

Performance analysis of Energy Efficient Ethernet on Video Streaming Servers

Antonio de la Oliva^a, Tito R. Vargas^b, Juan Carlos Guerri^b, José Alberto Hernández^a, Pedro Reviriego^c

^a*Universidad Carlos III de Madrid, Spain*

^b*Universitat Politècnica de València, Spain*

^c*Universidad Antonio de Nebrija, Spain*

Abstract

Current trends on traffic growth oversee a steady increase of video streaming services, and the subsequent development of the associated infrastructure to allocate and distribute such contents. One of the operational costs associated to this infrastructure is the power bill. Therefore any mechanism used to decrease it, reducing also the carbon footprint associated to it, is welcome. In this work we investigate the suitability of the recently standardized IEEE 802.3az Energy Efficient Ethernet (EEE) for video traffic generated by video-streaming servers. The conclusion of the analysis is positive about the achievable energy savings, due to the inherent features of traffic patterns of video-streaming servers which help reducing the number of transitions between active and low-power modes in EEE.

Keywords: Ethernet, Energy Efficiency, Video Streaming

1. Introduction

The Internet growth has also raised several concerns about its power consumption and the need to improve its energy efficiency [1]. It is estimated that the Internet accounts for tens of TWh, an amount that could be significantly reduced if energy efficiency policies were applied. One of the first

Email addresses: aoliva@it.uc3m.es (Antonio de la Oliva), tivarher@iteam.upv.es (Tito R. Vargas), jcguerri@dcom.upv.es (Juan Carlos Guerri), jahgutie@it.uc3m.es (José Alberto Hernández), previrie@nebrija.es (Pedro Reviriego)

efforts toward energy savings is the IEEE 802.3az (Energy Efficient Ethernet) standard [2] which reduces the energy consumption of Ethernet devices by introducing a low-power (or sleep) mode. As Ethernet is the dominant technology for wire-line LANs, with more than one billion devices already deployed, Energy Efficient Ethernet (EEE) is expected to enable large energy savings [3]. However, previous studies in simulation [4] and with actual power measurements [5] have shown that the performance of Energy Efficient Ethernet greatly depends on the traffic pattern, and may not be as efficient as originally expected. This is due to the large power-mode transition times; that is, most of the power required for the transmission of an isolated Ethernet frame is spent in activating the link and putting it back to the low-power mode, rather than on actual data transmission. On the contrary, when several Ethernet frames are transmitted back-to-back, such a burst of data shares the sleep-to-active mode transition overhead resulting in larger energy savings [6].

Such bursty traffic pattern has been observed in video-streaming traces, which suggests that the adoption of EEE in video-streaming servers may achieve large power savings. Hence, this article attempts to analyze and quantify the suitability of the Energy Efficient Ethernet standard for the service of real time video streaming in IPTV networks, using the H.264 codec which is becoming the most popular video codec [7]. Indeed, given the high popularity of video streaming services such as YouTube, Hulu or Netflix, Internet video is expected to comprise about half of the total traffic share in the Internet by 2016, according to latest estimates [8].

2. Background

2.1. Overview of Energy Efficient Ethernet

Concerning power consumption of legacy Ethernet equipment, two important facts have been observed:

1. Consumption increases with the link speed [5]. This is due to the increased hardware complexity that is required to achieve higher speeds.
2. Consumption is always maximum and does not depend on the traffic load because of the continuous transmission of physical layer signaling in the absence of user data [5]. This signaling is mandatory to keep the receivers aligned to the channel conditions.




