voipfuture

WHITEPAPER

2017

Will VoLTE Really Benefit from EVS?

Special Edition for the Mobile World Congress in Barcelona 27 Feb - 2 Mar 2017

Wendenstr. 4
20097 Hamburg
Germany

+49 40 688 90 01 0
info@voipfuture.com
www.voipfuture.com

The rise of voice over IP (VoIP) telephony services is opening up the opportunity for network operators to significantly improve the customer experience. Wideband and superwideband codecs, such as G.722, AMR-WB and EVS promise to boost the audio quality of VoIP services and improve customer satisfaction. Consequently, service providers and network operators are beating the drum for new high-quality audio services, which Skype, Facetime and other OTT services have been offering for years.

At first glance, the benefits are obvious. The main advantages of wideband over narrowband codecs typically cited by users are:

- clearer overall sound quality
- easier to recognize voices, distinguish confusing sounds and understand accented speakers
- reduced listening effort, resulting in increased productivity and lessened listener fatigue

While the theoretical benefits of wideband audio are evident, it remains to be seen whether the actual user experience in real networks can live up to the promise. This whitepaper reports on an analysis of call data collected in large VoIP networks, which suggests that customer satisfaction may depend on other factors.

Background

Voipfuture has performed an extensive analysis of quality-enriched call detail records (CDR) collected from Qrystal VoIP monitoring systems deployed in live networks. The aim was to understand the impact of call quality on call duration. The average call duration (ACD) is typically seen as a measure for the quality of a specific service or route and thereby implicitly also for the customer satisfaction.

DATA SET

The raw data set covers nearly 30 million international, domestic and mobile (VoLTE) calls. The CDRs were anonymized, i.e. calling and called party numbers were scrambled. The calls used at least 22 different types of codecs including G.711 (A-law/µ-law), G.729, G.722, G.723 and various modes of AMR-NB and AMR-WB. The majority of these codecs were however used so rarely that respective CDRs had to be excluded from the analysis. In summary, the following calls were not considered for the analysis:

- calls which did not use the top four most widely used codecs G.711 (A-law/µ-law), G.729, AMR-NB 12.2k or AMR-WB 12.65k
- calls which lasted less than 10 seconds or more than one hour
- calls where at least one call direction was impacted by duplicate packets

At the end of this filtering process the data body consisted of more than 16 million CDRs – the codec distribution is shown in the table below.

Codec	Net bitrate/kbit/sec	#CDRs	Percentage
G.711	64	10,998,417	67.3%
G.729	8	4,108,152	25.1%
AMR-NB 12.2k	12.2	1,044,612	6.4%
AMR-WB 12.65k	12.65	202,906	1.2%

Table 1: Codec distribution

VOIPFUTURE : WILL VoLTE BENEFIT FROM EVS?

Note that codecs can change in the course of a call. The CDRs provided by the Voipfuture Qrystal system list the main codec per direction, i.e. the codec that was used for most of the call duration.

CALL DETAIL RECORDS

This study uses call detail records produced by Voipfuture's Qrystal product [1], a non-intrusive VoIP monitoring system estimating the call quality of real calls. Among other things Qrystal analyses the flow of RTP packets, their interarrival times and the information contained in the IP, UDP and RTP headers. For every five second segment of each RTP flow Qrystal generates a quality summary, which includes several hundred metrics describing the segment [2]. This quality summary contains basic information about the source and destination IP addresses, VLANs and DSCP classes used, as well as detailed information about packet losses, packet interarrival times and estimated R-factor and MOS values. Qrystal also marks 'critical' five-second segments which suffer from burst loss or excessive jitter.

Qrystal Connect creates quality-enriched call detail records summarizing the five second data, e.g. by storing the minimum, average and maximum MOS value for each call direction. In addition, the critical quality ratio for each call direction is calculated, which is the ratio of critical five second segments over all segments. In total each CDR has more than 200 fields.

CALL DURATION

Call duration is the time difference between call establishment and call termination. Generally, it is affected by many factors, such as the calling and called party situation, the amount of information to be exchanged, social circumstances, gender and age of call parties and their nationalities. However, such details are not available for the CDRs underlying this study. The following section looks into the impact of some technical factors on the call duration.

