5G-HEART White Paper

31 March 2023

5G HEART
5G Health, Aquaculture and Transport Validation Trials

Project website: https://5gheart.org/

INSIDE THIS ISSUE

1. Main conclusions, lessons learned and recommendations

1.1 Healthcare vertical

Main conclusions from the healthcare vertical

A number of use cases were implemented and tested within the areas of remote support, advanced diagnostics and remote monitoring.

The main conclusion from the healthcare trials is that 5G technology can deliver sufficient performance to enable advanced e-health applications. It is also clear that 'one-size-does-not-fit-all'. E-health use cases vary from technologically mature, close to market cases, and others which are in early stages still requiring research and testing. Some e-health applications can be served by 4G technology with reasonable quality but will benefit and improve from 5G. Different use cases also pose diverse requirements on the technology. 5G-HEART was tasked to assess 5G as an enabler and assess the performance. The following 5G KPIs have been identified as the most important: *Throughput, latency, coverage, reliability, security, location,* and *energy efficiency/power consumption*.

Throughput is especially important to carry large volumes of data and HD video content, like in remote ultrasound video setups. *Latency* is important for the application of real-time cloud processing of video and data, like e.g., in the case of the wireless colonoscopy case. *Coverage* is important if the society wants to bring e-health to where people live and stay, especially in rural areas. *Reliability* ensures life-saving critical communications. *Security* is important for the patient and to conform with GDPR requirements. A *location* service adds to remote monitoring solutions to find patients in critical situations. Not least, *energy efficiency and power consumption* for securing sufficient lifespan of body-worn sensors, like vital-sign patches.

Lessons learned

Healthcare over 5G is not *only* about 5G. It became clear that the resulting end-to-end performance of e-health applications does depend on several factors which is not part of the 5G network performance. Many applications include video transmissions and user data analysis. This was clearly demonstrated in testing of remote learning, consultancy, and attendance, delivering a real-time video feed from an operating room to students in a classroom or at home. In this case, the video processing delay far overshadowed the network delay (latency).

5G Technology maturity did not progress as fast as believed when the project started in 2019. E.g., the project aimed to test both NSA and SA setups to compare the performance, however SA implementations were not generally available during the time-span of the project and was only possible in one test.

Recommendations for adopting 5G as the connectivity enabler

From the technological perspective, 5G has the potential to be the preferred technology for providing connectivity to critical verticals, because it has the potential to deliver an end-to-end solution with sufficient quality guarantees and performance. This is not obvious, because Wi-Fi/Internet based solutions will be able to do the same in many cases. At the end of 5G-HEART project, the public discussion seems to focus around 5G to be full of broken promises, not yet delivering on the vision (see: 5G-PPP), while operators in Europe are investing heavily in building and upgrading their networks. The payback of these investments is uncertain since the customers of today want more for less. For 5G to be adopted by verticals, even further investments must be made, but the financing model is yet unknown. Traditionally, network investments have been taken by the operators, which later get their ROI from increased revenues from the customers’ increased use. The investments needed for the vertical support will be different, and needs a new ecosystem approach (see e.g., EUCNC21).

Based on the learnings from 5G-HEART, where both technology and business were trialled and studied, some recommendations can be given:

A new ecosystem must be created, based on partnering, and sharing the investment burden. Telcos cannot be expected to do all network investments to support healthcare use cases with special requirements.

Investments in coverage, capacity and capability cannot address single use cases. The industry needs to look across use cases and across verticals to build economy of scale with standardised solutions. (E.g., mission critical use cases from healthcare and defence might have common requirements).
Both these dimensions are needed to get traction on digitizing healthcare to be more that common electronic journals and administrative systems. The uptake also depends on the equipment vendors, both from the mobile industry and the healthcare technology domain, and these actors might be reluctant to invest in their product development unless a larger e-health ecosystem is in place.

1.2 Transportation vertical

Main conclusions from the transportation vertical

A set of four representative use case categories (Platooning, Autonomous/assisted driving, Support for remote driving, Vehicle data services) have been considered for the transport vertical sector. The main use case families are then divided into core and supplementary scenarios. By focusing the large-scale implementation and in-depth trialling activities to the core scenarios, the key 5G functionalities and KPIs of the transport vertical use cases have been investigated and validated.

The multiple vertical scenarios relied on dedicated 5G slices (eMBB, URLLC and mMTC) during trials. And almost all test facilities used were upgraded to support 3GPP Rel. 16, SA, by the time of final phase of trials.

The main conclusion from the trials is that 5G technology is indeed able to support the requirements of the applications considered in the transport vertical trials. Although some applications can be served by 4G technology with reasonable quality but even these will benefit from 5G, in terms of scalability and additional reliability and lower latencies provided by 5G.

Lessons learned and recommendations for adopting 5G as the connectivity enabler

5G Standalone network slicing is able to guarantee sufficient radio resources for the delivery of time critical data/traffic. Dynamic Radio Resource Partitioning was not available at the time of trials, but such measures promise fine control over the radio spectrum with slices and should be explored. Also, scaling of the slicing setup in large nation-wide commercial networks is not trivial. The requirements for deploying (priority) slicing for numerous services over a common core network were out of scope in this project and thus remain to be investigated.

