To share or not to share: reliability assurance via redundant cellular connectivity in Connected Cars

Emeka Obiodu, Abdullahi Abubakar, Aravindh Raman, Nishanth Sastry, Simone Mangiante

Abstract*—As adoption of connected cars (CCs) grows, the expectation is that 5G will better support safety-critical vehicle-to-everything (V2X) use cases. Operationally, most relationships between cellular network providers and car manufacturers or users are exclusive, providing a single network connectivity, with at best an occasional option of a back-up plan if the single network is unavailable. We question if this setup can provide QoS assurance for V2X use cases. Accordingly, in this paper, we investigate the role of redundancy in providing QoS assurance for cellular connectivity for CCs. Using our bespoke Android measurement app, we did a drive-through test on 380 kilometers of major and minor roads in South East England. We measured round trip times, jitter, page load times, packet loss, network type, uplink speed and downlink speeds on the four UK networks for 14 UK-centric websites every five minutes. In addition, we did the same measurement using a much more expensive universal SIM card provider that promises to fall back on any of the four UK networks to assure reliability. By comparing *actual performance* on the best performing network versus the universal SIM, and then *projected performance* of a two/three/four multi-operator setup, we make three major contributions. First, the use of redundant multi-connectivity, especially if managed by the demand-side, can deliver superior performance (up to 28 percentage points in some cases). Second, despite costing 95x more per GB of data, the universal SIM performed worse than the best performing network except for uplink speed, highlighting how the choice of parameter to monitor can influence operational decisions. Third, any assessment of CC connectivity reliability based on *availability* is sub-optimal as it can hide significant under-performance.

***Index Terms*—4G, 5G, 5G era, V2X, redundancy, reliability, resilience, network measurement, connected cars.**

I. INTRODUCTION

Cars of the future are destined to be connected. Analysys Mason report that 164 million passenger cars were connected by the end of 2018 (about 16.5% of the total) and forecast this will rise to 831 million by 2027 (about 56% of total) [1]. The use of connectivity in a ‘Connected Car (CC)’ is much broader than the cliched quest to support ‘driverless or autonomous cars’. In practice, cars use connectivity for safety and non-safety reasons, including infotainment, navigation and control. There are broadly four groups of CC Vehicle-to-Everything (V2X) use cases: Vehicle-to-Vehicle (V2V), Vehicle-to-Person (V2P), Vehicle-to-Infrastructure (V2I) and Vehicle-to-Network (V2N) [2]. The two competing CC connectivity technologies are the IEEE802.11p based V2X version - also referred to as Dedicated Short Range Communication (DSRC) - and Cellular V2X, the focus of this paper. We note that the ecosystem seems to be coalescing on cV2X as the consensus choice [3].

E. Obiodu is with King’s College London, UK. Email: chukwuemeka.obiodu@kcl.ac.uk. A. Abubakar and N. Sastry are with University of Surrey, UK. E-mail: a.abubakar@surrey.ac.uk and n.sastry@surrey.ac.uk. A. Raman is with Telefonica I+D, Spain. Email: aravindh.raman@telefonica.com. S. Mangiante is with Vodafone, UK. Email: simone.mangiante@vodafone.com.

In this paper, we explore the role of redundancy in providing QoS assurance for V2N use cases. We compare three options for providing multi-operator CC connectivity (Figure 1). Option 1 (Figure 1a) is the current default option with an exclusive contractual relationship between the network provider and the car manufacturer or user. Option 2 (Figure 1b) is a supply-side managed multi-operator connectivity option where the CC connects to a single service provider who then manages the actual network connectivity in the backend. The service provider is either a network operator who relies on ‘national roaming’ (i.e permitted to use other networks [4]) or a mobile virtual network operator (MVNO) who relies on wholesale deals with multiple operators. In Option 3 (Figure 1c), the decision on which network to use is repatriated to be managed by the demand-side. That is, the CC is connected to multiple operators and makes its decision on which to use.