Figure 1: Transitions between active and low-power modes in Energy Efficient Ethernet

Protocol	Min T_w (μs)	Min T_s (μs)	T_{Frame} (1500B) (μs)	η_{Frame}	T_{Frame} (150B) (μs)	η_{Frame}
100BASE-TX	30	200	120	34.3%	12	4.9%
1000BASE-T	16	182	12	5.7%	1.2	0.6%
10GBASE-T	4.48	2.88	1.2	14.0%	0.12	1.6%

Table 1: Wake up, sleep and frame transmission times for different speeds, as proposed in IEEE 802.3az

The latter behavior is very effective in ensuring high-speed communication, but it results in poor energy efficiency, especially when the link is lightly loaded. This is usually the case for links that interconnect end-user computers to a LAN, where there is more room for power savings [9].

To improve energy efficiency, the Energy Efficient Ethernet (IEEE 802.3az) standard [2] defines a low-power (also known as, low-power idle or sleep) mode that can be used when there is no data to transmit. The power consumption while in the low-power mode is expected to be much lower (typically 10%) of that in the active mode [4]. The operation of Energy Efficient Ethernet is illustrated in Fig. 1. In the figure, the following times are defined:

T_w denotes the wake-up time, that is, the time required to exit the low-power mode, and bring the link ready for transmission.

T_s refers to the sleep time, that is, the time needed to enter the sleep (low-power) mode, for instance, after the transmission of the last frame.

T_q denotes the maximum values for periods with no transmission. After a T_q with no activity, a refresh period T_r is needed.

T_r refers to the short refresh periods, that is, the signaling periods necessary to keep the receivers aligned.

However, EEE is not as efficient as it was originally expected because the wake-up and sleep-down times are too large with respect to the frame transition time. Table 1 summarizes the minimum values (that is, best cases) for T_s and T_w as proposed in the 802.3az standard, along with their transmission efficiencies η_{Frame} computed for long and short Ethernet frames as:

$$\eta_{\text{Frame}} = \frac{T_{\text{Frame}}}{T_w + T_{\text{Frame}} + T_s} \quad (1)$$

where T_{Frame} refers to the Ethernet frame transmission time (frame size divided by link rate). Thus, η gives the percentage of time spent on actual data transmission per wake-transmission-sleep cycle.

As noted from Table 1, the wake and sleep times are considerably high with respect to the frame transmission time T_{Frame} , especially for small frames. Hence, the energy consumed for the transmission of a single frame, assuming the Network Interface Card (NIC) is in the low-power mode, requires a total time of $T_w + T_{\text{Frame}} + T_s$, whereas only T_{Frame} is used for actual data transmission. The T_s and T_w clearly impose an energy overhead which, for the case of 1000BASE-T reaches up to 99.4% for small frames (150-byte long). Such an excessive energy overhead makes EEE very inefficient in scenarios where the traffic load is low and the average frame size is small, as shown in [4, 5].

However, large power savings may be achieved if a group of frames are transmitted together (i.e. back-to-back), as proposed in [6]. This way, the link is only awoken for the transmission of a large portion of data, hence the cost of waking up and sleeping down a link is shared between several frames, greatly improving its efficiency:

$$\eta_{\text{Burst}} = \frac{T_{\text{Burst}}}{T_w + T_{\text{Burst}} + T_s} \quad (2)$$

Here, $T_{\text{Burst}} = N \times T_{\text{Frame}}$, typically $N \geq 10$.

2.2. Traffic profile in IPTV networks

The forthcoming experiments consider the output of a video streaming server intended for IPTV scenarios over UDP transport. In such scenarios, the operator often uses MPEG-4 or H.264/AVC codecs [10] to generate either

