

Protocol Data Unit Set-based QoS-aware Downlink Scheduling for Extended Reality Applications in 6G

Moyukh Laha*, Gabriel-Miro Muntean*

*School of Electronic Engineering, Dublin City University, Ireland
{moyukh.laha, gabriel.muntean}@dcu.ie

Abstract—This paper presents a downlink scheduling framework for Extended Reality (XR) applications in 6G networks. Building on the 3GPP Release 18 Protocol Data Unit (PDU) Set-based quality of service (QoS) framework, we propose a *PDU Set Delay Tracking Queue (PDTQ)* mechanism that provides real-time visibility into each PDU Set’s remaining delay budget. Unlike existing approaches relying on indirect urgency metrics, PDTQ directly tracks PDU Set aging across queue positions for precise identification of delay-critical data. Leveraging this, we develop the *PDTQ-Guided XR (PiXR)* scheduler, an adaptive heuristic that allocates resources based on delay criticality and instantaneous channel conditions. Simulations show that PiXR achieves approximately 4–5× higher satisfied user counts than state-of-the-art PDU Set-aware schedulers under limited resources, while maintaining comparable or better performance under adequate conditions.

Index Terms—Extended reality, PDU Set aware scheduling, 6G.

I. INTRODUCTION

Extended Reality (XR) encompasses immersive technologies, including Virtual Reality (VR), Augmented Reality (AR), and Mixed Reality (MR), and is poised to revolutionize applications ranging from remote collaboration to industrial training. As a result, XR is identified as a critical service class for 6G networks [1]. However, XR presents unique challenges: it simultaneously demands high data rates for volumetric content, ultra-low latency for motion-to-photon responsiveness, and strict reliability to maintain immersion. This stringent coupling places XR at the intersection of enhanced Mobile Broadband (eMBB) and Ultra-Reliable Low-Latency Communication (URLLC) [1], making it particularly challenging to support.

A fundamental characteristic of XR traffic is that large application-layer frames (e.g., video frames) are fragmented into multiple interdependent Protocol Data Units (PDUs), all of which must be delivered within a tight delay budget to ensure frame decodability. The loss or delayed arrival of even a single PDU renders the entire frame unusable, wasting transmission resources and degrading user experience [2]. Recognizing this, 3GPP Release 18 introduced the PDU Set framework [3], [4], wherein each PDU Set represents all PDUs carrying a single media unit. This framework extends traditional IP-layer QoS with XR-specific parameters: the PDU Set Delay Budget (PSDB) defines the delivery deadline for the entire set [5]; the PDU Set Error Rate (PSER) bounds incomplete delivery rates; the PDU Set Importance (PSI) enables content-based prioritization; and the PDU Set Integrated Handling Indicator

(PSIHI) signals that partial delivery is useless, allowing early discard. Importantly, 3GPP defines the XR capacity as the number of XR users supported per cell [6], departing from traditional throughput-centric metrics. An XR user is deemed satisfied only if at least 99% of its PDU Sets are successfully received within the PSDB [2]—a new Quality of Service (QoS) metric that jointly targets delay and reliability. This binary satisfaction criterion requires schedulers that respect frame-level dependencies and delay constraints rather than maximization of aggregate throughput.

MAC-layer schedulers must therefore operate at PDU Set granularity and account for inter-packet dependencies with strict timing constraints. However, such PDU Set-aware schedulers remain scarce in the literature. Legacy schedulers, such as Proportional Fair (PF), treat packets independently and are fundamentally unsuitable for structured XR traffic. A few XR-aware approaches have been proposed: [7] models scheduling as a periodic Markov Decision Process but assumes periodic arrivals and uniform packet sizes, limiting practical applicability. [8] proposes Frame-Level Integrated Transmission, but the scheme lacks accurate delay tracking of the PDUs. Architectural enhancements, such as the application-aware MAC in [9] and AI-assisted provisioning in [10], improve user satisfaction but do not target PDU Set-level XR scheduling. PDU Set-aware scheduling has received limited attention. The heuristic in [2] computes priority for user i and PDU Set k as $m_{k,i} = e^{\alpha_{k,i}} / \beta_{k,i}$, where $\alpha_{k,i}$ denotes transmission progress and $\beta_{k,i}$ reflects normalized urgency. While this balances completion likelihood and temporal urgency, it lacks explicit delay tracking. Without real-time visibility into remaining delay budgets, the scheduler cannot differentiate truly delay-critical sets from those with sufficient slack. This limitation becomes severe under tight PSDBs or bandwidth constraints. The DROX scheduler [11] introduces Delay Tracking Queues, which were subsequently extended to multisensory scenarios in [12]. DROX formulates scheduling as a Mixed Integer Linear Program (MILP) but assumes complete knowledge of future traffic and channel states. This assumption leads to NP-hard complexity, rendering real-time operation per Transmission Time Interval (TTI) infeasible.

