IN THE PRESS

Learnings from Lockdown: Sustainable Streaming

By Maria Ingold, Board Advisor to iSIZE Technologies

Fifty years ago, back when my father built our cabin off-grid by hand, sustainability was called environmentalism and considered hippy, not hip. In that time **global energy consumption** has increased 173%, in an ever-upward trend – until COVID-19. As of the 28th of April, 54% of the global population was in some form of lockdown. **Global energy demand declined 3.8% in Q1 2020**, with full-lockdown countries experiencing an average 25% decline in energy demand per week, and those in partial lockdown 18%.

Paul Massara, former CEO of npower and fellow Board Advisor to **iSIZE**, who deliver machine learning bitrate and energy reduction and perceptual quality enhancement for video, notes that, “At the same time, global carbon use has reduced around 5% as economies have slowed and airplanes have remained grounded. And yet if we are to hit our net zero targets and keep global temperature rises to less than 2%, we require a 7% year on year reduction in carbon, year in year out. The challenge is to achieve such carbon reductions without a crashing of the world economy.”

Impact of video streaming on sustainability

Traffic

Changes during lockdown could result in a new way of working. MIT discovered that 48.7% of US workers worked from home after COVID-19 (14.6% had already been working from home.) Global Workplace Analytics noted that globally 77% of white collar workers were now working from home full time (compared to 9% before.) The percentage of those who would like to work from home at least 1 day a week increased from 31% to 77%. Indeed, companies like Twitter have already approved permanent remote working, which will result in a very different use of transport and need for office space, as well as work-life balance and use of bandwidth.

With COVID-19, all the internet traffic usually across enterprise, education, consumer, public WiFi and, to a lesser extent, mobile and satellite, was suddenly consolidated to a consumer network, monitorable by Sandvine. Sandvine analysed 1 February to 19 April under these unique circumstances. They discovered that globally traffic grew by 38%, with upstream increasing by 121% and then plateauing, and downstream by 23% and rising. It makes sense that upstream grew to accommodate home workers, but steady downstream growth says they're consuming more. Video is still almost 60% of traffic, but Sandvine predicts it could have reached 70% if the major streamers hadn't reduced bitrates as requested by the European Commissioner.

Data Rates

Some customers noticed the reduction in bitrates, especially for AppleTV+ which reduced resolutions to 670 pixels high and compressed heavily with blocky artefacts. Netflix and AppleTV have mostly returned to normal, and YouTube was supposed to stop defaulting to SD after a month.

For an idea of Netflix's typical data usage, Standard Definition (SD) is 0.7 GB per hour (1.6 Mbps), High Definition (HD) is up to 3 GB (6.7 Mbps) and Ultra High Definition (UHD) is 7 GB per hour (15.6 Mbps). Netflix is currently taking up 11.42% of global traffic, which, while a slight decline in share, is still an increase in global traffic. To cope with the overall rise in traffic, Netflix added four times the normal capacity in Internet Service Providers (ISPs) in April.

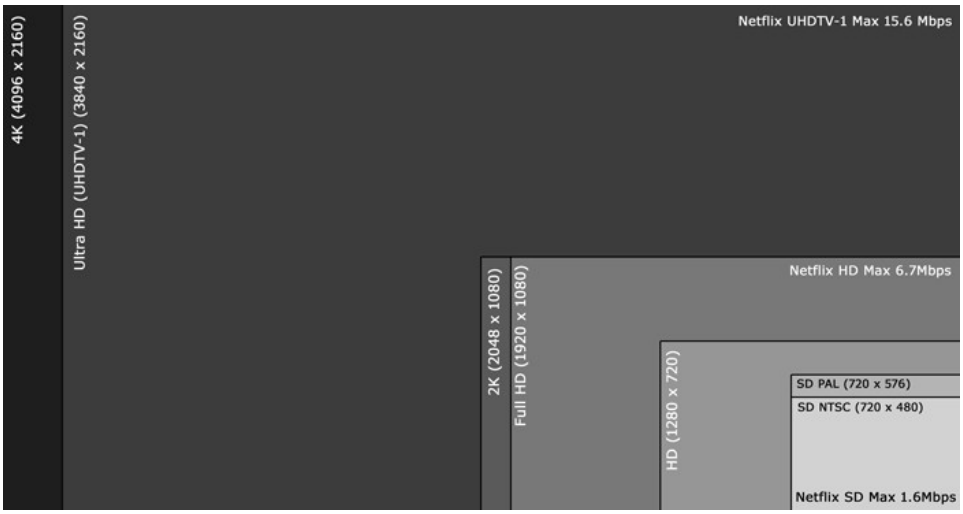

FIGURE 1. Standard resolutions with maximum Netflix data rate.