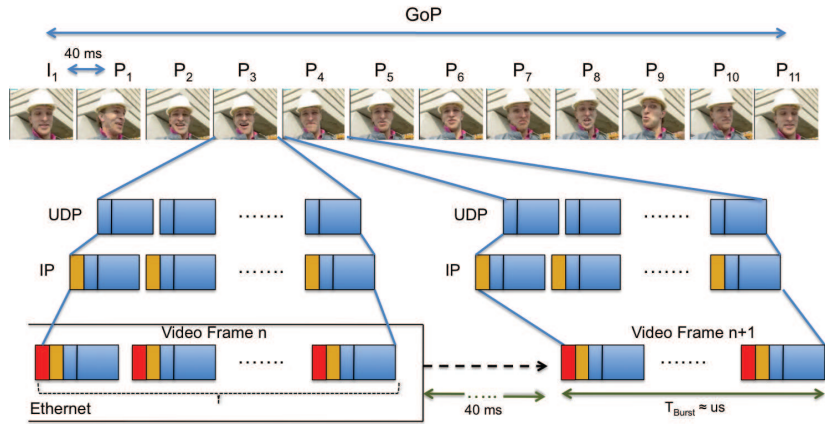


Figure 2: Top-down Video Frame Packetization

fixed or variable-quality (Q) Standard or High Definition content. The fixed- Q case aims at maintaining the quantization level at the expense of producing more traffic when the video sequence contains several moving objects. On the contrary, the variable- Q case produces an almost constant bit rate traffic pattern, however at the expense of having quality degradation at scenes with a large number of moving objects.

Fig. 2 shows the video-encoding and frame packetization process. Essentially, the codec generates 12-frame Group of Pictures (GoP) every 480ms, where frames can be of type I (Intra), P (Predictive) or B (Bidirectional) [11]. Typical examples of GoPs are IPPPPPPPPPPP and IBBPBBPBBPBB. The I frames are usually larger than the P and B frames, and therefore they produce larger bursts of Ethernet frames (see $E(B_I)$ and $E(B_P)$ and $E(B_B)$ in Tables 2 and 3). Additionally, because there are 11 times more P/B frames than I frames, the average burst size $E(B)$ is usually closer to $E(B_P)$ and $E(B_B)$ than to $E(B_I)$. Table 2 summarizes the main features of three different video traces (Movie, Sport and News) of about two minutes of duration encoded with the H.264 standard and GoP structure IPPPPPPPPPPP. The upper part of the table considers variable- Q and fixed resolution with nearly constant bitrate: 640x480 (2 Mbps), 720x480 (4 Mbps), 1280x720 (10 Mbps) and 1920x1080 (14 Mbps) [12]. The lower part of the table considers the same videos with fixed Quantization level values, namely 31, 16, 8 and 2 respectively. In this case, the bitrates experience more variability than be-

Variable Q (without B frames)					
Trace	Resolution (pixels)	Avg. Bitrate (Mbps)	$E(B_I)$ (Pkts)	$E(B_P)$ (Pkts)	$E(B)$ (Pkts)
Movie	640x480	2.41	27.69	7.38	9.28
Movie	720x480	3.91	48.87	11.44	14.74
Movie	1280x720	10.18	93.48	32.30	37.63
Movie	1920x1080	14.36	110.07	47.46	52.90
Sport	640x480	1.69	14.78	5.82	6.65
Sport	720x480	3.91	33.40	12.89	14.75
Sport	1280x720	10.22	68.77	34.66	37.76
Sport	1920x1080	13.45	74.41	47.07	49.56
News	640x480	2.07	13.09	7.00	8.01
News	720x480	3.67	28.77	11.49	13.86
News	1280x720	9.36	48.25	32.28	34.62
News	1920x1080	12.20	54.74	43.28	44.97
Fixed Q (without B frames)					
Trace	Quantizer Q	Avg. Bitrate (Mbps)	$E(B_I)$ (Pkts)	$E(B_P)$ (Pkts)	$E(B)$ (Pkts)
Movie	31	2.88	24.78	9.53	10.98
Movie	16	3.99	36.75	12.93	15.02
Movie	8	5.52	49.37	17.63	20.62
Movie	2	14.39	101.25	46.99	52.96
Sport	31	1.68	13.10	5.97	6.62
Sport	16	2.54	21.12	8.71	9.75
Sport	8	4.00	29.47	13.61	15.07
Sport	2	13.45	74.41	47.07	49.56
News	31	2.14	12.26	7.50	8.28
News	16	3.87	29.47	13.19	14.61
News	8	4.35	22.48	15.16	16.33
News	2	4.96	21.55	17.98	18.59

Table 2: Video traces for fixed bitrate and variable quantizer (Up) and variable bitrate and fixed quantizer (Down). GoP structure IPPPPPPPPPP.