In some use cases, it was observed that when it comes to the one-way latencies the main challenges with the current 5G technologies are in the UL direction, with jitter remaining high. The impact of jitter variations was seen to be more severe than end-to-end delay, directly impacting the stability and consistency e.g., of the remotely located operator’s interaction and control over the vehicle, in the case of Tele-operated service, requiring communication with remote operation centre.

In other cases, adequate service quality by a single 5G cells was limited to low number of simultaneous users. For latency and safety critical scenarios, use of 5G SA with dedicated slices is highly recommended, to provide the reliability necessary for the network to host such applications.

The findings clearly indicate that even though 5G clearly provides significantly better performance than 4G and makes network edge-based deployments of latency-sensitive services feasible, more effort needs to be put into the optimisation of UL scheduling algorithms (particularly for continuous low data rate packet streams). As the packet data scheduling algorithms in the current 5G equipment are still more geared towards DL dominated eMBB traffic, increasing support for URLLC and mMTC use cases in the most recent 3GPP releases should alleviate the problem when the related functionalities come into large-scale use also in commercial network deployments.

1.3 Aquaculture vertical

Main conclusions from the aquaculture vertical

The aquaculture vertical consisted of five scenarios which provide the benefits of using the 5G technology along with the artificial intelligence (AI) and machine learning (ML) techniques, as they are outlined below.

Scenario 1: Sensory data monitoring involved the use of IoT gateways to collect data from the sensors used to monitor the sea environment with regards to the water quality. A high number of IoT devices can be supported from 5G in this scenario. The IoT gateway is used to monitor the quality of the water which is important for the feed barge to deliver bait accurately according to the water quality.

Scenario 2: Camera Data Monitoring involved the use of underwater cameras to transmit video data through the network. The high uplink data rate when the number of underwater cameras increase can be supported by 5G in this scenario.

Scenario 3: Automation and actuation functionalities incorporated an underwater drone to monitor the fence of the cage. Fish net monitoring with an underwater robot and or drone is important for a fish farmer to avoid fish escape from the cage in case there is a damage to the net which needs to observed in real time. This monitoring requires video observance of the fish net environment and contributes to the uplink data rate requirements.

Scenario 4: Edge and cloud based computing considered the use of edge equipment for the processing of large amount of data. Edge computing provides a smart solution on the edge of the network close to where the data is produced, collected, analyzed and processed.

Scenario 5: Wireless Communication involved the communication between the cages instead of sending the data to the central or edge nodes. The use of wireless technology is preferable than the cabled underwater one which can lead to breakage and damage.

Lessons learned

During the project a single gateway was used in Scenario 1 therefore the data rate was small and 5G network capability could not be assessed. As the number of cages increase to 30 which is the case of the fish farm of Skironis in Greece, then a larger number of sensors are incorporated in the sea water for which the use of 5G is seen to be important.

The data rate used by three cameras in Scenario 2 was 15 Mbps. When the number of cages and cameras increase then a 5G network is needed. The video information transmitted in real time to a data center is used for the fish behavior/analysis, and biomass estimation which is important to be provided on time to avoid health deterioration of fish.

The IoT system that is used in aquaculture Scenario 3 is based on a large number of data that needs to be uploaded as unified at bursts, and processing and analysis of data will be centralized. Hence, uploading all the data to the cloud will consume a lot of network and cloud resources. If the network is unstable, it is not possible to process data and control the equipment in real time. If there is network instability, there is risk that data can be lost and this needs to be avoided in a real production.

Local decision making and event processing in Scenario 4 can meet the stringent demands for KPI requirements, of network resources and security challenges. Machine learning at the edge can even improve the stability of the data transmitted to the network and improve the performance of the network.

Communication between the cages in Scenario 5 can play an important role in reducing the vast amount of data that needs to be processed locally instead of sending them to the cloud. The closer processing of data can result to higher stability issues and faster reaction times.

Recommendations for adopting 5G as the connectivity enabler

For Scenario 1, when a large number of sensors is used then 5G is preferred to be utilized.

In Scenario 2, 5G solution is needed where high-quality videos can be streamed from all cages at the same time. A fish recognition and peculiar fish movements in the sea and the data streams back to the management systems can be reality with the use of 5G. The use of 5G enables high quality cameras and data feeds that power the technology of the fisheries.

An underwater drone used in Scenario 3 can be a perfect instrument for monitoring the whole fish cage including fish and cage net. A high resolution camera can improve the images that are taken not only from one cage but from multiple ones as it is the case for many fish farms.

In Scenario 4, edge computing can be used to reduce the amount of data that needs to be transferred from the fish farm to the processing center. AI and ML also has proven to provide good results for pellet detection and production improvement. With more computational power at the edge resource facilities, it is possible to store and analyze local monitoring data for faster reaction time to manage changes in the environmental conditions or improve feeding strategy. The local edge equipment can provide much faster feedback compared to performing the operations in the central cloud and communicating instructions back to the edge side. With edge computing techniques, it is possible to build a smart infrastructure to introduce AI and ML techniques that will optimize feeding process or reduce cost by minimizing human error and reacting faster to machine failures.

