NR-U and Wi-Fi Coexistence in sub-7 GHz Bands: Implementation and Evaluation of NR-U Type 1 Channel Access in ns-3

George V. Frangulea
Philippos Assimakopoulos
University of Kent
Canterbury, United Kindom
{gvf3,p.asimakopoulos}@kent.ac.uk

Biljana Bojović
Sandra Lagén
Centre Tecnològic de Telecomunicacions de Catalunya (CTTC/CERCA)
Barcelona, Spain
{biljana.bojovic,sandra.lagen}@cttc.es

ABSTRACT

This paper addresses the challenge of optimising coexistence between 5G New Radio Unlicensed (NR-U) and Wi-Fi networks in the unlicensed spectrum at the sub-7 GHz bands. Leveraging the latest 3GPP standard TS 37.213, we align the listen-before-talk procedure with the latest standardisation, including implementation improvements. Through simulations, we demonstrate the advantages and limitations of the Type 1 channel access procedure. This work brings valuable insights, proposes solutions, and sets the groundwork for an NR-U extension crucial for future research. In particular, we evaluate the interplay between the NR-U numerologies and the channel access procedure set with different priorities.

CCS CONCEPTS

• **Networks** → **Network simulations**; **Mobile networks**.

KEYWORDS

ns-3, NR-U, LBT, Coexistence, Unlicensed.

ACM Reference Format:
George V. Frangulea, Philippos Assimakopoulos, Biljana Bojović, and Sandra Lagén. 2024. NR-U and Wi-Fi Coexistence in sub-7 GHz Bands: Implementation and Evaluation of NR-U Type 1 Channel Access in ns-3. In *2024 Workshop on ns-3 (WNS3 2024), June 05–06, 2024, Barcelona, Spain.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3659111.3659114

1 INTRODUCTION

With the ever-increasing demand for network connectivity, the evolution to 5G must exceed the performance benchmarks of previous cellular generations. Achieving high capacity, data rates, and low latency poses a multifaceted challenge in meeting the dynamic requirements of modern communication systems. As part of its initiative to augment 5G cellular operations into unlicensed spectrum, the 3rd Generation Partnership Project (3GPP) initially focuses on utilising the Unlicensed National Information Infrastructure (UNII) bands, specifically those at 5 GHz and 6 GHz as well as millimeter-Wave (mmWave) bands at 60 GHz, through the so-called New Radio Unlicensed (NR-U) technology. Wireless systems have the flexibility to function across unlicensed bands, given their adherence to spectrum regulations. These regulations aim to guarantee the harmonious coexistence of various wireless technologies operating at the same unlicensed band. Considering the widespread deployment of Wi-Fi networks, facilitating harmonious coexistence between 5G NR-U and Wi-Fi is a fundamental requirement.

The 3GPP introduced various solutions, including Long Term Evolution Unlicensed (LTE-U) [12] and LTE Licensed Assisted Access (LTE-LAA) [1]. LTE-U employed a cycling-based scheme called Carrier Sensing Adaptive Transmission (CSAT) [11] to ensure equitable coexistence with Wi-Fi networks, albeit limited to regions where Listen Before Talk (LBT) is not mandatory. Seeking enhanced adaptability, 3GPP incorporated an initial load-based (LB) LBT for LTE-LAA detailed in TR 36.889 Release (Rel) 13 [1] (referred to as Cat4 (Rel13) in this work), akin to IEEE 802.11's Carrier Sense Multiple Access (CSMA) based Distributed Coordination Function (DCF) [9]. The initial LTE-LAA LBT procedure was based on Option B in clause 4.8.3.2 European Telecommunications Standards Institute (ETSI) EN 301 893 V1.1.1 [7]. However, this procedure became non-standardised with the release of ETSI 301.893 V2.1.1 in 2017 [8]. Consequently, 3GPP redefined the LTE-LAA LBT in TS 36.213 [2], aligning it with the latest ETSI 301.893 standard. The 3GPP provides the standard for NR-U in TR 38.889 [3], where the LTE-LAA LBT procedure is mentioned as the baseline procedure for the 5 GHz bands and a starting point for the 6 GHz band. In 2018, 3GPP initiated TS 37.213, covering both NR-U and LAA LBT procedures (note that in TS 37.213, the term LBT is not used, only the term Channel Access Procedure (CAP) is used) for shared spectrum channel access.

Despite the fact that the LTE-LAA LBT procedures (e.g., Cat4 (Rel13)), as defined in TR 36.889 [1], are not compliant with the current ETSI EN 301 893 standards, a significant portion of current research continues to rely on outdated LBT procedures. These procedures can result in unfair coexistence towards Wi-Fi networks, leading to unfair channel access and often hindering effective data transmission for Wi-Fi networks. This issue will be discussed further in Section 4. This underscores the need for new studies that implement the standard compliant CAPs. The ns-3 5G-LENA simulator [13], integrating the open-source NR-U module, facilitates comprehensive research by providing a complete NR-U operational framework based on LTE-LAA LBT (Rel13). This module, developed in collaboration with InterDigital, draws from the author's experience with Wi-Fi Alliance Spidercloud Wireless during the unlicensed LTE development phase. This work resulted in an upgraded

WNS3 2024, June 05–06, 2024, Barcelona, Spain

ACM ISBN 979-8-4007-1763-5/24/06
https://doi.org/10.1145/3659111.3659114

NR-U module that builds upon the LTE-LAA Rel(13) LBT framework [1]. This work brings significant contributions by upgrading the LTE-LAA LBT (Rel13) procedure to conform with the latest 3GPP standard [5], accompanied by comprehensive simulation results elucidating the benefits and constraints of the introduced enhancements. Furthermore, the paper identifies potential challenges and offers solutions, thereby providing valuable insights for developing an NR-U module that is conducive to future research endeavours.

The remainder of the paper is structured as follows. Section 2 presents the LBT procedures as defined in TS 37.213 [5]. Section 3 describes the original and the updated LBT-based CAP implementation in ns-3 as well as some other aspects important for this work (e.g., NR frame structure and drawback of Type 1 CAP). In Section 4, we discuss the performance of various LBT procedures that have been described and implemented. Finally, in Section 5, we highlight the lessons learnt and propose future work ideas.

2 3GPP TS 37.213 CHANNEL ACCESS PROCEDURES FOR SHARED SPECTRUM

In TS 37.213 [5], 3GPP provides the specifications for the physical layer procedures for shared spectrum channel access used by LAA or NR-U systems. These procedures fall into two categories: asynchronous and synchronous channel access. These are known as load-based and semi-static, respectively, where semi-static is also known as frame-based (FB) LBT. The asynchronous LBT must be used in cases where long-term assurance of the absence of Wi-Fi or any other asynchronous LBT-based nodes cannot be guaranteed. When the absence of synchronous LBT-based nodes is guaranteed, FB-LBT can be implemented, achieving a frequency reuse factor of 1 and lowering channel access complexity due to the absence of the need for random back-off. However, the FB-LBT system requires synchronisation of NR-U base stations Next Generation Node B (gNBs) and has the drawback of fixed overhead within a frame during idle time.

