Can 5G NR Sidelink communications support wireless augmented reality?

Ashutosh Srivastava*[†], Qing Zhao*, Yi Lu*, Ping Wang*, Qi Qu*, Zhu Ji*, Yee Sin Chan*, Shivendra S. Panwar[†]

*Meta Platforms, Inc., USA, {lulouis, pingwang, qqu, zhuji}@meta.com
[†]New York University, New York, USA, {as12738, sp1832}@nyu.edu

Abstract—Smart glasses that support augmented reality (AR) have the potential to become the consumer's primary medium of connecting to the future internet. For the best quality of user experience, AR glasses must have a small form factor and long battery life, while satisfying the data rate and latency requirements of AR applications. To extend the AR glasses' battery life, the computation and processing involved in AR may be offloaded to a companion device, such as a smartphone, through a wireless connection. Sidelink (SL), i.e., the D2D communication interface of 5G NR, is a potential candidate for this wireless link. In this paper, we use system-level simulations to analyze the feasibility of NR SL for supporting AR. Our simulator incorporates the PHY layer structure and MAC layer resource scheduling of 3GPP SL, standard 3GPP channel models, and MCS configurations. Our results suggest that the current SL standard specifications are insufficient for high-end AR use cases with heavy interaction but can support simpler previews and file transfers. We further propose two enhancements to SL resource allocation, which have the potential to offer significant performance improvements for AR applications.

Index Terms—Augmented reality, Sidelink, 5G NR, D2D, 3GPP, Resource allocation

I. INTRODUCTION

Augmented reality (AR) involves overlaying virtual information on top of a user's live 3D view of the real world. High-quality AR experiences can be delivered through AR glasses. AR glasses will enable applications that traditional handheld devices, like smartphones, cannot realize. These include many flavors of use cases such as 3D holographic calling, real-time sharing and overlay of virtual objects, real-time frictionless internet browsing, and messaging [1]. AR applications require advanced processing and computation to estimate the user's location under 6DoF motion, identify objects in their field-of-view (FoV), and render the virtual content on top of their real-world view. The AR glasses may also need to maintain a reliable, low latency, and high-speed internet connection to support some use cases. However, for the best user experience, the wearable AR glasses must have a small form factor and long battery life, which limits their computation capacity. In such a scenario, we may need a companion device tethered to the glasses to perform the bulk of the processing. The companion device, such as a smartphone with cellular connectivity, serves as the bridge between the AR glasses and the internet.

A wired connection between the AR glasses and the companion device provides high-speed connectivity but restricts the user's mobility and degrades their comfort level. For a seamless experience, the AR glasses must have a wireless
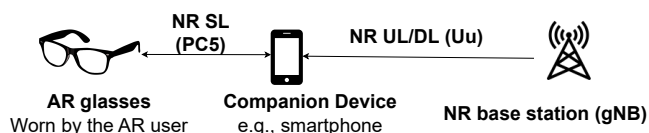

Fig. 1: AR user's glasses connected to the cellular network via a wireless SL (D2D) connection to a companion device

connection. This device-to-device (D2D) wireless connection between the glasses and the companion will be a critical component of the AR ecosystem. In many scenarios, we need low-latency D2D transfers to ensure a sufficient latency budget for the other components of the AR processing pipeline. The D2D link must also operate within the thermal constraints of the AR glasses, requiring sufficient duty cycling. Possible candidate technologies for this D2D wireless link are: Bluetooth low energy (BLE), Wi-Fi, and cellular D2D, i.e., sidelink (SL) [2].

In this work, we study the feasibility of using SL to support the D2D connection in AR. The model for this is shown in Fig. 1. The 3rd Generation Partnership Project (3GPP) first standardized D2D communication for proximity services in its Release 12 and has since updated it in later releases, including the introduction of SL as the D2D interface of 5G NR in 3GPP Rel. 16 (2020). Although, as of Rel. 17, SL is primarily designed for V2X use cases, it will likely be enhanced for AR/VR in the future (Rel. 18 and beyond). This paper conducts an in-depth system-level evaluation of the feasibility of 5G NR SL for supporting AR use cases, motivated by the increasing interest of significant industry players in SL technology [3]. We believe we are the first to do such an analysis, which differs from previous simulation-based studies on LTE C-V2X and 5G NR V2X [4]–[8], that primarily focused on V2X deployments and traffic conditions [9]. While some recent works [10], [11] discuss the potential of SL for XR, they do not present a complete performance evaluation.

The **key contributions** of this work are as follows:
We develop an abstracted use case framework that involves a connection between the AR glasses and the companion device. We provide a semi-quantitative analysis of the supportability of a spectrum of AR use cases with SL. For this purpose, we design a system-level simulation setup in MATLAB, which models the PHY layer structure and MAC layer protocols of NR SL. We evaluate SL performance in deployments representative of typical AR use case conditions with reasonable assumptions. 3GPP standard channel models, modulation and

coding schemes (MCS), and antenna configurations are used along with the link layer abstraction model from [12] to ensure that our setup aligns with realistic wireless channel conditions. We identify key limitations that hinder the feasibility of a SL-based D2D connection for AR glasses. Finally, we propose two promising solutions for SL resource allocation: multiple active reservations and full-duplex sensing, which can significantly enhance SL performance for AR use cases.