To fill this research gap, we propose a *PDU Set Delay Tracking Queue (PDTQ)* mechanism that maintains per-user, per-PDU Set queues indexed by remaining delay budget, enabling real-time visibility into PSDB violation risks. Leveraging this structure, we design the *PDTQ-Guided XR (PiXR) Scheduler*, which integrates delay criticality with instantaneous channel

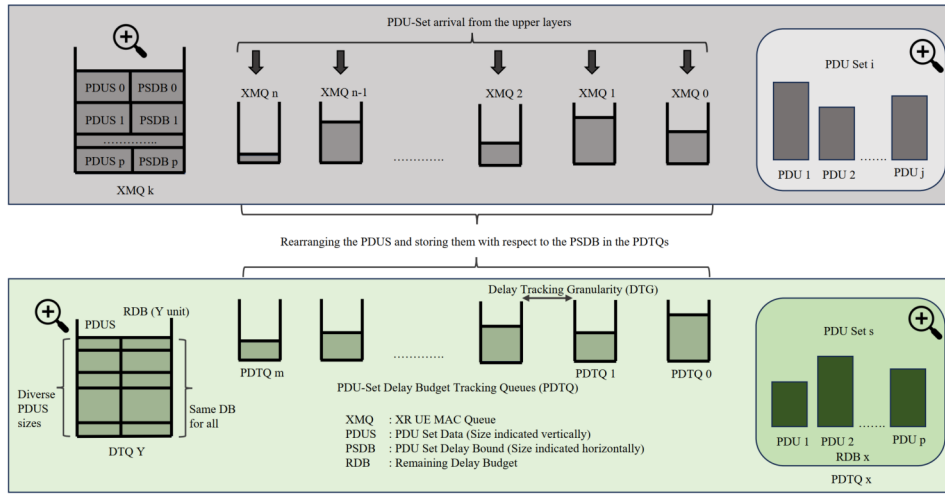


Figure 2. PDU Set Delay Tracking Mechanism: PDU Sets from the upper layers come to the individual XR user MAC queue, which are then placed to the appropriate PDTQs based on their PSDB. A random XR MAC queue, one PDTQ content (same RDB for all the PDUs), and PDU Sets are zoomed in.

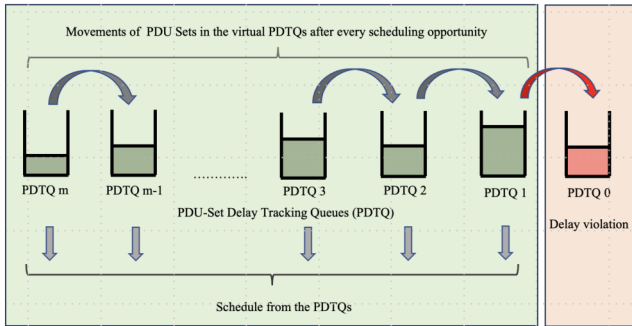


Figure 3. PDU Set movements through the PDTQs: precise delay tracking.

