

CONVINcE: Towards Power-Optimized Video Distribution Networks

Raoul Monnier¹, Adrian Popescu^{2,*} and Rickard Ljung³

¹Thomson Video Networks, 33571 Cesson Sévigné, France

²Blekinge Institute of Technology, Faculty of Computing, 371 79 Karlskrona, Sweden

³Sony Mobile Communications, Mobilvägen 4, 221 88 Lund, Sweden

*Corresponding author: adrian.popescu@bth.se

Abstract—CONVINcE is a CELTIC-PLUS project dedicated to minimizing the power consumption in IP-based video distribution networks, from the headend to the terminal. The entire video distribution chain is considered in the project, covering a wide range of entities involved in this process. Examples of these entities are headend, edge cloud, Content Distribution Network (CDN), core backbone network, Radio Access Network (RAN) as well as fixed and mobile terminals. Related to this, one of the most difficult research questions is regarding the provision of minimum end-to-end power consumption for video streams combined with the best possible Quality of Experience (QoE) obtained at the terminal. It requires solving a number of sophisticated research questions, among them modelling and optimization problems.

The paper first provides an overview of video distribution networks, which is followed by a short presentation of the CONVINcE project. The second part is focused on the optimization concept suggested to provide both minimum end-to-end power consumption and the best possible QoE at the terminal.

Keywords—IP-based video distribution networks; energy saving; performance optimization

I. INTRODUCTION

The creation, distribution and delivery of video contents are a multi-dimensional process that refers to the acquisition of video sources, content production and packaging as well as the distribution to customers. IP-based video networks are one category of distribution networks used for the transfer of video signals of different categories. Today, the Internet is under very high pressure to satisfy huge demands for bandwidth increase. This is because of factors like exponential growth in video demands and the appearance of new bandwidth-demanding video streaming services such as Netflix. This process is further complicated because the Internet has also democratized the process of creation, distribution and sharing of user-generated video contents by services like YouTube. Finally, one should also mention that the situation today is even more aggravated by the adoption of Ultra High Definition (UHD) video formats needing more bandwidth.

To cope with the large demand for bandwidth growth, the shift to wireless communications as well as to solve other issues related to the current Internet (e.g., security), new architectures for the future Internet have been proposed and prototyped. These are architectures designed to support a wide range of services, but where no particular focus is laid on video distribution and solving the particular problems associated with this. Examples include Content-Centric Networks (CCN) [3] or content-based extensions to Software-Defined Networking (SDN) [4]. The conclusion is that there is a need for sustained research and development activities on IP-based video distribution networks. Questions like how well the future Internet architecture facilitates video delivery, how to provide real-time guarantees for live and interactive video streams, what are the subjective and objective metrics for performance measurement and how to reduce the energy associated with video distribution, need to be answered.

The research project “Consumption Optimization in Video NETworks (CONVINcE)” is a three years CELTIC-PLUS project. The goal is to study and to analyze the energy consumption in IP-based video distribution networks in order to propose solutions to decrease this consumption.

The paper is organized as follows. Section 2 is dedicated to a short presentation of video distribution networks and the problems associated with this. Section 3 presents the CONVINcE project. Section 4 presents the high-level architecture model of CONVINcE. Based on that, Section 5 reports a new method suggested for performance optimization in IP-based video distribution networks. Finally, Section 6 concludes the paper.

II. VIDEO DISTRIBUTION NETWORKS

Two categories of video distribution networks can be distinguished. The first one regroups the broadcast networks addressing simultaneously millions of end users. These are terrestrial, satellite and, to some extent¹, cable television networks. The second category, broadband networks, is based on the IP technology. The broadband networks are

¹ Cable networks carry both Broadcast signals (e.g. DVB-C technology) and broadband signals (e.g. DOCSIS technology)

targeting individual users, even if, in some cases, multicast and broadcast mechanisms can turn them into networks addressing several users at a time.

