The Case for Boosting Mobile Application QoE via Smart Band Switching in 5G/xG Networks

Ahmad Hassan
University of Southern California
ahmadhas@usc.edu

Wei Ye
University of Minnesota, Twin Cities
ye000094@umn.edu

Anlan Zhang
University of Southern California
anlanzha@usc.edu

Jason Carpenter
University of Minnesota, Twin Cities
carpe415@umn.edu

Ruiyang Zhu
University of Michigan
ryanzhu@umich.edu

Shuowei Jin
University of Michigan
jinsw@umich.edu

Z. Morley Mao
University of Michigan
zmao@umich.edu

Feng Qian
University of Southern California
fengqian@usc.edu

Zhi-Li Zhang
University of Minnesota, Twin Cities
zhzhang@cs.umn.edu

ABSTRACT

5G and future 6G networks support diverse combinations of access technologies, architectures, and radio frequencies, with each combination termed as a "band" henceforth. Through comprehensive measurements in 12 cities across 5 countries, we experimentally show that operator-configured default bands are often highly suboptimal, particularly under mobility. We then propose *smart band switching*, where a UE's band can be dynamically changed to improve the network performance and boost the application QoE. We discuss challenges, opportunities, and design choices for building a practical smart band switching system. We further develop preliminary UE-side band-switching logic on commodity smartphones, and evaluate it on commercial 5G networks.

CCS CONCEPTS

• **Networks** → **Mobile networks**; *Network measurement*; *Network performance analysis*;

KEYWORDS

5G, xG, Band Switching, Band Selection, Network Measurement, Mobile QoE, Band, Radio Bands

ACM Reference Format:

Ahmad Hassan, Wei Ye, Anlan Zhang, Jason Carpenter, Ruiyang Zhu, Shuowei Jin, Z. Morley Mao, Feng Qian, and Zhi-Li Zhang. 2024. The Case for Boosting Mobile Application QoE via Smart Band Switching in 5G/xG Networks. In *The 25th International Workshop on Mobile Computing Systems and Applications (HOTMOBILE '24), February 28–29, 2024, San Diego, CA, USA.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3638550.3641132

∗ Work done while at University of Minnesota – Twin Cities.

HOTMOBILE '24, February 28–29, 2024, San Diego, CA, USA

ACM ISBN 979-8-4007-0497-0/24/02.
https://doi.org/10.1145/3638550.3641132

1 INTRODUCTION

5G New Radio (NR) supports a wide spectrum of frequency bands, including Low-Band (<1GHz), Mid-Band (1–6Ghz), and High-Band (mmWave, 24–40Ghz) to enable various use cases. These bands are combined with different architectures, such as standalone mode (SA) and non-standalone mode (NSA). In addition, 5G and traditional LTE technologies will coexist for a very long time. This is in sharp contrast to 3G/4G/LTE networks that only employ a limited range of radio frequencies (mostly Low-Band) in a monolithic architecture. These bands will become even more diverse as we move towards next-generation (nextG/xG) cellular networks [5, 9]. Besides, cellular operators often employ carrier aggregation to boost the network bandwidth. In the remainder of this paper, for simplicity, we use the term "band" to refer to a cellular data channel configured with *a combination of technology, frequency, architecture, and carrier aggregation*, and collectively refer to them as *multiband*.

Co-existence and availability of multiband access pose a fundamental research question: *How can one intelligently leverage these diverse bands to best support applications?* In today's networks, when a user equipment (UE) attaches to the network, the base station will configure and select a *"default"* band regardless of application requirements. Does this band provide the best performance for the application? Will any of the bands outperform other bands consistently, especially when the user is on the move? If not, how can one smartly and dynamically switch among the available bands to deliver the best application performance?

Measurements in the Wild. Despite various studies characterizing 5G networks [17, 19], there is, to our knowledge, no work focusing on understanding how band switching impacts 5G performance and the application quality-of-experience (QoE). To bridge this gap, we conduct extensive measurements in 12 cities across 5 countries in North America and Europe using several 5G smartphone models, covering a total travel distance of 520 km+. Our measurements reveal the *wide availability and heterogeneity* of 5G bands in the wild. For 95% of the time during our experiments, a UE can access 5-6 bands encompassing 5G-SA, 5G-NSA, and LTE; the available bands are generally stable spatially and temporally, based on our continuous measurement over 13 months. More surprisingly, we discover the suboptimality of the default band switching

Table 1: **Statistics of the data collected using five commercial 5G operators in the US and Europe.**

# of Locations/Cities/Countries	90+/12/6
# of unique mobility scenarios	6
# of 5G smartphones (and models)	8 (4)
# of *band mapping* tests	300+
Cumulative time of network traces	2156 mins+
Total distance walked/travelled	520 km+

schemes employed by major 5G operators: when a UE is stationary, the median downlink throughput gap between the operator's selected band (referred to as the *default* band) and the throughput-wise best band is 34 mbps, whereas the gap increases to 64 mbps when the UE moves at walking speed.

