

Towards Energy Benchmarking for Green Video

Oche Ejembi

School of Computer Science
University of St Andrews, Fife, UK
Email: ooe@st-andrews.ac.uk

Saleem N. Bhatti

School of Computer Science
University of St Andrews, Fife, UK
Email: saleem@st-andrews.ac.uk

Abstract—Digital video is responsible for the largest proportion of traffic on the Internet today – upto $\sim 70\%$. However, very little published research has examined the energy impact of this growing traffic type on a global scale (on client systems, servers and in the network). We summarise results and lessons learned from our measurement-based experiments on the energy use of digital video. By providing users with appropriate information and feedback, we could enable changes in user behaviour to save energy during use of digital video. We discuss the ongoing development of our benchmark tool which generates information on energy usage for users or other interested stakeholders.

I. INTRODUCTION

The growing energy usage of ICT systems is a source of concern for several stakeholders – researchers, industry practitioners, governments and end users. Some estimates suggest that the Internet is responsible for up to 4% of the world's CO₂ emissions – similar to proportions due to the Aviation industry. Considerable research has been focussed on improving energy-efficiency and energy-awareness for various aspects of ICT. Meanwhile, digital video is responsible for up to 70% of traffic on the Internet today [1]. Thus, a significant proportion of the global ICT energy usage will be spent on operations related to video playback, such as capture, encoding, transmitting, decoding (playback) and storage. These operations are generally more resource-intensive than other popular ICT use-cases such as text or still-image based applications.

Several studies have investigated energy usage by ICT equipment – for example, focusing on networking devices & systems, servers and data centres. However, very little work has focused on digital video – especially in a *global* context. The energy-related work that does look at video, is usually focused within the context of mobile and/or hand-held devices, for example [2]. The motivation behind this is to prolong the usage of the system with the energy held on the rechargeable battery. Anecdotally, mobile device users will confirm that their batteries will drain noticeably faster while watching video, when compared to other daily uses of their device [3]. This is despite the fact that modern mobile devices will typically be optimised with specialised hardware for efficient video playback. As these mobile device users have a limited power supply on their device, they will often change their behaviour and usage patterns to extend battery life. An interesting way to regard this phenomena is to consider this limited power supply as an implicit *incentive* to save energy.

Non-mobile users of digital video, especially those in the developed world, may not be as concerned about energy use. These users will typically get their electricity supply directly from mains power. However, this means that there may be a lack of awareness and concern about how much energy is being used. As the popularity of video continues to increase, so will the energy cost and any associated carbon emissions. There is the danger that the energy consumption by digital video on these non-mobile systems could become a significant drain on future energy usage. Mains power can be relatively expensive and inconsistent in supply in the developing world. In such regions, the use of video is important due to high levels of (functional) illiteracy, and can have far-reaching societal impacts (also in developed regions) with applications for e-Education, e-Agriculture, e-Health and other e-* applications.

Through our work, we wish to provide insights on the scale of global energy consumption due to video delivered over the Internet. In Section II, we briefly describe our experiments in which we measured possible energy savings by simple application layer modifications. Our results have inspired us to develop the concept for a green benchmarking tool for video, which we present in Section III.

II. EXPERIMENT METHODOLOGY AND EVALUATION

In [4], we performed an empirical investigation of the system resource and associated energy usage of decoding and encoding video when using several popular video codecs and picture sizes (resolutions). Some of the codecs we investigated include FLV, H.264, H.265 (HEVC), VP8 and VP9. We found that there are significant differences in the amount of energy required to decode/playback the same video using these various codecs. For example, for one 1080p, 2-minute video sample, the FLV codec consumed approximately 14 Joules per second of video playback, H.264 consumed ~ 25 J/s and H.265 consumed ~ 54 J/s. While these values may not seem large for a single flow, when they are considered on a global scale, they aggregate very significantly. For instance, YouTube viewers watch 72 billion hours of video annually. If we make a simple Fermi estimate, using the difference between streaming all YouTube video at 1080p FLV and 1080p H.265, we can see that there is a possible difference of 2880 GWh a year – enough energy to power over 750,000 homes in the United Kingdom. Of course in reality, the usage patterns and actual energy use will vary. However, we wish to show the sheer magnitude of the energy usage globally and the possibilities

for significant savings. For future work, more comprehensive studies are needed on a wider range of hardware to get more accurate estimates of energy usage. We briefly discuss how this might be enabled through the use of a platform-independent benchmark and measurement tool in Section II




Fig. 1. Mean system-wide energy usage by decoding two minutes of a test video “Tears of Steel” using various video codecs (from [4]).