For evaluation, we conducted field measurements on a 380 kilometer stretch of major and minor roads in the UK on 15 November 2020. This included 192 kilometers on the M25, the UK’s busiest motorway, and another 188 kilometers on rural roads in the Hertfordshire and Bedfordshire countryside. We used a specially designed *Cell_Perf* Android app on four Xiaomi M1 4G devices connected to each of the four cellular networks, to capture round trip times (RTT) to 14 UK-specific websites. The app also captures the uplink/downlink speeds, page load times (PLT), TCP packet loss, jitter, network type and base station ID. In addition, we ran the app on another smartphone with a universal SIM card that has been configured to fall back on any of the four UK networks. We designate the four UK networks providers as NP 1 - 4 and the universal SIM as Supply-Side-Managed Multi-operator Connectivity (SSM-MoC). Each of NP 1-4 provides insight on Option 1, SSM-MoC is an example of Option 2 and any post-measurement combination of the measurement data for NP 1 - 4 would give insights on Option 3. We describe the Option 3 possibilities as Demand-Side-Managed Multi-operator Connectivity (DSM-MoC) and explore DSM-MoC for a combination of two, three and four network operators.

Our main contribution is to show that, contrary to common practice today (especially using ‘national roaming’ [4]) DSM-MoC will deliver superior performance over the best-performing single-operator option - up to 28 percentage points (pp) for some parameters. In our second contribution, we

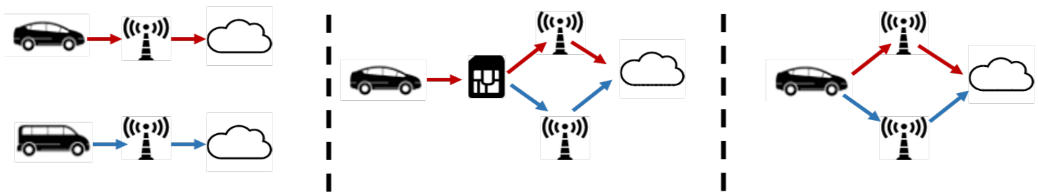

Fig. 1: Vehicle-to-Network (V2N) connectivity options for Connected Cars (CCs). While most existing CC contractual relationships with cellular network operators are based on Option (a), some MVNOs are stepping in to offer Option (b). But would Option (b) provide better reliability and QoS assurance than the hardly used Option (c)?

show that, despite costing 95x more per 1GB of data, the SSM-MoC option can significantly underperform - in this case, it had worse performance than NP 1 - 4 on RTT, PLT, packet loss. Our final contribution is to provide empirical support that assessing reliability based on *availability* can hide significant under-performance that would have been clearer if a performance-constrained reliability benchmark is used [5].

The remaining of this paper is organized as follows. Section II provides a background and related works for of CC connectivity. Section III describes the setup, routes, equipment and app used for our field measurements. In Section IV, we provide extensive analysis of the results from our study and follow this up with a discussion of our observations in Section V. Section VI identifies the further research direction for our work and some concluding remarks.

II. Background

While the current reliability of 2G/3G/4G is satisfactory for many V2X use cases, the expectation is that the more reliable 5G is needed to support safety-critical use cases [6]. Given their mobility, CCs are subjected to significant geographical variation in connectivity performance as they move from one base station to the other. [7] notes that this reality means that, compared to smartphones, there is only a small window to deliver large volumes of data, making it imperative to assure reliability in that small window. There are generally five approaches to improving and assuring the reliability of CC connectivity: (i) the use of Sidelink, introduced in 3GPP Release 16, for near field communications and which is ideal for V2V/V2P/V2I scenarios [8], (ii) improvements in the underlying reliability of the network to support V2N use cases (e.g. moving from 4G to 5G) [9], (iii) adoption of prioritisation for different service classes (e.g. via 5G Network Slicing) [10], [11], (iv) use of bespoke/private networks on the roads [12] and (v) via multi-operator connectivity or national roaming agreements on public cellular networks [13]. This paper focuses on (v).

III. Measurement parameters & setup

We conducted field measurements on a stretch of roads in South East UK totaling 237 miles (380 kilometers) on 15 November 2020. This included the entire length of the M25, the 120 miles (192 kilometers) orbital/ring road around London - the busiest motorway in the UK - and another 117 miles (188 kilometers) drive through rural roads in the Counties of Hertfordshire and Bedforshire.