The rest of the paper is organized as follows: In section II, we provide a basic introduction to 5G NR SL, focusing on its PHY structure and MAC layer resource scheduling. The methodology of our analysis is described in Sec. III. We then present the results of this study in Sec. IV. Finally, Sec. V concludes our study and presents future research directions.

II. 5G NR Sidelink: Introduction

We start with a basic introduction to 5G NR SL. This will serve as the background for the subsequent sections.

A. SL transmissions at the PHY layer

NR SL supports unicast, multicast, and broadcast transmissions. The communication between the AR glasses and the companion device is unicast. We assume that SL operates in the sub-7 GHz frequency bands (FR1).[1] SL uses the orthogonal frequency division multiplexing (OFDM) waveform with cyclic prefix (CP). In the frequency domain, 12 contiguous OFDM subcarriers form a physical resource block (PRB). The minimum resource allocation unit in frequency is known as a *subchannel* which can comprise 10 to 100 PRBs. In the time domain, a single SL radio frame is of duration 10 ms and is further divided into 10 subframes (1 ms each). Each subframe has multiple time *slots*, the minimum unit of resource allocation in time. SL supports NR's scalable numerology, with slot durations depending on the OFDM sub-carrier spacing (SCS). For FR1 SL, SCS values of $\{15, 30, 60\}$ kHz are allowed, corresponding to slot durations of $\{1, 0.5, 0.25\}$ ms, respectively. One time slot contains 14 OFDM symbols. One SL transmission instance, known as a transport block (TB), can comprise multiple contiguous subchannels in the same slot.

A SL TB comprises control information followed by data. Sidelink control information (SCI) is sent in two parts: first stage and second stage SCI (SCI_1 and SCI_2). SCI_1 is sent over the physical SL control channel (PSCCH) and can be decoded by any other SL UE if received without error. SCI_2 and the data are sent over the physical SL shared channel (PSSCH) and can only be decoded by the intended receiving UE. We illustrate this in Fig. 2 where TX1-RX1 is an ongoing SL transmission. RX1 can receive and decode the entire TB, while other nearby SL UEs (e.g., TX2) can only decode SCI_1. SCI_1 includes information about the resources reserved by TX1, and its resource reservation interval (RRI), i.e., the transmission periodicity. A reserved resource refers to a set of contiguous subchannels in the frequency domain. If TX2 receives an SCI
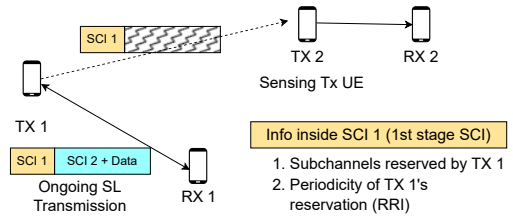
Fig. 2: SCI and data transmissions in NR SL

from TX1 in slot x, indicating that TX1 uses resource S with an RRI of y, then it knows that TX1 will also occupy the same resource in slots $x + y$, $x + 2y$ and so on.

B. SL resource allocation procedures at the MAC layer

Mode 1: This is the centralized resource allocation mode of SL where the NR gNB (base station) controls the scheduling and allocates resources to the SL UEs for their D2D transmissions. Mode 1 requires the SL Tx UE to be in coverage of an NR gNB. To support Mode 1, the AR glasses may rely on the companion if they cannot directly connect to the gNB (possibly due to hardware/battery constraints).

Mode 2: In Mode 2, each SL Tx UE performs a distributed sensing-based procedure to select resources for its transmissions. We consider the semi-persistent scheme (SPS) of Mode 2, whereby the SL Tx UE uses a newly reserved resource for its next RRC[2] consecutive transmissions in periodic time slots.

When selecting a new resource, the SL Tx UE defines a time window (**selection window**) that starts soon after the latest packet arrival time and lasts up to the packet delay budget (PDB). Step 1 of Mode 2 is the exclusion of resources already reserved by other SL Tx UEs based on the sensing information (SCIs) collected in a recent history of time known as **sensing window**. A SCI_1 from another UE is considered for resource exclusion only if its reference signal received power (RSRP) at the sensing UE is above a certain pre-configured threshold.

Step 2 involves randomly selecting a suitable resource among the available candidate resources after Step 1. Once a resource is reserved, it can be used by the SL Tx UE over periodic time slots for its next RRC transmissions. Once this reservation expires, i.e., the resource counter goes to 0, the Tx UE can perform a new resource selection with probability p_{change} or can keep the same resources for another RRC transmissions with probability $1 - p_{change}$.[3]

Our study includes an analysis and comparison of Mode 1 vs. Mode 2, which is elaborated in Sec. III-G.

III. Methodology: System Model, Assumptions and Simulation Setup

In this section, we describe our assumptions, system model, and simulation setup for evaluating the feasibility of using SL for AR. Our goal is to provide directional insights that can serve as a framework for enhancing the technology, despite the possibility that some assumptions may differ in real scenarios.