Broadcast networks are, due to their topology – a single transmitter for millions of receivers – the most power-efficient way to distribute video to a large amount of people. Broadband networks are more flexible and allow more easily the deployment of new video services, especially if the Internet is the basis for the deployment of these services. Unfortunately, they can also be considered as a victim of their own success: some researchers recently warned that the Internet could collapse: “The web could reach its limit in just eight years and use all of Britain’s power supply by 2035” [5]. There is thus urgency to propose solutions to reduce the consumption of IP-based networks.

IP-based video distribution networks started with IPTV services. It is well known that IP protocols provide a “best-effort” service, where the Quality of Service (QoS) is far from the QoS of a video broadcast network (in terms of high availability, constant bandwidth and delay). To offer customers a QoS close to this one, Telcos have built managed video-over-IP networks where one can control the available bitrate, the maximum delay and jitter.

Later on, the “Over The Top” (OTT) concept emerged to provide, in a first step, Video On Demand (VOD) services. Here, there is no dedicated broadband network for the video service. As of today, many OTT services are using the Internet together with Adaptive Bit Rate (ABR) protocols to align the quality of the video to the available instantaneous bandwidth. These services have become very popular, and famous OTT service providers such as YouTube and Netflix have appeared. The consequence was that, in order to satisfy millions of users, it was necessary to increase the network throughput but also to build Content Distribution Networks (CDN) consisting of many servers cleverly distributed over countries to facilitate simultaneous access to the content. All that was done at the price of increasing the power consumption of the video networks!

IP-based video distribution networks are today facing the challenge of introduction of UHD services. UHD does not mean the sole increase of the picture resolution: “4K”, which is 4 times the number of pixels of a High Definition (HD) video, or “8K” which is the format Japan intends to introduce in 2020 with 16 times the number of pixels of an HD video. The UHD project also advances improvements in other dimensions. The first one is the temporal dimension with Higher Frame Rates (HFR): 100 or 120 frames per second. A better rendering of the luminance (darker black and brighter white) is also offered as well as a better representation of the colours: Wide Colour Gamut (WCG). As a result of the introduction of UHD services, there is a significant increase of the bitrates, especially for high resolution videos. Fortunately, progress in video coding tempers this increase. For instance, High Efficient Video Coding (HEVC) allows, at a constant picture quality, in the order of 50% improvement of the bitrate versus the State-Of-The-Art Advanced Video coding (AVC), even more for

very high resolutions [6]. Additional gains are even expected further to researches in the area of video coding.

III. CONVINCe


The concept of the CONVINCe project grew out of two observations: the Internet’s carbon footprint is expected to exceed the one of air travel by a factor of two by 2020 [7] and also that the Internet traffic is driven by video [1]. The sum of all forms of IP video will be 80% of total IP traffic in 2019. Another important element is the economic aspect of the problem: the price of electricity is increasing and will still increase in the coming years.

CONVINCe addresses the challenge of reducing the power consumption in IP-based video networks with an end-to-end approach, from the headend, where contents are encoded and streamed, to the terminals where they are consumed, embracing the Content Delivery Networks (CDN) and the core and access networks. IP-based networks are the target of the project, but some results on headends and terminals can however be applied to broadcast networks as well.

The project tackles a wide range of areas and technologies, encompassing the following components:

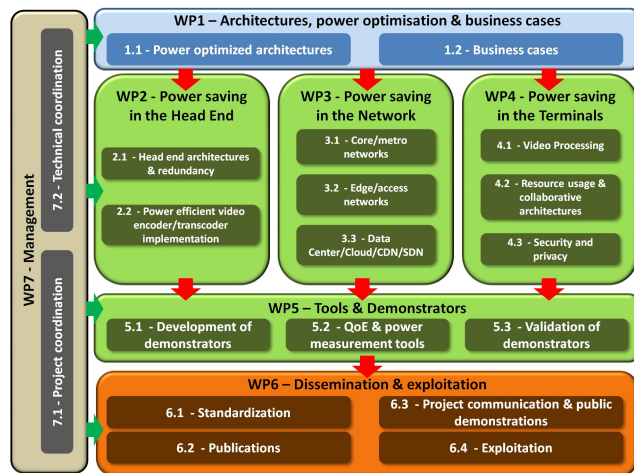
- Video encoding and transcoding,
- Adaptive bit streaming,
- Core/metro networks,
- Access networks,
- Content Delivery Networks (CDN),
- Routing protocols,
- Fixed and mobile terminals.

Partners’ efforts concentrate on architectures, hardware and software design, protocols and basic technologies in the devices. In parallel to the activities focused on optimizing the power consumption in each part of the system, the project runs transversal activities on software best practices and eco-design and also on power & QoE measurements. The project also considers the use of new technologies recently suggested for energy savings in the form of Software Defined Networking (SDN) and Network Function Virtualization (NFV).



End-to-end approach of the CONVINCe project

The CONVINCe project is built on three strong pillars (work packages) investigating power saving in the headend, in the networks and in the terminals. In order to ensure an end-to-end approach for energy saving, another work package is fully dedicated to system architecture, power optimization & related business cases. Results coming from these four work packages are providing inputs to a fifth work package in charge of building demonstrators and developing new tools for QoE and power measurements. Finally, results of the project are disseminated and exploited through standardization, publication and demonstration activities.



Organization of the CONVINCe project

CONVINCe is a three years CELTIC-PLUS³ project labelled in November 2013 and kicked-off in September 2014. The consortium is composed of fifteen partners across four countries (France, Finland, Sweden and Turkey). It is coordinated by Thomson Video Networks.

After one year and a half of work, the first results of the project are emerging, showing substantial energy savings in several areas such as:

- Video transcoding: 60% saving thanks to the use of GPU-assisted CPUs,
- Wireless multimedia sensors: 50% saving with a two-tier architecture approach,
- SDN-based data centers: 40-45% gains, up to 69%, with geographical load balancing,
- LCD displays: 15% saving thanks to LED bar optimization and light guide plate design.

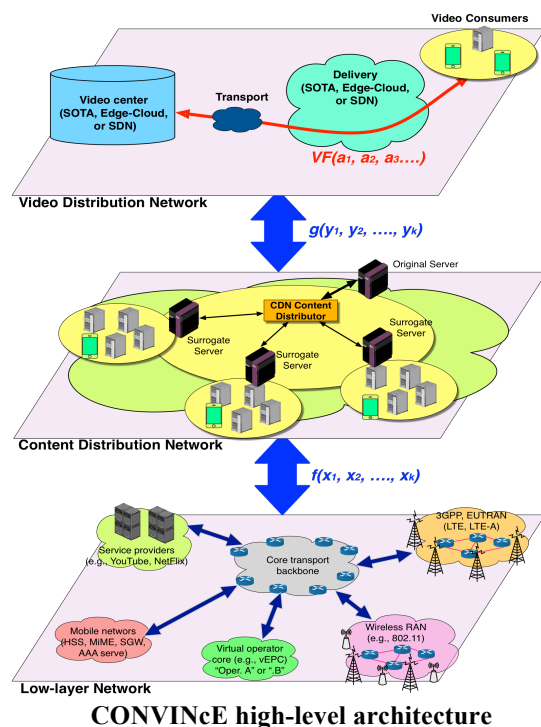
In the second part of the project, while carrying on their work on optimizing the consumption of individual pieces of the system, partners will work together to integrate their improvements in demonstrators in order to show end-to-end power savings, which will not be just the addition of individual savings. Partners are confident in the global

³ More information on CELTIC-PLUS : <https://www.celticplus.eu/>

energy saving of some architectural approaches, which are studied in the project, such as the edge cloud architecture.