B. Algorithm Description

PiXR operates in two modes based on system load, as detailed in Algorithm 1. Phase 1 analyzes PDTQs for each user i , computing three key metrics. First, $\rho_i^{\text{crit}} = \lceil Q_{i,1}/R_i \rceil$ estimates the RBs needed to clear PDTQ-1, where $Q_{i,1}$ is the buffered data in PDTQ-1 and R_i is the user's instantaneous rate per RB. Second, $\rho_i^{\text{urg}} = \lceil (Q_{i,1} + Q_{i,2})/R_i \rceil$ estimates the RBs needed to clear both PDTQ-1 and PDTQ-2. This two-tier classification is critical: clearing only PDTQ-1 merely defers failure, whereas clearing both ensures short-term survivability. Third, the pressure score $P_i = \sum_{k=1}^{|\mathcal{K}|} w_k \cdot Q_{i,k} / \sum_{k=1}^{|\mathcal{K}|} Q_{i,k}$ uses exponentially decaying weights— $w_1 = 10$, $w_2 = 5$, $w_3 = 2$, $w_k = 1$ for $k \geq 4$ —reflecting diminishing urgency with greater delay slack.

Phase 2 determines the operating mode using the scarcity ratio $\Gamma = \sum_{i \in \mathcal{N}} \rho_i^{\text{crit}} / C$, where C is the total available RBs. If $\Gamma > 1.5$, critical demand exceeds capacity by over 50%, triggering scarcity mode. PiXR then performs strategic user selection by computing saveability $\sigma_i = P_i / \rho_i^{\text{urg}}$ for users satisfying $\rho_i^{\text{urg}} \leq 0.3C$ (i.e., not overly resource-intensive), and greedily selecting users in descending σ_i until $\sum_{i \in \mathcal{S}} \rho_i^{\text{urg}} \leq C$. This ensures selected users receive sufficient resources to avoid

failures.

Under normal mode ($\Gamma \leq 1.5$), PiXR computes channel-aware priority $\Pi_i = P_i \cdot \sqrt{R_i/\bar{R}}$, where \bar{R} is the mean channel rate. The square root balances exploitation and fairness: linear weighting favors strong channels excessively, while channel-blind priority wastes resources on poor links. Resources are allocated in three passes: PDTQ-1 by descending Π_i , then PDTQ-2, and finally lower-priority queues, ensuring critical deadlines are met first.

PiXR's effectiveness stems from three principles. First, *direct urgency measurement* via PDTQ position captures actual deadline pressure—e.g., 1000 bits in PDTQ-1 require immediate action, unlike the same volume in PDTQ-5. Second, *adaptive mode switching* avoids universal failure under overload by sacrificing unsalvageable users. Third, *saveability assessment* balances urgency with resource efficiency, prioritizing users who are both urgent and have better channels to serve.

Note that the scheme's hyperparameters (e.g., weights, $\Gamma = 1.5$) are manually tuned specifically for the simulation environment. These may be adapted dynamically via machine learning, which we leave to future work.

C. Computational Complexity

PiXR maintains a complexity of $O(|\mathcal{N}| \cdot (|\mathcal{K}| + \log |\mathcal{N}|))$, suitable for real-time deployment at sub-millisecond timescales (\mathcal{N} is the number of users; \mathcal{K} is the maximum PDTQ index). PDTQ analysis requires $O(|\mathcal{N}| \cdot |\mathcal{K}|)$ operations, sorting requires $O(|\mathcal{N}| \log |\mathcal{N}|)$, and allocation requires $O(|\mathcal{N}|)$ in both modes. This complexity is comparable to traditional schedulers like PF while delivering substantially superior performance.

IV. PERFORMANCE EVALUATION

A. Simulation Environment

We implement a 5G NR simulation framework in Python to assess PiXR performance. The evaluation environment consists of a single gNB serving multiple XR users, with scheduling