Replacing the fibre cables with a 5G solution as in Scenario 5 would be of a great benefit. Seen from a fish farmers point of view, the less physical equipment on the fish farms the better. It would be especially beneficial to remove the fibre cables as they are exposed to being damaged by boats or activities on the farm. In that case there is a risk of losing the connection between the sensors and cameras on the cages and the feeding barge.

2. 5G Technology Evaluation

5G-HEART assessed the performance of the 5G technology spanning from a single stand-alone vertical and application KPI assessment to concurrent trials of co-located verticals and applications. The methodology that was used for each experiment is described and the target KPI metrics used in the analysis as well as the measurement data collected are provided (see: 5G-HEART deliverables). The results were analyzed for each trialed scenario and the outcomes from the final trials for the 5G technologies validation are described in the tables below.

Healthcare use cases

Table1: Outcomes from Healthcare Use Case.

KPIs	Targets	Test cases	Achievement	Comments
User experienced DL throughput (Mbps)	Medium (> 10 Mbps)	H1C, H1B-TC1, H3C	PASSED. User experienced DL throughput was more than 200 Mbps for H1B and H1C and ~ 134 Mbps for H3C	
	High (> 60 Mbps)	H1A	PASSED for all network types. In H1A it was ~ 60 Mbps for 4G and > 600 Mbps for 5G.	
User experienced UL throughput (Mbps)	Medium (> 10 Mbps)	H1C, H1B-TC2, H3C	PASSED. User experienced UL throughput was more than 30 Mbps for H1B (test case 2), H1C and H3C.	
	High (> 60 Mbps)	H1A, H1B-TC1	PASSED for 5G SA. In H1A the user experienced UL throughput was very low for 4G, while for 5G it was ~50 Mbps for 5G NSA and ~ 63 Mbps for 5G SA. In H1B (test case 1), the user experienced UL throughput was up to 70 Mbps, which is close to the target value.	
Broadband connectivity/peak data rate (Mbps)	Low (> 60 Mbps)	H1A	PASSED. For 5G SA the Broadband connectivity was more than 60 Mbps, while for 5G NSA it was 50 Mbps	
	High (> 100 Mbps)	H1B-TC2, H1C	PASSED. In H1B the broadband connectivity was ~ 784 Mbps for DL and 85.8 Mbps for UL, while in H1C it was 734 Mbps	

Latency requirements (ms)	Low (< 20 ms)	H1B-TC1	PARTIALLY PASSED. When measured from the application the E2E latency was ~ 20 ms, but with Speedtest it was 40 ms.	
	Medium (< 100 ms)	H3C	PASSED. The average latency was 27 ms	
	High (< 200 ms)	H1A, H1B-TC2, H1C	PASSED. In H1A and for 5G networks the latency was close to but below 200 ms, while in H1B (in the 2nd scenario) it was around 7 - 11 ms, which is well below the threshold. In H1C the latency was 350 ms, because communication path includes best-effort network segment (Internet). However, for this scenario the achieved latency of 350 ms was adequate.	
Reliability (%)	High (> 99.99%)	H3C	PASSED. The average service reliability was 99.994%	
Mobility (km/h)	N/A			Not of interest to the investigated scenarios
Location accuracy (m)	Medium (≤ 25 m)	H3B	PARTIALLY PASSED: PASSED in laboratory and indoor environment (1 - 10m), NOT PASSED in outdoor environment (30 m)	
Connection/device density (devices)	N/A			Not of interest to the investigated scenarios
Interactivity (transactions/s)	N/A			Not of interest to the investigated scenarios
Area Traffic capacity (Mbps/m2)	N/A			Not of interest to the investigated scenarios
Security/Privacy (Private/Public)	N/A			Not of interest to the investigated scenarios
Energy consumption (mWh)	Medium (~3 mWh)	H3A	PASSED. The measured energy consumption was less than 3.5 mWh	
Coverage gain (dB)	Medium (15 - 20)	H3A	NOT PASSED. The gain was below 15 dB	

Transportation use cases

Table 2: Outcomes from Transport Use Case.

