5G Multicast-Broadcast On Air: An open SDR Testbed and Experimental Characterization of MBS

Rajendra Paudyal✉
George Mason University
Fairfax, Virginia, USA
rpaudyal@gmu.edu

Al Nahian Bin Emran
George Mason University
Fairfax, Virginia, USA
abinemra@gmu.edu

Rajendra Upadhyay
George Mason University
Fairfax, Virginia, USA
rupadhya@gmu.edu

Arupjyoti Bhuyan
Idaho National Laboratory
Idaho Falls, Idaho, USA
arupjyoti.bhuyan@inl.gov

Duminda Wijesekera✉
George Mason University
Fairfax, Virginia, USA
dwijesek@gmu.edu

Abstract

3GPP Release 17 integrates multicast-broadcast services (MBS) into the 5G system, yet open and instrumentable over-the-air implementations remain scarce. We present a 5G MBS testbed combining a 5G-MAG core, an MBS-enabled OCUDU gNB, and modified OpenAirInterface UEs operating over software-defined radios. The prototype implements broadcast-session signaling, multicast N3mb ingest, Group Radio Network Temporary Identifier (G-RNTI) addressed scheduling, blind UE decoding, and application delivery. In a 30-minute point-to-multipoint (PTM) trial, 1,801 sequence-numbered text messages were broadcast at one message per second. The first 29 were transmitted while the UEs were still completing registration and multicast group joining, and are therefore outside the observable interval rather than radio losses. Over the remaining interval, both UEs recovered all 1,772 messages with byte-exact payloads: 100% delivery over the observable interval and 98.389% of all messages transmitted. This open-source platform provides a reproducible foundation for over-the-air 5G MBS experimentation and future research.

CCS Concepts

• **Networks → Mobile networks**.

Keywords

5G, MBS, PTM, software-defined-radio, testbed, over-the-air

WiNTECH '26, Austin, TX, USA

ACM ISBN 979-8-4007-2879-2/2026/10
https://doi.org/10.1145/3831662.3844176

ACM Reference Format:
Rajendra Paudyal, Al Nahian Bin Emran, Rajendra Upadhyay, Arupjyoti Bhuyan, and Duminda Wijesekera. 2026. 5G Multicast-Broadcast On Air: An open SDR Testbed and Experimental Characterization of MBS. In *20th ACM Workshop on Wireless Network Testbeds, Experimental evaluation & Characterization (WiNTECH '26), October 26–30, 2026, Austin, TX, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3831662.3844176

1 Introduction

Cellular networks primarily rely on unicast delivery, which requires separate on-air radio resources for each user receiving their content. In this unicast mode, as the number of users increases, radio resource utilization and network packet volume grow approximately linearly, even if all users request the same information. To improve spectral efficiency and preserve network bandwidth consumption for one-to-many communications, 3GPP Release 17 introduced MBS, which enables a single on-air transmission to serve multiple receivers simultaneously. MBS succeeds LTE's enhanced Multimedia Broadcast Multicast Service (eMBMS) while redesigning it as a native capability of the 5G system. The MBS is specified in the 3GPP technical standard documents [1–7].

However, the 5G MBS is still under development and is not yet available as an end-to-end working system in an open-source ecosystem. Large-scale experimental platforms [9, 10] provide the radio infrastructure for 5G research but do not supply an MBS path. Open-source 5G software stacks [19–21, 26] and an MBS-capable core [8] likewise provide the necessary components in isolation, with no radio path that joins them. To the best of our knowledge, no open platform lets these components operate together and be measured end-to-end. This limits the questions we can study. A PTM channel provides neither HARQ nor ARQ, so scheduler conflicts, lost session contexts, and radio-thread starvation propagate directly to the application without being masked by retransmission. Determining which of these dominates in practice requires a working implementation on real hardware; simulation abstracts away the host and radio behavior, and closed prototypes do not permit the necessary instrumentation. This paper presents such a system. It carries an end-to-end 5G MBS multicast session over the air using SDR hardware and modified open software as defined in Fig. 1. An application function (AF/AS) in a core derived from Dockerized 5G-MAG [8] requests a session; the MB-SMF allocates a TMGI; the AMF signals the gNB over Next Generation Application Protocol (NGAP); and the MB-UPF forwards the multicast payload over a shared N3mb multicast GTP-U tunnel. The OCUDU gNB is extended with an MBS control plane, a multicast user-plane receiver, and a PTM MAC scheduler that transmits the stream on a G-RNTI-addressed Physical Downlink Shared Channel (PDSCH) through a USRP B210 on band n78. Two OpenAirInterface (OAI) UEs [19], each with its own B210, first complete standard 5G AKA registration, then blind-decode the same G-RNTI transport blocks, reassemble RLC-UM, and hand the recovered stream to the application. The same cell serves concurrent per-UE unicast, so PTM and point-to-point (PTP) run side by side. Two services run on this chain today: a sequence-numbered message service that sent one message per second, the shape of an alerting or notification service, and a live MPEG-TS video stream delivered to unmodified VLC players on both UEs.