Algorithm 1 PiXR: PDTQ-Guided XR Scheduler

```

1: Input:  $\mathcal{K}$  (PDTQ indices),  $\mathcal{N}$  (users),  $C$  (max RBs)
2: Input:  $Q_{i,k}$  (PDTQ state),  $R_i$  (channel rate)
3: Output:  $\{x_i\}_{i \in \mathcal{N}}$  (RB allocation per user)
4: Phase 1: PDTQ Analysis
5: for each user  $i \in \mathcal{N}$  do
6:    $\rho_i^{\text{crit}} \leftarrow \lceil Q_{i,1}/R_i \rceil$  ▷ RBs for PDTQ-1
7:    $\rho_i^{\text{urg}} \leftarrow \lceil (Q_{i,1} + Q_{i,2})/R_i \rceil$  ▷ RBs for PDTQ-1,2
8:    $P_i \leftarrow \sum_{k=1}^{|\mathcal{K}|} w_k \cdot Q_{i,k} / \sum_{k=1}^{|\mathcal{K}|} Q_{i,k}$  ▷ Pressure score
9: end for
10: Phase 2: Strategy Selection
11:  $\Gamma \leftarrow \sum_{i \in \mathcal{N}} \rho_i^{\text{crit}} / C$  ▷ Scarcity ratio
12: if  $\Gamma > 1.5$  then
13:   Scarcity Mode:
14:    $\mathcal{S} \leftarrow \emptyset$ 
15:   for each user  $i$  with  $\rho_i^{\text{urg}} \leq 0.3C$  do
16:      $\sigma_i \leftarrow P_i / \rho_i^{\text{urg}}$  ▷ Saveability
17:   end for
18:   Sort users by  $\sigma_i$  descending
19:   Greedily add users to  $\mathcal{S}$  until  $\sum_{i \in \mathcal{S}} \rho_i^{\text{urg}} \leq C$ 
20:   Allocate RBs to users in  $\mathcal{S}$  only
21: else
22:   Normal Mode:
23:   for each user  $i \in \mathcal{N}$  do
24:      $\Pi_i \leftarrow P_i \cdot \sqrt{R_i/\bar{R}}$  ▷ Channel-aware priority
25:   end for
26:   Sort users by  $\Pi_i$  descending
27:   Allocate RBs: Pass 1 (PDTQ-1)  $\rightarrow$  Pass 2 (PDTQ-2)
    $\rightarrow$  Pass 3 (rest)
28: end if
29: return  $\{x_i\}_{i \in \mathcal{N}}$ 

```

decisions made at the MAC layer according to the parameters outlined in Table I, in accordance with 3GPP standards. Users are randomly distributed within the cell coverage area, and each 1-second simulation is averaged over 25 independent runs with varied random seeds. Benchmark comparisons include conventional schedulers—*Proportional Fair (PF)*, *Round Robin (RR)*, and *Maximum Rate (MR)*—as well as the PDU Set-aware heuristic, denoted as *PSH* [2].

Traffic generation follows the 3GPP XR model [14], representing video as pseudo-periodic frames with sizes and arrival jitter sampled from truncated Gaussian distributions to capture real-world variability. Performance assessment relies on two metrics: (i) *Satisfied User Count*—the number of users with 99% of PDU Sets delivered within their PSDB; and (ii) *Delay-Reliable Throughput (DRLT)*—the fraction of PDU Sets delivered within PSDB constraints, averaged over all users, indicating application-layer usability.

B. Results and Discussions

We now examine the impact of various parameters on the performance of the proposed PiXR scheduling scheme in comparison to existing approaches.

Table I
SIMULATION AND XR TRAFFIC PARAMETERS

Parameter	Value	Parameter	Value
Carrier freq.	3.5 GHz	Channel model	3GPP 38.901 [15]
BS max tx power	43 dBm	UE tx power	23 dBm
Data rate (d_r)	30, 60 Mbps	Frame rate (f_r)	60, 120 fps
Numerology (μ)	1	Max bandwidth	100 MHz
Delay bound	5–15 ms	Target reliability	99%
Mobility model	Stationary users	Mean DU size (\bar{S})	$\frac{d_r \times 10^6}{f_r \times 8}$ [Bytes]
Std. dev. DU size (σ_S)	$0.105 \times \bar{S}$	Packet truncation range	$[0.5 \times \bar{S}, 1.5 \times \bar{S}]$
Mean inter-arrival time (\bar{T})	$\frac{1}{f_r}$ [s]	Jitter mean (\bar{J})	0 ms
Jitter std. dev. (σ_J)	2 ms	Cell range (R)	500m

1) *Effect of Input XR Users:* Fig. 4 compares the scheduler performance across varying user loads. The left plot shows satisfied user counts and the right plot presents DRLT.