Smart Band Switching: Challenges and System Design. Motivated by the above measurements, we propose *smart band switching*, where a UE's band can be dynamically changed to improve the network performance and boost the application QoE. We highlight the challenges in realizing a practical band swtiching system, such as UE's incapability of accessing multiple bands simultaneously, non-trivial swtiching overhead, and the gap between band switching and applications' QoE requirements. We also compare in detail three high-level designs: (1) *on-device band switching*, where each UE individually initiates the switching based solely on its local knowledge; (2) *RAN-based switching*, where the Radio Access Network (RAN) dictates the switching through its global knowledge; and (3) a *collaborative scheme* where the RAN and UEs jointly make band switching decisions. Given the pros and cons of the above designs, we take the collaborative scheme as the basic architecture. Leveraging 5G QoS framework [3] and recent innovations in 5G RAN intelligence [2, 8], we further sketch a *QoE policy-based mechanism* to enable band switching. A policy-based approach will offer flexibility to define fine-grained QoE goals and handle conflict mitigation with operator-configured (*default*) band switching policies.

Preliminary Proof-of-Concept (PoC) and Real-world Evaluations. To demonstrate the potential benefits of our proposal, we develop preliminary UE-side band-switching logic on commodity smartphones (and leave the RAN-side coordination as future work). Our PoC implementation strategically balances the critical trade-off between exploration (switching to a new band) and exploitation (staying on the current band). The band performance prediction is achieved through a forecast engine that employs historical measurements to predict each band's performance. We then evaluate our PoC on commercial 5G networks. Compared to using the default band, smart band switching boosts application QoE (Video On-Demand Streaming and HTTP File Download) by 19.8% to 37% in live 5G networks. Under diverse mobility scenarios, it improves the throughput by 30% to 190% (average 99%) (§4.2). We envision that by further incorporating the RAN-side knowledge, the benefits of smart band switching would be even more significant.

To summarize, this positioning paper makes three contributions. First, motivated by our field measurements, we propose smart band switching as a new dimension for application performance optimization. Second, we outline key challenges, design choices, and research directions for developing a practical band switching system. Third, our PoC implementation and evaluations on commercial UEs and 5G networks demonstrate the potential of our approach.

2 MOTIVATING BAND SWITCHING THROUGH AN EXPLORATORY STUDY

We conduct extensive experiments in the wild to: (i) verify the spatial and temporal stability of bands (§2.2), and (ii) measure and quantify band switching gains (§2.3).

2.1 Measurement Setup & Tools

Location, Operator, Band, and Technology. We conducted our measurement study in 12/5 cities/countries across multiple regions in Europe and the US. Our collected dataset spans *six* mobility scenarios, including stationary, walking, driving, light rail, public bus, and indoors. We selected a major commercial cellular operator (T-Mobile) in the US for our experiments. To the best of our knowledge, T-Mobile supports the most diverse bands among all US cellular operators. Specifically, T-Mobile deployed their cellular services using 4G/LTE, NSA-5G, and SA-5G in the Low-Band and Mid-Band radio frequency spectrum (600–2500 MHz). To collect data in Europe, we used countries' local cellular operators (Vodafone, Telekom, SFR, and Orange). European operators also support 4G/LTE and NSA-5G in Low-Band and Mid-Band range (800–3300 MHz); however, only Vodafone offers SA-5G. Note that our insights will also be valid for bands not explored in this study, *e.g.,* mmWave. Key statistics of the collected dataset are summarized in Table 1.

Measurement Tool. We developed an Android application to characterize the network performance of bands. Using this application, we saturate the uplink and downlink channels of the device with UDP packets. The uplink/downlink sending rate is set to 100/700 Mbps (*i.e.,* slightly higher than Mid-Band's capacity). We used a university-hosted server with 4Gbps+ network bandwidth. Hence, the Internet was not a bottleneck. Additionally, we capture other key information such as geolocation, mobility speed, 4G/5G signal strength information, *etc.*, using standard Android APIs.

