TAKING STEPS TOWARD GREENER STREAMING

P. Angot[1], V. Lepec[1], A. Nochimowski[1], P. Gonon[2], C. Thienot[3], and J.C. Vargas-Rubio[3]

[1]Viaccess-Orca, France, [2]Celad for Viaccess-Orca, France and [3]Enensys Technologies, France

ABSTRACT

With consumer IP video traffic representing 84% of all consumer IP traffic in 2021 (up from 79% in 2016) (1), the notion of Corporate Social Responsibility (CSR) is garnering increasing attention from video service providers. The energy consumption of their service deployments needs to be, at the very least, mastered and minimized. At best, it has to achieve the carbon footprint neutrality already publicly committed to by some of the industry's key players (2).

This paper will present some of the work conducted as part of the New vidEo STandards for Enhanced Delivery (NESTED) collaborative project (3) with regards to video distribution over 5G networks. A key focus of the paper will be on the device-side contribution of next-generation codecs compared with legacy ones in terms of energy consumption. The paper will also provide a comparison of the simulations performed on unicast and multicast 5G network architectures.

INTRODUCTION

Over the past few years video traffic has grown steadily. What was limited only to the consumption of live and video-on-demand (VOD) services on a managed network from a set-top-box at home 10 years ago has turned into access to the same services from everywhere and on any type of terminal with the advent of over-the-top (OTT) video streaming and adaptive bit rate (ABR) technology. In this context, where users see their experience continuously enriched by new functionalities such as targeted advertising, multiview, watch party, or virtual reality-enhanced/360-degree video, it seems obvious that the share of video traffic as a part of overall internet usage will continue increasing. As a consequence, the question of optimizing the carbon footprint of the end-to-end technical chain enabling such advanced experiences arises. New generations of codecs are emerging, and new broadcast technologies are being implemented in order to cope with this demand of richer video experiences and the related increase in video traffic. But, what is the environmental cost of this unending race for innovation?

In this paper, we present some of the findings from our analysis related to energy consumption in the context of video content viewership from a smartphone (endpoint-side) perspective. Such findings are nurtured by the research conducted in the context of the NESTED collaborative project with the support of France's Brittany region (3) in relation to video distribution over 5G networks based on state-of-the-art standards such as Common Media Application Format (CMAF) and Dynamic Adaptive Streaming over HTTP (DASH).

Figure 1 provides a simplified view of the various technology layers involved in video content streaming. This paper focuses primarily on the endpoint side of the video content, network/transport and physical/radio layers, which proves critical in the rendering of the end user's perceived quality of experience (QoE). It notably leverages some of the empirical (when possible) and theoretical measurements performed around Viaccess-Orca's (media player) and Enensys Technologies' (multicast gateway) endpoint components. The results shared in the paper should be considered as work-in-progress and will be further enriched in the context of the two-year long NESTED project but will in any case constitute a partial view of a wider and more complex subject (especially when adding the head-end part into the equation).