Our measurements give the following results. The scheduler places multicast transmissions in ten of the twenty slots per radio frame, with transport blocks of up to 201 B, corresponding to a nominal capacity of 1.61 Mb/s §5.2. We note that this figure is an analytical envelope; the measured traces remain below it and do not establish a saturation point. Once multicast group joining was completed, the message service delivered byte-exact payloads to both listening UEs, with RLC backlog bounded at one Service Data Unit (SDU) and no PDCP discards §5.3. Video traverses the same path but does not yet play smoothly §5.4. The failures that most often invalidated a run originated outside the radio link: CSI-RS collisions, loss of MBS context on gNB restart, and TMGI pool exhaustion §6. This paper contributes the following:

• **A first-of-its-kind testbed implementation (§4):** The first publicly documented end-to-end, over-the-air 5G MBS system: open core, SDR-based gNB with a new PTM data path, multiple SDR UEs, a working message multicast service.

• **An experimental characterization (§5):** The first measured operating envelope of NR PTM on real radios: message delivery ratio, per-slot transport-block distribution, and user-plane arrival validation.

• **Experimental lessons that transfer (§6):** Three pitfalls with signatures and fixes: MTCH/CSI-RS collisions and a scheduler guard, MBS sessions lost on gNB reconnect, and TMGI pool exhaustion.

• **Reproducibility:** An open-source release of the complete testbed, including the MBS-enabled gNB and UE, core deployment scripts, configuration and runbooks, and logs. This helps the community reproduce our results, run its own over-the-air 5G MBS experiments, and further expand the work.

We are careful about the originality claims. The MBS core functions come from the 5G-MAG project [8]. Our contribution is the MBS implementation on the OCUDU gNB and the OAI UE, the integration of these components into a working over-the-air system, and the experimental results. The artifacts link: https://github.com/rajendra1124/e2e5Gmbs

2 Background

2.1 Service model

The 5G MBS radio interface supports point-to-point (PTP) and point-to-multipoint (PTM) delivery [3, 6]. In PTP, MBS data uses a UE-dedicated bearer scheduled with the UE's C-RNTI. In PTM, the multicast traffic channel (MTCH) is mapped to a PDSCH scheduled with a G-RNTI, allowing multiple UEs to decode the same transport blocks. The MCCH provides the configuration needed to receive the PTM transmission. Multicast requires no per-UE RAN session state, whereas multicast uses explicit UE joins and may include additional reliability procedures. Our testbed currently implements multicast service delivery.

2.2 Network Anatomy

Release 17 adds MBS-specific functions to the 5G core [2, 3]. The AF/AS requests an MBS session and provides the media, while the MB-SMF assigns a Temporary Mobile Group Identity (TMGI) and manages the session. The MB-UPF receives media over N6mb and forwards it to one or more gNBs over N3mb, which carries GTP-U traffic and may use IP multicast. The AMF signals the session to the NG-RAN through the NGAP `BroadcastSessionSetup` procedure. The gNB then installs the MBS radio context and schedules the MTCH using the configured MCCH and G-RNTI.

2.3 PTM complexity

The PTM path in our implementation provides no per-UE feedback: it uses neither HARQ ACK/NACK nor group-specific CSI reporting, and RLC operates in unacknowledged mode. Reliability therefore depends on robust modulation and coding, repetition, or application-layer FEC [11]. Because all receivers share the same transmission parameters, the MCS must accommodate the receiver with the weakest channel. We consequently use QPSK with a low code rate. These constraints guide both the implementation (§4) and its evaluation (§5).

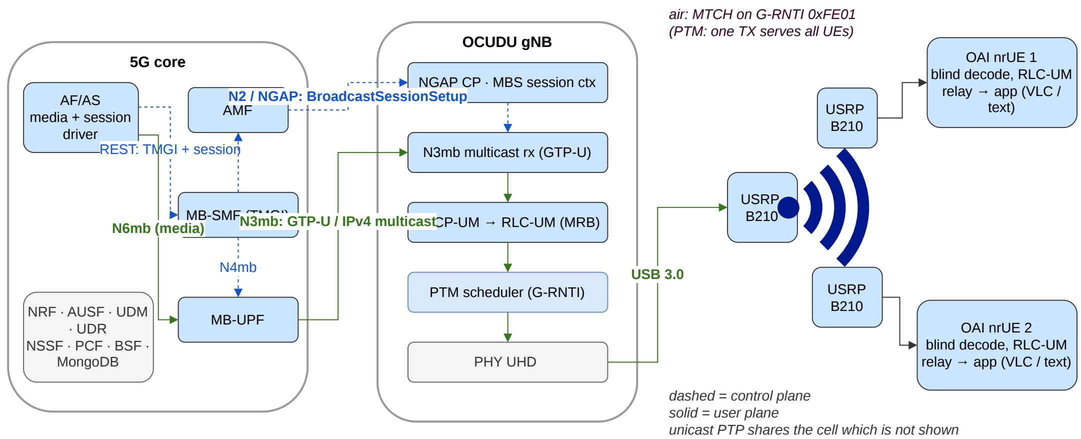

Figure 1: The end-to-end SDR MBS testbed. One media stream, one group-addressed radio transmission, two independent SDR receivers.