Data Collection. We used multiple smartphone models to minimize the impact of smartphones' diversity: Samsung Galaxy S10 (S10), S20 Ultra (S20U), S21 Ultra (S21U), and S22+ (S22+). These phones have diverse radio capabilities. We conducted two types of experiments in our study: *band mapping* (referred to as D1 afterward) and *performance characterization* (D2). To *map* the footprint of multiple bands, we carried out (i) stationary experiments at 90+ distinct locations[1] across the US and Europe (D1.1), (ii) we set out on a measurement campaign spanning more than 13 months (D1.2). Specifically, we conducted experiments at four fixed locations (university campus, suburban residential area, downtown plaza, and airport) in a large US city. Using ADB scripts, we conducted experiments by cycling through all 4G and 5G band settings. We repeated all the stationary experiments 3x at each location. For D2, we configured the five most frequently seen bands (including the default band setting), in D1, on five S21U phones given S21U's wide band coverage. To simultaneously characterize the performance of bands at any location, we placed all five devices (one for each band setting) side-by-side and collected data concurrently. We also performed benchmarking experiments to confirm that (i) band settings on phones do not interfere, and (ii) the same band setting on different phones offers similar performance.

[1]We conducted experiments at diverse locations (*e.g.,* airports, university campuses, tourist attractions, *etc.*), not in close proximity to each other.

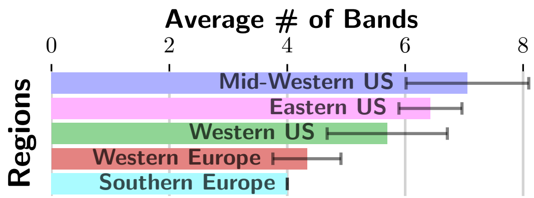

Figure 1: Availability of bands across different regions of the US and Europe.

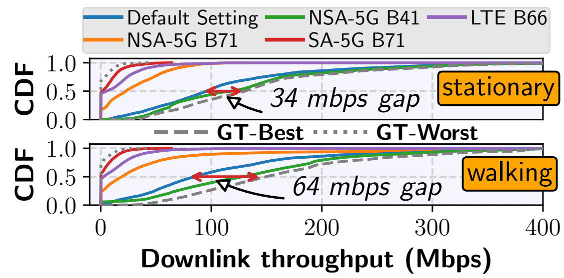

Figure 2: Network throughput comparison across most frequently detected bands.

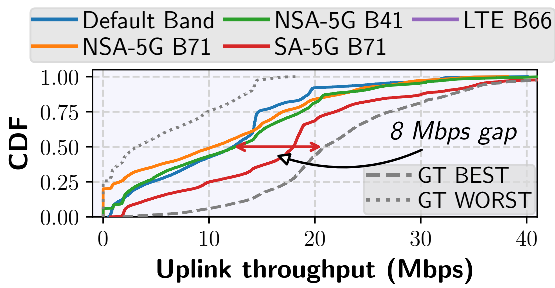

Figure 3: Network throughput comparison across most frequently detected bands for the *walking* user.

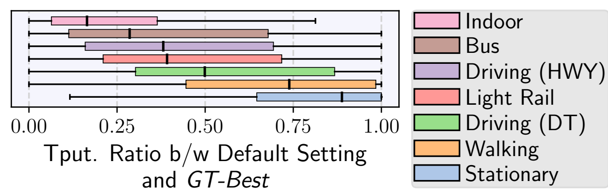

Figure 4: Quantifying the downlink throughput gap for various mobility scenarios.

2.2 Spatial and Temporal Stability of Bands

We use D1.1 and D1.2 datasets to study the spatial and temporal stability of bands, respectively, which are important prerequisites for realizing multiband gains. We make two observations. **(i)** The multiband access is spatially stable, and the bands are abundant. Fig. 1 plots the number of available bands in different regions across the US and Europe. We find that, at any location, more than 3 bands can be accessed ~94% of the time. Additionally, mobile devices can leverage 3-6 bands in the median case for different locations, and the number can be as large as 8 at some locations. Our analysis points towards a higher availability of bands across the US as compared to Europe. The disparity is caused by the lower availability of SA-5G and NSA-5G Low-Band in Europe: only one European operator has deployed SA-5G, while NSA-5G Low-Band is not found in our data. On the other hand, operators in the US extensively use NSA-5G Low-Band to offer 5G services [17]. **(ii)** After conducting experiments at regular intervals across a year, we also verify the temporal stability of bands. The set of bands always remained the same except when T-Mobile added a new SA-5G band at two of the four locations. At all locations, we observed 5-6 bands ~95% of the time over 13 months.