KPIs	Targets	Test cases	Achievement	Comments
User experienced DL throughput (Mbps)	Medium (1 - 5 Mbps)	T3S1	PARTIALLY PASSED. For T3S1 the user experienced DL throughput was less than 1 Mbps, as only the remote control commands, which are very small payloads, were transmitted.	The achieved maximum performance mostly exceeded the target values for the vertical. The lower threshold for the measured values in T2S3 was caused by the tested service, not by the 5G DL. It is interesting to note that in T4S4, the requirement was achieved only for 5G network.
	High (> 10 Mbps)	T2S1, T2S3, T2S4	PASSED. For T2S1 the user experienced DL throughput was 412 Mbps, while for T2S4 it was more than 500 Mbps. However, for T2S3 the throughput was well below 10 Mbps due to the updated application	
	Ultra high (> 100 Mbps)	T4S2, T4S4	PASSED. For T4S2 the user experienced DL throughput was 460 Mbps, while for T4S4 the average throughput was 223 Mbps	
User experienced UL throughput (Mbps)	High (> 10 Mbps)	T2S1, T2S3, T2S4, T3S1 T4S1	PARTIALLY PASSED. For T2S1, T2S4, T3S1 and T4S1 the user experienced UL throughput was always more than 30 Mbps. For However, for T2S3 the throughput did not exceed 7 Mbps.	The achieved maximum performance was in the range of the composite target values for the vertical. The lower threshold for the measured values was caused by the tested service, not by the 5G UL. The upper threshold for the target values was not achieved with the tested services.
	Ultra high (> 100 Mbps)	T4S6	NOT PASSED. For T4S6 the user experienced UL throughput was 80.5 Mbps	
Broadband connectivity/peak data rate (Mbps)	High (> 100 Mbps for the DL and < 100 for the UL)	T2S4, T4S1	PASSED. The maximum throughput was more than 60 Mbps for the UL and more approximately 700 Mbps for the DL.	The achieved maximum performance was in the range of the target.
	Ultra high (> 1000 Mbps)	T4S5	PASSED. The maximum average throughput was 1018.2 Mbps.	

Latency requirements (ms)	Ultra-low (5 ms)	T4S4	NOT PASSED. The latency was 0.3 seconds for 4G and and it was improved to 8.3 ms for 5G-SA	The achieved maximum performance was in the range of the target values for the vertical. The lower threshold for the E2E/RTT target values was not achieved with the tested services. Even though the strict industrial target values were not fully achieved, in most cases the latency performance still enabled the service to be efficiently used on top of 5G.
	Low (< 10 ms)	T2S4	PARTIALLY PASSED. In most of the trials, the E2E latency was around 10 ms, but less than 20 ms for T2S4	
	Medium (< 20 ms)	T3S1, T4S3, T4S5	PARTIALLY PASSED. For T4S5, the RTT times varied from 8.9 - 23.6 ms, according to packet size. However, for T3S1, the measured DL latency was between 20-40 ms (remote control commands), while for T4S3 it was even higher	
	High (< 100 ms)	T2S1, T2S3, T4S1, T4S2	PASSED. The measured latency was 13.4 ms for T2S1, 44.4 ms for T2S3, up to 12.2 ms for T4S1 and 6.6 for T4S2	
Reliability (%)	High (> 99.99%)	T2S4	NOT PASSED: For T2S4 the reliability was less than 99.99%.	The achieved maximum performance was mostly in the range of the target values for the vertical. Reliability targets with latency limits were not achieved due to large jitter.
	Ultra-high (>= 99.999%)	T2S3, T3S1, T4S1, T4S2	PARTIALLY PASSED. For T3S1 the reliability was 100% (no packet losses), while for T4S1 and T4S2 it was also more than 99.999%. However, for T2S3 the reliability was a bit lower, specifically ~98.5%.	
Mobility (km/h)	Medium (50-200 km/h)	T2S3, T2S4, T3S1	PASSED. The achieved mobility in T2S3 was up to 100 km/h. The vehicle speed in T2S4 was limited to 40 km/h and in T3S1 to 25 km/h due to the testing environment.	The maximum vehicle speed during the trials was in the range of the composite target values for the vertical. The upper threshold for the measured values was caused by the speed limits

				on public roads, not by the 5G network.
Location accuracy (m)	N/A			The test facilities utilised in the trials were not able to support the testing of this KPI.
Connection/device density (devices)	Medium (> 430 devices)	T2S4	NOT PASSED. The achieved connection (device) density was 10 devices/km2.	The reasons for not achieving the device density requirements are explained in respective sections
	High (> 500 devices)	T4S1, T4S2	NOT PASSED. The achieved connection (device) density was 8 devices/km2 for T4S1 and up to 60 devices/km2 for T4S2.	
Interactivity (transactions/s)	Medium (> 100 transactions/s)	T2S3, T4S1	PARTIALLY PASSED. The measured interactivity was 7000 transactions/s for T4S1, but only 26.7 transactions/s for T2S3	The reasons for not achieving the interactivity requirements are explained in respective sections
	High (> 1000 transactions/s)	T2S4	NOT PASSED. The measured interactivity was 90 transactions/s.	
Area Traffic capacity (Mbps/m2)	Low (> 0.0043 Mbps/m2)	T2S4	NOT PASSED. The resulting area traffic capacity in the final trial setup was 0.0008 Mbps/m2	The achieved maximum performance was mostly in the range of the target values for the vertical.
	High (> 0.005 Mbps/m2)	T4S1,T4S2	PASSED. The resulting area traffic capacity in the final trial setup was 0.011 Mbps/m2 for T4S1 and 0.36 Mbps/m2 for T4S2	
Security/Privacy (Private/Public)	N/A			The test facilities utilised in the trials were not able to support the testing of this KPI.
Network slice instantiation /termination time	Low (< 90 sec)	T4S5	PASSED. The mean service creation time was 53.6 sec	The achieved maximum performance exceeded the target values for the vertical.

Aquaculture use cases

Table 3: Outcomes from Aquaculture Use Case.

