A Large-Scale Study of the Potential of Multi-Carrier Access in the 5G Era

Fukun Chen[1⋆], Moinak Ghoshal[1*], Enfu Nan[1], Phuc Dinh[1], Imran Khan[1], Z. Jonny Kong[2], Y. Charlie Hu[2], and Dimitrios Koutsonikolas[1]

[1] Northeastern University, USA
{chen.fuk,ghoshal.m,nan.e,dinh.p,khan.i,d.koutsonikolas}@northeastern.edu
[2] Purdue University, USA {kong102,ychu}@purdue.edu

Abstract. Despite much promise, numerous recent studies have shown that 5G coverage remains sporadic and its performance is often suboptimal, leading to degraded QoE for high-throughput, low-latency applications. While two alternative approaches to multi-carrier access, link selection and link aggregation, have been proposed and shown to potentially enhance performance, their actual performance benefit in real-world 5G deployments remains unclear. In this paper, we conduct an extensive measurement campaign involving multiple cross-country driving trips spread out over 12 months and covering a total of 8k+ Km, while simultaneously measuring the performance of the three major US mobile operators, and explore the potential of multi-carrier access to improve throughput and latency. Our measurements show that there exists a substantial amount of diversity in cellular network performance across different operators at a given time and location in terms of both throughput and RTT. Our trace-driven analysis shows that multi-carrier access techniques have the potential to provide significant performance gains over the single-path throughput/RTT of the worst-performing operator.

1 Introduction

In recent years, an emerging class of mobile applications, including Augmented Reality (AR), Virtual Reality (VR), Mixed Reality (MR), Connected Autonomous Vehicles (CAVs), and the Metaverse, has gained significant traction. These applications demand exceptionally high bandwidth and low latency from underlying wireless networks. 5G NR, rolled out in 2019, holds the promise to meet these demands unlike its predecessor (LTE), through advancements such as higher-order modulation, beamforming, massive MIMO, and wider channels.

However, recent measurement studies [14,15,29,39,11] have shown that 5G coverage remains sporadic and performance often falls short of expectations, resulting in a degraded Quality of Experience (QoE) for these demanding applications. A promising solution to enhance performance is leveraging multi-carrier access in mobile networks. Broadly, this can be approached in one of two ways: (a) selecting the optimal link based on real-time performance [35,20,5],

⋆ Fukun Chen and Moinak Ghoshal are co-primary authors.

or (b) aggregating multiple links (e.g., via a multipath transport protocol such as MPTCP [24] or MPQUIC [23]) to achieve higher overall throughput. Either approach comes with its own trade-offs. For example, while link aggregation can boost throughput, it may underperform compared to the single-path performance, especially in cases of very heterogeneous links [19,8]. Conversely, selecting a single best link offers potential performance gains in such scenarios but introduces overhead related to frequent link switching in dynamic network conditions [35,20], which may impact application performance, while its maximum possible gain is theoretically lower than that of a multipath solution. Which approach achieves higher performance depends on the diversity and dynamics of the per-link performance in real-world deployments.

Previous research efforts [31,9,20,42,35,41] have primarily focused on designing systems that leverage link aggregation or link selection to improve application performance. However, these approaches have been evaluated in limited settings only. To date, there has been no large-scale real-world study measuring the performance of different carriers, quantifying their diversity in coverage and performance, and analyzing the potential of multi-carrier access approaches in real-world cellular deployments.

To bridge this gap, in this paper, we conduct an extensive measurement campaign involving three cross-country driving trips (Los Angeles to Boston, Boston to Chicago, and Boston to Atlanta) covering a total of 8k+ Km spread out over a 12-month period (August 2022 to August 2023), while *simultaneously measuring the performance of the three major US mobile operators*. The goal of our study is twofold: (a) to analyze the operator diversity in cellular network performance at a given location and time and (b) to gain critical insights into the potential gains of the two multi-carrier access approaches – link selection and link aggregation – in real-world environments. *Our dataset is publicly available at [1].* The main findings of our study are as follows:

• Our dataset shows a substantial amount of diversity in cellular network performance across different operators at a given time and location in terms of both throughput and RTT, suggesting that the overall network performance can significantly benefit from multi-carrier access solutions.

• Our trace-driven analysis of the theoretical gains of the two multi-carrier access techniques – link selection and link aggregation – assuming optimal link selection/multipath scheduling algorithms, shows that the multi-carrier access schemes provide a median relative throughput gain of 15-309% (DL: 15-111 Mbps, UL: 1-9 Mbps) and a median RTT improvement of 8-20% (5-12 ms) compared to the single-path throughput/RTT of the worst-performing operator.

• The best-operator coherence time – the time during which one operator continuously outperforms another operator – is in the order of several tens of s in the median case, suggesting that link selection techniques can be effective in real-world scenarios without suffering from high switching overhead.

• The probability of handovers to happen at the same time for an operator pair is nearly 0, suggesting that link aggregation can effectively mask handover-induced communication disruptions in real-world cellular network deployments.

2 Related work

5G measurements. Since the rollout of 5G in 2019, a large number of studies [26,25,28,15,13,37,7,27,14,10,12,39,17,40,33,36,11,22] have measured various aspects of 5G performance including coverage, network performance, carrier aggregation, resource allocation, handovers, beamforming, and application QoE. To the best of our knowledge, there is no other large-scale measurement work studying the performance diversity among different cellular operators and evaluating the benefits and trade-offs of multi-carrier access in 5G networks.