The differences in energy usage among codecs are even more pronounced when considering video encoding. Due to several factors, including compression effectiveness and maturity of coding algorithms and software, the amount of time required to complete an encoding task can vary significantly. At default settings, we found that encoding the same video 2-minute, at the same bit rate and picture size can take between 0.2 seconds (using FLV) and 2 hours (using HEVC). Of course, other things being equal, the longer the amount of time required to finish an encode, the more energy is used. These results are perhaps more relevant to video on demand (VoD) providers who are involved in encoding large batches of video into various target formats. However, this could also be relevant to end users involved in real-time video communication, as encoding will take place simultaneously with decoding in this use case.

We have also performed another set of experiments on the same hardware set-up using the popular Netflix service as a case-study. Netflix is a premium VoD service available in around 40 countries, with plans to extend services worldwide to 200 countries by 2017. Netflix is mainly delivered via the Microsoft Silverlight browser plug-in which supports the H.264 codec. In our experiment, we found that savings in energy usage of up to 40% are achievable just by selecting picture quality settings i.e choosing a low quality stream over a high quality one. Netflix offers users an incentive to use low quality by offering cheaper plans for lower quality streams, which also saves energy (though not an intended goal for Netflix) However, it is not clear to the naive user that savings in energy are possible by choosing different quality settings.

From our experiments, we have seen that it is possible to achieve energy savings by making simple adjustments using application layer mechanisms (e.g changing codec, picture size or quality). In reality, users would have to affect these adjustments, directly or through preference settings in their

applications. To enable this, they need to be made aware of their energy usage and possible means to reduce it. They may also need to be given incentives to encourage them to save energy [5].

VoD providers should also have a role to play. We can see from Figure 1, how various codecs have different energy usage profiles. The H.264 codec is widely used today. However, older codecs like FLV and MPEG-4 use slightly less energy. Although impractical, there would be significant savings globally if all of Netflix’s video were encoded in FLV, for example. However, this does not consider the Quality of Experience (QoE) of the video. In the same light, UHD and 4K video offerings (4096 x 2160 picture size) are now being rolled out by major providers like Netflix and YouTube. This large picture size is only supported by newer codecs like H.265 and VP9. We can already see from Figure 1, that these codecs (H.265 especially) consume more energy than all other codecs, even at the same picture size, bitrate and quality. While we have not run any experiments on 4K video, it is not unreasonable to assume that the resource costs could scale linearly as well i.e. 4 times the amount of energy could be required to support these video codecs.

III. ONGOING WORK: A GREEN BENCHMARK FOR VIDEO

The experiments described in Section II, were performed on a small set of hardware configurations. To gain a more comprehensive picture, it is necessary to investigate energy usage on a wider, more diverse array of hardware. This includes other desktop systems, laptops, mobile devices and even smart TVs and set-top boxes. This motivates the need for a cross-platform benchmark tool for video.

A motivating example is the Standard Performance Evaluation Corporation (SPEC) SPECpower_ssj2008¹ application. This is a throughput-based benchmark and assesses power usage in relation to the processing capabilities of a server system. The ssj2008 benchmark is written in Java and places an artificial load on the server by emulating e-Commerce type transactions on a the system under test (SUT).

We present a high-level concept for our cross-platform benchmark tool for video inspired by our experiments, *vEQ-benchmark*. Targeted at client systems initially (client decoding events – playback – for video far outnumber encoding events), the execution of this benchmark will ultimately result in the creation of a system file containing information about resource usage and performance of a SUT for video playback. The benchmark shall measure the performance and resource usage of the SUT while playing back several reference video samples. These reference video samples will be available on a publicly accessible repository. Additionally, the samples would have been transcoded into several different popular formats and resolutions, representative of the most popularly used formats available. Streaming online video content will also be supported, thus incorporating network usage into its benchmarking considerations.

¹http://www.spec.org/power_ssj2008/

The energy usage of the SUT can be retrieved from several sources. An example is an external power meter. SPEC, for example maintain a list of acceptable (but relatively expensive) power analysers for use with their SPECpower_ssj2008. In our experiments, we have used a cheap consumer power meters in our experiments, modified at low cost with a wireless transceiver for data collection. Another source of energy information is battery usage information on some mobile devices. It may also be possible that future systems may have energy monitoring functionality built in, e.g. via the ACPI API ², as is already the case with Apple System Management Control (SMC) on modern Apple computers [5]. We envisage that there will be practical and engineering considerations for system energy measurements (e.g. non-linear drain observed on batteries, accuracy and precision considerations when using various measurement hardware, etc.). However, these can be addressed individually in experiment designs or according to the use-case scenarios of the benchmark. A schematic diagram for the functions and data flows in the tool is in Figure 2.