Factors Influencing Call Duration

A previous study [3] has confirmed a general dependency of call quality on the average call duration in mobile networks. However, the study does not provide details on how different aspects of "call quality" affect the call duration. A service provider obviously has no influence on the gender, nationality and circumstances of the call parties, but can only control the service's technical aspects, i.e. the codec choice and the network performance.

DEPENDENCY ON CODEC QUALITY

The average call duration over the entire (filtered) data set is 220 seconds. The table below shows how the average call duration depends on the calls' main codec type.

Codec	Max. R-Factor	Max. MOS_{NB}	Average call duration / seconds
G.711	93	4.41	237
AMR-WB 12.65	87	4.26	217
AMR-NB 12.2	83	4.13	209
G.729	82	4.10	180

Table 2: Best quality achieved and average call duration per codec

The relative impact of codec quality on the average call duration

Fig. 1: Codec quality and ACD (G.711 is set as 100%)

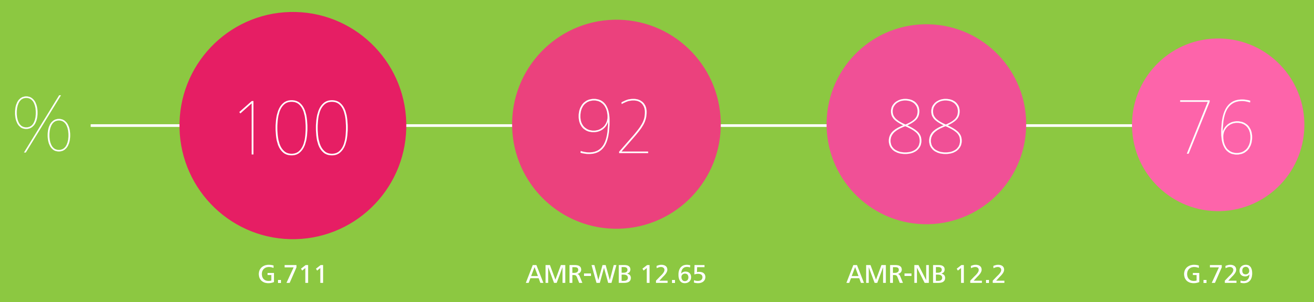

Higher levels of jitter and packet loss lead to lower call durations

Fig. 2: Call duration (seconds) and network performance (CDR classes)

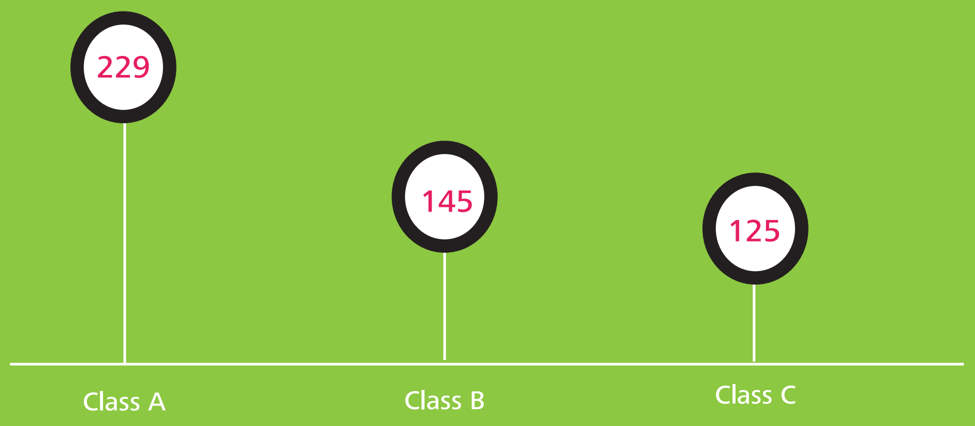

The data clearly shows that the best possible user experience that can be achieved by a codec, i.e. in the absence of critical loss and jitter, directly correlates with the average call duration. The better the codec's R-factor/MOS, the longer the average call duration.

Figure 1 shows the relative impact of codec quality on the average call duration. The ACD of G.711 is set as 100%.

This finding supports the introduction of wideband and superwideband codecs and could be seen as justification for associated investments – if networks were perfect.