Multi-carrier access in cellular networks. Most works with respect to multi-carrier access in cellular networks focus on systems apsects [18,42,31,9,41,30,20,35]. These works propose protocols for link selection [20,35] or multipath schedulers [18,42,31,9,41,30] for MPTCP or MPQUIC to improve throughput [30], latency [18], or application QoE [42,21,31,9,41,30], and most of them have performed evaluation in limited settings. These works are orthogonal to ours, which focuses on quantifying the diversity in cellular network performance and the potential benefits of multi-carrier access techniques over cellular networks in the 5G era. The work in [38] evaluates the performance of several MPTCP schedulers over one 5G and one LTE or WiFi link using simulations while the work in [19] evaluates the performance of MPTCP over LTE networks in high-speed trains. The works closest to ours are [2,3,34], which similar to this work, evaluate the potential gains of multi-carrier access, but *in the pre-5G era*. Additionally, these works focus on critical intelligent transportation services (ITS), evaluating network availability and transfer time of critical ITS messages, while our work evaluates throughput and RTT, two metrics that are applicable to a variety of applications. To our best knowledge, our work is the first to empirically evaluate the potential of both multi-carrier access techniques over real-world 5G deployments, featuring a mix of 5G and LTE links, via a large-scale measurement campaign comprising multiple cross-country drives.

3 Methodology

We conducted experiments over the three major U.S. mobile operators: Verizon, T-Mobile, and AT&T. We used Samsung Galaxy S21 Ultra smartphones as User Equipment (UE) and utilized AWS EC2 to measure downlink/uplink (DL/UL) throughput and ICMP-based PING to measure round-trip time (RTT). All our experiments were conducted with 3 UEs, each connected to a different operator, to *concurrently* measure the performance over the three operators. More details about the methodology can be found in §A.

4 Results

4.1 Coverage and single-operator performance

We begin by taking a look at the overall cellular technology coverage and the corresponding network performance for all three operators. Fig. 1a shows the cellular technology coverage for all three operators. 5 years after 5G rollout, we observe a variagated distribution of different cellular technologies among the three operators. While T-Mobile exhibits an impressive 5G coverage ($\geq 97\%$),

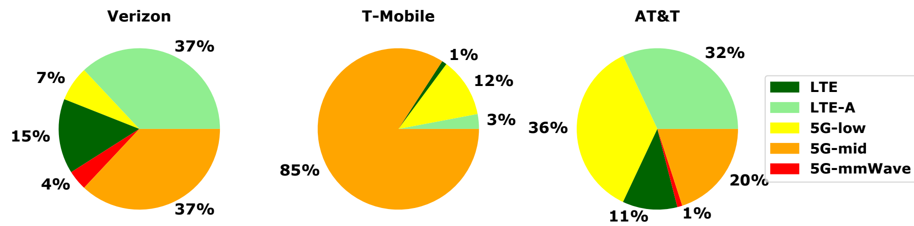

(a) Overall cellular coverage.

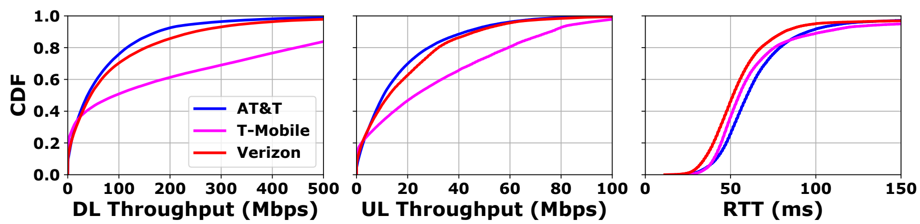

(b) DL throughput, UL throughput and RTT.

Fig. 1: Operator-wise coverage and performance comparison.

for Verizon and AT&T, the overall 5G coverage is much lower, 48% and 57%, respectively. Further breakdown of the 5G bands into low-speed 5G (5G-low) and high-speed 5G (5G-mid, and 5G-mmWave), shows that for T-Mobile and Verizon, the majority of the 5G coverage is contributed by high-speed 5G: 85% and 41%. On the other hand, for AT&T, 5G-low dominates the overall 5G coverage (36%).

Consequently, in Fig. 1b, we observe that T-Mobile provides the best performance in terms of DL and UL throughput in the median and the 75-th percentile. On the other hand, Verizon provides the lowest RTT among the three operators. One reason for this is the use of the Amazon Wavelength edge servers, which are accessible only via the Verizon network, in select cities. Additionally, operators are often times known to not elevate the cellular technology from LTE to 5G in the absence of heavy traffic [14], which is a possible reason for the T-Mobile to not achieve the best RTT even though it has the highest 5G coverage.

Overall, *we observe large disparities among the three operators in terms of both coverage and performance (throughput and RTT). These observations motivate the need for multi-carrier solutions to enhance mobile network performance.* In subsequent sections, we take a deeper look in quantifying the disparity in performance across operators at a given time and location, and theoretically analyze the efficacy of multi-carrier solutions.

4.2 Operator diversity in cellular network performance

In this section, we take a deeper look at the diversity in cellular network performance at a given location and time. We consider only the throughput (RTT) samples collected concurrently for any pair of operators, and calculate the throughput (RTT) difference for each operator pair.

Throughput diversity. Fig. 2a shows the CDF of the throughput difference at a given location and time for each pair of operators. We observe substantial

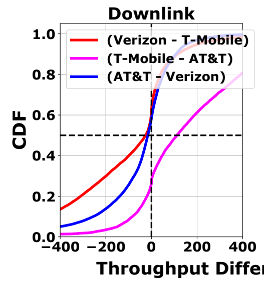

(a) CDF of throughput difference for different operator pairs.

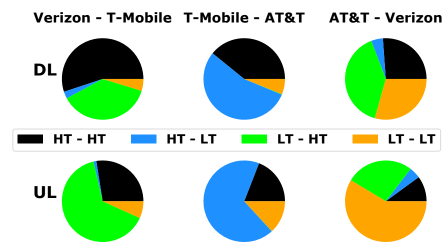

(b) Technology-breakdown for different operator pairs.