3 Related Work

LTE eMBMS and Rel-17 MBS: LTE eMBMS provided PTM delivery through a largely separate subsystem comprising the BM-SC, MBMS-GW, MCE, and MBSFN operation [17]. Its practical adoption was constrained by specialized infrastructure, architectural complexity, and deployment cost [27]. Release-17 MBS instead integrates multicast and broadcast into the service-based 5G architecture [3, 25]. Prior work describes its architectural and media-delivery extensions [23], evaluates mixed unicast-multicast operation through simulation [13], develops RAN designs for terrestrial broadcast [24], and surveys integration with non-3GPP broadcast systems [16]. These studies establish the design foundations of 5G MBS but do not provide an open, instrumentable over-the-air implementation.

Closed-platform prototypes: Commercial Rel-17 MBS prototypes have demonstrated cellular and high-tower deployments, capability verification, and multimedia delivery using commercial network equipment and handsets [22, 28, 29]. Although these trials establish deployment feasibility, their closed implementations prevent independent modification and cross-layer instrumentation of the control, protocol, scheduling, and radio paths.

Open experimental platforms: The 5G-MAG tools provide an Open5GS-derived MBS core, including MB-SMF, MB-UPF, and AF/AS components [8, 18], while srsRAN and OpenAirInterface provide extensible open-source NR stacks [14, 15, 19, 26]. General SDR platforms such as POWDER and Colosseum support experimental 5G research but do not supply an end-to-end MBS path [9, 10]. Our work closes this gap by integrating MBS session handling, multicast N3mb ingest, PTM scheduling, G-RNTI-addressed over-the-air delivery, and UE reception. To our knowledge, it is the first open and instrumentable end-to-end platform exposing this complete Release-17-aligned path on real radios.

4 Testbed Implementation

4.1 Hardware and software

Figure 1 shows the architecture of the 5G MBS system. Three commodity x86 hosts and three USRP B210s implement the whole network. The radio and MBS parameters used in the experiments are summarized in Table 1. **Core:** A Dockerized 5G-MAG/Open5GS-derived core [8]: NRF, MBS-capable AMF, AUSF, UDM, UDR, NSSF, PCF, BSF, MB-SMF, MB-UPF, a MongoDB subscriber store, and an AF/AS container that both drives the MBS session API [2] and sources media. Services are pinned to fixed IP addresses on a single Docker network so that addressing remains stable across restarts. This removed a recurring source of run-to-run failures during our experiments. **RAN:** OCUDU, a monolithic CU/DU [26], runs on the host (not containerized) with real-time tuning: isolated cores, `SCHED_FIFO` radio threads, and deep C-states disabled. The front end is a USRP B210 on USB 3.0 [12]. Band n78, 30 kHz SCS, 51 PRB (20 MHz), TDD, single layer.

UEs: Two OAI `nr-uesoftmodem` instances [19], each on its own host and B210. Each UE performs standard 5G AKA registration and carries its own unicast PDU session. For MBS, each UE blind-decodes the shared G-RNTI, reassembles RLC-UM, and forwards recovered MTCH SDUs over UDP to a local relay that feeds an unmodified VLC player. **Multicast sources:** Two senders drive the session from the AF/AS. A *message sender* emits one sequence-numbered text message per second ("`testmsg` N"), with N beginning at 0, both a service in its own right and our end-to-end reliability instrument, because every missing message N at a UE is a loss (no retransmission). A *video pipeline* (`ffmpeg`, H.264/AAC

Table 1: Testbed parameters

Parameter	Value
Band / duplex	n78 / TDD
Numerology / SCS	$\mu = 1$ / 30 kHz
Bandwidth	20 MHz (51 PRB)
DL center frequency	3489.42 MHz
Antennas	1T1R, single layer
MTCH PDSCH allocation	≈24 PRB (CORESET#0-sized)
MCS (MTCH)	QPSK, MCS 4 ($R \approx 0.30$)
DCI / search space	format 1_0, common SS, AL 4
G-RNTI / MCCH-RNTI	0xFE01 / 0xFFFF
RLC mode / PDCP SN	UM / 12 bit
N3mb G-PDU payload	1316 B (7 × 188)
SDR / sample rate	USRP B210 / 23.04 MS/s
Fronthaul	USB 3.0

in MPEG-TS at a controlled mux rate, swept over roughly 0.18 to 5 Mb/s) stresses the same path at the media rate. Either source is paced into GTP-U G-PDUs of $7 \times 188 = 1316$ payload bytes toward the MBS multicast group and TEID.

4.2 Control plane: NGAP multicast sessions at the gNB

A multicast session begins in the core, not in the gNB. The application function requests the MB-SMF to create an MBS session. The MB-SMF assigns the session its identity, the Temporary Mobile Group Identity (TMGI). The AMF then informs the gNB using the NGAP procedure `BroadcastSessionSetupRequest`. The existing OCUDU gNB has no support for this procedure, and we implemented it. Now the gNB decodes the request, installs the session, replies with `BroadcastSessionSetupResponse`, and also handles the corresponding release procedure that tears the session down. Figure 2 shows that after gNB attachment and UE registration, the AF/AS creates a multicast session through the MB-SMF, and the AMF signals the session to the gNB over NGAP. Inside the setup request, the core sends the gNB where the multicast data will arrive; this is the shared N3mb tunnel, an IP multicast address, and a tunnel identifier (TEID) on which the MB-UPF will send the user plane traffic.