This section examines asynchronous LBT procedures, which, as per TS 37.213, are referred to as Downlink (DL) or Uplink (UL) CAPs [5].

2.1 Type 2 Channel Access Procedures

Type 2 CAPs are used when the duration spanned by the sensing slots over which the channel must be sensed as idle preceding UL or DL transmissions is deterministic. Type 2 CAPs are further classified into three subcategories: Type 2A, 2B, and 2C.

Type 2A CAP is employed by a base station when the acquired channel is intended for sending discovery burst signals, discovery bursts signals multiplexed with non-unicast data or following transmission(s) of a UE after a gap of $25\mu s$ in a shared Channel Occupancy Time (COT) obtained through the Type 1 CAP. Type 2A is used by a UE as indicated by the serving base station. This procedure is deemed successful if the channel remains idle during a sensing interval of $T_{short_{dl}} = 25\mu s$, where $T_{short_{dl}}$ comprises a $16\mu s$ duration T_f with one sensing slot at the beginning, immediately followed by another sensing slot T_{sl}. The interval $T_{short_{dl}}$ is considered idle if both sensing slots are sensed as idle. The same applies for $T_{short_{ul}}$ at the UE side for UL CAP.

Table 1: Channel Access Priority Class (p) for UL

p (p)	m_p	$CW_{\min,p}$	$CW_{\max,p}$	$T_{\mathbf{mcot},p}$ [ms]	Allowed CW_p sizes
1	2	3	7	2	{3,7}
2	2	7	15	4	{7,15}
3	3	15	1023	6 or 10	{15,31,63,127,511,1023}
4	7	15	1023	6 or 10	{15,31,63,127,511,1023}

Table 2: Channel Access Priority Class (p) for DL

p (p)	m_p	$CW_{\min,p}$	$CW_{\max,p}$	$T_{\mathbf{mcot},`p}$ [ms]	Allowed CW_p sizes
1	1	3	7	2	{3,7}
2	1	7	15	3	{7,15}
3	3	15	63	8 or 10	{15,31,63}
4	7	15	1023	8 or 10	{15,31,63,127,511,1023}

Type 2B and Type 2C CAP applies before transmissions that are performed after a gap of $16\mu s$ or up to $16\mu s$, respectively within a shared COT. Type 2B CAP is considered successful if the channel is sensed as idle within a duration of $T_f = 16\mu s$, where T_f includes a sensing slot within the last $9\mu s$ of T_f. Type 2C requires no sensing and can only be used for transmissions under $584\mu s$.

2.2 Type 1 Channel Access Procedure

Figure 1 presents the flow diagram illustrating Type 1 CAP based on TS 37.213 Rel18 [5]. The red highlight shows the option defined for Cat4 (Rel13) in TR 36.889 [1], which was then removed in accordance with ETSI EN 301 893. To study the impact of this change, the possibility to activate or deactivate this option was left (more details in Section 3). The green highlight covers the mechanisms implemented during this work, which were required to achieve the Type 1 CAP.

Type 1 CAP procedure is employed by the initiating node (for NR-U networks, it is either gNB or UE) when the duration of sensing slots preceding UL or DL transmissions is random. This procedure is defined for DL and UL and is differentiated by their Channel Access Priority Class (CAPC) parameters. For the remainder of this paper, CAPC will be referred to as p. Type 1 UL CAP is designed to be more stringent in allowing access to the channel. This is done during the defer duration T_d and the random backoff procedure. The defer duration is defined as $T_d = T_f + m_p \times T_{sl}$, where T_f is equal to $16\mu s$ and encompasses a sensing slot T_{sl} at the start. The m_p parameter depends on the priority p used and the direction of transmission (see Table 1 and Table 2, for UL and DL CAPC, respectively). This procedure is considered successful if the channel is sensed as idle during all sensing slots (i.e., energy sensed on the channel during T_{sl} is below the energy detection threshold (X_{Thresh}) for at least $4\mu s$ during the duration of each sensing slot). Table 1 shows that for the UL priority *p1* and *p2*, the m_p values are larger than DL *p1* and *p2* defined in Table 2. This makes the User Equipment (UE) access to the shared channel more restrictive. The same is valid for priority *p3* UL, where the Contention Window size (CWS) is the same as priority *p4*.

The Type 1 CAP is deemed successful if the channel is sensed as idle during the T_d duration, and the counter N reaches zero, as