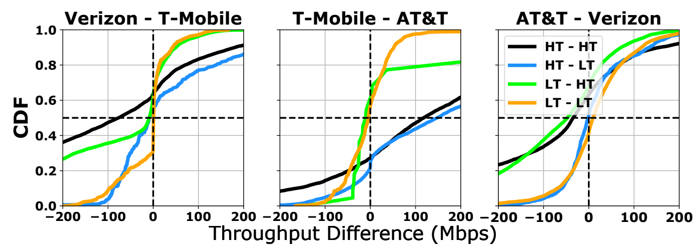

(c) CDFs of DL throughput difference for each operator pair broken down by technology.

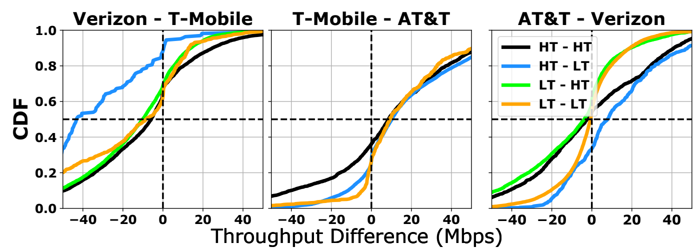

(d) CDFs of UL throughput difference for each operator pair broken down by technology.

Fig. 2: Operator-wise throughput diversity for concurrent tests.

diversity at a given location and time across operators in both traffic directions. The throughput difference between any pair of operators is often higher than 200 (25) Mbps and sometimes exceeds 400 (50) Mbps in the DL (UL) direction.

Recall from Fig. 1a that there exists a large variability in terms of cellular technology coverage across the three operators and the disparity in performance among different operators could be because of different operators using different cellular technology. To understand the impact of different cellular technologies on such diverse performance among different operators, we break down the throughput samples pairs for each pair of operators into four bins based on technology used by the two operators: HT-HT, HT-LT, LT-HT, LT-LT, where HT represents high-speed 5G (5G-mmWave and 5G-mid band), and LT represents LTE, LTE-A, and 5G-low. We classify 5G-mid band and 5G-mmWave as HT, because these are the two 5G technologies that can achieve Gbps throughput, in contrast to 5G low and LTE. This classification is also followed by the three major cellular providers in the US, which classify their 5G high throughput technology as a mix of 5G-mid band and 5G-mmwave, for example, 5G Ultra-Wideband for Verizon, 5G Ultra-Capacity for T-Mobile, and 5G+ for AT&T.

Fig. 2b plots the distribution of the 4 bins for each operator pair in the DL and UL direction. We observe no single bin dominating for all operator pairs and traffic directions. For Verizon-T-Mobile, a significant portion of the samples consist of HT-HT and LT-HT combinations in the DL direction, while LT-HT dominates in the UL direction, thanks to the significantly higher T-Mobile coverage in the 5G midband (Fig. 1a). Due to the same reason, in the case of T-Mobile-AT&T, the majority of samples belong to the HT-LT bin. On the other hand, for AT&T-Verizon, we see a different trend, with LT-HT and LT-LT being the most prevalent combinations in the DL direction and LT-LT dominating in the UL direction.

Figs. 2c, 2d plot the CDFs of throughput difference for each bin. For the HT-HT bin, T-Mobile outperforms Verizon and AT&T in both traffic directions most of the time. Additionally, Verizon HT performs slightly better in contrast to AT&T's HT, especially in the DL direction, probably because of better 5G-mmWave coverage (Fig. 1a). In contrast, for the LT-LT bin, for all operator pairs in the DL direction and AT&T-Verizon in the UL direction, the throughput difference curve is symmetric around 0, indicating roughly equal probability for one operator to outperform the other one. On the other hand, T-Mobile outperforms the other two operators in the UL direction. Overall, *the throughput difference between any pair of operators can be very large (more than 200 Mbps for HT-HT in the DL direction, and several tens of Mbps for HT-HT in the UL direction and for LT-LT in both directions), indicating substantial performance diversity across operators even for the same technology.*

Regarding the two heterogeneous bins (LT-HT, and HT-LT), one would expect the operator using an HT technology to always outperform the operator with an LT technology. However, our results, aligned with the findings of two previous works [14,17], show that this is not always the case. For example, in the case of T-Mobile-AT&T, where more than 50% of the samples are HT-LT in both directions, 20-30% of the time the HT throughput of T-Mobile is lower than the LT throughput of AT&T. Similarly, for Verizon-T-Mobile, where around 40% and 70% of the DL and UL samples respectively are LT-HT, Verizon's DL and UL LT throughput is higher than T-Mobile's HT throughput around 30% of the time. Overall, *our results show that connecting to an HT technology does not always guarantee better throughput compared to an LT technology, which further contributes to the diversity across operators in terms of performance. Consequently, a simple link selection scheme that always prefers an HT network over an LT one does not always work in practice.*

RTT diversity. Next, we look at the performance diversity in terms of RTT in Fig. 3a. Similar to throughput, we observe significant diversity in terms of RTT across operators at a given time and location. Fig. 3b, which breaks down the RTT difference samples into the 4 bins (HT-HT, HT-LT, LT-LT, LT-HT), shows

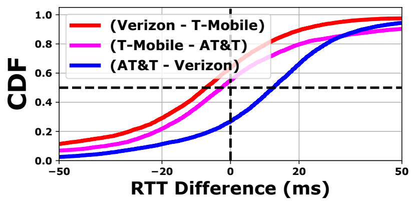

(a) CDF of RTT difference for different operator pairs.