Variable Q (with B frames)						
Trace	Resolution (pixels)	Avg. Bitrate (Mbps)	$E(B_I)$ (Pkts)	$E(B_P)$ (Pkts)	$E(B_B)$ (Pkts)	$E(B)$ (Pkts)
Movie	640x480	2.24	27.84	9.46	5.44	8.65
Movie	720x480	3.92	46.48	17.89	8.84	14.80
Movie	1280x720	10.24	94.95	49.30	24.85	37.82
Movie	1920x1080	13.46	110.09	72.53	31.70	49.59
Sport	640x480	1.78	14.42	8.37	5.39	6.99
Sport	720x480	3.92	33.50	19.25	10.51	14.77
Sport	1280x720	10.03	69.74	55.85	25.60	37.09
Sport	1920x1080	10.68	72.06	61.21	26.85	39.46
News	640x480	1.84	13.54	6.37	5.69	7.20
News	720x480	3.66	27.59	16.23	9.55	13.86
News	1280x720	9.30	55.44	46.44	24.77	34.44
News	1920x1080	10.66	58.78	54.23	28.86	39.38
Fixed Q (with B frames)						
Trace	Quantizer Q	Avg. Bitrate (Mbps)	$E(B_I)$ (Pkts)	$E(B_P)$ (Pkts)	$E(B_B)$ (Pkts)	$E(B)$ (Pkts)
Movie	31	2.40	25.06	11.87	5.83	9.23
Movie	16	3.44	36.26	16.33	8.49	13.02
Movie	8	5.28	50.97	23.29	13.91	19.74
Movie	2	9.01	73.31	38.72	25.35	33.36
Sport	31	1.78	12.96	8.44	5.60	6.99
Sport	16	2.63	20.65	12.27	7.83	10.06
Sport	8	4.21	29.46	19.89	12.42	15.83
Sport	2	7.26	46.49	34.28	21.58	26.96
News	31	1.88	12.72	7.41	5.69	7.34
News	16	3.00	24.33	13.19	8.16	11.42
News	8	4.28	22.62	18.62	13.27	16.11
News	2	7.19	34.81	31.75	22.65	26.74

Table 3: Video traces for fixed bitrate and variable quantizer (Up) and variable bitrate and fixed quantizer (Down). GoP structure IBBPBBPBBPBB.

fore with respect to the video-clip type. Basically, keeping a certain quality resolution for a sport event requires a higher bitrate value than for a movie, and much more than for the news TV show. Table 3 shows trace statistics for the same video clips and encoding settings, but with GoP structure IBBPBBPBBPBB.

For high-resolution videos, the I frames comprise about 50-100 Ethernet frames whereas the P/B frames contain 20-40, depending on the encoding settings. Thus, every I/P/B frame produces a burst of Ethernet frames which are then sent back-to-back to the NIC. Such traffic generation pattern is very suitable for EEE since the link only goes to sleep after the whole video frame is transmitted, thus sharing the cost of waking up and sleeping down the link very much like in burst transmission EEE. Indeed, it is the video traffic source which generates the data bursts, and no burst-assembler is needed at the NIC, which brings the energy efficiency benefits of burst transmission EEE without the delay penalty.