2.3 Potential Gains from Band Switching

We utilize D2 to investigate three questions: **(i)** Is there band heterogeneity in the wild? **(ii)** If yes, how much performance gain can multiband access yield? **(iii)** What are the different factors that dictate the level of performance gain?

To answer **(i)** and **(ii)**, we use a subset of our walking data where we walk around a 1.4 kms rectangular loop on the university campus, covering an area of ~0.1 km^2. We also perform stationary experiments at multiple locations across the loop to fully understand the impact of mobility on bands' performance. Fig. 2 compares the downlink throughput for bands across different technologies and architectures. We plot the CDF of throughput achieved by all the bands. The *GT-Best* and *GT-Worst* lines are the highest and lowest throughput experienced by all bands at any given time, respectively. Each label represents the primary frequency band, and secondary (aggregated) carriers are ignored for brevity.

There are two key takeaways from the results: (i) The default band switching strategies employed by the operators are suboptimal, leaving considerable room for improvement. Fig. 2 shows that the median gap between the throughput of the default band and *GT-Best* band is 34 mbps when the UE is stationary. To compare, the median downlink throughput for the default band is 95 mbps and 84 mbps for stationary and walking, respectively. The gap increases to 64 mbps in the walking case. An *optimal* band switching policy can effectively boost network throughput (or improve application QoE) if the mobile device can dynamically switch bands based on their performance. (ii) Uplink and downlink do not necessarily share the same *optimal* band (see Fig. 2 and 3 for comparison). For example, NSA-5G B41 offers the highest downlink throughput in our walking experiments; however, its uplink performance is considerably worse than SA-5G B71. The *optimal* bands are different because operators typically prioritize the downlink channel [14], use different antenna technologies for uplink *vs.* downlink[12], and employ low-power-consuming bands for uplink transmission [15].

We use the full might of dataset D2 to answer the third question. Our analysis suggests that *mobility speed* is critical in determining the performance gains during band switching. For our walking loop, NSA-5G B41 almost always (98% of the time) offers the highest throughput when the user is static (Fig. 2). However, NSA-5G B41 has the highest throughput 87% of the time when the UE moves at walking speed. We also characterize the performance gap for all mobility scenarios in D2. Fig. 4 illustrates that the throughput gap between the *GT-Best* band and the default band setting increases for complex mobility scenarios, primarily due to rapid channel fluctuations, blockages, number of users in the vicinity, *etc.* This ultimately means that multiband access can be particularly useful for complex mobility scenarios.

The performance gains depend on the *application type* as well. Applications with higher band switching tolerance see significant QoE improvement, *e.g.,* the median HTTP file download time reduces by 25% to 37% in our real-world experiments (§4.2). In contrast, applications with lower switching tolerance show comparatively less improvement, *e.g.,* real-time point cloud video streaming sees an average 5.7% and 6.2% improvement in video quality and stall, respectively. The *band switching overhead* is another factor impacting the performance gain. The higher the interruption is, the lower the QoE improvement will be. From the control plane perspective, band switching takes ~147 ms in the median case. Although band switching usually does not break the network (TCP/UDP) connection, it causes a median (95th percentile) interruption of ~412 ms (604 ms) on the data plane, during which the UE cannot transmit/receive data. The data plane interruption is higher than the control-plane switching delay because applications typically require more time (100s of ms) to recover after the radio connection is reestablished, as shown by a previous study on cellular handovers [21]. Lastly, the *phone model* determines the performance gain, too, as devices with old cellular modems cannot support some newly deployed bands. S10, being a relatively old model, supported on an average 2-3 fewer bands than the newest S22+.

3 SMART BAND SWITCHING

Our findings in the last section highlight the feasibility and potential gains of multiband access. However, it is still unclear where exactly inside the radio network *smart band switching* can be implemented. Table 2 enlists three design alternatives – namely *on-device*, *RAN-based*, and *collaborative* – each with its own set of challenges and opportunities. Here, we discuss the challenges, opportunities, and design alternatives for a practical smart band switching system.