IV. HIGH-LEVEL ARCHITECTURE

High-level architecture of a given system is defined to be a general-purpose architecture that describes the system considered for study. The diversity of video streaming network solutions and components indicates the complexity existing in the design and development of architectural solutions for CONVINCe and the associated optimization solutions. To alleviate such complexity, a three layers high-level architecture is used, as shown in the figure below: Video Distribution Network (VDN), Content Distribution Network (CDN) and low-layer network (LLN), which refers to networking elements at layers L2 (Data Link Layer) and L3 (Network Layer) in the TCP/IP protocol stack.



Three categories of low-layer network architectural solutions are considered in CONVINCe. These are the non-cloud based architecture, the edge-cloud based architecture and the SDN/NFV-based architecture. The non-cloud based architecture is the State-Of-The Art (SOTA) architecture and it is used as the reference to measure improvements in power consumption. It relies on a classical headend making use of dedicated equipment for encoding and transcoding. Given the popularity of architectures existing today, in the first phase of the project, the focus is laid on the edge-cloud based architecture. SDN/NFV-based architecture will be studied in a second step.

Four application scenarios are considered in CONVINCe: On-Demand Video Streaming; Live Video Streaming; Cloud Gaming and Camera-Based Sensor Network for video surveillance.

VDN is highly related to the formal representation of video flows considered in different CONVINCe application scenarios. A video flow is defined to be an end-to-end (e2e) data transmission from the video streaming source side (e.g., headend) to the video consumer side (e.g., terminal). This process is highly related to the application scenario and it can be described as a function dependent on elements like the application (e.g. On-Demand Video Streaming, Live Video Streaming), the network solution (e.g. edge-cloud based, SDN) and video flow elements (e.g. Video Data Center, consumer).

CDN is basically a large system of geographically distributed specialized cache servers deployed near access providers in Data Centers across the Internet, which are combined with the use of a special routing code to redirect the media requests to the closest servers. High-level network intelligence is used to improve the performance of delivering media content over the Internet. Requests for content are typically directed to servers that are optimal in some way. Media delivery can accordingly be done with reference to different parameters used in the process of optimization like highest possible throughput, minimum power, minimum cost.

The concept of distributing the content to cache servers located close to end users results in better performance with regard to maximizing the bandwidth, minimizing the latency and jitter, improving accessibility as well as better balance between the cost for content providers and the QoE for customers. The benefits of using CDN are substantial to end-users, content and application owners as well as network to service providers.

Finally, another important architectural element is the coordinating middleware. The importance of middleware is that it coordinates the operation of individual architectural elements towards minimizing the e2e power consumption combined with the best possible QoE performance obtained at the end user, like for instance the OpenStack API and the ApacheServiceMix [8, 9]. Multiple Application Programming Interfaces (API) are involved in serving different requests. Given the large number of networking elements involved in the CONVINCe architecture, the number of interacting APIs increases as well as the number of peers that an element may interact with. Furthermore, another serious difficulty is because the format (syntax and semantics) of the APIs is different among elements of different types. This creates the need for the presence of a so-called mediator that transforms the API calls from one format to another and coordinates interactions among different elements. Similar to the requirements regarding minimizing the energy consumption of different architectural elements, the middleware must be designed in a similar way, i.e. with regard to minimizing the energy consumption.

V. PERFORMANCE OPTIMIZATION


Performance optimization in an end-to-end perspective is dependent on the criteria adopted for optimization. On the one hand – specifically for the overall purpose of the CONVINCe project – an optimization in terms of total end-

to-end power consumption would be the selected criteria. Additional aspects need however to be taken into account, specifically in terms of system and end-user performance, the so-called Key Performance Indicator (KPI). These parameters can for instance be in terms of media delivery parameters like average latency and jitter, and also even as important for video distribution, the system data throughput and capacity.