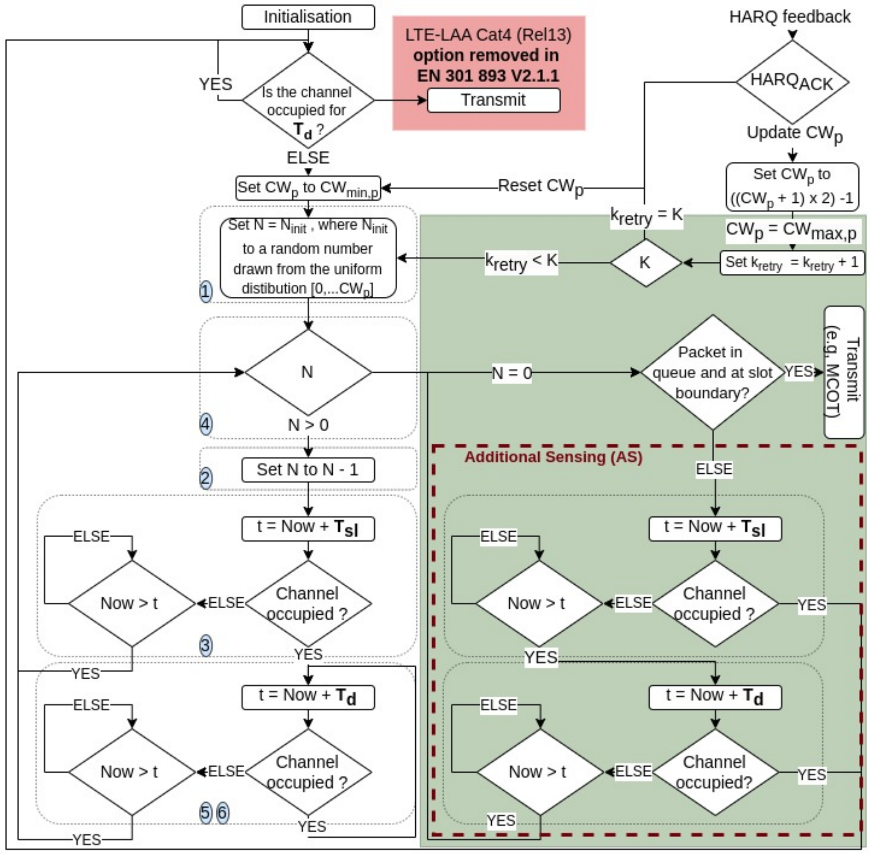

Figure 1: Type 1 CAP Flowchart based on the TS 37.213 Section 4.1.1

illustrated in step 4 of Figure 1. When a transmission is first initiated for a specific priority p, the CW_p is set to its minimum value, $CW_{\min,p}$. In Figure 1 step 1, the counter N is randomly chosen from the $[0, ..CW_p]$ range, which is followed by step 4. If the counter is randomly chosen to be zero and the initiating node has packets in the queue and is at the slot boundaries, it can access the channel for a specified Maximum Channel Occupancy Time (MCOT). If the counter is greater than zero, step 2 is executed, where the counter N is decremented if chosen to do so. Step 3 involves sensing the channel for a duration of T_{sl}. If the channel is sensed as idle during T_{sl}, step 4 is executed. If the channel is sensed as busy, the procedure enters a loop consisting of steps 5 and 6, repeatedly sensing the channel for an additional duration of T_d until the channel is sensed as idle and step 4 can be executed. Assuming that steps 1 to 6 have been completed, but the initiating node did not initiate a transmission at the beginning of the COT due to an empty transmission buffer or due to misalignment with the slot boundary, Additional Sensing (AS) is required. In the AS phase, the base station must ensure that the channel is sensed as idle during all sensing slots within the combined T_{sl} and T_d duration. If packets are present in the transmission buffer and there is an alignment with the slot boundary, the initiating node can proceed with transmission. Otherwise, the AS procedure is repeated. If the channel is sensed as busy during any of the sensing slots within the T_{sl} or T_d duration of the AS, the entire procedure must restart, starting with step 1, as soon as a T_d duration is sensed as idle. The functionality to enable or disable the requirement of an AS just before initiating a transmission was added during our work to study its impact. This was a highly important aspect when upgrading from Cat4 (Rel13) to Type 1 CAP due to the frame-based transmission nature of NR. Section 3 will elaborate on the advantages and drawbacks of incorporating this functionality, while Section 4 will delve into performance metrics illustrating the enhancements this brings to the coexistence scenario.

Upon successful initiation of a COT by the initiating node and availability of Hybrid Automatic Repeat Request (HARQ) feedback, the CW_p will be adjusted as illustrated in Figure 1. The CW_p is reset to its minimum value ($CW_{\min,p}$) under two conditions:

(1) For cases where transport block-based feedback is employed, the CW_p is reset to $CW_{\min,p}$ if at least one HARQ-ACK feedback is an acknowledgement (ACK).
(2) For cases where code block group-based feedback is employed, the CW_p is reset to $CW_{\min,p}$ if at least 10 % of the HARQ feedback must be ACKs.

In all other cases, the CW_p will be increased to the next higher value. When CW_p reaches the maximum value, the k parameter set by the initiating node will be used to decide if CW_p is maintained

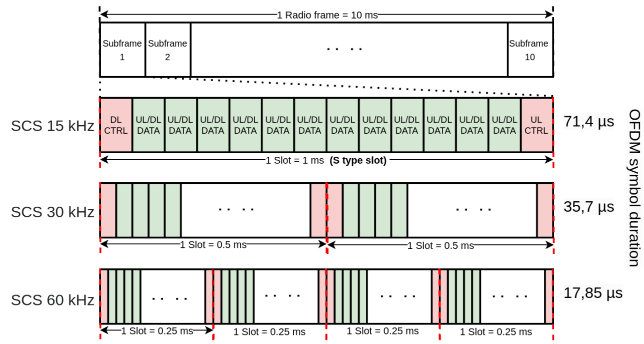

Figure 2: 5G NR Radio Frame Structure and Slot Configuration

or will reset to $CW_{min,p}$. The parameter k can take values within the $1, 2, ..., 8$ range and determines how many consecutive times $CW_p = CW_{\max,p}$ can be used for a given priority p.

3 NS-3 IMPLEMENTATION

Our NR-U module is built upon the 5G-LENA's NR [13]. The NR-U module incorporates crucial NR-U features such as various Channel Access Manager (CAM) implementations. Core functionalities of the NR module also support NR-U features like energy detection (ED) and unlicensed mode PHY state machine. This section highlights the upgrade from the Cat4 (Rel13) [1] to the latest Type 1 CAP (Rel18), showcasing advancements achieved through our work. Furthermore, the adaptation of Cat 2 and CAM to accommodate the implementation of Type 2 CAPs is discussed.

3.1 5G Radio Frame Structure

This work primarily focuses on the sub-7 GHz bands, utilising numerologies 0, 1, and 2 to achieve Subcarrier Spacings (SCS) of 15 kHz, 30 kHz, and 60 kHz, respectively. The NR module provides a standard-compliant implementation of the NR Time Division Duplexing (TDD) pattern scheduling. The slot type directly influences the slot structure at the symbol level. For this work, all slots utilised a Special-type (S-type) structure, as illustrated in Figure 2. This structure includes a DL symbol at the start of the slot for Downlink Control Information (DCI) transmission and a UL symbol at the end of the slot for Uplink Control Information (UCI). The remaining 12 Orthogonal Frequency-Division Multiplexing (OFDM) symbols within each slot are flexible (F) symbols, usable for either UL or DL data transmissions.

Another crucial consideration when selecting the SCS is the flexibility each numerology offers. The SCS is inversely proportional to the OFDM symbol duration, resulting in shorter transmission times for higher SCS. Higher SCS also lead to smaller Resource Block (RB) sizes. This relationship is significant, particularly in scenarios with high interference, as larger RBs correlate with lower error probabilities. Increasing the SCS may elevate the error probability due to smaller Transport Block Size (TBS). Figure 2 illustrates that each increase in the SCS doubles the number of slots within a subframe duration and the number of OFDM symbols per slot.

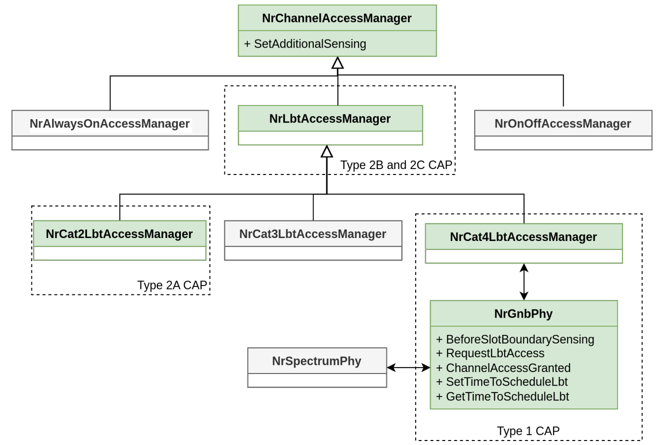

Figure 3: NR-U Type 1 CAM Class Diagram

3.2 Implementation of NR-U Channel Access Procedures

Various CAMs utilise the `NrChAccessManager` to control channel access. Based on the CAM algorithm outcome, the physical (PHY) layer is informed of channel availability. The existing NR-U model and its extensions are illustrated in Figure 3.