3.1 Challenges and Opportunities

There are several challenges in realizing a practical band switching system. *Single-band-access:* To begin, today's smartphones cannot access all available bands concurrently and probing each band sequentially yields potentially out-of-date measurements. *Switching overhead:* A device-triggered band switch incurs a median switching delay of 417 ms during which no data transfer can be performed. *Volatility of available bands:* Bands come and go as users move from one place to another. *Dynamics of individual band's performance:* A band's performance itself is dynamic, especially during mobility. *QoE-driven band performance prediction:* Predicting the performance of such dynamic and diverse bands to optimize QoE is challenging yet interesting. *Interference with network policies:* A greedy band switching scheme can lead to resource unfairness and spectral inefficiency in a shared radio resource environment. *Knowledge of QoE demands:* Smart band switching requires an accurate characterization of UE's application QoE requirements which can be challenging to obtain unless explicitly specified by the app.

Research Directions. Regardless of what design approach we choose, there are several questions that need careful investigation. First, assuming that we know the high-level application requirements (*e.g.,* bandwidth hungry *vs.* latency sensitive), how can we translate these requirements into QoE policies that are meaningful to RAN and the core network? Second, how should we combine QoE policies with operator-configured policies to co-optimize for application QoE, resource fairness, spectral efficiency, *etc.*? Third, how to enforce all these policies in real-time while making sure that the applications get a guaranteed QoE? Finally, how to detect the band with the highest QoE, as bands' performance depend on multiple factors outlined above (§2.3)?

Table 2: Comparison of different design alternatives for band switching. The ● and ○ marks qualitatively indicate high and low values, respectively.

Approach	Application Agnostic	Switching Overhead	Deployment Cost	Resource Unfairness
On-device	○	●	○	●
RAN-based	●	○	●	○
Collaborative	○	○	●	○

3.2 Solution Space

On-device Band Switching. To begin, one might argue that UEs, such as mobile devices, are in the best position to make smart band switching decisions as they can leverage *cross-layer* information as well as *user contexts* (*e.g.,* user mobility patterns). However, the device-initiated band switch will have a higher switching overhead than the network-initiated one. This is because a device-initiated band switch involves an additional step of letting the RAN know which band to use for the device. Additionally, a device-based approach cannot account for other users in the vicinity, which can lead to radio resource unfairness. It can also adversely impact spectral efficiency since UEs will only favor the high-performing bands. Nonetheless, an on-device solution has the lowest deployment cost among all possible options (see §4).

RAN-based Band Switching. To address the issues of an on-device band switching scheme, the base station can implement band switching to optimize network performance for all connected UEs. In doing so, the RAN can co-optimize for resource fairness, spectral efficiency, and UE's network performance. However, such a design will be oblivious to the QoE demands of individual UEs and can lead to the same performance issues as existing band switching strategies.

Collaborative Band Switching. A better approach is to facilitate collaboration among devices and RAN where the network utilizes the collective information (*e.g.,* application demands and user mobility patterns) from all connected UEs to make an optimal global decision. Although 5G offers a flow-level QoS framework [3], it is simply employed for resource allocation (*e.g.,* in MAC scheduling and/or network slicing) and not for deciding the actual band a UE will use. To this end, one can introduce *QoE policies* inside the network to enable smart band switching. We argue that policy-based collaborative band switching can explicitly boost mobile QoE while also ensuring fairness among users and low switching overhead. Moreover, a policy-based approach will offer flexibility to define fine-grained QoE goals and handle conflict mitigation with operator-configured band switching policies.

3.3 Solution Sketch

One way to realize a collaborative solution is by combining the 5G flow-level QoS framework with fine-grained radio network control to enable smart band switching. Mobile applications can utilize 5G QoS to communicate their QoE requirements which can be formulated into *QoE policies*. The legacy cellular architecture already supports policy-based rules for service data flow detection (*e.g.,* voice service *vs.* data traffic), policy enforcement, and flow-based charging [4]. To make sure the solution is 3GPP-compliant, the existing policy-based QoS framework (*e.g.,* 5G's Policy Control Function [1]) can be extended for smart band switching.