KPIs	Targets	Test cases	Achievement	Comments
User experienced DL throughput (Mbps)	Medium (> 1 Mbps)	A1S1, A1S2, A1S3	PASSED. For A1S1 the user experienced DL throughput was 161 Mbps, for A1S2 320 Mbps and for A1S3 132 Mbps. All these were performed in the Greek pilot	The scenarios support only UL traffic; therefore the measured value was taken with dummy data transmitted through the network
User experienced UL throughput (Mbps)	High (> 100 Mbps)	A1S1, A1S2, A1S3, A1S4	PARTIALLY PASSED. For A1S4, the user experienced UL throughput was about 150 - 350 Mbps. However, for A1S1, A1S2 and A1S3 the user experienced UL throughput was lower than the target value. All these were performed in the Greek pilot. When A1S1, A1S2 and A1S3 were run concurrently to load the network, the resulting traffic fluctuated around 64 Mbps, with momentary peaks as high as 256 Mbps, (created from additional mobile UEs).	For A1S1 the UL throughput was low because only one sensor was measuring. For A1S3, the maximum UL value was a lower than the target value because only one camera was used in this scenario. Similarly, for A1S2, only three cameras were used
Broadband connectivity/peak data rate (Mbps)	Low (> 1 Mbps)	A1S1	PASSED. For A1S1, the maximum throughput was 126 Mbps for UL and 812 Mbps for DL. (Norwegian pilot)	This target was achieved in the Norwegian site since the commercial access network has 100 MHz which is higher than the 50 MHz of the trial network in Greece
	High (> 200 Mbps)	A1S5	PASSED. For A1S5, the maximum throughput was around 680 - 812 Mbps. (Norwegian pilot)	
Latency requirements (ms)	Ultra-low (5 ms)	A1S5	NOT PASSED. The average latency was 15 ms for Norwegian pilot	The reasons for not achieving the low latency requirements are explained in respective sections
	Low (< 20 ms)	A1S1, A1S2, A1S3	PASSED. For Greek pilot, the average latency was ~11 ms for A1S1 and A1S3 and ~ 15 ms for A1S2	
Reliability (%)	High (> 99.99%)	A1S1, A1S2, A1S3	PARTIALLY PASSED: Generally, the reliability was around 99.9999%, but in some cases it dropped up to 74%. This happened because the sensors were overloaded with data	

Mobility (km/h)	N/A			Not of interest to the investigated scenarios
Location accuracy (m)	N/A			Not of interest to the investigated scenarios
Connection/device density (devices)	N/A			Only a small number of devices were used
Interactivity (transactions/s)	Low (> 1 transaction/s)	A1S1	PASSED. The measured interactivity was 27 transactions/s for A1S1 (Norwegian pilot)	
Area Traffic capacity (Mbps/m2)	N/A			Only a small number of devices were used
Security/Privacy (Private/Public)	Medium (Restricted)	A1S1, A1S2	PASSED	

3. Slicing solutions

5G-HEART investigated slicing solution in T4S5 use case: End-to-end (E2E) slicing. This section provides the results of the experiments conducted, in support of T4S5 scenario in the Surrey Platform. The experiments consist of the baseline throughput, service creation and RTT tests, conducted using the 5GNR Rel.16 SA network configuration, using a commercial handset. The tests aim at verifying the baseline slicing performance as well as throughput and latency of 5G NR SA configuration, over the 5GENESIS Surrey Platform.

Service instantiation and network slicing inside shared compute virtual infrastructures has been a key target, and in this respect, the Service Creation Time experiments are essential to evaluate the platform's performance and capabilities. The experiments provide metrics regarding the duration of the process, beginning from the moment a request is received by the Slice Manager, until the Mobile Network is fully operational and ready to accept UE's connections. Time records are reported for each one of the Service Creation stages (NS placement time, NS provisioning time, NS deployment time, Transport network configuration time (via WIM) and Radio configuration time) individually, that together constitute the overall "EtE slice deployment time (aka. "service creation time")". The collected results are illustrated in Figure 1.