In addition, it is worth noticing that the average burst size $E(B)$ (in terms of Ethernet frames) for GoP IPPPPPPPPPPP is slightly larger than for GoP IBBPBBPBBPBB. Hence, the use of B frames contributes to an increase in video compression, but has the disadvantage of a greater computational complexity than for GoPs without B frames. Concerning transmission, B frames need to be temporally stored until their I or P frame has been generated. For example, GoP $I_1 B_2 B_3 P_4 B_5 B_6 P_7 B_8 B_9 P_{10} B_{11} B_{12}$ would be transmitted in the following order: $I_1 P_4 B_2 B_3 P_7 B_5 B_6 P_{10} B_8 B_9 I_{13} B_{11} B_{12}$ since B_2 and B_3 depend on P_4 , B_5 and B_6 depend on P_7 , etc. Finally, since the video server must wait for the generation and codification of P_4 before encoding frames B_2 and B_3 , we can see from the traces that the video server actually transmits the group $P_4 B_2 B_3$ back-to-back, thus generating a burst of about $E(B_P) + 2E(B_B)$ Ethernet frames, which favours energy savings, as shown in the next section.

A final case of video encoding settings will be considered in the EEE performance analysis of the next section. Such a case considers H.264 Scalable Video Coding (aka H.264/SVC) that takes into account a different GoP structure built by hierarchical B frames. Such a hierarchical feature of H.264/SVC refers to the fact that B frames can be used to predict other B frames, in contrast with the classical H.264/AVC B frame prediction process where a B frame depends on the previous and next I/P frames. Essentially, the layer-scalable H.264/SVC codec offers temporal, spatial and quality scalability to perform the decoding process using a single base layer or the base layer plus several enhancement layers (see [7, 13] for further details). This allows to

remove some B frames without losing the ability to decode the video.

More specifically, the sample video Movie has been encoded using H.264/SVC with temporal scalability and GoP size 16 with hierarchical B frames (see Table 4 for further details), using the same software tool as in [7]. Several differences must be noted with respect to the H.264/AVC cases: Firstly, the GoP structure and size varies depending on the number of layers transmitted, going from G16B15 for the highest temporal scalability with all layers (5 in this case) down to G1B0 for the temporal base layer only, and so does the average GoP size.

Concerning transmission, given the hierarchical nature of the prediction of B frames, the encoded video is transmitted on a per-GoP basis, that is, all the frames of the same GoP are transmitted back-to-back every 680ms. For instance, the five-layered Movie trace would be transmitted in the following order: $I_0 I_{16} B_8 B_4 B_2 B_1 B_3 B_6 B_5 B_7 B_{12} B_{10} B_9 B_{11} B_{14} B_{13} B_{15} I_{32} B_{24} B_{20} \dots$. So, the whole GoP $I_{16} B_8 B_4 B_2 B_1 B_3 B_6 B_5 B_7 B_{12} B_{10} B_9 B_{11} B_{14} B_{13} B_{15}$ would be transmitted at once, again favouring energy savings. This shall be studied in the next section.

# of layers	Bitrate (Mbps)	$E(B_I)$ (Pkts)	$E(B_B)$ (Pkts)	$E(B)$ (Pkts)	GoP struct.
Five	2.51	46.70	10.79	13.05	$I_0 B_1 B_2 B_3 B_4 B_5 B_6 B_7 B_8 B_9 B_{10} B_{11} B_{12} B_{13} B_{14} B_{15}$
Four	1.94	46.70	13.94	18.05	$I_0 B_2 B_4 B_6 B_8 B_{10} B_{12} B_{14}$
Three	1.43	46.70	17.43	24.78	$I_0 B_4 B_8 B_{12}$
Two	1.02	46.70	20.55	33.66	$I_0 B_8$
Base	0.74	46.70	-	46.70	I_0

Table 4: Movie video trace encoded with H.264/SVC.