We used five Xiaomi Mi 4i devices (released April 2015) and connected to SIM cards for Operators NP 1 - 4 and SSM-MoC. The CPU of the devices is the Octa-core (4x1.7 GHz Cortex-A53 & 4x1.0 GHz Cortex-A53), running on the Qualcomm MSM8939 Snapdragon 615 (28 nm) chipset. The devices run Android version Android 5.0.2 (Lollipop) and support Android SDK version of up to API level 23. Each of the devices had a SIM card for operators NP 1 - 4 and SSM-MoC. While NP 1 - 4 SIM cards were picked up locally and the cost of data usage on them was approximately £1 / GB, the SSM-MoC SIM card had to be specially ordered and its aggregate cost of data was £95 / GB.

Our measurement app, the specially designed *Cell_Perf* Android app, is based on an adaptation of the Multiping-for-Android app [14]. It enables us to conduct 24 different measurements for 14 websites every 5 minutes. These include RTT on both TCP port 80 and TCP echo port 7, PLT, packet loss, jitter, uplink speed, downlink speed, network type (e.g. LTE, HSPA+ etc), base station ID, base station location, device CPU parameters, device RAM usage. The choice of websites was based on their Alexa Ranking as at December 2019 and on their perceived importance to the UK digital society. In total, we recorded 728 readings per device in 5 hours 20 minutes of measurement, a total of 3,640 measurement cycles and 87,360 individual readings. In table I, we summarise the benchmarks we are using to evaluate performance.

IV. Results

We firstly define R_s as the reliability that an active connectivity path exists between our device and the internet. We define Q_s as the performance constrained reliability when required R_s meets a specific performance indicator threshold. For NP 1 - 4 and SSM-MoC, Q_s is based on the field measurement on *each network*. For DSM-MoC options, Q_s is based on meeting the threshold on *any network*. We structure the rest of our analysis based on the reliability framework of network reliability, data plane reliability and performance reliability while comparing the actual performance of NP 1 - 4, SSM-MoC and 11 combinations of DSM-MoC.

A. Reliability gaps in overall performance

In Figure 2, we observe that there are significant performance gaps across all the four NPs and SSM-MoC on all the

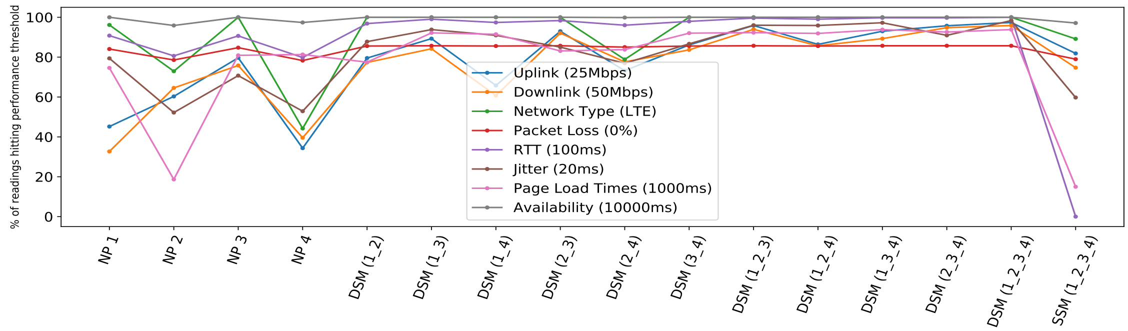

Fig. 2: Performance comparison for 8 measurement parameters: actual performance from 4 network providers (NP 1- 4) & Supply-Side Managed provider (SSM); projected performance for 11 combinations of Demand-Side Managed (DSM) setups. The clear out-performance of DSM shows that a multi-operator setup will provide better reliability and QoS assurance

TABLE I: Performance thresholds (Adapted from [10], [15])

Metric	Benchmark
Uplink speed	25 Mbps
Downlink speed	50 Mbps
Packet Loss	0%
Round Trip Times (RTT)	100 ms
Jitter	20 ms
Page Load Times (PLT)	1000 ms
Availability	RTT $>$ 10,000 ms

benchmarks. This observation makes it clear that any CC that is relying on any single NP is unlikely to experience consistent reliability or QoS assurance across time and in all locations. Out of the 3,640 measurements, 68% achieved a latency of 100ms, 54% for PLT of 1000ms, 60% for uplink speeds of at least 25Mbps, 57% for downlink speeds of at least 50Mbps and 80% were connected to LTE network. The cumulative performance data masks the diversity in performance on each of the NPs too. For example, the percentage of PLT measurements on NP 2 that were below 1000ms was only 19% compared to 75% for NP 1 and 81% for NP 3/4.