There are different elements of interest to optimize with regard to end-to-end perspective. An end user is expected to be much more interested in parameters like user terminal power consumption, playback latency and delivered data rate. On the other hand, a network operator is expected to mainly consider the total network capacity as well as the Quality of Service (QoS) parameters relevant for the particular networking entities. All in all, one can group a set of different KPIs into an end-to-end quality of experience metric.

Therefore, the performance optimization ends up in a trade-off, where the end-to-end power consumption can be optimized for a given set of quality of service and quality of experience metrics.

In the following, a particular networking scenario is considered to present the CONVINCe optimization concept. We consider a network chain composed of a headend, a core backbone network, an edge cloud, a wireless Radio Access Network (RAN) and a mobile terminal.



Typical networking scenario

Given a particular networking scenario, it has been observed that the most performance limiting networking segment is the wireless LAN terminal. This regards both the minimum throughput in the end-to-end networking chain as well as the largest power consumption compared to other networking entities [10].

The procedure for performance optimization is as follows:

- Define the block diagram and the component elements like access (eNodeB, HGW, ...), backhaul and transport (switching, ..), core network (element1, ..) and servers (Data Center, ..),
- Model and measure the power consumption of each element in a steady state and also when serving video flows,
- Model and measure the throughput provided by each network element,
- Select the network element(s) with the minimum throughput/channel in the end-to-end network chain,

- Do performance optimization in terms of both (terminal) QoE and minimum end-to-end energy consumption.

There are four scenarios considered in CONVINCe. The constrained optimization problem for the particular case of a Live Video Streaming scenario is as follows.

Given CONVINCe edge-cloud architecture with the network elements: headend, core backbone network, edge cloud, RAN and terminal,

Do Power minimization for individual entities in terms of relevant parameters: number of nodes weighted by the number of active CPUs (Edge Cloud); best route with regard to optimum packet forwarding (Internet); maximum ratio RF output power/total AC input power (RAN); minimum number of replicating proxies (edge cloud); minimum power amplifier consumption (HFC); minimum power consumption and maximum QoE (terminal),

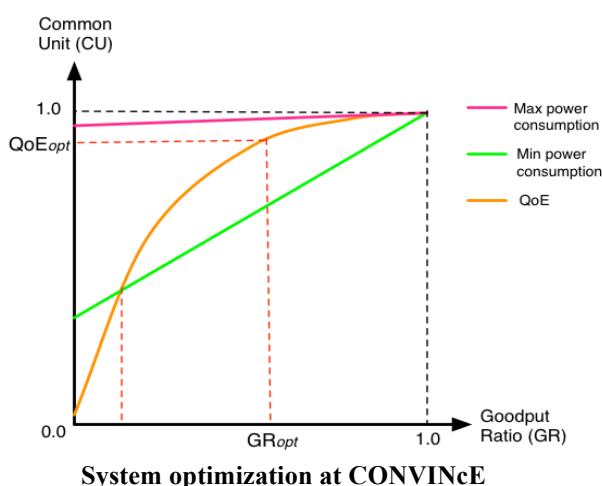
Also Measure the e2e throughput/channel for every individual entity in the e2e network chain,

Select the individual entity that shows the minimal e2e throughput/channel and denominate this throughput as being the system throughput for the whole system,

Optimize the performance in terms of minimum e2e power consumption and best QoE at terminal, under the condition of given system throughput.

Similar constrained optimization problems can be defined for the other scenarios as well.

The figure below shows the concept of system optimization for CONVINCe.



The y-axis represents the so-called parameter Common Unit (CU). This is an abstract parameter, which is used to represent both QoE and Power parameters, but with reference to their respective maximum values. In other words, this means that both parameters have a maximum value of 1.0 on the CU axis. Furthermore, the x-axis represents the so-called goodput ratio (GR), which is used to represent the channel throughput of the particular system, with reference to the maximum value of throughput. This further means that this parameter has a maximum value of 1.0 on the GR axis.