3. Evaluation of Energy Efficient Ethernet Performance in Video Streaming Servers

3.1. Transmission of a single H.264/AVC-encoded video

We have used the same simulation model of Energy Efficient Ethernet of previous studies [4] to evaluate the performance of EEE for each video trace (Movie, Sport and News with either variable or fixed Q and both GoP types) of Tables 2 and 3. Essentially, the traces provide both the frame arrival-time and size values, necessary to estimate the amount of time that the

Ethernet link spends in the sleep and active modes, as well as the transitions between them, for different link speeds: 100 Mbps, 1 Gbps and 10 Gbps. The simulation tool then estimates the average efficiency value (η) in each scenario as:

$$\bar{\eta} = \frac{\text{Total data transmission time}}{\text{Total time in active mode and transitions}} \quad (3)$$

which, as eq. 2, is near 100% when the number of T_w and T_s is small compared with the portion of time under actual data transmission.

The efficiency values η are then obtained for the following configurations:

Poiss This is the efficiency figure obtained after simulation as if packets were transmitted following a Poisson basis, rather than back-to-back per video frame. This simulation was performed assuming the same traffic load and packet sizes as in the original trace, but with exponentially distributed packet inter-arrival times with mean $1/\lambda$, where:

$$\lambda = \frac{\rho}{8E(P)/C}$$

for a given load ρ obtained from the trace as:

$$\rho = \frac{\text{Bitrate}}{C}$$

It is important to remark that the average Ethernet frame size in the video traces is $E(P) = 1370$ bytes, and C refers to the link capacity (100 Mbps, 1 or 10 Gbps).

Real This is the efficiency value for the transmission of the video trace over EEE as specified in the IEEE 802.3az standard. As shown, the results are far better than in the Poisson case. Essentially, EEE benefits from the back-to-back nature of packet arrivals at the NIC, which allows to share the energy cost of T_w and T_s per I, P or B frame.

GoP This case considers hypothetically that a whole GoP is aggregated before its transmission. We assume that the NIC aggregates the I frame and its 11 B or P subsequent frames and transmits them all at once. Such a per-GoP burstification is proposed since the video source is able to change its codification parameters every GoP period. When following this strategy, data bursts of about $12 \times E(B) \approx 135$ KBytes are transmitted every $12 \times 40 = 480$ ms.

Ideal Ideally, efficiency should be 100% if the T_s and T_w values were negligible with respect to the frame transmission times.

The simulated efficiency results are displayed in Figs. 3 and 4. For instance, let us take Fig. 3 case Movie, 100Base-Tx, resolution 640x480, I/P. This case considers the experiment of encoding the Movie clip using resolution 640x480 and no B frames. Table 3 first row states that the resulting packet trace has the following features: bitrate 2.41Mbit/s, 27.69 packets per I frame on average and 7.38 packets per P frame on average, 40ms inter-frame arrival time at the NIC. The packet trace is then transmitted by a 100Mbit/s Ethernet NIC employing EEE. The theoretical efficiency in this case would be (eq. (2)):

$$\eta_{real} = \frac{E(B)T_{frame}}{T_w + E(B)T_{frame} + T_s} = \frac{9.28 \times 120\mu s}{30\mu s + 9.28 \times 120\mu s + 200\mu s} = 0.82$$

which is the result shown in the Figure (grey bar). A Poisson traffic profile with the same bitrate results in a much lower efficiency, 30% in the figure, since the link is in most cases awoken per frame. Theoretically this would be:

$$\eta_{Poisson} \approx \frac{120\mu s}{30\mu s + 120\mu s + 200\mu s} = 0.34$$


Finally, the GoP burstification achieves nearly 100% efficiency since the whole GoP is transmitted at once every 480ms:

$$\eta_{GoP} = \frac{9.28 \times 12 \times 120\mu s}{30\mu s + 9.28 \times 12 \times 120\mu s + 200\mu s} = 0.98$$


which is consistent with the result depicted in the Figure.

Energy efficiency simulations have been performed for every single experimental case: video type and encoding settings. The following conclusions per link speed can be observed from the figures:


- In 100BASE-TX, the efficiency η is in the order of 10-40% for Poisson traffic. This value rises to nearly 70-90% for video traffic due to the back-to-back nature of video traffic, and to nearly 100% when GoP burstification is applied.



