8] S. Kim *et al.*, "Latency enhancement for eht," *document IEEE 802.11-19/1524r0*, Sep. 2019.
- [9] A. Marshall *et al.*, "Providing qos for networked peers in distributed haptic virtual environments," *Advances in Multimedia*, 2008.
- [10] E. Steinbach *et al.*, "Haptic codecs for the tactile internet," *Proceedings of the IEEE*, vol. 107, no. 2, pp. 447–470, 2018.
- [11] V. Gokhale *et al.*, "On qos-compliant telehaptic communication over shared networks," *Computer Networks*, 2019.
- [12] Y. Daldoul *et al.*, "Performance evaluation of ofdma and mu-mimo in 802.11 ax networks," *Computer Networks*, vol. 182, p. 107477, 2020.
- [13] "Tactile internet over wifi-7 simulator [online]," <https://github.com/VinGok/Tactile-WiFi7>, accessed: 2021-08-06.
- [14] D. Magrin *et al.*, "Validation of the ns-3 802.11 ax ofdma implementation," in *Proceedings of the Workshop on ns-3*, 2021.
- [15] V. Gokhale *et al.*, "Tixt: An extensible testbed for tactile internet communication," *IEEE Internet of Things Magazine*, vol. 3, no. 1, pp. 32–37, 2020.