With regard to the figure above, the performance optimization problem is defined as follows.

Given CONVINCe edge-cloud architecture with the network elements: headend, core backbone network, edge cloud, RAN and terminal,

Do Power minimization for individual entities,

Also Measure the e2e throughput/channel for every individual entity in the e2e network chain; also measure the QoE at the terminal,

Select the individual entity that shows the minimal e2e throughput/channel and denominate this throughput as being the system throughput for the whole system,

Plot the function $CU = f(GR)$,

Select the optimal value for this function; select accordingly the values for QoE_{opt} and GR_{opt} ,

Subject to Specific individual requirements for power minimization of entities in terms of, e.g., minimum power consumption for headend \cap minimum power consumption in Edge Cloud \cap best route (Internet) \cap minimum power consumption in RAN.

The question of selecting the optimal value for the function $CU = f(GR)$ is left open for future investigations and researches. The goal for QoE is towards $GR = 1$, whereas the goal for power consumption is towards $GR = \min$. The figure above shows the particular case of using a simple arithmetic average for the value of GR_{opt} ($GR_{opt} = 0.5$) in the case of minimal power consumption (green line). Based on this, a particular value is obtained for QoE_{opt} , which is acceptable, as observed in the figure. QoE is almost maximum in this case, which corresponds to $CU = 1$. Furthermore, it is also observed that in the case of maximum power consumption (red line) the optimal solution is in form of $GR_{opt} = 1$.

The analytical results will be further validated by future measurement experiments done in the last segment: RAN to terminal.

VI. CONCLUSIONS

The paper introduced the CELTIC-PLUS project called CONVINCe where the goal is to minimize the power consumption in IP-based video distribution networks, from the headend to the terminal. The entire video distribution chain is considered in the project, which covers a wide range of entities involved in the process of video distribution. The first part of the paper is dedicated to explain the motivation for the CONVINCe project. A short presentation of the project is done and, connected with this, the high-level architecture for CONVINCe is described. Finally, the last part of the paper describes the optimization concept suggested to obtain minimum end-to-end power consumption and best QoE performance at the end-user terminal.

Future theoretical work in the project is about developing and testing the optimization algorithms as well as developing the associated theoretical models for CONVINCe.

ACKNOWLEDGMENT

This work is supported by the European Celtic-Plus project CONVINCe and was partially funded by Finland, France, Sweden and Turkey

REFERENCES

- [1] CISCO, "Cisco Visual Networking Index: Forecast and Methodology, 2014.2019", http://www.cisco.com/c/en/us/solutions/collateral/service-provider/ip-ngn-ip-next-generation-network/white_paper_c11-481360.html
- [2] CONVINCe: Consumption Optimization in Video Networks, CELTIC-PLUS project, https://bscw.celticplus.eu/pub/bscw.cgi/d985/CONVINCe-start_lq.pdf
- [3] Kurose J., Technical Perspective: Content-Centric Networking, Communications of the ACM, Vol. 55, No. 1
- [4] Software Defined Networks, <http://sdn.ieee.org>
- [5] Is the internet on the brink of collapse? <http://www.dailymail.co.uk/sciencetech/article-3064915/The-Internet-reach-limit-just-eight-years-warn-engineers.html#ixzz3yIBxo681>
- [6] Video Quality Evaluation Methodology and Verification Testing of HEVC Compression Performance <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=7254155>
- [7] ICT Sector Accounts for 2 Percent of Global Carbon Emissions: <http://www.natureworldnews.com/articles/467/20130107/ict-sector-account-2-percent-global-carbon.htm>
- [8] OPENStack API, <http://developer.openstack.org/api-ref.html>
- [9] Apache ServiceMix, <http://servicemix.apache.org>
- [10] Veerecken W. and Teo Y.M., "Power Consumption in Telecommunication Networks: Overview and Reduction Strategies", IEEE Communications Magazine, June 2011