(a) Movie




(b) News




(c) Sport


Figure 3: Efficiency for Variable Q



(a) Movie



(b) News



(c) Sport

Figure 4: Efficiency for Fixed Q

- In 1000BASE-T, η is nearly 5-10% for Poisson traffic, and rises to 40-75% for video traffic (depending on the video type and encoding settings), and to 90% or above when GoP-burstification is applied.
- In 10GBASE-T, η is around 10-15% for Poisson traffic, and rises to 60-80% for video traffic, and then to 95% when GoP-burstification is applied.

Hence, the back-to-back nature of video-streaming traffic allows a substantial energy efficiency increase in all cases, and particularly in the case of dynamic scenes (movie and sports) and/or high bitrate. Therefore, EEE allows a substantial power consumption decrease in video-streaming sources since the cost of waking up and sleeping down a link is shared per I, P or B frame, rather than per packet.

Concerning the GoP type, the use of B frames achieves an extra efficiency of 5-10% for 100BASE-TX, 20-25% for 1000BASE-T and about 20% for 10GBASE-T with respect to GoP type IPPPPPPPPPPP. Such a performance increase for GoP type IBBPBBPBBPBB is the consequence of larger data bursts than for GoP type IPPPPPPPPPPP, since the traces reveal that frames are transmitted together as a group of two B and one I/P every 120ms (instead of one video frame every 40ms). Therefore, the use of B frames is highly recommended for energy efficiency purposes.

3.2. Transmission of a single H.264/SVC-encoded video

Fig. 5 shows the energy efficiency values for the H.264/SVC-encoded video traces over 100BASE-TX, 1000BASE-T and 10GBASE-T, also taking into account each temporal scalability layer where EL and BL stand for Enhancement Layer and Basic Layer respectively in the figure. Essentially, given the hierarchical nature of B frames, the whole GoP is transmitted in a single burst, thus producing very-high efficiency values, similar to those obtained in the hypothetical GoP burstification case of H.264/AVC. In addition, we observe that removing different enhancement layers reduce the average burst size, and its energy efficiency consequently (smaller bursts). The results obtained show that the H.264/SVC encoding settings provide better energy efficiency values than H.264/AVC while increasing the compression ratio.

3.3. Analysis for a large number of video streaming sources

The efficiency values for single video transmission provides a good starting point concerning the suitability of EEE for video-streaming traffic. However,




Figure 5: Efficiency for H.264/SVC encoded videos

video-streaming servers usually dispatch many videos in parallel rather than a single one. In light of this, Figs. 6(a) to 6(c) show the power consumption figures when serving multiple videos over the same EEE NIC, at different link speeds (the results presented are average values for 100 simulations). The simulations are performed in scenarios that combine different link speeds and bitrates for variable Q, and GoP type IPPPPPPPPPPP. We assume that the starting times of the videos are uncorrelated, since these come from different user requests.

It is worth remarking that one video flow at 640x480 resolution has an average bitrate of 2Mbit/s (see Table 2) which, at a 100Mbit/s link speed represents 2% of the total link utilization. Thus, about 50 videos need to be randomly aggregated to achieve a simulation scenario of nearly 100% load at 100Mbit/s. Similarly, about 500 aggregated videos are needed to achieve 100% load at 1Gbit/s and 5000 videos at 640x480 resolution are needed for 100% load at 10Gbit/s.

As shown, the Poisson traffic pattern shows very poor results as already reported in previous studies [4]. However, the transmission of multiple videos in parallel is very close to the ideal behavior, especially at 100 Mbps and 10 Gbps and for high resolution values (720p and 1080p). Again, the reason for this is that low-resolution videos generate short-size bursts, which translates into worse efficiency values than higher-resolution videos, and results in a clear deviation from the ideal case.