Once the session is accepted, the rest of the setup is internal to the gNB. The CU-CP records the session in its session table, keyed by the TMGI. It then runs the standard E1AP procedure to create the multicast bearer context, followed by the F1AP procedure that creates the MBS context in the DU and binds the F1-U user-plane path. Finally, the DU manager hands the result to the MAC layer and its scheduler: the MCCH and MTCH channels are configured with their RNTIs (0xFFFF and 0xFE01), and the scheduler installs the

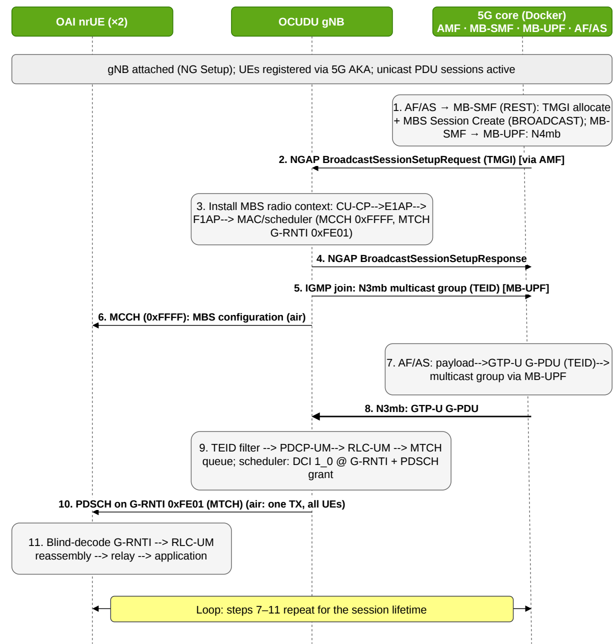

Figure 2: End-to-end MBS session establishment and multicast delivery flow.

MBS radio context. From this time onwards, the gNB log reports `MBS radio context installed`, and the cell is ready to transmit data to PTM. Although our gNB runs as a single process, we deliberately exercise the standard E1AP and F1AP interfaces internally, so that the same code would work in a disaggregated CU/DU deployment.

4.3 User plane: Multicast N3mb Ingest

N3mb is the user-plane path between the MB-UPF and the gNB: the multicast payload arrives as GTP-U packets carried over UDP. In Release-17 MBS, this tunnel is *shared*. Unlike a unicast N3 tunnel, which belongs to one UE, the N3mb tunnel belongs to the session, and the core may deliver it to an IP multicast group so that many gNBs can receive the same stream at once. Because the existing OCUDU gNB does not have a way to receive such traffic, we built a dedicated entry component (`mbs_ngu_receiver`) to receive the multicast stream. When the multicast session is set up, the receiver opens a UDP socket and joins the session's IPv4 multicast group using IGMP. The join is issued on the active multicast-capable network interface of the host. In our testbed, the core runs in Docker, so the multicast stream arrives on a Docker bridge interface. Each received packet is then checked before it is trusted. The receiver validates the GTP-U header, accepts only user-data packets, and verifies that the TEID matches the one assigned to this session. Packets belonging to any other tunnel that happens to share the same group

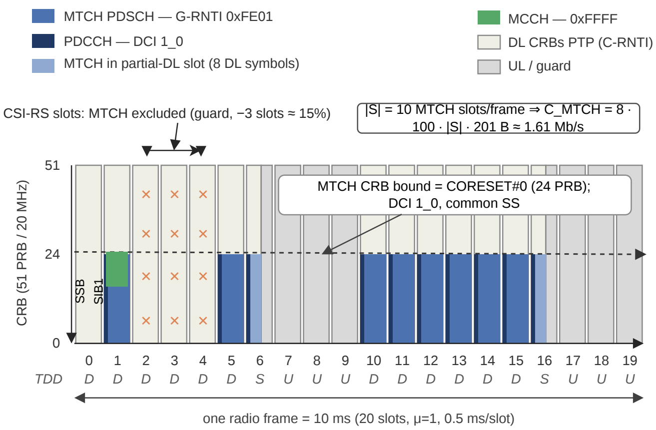

Figure 3: MTCH allocation in one 10-ms NR frame. Ten CSI-RS-safe MTCH slots with a measured 201-B maximum TBS give a nominal 1.61-Mb/s PTM ceiling; remaining resources support unicast PTP.

and port are dropped. Only after these checks, the OCUDU gNB receiver strips the tunnel header and hands over the inner payload to the MTCH data path.

Multicast delivery on N3mb extends the same resource-sharing property to the core network. On the air interface, a single PDSCH transmission serves all listening UEs. On N3mb, a single transmission by the MB-UPF serves all listening gNBs, so the incremental core-network cost of adding a base station to the multicast/broadcast area is comparable to the incremental cell cost of adding a UE. These savings would compound in a multi-cell deployment, although our testing was limited to a single gNB. Finally, for cores that deliver N3mb over plain unicast instead of multicast, the same ingress component accepts a unicast tunnel as a fallback, so the gNB interoperates with both delivery styles.