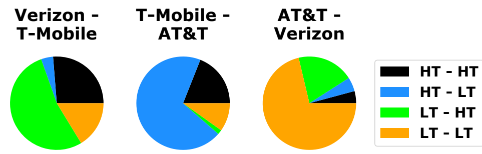

(b) Technology breakdown for different operator pairs.

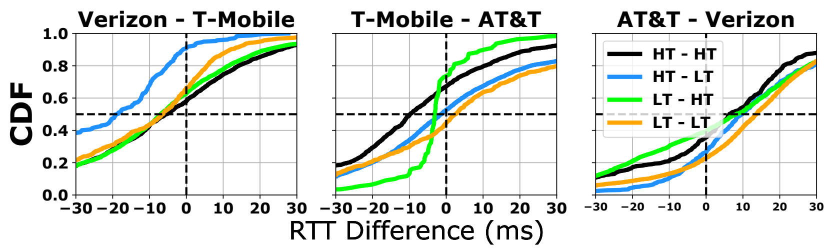

(c) CDFs of RTT difference for each operator pair broken down by technology.

Fig. 3: Operator-wise RTT diversity for concurrent tests.

that the dominant bin is LT-HT for Verizon-T-Mobile, HT-LT for T-Mobile-AT&T, and LT-LT for AT&T-Verizon, suggesting that two operators (Verizon, AT&T) often do not upgrade the UE to 5G midband/mmWave for low-rate ICMP traffic. Further, Fig 3c shows that at least 55-60% of the time (90% in some cases like Verizon-T-Mobile HT-LT) Verizon outperforms the other two operators across all bins. For T-Mobile-AT&T, T-Mobile performs better than AT&T 60% of the time in the HT-HT case, whereas for LT-LT and for HT-LT, the RTT difference curves are symmetric around 0 in the LT-LT and HT-LT cases. Interestingly for the LT-HT bin, T-Mobile outperforms AT&T 75% of the time. Overall, *our results show large RTT disparities, often more than 20 ms for all three operator pairs, which can significantly affect the QoE of latency-critical applications (AR, VR, CAVs). For example, an edge/cloud-assisted AR app that offloads frames to a server to perform object detection, requires an and-to-end latency of 33 ms. Additionally, similar to the throughput results, selecting an HT network does not always guarantee lower RTT.*

4.3 Potential benefits of multi-carrier cellular access

Our results in the previous section showed very large diversity in terms of network performance across different operators at a given location and time, suggesting that applications can benefit significantly from multi-carrier access techniques. In this section, we use our dataset to explore the theoretical (upper bound) benefits and trade-offs of the two multi-carrier access techniques – link selection and link aggregation.

We assume dual-sim smartphones capable of simultaneously using two different cellular networks.[3] Since our goal is to investigate the maximum possible

[3] Using two sim cards on smartphones is more practical than using three sims.

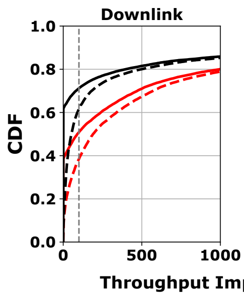

(a) T-Mobile-Verizon (Relative).

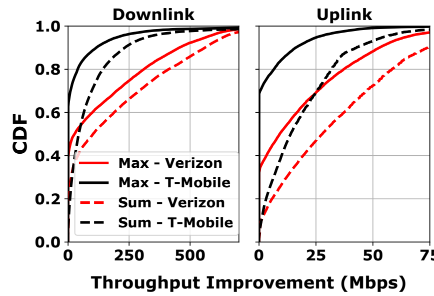

(b) T-Mobile-Verizon (Absolute).

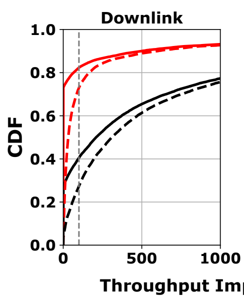

(c) AT&T-T-Mobile (Relative).

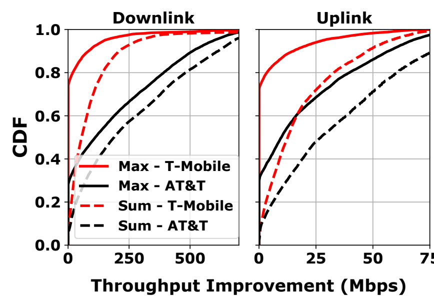

(d) AT&T-T-Mobile (Absolute).

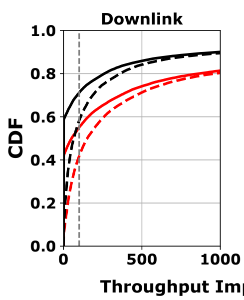

(e) AT&T-Verizon (Relative).

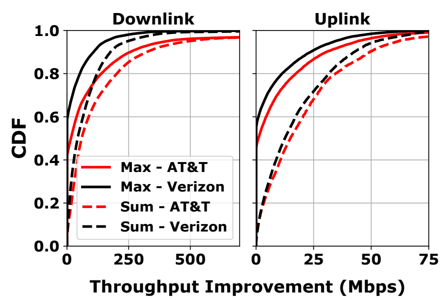

(f) AT&T-Verizon (Absolute).

Fig. 4: CDFs of throughput improvement of Max and Sum over a single operator for all operator pairs. The grey dashed line shows the 100% improvement mark.

gain of the two aggregation techniques over the performance of a single network, we assume that the link selection or link aggregation is performed by two `oracles`. In the case of link selection, the oracle (Max) always chooses the best network in terms of performance at a given time, and we ignore the network switching overhead; hence, the Max throughput is always equal to the single-path throughput over the best network. In the case of link aggregation, the oracle (Sum) performs optimal packet scheduling over the two links and we ignore the impact of congestion control; hence the Sum throughput is always equal to the sum of the single-path throughputs over the two networks.