The `NrChAccessManager` base class specialises in three derived classes:

- `NrAlwaysOnAccessManager`: This CAM class ensures constant access to the channel, bypassing the LBT procedure and emulating NR-like behaviour.
- `NrOnOffAccessManager`: Operating in a duty-cycled manner, this CAM class alternates between allowing and preventing transmissions on the channel.
- `NrLbtAccessManager`: Implements the 3GPP LBT-based procedures (now referred as CAPs).

Among the available CAM types, only the `NrLbtAccessManager` and its specialisations implement the CAPs for shared spectrum and use the channel state information required to execute the CAP algorithm, i.e., uses the information whether the channel is idle or busy to perform defer, backoff, additional sensing, etc. This information is obtained from the attached instance of `NrSpectrumPhy`, which implements the NR-U state machine, i.e., when the device is not transmitting, it switches to the sensing mode. The `NrLbtAccessManager` specialisations were upgraded and re-used in the following way:

- `NrCat2LbtAccessManager`: CAP without random back-off in which the sensing period is deterministic was re-used to implement Type 2A CAP. In contrast, Type 2B and 2C CAPs were integrated in `NrLbtAccessManager`.
- `NrCat3LbtAccessManager`: CAP with random random-back-off with a fixed-size CWS was removed as per TS 37.213 [5].
- `NrCat4LbtAccessManager`: Cat4 (Rel13) [1] with random backoff implementation described in [6] was upgraded to achieve Type 1 CAPs as per TS 37.213 (Rel18) [5].

3.2.1 NR-U CAPs based on 3GPP TR 36.889. Figure 4 (a) illustrates a DL transmission based on LTE-LAA (Rel13) [1], where the channel

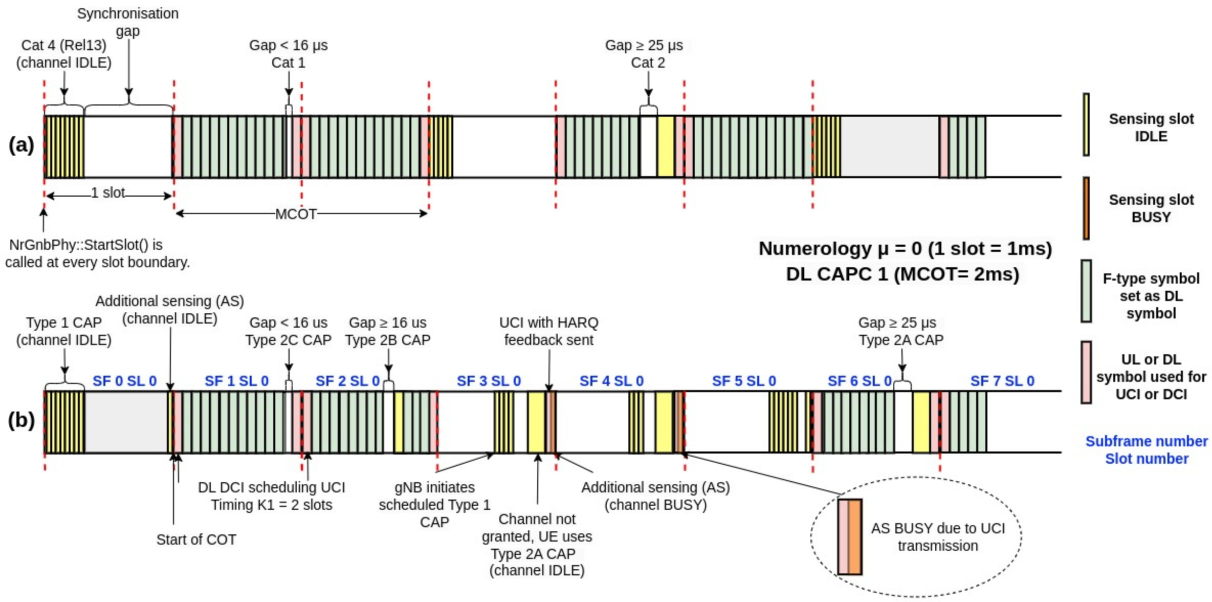

Figure 4: Old and New NR-U CAPs Implementations: (a) the Non-standardised LTE-LAA LBTs (3GPP LTE-LAA TR 36.889) and (b) the NR-U CAPs as per TS 37.213

is acquired through Cat4 (rel13) and shared through Cat 1 and Cat 2.

At the base station PHY layer, at every slot boundary, during `StartSlot`, the system checks for scheduled data or control signalling. If scheduling exists and the channel is not granted, the PHY layer calls `RequestAccess` using `NrCat4LbtAccessManager` CAM. During the Cat4 (Rel13) procedure, if the channel is sensed as idle, the CAM indicates this via the `ChannelAccessGranted` function, granting access for MCOT duration depending on the p used. Subsequently, transmission can occur at the next slot boundary without requiring further sensing. The time from the end of Cat 4 (Rel13) to the slot boundary is shown in Figure 4 (a) as the synchronisation gap. Transmission proceeds if data is scheduled within the MCOT granted. Otherwise, the gNB queues back the allocation without requesting a new slot from the MAC. The rest of the transmission depicts scenarios with gaps within the MCOT and the procedures used in such cases. For gaps smaller than 16 μs, Cat 1 LBT is used (no sensing) to continue COT sharing. The most significant gap allowed between DL and UL transmissions is $25\mu s$, requiring Cat 2 procedure to be performed.

3.2.2 NR-U CAPs based on TR 37.213. In addition to implementing the CAPs, this work introduces a scheduling technique for the Type 1 CAP, referred to as Scheduled Type 1 CAP. This technique minimises the synchronisation gaps depicted in Figure 4, which can vary from 0.9 to 0.2 milliseconds (ms) depending on the SCS used. Substantial gaps in the time domain undermine the effectiveness of employing a load-based LBT procedure, as the extended duration renders the LBT outcome invalid.