4.4 MAC: a PTM scheduler

The MAC scheduler decides every 0.5 ms which transmissions use the available radio resources. We extended the normal unicast scheduler with a PTM branch for MBS, as shown in Fig. 3. Before scheduling a broadcast grant, the gNB checks that the slot is downlink-capable, that MCCH or MTCH is due, and that the slot does not overlap with CSI-RS resources. Broadcast control is sent in the common search space over CORESET 0 using DCI format 1_0. The DCI CRC is scrambled with the group RNTI, G-RNTI 0xFE01, so all MBS-capable UEs can detect the same grant and decode the same PDSCH. We use aggregation level 4 to improve control-channel robustness. The scheduler then computes the transport block size following TS 38.214 [1]. In our configuration, the CORESET#0, limited allocation supports up to 201 bytes per slot with QPSK and MCS 4. This allows one

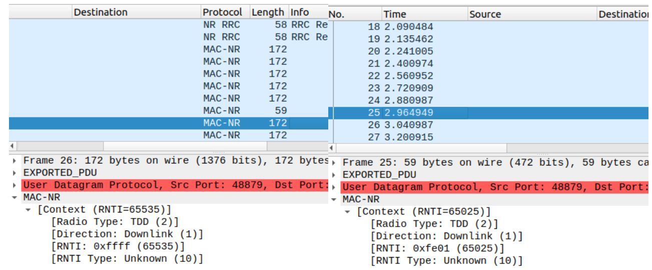

Figure 4: MCCH and MTCH allocations visible in the captured packets.

PTM transmission to serve all listening UEs, while the remaining radio resources continue to support normal unicast traffic.

4.5 UE: decode to media player

The UE follows the PTM scheduling design in §4.4. Since multicast data is not scheduled per UE, each receiver monitors the common search space over CORESET#0. We extend the OAI UE to blind-decode DCI candidates and check them against the G-RNTI, 0xFE01. When a matching DCI is found, the UE decodes the corresponding PDSCH and forwards the recovered transport block to the upper layers.

The received stream is reassembled using RLC-UM, so there are no retransmissions. Complete MTCH SDUs are forwarded over UDP to a lightweight relay on the UE host. The relay delivers text messages directly and sends MPEG-TS video to an unmodified VLC player. Using VLC gives us a realistic end-to-end test of the media path. The text service works reliably, while video also reaches the receiver but may pause because of packet loss and arrival jitter. These effects are small for low-rate text traffic but are more visible in continuous video playback.

5 Experimental Characterization

5.1 Methodology

All measurements come from instrumented logs on the running system, extracted by a released script into CSVs. On the gNB, we log, per slot, every MTCH grant (RNTI, TBS, slot index), N3mb accepted-PDU counters, RLC queue depth, PDCP discard events, and RF under-run events. Figure 4 shows captured packets confirming that MCCH and MTCH transmissions are genuinely on air. On each UE we log G-RNTI decode matches and delivered MTCH SDU counts. The sequence-numbered message service adds an application-level reality of per-UE delivery ratio and loss-gap lengths, with no retransmission machinery to mask anything. Three analytical envelopes anchor the measurements. With f frames/s, S

MTCH-eligible slots per frame, and $L_{\max}$ the maximum per-slot MTCH transport block in bytes, the air-capacity ceiling is

$$C_{\text{MTCH}} = 8fSL_{\max}. \tag{1}$$

With G-PDU payload b bytes, offered bitrate B implies an N3mb arrival rate

$$\lambda = B/(8b). \tag{2}$$

And with PDCP SN size s (window $W = 2^{s-1}$ SDUs) and mean SDU size $\bar{L}$, a total radio stall beginning at t_0 first drops an SDU.

$$t_{\text{drop}} = t_0 + 8W\bar{L}/B. \tag{3}$$

A fourth envelope constrains the platform itself: at a complex sample rate f_s with w wire bytes per sample, each direction of the SDR fronthaul needs $8f_s w$ bit/s, and duplex operation must fit inside the usable capacity of the host-radio link:

$$8f_s w \cdot 2 \leq C_{\text{link}}. \tag{4}$$

5.2 Group channel capacity

The measured per-slot MTCH TBS reaches 201 bytes on the ≈24-PRB, QPSK MCS-4 allocation. The TS 38.214 envelope for the same configuration gives ≈241 bytes, the difference being DM-RS overhead and TBS quantization. The model is an upper envelope, and the measured distribution sits properly beneath it. After the CSI-RS guard, $S \approx 10$ slots per 10 ms frame carry MTCH. Equation (1) then predicts $C_{\text{MTCH}} \approx 8 \cdot 100 \cdot 10 \cdot 201 \approx 1.61$ Mb/s. Measured behavior confirms streams at 160 kb/s and 1.2 Mb/s run drop-free for their full duration (zero PDCP discards, bounded RLC backlog). Figure 5 compares the analytical envelopes with the measurements. For Eq. (1), the scheduler trace directly establishes $S = 10$ used MTCH slot indices per 10-ms frame and a maximum observed transport-block size of $L_{\max} = 201$ B. Substitution gives the nominal full-occupancy envelope $C_{\text{MTCH}} = 8(100)(10)(201) = 1.608$Mb/s. This ceiling reflects our specific configuration rather than a limit of the architecture. Wider PRB allocations, higher-order modulation, and a larger number of eligible slots would each increase the value given by Eq 1.