Traditional schedulers fail across all loads, satisfying at most 1 user. Lacking PDU Set awareness, they deliver partial PDU Sets useless at the application layer. Additionally, without delay tracking mechanisms, substantial data arrives after PSDB expiration, compounding their failure. This fundamental mismatch renders them unsuitable for XR traffic.

At low to moderate load (5-10 users), PSH and PiXR perform comparably. With adequate resources, both operate in normal mode. Critical divergence arises under high load. At 15 users, PSH collapses to 2-3 satisfied users while PiXR maintains 6-8 users—a 3-4x improvement. At 17 users, PSH satisfies fewer than 1 user, while PiXR serves approximately 5 users. This gap reveals PiXR’s core advantage: direct delay tracking through PDTQ structure versus indirect urgency estimation of PSH. In overloaded conditions, PSH still spreads resources across all users, resulting in universal partial service and universal failure. PiXR’s PDTQ structure provides real-time visibility into each PDU Set’s delay criticality. When resource scarcity is detected under high load, PiXR switches to strategic abandonment—computing saveability, concentrating resources on targeted users, and deliberately sacrificing resource-inefficient ones. This binary decision (fully satisfy a subset versus partially serve all) maximizes satisfied user count under constraints. This is in accordance with the 3GPP which defines XR capacity as the number of XR users supported per cell [6].

Delay-reliable throughput (DRLT) shows similar trends. Under low load, PSH and PiXR perform comparably; under high load, PiXR achieves approx. 10% higher DRLT. The difference is smaller than the satisfied user count because these results average over all users. PiXR concentrates resources to provide 99% reliability to selected users while sacrificing others (evident in the left plot), while PSH degrades everyone uniformly, achieving similar average throughput but satisfying

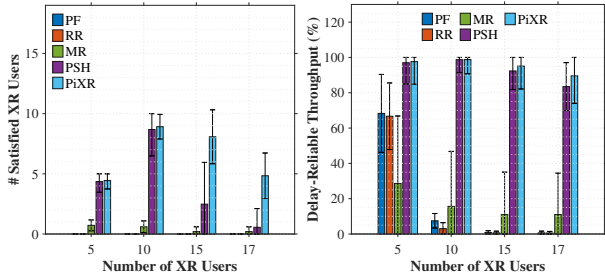


Figure 4. Effect of the number of input XR users: $d_r = 30\text{Mbps}$, $f_r = 60\text{fps}$, PSDB = 5ms, Target Reliability = 99%, $f_c = 3.5\text{ GHz}$, FR1, BW = 100MHz, $\mu = 1$, $R = 500\text{m}$, Mobility: Stationary.

no users. According to 3GPP standards, the objective is to maximize the satisfied user count [6], which is exactly the focus of this work.

These results inform admission control policies. At 15 users, PiXR satisfies 6-8 users while PSH collapses. An admission controller using PiXR feedback could proactively reject a few users at session establishment rather than accepting all and satisfying none, transforming uncontrolled degradation into managed capacity with guaranteed QoS. PSH’s graceless degradation provides no clear admission signal, making capacity planning infeasible.

2) *Delay Statistics Analysis*: Fig. 5 presents delay distributions for satisfied users under moderate (10 users) and high (17 users) load. Note that these statistics reflect only users meeting satisfaction criteria—not all users, PF and RR could not support even a single user; therefore their delay statistics are not shown in the plot. Consequently, PSH and PiXR show similar delay characteristics among their respective satisfied users, although PiXR serves 4-5 times more satisfied users under heavy load (as shown in Fig. 4).

At 10 users, tight distributions indicate adequate resources: MR maintains a median $\sim 0.5\text{ ms}$ (serving only the best-channel users), while PSH and PiXR exhibit medians of $\sim 0.5\text{-}0.6\text{ ms}$ with means of $\sim 0.9\text{-}1.0\text{ ms}$. At 17 users, a divergence emerges. PSH widens significantly with a median of 1.5 ms, a mean of 1.8 ms, and Q1-Q3 spanning 1.0-2.5 ms. PiXR maintains tighter control: median 1.0 ms, mean 1.4 ms, Q1-Q3 of 0.5-2.0 ms. The lower mean delay and narrower spread demonstrate that strategic abandonment not only increases the satisfied user count, but also ensures predictable, consistent service for selected users. In summary, PiXR achieves 4-5 times higher satisfied user count while maintaining comparable or better delay performance.