It is clear from Figure 2 that a DSM-MoC setup that is able to switch between NPs will always provide a much higher reliability and QoS assurance than any single-operator option or the SSM-MoC option. We show that if a four NP, DSM-MoC setup (i.e DSM-MoC 1/2/3/4) is implemented, overall system Q_s for RTT will improve to 100%. For jitter, Q_s improves to 98%, for uplink speed 97%, for downlink speed 96%, for PLT 94%, and for packet loss 86%. The delta between DSM-MoC 1/2/3/4 versus the best single-operator NP 3 is as high as 28 percentage points (pp) for jitter, 20pp for downlink speed, 18pp for uplink speed and 9pp for RTT.

B. Network Reliability: Uplink/Downlink Speed

We measured uplink/downlink speeds using Android's NetworkCapabilities API. The *getLinkDownstreamBandwidthKbps()* and *getLinkUpstreamBandwidthKbps()* methods return the speed of the connection in kbps. Figures 3 and 4 show the uplink and downlink speeds achieved on the five NPs. On Uplink, the proportion of measurements that achieved the 25Mbps uplink speed benchmark was 82% for SSM-MoC, 80% for NP 3, 60% for NP 2, 45% for NP 1 and 34% for NP4. This reflects on the median uplink speeds achieved too: 25.5Mbps for SSM-MoC, 25.4Mbps for NP 2/3, 25Mbps for NP 1 and only 8.5Mbps for NP 4. On downlink, the proportion of measurements that achieved the 50Mbps benchmark was 76% for NP 3, 75% for SSM-MoC, 65% for NP 2, 40% for NP 4 and 33% for NP 1. Crucially, the median shows a slightly different hierarchy as SSM-MoC has the fastest downlink median speed of 50.7Mbps, followed by NP 2/3 at 50.5Mbps and NP 1 at 50.0Mbps.

We draw three insights here. First, as Figure 2 shows for uplink speeds, the actual single-operator and projected DSM-MoC performance progressively improved as was predicted in Figure 1: that is, Option 1 $<<$ Option 2 $<<$ Option 3. Second, as speed is the defining benchmark for most internet measurement studies, it is not surprising that SSM-MoC has optimised for uplink/downlink speeds. This is a crucial finding as it suggests that SSM-MoC's architectural and operational setup has been optimised to deliver fast speeds to its customers. In doing so, SSM-MoC can state that it is delivering a superior performance compared to using any of NP 1 - 4. Third, as we explore in Section V-A, SSM-MoC had the overall worst performance for many of the parameters, showing that a fast speed is not sufficient to deliver the best outcomes for a CC. This is in line with the warning in [16] that speed is not everything and faster speeds may not always bring a better experience to end users.

C. Network Reliability: Network Type

Figure 5 confirms that we were mostly connected to 4G along our driving route. We measure Network Type using the 'Network_Type' API in Android's TelephonyManager package. For 4G readings, the API returns 13. For 3G readings, it returns 15 for HSPA+, 10 for HSPA, 8 for HSDPA or 3 for UMTS. For 2G readings, it returns either 2 for EDGE or 1 for GPRS. Overall, we note that the proportion of readings reporting a value 13 for LTE was 100% for NP 3, 96% for NP 1, 89% for SSM-MoC, 73% for NP 2 and 44% for NP 4. We draw two insights here. First, given the performance of NP 3, a CC connectivity that relies on any multi-operator combination that includes NP 3 would have achieved 100% compliance. Second, Figure 5 illustrates that while the newest cellular generation is expected to provide the best connectivity experience, earlier generations can occasionally be better: NP

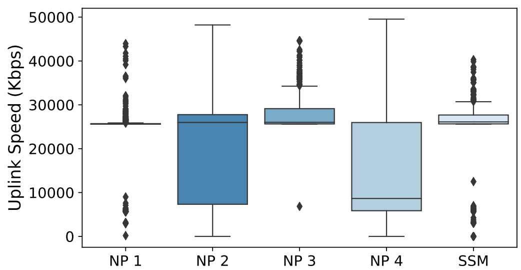

Fig. 3: Uplink analysis: the only measurement that conformed to the theoretical expectation that an SSM-MoC option should outperform a single operator setup.