5.3 One transmission for Multiple UEs

The message sender emits "`testmsg` N" once per second. Each UE logs every recovered message. Because PTM has no HARQ and no ARQ, the per-UE delivery ratio measures the whole pipeline. Every missing sequence number is an end-to-end loss. Figure 6 characterizes the text multicast service using a test trial. In the gNB, the apparent delay from the AF transmission timestamp to the first MTCH scheduling event has a median of 0.797 ms, with p95 and p99 values of 2.084 ms and 2.174 ms, respectively. These values assume a common clock reference and do not represent end-to-end UE latency. The scheduler carries 1,194 messages in one transport block and fragments 607 messages across two transport blocks. In the independent UE trial, the observable interval spans sequence numbers 29 to 1800. The preceding 29 messages were transmitted before the UEs completed registration and multicast group joining, so they were not available for reception and are excluded from the delivery ratio rather than counted as radio losses. Over the observable interval, both UEs recovered all 1,772 messages, with every payload byte delivered at exactly 100% delivery over the interval, corresponding to 98.4% of the 1,801 messages transmitted. This operating point is relevant beyond our demonstration. Public-safety alerting, C-V2X notification, and IoT group command are low-rate services with strict reliability requirements, which is the regime for which a feedback-less multicast channel at a conservative MCS is best suited.

5.4 Media-rate PTM: video stream

Live MPEG-TS video traverses the same chain to unmodified VLC players on both UEs. However, playback is not smooth, and our runs exhibit recurring stalls and decode artifacts. We report this as-is and treat it as a measurement subject. Three conditions constrain the experimental outcomes, and violating any one of them dominates: (i) pace the sender at the encoded media rate—over-pacing a low-rate clip causes PCR discontinuities and player stalls with zero radio loss; (ii) the offered rate must remain below the bound of Eq. (1); increasing it toward 5 Mb/s caused delivery to fail, since a fixed group allocation does not adapt to higher offered load; and (iii) the fronthaul rate must satisfy Eq. (4). Satisfying all three constraints is necessary but not sufficient for smooth playback, and identifying the remaining cause requires further measurement.

6 Lessons Learned

Three operational failures repeatedly compromised otherwise plausible experiments. Each has a recognizable signature and a corresponding testbed-level safeguard. These lessons illustrate that reproducible MBS experimentation depends not only on protocol implementation but also on host configuration, radio-resource isolation, and session lifecycle management.

L1: Isolate feedback-less transmissions from deterministic resource conflicts. PTM provides no per-UE HARQ or ARQ feedback, so control or data collisions may appear only as unexplained receiver failures. We therefore exclude MTCH scheduling from the configured CSI-RS slots $\{2, 3, 4\}$. The archived scheduler trace confirms that no MTCH grants occur in these slots. This guard removes three of 13 otherwise eligible MTCH opportunities 23% of the pre-guard opportunity set, or 15% of all frame slots. A guard-off control

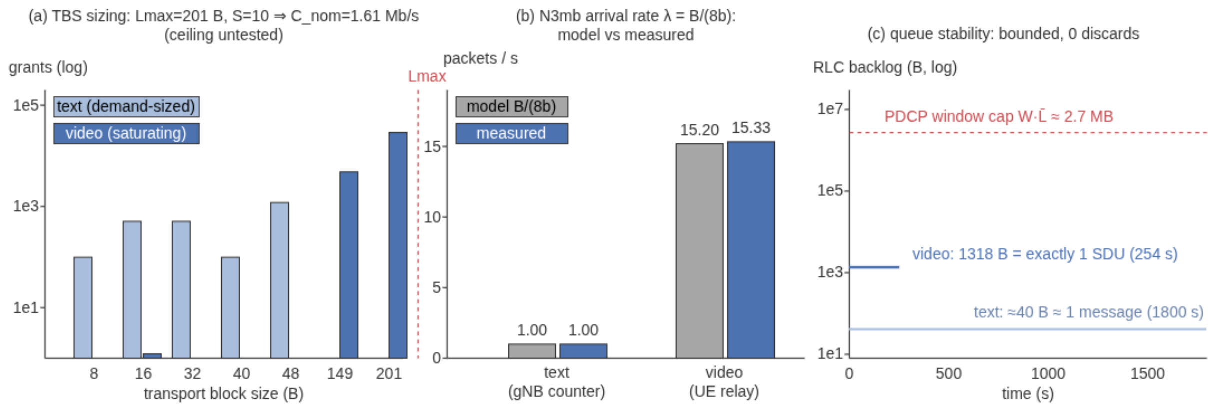

Figure 5: Analytical envelopes and measured traces for MTCH capacity, N3mb arrival rate, and RLC queue stability.