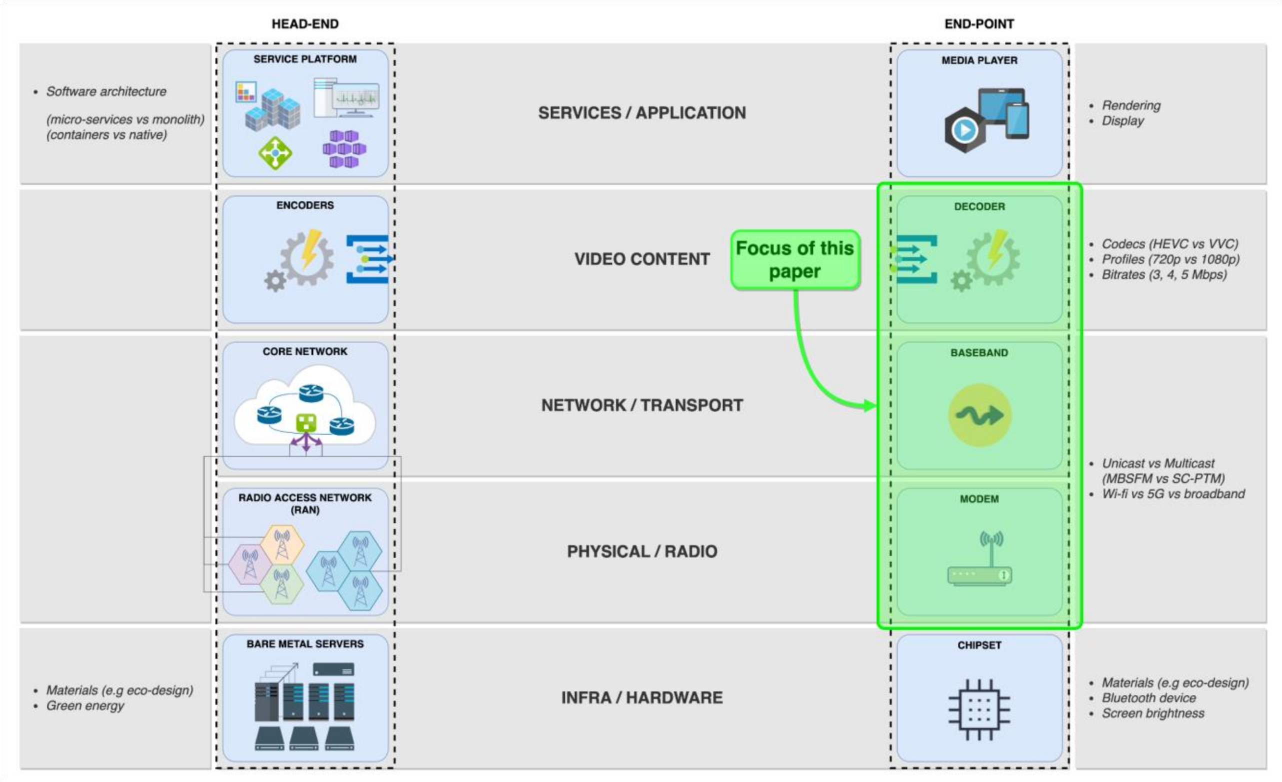

Figure 1 - Overview of the video streaming end-to-end layers

With regards to the physical/radio and network/transport layers, in the absence of empirical measurements exploitable at time this paper is written, we present the results of a theoretical study related to the impact of 5G broadcast network solutions on a smartphone's energy consumption. The second part details the results of empirical measurements centered on the High Efficiency Video Coding (HEVC) and Versatile Video Coding (VVC) codecs and the associated profiles (video content layer). Finally, we try to apply these results (and draw some — partial — conclusions) to a concrete case study: users concentrated in a stadium watching the highlights of the event they are attending.

THE USE OF MULTICAST-BROADCAST TECHNOLOGIES CAN SIGNIFICANTLY REDUCE ENERGY CONSUMPTION AT THE ENDPOINT

Energy consumption depends on many system parameters: bitrate, transmission time, transmission power, bandwidth, transmission mode, and network deployment, among others. Providing a model that considers all these parameters is not an easy task. Authors in (4), proposed a scheduling algorithm to reduce UE energy consumption when receiving in MBSFN mode. They showed that MBSFN helps reduce energy consumption on the UE compared to unicast. However, they consider the classical regular lattice for the base stations, and a model for other broadcast approaches (e.g., SC-PTM) is not provided. Authors in (5) consider a scenario in which a cell transmits data to Machine-Type Communication (MTC) devices using multicast transmission. Devices with better channel conditions receive the data from the base station and act as relays for devices at the cell border, known as Device to Device (D2D) communication. Furthermore, they assume the relays are synchronized and can perform SFN transmission. They propose an algorithm to determine the subgroup of devices that can receive data directly from the base station, the subset of relay nodes and the Modulation and Coding Scheme (MCS) for the multicast transmission and the D2D SFN transmission in order to improve the network energy efficiency. However, (5) does not address broadcast transmission in multiple cells (MBSFN) or Human Type Communication (HTC), and their approach is based on simulations.

We believe a more systematic approach to broadcast transmission is needed, as broadcast represents a promising solution to reduce energy consumption in situations where many users consume the same content at the same time (such as sport events or group communications in mission-critical scenarios). In (6) and (7), we show that broadcast transmission can reduce bandwidth usage compared with unicast. Our results motivated a study (8) in which we compare broadcast and unicast (UC) in terms of endpoint energy consumption.

Multicast-Broadcast Single-Frequency-Network (MBSFN) and Single-Cell Point-to-Multipoint (SC-PTM) are the two broadcast technologies standardized by the Third Generation Partnership Project (3GPP). SC-PTM refers to broadcast transmission in only one cell. Users in the cell interested in the broadcast content receive the same information at the same bit rate. The bit rate is set based on the user with the lower Signal-to-Interference-plus-Noise Ratio (SINR) in the cell. This is different from unicast in which each endpoint receives at a bit rate according to its own SINR as a result of link adaptation techniques. However, in terms of interference, SC-PTM and unicast are similar because neighboring cells generate interference power. On the other hand, MBSFN transmission consists of a group of synchronized cells, called MBSFN area, that transmit the same information at the same time to the users demanding the broadcast content. The transmission bit rate aims to cover the users with the lowest SINR in the whole service area. In MBSFN mode, interference is generated only by the cells outside the MBSFN area. See Figure 2.