DEPENDENCY ON TRANSPORT QUALITY

Live networks always introduce a certain level of packet loss and jitter to the traffic, with negative impact on the voice quality. This section investigates the effect of network performance and transport quality on the average call duration.

Qrystal marks any five second RTP time slice as critical, if it exhibits packet burst loss or excessive jitter, that will likely lead to effective packet loss. The CDRs created by Qrystal Connect include the critical minute ratio for each direction of a call. This is the ratio of critical five second segments over all RTP stream segments and thereby a good indicator for the amount of transport quality impairments affecting a media stream.

The CDRs were assigned to one of three classes A, B and C based on the critical minute ratio of their worst stream. No distinction was made between codecs. Class A contains calls which experienced a good network performance, while class C contains all calls which suffered from a bad network performance. The classes are defined in the table below.

Class	Critical Minute Ratio
Class A (good network performance)	0% – 15%
Class B (tolerable network performance)	15% – 30%
Class C (bad network performance)	30% – 45%

Table 3: CDR classes

The average call durations for both classes and each of the investigated codecs is shown in Figure 2. It clearly shows that higher levels of jitter and packet loss lead to lower call durations.

It should be noted that the critical minute ratio reveals more about the distribution of impairments over the entire call duration than about a call's overall level of packet loss and jitter. For example, a critical minute ratio of 100% may apply to RTP streams which lose two packets every five seconds. This corresponds to a packet loss ratio of less than 1%, which is not a lot by conventional wisdom.

The above has established that bad transport quality clearly reduces the average call duration. The next question is which levels of packet loss and jitter counterbalance the choice of a better codec. Figure 3 answers this question for the G.711 codec, which is the codec with the highest R-Factor and MOS under investigation.

Which levels of packet loss and jitter counterbalance the choice of a better codec

Fig. 3: Critical minute ratios at which G.711 ACD drops below worse codecs

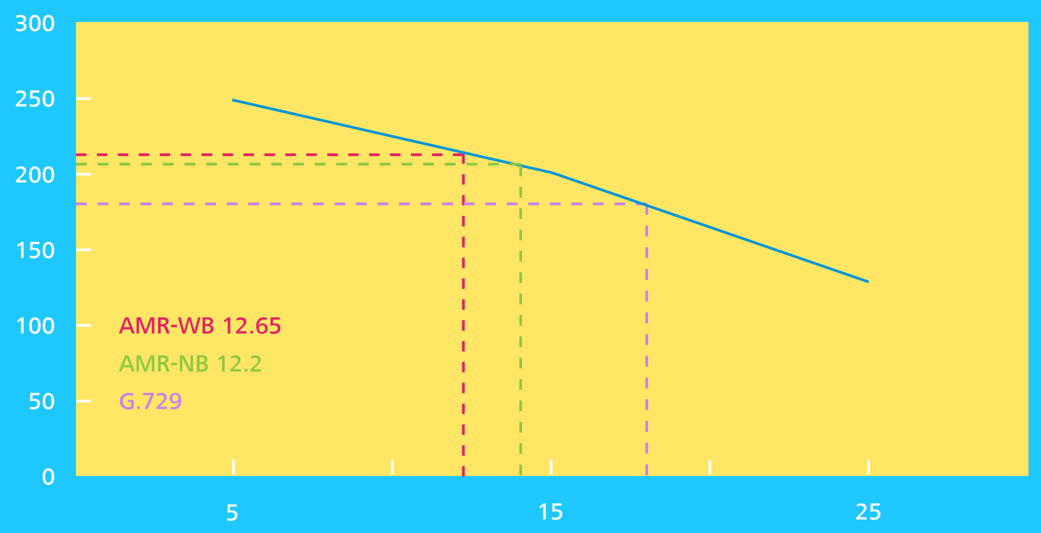

Quality of RTP traffic: LTE against fixed line network

Fig. 4: Share of streams with no critical impairments (%)

