



TV on the move: How the growth in Internet streaming influences the video quality on your mobile device.

Jason Quinlan

Mobile and Internet Systems Laboratory, Department of Computer Science,
UCC

Any sufficiently advanced technology is indistinguishable from Magic (Arthur C. Clark)

Introduction

Every day, millions of people logon to the Internet to view their favorite TV show on Netflix, or similar streaming services, or to watch the latest viral video on YouTube. Two things are paramount, 1) that they receive the best streaming quality available, and 2) the video starts to play as quickly as possible. There is nothing worse than a video that stops and starts, takes forever to view or constantly changes between viewable qualities (resolutions). Due to our limited download speeds (bandwidth), in most houses it is not uncommon to hear “Stop downloading, I’m trying to watch something on Netflix”.

When we couple this rise in online streaming with the growing number of portable devices (smart phones, tablets, laptops) we see an ever-increasing demand for high-definition online videos while on the move. This demand for mobile streaming highlights the need for adaptive video streaming schemes that can adjust to available bandwidth, where the cellular or Wi-Fi network can limit the quality of the video streaming, and can provide us with graceful changes in video quality, all while increasing our viewing satisfaction. This is the focus of my research. To date my research colleagues and I have developed three new schemes and have a patent pending.

How far we have come

It is a cold Friday night in December 1983, three weeks before Christmas. In the living room, adults and children are huddled round an impossibly large 21” color television, its soft irradiating glow adding to the festivities. To place the size of the TV in context, it is marginally smaller than a Ford Fiesta. The children are waiting for “The Late Late Toy Show” to begin, while the adults are hoping no new toy will be asked for, which would upset the letter already sent to Santa. Happy happy memories... If you were lucky, or

wealthy enough, you had a second TV in an adjacent room, possibly a 14" portable. Again let me also place the 14" portable in context. On average it would take four rather strong children to move this 14" portable, that said it was easier to move the 14" compared to the Ford fiesta in the living room.

In this nostalgic age, there were only two concerns 1) that you did not fall asleep, thus missing the chance to talk with your friends about all the fantastic toys you saw and 2) that the Electrical Supply Board of Ireland was able to cope with 2.5 million kettles being switched on at the same time, during the ad breaks.

Advances in modern Internet and video streaming technologies permitted "The Late Late Toy Show" of 2013 to be streamed live to all corners of the globe. The video received almost 120,000 stream requests in two days as well as over 1.4 million Irish viewers during its broadcast. Meanwhile the benefits of Sky+, TiVo and similar techniques, permit us to pause, record and playback live TV broadcasts. We no longer need to stay awake during the show or overload the national grid during ad breaks. Modern lightweight devices such as smartphones, laptops and tablets now allow us to bring the TV with us as we move, thus introducing an age of mobile video. We can now watch episodes of our favorite TV program when we want, where we want and on what device we want.

But the rise in the number of people using mobile video and the increasing capabilities of mobile devices, soon to reach screen sizes upwards of 13", leads to an ever-growing demand for high quality videos. This equate to large downloads and thus congestion, or blockages, on the network. This ultimately leads to unpleasant video streaming quality. Before we go any further, let's review a little bit of the technology behind video streaming.

Techie bit

Two major technologies assist your device in viewing video streaming. The first is the Internet and associated Wi-Fi networks and the second is the video streaming technology itself.

Internet

The Internet is primarily composed of server computers, connected by a global network of cables and wireless connections, known as links. As the demand on the Internet has increased, the interconnected copper cables have been replaced with super fast fibre optic cables. These cables provide the "backbone" of the Internet and it is these cables that arrive at your home, your business and to the cellular towers you seen strewn over our cities and counties. The wireless links provide the "last hop" from the cable network to

your device. In our homes we use Wi-Fi to connect our devices to the Internet, while we use cellular network technologies like 3G, and soon to be 4G, when we are on the move.

When information or video data is sent from one of the servers to your device, it is normally too large to send all at once. Therefore, the server will cut the data into pieces and each of these pieces will be sent, or transmitted, as a packet. Thus each packet contains a little bit of the information needed by your device to view your requested video. You can think of these packets as dominos. If a domino is missing, then the next domino can't fall until you push it again and it is similar with video packets. If a packet is lost, the information it contains is lost and until your device tells the server to send the packet again, your video will pause and wait until the packet successfully arrives before the video starts to play again.