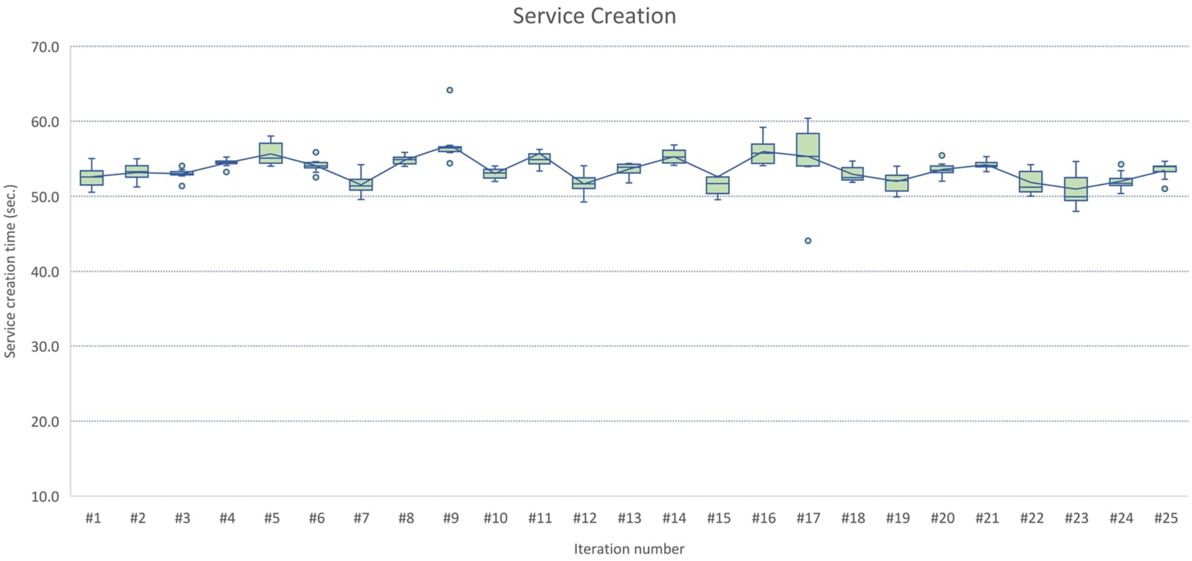

Figure 1. 5G Slice Service Creation Time test results.

The purpose of this RTT test case is to assess the end-to-end RTT & jitter from a 5G UE client to a server over a 5GNR SA Rel.16 mobile network. The mean RTT measured is 12 ms. It should be noted that the measured RTT values are lower than the respective values in similar setups of the other 5GENESIS platforms. This is a result of the fact that the 5G RAN components in the Surrey Platform are connected to the SDN switch in the core network via fast and reliable fibre links. Moreover, the Surrey Platform 5G network consists of powerful fast performing servers that are able of performing all required tasks in a computationally efficient manner, thus reducing the resulting latencies. The results of the round-trip-time tests are shown in Figure 2.

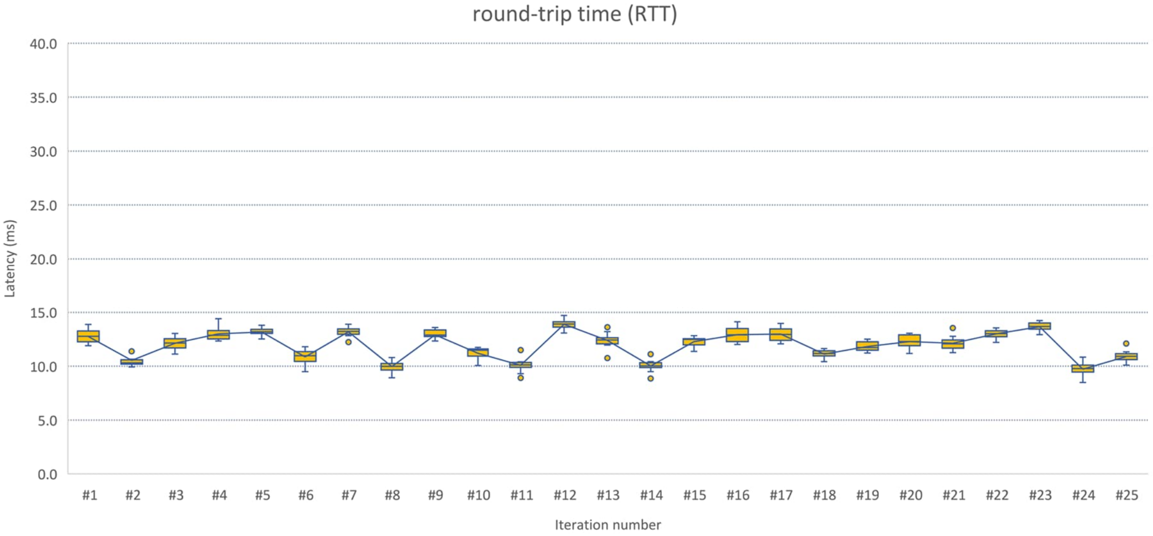

Figure 2. Round Trip Time test results

In the UL direction, the throughput of both UDP and TCP protocols was measured. For UDP, the 95th percentile throughput was 157 Mbps while for TCP the respective throughout value was 148 Mbps. The results of the UL throughput tests are illustrated in Figure 3 (UDP) and Figure 4 (TCP). In the case of the TCP throughput, it can be seen that there is a somehow noticeable reduction in throughput (compared to expected theoretical) and increased variance of the results. This is due to the observed TCP-Acks and retransmissions during the time of testing.

The throughput is measured using Iperf3 agents at the edge node and the UE, in ideal line-of-sight conditions. The distance between the gNB and the UE is approx. 20m. The UDP UL throughput measurement is on average 137 Mbps (95^{th}-%ile = 157 Mbps), close to the one theoretically expected PHY-layer mean throughput (i.e. 205-240 Mbps, @ SINR: 30-35 dB) for the given RAN (TDD frame structure) configuration and after allowing for the higher-layer overheads.

In the DL direction, the 95th percentile throughput for UDP is 1108 & for TCP is 748 Mbps.