3) *Effect of System Bandwidth*: Figure 6 examines the scheduler performance under bandwidth constraints with 10 XR users. The stark performance contrast between 50 MHz and 100 MHz reveals critical insights into resource scarcity handling.

At low bandwidth (50 MHz), the gap between PiXR and competing schemes becomes pronounced. Traditional schedulers satisfy fewer than 1 user with throughput below 20%. PSH collapses entirely: 0 satisfied users despite achieving $\sim 75\%$ DRLT. This counter-intuitive result exposes PSH’s

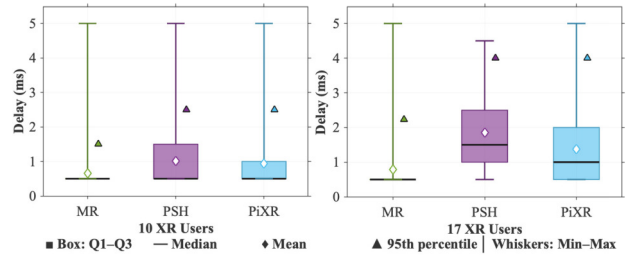


Figure 5. Detailed delay analysis of the satisfied XR users: $d_r = 30\text{Mbps}$, $f_r = 60\text{fps}$, PSDB = 5ms, Target Reliability = 99%, $f_c = 3.5\text{ GHz}$, FR1, BW = 100MHz, $\mu = 1$, $R = 500\text{m}$, Mobility: Stationary.

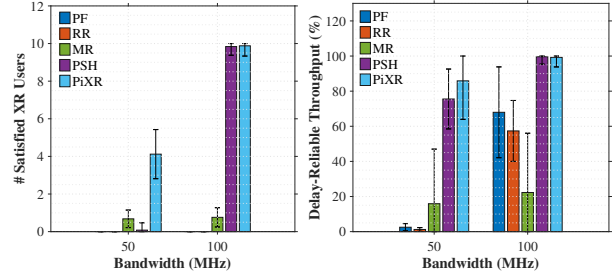


Figure 6. Effect of system bandwidth: $N = 10$, $d_r = 30\text{Mbps}$, $f_r = 60\text{fps}$, PSDB = 10ms, Target Reliability = 99%, $f_c = 3.5\text{ GHz}$, FR1, $\mu = 1$, $R = 500\text{m}$, Mobility: Stationary.

fundamental flaw: uniform resource spreading delivers partial PDU Sets to all users, satisfying none. PiXR satisfies 4 users by entering scarcity mode, sacrificing 6 users to fully serve 4. This validates the binary decision principle: when universal service is infeasible, strategic selection maximizes the satisfied count. At adequate bandwidth (100 MHz), PSH and PiXR perform similarly with PiXR maintaining slight advantage, while traditional schedulers remain ineffective, with at most 1 satisfied user for MR.

4) *Effect of Data Rate and Frame Rate*: Figure 7 evaluates performance under varying traffic intensity with 10 users. The configurations represent light load (30 Mbps, 60 fps) and heavy load (60 Mbps, 120 fps).

At light load, PSH and PiXR perform similarly. With adequate capacity relative to demand, both schedulers operate in normal priority-based modes, explaining their comparable performance. Traditional schedulers fail despite low load, reaffirming their unsuitability for XR traffic. At heavy load, dramatic divergence emerges. PSH collapses to fewer than 1 satisfied user with DRLT degrading to 65%, while PiXR maintains 3-4 satisfied users at 85% DRLT—a 4x improvement. The elevated frame rate combined with doubled data rate pushes the system to high load conditions. PiXR detects scarcity and immediately enters strategic abandonment, concentrating resources on 3-4 targeted users, while PSH fails to do so, thereby supporting no users. These results demonstrate PiXR’s robustness to traffic intensity scaling—critical for supporting emerging high-fidelity XR applications at 120+ fps.