In 2017 when Netflix had 117.58 million subscribers, its users streamed 140 million hours per day, or 1 hour 11 minutes per day per user. In 2019, 167 million subscribers watched an average of 2 hours per day. As of Q1 2020 Netflix has 182.86 million paid subscribers. 182.86 million people watching two hours per day at 3 GB per hour is 400 exabytes of data per year.

By the beginning of May, mobile video traffic for Disney+ across North America and Europe reached 7 exabytes per month, representing 1.2%-2.2% of all mobile video traffic. Netflix is at 7%-15% which is in the region of 40 exabytes per month. That's 480 exabytes per year just for mobile in North American and Europe.

Given that, 400 exabytes of data per year is low. Perhaps we really are at 3.2 hours of viewing per day as has been projected due to COVID-19. That's 640 exabytes per year, which still seems low.

Data rate	GB per year at 3.2 hours/day	Cost per year at .0025 USD per GB
SD (.7 GB per hour)	149,506,336,000	$373,765,840
HD (3 GB per hour)	640,741,440,000	$1,601,853,600
UHD (7 GB per hour)	1,495,063,360,000	$3,737,658,400

TABLE 1. Netflix data rate, GB per year and cost per year for 182.86 million subscribers watching 3.2 hours per day.

YouTube uses 1.1Mbps for 480p SD, 2.5Mbps for 720p HD, 5Mbps for 1080p HD and 20Mbps for UHD. YouTubers watch 1 billion hours per day, nearly twice the 585 million hours (at 3.2 hours per day) that Netflix users watch.

And UHD isn't the end of the story. YouTube's traffic, at 15.94%, is now more than Netflix. YouTube will also be the first major streaming provider to offer 8K on 8K TVs which support Alliance for Open Media's AV1 hardware decoding. 8K, or UHDTV-2, is 7680 x 4320 pixels.

Energy

Netflix used 451,000 megawatt hours (MWh) in 2019, an 84% increase over 2018, compared to a 20% user growth. 94,000 MWh are direct energy use and composed of their offices, studios and telecoms that form their Content Delivery Network (CDN). 357,000 MWh are indirect energy use which includes partnerships such as Amazon Web Services, Google Cloud and the caching servers they put into ISPs. This will have increased with the quadrupling of capacity in ISPs in April.

The exact sustainability hit of streaming video is under debate, however, with the July 2019 The Shift Project's "Climate Crisis: The Unsustainable Use of Online Video" results being contested by George Kamiya, a digital/energy analyst from the International Energy Agency (IEA) whose article came out in March 2020, a month after Netflix's impact report. He says The Shift Project's figures imply that Netflix streaming consumes 370 terawatt hours (TWh) per year, 800 times higher than what Netflix confirmed above. He further notes The Shift Project's numbers show 1.6Kg CO2e per half-hour of Netflix content, which IEA estimates to be closer to 0.028-0.057Kg CO2e. He says they overestimate bitrate, CDNs and data transmission networks, but underestimate energy consumption of devices.

Contributors to streaming video energy consumption

Preparing

Source content is normally encoded into a high-resolution master format and transcoded into variants. Variants are tailored for the device and service level of the consumer, and typically use "bitrate ladders" of low to high bitrates and resolutions. This could range from SD for a mobile on 3G to UHD for a TV on WiFi.

Video is compressed by an encoder (usually hardware) and decompressed by a decoder (usually software but hardware can be used to optimise). Codecs trade-off between ubiquity, compression (or bits saved) and the time and power required to encode. To further reduce the bitrate and increase perceptual quality, a precoder, like iSIZE's BitSave, may be inserted before the encoder.

Videos are usually stored, and archived, with technical metadata in a Media Asset Management (MAM) system. Metadata, such as title, synopsis, trailers and images are added to a Content Management System (CMS). The MAM and CMS can be combined alongside workflow management and other tools as a Software as a Service (SaaS) with content assets stored in the cloud.

Distributing

Content is usually delivered to a CDN and pushed to its edge servers. Akamai, one of the leading CDNs, has 250,000 edge servers deployed in thousands of locations to cache content at one network hop from 90% of the world's users. All of these servers sit in data centres. In 2018 data centres used about 200TWh, or 1% of global electricity.

Because of the volume of its traffic, Netflix created thousands of Open Connect caches to sit within Internet Service Providers (ISPs) to increase efficiency and reduce the overall demand on upstream network capacity.