Figure 4 (b) illustrates the NR-U transmission using the CAPs defined in TS 37.213 [5], using the Scheduled Type 1 CAP for COT initiation with priority $p1$ and SCS of 15 kHz. For subframe (SF) 0, slot (SL) 0, the `StartSlot` function is called to check for scheduled data and request the channel if necessary. If the channel was sensed as idle during the backoff duration of the Type 1 CAP, the CAM calls the `ChannelAccessGranted` function. However, unlike Cat4 (Rel13), the channel is not set as granted directly. Instead, the synchronisation gap is calculated and compared with a preset threshold. Suppose the gap is lower than the threshold. In that case, the channel is set as granted, indicating the MCOT based on the priority p set for the Type 1 CAP. Otherwise, the threshold is used to schedule `BeforeSlotBoundarySensing`, and the `SetTimeToScheduleLbt` function is called to store the synchronisation gap for further scheduling of the Type 1 CAP. Within `BeforeSlotBoundarySensing`, `SetAdditionalSensing` is called to indicate that the PHY layer enters the Type 1 CAP to perform AS. If the AS is sensed idle, the `ChannelAccessGranted` class is called again, which will align with the slot boundary, allowing the channel to be granted and the transmission to proceed. For future subframes and slots when `StartSlot` is invoked while the channel is not granted, and there is scheduled data for the next slot, `RequestLbtAccess` is scheduled using `GetTimeToScheduleLbt`. The function `GetTimeToScheduleLbt` returns a random time duration within 90% and 100% of the gap stored with `SetTimeToScheduleLbt`. This randomisation of the scheduling time aims to mitigate a known issue in NR-U, which is synchronised channel access, which can lead to increased packet collisions. Additionally, scheduling the Type 1 CAP close to the slot boundary enhances the likelihood of immediate transmission. This approach is feasible because the NR module assumes a two-slot delay for MAC-to-PHY processing by default [11]. Thus, it is possible to anticipate if data is scheduled for the next slot and schedule the procedure accordingly.

Next, the intra-Radio Access Technology (intra-RAT) recurrent Type 1 CAP blocking issue will be detailed using Figure 4 (b). In SF1 SL0 and SF2 SL0, the DL symbol carries DCI containing the $K1$ parameter, indicating to the UE the allocated slot for UCI transmission. Since the $K1$ parameter can vary from zero to four slots, it can not be guaranteed that the channel will be granted on the respective slots. This uncertainty arises when the MCOT duration is reached, requiring the release of the channel. In this example, each $K1$ parameter indicates two slots, meaning that UCI must be sent in SF3 SL0 and SF4 SL0. Since MCOT duration was reached at the end of SF2 SL0, the expected UE behaviour is to utilise Type 2A CAP to acquire the channel to send UCI. However, if both the gNB and UE sense the channel as idle in SF3 SL0 and SF4 SL0, the transmission of the UE UCI will inadvertently block the serving gNB during the AS. This issue is further compounded in systems where gNB serves multiple UEs. This challenge is particularly pronounced when UEs lack fast processing capabilities to transmit UCI within the remaining COT.

3.3 Buffer Occupancy

During our work, we created a new trace to measure the Buffer Occupancy (BO). 3GPP defined BO as the sum of the period of time in which the UE has data to transmit, including retransmissions. This trace was implemented in the Radio Link Control (RLC) layer, `LteRlcUm`. At the RLC, the function `EvaluateBufferOccupancy` was created to calculate the BO. This function is scheduled at the end of the simulation. During run time, we use counters to measure the time durations in which the transmission queue of a UE was non-empty. When `EvaluateBufferOccupancy` is called, we divide the stored time in which the buffer was non-empty by the time passed since it was last scheduled. Thus, we acquire the BO percentage per gNB.

4 NR-U AND WI-FI EVALUATION CAMPAIGN

This section presents a simulation campaign that evaluates the NR-U using SCS of 15, 30 and 60 kHz, defined for sub-7 GHz in coexistence with Wi-Fi 6 (also known as 802.11ax).

4.1 Scenario

For this work, the NR-U and Wi-Fi networks are set for DL transmission from a serving gNB/AP to two receiving UEs/STAs. The network is non-saturated and is modelled using the parameters from Table 3. The saturation level within the NR-U network is impacted by the priority p set for the Type 1 CAP. For example, using priority $p4$ can result in a double BO level compared to $p3$ due to their maximum backoff durations (priority $p4$ can reach a backoff duration of up to 9.206 ms while priority $p3$ can reach up to 0.567 ms). The traffic used is a 3GPP-defined generic video type of data [10] of 20 Mbps data rate and a frame generation rate of 60 frames per second.

The layout scenario used to study the CAP procedures and their impact on the coexistence between the NR-U and Wi-Fi networks is defined in [4] as Indoor-B. This layout is modelled as an office box measuring 40 meters by 40 meters, with one base station per operator at a three meter height. Base stations are randomly deployed within a virtual box measuring 10 meters by 10 meters, ensuring a minimum distance of two meters between them. The UEs are randomly deployed within the office box at a height of one meter.

Table 3: Simulation Parameters

Parameter	Value (baseline)	Value (tested variations)
Deployment and propagation parameters:		
Channel model	3GPP Indoor Hotspot Mixed Office	
Deployment	2 operators, 1 BSs and 2 users per BS (Indoor-B)	
gNB/AP/UE/STA height	3m/3m/1m/1m	
Channel bandwidth	20 MHz	
Central frequency	5.22 GHz	
Traffic parameters:		
Direction	Downlink	
Traffic type	3gppGenericVideo (R:20 Mbps, Fps: 60)	
Device parameters:		
gNb/AP; UE/STA antennas	2 x 2; 1 x 2	
gNb/AP; UE/STA transmit power	24 dBm; 18 dBm	
gNB/AP Noise figure	5 dB	
UE/STA Noise figure	9 dB	
NR-U parameters and functionalities:		
Frame structure	$\mu = 0$	$\mu = 1, \mu = 2$
MAC	Scheduled-based	
gNB CAP	Type 1 CAP	
UE CAM	Type 2A CAP	
RLC mode	RLC-UM	
gNB / UE ED threshold	-72 dBm	
CAPC	CAPC4	CAPC3
WiFi 11ax parameters and functionalities:		
Frame aggregation	A-MPDU with size: 9000 bytes	
DCF parameters:	same CWS range as CAPC4 or CAPC3, TxOp ns-3 default	
CCA-CS	-82 dBm	
CCA-ED	-62 dBm	

4.2 Impact of NR-U Numerology

The remaining sections present Cumulative Distribution Functions (CDFs) derived from data collected over 79 independent simulation runs. During one simulation, the application runs for 10 seconds, resulting in 600 packets being sent, with an average of 41700 bytes per packet. Three figures are displayed for each output metric, illustrating results for each SCS used in NR-U.

The CAPs evaluated (more details in Section 3, 2) in this section are as follows:

- Cat4 (Rel13) is the LTE-LAA LBT from the 3GPP TR 36.889 [1]. This CAP does not implement mandatory backoff or AS.
- Type 1 CAP (no AS, not scheduled). This variation of Type 1 CAP does not implement AS or procedure scheduling. This CAP does enforce mandatory backoff.
- Scheduled Type 1 CAP (no AS). This procedure is scheduled before the slot boundary and includes mandatory random backoff. This CAP does not implement AS.
- Scheduled Type 1 CAP follows the latest standard CAP for shared spectrum TS 37.213 [5]. This procedure implements our technique of scheduling the CAP before the slot boundary.