---

[1]Optimized SL design for mmWave frequencies (FR2 bands) is yet to be addressed by 3GPP, with only some ongoing research [13], [14]. Hence, this work does not cover SL over FR2, and its analysis is left for future work.

[2]RRC stands for the resource re-selection counter.
[3]For more details on Mode 2, please refer to [15].

A. What spectrum should be used for the AR SL?

There are three potential spectrum options for the glasses-to-companion SL communication: dedicated licensed spectrum (FR1) for XR SL, licensed spectrum (FR1) shared with other cellular traffic, or shared unlicensed spectrum over the 5-7 GHz bands. Each option has its pros and cons, and it is unclear which one the industry will adopt. Here, we make two simplifying assumptions: (1) the AR SL operates in a shared sub-7 GHz spectrum that may be either licensed or unlicensed, with no interference between the glasses-companion SL and the companion-gNB cellular link, as they operate on different bands, and (2) the companion device has two radios, one for the 5G NR Uu link to the gNB and one for the NR SL.

B. Modeling external interference

AR SL operating on either licensed or unlicensed spectrum will experience interference from other cellular or WiFi/NR-U users. To consider such interference, we assume that external traffic occupies $X\%$ of the available airtime, i.e., for each time slot, there is an independent probability of $X/100$ that it will be occupied by an external user, making it unavailable for AR users to use for SL transmission. If an AR user had already scheduled a transmission during such an occupied slot, they must postpone it to the next reserved slot. The underlying assumption here is that AR users can sense the channel for external interference to detect slots that are already occupied. SL devices already use channel sensing to detect ongoing SL transmissions in the network and decode their SCIs. To detect any external interference occupying the channel, a simple energy-based detection threshold similar to the one used in Wi-Fi systems can be employed.

C. Deployment Scenario

We design a deployment scenario that can represent several indoor AR use cases, such as an office space, a single floor of a multi-family apartment complex, an AR gaming arcade, or an exhibition/museum space with visitors wearing AR glasses. We model AR users as a pair of SL UEs (companion device and AR glasses) deployed in a $20\,\mathrm{m} \times 20\,\mathrm{m}$ 2D square grid with random locations and orientations. We vary the number of users inside the grid from 1 to 20 to cover scenarios with light (1 to 5), moderate (5 to 10), or heavy (5 to 20) user density. The glasses-to-companion distance is uniformly distributed between $1\,\mathrm{m}$ to $2\,\mathrm{m}$ for short-range D2D communication.

D. System Model

We use the 3GPP indoor office channel model for our study. The path loss is modeled by the following equation:

$$PL = 32.4 + 17.3 \log_{10} d + 20 \log_{10} f_c,$$

where, d is the 2D Euclidean distance between the transmitter and the receiver, and f_c is the carrier frequency (6 GHz). We assume a constant transmit power ($P_{tx} = 14$ dBm) and that omnidirectional beams are used for transmission/reception.

The received signal power is modeled as:

$$P_{rx} = G_{rx} * RL_{rx} * PL * G_{tx} * RL_{tx} * P_{tx}$$

where, G_{tx} and G_{rx} are the transmit and receive antenna gains. RL_{tx} and RL_{rx} are the transmit and receive RF losses. PL is the pathloss as explained above. We neglect shadowing and body blockage effects, with a worse performance expected, if included. We aim to identify key SL design bottlenecks rather than perform rigorous quantitative analysis, which requires more precise channel modeling and is left for future study. We consider a network with N SL Tx-Rx pairs, i.e., the AR users. Based on each user's MAC layer scheduling, only a subset of the N users perform SL transmission in a given time slot and frequency subchannel (active users). Let the number of such "active" SL pairs be N_{active}. Let $P_{rx}^{i,j}$ be the received signal power at RX_i from a transmission initiated at Tx_j, where P_{rx} is as defined above. Let P_{noise} be the noise power, then the signal-to-interference plus noise ratio (SINR) of the transmission of the $i - th$ Tx-Rx pair is calculated as:

$$SINR = \frac{P_{rx}^{i,i}}{\sum_{j \neq i, i=1}^{i=N_{active}} P_{rx}^{i,j} + P_{\mathrm{noise}}}$$

E. Traffic model

3GPP recommends a quasi-periodic traffic generation model for XR apps, where a burst of traffic is generated at the beginning of each period, or "traffic interval", with a rate of 1/T Hz [11], [16]. Our simulations use a default rate of 45 Hz, which can range from 30-120 Hz. The burst includes both glasses-to-companion and companion-to-glasses traffic, with the amount depending on the use case, e.g., a 20 Mbps requirement results in 55kB of traffic every 22 ms. Traffic generated at the start of an interval must be transferred within the same period, imposing a latency limit on every SL transmission. A packet is dropped if it cannot be scheduled within this limit. Promptly completing the SL transfer allows the AR glasses to enter sleep mode and save power, as discussed in [17].