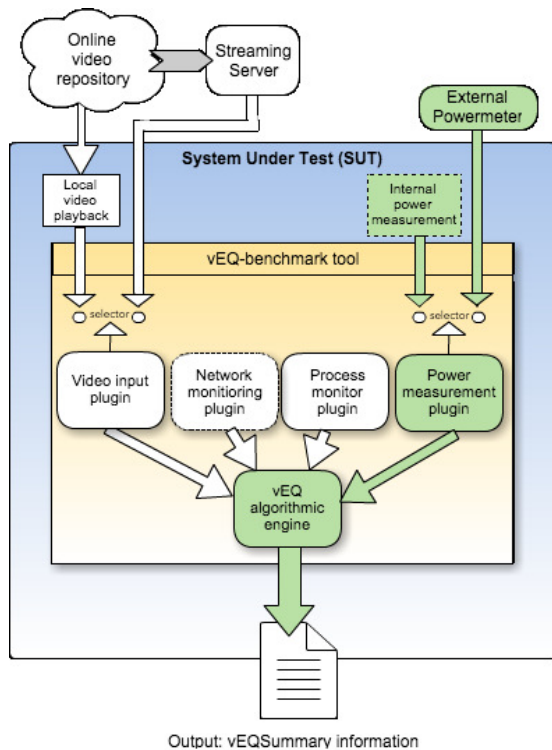


Fig. 2. A high-level schematic diagram for vEQ-benchmark showing the various functional blocks. The arrows depict data/information flows.

The energy and system resource usage of a given configuration will be measured during the benchmarking exercise. This, along with other relevant information will be saved in a summary file which we shall call a vEQ-summary file. This extra information would include details on codec type, hardware availability and QoE metrics. These vEQ- summary files will serve as comparable statistics across different devices and platforms.

²<http://www.uefi.org/acpi/specs>

There will be many benefits from an open benchmarking tool. Firstly, with a wide enough sample set (achievable through crowd-source experiments) we may more accurately assess the worldwide energy usage of video. Secondly, stakeholders (users, manufacturers, policy makers etc.) can assess the suitability and sustainability, of their device for video playback. This could, for example, inform the decision on what kind of hardware to purchase for a given use-case. Alternatively, based on the contents of the vEQ-summary file, a video codec or video component (such as a browser plugin) could, along with information about the video, provide an interface to the user to inform them about the energy usage and other information for a specific video file. For future work we shall consider how to enable automatic adaptation for client systems based on historical benchmark information and current usage, towards enabling energy awareness in digital video usage.

IV. FUTURE WORK AND CONCLUSION

We have presented a brief overview of our work which investigates energy usage and awareness in digital video. We discussed the promising results from our empirical experiments which have inspired us to begin the development of a system benchmark tool for video. We believe that there are several avenues for future research based for energy usage in digital video. One avenue is in investigations into consumer behaviour, and incentivising energy-efficiency for digital video. Other offshoots from our research could optimising or re-engineering video software, applications and hardware for energy efficiency. An industrial application of our studies is possible for Video-on-Demand service producers - making them conscious of the energy usage of their services at both server and client sides, and providing suitable information to users and consumers of video.

ACKNOWLEDGEMENTS

This work was partly supported by the IU-ATC project, funded by grant EP/J016756/1 from the Engineering and Physical Sciences Research Council (EPSRC). Oche Ejembi is funded by the Scottish Informatics and Computer Science Alliance (SICSA). Special thanks to Jim Park (University of St Andrews) for his technical assistance.

REFERENCES

- [1] Cisco Systems, "Cisco Visual Networking Index 2014," Jun 2014. [Online]. Available: <http://goo.gl/3gZy0z>
- [2] S. Ickin, M. Fiedler, and K. Wac, "Demonstrating the stalling events with instantaneous total power consumption in smartphone-based live video streaming," in *SustainIT 2012*, Oct 2012, pp. 1–4.
- [3] D. Ferreira, A. Dey, and V. Kostakos, "Understanding human-smartphone concerns: A study of battery life," in *Pervasive Computing*, ser. Lecture Notes in Computer Science, 2011, vol. 6696, pp. 19–33.
- [4] O. Ejembi and S. N. Bhatti, "Help save the planet: Please do adjust your picture," in *MM2014 - 22nd ACM Intl. Conf. on Multimedia*, Nov 2014.
- [5] Y. Yu and S. N. Bhatti, "The Cost of Virtue: Reward As Well As Feedback Are Required to Reduce User ICT Power Consumption," in *e-Energy 2014 - 5th ACM Intl. Conf. on Future Energy Systems*, Jun 2014.