Given a QoE policy, a band switching system can be designed on top of recent innovations in 5G radio network control. O-RAN's RAN Intelligent Controller (RIC) [8] and MEC's Radio Network Information Service (RNIS) [2] add programmability to the RAN, providing a foundation to build intelligent systems on. For instance, we can enable policy orchestration and enforcement via OpenRAN's RIC-based network automation tools, such as xApps and rApps [6, 7]. To control the radio network, the RIC comes in two forms: near-real-time (Near-RT) – from 10 ms to 1 sec – and non-real-time (Non-RT) – more than 1 sec. Non-RT RIC takes a helicopter view of the network and can be used for policy orchestration via rApps. On the other hand, Near-RT RIC can employ xApps for policy enforcement to handle UE's fine-grained QoE demands.

4 POC AND EVALUATION

To understand the benefit of multiband access for typical mobile apps and avoid the huge cost of setting up a radio network, we develop preliminary UE-side band-switching logic to roughly estimate band switching gains. This section briefly overviews our proof-of-concept (PoC) solution, followed by its evaluation on commercial 5G networks.

4.1 System Overview

Our PoC solution has three building blocks: *Decision Framework*, *Forecast Engine*, and *Band Switcher*. The *Decision Framework* systematically explores new bands to capitalize on multiband access and exploits the predicted *best* band to improve performance over the default band setting. The exploration is context-aware, depending on contextual information (*e.g.,* application type and mobility level). The band performance prediction is achieved through *Forecast Engine*. It passively obtains network measurements and predicts long-term performance with a lightweight filter-based approach. Finally, the *Band Switcher* executes the band switch.

Decision Framework can decide to use the same band or switch to another one once the band performance predictions are received from the *Forecast Engine*. The switching decision depends on the state PoC is in. Recall from §2.1 that today's mobile devices can only access a single band at any time. On one hand, band switching is indispensable for exploring new bands; on the other hand, it incurs data interruption that can reduce performance gains. We formulate two states to tackle this exploration-exploitation dilemma: *exploit* and *explore*. In the *exploit* state, the system simply keeps using the predicted *best* band while measuring its performance. In the *explore* state, we measure the performance of potentially suboptimal (≠ *best*) bands that have not been used for a long time. PoC systematically explores all bands to minimize the gap between *GT-Best* band and active band.

Forecast Engine measures and predicts the network throughput of all available bands. At each time step, we combine all the samples (*i.e.,* bytes sent/received at the cellular interface) to get measured throughput E_b^t for band b. PoC uses E_b^t to predict band b's throughput μ_b^{t+1} for the next time step. We design a lightweight filter-based approach, *i.e.,* Dual-filter Exponential Weighted Moving Average (D-EWMA), for the throughput prediction task. Since the proper value of the α parameter in vanilla EWMA depends on the throughput fluctuation level, we apply another EWMA to smooth out the throughput variation. D-EWMA is lightweight, can effectively predict long-term performance, tolerate measurement noises, and take into account the throughput variation.

Implementation. For evaluation, we implement PoC on commodity Android smartphones. PoC runs as a background service, and uses `TrafficStats`, `Location`, and `Activity Recognition` APIs [10] to monitor the network traffic, UE's moving speed, and currently running app. We utilize Android's `TelephonyManager` API [11] and Samsung's special access code (*#2263#) to execute a band switch.

4.2 Quantifying Band Switching QoE Gains

We consider three mobile applications: File Download, VoD Streaming, and Point Cloud Streaming. We use two S21U devices, one running PoC and the other default band. This setup allows us to compare the two settings side-by-side. Overall, we collect 5+ hrs of data for each application.

HTTP File Download. We repeatedly download a 256 MB file from the server on both phones. Fig. 5(a) illustrates the summary of our results. Compared to the default band setting, PoC results in 37% and 25% lower median file download time for driving and walking, respectively. Recall from §2.3 that relatively more multiband gains are available for complex mobility scenarios such as driving. Therefore, PoC offers higher improvement in such scenarios. We also plot a representative trace of PoC's performance in Fig. 5(b) and compare it with the default band setting. During the 5 min timeline, PoC spends 4% (12 secs) of the time in the *explore* state, triggers 9 band switches, and downloads 67% more data than the default band setting.

VoD Streaming. Our VoD streaming experiments use a Dash.js [13] player to download and play the video. We test buffer-based (BOLA [20]) and rate-based [18] adaptive bitrate (ABR) algorithms due to their popularity. Fig. 6 shows that, for buffer-based ABR, band switching improves the average video bitrate by 20.1% and reduces absolute stall percentage by 0.4% and 6.4% for walking and driving, respectively, compared to using the default band. Likewise, PoC improves video bitrate by 19.8% and decreases stall by 1.1% to 7.2% for the rate-based ABR.