F. Bi-directional traffic support

In the unicast operation of SL Mode 2, the Tx UE performs resource selection for its transmissions to its corresponding Rx UE. However, in AR, the SL traffic is bidirectional, where both the AR glasses and the companion can act as the Tx or Rx UE. We consider two options: (1) the companion device performs sensing and resource allocation and informs the AR glasses through SCI, and (2) independent sensing and resource allocation by both devices. We select Option 1, where the AR glasses do not have to perform sensing and resource selection, saving energy. Coordination between the companion device and the glasses for resource allocation should be feasible in practice. Option 2 will result in performance degradation due to collisions as both devices compete for channel access.

G. SL resource allocation: Mode 1 vs. Mode 2

In Mode 1 (centralized), the gNB controls the scheduling and has more information about the network than the SL UEs

performing sensing-based resource selection in the case of Mode 2. A comparative analysis of both modes is desirable. Before describing our simulation approach, we first list the following key features of Mode 1:

- The gNB uses its knowledge of future resource allocations for all SL users to prevent collisions among them.
- In Mode 2, the SL Tx UE cannot sense the channel in a time slot in which it is also transmitting data unless the device can support full duplex operation. Mode 1 **does not** have this constraint, as the gNB performs the resource allocation. The gNB has all the information on the resources allocated to SL users.
- The gNB can utilize the location information of all SL UEs to efficiently reuse resources among spatially separated SL Tx-Rx pairs.

In our simulation, we model Mode 1 by augmenting our Mode 2 implementation with a "genie" assisted approach. The genie provides all SL Mode 2 UEs with additional information about the network, which is used to emulate the aforementioned features unique to Mode 1. This allows us to approximate Mode 1 performance. More details of this approach are provided in [18], where it was used to get an upper bound for SL Mode 2 performance. We compare Mode 1 and Mode 2 performance results in Section IV.

H. MCS configuration and Link adaptation

We use three representational MCS values (Table I) from Table 5.1.3.1-1 of 3GPP standard 38.214 [19], with MCS 19 for high data rates and MCS 4 for reliable data transfer in poor link conditions. The companion estimates interference using the SCIs it received during sensing. It calculates the expected SINR for its SL transmission over the selected resources and maps the SINR to a block error rate (BLER). MCS is then adapted to meet a desired BLER target (1% by default[4]).

For the SINR to BLER mapping, we use the physical layer (PHY) abstraction model from [12], which captures the 5G NR specifications related to low-density parity check (LDPC) channel coding, such as LDPC base graph selection and code block segmentation. It uses link-level simulations and the exponential effective SINR mapping (EESM) method to obtain the SINR vs. BLER data. For more details, you may refer to the original work in [12].

MCS Index	Modulation Order	Code Rate ($R \times 1024$)	Spectral efficiency	Data rate at 100 MHz BW (Mbps)
4	2	308	0.6016	60.16
11	4	378	1.4766	147.66
19	6	517	3.0293	302.93

TABLE I: MCS configurations

---

[4]The BLER target is device vendor implementation dependent. A higher value may be configured if it can be tolerated by the target AR use cases.

I. SL Simulation parameters

Our simulation parameters comply with 3GPP specifications and are shown in Table II[5]. For details on SL Mode 2 parameters, please refer to [15]. The reservation size equals the number of frequency sub-channels, optimizing the AR D2D transfer latency by using the entire bandwidth in a single slot.

Simulation Parameter	Value
Available bandwidth (MHz)	20/40/60/80/100
Sub-carrier spacing (SCS)	30 kHz
Slot Time	0.5 ms
Subchannel size	15 PRBs
Resource Reservation Interval (RRI)	2 ms
Resource re-selection counter	$\sim U[25, 75]$
Probability of changing resources in Mode 2	0.5
RSRP threshold in Mode 2	-96 dBm
Mode 2 sensing window	100 ms

TABLE II: SL simulation parameters

J. Performance evaluation metrics

Packet reception ratio (PRR): Each SL packet (TB) can have three possible outcomes: (1) *success*: Successful reception at the Rx [6], (2) $Rx_{failure}$: Packet reception error at the Rx side, and (3) *drop*: Packet is dropped at the Tx side due to exceeding its deadline, i.e., 22 ms latency limit. A user's PRR is the ratio of its *success* packets to its total number of transmissions.

Latency for a SL packet is $T_{reception} - T_{arrival}$, the time difference between its reception at the Rx and generation at the Tx. AR traffic is generated every 22 ms, and the total "on-time" of the glasses-companion link is determined by the latency of the last SL transmission of an AR user in each traffic interval. This is the key metric since it determines whether the SL connection can satisfy the AR glasses' thermal constraints and whether the bursty AR traffic can meet its total latency requirements. In our results, we refer to this metric simply as "Latency".

K. How do we judge AR use case satisfaction ?

Without limiting ourselves to specific product specifications, we use a semi-quantitative approach to evaluate the feasibility of SL for AR. We use two criteria - end-to-end latency impact and thermal limit of the AR glasses. AR involves multiple processing steps [7], and each step incurs latency, including the bidirectional glasses-to-companion SL transfer. Previous studies indicate that wireless transfer latency requirements for AR range from 5-50 ms, depending on the specific use case [20]. We classify use cases into three broad categories:

Class A- Interaction-based use cases with world-locked rendering (WLR): The user interacts with the AR content in

---

[5]Our slot time is 0.5 ms, and SL UEs follow the half-duplex sensing constraint [15]. Under these assumptions, an RRI of 1 ms would restrict the UE's choice of resources, making the use of sensing-based Mode 2 meaningless and possibly causing high packet errors. Thus, we select $RRI = 2$ ms, which is the next minimum permissible value, to minimize SL transfer latency.