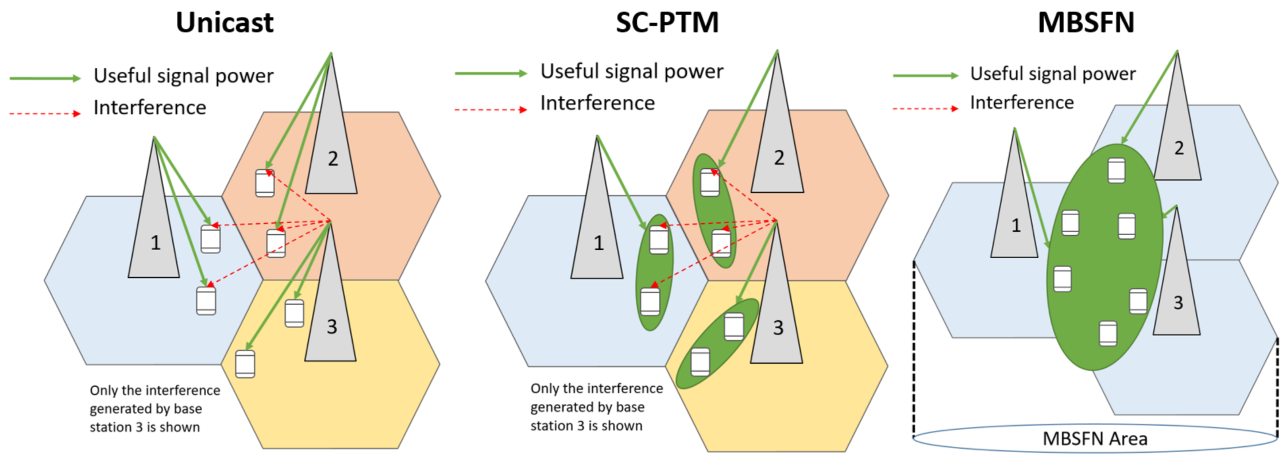

Figure 2 - Unicast and broadcast transmission modes

The study in (8) provides an analytical model to calculate the energy consumption on the user equipment (UE) (i.e., endpoint) side and the base station (BS) side in unicast, MBSFN and SC-PTM transmission modes. It also presents how to calculate the number of users per cell to switch from unicast to MBSFN or SC-PTM in order to reduce energy consumption. In the following two subsections we present and analyze the results obtained in (8).

System Model

Power consumption

The UE power consumption (P_{ue}) is calculated as $P_{ue} = k + p_{1bb} C_{Rx}$, where k and p_{1bb} are constant values given in (8) and Table 1 below. C_{Rx} is the downlink (DL) bitrate in Mbps. This model assumes a constant value for the UE downlink power level.

Energy consumption

For the energy consumption model, we consider that omnidirectional base stations with the same configuration are randomly located according to a Poisson Point Process (PPP) of density λ_{BS}. We consider the propagation effect of path loss and fading. Then, consider a service area in

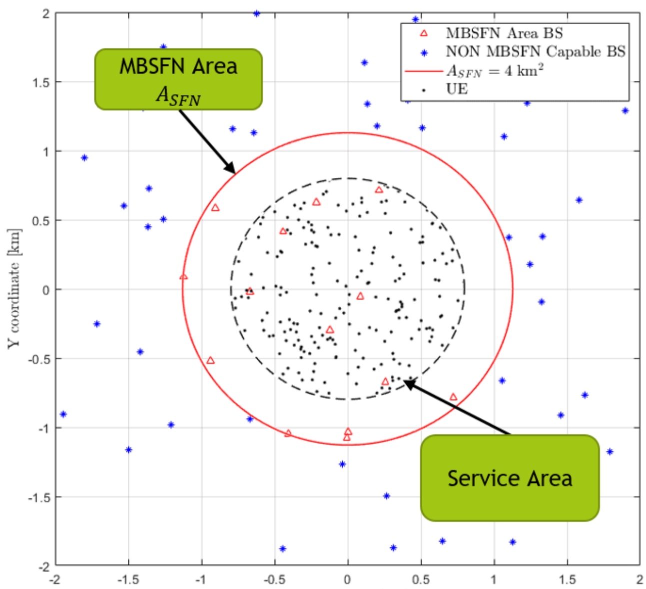

Figure 3 - Model for the energy consumption in MBSFN mode

which the users are located according to a PPP of density λ_{UE}. In MBSFN, this service area is a disk smaller than the MBSFN area and centered on it, shown in Figure 3. In MBSFN mode, the minimum SINR to cover (S_{mSFN}) sets the transmission bitrate. This minimum SINR is usually perceived by users close to the border of the MBSFN area since they receive a high interference power. Therefore, we fix the radius of the service area as ξR_s, with $0 \leq \xi \leq 1$ and calculate the probability of coverage for a user at the border to determine the bitrate of the MBSFN transmission, shown in Figure 4.