Real-time Point Cloud Video Streaming. We replicate ViVo [16], a (3D) point cloud streaming system, and adapt it to a live system to demonstrate PoC's ability to improve the real-time application's performance. Fig. 7 illustrates the quality level (corresponding to 5 point-cloud density levels) and frame delay of the point cloud stream. As shown, PoC has 5.7% higher average quality and 6.2% lower average frame delay than the default band setting.

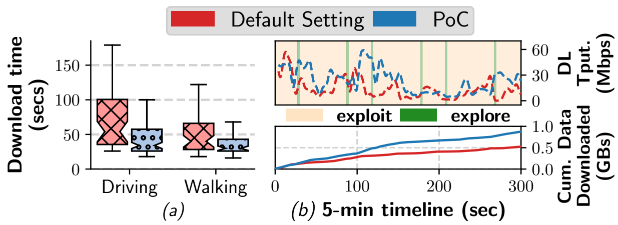

Figure 5: Performance comparison b/w PoC and the default band setting for HTTP File Download.

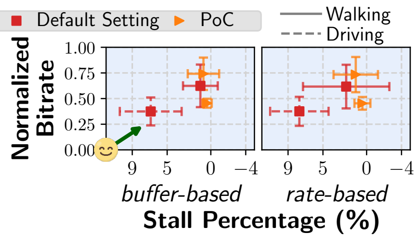

Figure 6: Performance comparison b/w PoC and the default band setting for VoD Streaming.

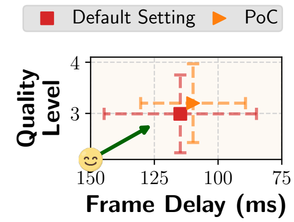

Figure 7: Performance comparison for Point Cloud Streaming.

Performance under Mobility. To explain how band switching behaves under different mobility conditions, we run TCP Bulk Transfer (iPerf3) for all six mobility scenarios separately. Our analysis shows that band switching achieves 30% (*walking*) to 190% (*indoors*) higher throughput depending on the mobility level (average 99% across all scenarios).

Applicability of Evaluation Results. PoC incurs a high switching overhead and provides a lower bound of performance gains offered by band switching. We believe that a *collaborative* approach will result in an even higher QoE boost while addressing all the challenges highlighted in §3.1.

5 LIMITATIONS AND FUTURE WORK

First, since we are not collaborating with cellular operators, our study cannot reveal the impact of band switching on RAN, such as signaling load, resource fairness, and spectral efficiency. Second, the PoC uses throughput as the key performance metric when making band switching decisions; other metrics (*e.g.,* latency and energy efficiency) can be incorporated into our future work. Lastly, evaluating the performance of all design alternatives in §3 will render valuable insights for the community and mobile stakeholders.

ACKNOWLEDGMENTS

We thank anonymous reviewers for their suggestions and feedback. This research was in part supported by NSF under Grants CNS-1915122, CMMI-2038559, CNS-2154078, CNS-2128489, CNS-2220286, CNS-2220292 and CNS-2321531, and a Cisco University Research Grant.

REFERENCES

[1] 2019. *ETSI 123 503: Policy and charging control framework for the 5G System (5GS).* https://www.etsi.org/deliver/etsi_ts/123500_123599/123503/15.06.00_60/ts_123503v150600p.pdf

[2] 2019. *Multi-access Edge Computing (MEC); Radio Network Information API.* https://www.etsi.org/deliver/etsi_gs/MEC/001_099/012/02.01.01_60/gs_mec012v020101p.pdf

[3] 2020. *ETSI 123 501: System architecture for the 5G System (5GS).* https://www.etsi.org/deliver/etsi_ts/123500_123599/123501/16.06.00_60/ts_123501v160600p.pdf

[4] 2021. *ETSI 129 513: Policy and Charging Control signalling flows and QoS parameter mapping.* https://www.etsi.org/deliver/etsi_ts/129500_129599/129513/16.06.00_60/ts_129513v160600p.pdf

[5] 2021. *FCC Seeks Comment on Maximizing Efficient Use of 12 GHz Band.* https://www.fcc.gov/document/fcc-seeks-comment-maximizing-efficient-use-12-ghz-band-0