[6]Our simulator maps SINR values to BLER [12] and determines the success of a TB's reception by flipping a biased coin with the BLER as the probability.

[7]This includes data capture using sensors, computation-heavy tasks like simultaneous localization and mapping (SLAM), hand tracking and rendering, the wireless D2D transfer, and the final display of the rendered object(s).

real time. WLR is supported, where content is placed at a fixed position in the environment, for comfortable user experience [21], [22]. Some examples of these use cases include holographic AR board games, online AR shopping, and AR-guided remote assistance for industry services [1]. WLR requires advanced processing, resulting in a stricter latency budget. We set the round-trip D2D wireless latency target for Class A to be around 15-20 ms.

Class B - Interaction-based use cases without WLR: The latency requirements are around $30 - 40$ ms. An example use case for this category is 3D AR video calling.

Class C - Simple preview, file transfer, or messaging: Less latency-sensitive use cases, with 100ms - 1sec. requirements. Examples include previewing 3D models of new products, 3D image/video sharing, and instant messaging using AR prompts or virtual keyboard [1], [23].

SL must operate within the thermal constraints of the AR glasses, which limits the duty cycle of the wireless radio. Specifically, if the thermal limit is $X\%$, all SL transfers must be completed within $(X/100) * T$ ms of a periodic traffic interval of T ms. The radio will then enter sleep mode for the remaining time to save power. We set the thermal limit to 50% as a rough anchor point, recognizing that its accuracy may depend on product specifications and design choices.

IV. RESULTS

For each result, we ran 20 simulations, each with a different random distribution of UEs for the deployment scenario described in Sec. III-C. Each simulation run lasted for 10 seconds. The results represent the system-wide average PRR and latency for all AR users and across the 20 distributions.

A. Baseline performance @ 100 MHz BW, no BG interference

Fig. 3 shows the performance of the AR glasses' SL connection in a multi-AR user scenario. The available bandwidth is 100 MHz, with no background interference. We simulate SL Mode 2 and our genie-assisted approximation of Mode 1 from Sec. III-G. With a traffic load of $5 + 5$ Mbps[8], SL can support up to 15 users with high reliability and latency below 10 ms. This meets all use case requirements and stays within the thermal limit [9]. However, as the traffic load increases, the PRR goes down, and latency increases, degrading performance and exceeding the thermal limit. With $15+15$ Mbps traffic, the requirements of class A use cases are not met, and only class B and C use cases can be supported for less than 10 users. With even more traffic ($30+30$), only class C use cases can be supported by allowing SL TBs to be queued for transmission in subsequent intervals. We also see that Mode 1 and Mode 2 have similar latency, while Mode 1 shows better reliability

[8]$5 + 5$ denotes glasses-to-companion (G2C) and companion-to-glasses traffic. Our results with symmetric traffic still provide a representative view of performance with asymmetric traffic and the same total load. As the G2C link is usually less loaded, our choice can represent the worst-case scenario (for the power-constrained AR glasses) under the same total load.

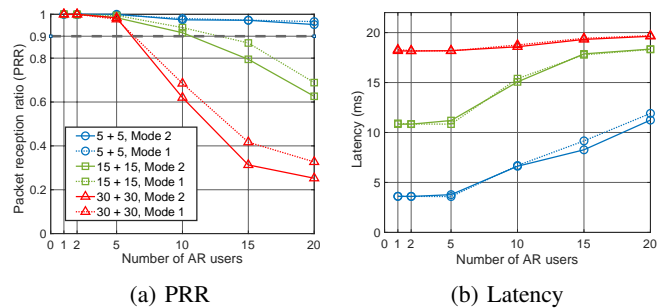[9]The thermal limit (maximum ON duration) is 50% of $T(22ms) = 11ms$

Fig. 3: Performance of the AR glasses' SL connection at varying traffic loads - 100 MHz BW, no external interference

under heavy traffic with many users [10]. Despite this, the trends for both schemes are similar, indicating that our results for Mode 2 can be a good indicator of Mode 1 performance. Hence, the results presented in subsequent sections will only consider Mode 2 unless specified otherwise.

B. Impact of available bandwidth and external interference

In **Fig. 4**, we look at how the amount of bandwidth (BW) available for SL impacts performance. A lower BW implies more congestion in the network, leading to lower PRRs and higher latencies. At a traffic load of $5 + 5$ Mbps, at least 60 MHz of BW is required to support all use cases for up to 10 users (Fig. 4a and Fig. 4b). Up to 5 users can be supported with 40 MHz BW, and 20 MHz is insufficient for any use case due to the thermal limit. For $15 + 15$ Mbps, we need a minimum BW of 80 MHz to support all use cases for up to 5 users. So far, no external interference has been considered. In **Fig. 5**, we fix BW to 100 MHz and vary the external interference. For $5 + 5$ Mbps, up to 10% interference can be tolerated without significant degradation in KPIs, with higher tolerance if there are ≤ 5 users. For $15+15$ Mbps, only low interference levels ($0-10\%$) are acceptable, while high interference (40%) results in excessive SL latency beyond the thermal limit.