Considering this model, the total energy consumption for all the UE in UC mode is calculated as

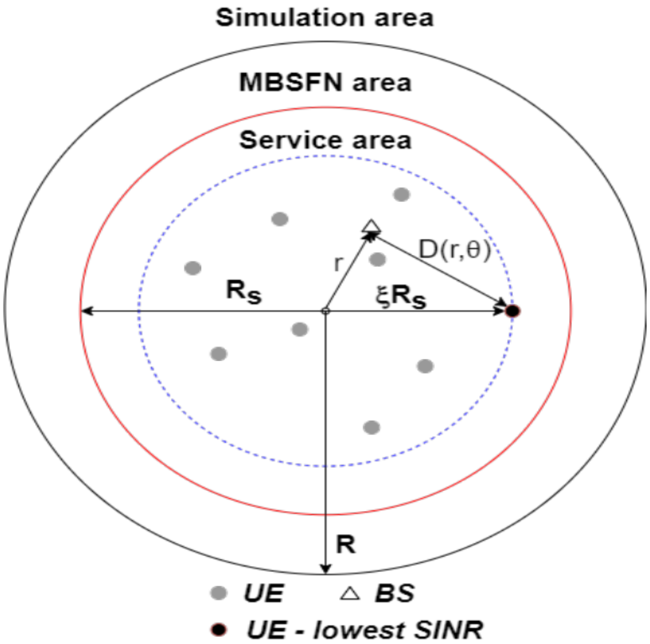

Figure 4 - Model for the MBSFN SINR distribution

$$E_{\mathrm{UE_{UC}}} = N_{\mathrm{u}} D_{\mathrm{Tx}} t_{\mathrm{UC}} \frac{N_{\mathrm{u}}}{N_{\mathrm{BS_{UC}}}} \left(K + p_{1_{\mathrm{BB}}} C_{\mathrm{UC}} \frac{N_{\mathrm{BS_{UC}}}}{N_{\mathrm{u}}} \right)$$
$$= {N_{\mathrm{u}}}^2 \frac{D_{\mathrm{Tx}} t_{\mathrm{UC}} K}{N_{\mathrm{BS_{UC}}}} + N_{\mathrm{u}} D_{\mathrm{Tx}} t_{\mathrm{UC}} p_{1_{\mathrm{BB}}} C_{\mathrm{UC}}.$$

where $N_{BS_{UC}}$is the number of BS transmitting in UC mode, D_{Tx} is the amount of data to transmit, t_{UC} is the time it takes to transmit one bit and N_u is the number of UE. t_{uc} is calculated using the SINR distribution, see (8) for further detail. The factor $\frac{N_{BS_{UC}}}{N_u}$ accounts for the fact that in unicast the bit rate is divided among the user in one cell.

A similar model is used to calculate the UE energy consumption in MBSFN and SC-PTM modes. In MBSFN

$$E_{\mathrm{UE_{SFN}}} = N_{\mathrm{u}} D_{\mathrm{Tx}} t_{\mathrm{SFN}} (K + p_{1_{\mathrm{BB}}} C_{\mathrm{SFN}})$$

where t_{SFN} is the time it takes to transmit one bit in MBSFN mode and C_{SFN} is the bit rate. Both parameters are calculated using the SINR distribution. Notice that in broadcast mode the same resources are used by all users; therefore, t_{SFN} and C_{SFN} do not depend on the number of users but on the minimum SINR to cover (S_{mSFN}).

The UE energy consumption in SC-PTM is calculated as

$$E_{\mathrm{UE_{SC}}} = N_{\mathrm{u}} D_{\mathrm{Tx}} t_{\mathrm{SC}} (K + p_{1_{\mathrm{BB}}} C_{\mathrm{SC}})$$

where t_{SC} is the time it takes to transmit one bit in SC-PTM mode and C_{SC} is the bit rate.

User Threshold

We define the user threshold as the number of users per cell from which the MBSFN or SC-PTM mode consumes less energy on the UE side than the UC mode.

Unicast to MBSFN

The user threshold to switch from UC to MBSFN is calculated as in (8):

$$U_{\mathrm{SFN_{UE}}} = \frac{N_{\mathrm{BS_{UC}}}}{N_{\mathrm{BS_{SFN}}}} \left[\frac{t_{\mathrm{SFN}}}{t_{\mathrm{UC}}} - \frac{p_{1_{\mathrm{BB}}}}{K} \left(C_{\mathrm{UC}} - \frac{1}{t_{\mathrm{UC}}} \right) \right]$$

Unicast to SC-PTM

In the case of SC-PTM, the user threshold is obtained as in (8):

$$U_{\mathrm{SC_{UE}}} = \frac{t_{\mathrm{SC}}}{t_{\mathrm{UC}}} - \frac{p_{1_{\mathrm{BB}}}}{K} \left(C_{\mathrm{UC}} - \frac{1}{t_{\mathrm{UC}}} \right)$$

Results

Results were obtained through analytical calculation using MatLab considering the system parameters presented in Table 1.