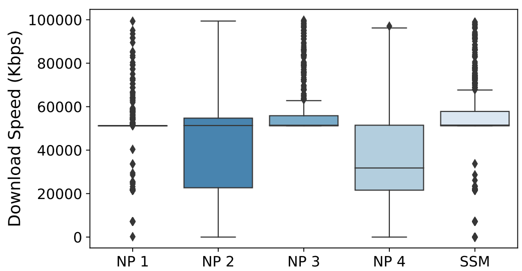

Fig. 4: Downlink analysis: Variation in performance for all NPs was at its highest for this parameter, strongly validating a multi-operator setup for QoS assurance.

3 had higher LTE rates yet SSM-MoC achieved better uplink and downlink median speeds.

D. Data Plane Reliability: Packet Loss

We measure packet loss by observing the losses recorded during our ping measurements. A lossy path will pose a huge safety risk for a CC, and will undermine many of the CC connectivity use cases. Figure 6 shows that SSM-MoC was the lossiest path, losing 100% of the packets in 18.3% of readings. This was followed by NP 2 at 18.0%, NP 4 at 17.6%, NP 1 at 14.7% and then NP 3 at 14.3%. The poor performance of SSM-MoC can be understood as it routes traffic via its cloud-based mobile core network, creating the longest physical path for traffic and increasing the chance for things to go wrong.

E. Performance Reliability: Latency & Jitter

The delay/latency in the path between a CC and a remote server is a critical indicator of reliability. We measure latency as the RTT recorded in accessing a website from our app. As some websites block ICMP ping by default, we send a 56 byte packet to both the TCP port 80 and TCP echo port 7 to assure a response from the remote server. We take the lesser of the two values as the lowest RTT measure. To calculate Jitter (i.e the variability in RTT) we repeat the RTT measurement four more times during each cycle and then calculate the standard deviation of the readings. Figures 7 and 8 show the RTT and jitter achieved on the five NPs in milliseconds (ms). The median RTTs were 38ms for NP 3, 41ms for NP 1, 42ms for NP 2, 53ms for NP 4 and 152ms for SSM-MoC. The median jitter readings were 11ms for NP 1, 15ms for NP 3/5, 18ms for NP 4 and 19ms for NP 2.

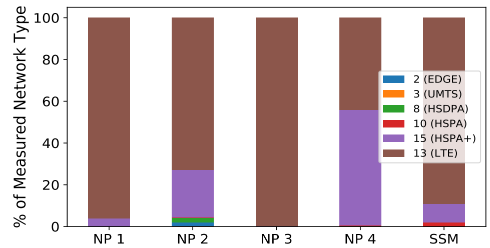

Fig. 5: Network Type analysis: All NP 3 measurements were on LTE, supporting the observation that NP 3 delivered the best overall performance for all five NPs.

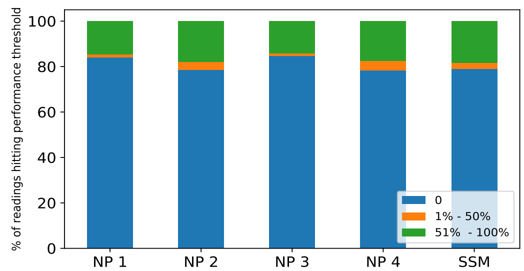

Fig. 6: Packet Loss analysis: About 80% of measurements, across all five NPs, had 0% packet loss

We draw two insights here. First, as Figures 7 and 8 show, and despite what the median values say, there were large numbers of outlier results, and huge variability in results, for each NP. This makes it challenging to rely on any of the NPs for ultra-low latency use cases in a CC environment. Second, our results on 4G suggest that latency will remain a concern in the 5G era, as 5G networks will not always deliver the hyped <10ms expected latency [6]. Our measurements were done in 2020 on 'mature' 4G networks as they have been deployed for over 10 years (i.e. deployed by 2010) and expected to deliver < 50ms latency. If similarly 'mature' 5G networks fail to deliver <10ms latency, then many of the V2X use cases that are predicated on low latency in 5G will be challenged.