5) *Effect of PDU Set Delay Budget*: Figure 8 examines PSDB sensitivity with 10 XR users at 50 MHz. The bandwidth

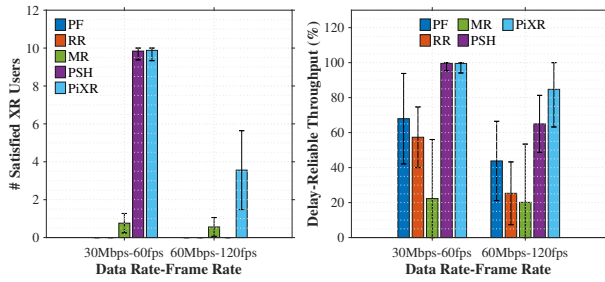


Figure 7. Effect of data rate and frame rate: $N = 10$, $d_r = 30, 60$ Mbps, $f_r = 60, 120$ fps, Target Reliability = 99%, PSDB = 10ms, $f_c = 3.5$ GHz, FR1, BW = 100MHz, $\mu = 1$, $R = 500$ m, Mobility: Stationary.

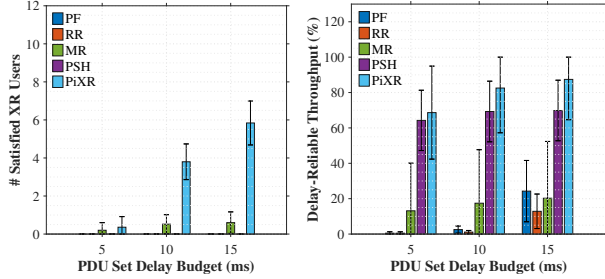


Figure 8. Effect of PSDB: $N = 10$, $d_r = 30$ Mbps, $f_r = 60$ fps, Target Reliability = 99%, $f_c = 3.5$ GHz, FR1, BW = 100MHz, $\mu = 1$, $R = 500$ m, Mobility: Stationary.

is deliberately chosen to create a resource-constrained scenario.

At stringent PSDB (5 ms) with 50 MHz bandwidth and 10 users, all schedulers struggle. When PSDB is relaxed to 10 ms, PiXR’s advantage emerges: 4 satisfied users (83% DRLT) versus ≤ 1 for PSH (70%). The PDTQ mechanism enables effective urgency differentiation through pressure scores, while PSH’s completion-ratio metrics remain ineffective. At further relaxed PSDB (15 ms), PiXR maintains 6 users while PSH still struggles to support even a single user. PSH can only support users when both bandwidth is sufficiently large and PSDB is relaxed; scarcity in either dimension significantly degrades its performance. The persistent gap despite relaxed budgets reveals a fundamental difference: PiXR’s PDTQ structure adapts effectively to the temporal evolution of PDU Sets, whereas PSH’s indirect metrics fail regardless of budget. Additionally, PiXR’s strategic abandonment of resource-inefficient users (essentially implementing implicit admission control by refusing service when infeasible) enables multifold improvements in the number of supported XR users.

V. CONCLUSION AND FUTURE WORK

This paper addressed the challenge of scheduling PDU Set-based XR traffic in 6G networks under stringent delay and reliability constraints. We proposed the PDU Set Delay Tracking Queue (PDTQ) mechanism that maintains real-time visibility into each PDU Set’s remaining delay budget, enabling accurate identification of delay-critical traffic. Building on this structure, we developed PiXR, an adaptive heuristic that operates in dual modes: priority-based allocation under normal conditions and strategic user abandonment under resource scarcity. Extensive simulations demonstrate that PiXR achieves

4-5 times improvement in satisfied user count compared to the state-of-the-art PDU Set-aware heuristic under resource crunch conditions, while maintaining comparable to slightly better performance under adequate resources.

Future work will extend PiXR to incorporate the PDU Set Importance (PSI) for video coding-aware prioritization.

ACKNOWLEDGMENTS

This work was funded by Taighde Éireann – Research Ireland (Grant No. 13/RC/2094_2, 21/FFP-P/10244 (FRADIS)) and co-funded by the European Union (SyMeCo, Grant Agreement No. 101081459). The views expressed are solely those of the author(s). Neither the European Union nor the European Research Executive Agency can be held responsible for them.