Parameter	Value
Simulation area radius (R)	$10km$
System bandwidth (W)	$10MHz$
Carrier frequency (f_c)	$700MHz$
BS transmission power (P_{Tx})	$43dBm$
Noise power (N)	$-95dBm$
Path loss exponent (α)	3.76
Path loss factor (k)	1.3804×10^{-12}
Target coverage probability	95%
BS power consumption (P_{BS})	$225W$
K	1.9071
p_{1BB}	2.89×10^{-3}

Table 1 - Simulation parameters

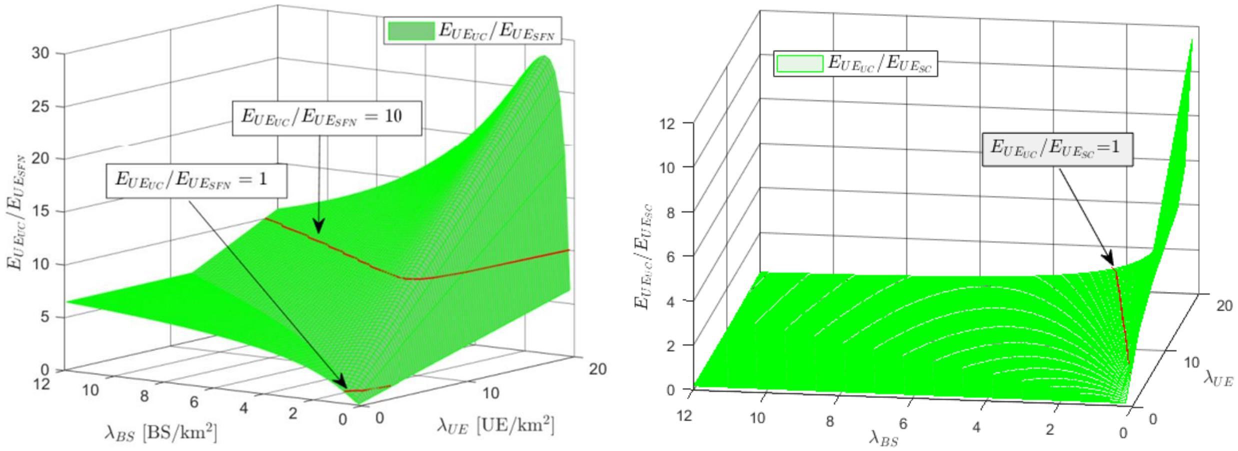

Figure 5 - Ratio between the UE energy consumption in UC and MBSFN modes (left). Ratio between the UE energy consumption in UC and SC-PTM modes (right).

The ratio between UE energy consumption in unicast mode ($E_{UE_{uc}}$) and MBSFN mode ($E_{UE_{SFN}}$) is presented in Figure 5 (left). Only when $\lambda_{BS} < 0.8BS/km^2$ and $\lambda_{UE} < 3.4UE/km^2$ the UE energy consumption is lower in UC than in MBSFN. Apart from that, we see that in most cases the UEs energy consumption is lower when receiving in MBSFN than in UC. This is because the UEs perceive a higher SINR in MBSFN mode due to a reduced interference and therefore they receive at a higher bit rate and transmission is faster.

Figure 5 (right) presents the ratio between the UE energy consumption in UC mode ($E_{UE_{UC}}$) and SC-PTM mode ($E_{UE_{SC}}$). We see that in most cases UE consumes more energy in SC-PTM mode than in UC mode. However, the UE consumes less energy when transmitting in SC-PTM mode when λ_{BS} is low and the user density is high. In SC-PTM, the total UE energy consumption increases linearly with the UE density, and it is independent of the BS density because the more BS the less UE per BS but at the same time, the higher the interference. On the other hand, in UC mode, the UE energy consumption increases exponentially (power of two) with the UE density, particularly with low BS densities.

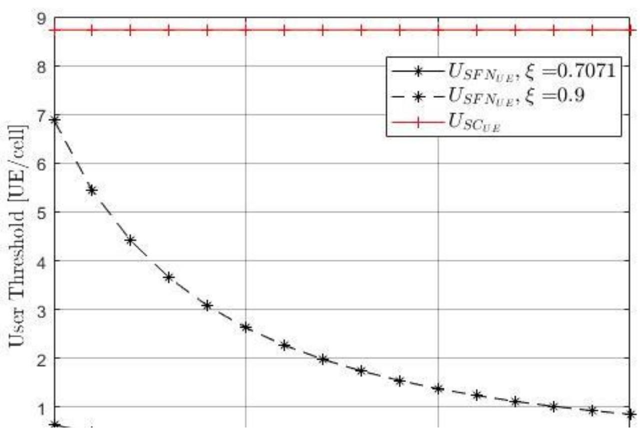

Figure 6 - User threshold to switch from unicast to MBSFN or SC-PTM in order to reduce energy consumption in the BS or UE.