Finally, like the diesel tank in your car, the number of seats in a plane and the legs on the Christmas turkey, the capacity, or capability, of each of the links on the Internet is limited. Only so many packets can be transmitted on a link at any point in time, which we call the maximum bandwidth of a link. Different links have different capacities. The limit on a copper cable is low, a fibre cable is high, with Wi-Fi and cellular technologies lying in between. Thus, irrespective of the number of servers transmitting packets on a specific link, once the link reaches capacity, known as congestion, the link has no other option than to drop, or discard, packets. If 100 packets will fit on a link at a given point in time, and 110 packets are transmitted, then 10 packets will be lost. Which 10 packets will be lost is unknown, as each link can select the packets to drop at random. Only your device will know if it has lost any packets, as the packets will need to be requested again from the Server.

Video Streaming

Each video clip is nothing more than a collection of images, known as frames, and it is the number of frames per second that provides the illusion of movement. Hence “moving pictures” as it was once know. As we have seen, packets from frames can be lost while being transmitted over the Internet. To counteract the effects of this loss, video streaming introduced the concept of “Group of Pictures” (GOP), which groups a number of adjacent frames together, and treats these as a fixed point in time, such that if one GOP incurs network loss, this loss will not affect any other GOP. This gives your device the option of requesting the lost packets or moving onto the next GOP and ignoring the loss. Neither of these options is beneficial to the person watching the video, as requesting lost packets pauses the video, while moving to the next GOP makes the video very jumpy, plus an important part of the video may be lost. Imagine missing “Luke, I am your . . .”.

My Research

As we have seen the Internet has a finite limit on the bandwidth available to the interconnecting links between servers and devices, while congestion on the links leads to packet loss. The overwhelming growth in video streaming, supported by increasing numbers of portable devices, further exasperates this scenario. Video streaming has introduced a number of mechanisms to recover from the packet loss that occurs, but loss of any kind equals unpleasant video streaming quality. What we need is a way to reduce the quality of a video when loss occurs, thus providing a means of directly reflecting the level of loss in the viewable quality of the video. In addition, when a lot of people are watching the same video clip, as happened with the “Toy Show”, we want to transmit only one stream to all the different devices.




Figure 1: A single SVC stream being transmitted from a cellular tower to three devices each with different capabilities. Image: Jason Quinlan.

To achieve these goals we can use an existing technology called “Scalable Video Coding (SVC)”. An SVC stream is composed of layers, where the viewable quality is dependent on the number of layers received at a device. As illustrated in Figure 1 one SVC stream, which contains three layers, is transmitted over the Internet to three devices connected to a cellular tower. Each device selects a different number of layers dependent on the capabilities of the device or the capacity of their respective link. This will allow each device select the correct quality for both the device and the link, plus it reduces congestion on the network as only one stream is transmitted for all devices rather than one stream for each device.

Unfortunately SVC does not cope very well with packet loss. With the exception of the lowest quality in the stream, called the “base layer”, each of the layers in SVC is dependent on at least one other layer. Thus to increase viewable quality, a device needs to receive a number of layers with no packet loss. As I explained in the Internet section, at present

packet loss occurs randomly and because of this there is no guarantee that a complete layer will arrive at a device.

This is where my research begins. I like all the benefits that SVC provided, but it had a few flaws. As part of my work, I proposed a few changes that better reflected what I saw as the underlying goal of SVC. I changed the way the server cut the data into pieces. In my design each packet would contain a piece of information from every SVC layer and every frame per GOP, thus reducing loss to a little bit from every layer rather than all of the loss from one layer or GOP. I also created a new technique where I added a little bit of data to every layer, which would help us to recover from packet loss. The size of the helper data per layer would be dependent on how important the layer was, i.e. the base layer would have the biggest amount of helper data. This allows viewable quality to be dependent on the quantity of packets lost and not on what was lost.

I hope that these small steps, as well as my other research, will help companies design future streaming technologies that will provide all of us with better quality streaming videos for years to come.

Jason J. Quinlan is a PhD student in the Mobile and Internet Systems Laboratory (MISL), Department of Computer Science under the supervision of Prof. Cormac J. Sreenan and Dr. Ahmed H. Zahran. He would like to acknowledge the support provided by the Science Foundation Ireland (SFI) and by the National Telecommunication Regulation Authority (NTRA) of Egypt. He is forever indebted to his supervisors, Cormac Sreenan and Ahmed Zahran, his MISL colleagues, Ilias, Tony, Nashid, Lau, Paul, Mary, Lanny, Dapong, Hazzaa, Thuy, Xiuchao, Neil, to name but a few, and his family, Teresa and Jack, for their help and support during his journey.