C. Proposed enhancements to SL resource allocation

In the results so far, we observe high SL transfer latencies in many scenarios, particularly the ones with high traffic loads, limited bandwidth, and external interference. The main bottleneck here is the periodic nature of resource scheduling in the semi-persistent scheme (SPS) of SL Mode 2, which is not well suited to the bursty nature of AR traffic. With a scheduling periodicity (RRI) of 2 ms, the latency of the last TB sent in a traffic interval, i.e., the overall SL transfer latency for the AR user becomes high. This is illustrated in Fig. 6 (A). 3GPP has also defined a "dynamic scheme", wherein the SL Tx UE selects a new time-frequency resource for each TB. This is more suited for aperiodic traffic. However, if all SL devices in a network use the dynamic scheme, the channel access becomes equivalent to random resource selection, degrading reliability. As discussed in Note 5, an RRI lower than 2 ms is

[10]Our results show the user plane latency but do not include the control plane latency overhead needed for exchanging resource allocation information between the gNB and the SL UE in Mode 1. This is left for future work.
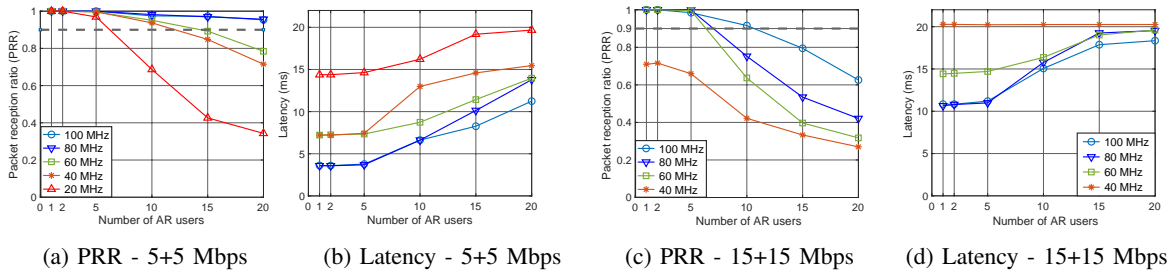
(a) PRR - 5+5 Mbps   (b) Latency - 5+5 Mbps   (c) PRR - 15+15 Mbps   (d) Latency - 15+15 Mbps

Fig. 4: Impact of available bandwidth on SL performance ($X = 0$ %)

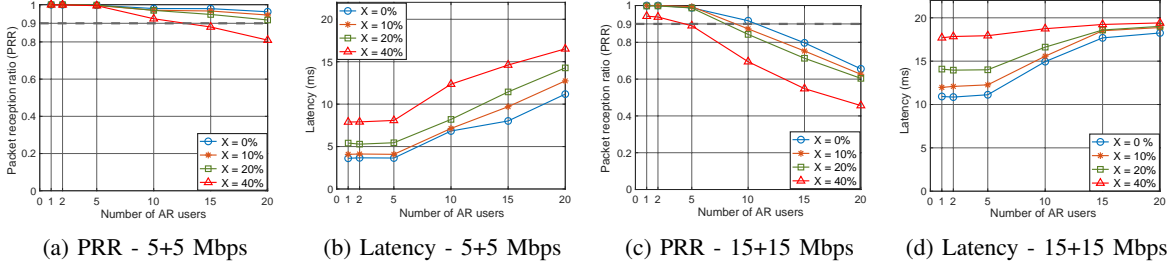(a) PRR - 5+5 Mbps   (b) Latency - 5+5 Mbps   (c) PRR - 15+15 Mbps   (d) Latency - 15+15 Mbps

Fig. 5: Impact of external interference on SL performance ($BW = 100$ MHz)

not feasible for our system. We propose two enhancements to SL Mode 2 that retain the benefits of sensing-based channel access while being more suitable for AR traffic:

(1) Use of **multiple active SL reservations (MAR)**: Instead of selecting a single periodic resource with $RRI = 2$ ms (periodicity of 4 time slots) as in Fig. 6 (A), we now allow a SL Mode 2 scheduling UE (the companion device) to select multiple resources in consecutive time slots. This allows an individual AR user's D2D transfer to complete quickly, reducing latency. In addition, the SL radio on the glasses will be active for a much shorter duration and can go to sleep after the transfer completes, thus alleviating the thermal impact issue. The MAR solution is depicted in Fig. 6 (B), where users can now use consecutive time slots to complete their SL transmissions. The RRI is set to 22 ms, corresponding to the traffic interval duration. Thus, a single UE can now select a "batch" of resources in consecutive time slots and use the same batch in longer periodic intervals to meet its traffic requirements. We simulate this scheme and observe its performance in **Fig. 7** for $15 + 15$ Mbps traffic. Our **baselines** for comparison are the conventional SL Mode 1 and Mode 2 schemes with periodic reservations. We see that SL Mode 2 with MAR has a significantly reduced latency; however, the PRR of Mode 2 + MAR is much lower than the baseline Mode 2 scheme [11]. With Mode 1, we observe that the MAR scheme can offer both low latency and high reliability.