The user thresholds to switch from UC to MBSFN or SC-PTM in order to reduce energy consumption in the UE are presented in Figure 6. We consider $\lambda_{BS} = 4BS/km^2$ and MBSFN areas ranging from $3.1km^2$ ($R_s = 1km$) up to $19.6km^2$ ($R_s = 2.5km$). The service areas have half the size ($\xi = 0.7071$) and 80% the size ($\xi = 0.9$) of the MBSFN area. First, notice that $U_{SC_{ue}}$ is equal to 8.7 UE/cell. This value does not depend on λ_{UE} or λ_{BS}. On the other hand, the user thresholds for MBSFN decrease with the MBSFN area size. The bigger the MBSFN area, the higher the SINR and therefore the lower the UE energy consumption in MBSFN. We see as well that the size of the service area has an important impact on the user threshold, e.g., if $R_s = 1km$, $U_{SFN_{UE}}$ increases from 0.6 UE/cell to 6.9 UE/cell when increasing ξ from 0.7071 to 0.9. This is due to the higher interference perceived by the users at the border of the service area.

In conclusion, the analytical approach proved that the user threshold to reduce UE energy consumption is fairly low in most cases. Furthermore, results show that MBSFN helps to reduce UE energy consumption in a more effective way than SC-PTM since the user thresholds for MBSFN are lower.

THE IMPACT OF VIDEO DECODING

This part of the study was carried out with the Android application implemented within the NESTED project. This application embeds the Viaccess-Orca Software Development Kit (SDK) player for video content playback. This same SDK integrates the OpenVVC library developed by NESTED project partner French research institute IETR (Institut d'Electronique et des Technologies du numéRique) for the decoding of VVC streams (software decoding), while it relies on the chipset capacity for HEVC decoding (hardware decoding). The test's objective is to estimate the application's power consumption over each 30-minute video sequence.

Test Materials and Description

Each test was done on the same device (Samsung S20FE 5G - Qualcomm Snapdragon 865 - 128 GB) by playing the same video sequence encoded by NESTED project partner Ateme's Titan Live equipment using different profiles in HEVC and in VVC standards and packaged in DASH. The Android application is consuming video from an origin server on a local network.

The device's battery is fully charged at the beginning of each test, and battery statistics are reset.

ID	Codec	Resolution	Frame rate	Bit rate
OLS.HEVC.720p.3M	HEVC	720p	50 fps	3 Mbps
OLS.HEVC.720p.4M	HEVC	720p	50 fps	4 Mbps
OLS.HEVC.720p.5M	HEVC	720p	50 fps	5 Mbps
OLS.HEVC.1080p.5M	HEVC	1080p	50 fps	5 Mbps
OLS.VVC.720p.3M	VVC	720p	50 fps	3 Mbps

OLS.VVC.720p.4M	VVC	720p	50 fps	4 Mbps
OLS.VVC.720p.5M	VVC	720p	50 fps	5 Mbps
OLS.VVC.1080p.5M	VVC	1080p	50 fps	5 Mbps

Table 2 - Test video sequences

Protocol and Measurement Tool

We use Battery Historian tool along with Android Debug Bridge (adb) *dumpsys* commands. The Battery Historian tool provides insight into a device's battery consumption over time. It visualizes power-related events from the system logs in an HTML representation.

For each video sequence, we run the following test protocol:

- Reset all battery statistics using *adb shell dumpsys batterystats --reset* command
- Unplug the device and launch the player to play a video sequence
- At the end of the playback, plug in the device to get battery statistics using *adb shell dumpsys batterystats* and *adb bugreport* commands
- Open up the dumpsys zip file into Battery Historian web application
- Filter out with the player app
- Build a custom graph with the following parameters: battery level, temperature, top app and voltage.

Then for each battery level stage, we compute the power consumption using $P = U \times I$ law with the battery discharge rate and the average battery voltage (4200mV). This computed power consumption is then reported to one hour and compared with other test sequences. Table 3 below summarizes the test results of a "codec-centric" comparison (i.e., HEVC vs VVC) and of "content quality-centric" comparisons (different profiles for a given codec).

For the codec-centric comparison, we compare VVC 1080p 5M sequence to a HEVC 1080p 5M reference sequence. Then for the HEVC content quality-centric comparison, we compare all HEVC sequences to a HEVC 720p 3M reference sequence. Finally, for the VVC content quality-centric comparison, we compare all VVC sequences to a VVC 720p 3M reference sequence.