4.2.1 Transport Block and MAC Protocol Data Unit Loss. Figure 5 illustrates the CDF of lost (corrupted) TBs across each independent simulation run. For the aforementioned packet size of around 41700 bytes and SCS values of 15 kHz, 30 kHz, and 60 kHz, with ideal SINR conditions and modulation and coding scheme (MCS) of 27, the TBS are 12976, 6488, and 3182 bytes, respectively.

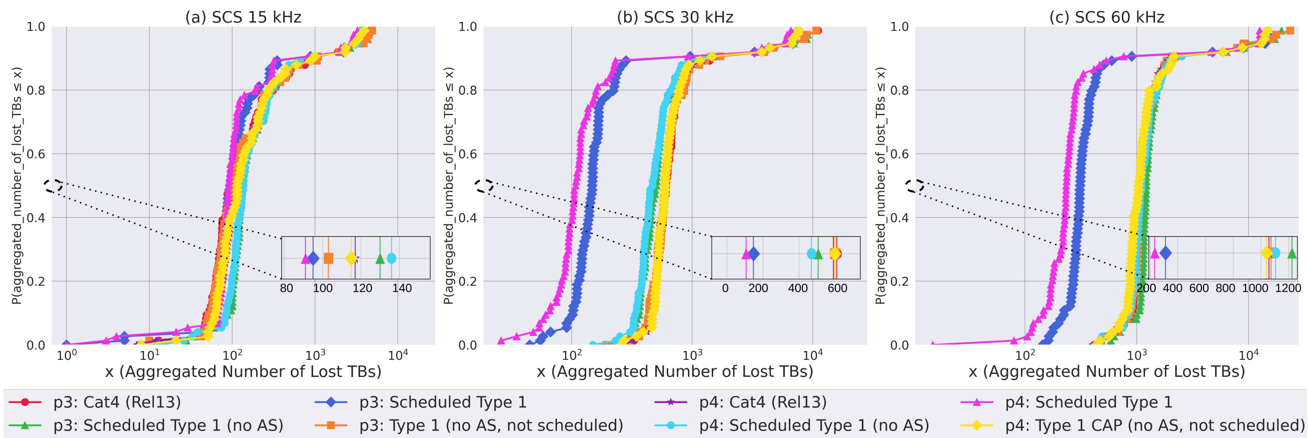

Figure 5: Number of TBs Lost at NR-U for Various CAPs Set with Priority *p3* or *p4* Evaluated at: (a) SCS 15 kHz, (b) SCS 30 kHz and (c) SCS 60 kHz in Coexistence with Wi-Fi 11ax

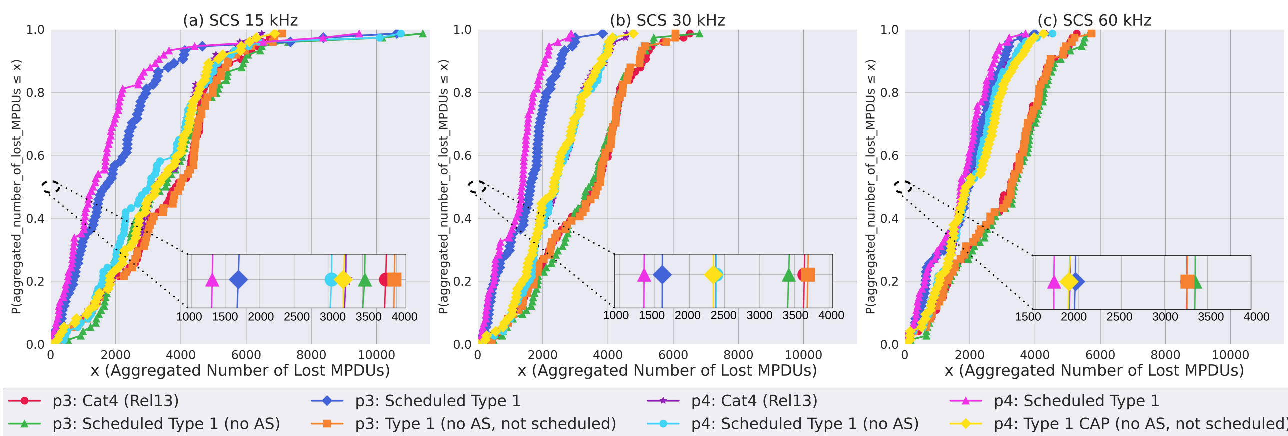

Figure 6: Number of MPDUs Lost at Wi-Fi 11ax when Coexisting with NR-U Utilising Various CAPs Set with Priority *p3* or *p4* at: (a) SCS 15 kHz, (b) SCS30 kHz and (c) SCS 60 kHz

Consequently, transmitting a single packet requires approximately 3.2, 6.4, and 13.1 TBs (the number of TBs used to send one packet is subject to variation based on channel conditions). In scenarios where the CAP fails to avert collisions over a specified duration, the number of lost TBs during that duration scales according to the SCS used.

The Scheduled Type 1 CAP emerges as the top performer in minimising TB loss across all SCS values of 15, 30, and 60 kHz, as illustrated in Figures 5 (a), (b), and (c) respectively. Among the tested priorities *p3* and *p4*, the latter proves superior interference prevention, owing to its larger maximum CWS, evident even across the CAPs that lack certain sensing mechanisms. Notably, the Scheduled Type 1 CAP set with priority *p3* achieves close performance to the CAPs set with priority *p4*, primarily due to incorporating the AS mechanism. Analysing the 50^{th} percentile of the CDF, we compare the most restrictive (Scheduled Type 1 CAP) with the least restrictive (Cat4 (Rel13)) CAP for SCS 15, 30, and 60 kHz. The percentage change between the former and the latter CAP is used for this comparison. For priority *p3* or *p4* at SCS 15, 30 and 60 kHz, Scheduled Type 1 CAP outperforms Cat4 (Rel13) by 0% or 28%, 311% or 446% and 239% or 354% respectively. Similar results were observed at the above 98^{th} CDF region, which shows how the network will perform in saturation scenarios. This region shows a 25% performance difference between Scheduled Type 1 CAP and Cat4 (Rel13) set with priority *p3* at SCS 15 kHz, whereas in the median, the CAPs achieved the same performance.