Fig. 4 compares the relative and absolute throughput gain of Max and Sum over single-path throughput for all operator pairs. We plot both the absolute and relative throughput gains, because the relative gain can sometimes be misleading. For example, improving throughput from 1 Mbps to 2 Mbps and from 100 Mbps to 200 Mbps both correspond to a 100% relative gain, but the latter is much more important in practice.

First, in Figs. 4a, 4c, 4e, we observe significant relative throughput gains ((red or black solid lines)) of Max over the worst-performing operator, ranging from 51-208% and 15-91% in the median case and from 538-868% and 169-330% at the 75-th percentile, in the DL and UL direction, respectively. These relative gains translate to an absolute throughput gain of 15-111 Mbps (1-9 Mbps) in the median case, and up to 339 Mbps (33 Mbps) at the 75-th percentile, for the DL (UL) direction (Figs. 4b, 4d, 4f), highlighting the efficacy of link selection as a multi-carrier access scheme, especially in the DL direction. In general, the benefit of Max is higher in cases where there is a high degree of heterogeneity in the performance of the two operators. For example, in Figs. 4a and 4c, where T-Mobile is one of the operators, the gain of Max over the worst performing operator (Verizon or AT&T) is higher compared to the more homogeneous operator pair (Verizon and AT&T) in Fig. 4e.

Next, we study the benefits of Sum (red and black dashed lines in Fig. 4). We observe that there is an additional gain in using Sum over Max in all cases. For example, in Fig. 4b the median DL throughput improvement of Max over Verizon is 26 Mbps (red solid line) while aggregating Verizon and T-Mobile links provides a throughput improvement of 102 Mbps. Overall, the relative gain of Sum over a single operator ranges from 31-309% in the median case and 113-967% at the 75-th percentile for DL and from 52-191% in the median case and 113-433% at the 75-th percentile for UL. Additionally, Sum offers significant gains over both operators, while Max mainly improves the performance over the worst of the two operators. For example, Max improves the performance over the best-performing operator (T-Mobile in Figs. 4a and 4c, Verizon in Fig. 4e) in the DL direction only 25-49% of the time, while Sum improves the performance over the same operator 75-90% of the time.

We also quantify the performance gain of multi-carrier access with respect to the RTT. Here we consider an `oracle` Min that chooses the link with the minimum RTT.[4] Fig. 5 plots the CDF of absolute and relative improvement of Min over a single operator for all operator pairs. The median relative RTT improvement over the worst-performing operator ranges from 8-20% translating to an actual improvement of 5-12 ms. Further, the absolute gain over the worst-performing operator at the 75-th percentile is higher than 15 ms for all three operator pairs. As we noted in §4.2, this is a significant improvement for latency-critical applications such as AR, VR, and CAVs, which require end-to-end latencies in the order of a few tens of ms.

Although our analysis assumes an `oracle` and disregards link switching overhead and the impact of congestion control, the results in this section clearly indicate that both multi-carrier access techniques have the potential to provide substantial performance benefits over cellular networks in the 5G era.

[4] In practice, this can be achieved by sending every packet over both networks.

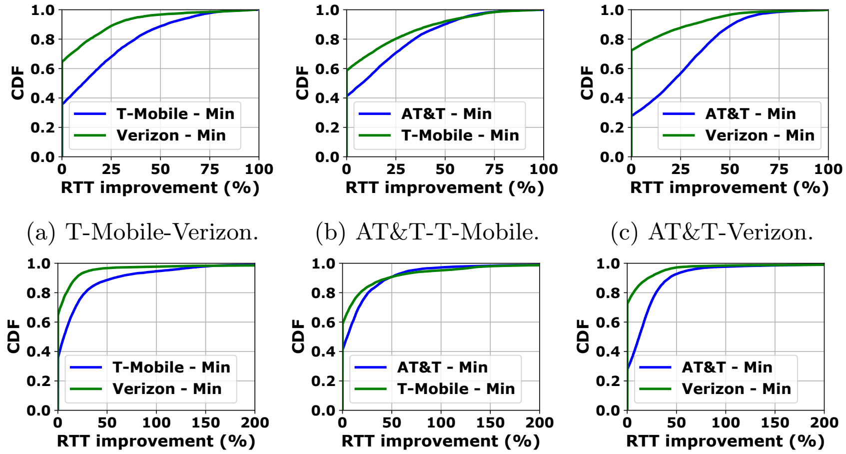

(a) T-Mobile-Verizon. (b) AT&T-T-Mobile. (c) AT&T-Verizon.

(d) T-Mobile-Verizon. (e) AT&T-T-Mobile. (f) AT&T-Verizon.

Fig. 5: CDFs of relative ((a), (b), (c)) and absolute ((d), (e), (f)) RTT improvement of Min over a single operator for all operator pairs.

4.4 Practical considerations

To further understand the potential of the two `oracle` systems, Sum and Max, in real world scenarios, in this section, we investigate two practical issues: (a) **Best-operator coherence time**: the time during which one operator continuously outperforms another operator, and (b) **Handover distribution for each pair of operators**: the handover occurrence times for two operators.

(a) Best-operator coherence time. For Max to be effective in practice, one operator should consistently offer better performance than the other operator for a substantial amount of time, as rapid switching between the two operators could incur high switching overhead and negate the potential gains. In Fig. 6, we plot the CDF of the duration during which the throughput of one operator is consistently higher than the throughput of the other operator by a threshold Θ for all three operator pairs and different values of Θ. We observe that, for $\Theta = 0$, the coherence time of the best operator is very high – 12/37/76 s in the DL direction and 6/19/47 in the UL direction at the 25-th/50-th/75-th percentile. With higher thresholds, the coherence time naturally decreases, however, we still observe high coherence time at the 75-th percentile. For example, in the DL direction, 20% of the time the throughput of a given operator remains higher than the throughput of the other operator by at least 100 Mbps for durations longer than 30 s. These results suggest that link selection as a multi-carrier access technique is likely to perform well in practice.