Results

Test stream ID	Reference stream ID	Duration	Energy ratio
1 - CODEC CENTRIC			
OLS.VVC.1080p.5M	OLS.HEVC.1080p.5M	30 min	2.28
2 - HEVC CONTENT QUALITY CENTRIC			
OLS.HEVC.720p.4M	OLS.HEVC.720p.3M	30 min	1.00
OLS.HEVC.720p.5M	OLS.HEVC.720p.3M	30 min	1.01
OLS.HEVC.1080p.5M	OLS.HEVC.720p.3M	30 min	1.02

3 - VVC CONTENT QUALITY CENTRIC			
OLS.VVC.720p.4M	OLS.VVC.720p.3M	30 min	1.17
OLS.VVC.720p.5M	OLS.VVC.720p.3M	30 min	1.17
OLS.VVC.1080p.5M	OLS.VVC.720p.3M	30 min	1.50

Table 3 - Energy ratio results

As a conclusion, we notice that a VVC software decoder consumes two times more energy than a HEVC hardware decoder. While this order of magnitude seems relatively aligned with other such codec studies (9), the ratio view can be nuanced with the global power consumption view of the device app. Indeed, for 10 hours of video streaming playback, our measurements show that the overconsumption generated by VVC software decoding compared with HEVC hardware decoding is in the same range as one hour of a low energy bulb consumption.

Furthermore, we notice that increasing quality has barely any impact on HEVC hardware decoder consumption, whereas for a VVC software decoder the impact is not negligible.

CASE STUDY: PREMIUM CONTENT CONSUMPTION OVER 5G IN A STADIUM

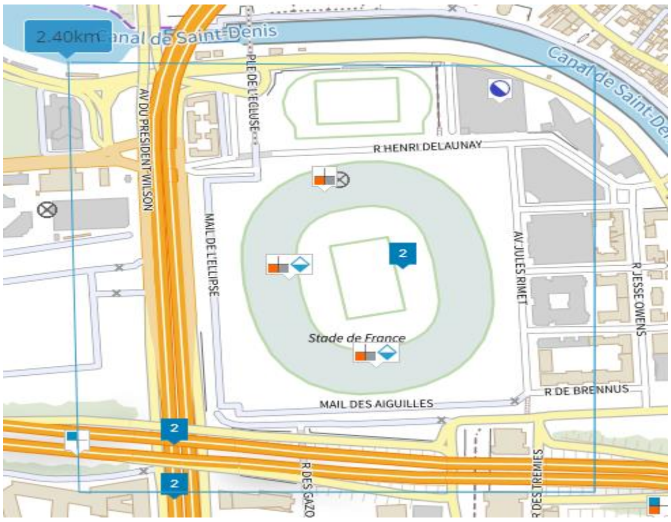

Figure 7 - Stade de France in Cartoradio

Using *The Stade de France* as an example, the stadium has a surface of approximately 0.4 km². This is the size of the service area. To adjust to our model, $\xi R_s = 0.36km$. We can consider two scenarios: $\xi = 0.7071$ and $\xi = 0.9$. It is, R_s values of $R_s = 0.36/0.7071 = 0.51km$ and $R_s = 0.36/0.9 = 0.4km$. Therefore, we consider two MBSFN area sizes, 0.82 km² and 0.5 km²

Since BS are located following a PPP, the probability of not finding a BS inside the MBSFN area A is given by $e^{-\lambda_{BS}A}$. Considering $A = 0.5$, the lowest value of λ_{BS} that we can consider to guarantee (with 95% certainty) that there is at least one BS in the MBSFN area, and therefore the correctness of the results, is $\lambda_{BS} = -ln(1-0.95)/0.5 = 6$ BS/km². Then, we consider $\lambda_{BS} \geq 6$ BS/km². *The Stade de France* can receive at most 90,000 spectators. We assume that at some point 1% of them want to watch some highlights of the game (e.g., replay of a goal) at the same time on their smartphones, which represents 900 UE. This is equivalent to a user density $\lambda_{UE} = 900/0.4 = 2250$ UE/km².

Under these assumptions we can provide a theoretical estimation of the energy savings provided by broadcast transmission on the UE side (Network/Transport and Physical/Radio layers). Results are shown in Table 4. We can appreciate how MBSFN and SCPTM help reduce UE energy consumption in almost all cases. Only when the service area is 80% the size of the MBSFN area and the BS density is low, UC performs better than MBSFN.

With regards to the video content layer, considering the Android application, we can compare what will be the power consumption of all devices consuming the video encoded in HEVC 1080p at 5 Mbps to all devices consuming the video encoded in VVC 1080p at 5 Mbps. This comparison is presented in Table 5.