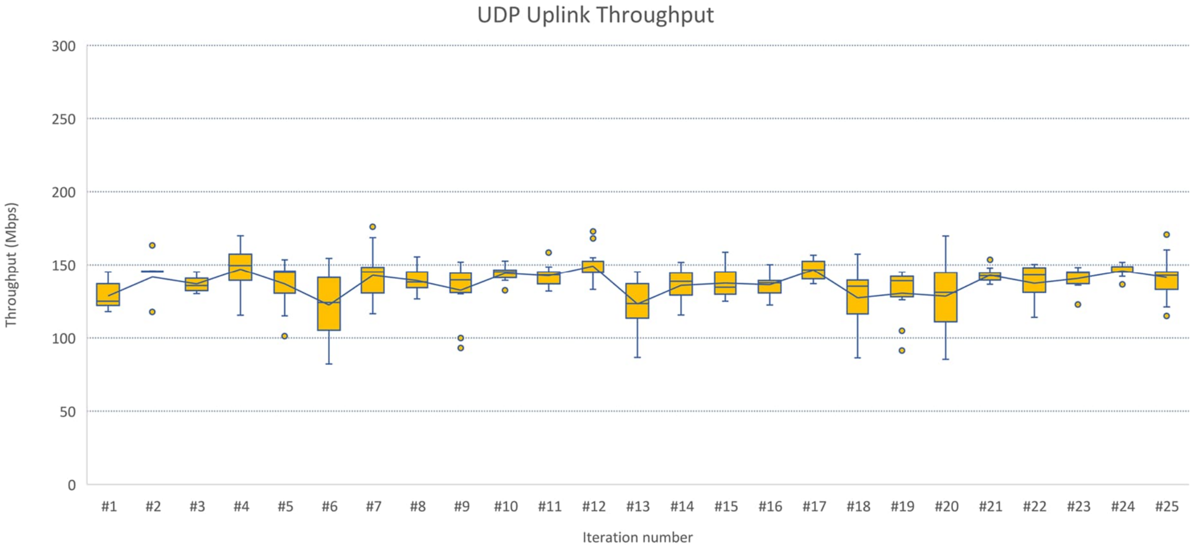

Figure 3: Throughput test results (Uplink, UDP)

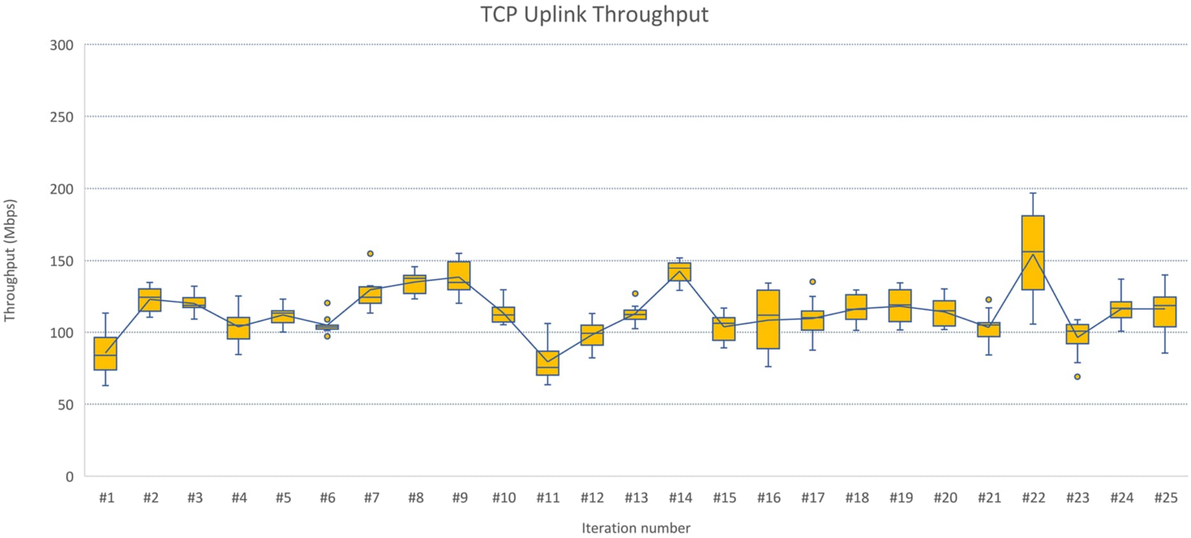

Figure 4: Throughput test results (Uplink, TCP).