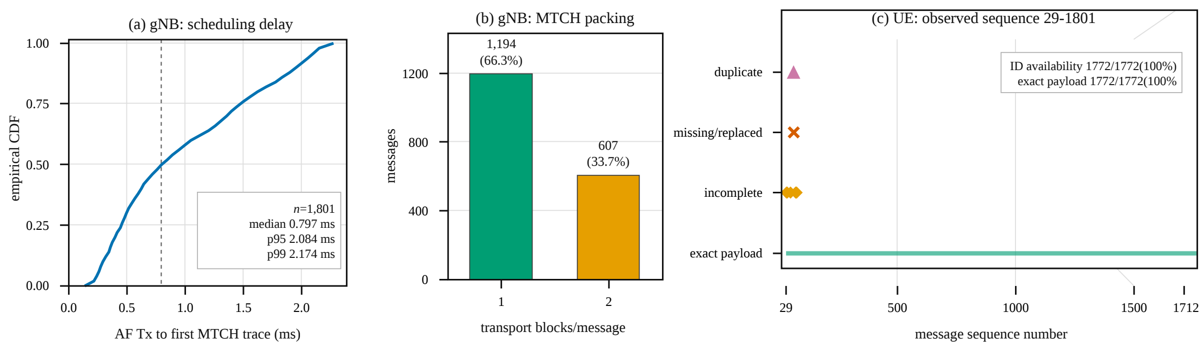

Figure 6: Thirty-minute text multicast trial: scheduling delay, MTCH packing, and byte-exact reception of all 1,772 observable messages.

experiment is still required to quantify the corresponding loss reduction.

L2: Treat RAN-side MBS context as soft state. Following a gNB restart, registration and unicast service may recover while multicast traffic remains absent. The core can retain the MBS session and continue sending data even though the gNB has lost its local radio context. Our tooling therefore reactivates the multicast session after every gNB connection or restart. Session restoration is therefore handled as an explicit step in our run procedure.

L3: Monitor session identifiers as well as packet delivery. Long-running campaigns can fail through control-plane state exhaustion. Unreleased sessions in the AF-side workflow accumulated TMGIs until the MB-SMF identifier pool could no longer create new sessions. We monitor both active-session and identifier counts and reset the affected functions when stale state accumulates. A complete solution requires a reliable session-release path. In the meantime, monitoring identifier allocation prevents this failure from being misdiagnosed as a radio or software problem. These constraints are not addressed by the functional MBS specifications, yet they determine whether an over-the-air experiment produces valid results. We include detection and recovery procedures for all three in the released artifact so that other groups can reproduce our runs without rediscovering them.

7 Conclusion

We presented an open, end-to-end, over-the-air implementation of Release-17-aligned 5G MBS, integrating an MBS-capable core, an SDR-based gNB with a PTM data path, and two SDR UEs receiving a shared group-addressed transmission alongside ordinary unicast service. In a 30-minute trial, the application function sent 1,801 sequence-numbered messages at one message per second. The first 29 were transmitted during UE registration and multicast group joining; over the subsequent interval, both UEs recovered every message with byte-exact payloads. The same path delivered live MPEG-TS video to both receivers, although VLC playback was not consistently smooth, exposing the importance of media pacing, timestamp continuity, and receiver buffering.

Scheduler measurements identified ten used MTCH slot indices and transport blocks up to 201 B, defining a nominal full-occupancy envelope of 1.61 Mb/s rather than a measured saturation limit. Together with the observed bounded queues and SDR fronthaul requirements, these results establish both the capabilities and current limits of the platform. By releasing the implementation, configurations, and measurement workflow, we provide a reproducible foundation for future research on reliable, scalable, and media-aware 5G PTM delivery. The platform also enables future over-the-air studies of blind link adaptation, application-layer FEC, PTM/PTP switching thresholds, multi-gNB synchronization, mobility, and coverage-edge receivers.

References

[1] 3GPP. 2020. *NR; Physical layer procedures for data.* Technical Report TS 38.214. 3rd Generation Partnership Project. Release 15.

[2] 3GPP. 2022. *5G System; Multicast-Broadcast Session Management Services; Stage 3.* Technical Report TS 29.532. 3rd Generation Partnership Project. Release 17.

[3] 3GPP. 2022. *Architectural enhancements for 5G multicast–broadcast services.* Technical Report TS 23.247. 3rd Generation Partnership Project. Release 17.

[4] 3GPP. 2022. *NR; Medium Access Control (MAC) protocol specification.* Technical Report TS 38.321. 3rd Generation Partnership Project. Release 17.

[5] 3GPP. 2022. *NR; Multiplexing and channel coding.* Technical Report TS 38.212. 3rd Generation Partnership Project. Release 17.

[6] 3GPP. 2022. *NR; NR and NG-RAN Overall description; Stage-2.* Technical Report TS 38.300. 3rd Generation Partnership Project. Release 17.

[7] 3GPP. 2022. *NR; Radio Resource Control (RRC) protocol specification.* Technical Report TS 38.331. 3rd Generation Partnership Project. Release 17.

[8] 5G-MAG. 2025. Reference Tools: MBS examples (rt-mbs-examples). https://github.com/5G-MAG/rt-mbs-examples.

[9] Leonardo Bonati, Pedram Johari, Michele Polese, Salvatore D'Oro, Subhramoy Mohanti, Miead Tehrani-Moayyed, Davide Villa, Shweta Shrivastava, Chinenye Tassie, Kurt Yoder, Ajeet Bagga, Paresh Patel, Ventz Petkov, Michael Seltser, Francesco Restuccia, Abhimanyu Gosain, Kaushik R. Chowdhury, Stefano Basagni, and Tommaso Melodia. 2021. Colosseum: Large-Scale Wireless Experimentation Through Hardware-in-the-Loop Network Emulation. In *2021 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN).* 105–113. doi:10.1109/DySPAN53946.2021.9677430

[10] Joe Breen, Andrew Buffmire, Jonathon Duerig, Kevin Dutt, Eric Eide, Mike Hibler, David Johnson, Sneha Kumar Kasera, Earl Lewis, Dustin Maas, et al. 2020. POWDER: Platform for open wireless data-driven experimental research. In *Proceedings of the 14th International Workshop on Wireless Network Testbeds, Experimental evaluation & Characterization.* 17–24.