As a reference, we notice that, in terms of energy consumption, 900 smartphones consuming the VVC video reference stream during one hour approximately corresponds to powering an OLED UHD TV set for a full day.

MBSFN Area [km²]	**BS Density [BS/km²]**	$E_{UE_{UC}}/E_{UE_{SFN}}$	$E_{UE_{UC}}/E_{UE_{SC}}$	$U_{SFN_{UE}}$	$U_{SC_{UE}}$
0.5	6	0.0098	42.84	3×10^4	8.73
	8	3.68	32.12	61.76	8.73
	12	7.73	21.418	19.62	8.73
	16	9.14	16	12.43	8.73
0.82	6	48	42.8	3.9	8.73
	8	70.16	32.12	2	8.73
	12	92.92	21.41	1	8.73
	16	122.7	16.06	0.6	8.73

Table 4 - Energy savings results

Test description	**Test stream ID**	**Total power consumption (Wh)**
HEVC reference stream playback	OLS.HEVC.1080p.5M	1,300
VVC reference stream playback	OLS.VVC.1080p.5M	3,000

Table 5 - Power consumption for reference streams

CONCLUSION

This paper presented a set of results, which constitutes a first milestone toward a more comprehensive attempt to measure (and ultimately try to optimize) the environmental impact of video streaming, with a clear focus on the endpoint. The paper presented the results of an analytical approach to calculate the endpoint energy consumption related to the network/transport and physical/radio technology layers. Our case study clearly demonstrates the advantage of MBSFN and SC-PTM technologies compared with unicast in scenarios where multiple users request the same content at the same time (energy reduction of up to 122 times). We believe these theoretical results now need to be reinforced by empirical measurements.

With regards to the video content technology layer, our empirical measurements demonstrate that the choice of a codec is much more impactful in terms of environmental impact than a mere change in content quality/encoding profile. However, in absolute numbers, a software implementation of a next-generation codec (such as VVC) does not represent a critical degradation in terms of carbon footprint compared with hardware implementations of legacy codecs (such as HEVC). This could advocate for the deployment of services leveraging software implementations of next-generation codecs, even without waiting for the deployment of a critical mass of next-generation smartphones (i.e., hardware decoding-enabled. Delaying the replacement of a generation of smartphones can in itself be a huge source of energy savings. To be comprehensive, these codec considerations cannot neglect the potential impact on the perceived quality of experience. To what extent does reducing the carbon footprint of video streaming impact the perceived QoE? Additionally, each one of the technology layers considered will require a wider analysis that must encompass the head-end, on top of the sole endpoint. The authors hope they will be able to provide additional insights on these directions soon, thereby shedding some light on the complex trade-offs video service providers will be facing, most likely at the behest of a customer base increasingly sensitive to sustainability aspects.

REFERENCES

1. CISCO, VNI Complete Forecast Highlights, 2021.

2. https://www.orange.com/en/oranges-commitment-environment

3. https://www.ateme.com/press/ateme-leads-nested-5g-consortium-to-provide-energy-efficient-solution-for-streaming-over-5g/

4. Chaoan Wu, Xuekang Sun, and Tiankui Zhang. “A Power-Saving Scheduling Algorithm for Mixed Multicast and Unicast Traffic in MBSFN,” 2012 Computing, Communications and Applications Conference, Jan. 2012, pp. 170–174.

5. Sara Pizzi et al., “Energy-Efficient Multicast Service Delivery Exploiting Single Frequency Device-To-Device Communications in 5G New Radio Systems,” Sensors 18.7, July 9, 2018.

6. Juan Vargas, Cedric Thienot, Christophe Burdinat and Xavier Lagrange, “Broadcast-Multicast Single Frequency Network versus Unicast in Cellular Systems,” 2020 16th

International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob), Oct. 2020, pp. 1–6.

7. Juan Vargas, Cedric Thienot, and Xavier Lagrange, “Analytical Calculation of the User Threshold for the Switching between Unicast and Broadcast in Cellular Networks,” IEEE/IFIP Network Operations and Management Symposium (NOMS), April 2022.

8. Juan Vargas, Cedric Thienot, and Xavier Lagrange, “MBSFN and SC-PTM as Solutions to Reduce Energy Consumption in Cellular Networks,” 27th IEEE Symposium on Computers and Communications (ISCC 2022), June 2022.

9. https://youtu.be/oGSXDuNtkiU?t=6235

ACKNOWLEDGMENTS

The authors would like to thank the Ateme, IETR, and Orange teams involved in the NESTED project for their important contributions to their work, and in particular Ateme’s T. Biatek, C. Burdinat and M. Raulet, Orange’s S. Kervadec, J. Lemotheux, P. Perrot and P. Crepieux and IETR’s PL. Cabarat and W. Hamidouche.