REFERENCES

- [1] C.-X. Wang *et al.*, “On the road to 6g: Visions, requirements, key technologies, and testbeds,” *IEEE Communications Surveys Tutorials*, vol. 25, no. 2, pp. 905–974, 2023.
- [2] P. Paymard, S. Paris, A. Amiri, T. E. Kolding, F. S. Moya, and K. I. Pedersen, “Pdu-set scheduling algorithm for xr traffic in multi-service 5g-advanced networks,” in *ICC 2024 - IEEE International Conference on Communications*, 2024, pp. 758–763.
- [3] S. Gül, S. Ahsan, S. Paris, and I. D. D. Curcio, “Pdu set-based qos handling in 3gpp: Release 18 overview and future directions,” *IEEE Communications Standards Magazine*, vol. 9, no. 3, pp. 104–110, 2025.
- [4] M. Gapeyenko, S. Paris, M. Isomaki, B. Yanakiev, A. Amiri, B. Sébire, J. Kaikkonen, C. Wu, and K. I. Pedersen, “Overview of nr enhancements for extended reality (xr) in 3gpp 5g-advanced,” *IEEE Access*, vol. 12, pp. 186 941–186 956, 2024.
- [5] 3GPP, “Service requirements for the 5G system; Stage 1,” 3GPP, Technical Specification TS-22.261, Dec. 2024, release-20.
- [6] 3rd Generation Partnership Project (3GPP), “Study on XR (Extended Reality) Evaluations for NR (Release 17),” 3GPP, 3GPP Technical Report TR 38.838, Dec. 2021, available at: https://www.3gpp.org/ftp/Specs/archive/38_series/38.838/.
- [7] X. Zhao, Y.-J. A. Zhang, M. Wang, X. Chen, and Y. Li, “Online multi-user scheduling for xr transmissions with hard-latency constraint: Performance analysis and practical design,” *IEEE Transactions on Communications*, vol. 72, no. 7, pp. 4055–4071, 2024.
- [8] E. Chen, S. Dou, S. Wang, Y. Cao, and S. Liao, “Frame-level integrated transmission for extended reality over 5g and beyond,” in *2021 IEEE Global Communications Conference (GLOBECOM)*, 2021, pp. 1–6.
- [9] S. Dutta, D. Roy, M. Laha, and G. Das, “Application-aware mac scheduling for xr over epon-based 5g-backhaul,” *Optical Fiber Technology*, vol. 98, p. 104545, 2026. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1068520025004201>
- [10] M. Laha, D. Roy, S. Dutta, and G. Das, “Ai-assisted improved service provisioning for low-latency xr over 5g nr,” *IEEE Networking Letters*, vol. 6, no. 1, pp. 31–35, 2024.
- [11] M. Laha, G. Das, and G.-M. Muntean, “Delay-reliability aware optimal downlink scheduling for extended reality applications in 6g,” in *IEEE INFOCOM 2025 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 2025, pp. 1–6.
- [12] M. Laha, G. Das, and G. Muntean, “Delay-reliability aware downlink scheduling for multi-sensory immersive extended reality in nextg,” *IEEE Journal on Selected Areas in Communications*, pp. 1–1, 2026.
- [13] A. Cedergren, F. de Laval, S. Falahati, L. Falk, D. H. Kang, R. S. Karlsson, Y. Lyazidi, A. Mirzaei, J. Olsson, J. L. Pradas, P. Schliwa-Bertling, N. Shi, B. Singh, R. Tano, and A. M. Voicu, “5g networks evolution for extended reality,” *IEEE Communications Standards Magazine*, vol. 8, no. 3, pp. 54–59, 2024.
- [14] “Traffic models and quality evaluation methods for media and xr services in 5g systems,” in *3GPP TR 26.926 V0.1.0*, Apr. 2021.
- [15] “Study on channel model for frequencies from 0.5 to 100 ghz,” in *(3GPP), TR 38.901, V16.1.0*, Jan.2020.