[11] Lorenz Minder et.al. 2011. RaptorQ Forward Error Correction Scheme for Object Delivery. IETF RFC 6330.

[12] Ettus Research. 2025. USRP B210 and the UHD driver. https://www.ettus.com.

[13] Eduardo Garro, Manuel Fuentes, Jose Luis Carcel, Hongzhi Chen, De Mi, Fasil Tesema, Jordi Joan Gimenez, and David Gomez-Barquero. 2020. 5G mixed mode: NR multicast-broadcast services. *IEEE Transactions on broadcasting* 66, 2 (2020), 390–403.

[14] Ismael Gomez-Miguelez, Andres Garcia-Saavedra, Paul D. Sutton, Pablo Serrano, Cristina Cano, and Doug J. Leith. 2016. srsLTE: an open-source platform for LTE evolution and experimentation. In *Proceedings of the Tenth ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation, and Characterization (WiNTECH '16).* ACM, 25–32. doi:10.1145/2980159.2980163

[15] Florian Kaltenberger, Aloizio P Silva, Abhimanyu Gosain, Luhan Wang, and Tien-Thinh Nguyen. 2020. OpenAirInterface: Democratizing innovation in the 5G Era. *Computer Networks* 176 (2020), 107284.

[16] Rashmi Kamran, Pranav Jha, Shwetha Kiran, Abhay Karandikar, Prasanna Chaporkar, Anindya Saha, and Arindam Chakraborty. 2022. A survey on multicast broadcast services in 5G and beyond. In *2022 National Conference on Communications (NCC).* IEEE, 344–349.

[17] David Lecompte and Frédéric Gabin. 2012. Evolved multimedia broadcast/multicast service (eMBMS) in LTE-advanced: Overview and Rel-11 enhancements. *IEEE Communications Magazine* 50, 11 (2012), 68–74.

[18] Open5GS. 2025. Open5GS: open-source 5G core. https://open5gs.org.

[19] OpenAirInterface Software Alliance. 2025. OpenAirInterface 5G NR softmodem. https://gitlab.eurecom.fr/oai/openairinterface5g.

[20] Rajendra Paudyal, Rajendra Upadhyay, Al Nahian Bin Emran, Brienne Douglas, Lisa Donnan, and Duminda Wijesekera. 2026. Pi5G: A lightweight 5G Testbed on a Raspberry Pi5. In *2026 22nd International Conference on DCOSS-IoT.* 1360–1365.

[21] Rajendra Paudyal, Rajendra Upadhyay, Al Nahian Bin Emran, and Duminda Wijesekera. 2025. gNB-based Local Breakout for URLLC in industrial 5G. In *2025 IEEE 30th International Workshop on CAMAD.* 1–6. doi:10.1109/CAMAD67323.2025.11229912

[22] Zeng Qingjun, Song Jian, Liang Xiangjun, and Pan Changyong. 2022. 5G MBS System Capability Verification in Commercial Scenario. In *2022 International Conference on Electrical Engineering and Photonics (EExPolytech).* IEEE, 29–32.

[23] Alberto Rico-Alvariño, Imed Bouazizi, Miguel Griot, Prasad Kadiri, Le Liu, and Thomas Stockhammer. 2022. 3GPP Rel-17 extensions for 5G media delivery. *IEEE Transactions on Broadcasting* 68, 2 (2022), 422–438.

[24] Mikko Säily, Carlos Barjau Estevan, Jordi Joan Gimenez, Fasil Tesema, Wei Guo, David Gomez-Barquero, and De Mi. 2020. 5G radio access network architecture for terrestrial broadcast services. *IEEE Transactions on broadcasting* 66, 2 (2020), 404–415.

[25] Vinay Kumar Shrivastava, Sangkyu Baek, and Youngkyo Baek. 2022. 5G Evolution for Multicast and Broadcast Services in 3GPP Release 17. *IEEE Communications Standards Magazine* 6, 3 (2022), 70–76. doi:10.1109/MCOMSTD.0001.2100068

[26] Software Radio Systems. 2025. srsRAN Project: open-source O-RAN CU/DU. https://gitlab.com/ocudu/ocudu.

[27] Domenico Striccoli, Giuseppe Piro, and Gennaro Boggia. 2019. Multicast and Broadcast Services Over Mobile Networks: A Survey on Standardized Approaches and Scientific Outcomes. *IEEE Communications Surveys & Tutorials* 21, 2 (2019), 1020–1063.

[28] Liang Xiangjun, Pan Changyong, and Zeng Qingjun. 2024. Multimedia service trials based on 5G NR MBS. In *2024 8th International Conference on Information, Control, and Communication Technologies (ICCT).* IEEE, 1–4.

[29] Qingjun Zeng, Shuang Li, and Jian Song. 2021. Field test for 5G NR multicast and broadcast services. In *2021 International Conference on Electrical Engineering and Photonics (EExPolytech).* IEEE, 129–133.