Figures 6 (a), (b), and (c) showcase the number of lost MAC Protocol Data Unit (MPDUs) observed at the Wi-Fi network when NR-U uses SCS of 15, 30, and 60 kHz, respectively. In line with the preceding discussion and findings, it is evident that the Wi-Fi network experiences fewer MPDUs lost when NR-U utilises more restrictive CAPs, such as those set with priority *p4* and the Scheduled Type 1 CAP. Looking at the 50^{th} CDF percentile in Figures 6 (a), (b), and (c), the following enhancement can be observed when using Scheduled

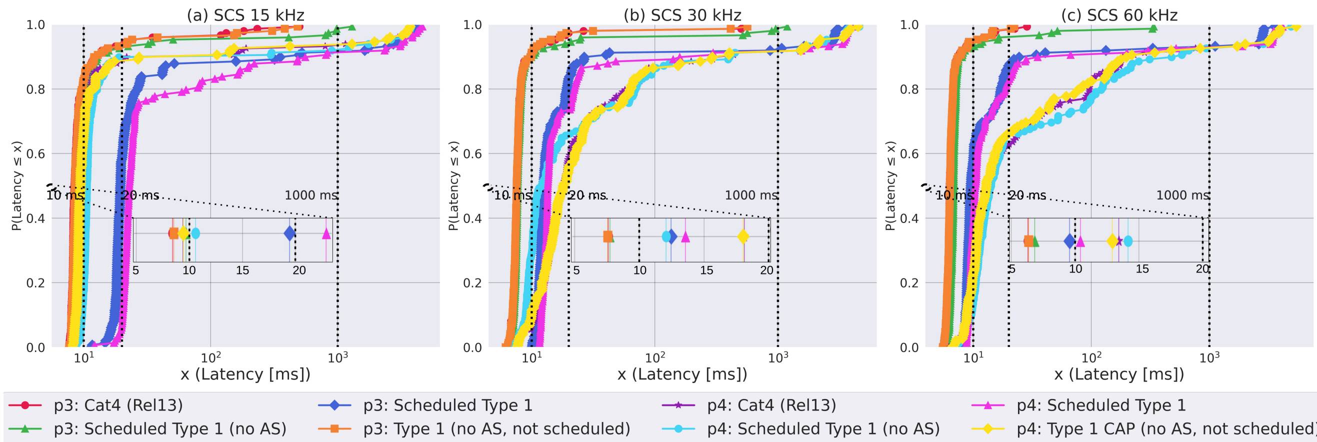

Figure 7: Latency Performance for NR-U for Various CAPs Set with Priority *p3* or *p4* at: (a) SCS 15 kHz, (b) SCS 30 kHz and (c) SCS 60 kHz in Coexistence with Wi-Fi 11ax

Type 1 CAP compared to Cat4 (Rel13). The former CAP outperforms the latter CAP (when set with priority *p3* or *p4*) by 128% or 150%, 124% or 76% and 66% or 11% in the number of lost MPDUs. An overall trend of decreased number of lost MPDUs can also be seen when increasing SCS. This indicates that shorter transmission durations of NR-U fosters more secure Wi-Fi transmissions.

4.2.2 Latency Performance. Figures 7 (a), (b), and (c) present the CDF of the NR-U latency when utilising various CAPs set with priority *p3* or *p4* at SCS of 15, 30, and 60 kHz, respectively. Each figure incorporates straight dotted lines facilitating visualisation when latency exceeds 10, 20, or 1000 ms. In line with the findings from Section 4.2.1, Figure 7 shows that the least restrictive CAPs set with priority *p3* (that provide faster access to the channel by differing less in the presence of interference) achieve the lowest latencies. The Scheduled Type 1 CAP achieves the worst latency performance due to the restrictive CAP design and the recurrent Type 1 CAP blocking discussed in Section 3.2.2.

Analysing the 50th CDF percentile, we observe latency changes upon increasing the SCS. For instance, transitioning from SCS 15 to 30 kHz, the Scheduled Type 1 CAP set with priority *p3* or *p4* experiences a latency reduction from 19.4 to 12.4 ms (36% decrease) or from 22.9 to 13.6 ms (40.7% decrease) respectively. The Cat4 (Rel13) set with priority *p3* or *p4* has latency changes from 8.43 to 7.6 ms (9.85% decrease) or from 9.4 to 18 ms (92.57% increase) respectively. The substantial latency improvements observed for the Scheduled Type 1 CAP set with priority *p3* or *p4* are attributed to the AS mechanism. Conversely, Cat4 (Rel13) set with priority *p4* experiences a significant degradation of 92.57% due to the absence of AS, exacerbating the delay resulting from the larger maximum CWS of priority *p4*. This proves that a larger maximum CWS is only efficient if it is ensured that the outcome of the CAP is still valid at the slot boundary (e.g., using AS). Similar latency trends are observed when increasing SCS from 30 to 60 kHz. The Scheduled Type 1 CAP, set with either priority *p3* or *p4*, undergoes latency reductions from 12.4 to 9.6 ms (23.3% decrease) or from 13.6 to 10.4 ms (23.3% decrease) respectively. Meanwhile, Cat4 (Rel13), set with priority *p3* or *p4*, experiences latency reductions from 7.6 to 6.3 ms (17.11% decrease) or from 18.1 to 13.4 ms (26% decrease) respectively. These results underscore that the Scheduled Type 1 CAP, set with either priority *p3* or *p4*, undergoes constant latency enhancements when increasing SCS, attributed to shorter transmission durations. Furthermore, even in the absence of AS, the Scheduled Type 1 CAP (no AS) demonstrates superior latency reduction compared to non-scheduled CAPs. This can be attributed to the shorter slot duration associated with higher SCS, thereby minimising the likelihood of invalidating the CAP outcome and increasing the probability of completing the CAP at the slot boundary.

Contrasting outcomes were observed within the Wi-Fi network, where the most restrictive CAPs and priorities used by NR-U allowed the Wi-Fi network to achieve lower latencies. Enforcing more sensing mechanisms within the NR-U CAPs invariably prevents concurrent transmission between NR-U and the Wi-Fi network, thus degrading NR-U latency while significantly enhancing Wi-Fi performance. For the 50th percentile of the CDF, the Wi-Fi network exhibited considerable latency improvements when NR-U employed the Scheduled Type 1 CAP instead of Cat4 (Rel13). When for the NR-U network, the former CAP or the latter CAP set with priority *p3*, Wi-Fi latency was of 10.15 or 48.6 ms (379% change), 9.6 or 102 ms (962% change) and 10.2 or 95.2 ms (824% change) (observed at NR-U SCS of 15, 30, and 60 kHz respectively). The improvements for priority *p4* were slightly smaller since priority *p4* already offers better protection to the Wi-Fi network through the larger maximum CWS.