(b) Handover distribution. Handovers in cellular networks are viewed as temporary disruptions in the communication as the phone's connection transfers from the old base station to the new one. Previous work [16] has shown that handovers in 5G networks can significantly affect the application QoE. Link aggregation techniques have the potential to mask such disruptions, by sending all the packets over one link when the other operator's link is temporarily disrupted due to a handover, if handovers for the two operators occur at different times,

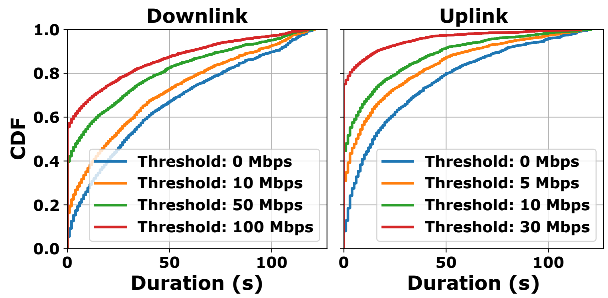

Fig. 6: Best-operator coherence time.

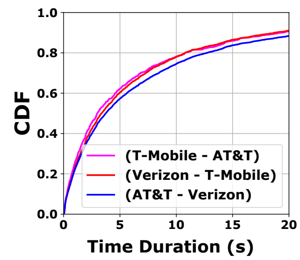

Fig. 7: Inter-operator handover intervals.

but offer no protection if handovers occur simultaneously (or very close to each other). In Fig. 7, we plot the CDF of the *inter-operator handover interval* for a given operator pair, defined as the time span between a handover in one operator's network and the nearest handover in the other carrier network. We observe that the probability of handovers to happen at the same time for an operator pair is nearly 0 and the median inter-operator handover interval is about 3 s. This finding is in line with the findings in [19], which studied the potential of MPTCP over multiple LTE networks in the case of high-speed rails. Given that the handover duration is typically shorter than 200 ms [16,14], *link aggregation shows high promise to mask handover-induced disruptions in real-world cellular network deployments.*

5 Conclusion

We performed the first large-scale empirical study of the potential gains of the two multi-carrier access approaches – link selection and link aggregation – in real-world environments in the 5G era via an extensive measurement campaign involving three cross-country driving trips covering a total of 8k+ Km spread out over a 12-month period, while simultaneously measuring the performance of the three major US mobile operators. Our measurements show that there exists a substantial amount of diversity in cellular network performance across different operators at a given time and location in terms of both throughput and RTT. Our trace-driven analysis of the theoretical gains of the two multi-carrier access techniques shows that they can provide significant gains over the single-path throughput/RTT of the worst-performing operator, especially in the DL direction. Additionally, our results show that the best-operator coherence time is in the order of several tens of s in the median case and the probability of handovers to happen at the same time for an operator pair is nearly 0. Although our analysis assumed optimal link selection/multipath scheduling algorithms and ignored link switching overheads and the impact of congestion control, a number of recent works [18,42,31,9,41,30,20,35] have proposed practical link selection/aggregation schemes that have been shown to work well under limited evaluation settings. Our work suggests that such schemes have a great potential and should be further explored to improve performance and QoE of bandwidth intensive and latency-critical 5G "killer" apps.

References

1. Dataset: A Large-Scale Study of the Potential of Multi-Carrier Access in the 5G Era. https://github.com/NUWiNS/pam2025-multi-carrier-dataset (2025)
2. Abdesslem, F.B., Abrahamsson, H., Ahlgren, B.: Measuring Mobile Network Multi-Access for Time-Critical C-ITS Applications. In: 2018 Network Traffic Measurement and Analysis Conference (TMA) (2018)
3. Abrahamsson, H., Abdesslem, F.B., Ahlgren, B., Brunstrom, A., Marsh, I., Björkman, M.: Connected vehicles in cellular networks: Multi-access versus single-access performance. In: 2018 Network Traffic Measurement and Analysis Conference (TMA) (2018)
4. XCAL Solo, https://accuver.com/sub/products/view.php?idx=11;
5. Ahuja, K., Singh, B., Khanna, R.: Network selection algorithm based on link quality parameters for heterogeneous wireless networks. Optik **125**(14), 3657–3662 (2014)
6. What is AWS Wavelength?, https://docs.aws.amazon.com/wavelength/latest/developerguide/what-is-wavelength.html
7. Caso, G., Rajiullah, M., Kousias, K., Ali, U., De Nardis, L., Brunstrom, A., Alay, O., Neri, M., Di Benedetto, M.G.: An Initial Look into the Performance Evolution of 5G Non-Standalone Networks. In: Proc. of IFIP/IEEE TMA (2023)
8. Deng, S., Netravali, R., Sivaraman, A., Balakrishnan, H.: WiFi, LTE, or Both? Measuring Multi-Homed Wireless Internet Performance. In: Proc. of ACM IMC (2014)
9. Dhawaskar Sathyanarayana, S., Lee, K., Grunwald, D., Ha, S.: Converge: QoE-driven Multipath Video Conferencing over WebRTC. In: Proc. of ACM SIGCOMM (2023)
10. Dinh, P., Ghoshal, M., Koutsonikolas, D., Widmer, J.: Demystifying Resource Allocation Policies in Operational 5G mmWave Networks. In: Proc. of IEEE WoWMoM (2022)
11. Fezeu, R.A.K., Carpenter, J., Fiandrino, C., Ramadan, E., Ye, W., Widmer, J., Qian, F., Zhang, Z.L.: Mid-Band 5G: A Measurement Study in Europe and US. In: Proc. of ACM SIGCOMM (2024)
12. Fezeu, R.A.K., Fiandrino, C., Ramadan, E., Carpenter, J., Chen, D., Tan, Y., Qian, F., Widmer, J., Zhang, Z.L.: Roaming across the European Union in the 5G era: Performance, challenges, and opportunities. In: In Proc. of IEEE INFOCOM (2023)
13. Fiandrino, C., Juárez Martínez-Villanueva, D., Widmer, J.: Uncovering 5G Performance on Public Transit Systems with an App-based Measurement Study. In: Proc. of ACM MSWiM (2022)
14. Ghoshal, M., Khan, I., Kong, Z.J., Dinh, P., Meng, J., Hu, Y.C., Koutsonikolas, D.: Performance of cellular networks on the wheels. In: Proc. of ACM IMC 2023 (2023)
15. Ghoshal, M., Kong, Z.J., Xu, Q., Lu, Z., Aggarwal, S., Khan, I., Li, Y., Hu, Y.C., Koutsonikolas, D.: An In-Depth Study of Uplink Performance of 5G MmWave Networks. In: Proc. of the ACM SIGCOMM 5G-MeMU Workshop (2022)
16. Hassan, A., Narayanan, A., Zhang, A., Ye, W., Zhu, R., Jin, S., Carpenter, J., Mao, Z.M., Qian, F., Zhang, Z.L.: Vivisecting mobility management in 5G cellular networks. In: Proc. of ACM SIGCOMM (2022)
17. Kousias, K., Rajiullah, M., Caso, G., Alay, O., Brunstorm, A., Nardis, L.D., Neri, M., Ali, U., Benedetto, M.G.D.: Coverage and Performance Analysis of 5G Non-Standalone Deployments. In: Proc. of ACM WiNTECH (2022)