F. Performance Reliability: Page Load Times

PLT reflects both the quality of the underlying connection and the design of the website that is being assessed. In this case, we use Android's webview to download the components of each of our 14 websites every 5 minutes. We record our start time and the finish time and calculate the PLT as the difference in milliseconds. For a CC that is assessing multiple services for different V2X use cases, this diversity of targeted websites gives good insight on expected performance. Figure 9 show that for the 1000ms benchmark, the proportion of readings achieving this was 81% for NP 3/4, 75% for NP 1, 19% for NP 2 and 15% for SSM-MoC. The performance of SSM-MoC is consistent with the observations on packet loss and RTT.

G. Base Station analysis

We use the CellID indicator in the Android TelephonyManager package to record the base station IDs along the 380km route. In total, the number of base stations we recorded was 55 for NP 3/4, 53 for NP 1 and 52 for NP 2/5. In terms of

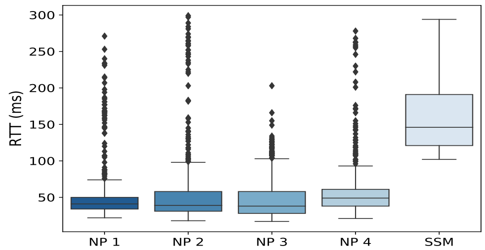

Fig. 7: RTT analysis: Large numbers of outlier results per NP.

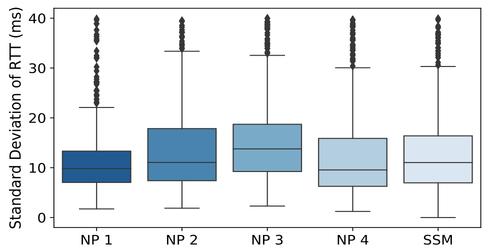

Fig. 8: Jitter analysis: Large swings in RTT measurements illustrate the challenges of providing QoS/reliability assurance.

similarity in the ID number, we note that there is only a 7% overlap between NP 2 and SSM-MoC. All other base station IDs were unique. We draw two insights here. First, as the four main UK NPs have extensive network sharing agreements that effectively split the country into two main mobile network blocks, the base stations could well be the same, or co-located on the same tower block, even if the IDs are different. Second, as Figure 10 shows, running a regression analysis between base station ID and all the measurement parameters, we note that uplink is the most correlated with base station changes. For a CC, this means that the uplink speed will most likely change during cell handovers.

V. Discussion

A. *Inferior performance of Universal SIM*

For a CC that wants to rely on a universal or super SIM connectivity, Figure 11 shows that the performance of SSM-MoC is demonstrably inferior. This is contrary to the theoretical expectation outlined in Figure 1. The SSM-MoC provider makes it clear that it routes cellular traffic through its own cloud-based mobile core network. By design, we would expect this setup to negatively impact latency-related measurements (as can be seen in Figures 7 for RTT, 8 for Jitter and 9 for PLT) as the route from our measurement device to the application server is elongated. However, on the purely network-dependent parameters (as can be seen in Figures 3, 4 and 5. for Network Type), our expectation was that SSM-MoC should be outperforming NP 1 - 4 as it is capable of switching host NPs to achieve better performance. This was only the case for Uplink/downlink speed (Figure 3) where SSM-MoC's 25.5Mbps and 50.7Mbps median was the highest. Unfortunately, our analysis of the root cause of this is limited given that SSM-MoC had the same base station ID with NP 2 for only 7% of readings (and none with NP 1/3/4).

The relative under-performance of SSM-MoC contradicts the expectation of using such setup (whether via an MVNO

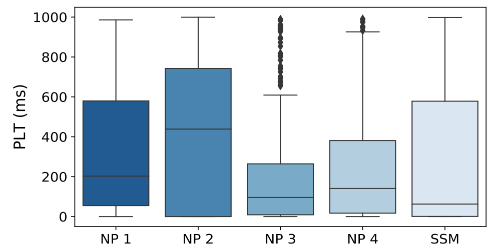

Fig. 9: Page Load Times (PLT) analysis.