4.2.3 Throughput Performance. Figures 8 (a), (b) and (c) depict the CDF illustrating the throughput achieved by the Wi-Fi network under various CAP configurations at NR-U across SCS values of 15, 30, and 60 kHz. Notably, the Scheduled Type 1 CAP demonstrates optimal throughput performance, with no significant degradation observed for priority *p3* or *p4*. The CAPs with fewer sensing mechanisms, such as Type 1 (no AS, not scheduled), Scheduled Type 1 (no AS), and Cat4 (Rel13) set with priority *p3*, exhibit increasing throughput degradation as NR-U SCS increases. This degradation

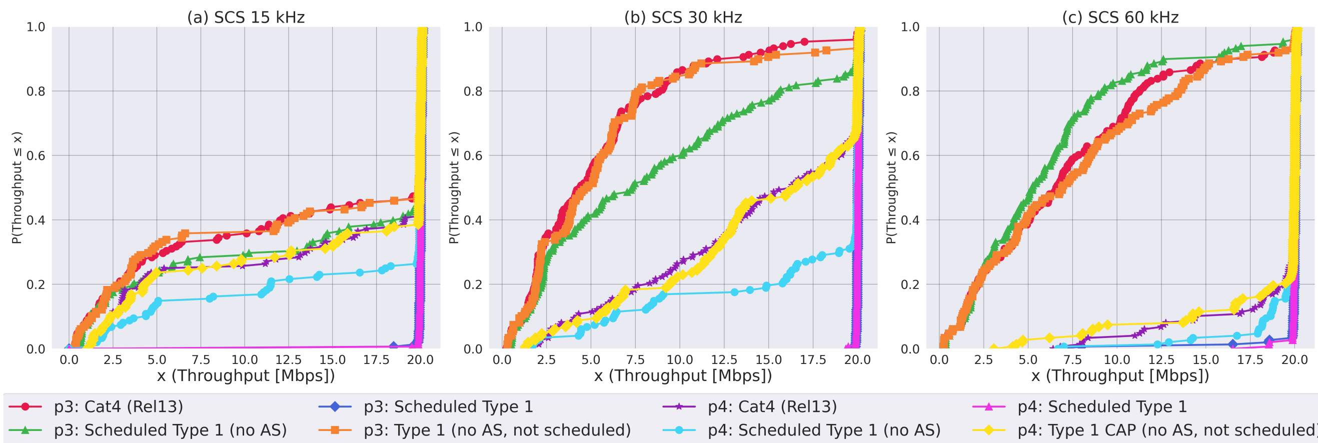

Figure 8: Throughput Performance of Wi-Fi 11ax when Coexisting with NR-U at Various CAPs Set with Priority *p3* or *p4* at: (a) SCS 15 kHz, (b) SCS 30 kHz and (c) SCS 60 kHz

is attributed to the heightened number of lost TBs associated with higher SCS, compounded by the smaller maximum CWS of priority *p3*. Conversely, the same procedures set with priority *p4* demonstrate significantly smaller degradation in Wi-Fi throughput, underscoring the effectiveness of a larger maximum CWS in mitigating simultaneous NR-U and Wi-Fi transmissions, even in procedures lacking certain sensing mechanisms.

Regarding NR-U throughput performance, the Scheduled Type 1 CAPs exhibit slight degradation at SCS 15 kHz at the region below the 50^{th} CDF percentile, which diminishes at higher numerologies. Specifically, at SCS 15 kHz, Cat4 (Rel13) outperforms the Scheduled Type 1 CAP at the 50^{th} CDF percentile, whereas at SCS 30 kHz and 60 kHz, the Scheduled Type 1 CAP surpasses Cat4 (Rel13). These results are consistent with the constant gains observed at higher numerologies when implementing the Scheduled Type 1 CAP.

5 CONCLUSION AND FUTURE WORK

The evaluation campaign conducted for NR-U and Wi-Fi coexistence provided valuable insights into the importance of implementing the 3GPP procedures described in TS 37.213 [5] for NR-U. In Section 4.2.1, significant improvements in the number of lost TBs or MPDUs were observed for both NR-U and Wi-Fi. Section 4.2.2 highlighted major latency improvements for the Wi-Fi network with the standardised procedure, albeit at the cost of latency degradation in the NR-U network. This degradation is due to the restrictive design of Type 1 CAP, to the recurrent intra-RAT Type 1 CAP blocking and synchronisation gaps discussed in Section 3.2.2. These issues underscore significant challenges in NR-U design that require addressing, which will be the focus of our future work. Additionally, the Wi-Fi throughput performance in Section 4.2.3 showed that Wi-Fi can achieve ideal throughput performance with only marginal throughput degradation in the NR-U network.

In light of the conclusions drawn in this work, the following objectives for future work can be outlined:

- Further research is needed for minimising NR-U latency in TDD-based transmission.
- Proper signalling must be implemented to avoid intra-RAT LBT blocking.
- Type 1 CAP optimisation for the slot-based transmission, including minimisation of the synchronisation gap effect, is needed.

ACKNOWLEDGMENTS

CTTC authors have received funding from Grant PID2021-126431OB-I00 funded by MCIN/AEI/10.13039/501100011033 and “ERDF A way of making Europe”, and Generalitat Catalunya 2021 SGR 00770.

REFERENCES

[1] 3GPP. 2015. Study on Licensed-Assisted Access to Unlicensed Spectrum. TR 36.889 (Rel. 13).
[2] 3GPP. 2016. Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Layer Procedures. 3GPP, TS 36.213, V13.1.1.
[3] 3GPP. 2018. Study on NR-based Access to Unlicensed Spectrum. TR 38.889 (Rel. 16), V16.0.0.
[4] 3GPP. 2020. Considerations on Performance Evaluation for NR in 52.6-71GHz. 3GPP TSG RAN WG1 Meeting 103-e,R1-2007943.
[5] 3GPP. 2024. Physical Layer Procedures for Shared Spectrum Channel Access. TS 37.213 (Rel. 18), V18.2.0.
[6] Biljana Bojović, Lorenza Giupponi, Zoraze Ali, and Marco Miozzo. 2019. Evaluating Unlicensed LTE Technologies: LAA vs LTE-U. *IEEE access* 7 (2019), 89714–89751.
[7] ETSI. 2015. Harmonized European Standard, “Broadband Radio Access Networks (BRAN); 5 GHz High Performance RLAN. EN 301 893, V1.1.1.
[8] ETSI. 2017. Harmonised Standard Covering the Essential Requirements of Article 3.2 of Directive 2014/53/EU. EN 301 893, V2.1.1.
[9] IEEE Std. 2020. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. IEEE Standard, 802.11-2012, Part 11.
[10] Sandra Lagen, Biljana Bojovic, Katerina Koutlia, Xiaodi Zhang, Ping Wang, and Qi Qu. 2023. QoS Management for XR Traffic in 5G NR: A Multi-Layer System View & End-to-End Evaluation. *IEEE Communications Magazine* (2023).
[11] Sandra Lagén, Lorenza Giupponi, Sanjay Goyal, Natale Patriciello, Biljana Bojović, Alpaslan Demir, and Mihaela Beluri. 2020. New Radio Beam-based Access to Unlicensed Spectrum: Design Challenges and Solutions. *IEEE Communications Surveys Tutorials* 22, 1 (Oct. 2020), 8–37.
[12] LTE-U Forum. 2015. LTE-U CSAT Procedure. TS V1.0.
[13] Natale Patriciello, Sandra Lagen, Biljana Bojovic, and Lorenza Giupponi. 2019. An E2E Simulator for 5G NR Networks. *Simulation Modelling Practice and Theory* 96 (Nov. 2019), 101933.