(2) **Full-duplex (FD) sensing**: For Mode 2 SL devices, half-duplex operation prevents sensing during transmission. There is no such constraint in Mode 1 (centralized), which does not rely on sensing. This constraint can be eliminated from Mode

2 if the SL device performing sensing has FD capabilities. In Fig. 7, we see that a combination of MAR and FD greatly improves SL Mode 2 reliability and latency performance.
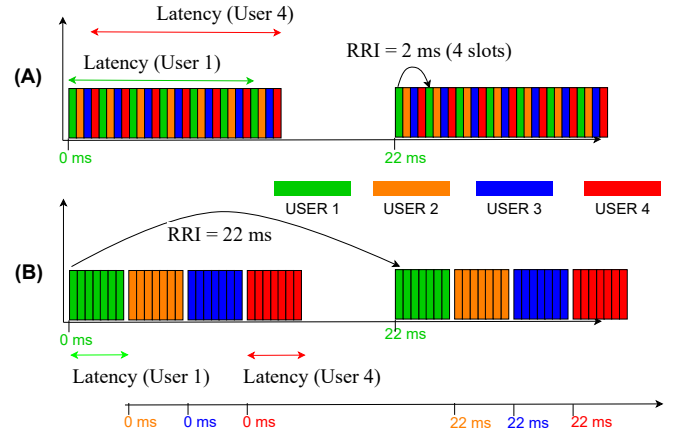

Fig. 6: Using multiple active reservations to reduce SL transfer latency for AR users, (A): Legacy SPS, (B): MAR scheme
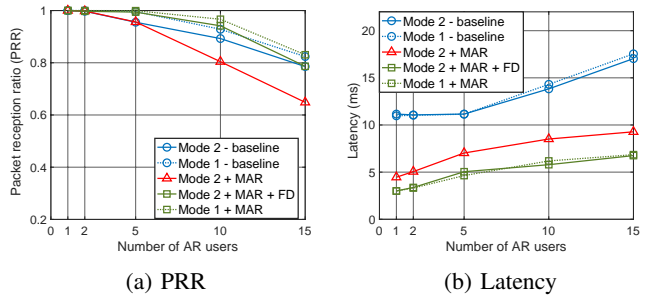

(a) PRR   (b) Latency

Fig. 7: Performance of the proposed enhancements - Multiple active reservations (**MAR**) and Full duplex sensing (**FD**)

[11] Persistent collisions can occur in the SPS scheme of Mode 2 [6]. This gets exacerbated with MAR, where users can hold on to a batch of reservations for a longer period. Mode 2 configurations, such as the RRC and p_{change}, can be optimized to better suit MAR. However, this is left for future work.

Implementing MAR should be feasible in practice; however, it remains to be seen whether its benefits can generalize to more use cases, thereby motivating efforts for standardization. The coexistence of MAR users with those using periodic SL reservations is an open question. FD sensing improved SL resource allocation for V2X in [24], and we demonstrated its benefits to SL for AR. However, designing practical FD radios with self-interference cancellation is an ongoing issue.

V. CONCLUSIONS

We analyze the feasibility of using sidelink (cellular D2D communication) for the wireless connection between a user's AR glasses and a companion device (such as their smartphone) used for computation offloading. We use system-level simulations to model the PHY and MAC layers of 5G NR SL and simulate realistic AR traffic and deployment scenarios. We find that the current SL design struggles to support advanced AR use cases, particularly those requiring real-time interaction, at medium to high traffic loads. The periodic scheduling of V2X SL does not align with AR glasses' latency and power consumption KPIs. To address these issues, we propose two enhancements to SL Mode 2: multiple active reservations (MAR) and full-duplex sensing. MAR enables SL to transmit continuously, reducing transfer latencies and power consumption. FD sensing ensures a highly reliable connection.

In future work, we intend to extend our simulations to model end-to-end AR use cases by incorporating the cellular link between the companion device and the base station. A more rigorous quantitative analysis will help guide industry-level product development. A theoretical framework to analyze SL performance for AR is a promising research direction.

REFERENCES

[1] 3GPP, "Extended Reality in 5G," 3GPP, TR 26.928, April 2022, v17.0.0.
[2] S.-Y. Lien *et al.*, "3GPP NR sidelink transmissions toward 5G V2X," *IEEE Access*, vol. 8, pp. 35 368–35 382, 2020.
[3] Q. Research, "How will sidelink bring a new level of 5G versatility?" September 2022. [Online]. Available: https://www.qualcomm.com/news/onq/2022/09/how-will-sidelink-bring-a-new-level-of-5g-versatility
[4] D. Wang, R. R. Sattiraju, A. Weinand, and H. D. Schotten, "System-level simulator of LTE sidelink C-V2X communication for 5G," in *Mobile Communication-Technologies and Applications; 24. ITG-Symposium.* VDE, 2019, pp. 1–5.
[5] A. Bazzi, A. O. Berthet, C. Campolo, B. M. Masini, A. Molinaro, and A. Zanella, "On the design of sidelink for cellular V2X: a literature review and outlook for future," *IEEE Access*, 2021.