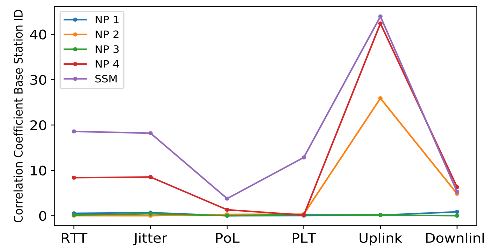

Fig. 10: RSQ - Correlation Coefficient between base station ID and individual parameters on NP 1 - 5. Uplink is the parameter that is most affected in a cell handover

or national roaming) and the justification for the hefty fees that come with it. We note that most announcements of contractual partnerships or pilots involving network providers and car manufacturers are exclusive and, we assume, under a presumption that there will be a fall back if needed. Furthermore, we note specifically that the SSM-MoC operator promotes it as being able to optimise on coverage, performance and price. Our data suggests that it has, instead, optimised mostly for speed to the detriment of other performance indicators.

B. *Optimal number of fallback networks*

The UK is among the few countries in the world with four network providers. In many countries, there are only three providers and occasionally, thanks to network sharing, only two actual physical networks. We see from Figure 12 that the worst 2-network fallback option would have under-performed the best performing single network (NP 3) for uplink/downlink speeds and network type. In contrast, the worst 3-network fallback option will consistently outperform NP 3. While this observation can not be generalised, it suggests that using DSM-MoC with only a 2-network fallback option cannot always provide assurances of superior performance, especially if those two networks are the worst performing. In which case, rather than rely on the two worst networks for reliability assurance, a CC manufacturer is better off striking an exclusive deal with the best-quality network provider.

C. *Availability vs Performance Constrained Reliability*

Figure 13 provides confirmation that assessing reliability of CC connectivity based on availability will always show a better outcome than an assessment based on a performance constrained benchmark (i.e $R_s >> Q_s$). Using latency as the basis of analysis, we assess availability (R_s) as an RTT of 10,000ms (10 seconds) - the Internet Control Message

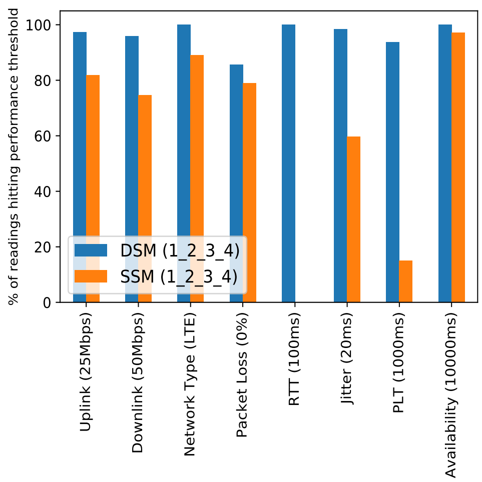

Fig. 11: Comparison of SSM vs DSM performance showing that SSM significantly under-performs a similar DSM option

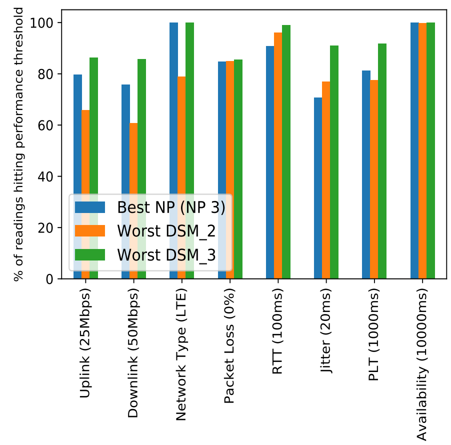

Fig. 12: Comparison of best NP vs the worst 2-network & 3-network DSM to understand optimal fallback options

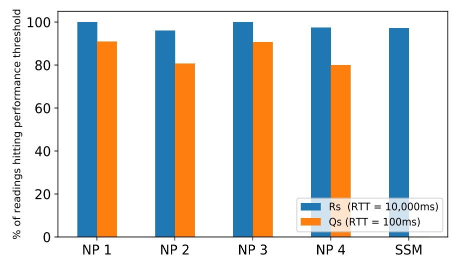

Fig. 13: Availability Analysis: Percentage measurements under 10,000ms vs 100ms. Benchmarks based on availability *R(s)* differ from performance constrained reliability *Q(s)*.

Protocol (ICMP) sessions default timeout - and a delay after which the user experience, even for mundane web browsing on smartphones, becomes unbearable [17]. For Q_s, we use 100ms which is required for competitive gaming [17], and is close to double the overall measured median RTT on NP 1 - 4. Figure 13 shows that for measurements on all five NPs, R_s is significantly better than Q_s. Accordingly, any official KPI or SLA that is benchmarked against availability R_s will indicate that the NP's network was performing well whereas a Q_s benchmark would have delivered a different outcome.