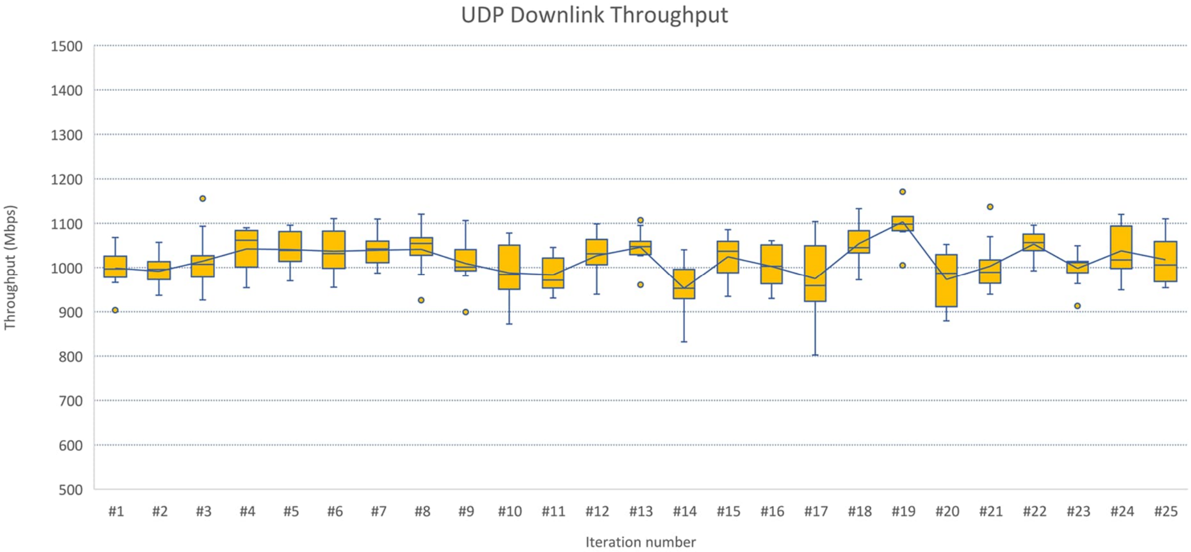

Figure 5. Throughput test results (Downlink, UDP).

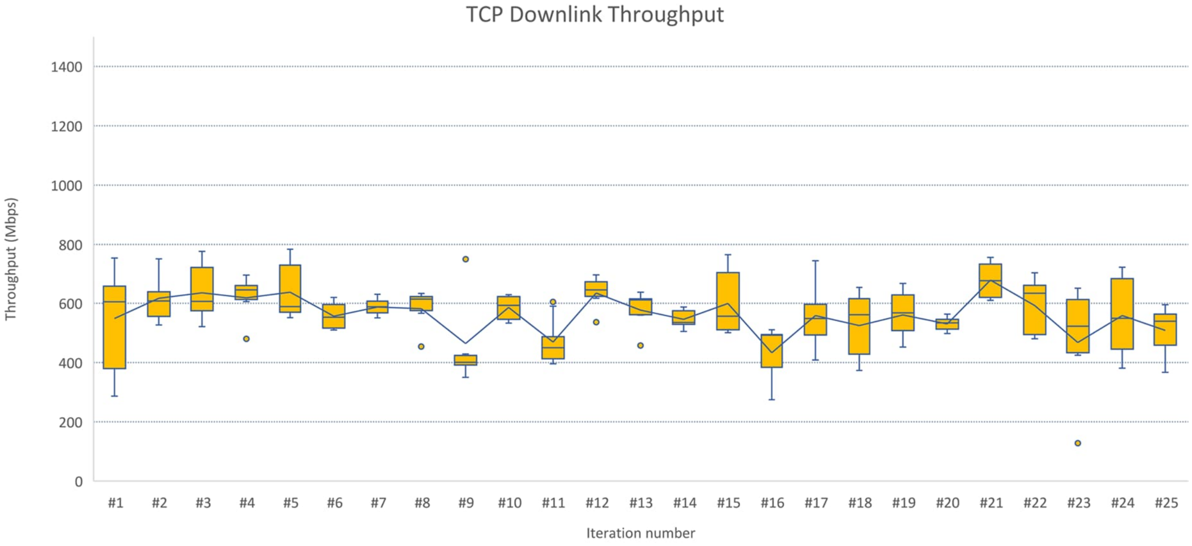

Figure 6. Throughput test results (Downlink TCP).

4. Summary

In this white paper, the main outcome, lessons learned, and recommendations of the 5G-HEART are presented. Overall, it was noticed that 5G technologies outperformed the 4G baseline, meeting most of the network requirements of the different vertical scenarios. The measured KPIs and test case parameters varied between the utilized test facilities, due to the different network and use case characteristics. Moreover, there were some configuration differences in the test networks due to the heterogeneous deployment and availability of 5G network components across the used facilities (4G, 5G-NSA, 5G-SA, commercial 5G). These differences highlight the heterogeneous nature of the requirements stemming from the different verticals towards the 5G networks. Finally, B5G/6G will provide cloud-computing capabilities within the radio access network (RAN) closer to the end users. Therefore, with the advent of this technology, which, among others, focuses on and encapsulates edge computing, the applications with lower latency requirements as well as higher device density (e.g. IoT) will benefit the most.

References

-5G-HEART deliverables: https://5gheart.org/dissemination/deliverables/

-5G PPP: 5G Vision. The 5G Infrastructure Public Privarte Partnership: the next generation of communications services and networks. 5G-PPP 2015. https://5g-ppp.eu/wp-content/uploads/2015/02/5G-Vision-Brochure-v1.pdf

-EUCNC 21: Brandsma, E., Hallingby, H. K., Lehne, P. H. A 5G health use case calling for ecosystem strategies: Resolving technology and business dependencies necessary to kick off the market. In Proceedings: The 30th European Conference on Networks and Communications (EuCNC), 8-11 June 2021, Porto (online). DOI: 10.1109/EuCNC/6GSummit51104.2021.9482489.