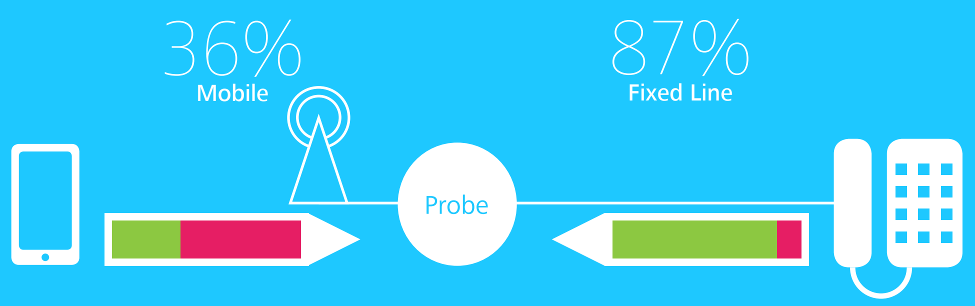

The red line in figure 3 shows that at a critical minute ratio of 12% G.711 only reaches the average call duration of AMR 12.65k under perfect conditions. At a critical minute ratio of 14% it degrades to the quality of AMR 12.2k and at 17% critical minutes G.711 goes down to the level of G.729.

Transport Quality in Fixed and Mobile Networks

One could argue that such levels of critical quality are not common in IP networks, which is mostly true for many fixed line services. However, for 4G mobile services using LTE and Wifi access packet loss and jitter are very common. Interferences on the air interface and contention through medium access control protocols prevent the timely arrival of packets and even leads to packet loss.

Figure 4 compares the quality of RTP traffic coming from a fixed network against the quality of traffic originating in an LTE network. The figure shows how RTP streams originating in fixed networks are much less impacted by critical quality than streams originating in the mobile network. The vast majority of fixed-originated streams, namely 87%, have absolutely no critical jitter or packet loss. For mobile-originated streams this is only true for 36% of streams. In contrast, 29% of all mobile-originated streams are critically impaired for up to 20% of their duration. This corresponds to the level of impairment which eliminates the benefits of better codecs as shown in Fig. 3. Only 9% of fixed-originated streams reach this impairment level.

The data covers an entire week and more than 0.5 million streams. Although this is just an example from one specific VoLTE service provider, it certainly shows a common behavior according to Voipfuture's experience.

Conclusion

This whitepaper has presented results from one of the largest studies on the connection between call quality and duration conducted so far.

THE MAIN FINDINGS ARE

1. **The average call duration depends on the codec quality – better codecs lead to longer call duration.**
2. **The average call duration depends on the transport quality – worse transport quality leads to shorter call duration.**
3. **The levels of transport quality typically observed in 4G mobile networks lead to a substantial reduction in the average call duration -- any benefits obtained from using high-quality codecs are quickly annihilated.**

Particularly the last finding is crucial to the introduction of Enhanced Voice Services in VoLTE and VoWifi networks. While this whitepaper did not analyze the actual performance of EVS in 4G networks, it is clear that the conditions in live mobile networks are so challenging that even AMR codecs struggle. Mobile operators may either hope that EVS is robust enough to deliver an excellent user experience or take action and optimize their networks.

In essence, if mobile operators neglect their network performance then any codec will give bad user experience. Even worse: the current marketing efforts are raising user expectations to a level that is calling for disappointment. Mobile operators therefore should focus on controlling their media quality and accompany the introduction of EVS with efforts to improve their RTP monitoring capabilities.*

REFERENCES

1 www.voipfuture.com
2 TM Forum GB934, “Best Practice: Voice over IP SLA Management – Application Note to SLA Management Handbook”, April 2008
3 J. Holub, J.G. Beerends, R. Smid, "A dependence between average call duration and voice transmission quality: measurement and applications," Wireless Telecommunications Symposium, 2004, pp. 75- 81, May 2004.

* Further details are presented in a paper that has been submitted to the IEEE IWQoS 2017 conference. This whitepaper is the first in a series of papers which aims to provide our customers with an in-depth understanding of how to improve the quality of experience.

About

Voipfuture is a premium voice quality monitoring vendor developing unique technology for assessing, aggregating, analyzing, and visualizing voice quality information. Voipfuture products offer a precise view on media and control plane to communication service providers, wholesalers and enterprises.

Since its launch, Voipfuture has been at the forefront of voice quality monitoring and continues to redefine Voice over IP by connecting their customers’ view on service quality with high resolution user experience.