Transmission networks, which transmit the bits of data, used about 260 TWh, or 1.1 of global electricity in 2018. Two-thirds of that was mobile networks.

Receiving

The Cisco Annual Internet Report (2018-2023) noted that there will be 29.3 billion networked devices by 2023, up from 18.4 billion in 2018, with 50% being Machine-To-Machine (M2M) connections, growing at 19% CAGR. Smartphones will grow second fastest at 7% CAGR, then Connected TVs and related devices at just under 6% CAGR. PCS continue to decline at 2%.

While smartphones make up the biggest percentage of consumer devices, mobile isn't where all video traffic is consumed. On Netflix, by month six, 70% of viewers watch on TV, 15% on laptops, 10% on mobile phones and 5% on tablets. With over 100 pay TV partnership deals and access to over 300 million global pay TV homes, TV is likely to continue to be important to Netflix.

Solutions to stream video sustainably

Efficient Bits

Encoding, transcoding, storing and distributing fewer bits helps reduce energy use and costs. When the Competition Commission asked for help reducing internet traffic, Netflix stopped using the top rung of each encoding bitrate ladder set up for SD, HD and UHD. YouTube changed its default from HD to SD. AppleTV+ used a highly compressed lower resolution. While these reduced bits, they also reduced quality and some viewers complained.

Keeping the perceptual quality the same or better, while reducing bits, can be achieved in a variety of ways. One is with better codecs. While an older codec like AVC (H.264) has the best availability across devices and platforms, HEVC (H.265) offers a 25-50% bitrate reduction on H.264. AV1 is even newer. It provides a 17% bitrate reduction on H.265 across entire bitrate ladders, but up to 30% for 1080p HD and 43% for UHDTV-1. It's slated to be a gamechanger for 8K encoding. The problem is the speed and power usage. AV1 is said to be 50-3000 times slower than HEVC and requires more powerful hardware to encode. That's being addressed, including with multi-dimensional parallelism (using many CPU cores simultaneously to process multiple parts of the encode). New codecs can take a while, so some improvement methods work with existing codecs. One way is during encoding, as in per-title encoding as created by Netflix in 2015.

Another way is by perceptually optimising the bits before they get to the encoder – precoding. I'm a Board Advisor for iSIZE, who have a codec-independent machine learning precoder called BitSave. iSIZE's precoding provides fewer bits to be encoded, which look perceptually the same, or in some cases better. iSIZE's BitSave combines two key things to reduce complexity, save bits and save energy: preprocessing and dynamic resolution selection. BitSave's consumer version, available as a SaaS service and API, both on bitsave.tech, includes preprocessing, which saves about 30% bitrate on average. The enterprise version, available on a trial basis for B2B users, includes both, so also enables up to a 5-fold (500%) reduction in the energy required by a video encoder.

I explain how BitSave preprocessing works in a previous article on opportunities for video streaming, including an independent analysis by streaming expert Jan Ozer, who finds the technology 'valid and valuable'. In summary, machine learning enhances areas in the frame that are important to the viewer and blurs areas that aren't. This reduces bitrate and improves perceptual metrics like VMAF, while balancing with fidelity metrics like PSNR and SSIM. For Full HD and UHD across a range of encoders – AVC (H.264), HEVC (2.65) and Google's VP9 – iSIZE BitSave preprocessing saves between 8% to 52% (an average of 30%) of the bitrate, and therefore an average of 30% of the energy to store and stream that piece of content.

Dynamic resolution scaling is used on top of preprocessing in the enterprise version. This intelligently downscales the pixel footprint going into any encoder. Some frames don't lose significant information when they are downscaled and then upscaled by a client. While iSIZE do provide an optional upscaler, players already automatically upscale based on the resolution and bitrate information presented in the DASH or HLS manifest file, so existing upscaling can be used with no change to the client.

If the aspect ratio is kept the same but the video's horizontal and vertical resolution is cut in half, then the new frame will only take up a quarter of the pixels. For instance, Full HD is a quarter the size of UHD. A quarter of the original size would result in significant reduction in the CPU cycles and energy consumption required, typically between a 2-fold to 3-fold reduction. And, at the same bitrate, more bits of data would represent each pixel, so quality loss could be mitigated.

To choose the best resolution, each frame is scaled to several resolutions and analysed to determine which one provides the best quality result. Netflix does this, but with two key differences: one, they use a 'brute force' approach to select the resolution, and two, they use linear filters to produce the actual resolution. iSIZE uses neural net filters to produce the optimal resolution, and finds that optimal resolution via a process it calls 'footprinting'. Footprinting does a potentially real time check on the rate and distortion of a set of resolutions, then selects the best using a mathematical optimisation approach. The energy required for footprinting is low and constant, and tests show iSIZE can achieve up to a 5-fold (500%) reduction in encoding time and its associated energy use.

iSIZE's preprocessing and dynamic resolution scaling are explained in detail in iSIZE's peer reviewed journal article, which will appear shortly in the IEEE Transactions on Circuits and Systems for Video Technology (arXiv preprint link).