[6] V. Todisco, S. Bartoletti, C. Campolo, A. Molinaro, A. O. Berthet, and A. Bazzi, "Performance Analysis of Sidelink 5G-V2X Mode 2 Through an Open-Source Simulator," *IEEE Access*, vol. 9, pp. 145 648–145 661, 2021.
[7] Z. Ali, S. Lagén, L. Giupponi, and R. Rouil, "3GPP NR V2X mode 2: overview, models and system-level evaluation," *IEEE Access*, 2021.
[8] M. M. Saad, M. T. R. Khan, S. H. A. Shah, and D. Kim, "Advancements in vehicular communication technologies: C-V2X and NR-V2X comparison," *IEEE Communications Magazine*, vol. 59, no. 8, pp. 107–113, 2021.
[9] 3GPP, " Study on evaluation methodology of new Vehicle-to-Everything (V2X) use cases for LTE and NR," 3GPP, TR 37.885, June 2019, v15.3.0.
[10] S. Saafi, G. Fodor, J. Hosek, and S. Andreev, "Cellular connectivity and wearable technology enablers for industrial mid-end applications," *IEEE Communications Magazine*, vol. 59, no. 7, pp. 61–67, 2021.
[11] V. Petrov, M. Gapeyenko, S. Paris, A. Marcano, and K. I. Pedersen, "Extended reality (XR) over 5G and 5G-advanced new radio: Standardization, applications, and trends," *arXiv preprint arXiv:2203.02242*, 2022.
[12] S. Lagen, K. Wanuga, H. Elkotby, S. Goyal, N. Patriciello, and L. Giupponi, "New Radio Physical Layer Abstraction for System-Level Simulations of 5G Networks," in *ICC 2020 - 2020 IEEE International Conference on Communications (ICC)*, 2020, pp. 1–7.
[13] S. Mura, F. Linsalata, M. Mizmizi, M. Magarini, M. N. Khormuji, P. Wang, A. Perotti, and U. Spagnolini, "Spatial-Interference Aware Cooperative Resource Allocation for 5G V2V Communications," in *2022 IEEE 95th Vehicular Technology Conference:(VTC2022-Spring)*. IEEE, 2022, pp. 1–6.
[14] A. Srivastava, S. Goyal, U. Salim, P. Liu, R. Pragada, and S. S. Panwar, "Overcoming Directional Deafness in High Frequency Sidelink Communications," in *2022 IEEE 95th Vehicular Technology Conference:(VTC2022-Spring)*. IEEE, 2022, pp. 1–7.
[15] M. H. C. Garcia *et al.*, "A Tutorial on 5G NR V2X Communications," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 3, pp. 1972–2026, 2021.
[16] 3GPP, "Study on XR (extended reality) evaluations for NR," 3GPP, TR 38.838, April 2022, v17.0.0.
[17] Y. Kim *et al.*, "UE Power Saving Techniques for Extended Reality (XR) Services in 5G NR Systems," in *2021 IEEE 32nd Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*. IEEE, 2021, pp. 1–7.
[18] A. Srivastava, S. Datta, S. Goyal, U. Salim, W. J. Hussain, P. Liu, S. Panwar, R. Pragada, and P. Adjakple, "Enhanced Distributed Resource Selection and Power Control for High Frequency NR V2X Sidelink," *IEEE Access*, 2023.
[19] 3GPP, "NR; Physical layer procedures for data," 3GPP, TS 38.214, September 2022, v17.3.0.
[20] F. Alriksson, D. H. Kang, C. Phillips, J. L. Pradas, and A. Zaidi, "XR and 5G: Extended reality at scale with time-critical communication," *Ericsson Technology Review*, vol. 2021, no. 8, pp. 2–13, 2021.
[21] Microsoft Corp., "Coordinate Systems," https://learn.microsoft.com/en-us/windows/mixed-reality/design/coordinate-systems.
[22] Y. Zhao, J. Stefanucci, S. Creem-Regehr, and B. Bodenheimer, "Evaluating Augmented Reality Landmark Cues and Frame of Reference Displays with Virtual Reality," *IEEE Transactions on Visualization and Computer Graphics*, vol. 29, no. 5, pp. 2710–2720, 2023.
[23] P. Jouet, V. Alleaume, A. Laurent, M. Fradet, T. Luo, and C. Baillard, "AR-Chat: an AR-based instant messaging system," in *2020 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*, 2020, pp. 153–157.
[24] C. Campolo, A. Molinaro, F. Romeo, A. Bazzi, and A. O. Berthet, "Full duplex-aided sensing and scheduling in cellular-V2X mode 4," in *Proceedings of the 1st ACM MobiHoc Workshop on technologies, models, and protocols for cooperative connected cars*, 2019, pp. 19–24.