18. Lee, H., Flinn, J., Tonshal, B.: RAVEN: Improving Interactive Latency for the Connected Car. In: Proc. of ACM MobiCom (2018)
19. Li, L., Xu, K., Li, T., Zheng, K., Peng, C., Wang, D., Wang, X., Shen, M., Mijumbi, R.: A measurement study on multi-path TCP with multiple cellular carriers on high speed rails. In: Proc. of ACM SIGCOMM (2018)
20. Li, Y., Deng, H., Peng, C., Yuan, Z., Tu, G.H., Li, J., Lu, S.: iCellular: Device-Customized cellular network access on commodity smartphones. In: Proc. of USENIX NSDI) (2016)
21. Lv, G., Wu, Q., Liu, Y., Li, Z., Tan, Q., Yang, F., Chen, W., Ma, Y., Guo, H., Chen, Y., Xie, G.: Chorus: Coordinating mobile multipath scheduling and adaptive video streaming. In: Proc. of ACM MobiCom (2024)
22. Moreno, J., Contini, M., Aguiar, A.: 5G NSA Performance: A Measurement Study. In: Proc. of IFIP WONS (2024)
23. Multipath QUIC, https://multipath-quic.org
24. MultiPath TCP, https://www.multipath-tcp.org
25. Narayanan, A., Ramadan, E., Carpenter, J., Liu, Q., Liu, Y., Qian, F., Zhang, Z.L.: A first look at commercial 5G performance on smartphones. In: Proc. of ACM WWW (2020)
26. Narayanan, A., Ramadan, E., Mehta, R., Hu, X., Liu, Q., Fezeu, R.A., Dayalan, U.K., Verma, S., Ji, P., Li, T., et al.: Lumos5G: Mapping and predicting commercial mmWave 5G throughput. In: Proc. of ACM IMC (2020)
27. Narayanan, A., Rochman, M.I., Hassan, A., Firmansyah, B.S., Sathya, V., Ghosh, M., Qian, F., Zhang, Z.L.: A Comparative Measurement Study of Commercial 5G mmWave Deployments. In: In Proc. of IEEE INFOCOM (2022)
28. Narayanan, A., Zhang, X., Zhu, R., Hassan, A., Jin, S., Zhu, X., Zhang, X., Rybkin, D., Yang, Z., Mao, Z.M., Qian, F., Zhang, Z.L.: A Variegated Look at 5G in the Wild: Performance, Power, and QoE Implications. In: Proc. of ACM SIGCOMM (2021)
29. Narayanan, A., Zhang, X., Zhu, R., Hassan, A., Jin, S., Zhu, X., Zhang, X., Rybkin, D., Yang, Z., Mao, Z.M., et al.: A Variegated Look at 5G in the Wild: Performance, Power, and QoE Implications. In: Proc. of ACM SIGCOMM (2021)
30. Ni, Y., Qian, F., Liu, T., Cheng, Y., Ma, Z., Wang, J., Wang, Z., Huang, G., Liu, X., Xu., C.: POLYCORN: Data-driven Cross-layer Multipath Networking for High-speed Railway through Composable Schedulerlets. In: Proc. of USENIX NSDI (2023)
31. Ni, Y., Zheng, Z., Lin, X., Gao, F., Zeng, X., Liu, Y., Xu, T., Wang, H., Zhang, Z., Du, S., Yang, G., Su, Y., Cai, D., Liu, H.H., Xu, C., Zhai, E., Ma, Y.: Cellfusion: Multipath vehicle-to-cloud video streaming with network coding in the wild. In: Proc. of ACM SIGCOMM (2023)
32. nuttcp - Network Performance Measurement Tool, https://www.nuttcp.net
33. Parastar, P., Alay, A.L.O.O., Caso, G., Perino, D.: Spotlight on 5G: Performance, Device Evolution and Challenges from a Mobile Operator Perspective. In: Proc. of IEEE INFOCOM (2023)
34. Rabitsch, A., Grinnemo, K.J., Brunstrom, A., Abrahamsson, H., Abdesslem, F.B., Alfredsson, S., Ahlgren, B.: Utilizing multi-connectivity to reduce latency and enhance availability for vehicle to infrastructure communication. IEEE Transactions on Mobile Computing **21**(5), 1874–1891 (2022)
35. Ramadan, E., Narayanan, A., Dayalan, U.K., Fezeu, R.A.K., Qian, F., Zhang, Z.L.: Case for 5G-aware video streaming applications. In: Proc. of ACM SIGCOMM 5G-MeMu (2021)