Consume Consciously

Consumers also have a role to play in reducing their energy consumption. Consuming consciously means using energy-efficient devices, efficient transmissions and reducing electronic waste.

A 50-inch LED television currently consumes five times as much electricity as a laptop and 100 times more than a smartphone. The type of display also effects efficiency. LED-backlit LCD TVs are more energy efficient then plasma TVs. OLED is more efficient than LCD, and microLED is more efficient than OLED. It also depends how old your TV is. Consumer Technology Association (CTA) showed that LCD TVs in 2015 consumed 76% less energy (per screen area) than in 2003. Furthermore, the CTA's sustainability work has reduced American set top box (STB) consumption by 39% since 2012, saving 29 million metric tons of CO2 emissions.

How content gets to the device also has an impact. Wireless and mobile are expected to make up more than 70% of Internet Protocol (IP) traffic in 2022, up from 50% in 2018. Streaming through 4G mobile networks consumes about four times as much electricity as WiFi, but 4G can be more than 50 times as energy efficient as 2G.

Newer devices are generally more energy efficient. The CTA notes that even with a 21% increase of electronic devices in homes there is a 25% reduction in home energy consumption. Production, however, still has an impact. Apple has taken significant steps to make energy-efficient products with renewable or recycled materials and renewable energy, but the production of an iPhone 11 contributes to 79% of its carbon emissions. Use contributes to 17%, transport to 3% and end-of-life processing for less than 1%. Electronic waste is a growing problem too with 50 million tonnes produced each year, amounting to $62.5 billion in valuable materials lost globally. Harvesting would generate fewer CO2 emissions than mining.

Use Renewables

As a result of COVID-19, coal, oil, gas and electricity demand have all dropped, with electricity demand decreasing by 20% or more during full lockdown. Increases in residential demand are strongly outweighed by the reduction of commercial and industrial operations. The impact enabled Britain to shut down its four remaining coal-fired plants in April. Renewables is the most resilient and the only one to see growth in demand – a 1.5% increase across all sectors (up 3% in electricity generation to a nearly 28% share) year-on-year in Q1 2020.

Renewables are likely to remain the only growth area, up 1% across all sectors and 5% in electricity generation in 2020. More wind, hydropower and solar projects are underway, with low operating costs, priority in the grid, and they don't have to adjust output to match demand. So if electricity demand decreases, renewables end up with a higher share in the electricity generation mix. As a result, global CO2 emissions are expected to decline throughout 2020 by 8% to 2010 levels. Unfortunately, recovery from every previous crisis has immediately rebounded CO2 emissions, including the highest ever year-on-year increase in 2010.

Luckily, many large digital companies are reducing their energy and carbon footprints. CDN, Limelight Networks, announced in May 2020 that even with a 50% traffic increase over the last year, it had increased the average amount of data delivered per Mbps per watt by almost 80%. It did this by switching to next-generation server hardware and software that uses less energy. Streaming Media's 'Greening of Streaming' initiative also notes that Limelight is proactively selecting data centres based on access to renewable energy.

Netflix ensured that 100% of its estimated direct and indirect non-renewable power use was matched with renewable energy certificates and carbon offsets in 2019. Furthermore, Google matched 100% of its electricity consumption with purchases of renewable energy and Microsoft intends to be carbon negative by 2030 and remove its entire impact by 2050. The good news is these aren't the only large companies to reduce their energy and carbon footprint, but individual corporate effort is only part of global structural change.

Summary

Demand for data will only grow. Videos, games and social sharing already account for 80% of internet traffic. The challenge will be to find new technologies that can help us grow the economy and also reduce energy and carbon. Ultimately we need many more solutions such as iSIZE's if we are to reduce carbon emissions and bring climate change under control.

Maria Ingold is a board-level strategic and technical innovator and Chief Technology Officer (CTO) with a nearly 30-year track record in successfully deploying emerging visual technology. She is a board adviser to iSIZE. As the former CTO of a Disney / Sony joint venture, she delivered one of the most successful video on demand (VOD) movie services in Europe from launch, which was profitable from year one. Maria is also a global public speaker (200+ times), writer and BAFTA (British Academy of Film and Television Arts) voting member.

Read the article on linkedIn here >