[6] 2021. *Near-RT RAN Intelligent Controller.* https://docs.o-ran-sc.org/en/cherry/projects.html#near-realtime-ran-intelligent-controller-ric

[7] 2021. *Non-RT RAN Intelligent Controller.* https://docs.o-ran-sc.org/en/cherry/projects.html#non-realtime-ran-intelligent-controller-nonrtric

[8] 2021. *O-RAN ALLIANCE Introduces 48 New Specifications.* https://www.o-ran.org/blog/o-ran-alliance-introduces-48-new-specifications-released-since-july-2021

[9] 2022. *FCC, NTIA Establish Spectrum Coordination Initiative.* https://www.ntia.gov/press-release/2022/fcc-ntia-establish-spectrum-coordination-initiative

[10] 2023. *Android API reference.* https://developer.android.com/reference

[11] 2023. *Android TelephonyManager.* https://developer.android.com/reference/android/telephony/TelephonyManager

[12] 2023. *Compar. SISO and MIMO Hardware.* https://arednmesh.readthedocs.io/en/latest/arednHow-toGuides/siso-mimo.html

[13] 2023. *DASH.js framework.* https://github.com/Dash-Industry-Forum/dash.js/

[14] 2023. *Fast Internet Download Speeds Aren't the Only Thing That Matters. Here's Why.* https://www.cnet.com/home/internet/fast-internet-download-speeds-arent-the-only-thing-that-matters-heres-why/

[15] Moinak Ghoshal, Z. Jonny Kong, Qiang Xu, Zixiao Lu, Shivang Aggarwal, Imran Khan, Yuanjie Li, Y. Charlie Hu, and Dimitrios Koutsonikolas. 2022. An In-Depth Study of Uplink Performance of 5G MmWave Networks. In *Proceedings of the ACM SIGCOMM Workshop on 5G and Beyond Network Measurements, Modeling, and Use Cases (5G-MeMU '22).* Association for Computing Machinery, New York, NY, USA, 29–35. https://doi.org/10.1145/3538394.3546042

[16] Bo Han, Yu Liu, and Feng Qian. 2020. ViVo: Visibility-Aware Mobile Volumetric Video Streaming. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking (MobiCom '20).* Association for Computing Machinery, New York, NY, USA, Article 11, 13 pages. https://doi.org/10.1145/3372224.3380888

[17] Ahmad Hassan, Arvind Narayanan, Anlan Zhang, Wei Ye, Ruiyang Zhu, Shuowei Jin, Jason Carpenter, Z. Morley Mao, Feng Qian, and Zhi-Li Zhang. 2022. Vivisecting Mobility Management in 5G Cellular Networks. In *Proceedings of the ACM SIGCOMM 2022 Conference (SIGCOMM '22).* Association for Computing Machinery, New York, NY, USA, 86–100. https://doi.org/10.1145/3544216.3544217

[18] Zhi Li, Xiaoqing Zhu, Joshua Gahm, Rong Pan, Hao Hu, Ali C. Begen, and David Oran. 2014. Probe and Adapt: Rate Adaptation for HTTP Video Streaming At Scale. *IEEE Journal on Selected Areas in Communications* 32, 4 (2014), 719–733. https://doi.org/10.1109/JSAC.2014.140405

[19] Arvind Narayanan, Xumiao Zhang, Ruiyang Zhu, Ahmad Hassan, Shuowei Jin, Xiao Zhu, Xiaoxuan Zhang, Denis Rybkin, Zhengxuan Yang, Zhuoqing Morley Mao, Feng Qian, and Zhi-Li Zhang. 2021. A Variegated Look at 5G in the Wild: Performance, Power, and QoE Implications. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference (SIGCOMM '21).* Association for Computing Machinery, New York, NY, USA, 610–625. https://doi.org/10.1145/3452296.3472923

[20] Kevin Spiteri, Rahul Urgaonkar, and Ramesh K. Sitaraman. 2020. BOLA: Near-Optimal Bitrate Adaptation for Online Videos. *IEEE/ACM Trans. Netw.* 28, 4 (aug 2020), 1698–1711. https://doi.org/10.1109/TNET.2020.2996964

[21] Shichang Xu, Ashkan Nikravesh, and Z. Morley Mao. 2019. Leveraging Context-Triggered Measurements to Characterize LTE Handover Performance. In *Passive and Active Measurement*, David Choffnes and Marinho Barcellos (Eds.). Springer International Publishing, Cham, 3–17.