36. Rochman, M.I., Sathya, V., Fernandez, D., Nunez, N., Ibrahim, A.S., Payne, W., Ghosh, M.: A comprehensive analysis of the coverage and performance of 4G and 5G deployments. Computer Networks **237** (2023)
37. Rochman, M.I., Ye, W., Zhang, Z.L., Ghosh, M.: A Comprehensive Real-World Evaluation of 5G Improvements over 4G in Low- and Mid-Bands. Tech. Rep. 2312.00957[cs.NI] (2023)
38. Wu, H., Caso, G., Ferlin, S., Alay, O., Brunstrom, A.: Multipath Scheduling for 5G Networks: Evaluation and Outlook. IEEE Communications Magazine **59**(4), 44–50 (2021)
39. Xu, D., Zhou, A., Zhang, X., Wang, G., Liu, X., An, C., Shi, Y., Liu, L., Ma, H.: Understanding operational 5G: A first measurement study on its coverage, performance and energy consumption. In: Proc. of ACM SIGCOMM (2020)
40. Yang, X., Lin, H., Li, Z., Qian, F., Li, X., He, Z., Wu, X., Wang, X., Liu, Y., Liao, Z., et al.: Mobile access bandwidth in practice: Measurement, analysis, and implications. In: Proc. of ACM SIGCOMM (2022)
41. Zheng, Z., Ma, Y., Liu, Y., Yang, F., Li, Z., Zhang, Y., Zhang, J., Shi, W., Chen, W., Li, D., An, Q., Hong, H., Liu, H.H., Zhang, M.: XLINK: QoE-driven multi-path QUIC transport in large-scale video services. In: Proc. of SIGCOMM (2021)
42. Zhou, Y., Wang, T., Wang, L., Wen, N., Han, R., Wang, J., Wu, C., Chen, J., Jiang, L., Wang, S., Liu, H., Xu, C.: AUGUR: Practical mobile multipath transport service for low tail latency in Real-Time streaming. In: Proc. of USENIX NSDI (2024)

Appendix

A Methodology

5G Carriers, User Equipment (UE), and Servers. We obtained multiple unlimited data plans from the three major U.S. service providers, Verizon, T-Mobile, and AT&T, to conduct our experiments. All three carriers have deployed 5G low-band, mid-band, and high-band (mmWave) services across the country, predominantly co-existing with LTE via the Non-Stand Alone (NSA) architecture. T-Mobile is the only U.S. carrier currently offering both NSA and Standalone (SA) 5G services. We used Samsung Galaxy S21 Ultra smartphones as our UEs, which support most of the 5G frequency bands available in the U.S. Since commercial off-the-shelf smartphones do not provide access to low-level KPIs such as handover details or 5G frequency bands, we used Accuver XCAL Solo [4] devices, which tap into the Qualcomm diagnostic (diag) interface to capture signaling messages exchanged between the UE and the base station. For server-side infrastructure, we utilized AWS EC2 cloud instances for tests with all three operators, as well as AWS Wavelength [6] edge servers for specific Verizon tests conducted in selected cities along our driving routes.

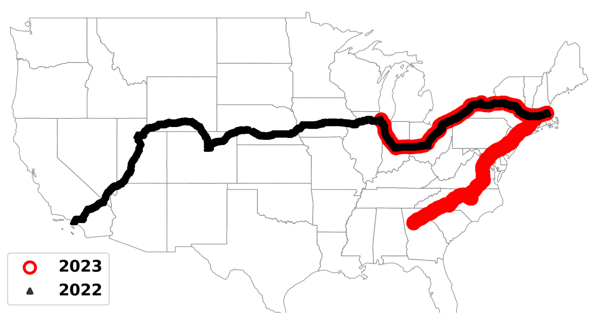

Fig. 8: Measurement routes.

Measurement routes. We conducted three measurement campaigns spread out over a 12-month (August 2022 to August 2023) and covering over 8,000 km of driving across the US. The routes of these measurement campaigns are shown in Fig. 8. In August 2022, we performed a cross-country drive from Los Angeles to Boston, crossing Las Vegas, Salt Lake City, Denver, Omaha, Des Moines, Chicago, Indianapolis, Columbus, Cleveland, and Buffalo. In May 2023 (05/13/2023 - 05/16/2023), we conducted a driving trip from Boston to Chicago, covering the same major cities of the 2022 trip along the route. In August 2023 (08/05/2023 - 08/08/2023), we conducted a second driving trip from Boston to Atlanta, passing through Providence, New Haven, New York City, Baltimore, Washington, Richmond, Durham, and Charlotte. All measurement data reported in this study were collected on interstate highways, in suburban areas, and within city limits. Additional measurements were conducted within several major cities encountered along the routes. The vehicle speed ranged from 5 miles per hour in cities to 100 miles per hour in inter-state highways.

Experimental Methodology. We conducted three types of tests: backlogged downlink and uplink throughput tests, both using nuttcp [32], and round-trip time (RTT) tests using ICMP-based ping, in a round robin fashion throughout the driving trips. The nuttcp tests lasted for 30-180 seconds each while the RTT tests ran for 20-30 seconds each. The throughput and RTT samples were logged at intervals of 500 ms and 200 ms, respectively.