VI. Conclusion and further work

Our work has investigated how to improve connectivity reliability and provide QoS assurance for CCs to support a plethora of V2X use cases. We show that using multi-operator redundant connections, especially when managed from the demand-side by the CC, will deliver superior performance. We show also that relying on a supply-side managed multi-operator setup is risky as these can be significantly expensive (e.g 95x more per 1GB) and yet deliver inferior performance. Finally, we provide empirical support that any assessment of reliability that is based on network *availability* is unlikely to provide QoS assurance as it exaggerates the actual system performance. While our work is based on actual measurements on single networks or using a universal SIM, we have projected the performance of the demand-side managed multi-operator option. Accordingly, our next research direction is to design and implement a system that does the actual switching in real time to evaluate how it will perform in the field.

References

[1] Michele Mackenzie. The automotive sector will drive the demand for 5G. *Analysys Mason*, 2019.

[2] 5G Infrastructure Association. 5G automotive vision. *Oct 20th*, 2015.

[3] Tien Viet Nguyen, Patil Shailesh, Baghel Sudhir, Gulati Kapil, Libin Jiang, Zhibin Wu, and et al. A comparison of cellular vehicle-to-everything and dedicated short range communication. In *2017 IEEE Vehicular Networking Conference (VNC)*, pages 101–108. IEEE, 2017.

[4] Milan Gligorijević, Milica Bošković, and Aleksandar Maksimović. The role and importance of the "national roaming" in emergencies. *Nauka, bezbednost, policija*, 22(1):111–126, 2017.

[5] SH Lee. Reliability evaluation of a flow network. *IEEE Transactions on Reliability*, 29(1):24–26, 1980.

[6] Emeka Obiodu. The 5G Guide: a reference for operators. *GSM Assocation*, April 2019. [Available online] https://www.gsma.com/wp-content/uploads/2019/04/The-5G-Guide_GSMA_2019_04_29_compressed.pdf; [Accessed] 11/12/2020.

[7] Carlos E Andrade, Simon D Byers, Vijay Gopalakrishnan, Emir Halepovic, David J Poole, and et al. Connected cars in cellular network: a measurement study. In *Proceedings of the 2017 Internet Measurement Conference*, pages 235–241. ACM, 2017.

[8] Rafael Molina-Masegosa and Javier Gozalvez. LTE-V for sidelink 5G V2X vehicular communications: A new 5G technology for short-range vehicle-to-everything communications. *IEEE Vehicular Technology Magazine*, 12(4):30–39, 2017.

[9] Qian Clara Li, Huaning Niu, Apostolos Tolis Papathanassiou, and Geng Wu. 5G network capacity: Key elements and technologies. *IEEE Vehicular Technology Magazine*, 9(1):71–78, 2014.

[10] Claudia Campolo, Antonella Molinaro, Antonio Iera, and Francesco Menichella. 5G network slicing for vehicle-to-everything services. *IEEE Wireless Communications*, 24(6):38–45, 2017.

[11] Emeka Obiodu and Nishanth Sastry. From ATM to MPLS and QCI: The Evolution of Differentiated QoS Standards & Implications for 5G Network Slicing. *IEEE Communications Standards Magazine*, 4(2):14–21, 2020.

[12] Tom Lecklider. Autonomous mining equipment years ahead of car development. *RN (editor)*, 2018.

[13] Syed S Husain, Andreas Kunz, Athul Prasad, and Emmanouil et al Pateromichelakis. Ultra-high reliable 5G V2X communications. *IEEE Communications Standards Magazine*, 3(2):46–52, 2019.

[14] Softgearko. Github/MultiPing-for-Android, 2015. [Accessed] 11/12/2020; https://github.com/softgearko/MultiPing-for-Android.

[15] Ofcom. State of the mobile network experience. Technical report, 2019.

[16] Volker Stocker and Jason Whalley. Speed isn't everything: a multi-criteria analysis of broadband access speeds in the UK. 2016.

[17] Djuro Mirkovic, Grenville Armitage, and Philip Branch. A survey of round trip time prediction systems. *IEEE Communications Surveys & Tutorials*, 20(3):1758–1776, 2018.