Figure 6: Consumption vs load when multiple flows are transmitted over EEE: (a) 100BASE-TX, (b) 1000BASE-T, (c) 10GBASE-T

4. Conclusions

This work has studied the behavior and performance of Energy Efficiency Ethernet in video-streaming servers, showing good results in terms of energy savings. Essentially, the back-to-back transmission nature of the video-streaming sources results in the transmission of a large portion of data per wake-up and sleep-down cycle. Indeed, the number of power-mode transitions is reduced with respect to other traffic patterns, and the energy overhead caused by switching between the active and low-power sleep mode is minimized. The experiments have further revealed that the use of B frames produce larger transmission bursts than those without B frames, hence achieving

larger efficiency values. Furthermore, it has been shown that the H.264/SVC codec not only increases the compression ratio with respect to H.264/AVC but is also more energy efficient since it generates GoP-size data bursts.

Such behavior is quite different from that observed in previous studies for other applications, where EEE showed worse performance. Hence, video-streaming servers do not need any further technique to improve Energy Efficient Ethernet such as the aggregation of packets before transmission [6], since the video-streaming codecs already implement this functionality as part of its encoding process.

It is also worth remarking that this study has focused on UDP-based video streaming, where it is reasonable to assume that the original packet timing is preserved for transmission. Future work will address the case of TCP-based streaming, where the frame transmission timing is expected to be throttled by the TCP congestion-control mechanisms, causing a significant power efficiency decrease. In addition to this, the TCP acks sent back to the video source are expected to further degrade the efficiency of EEE direction.

Acknowledgments

Part of the research leading to these results received funding from the European Community's Seventh Framework Programme (FP7-ICT-2009-5) under grant agreement no. 258053 (MEDIEVAL project). Additionally, the authors would like to acknowledge the support to this work by the CAM-UC3M Greencom Research Grant (under code CCG10-UC3M/TIC-5624), the FIERRO Spanish project (TEC2010-12250-E) and the Google Research Award "New Protocol Semantics and Scheduling Primitives for Energy Efficiency: Burst Coalescing at the Link and Application Layers".

References

- [1] M. Gupta, S. Singh, Greening of the Internet, in: Proc. ACM SIGCOMM, 2003, pp. 19–26.
- [2] IEEE Std 802.3az Energy Efficient Ethernet-2010 (2010).
- [3] K. Christensen, P. Reviriego, B. Nordman, M. Bennett, M. Mostowfi, J. Maestro, IEEE 802.3 az: the road to energy efficient ethernet, IEEE Communications Magazine 48 (11) (2010) 50–56.

- [4] P. Reviriego, J. A. Hernández, D. Larrabeiti, J. A. Maestro, Performance evaluation of Energy Efficient Ethernet, *IEEE Communications Letters* 13 (9) (2009) 697–699.
- [5] P. Reviriego, K. Christensen, J. Rabanillo, J. Maestro, An Initial Evaluation of Energy Efficient Ethernet, *IEEE Communications Letters* 15 (5) (2011) 578–580.
- [6] P. Reviriego, J. A. Maestro, J. A. Hernández, D. Larrabeiti, Burst transmission for Energy Efficient Ethernet, *IEEE Internet Computing* 14 (4) (2010) 30–37.
- [7] P. Seeling, M. Reisslein, Video transport evaluation with H.264 video traces, *IEEE Communications Surveys Tutorials* PP (99) (2011) 1–24. doi:10.1109/SURV.2011.082911.00067.
- [8] C. W. Paper, Cisco Visual Networking Index: Forecast and Methodology, 2011-2016, Tech. rep., Cisco Systems inc. (2012).
- [9] A. Odlyzko, Data networks are lightly utilized, and will stay that way, *Review of Network Economics* 2 (3) (2003) 2.
- [10] H. Kalva, The H.264 video coding standard, *IEEE Multimedia* 13 (4) (2006) 86–90.
- [11] P. Seeling, M. Reisslein, B. Kulapala, Network performance evaluation using frame size and quality traces of single-layer and two-layer video: A tutorial, *IEEE Communications Surveys & Tutorials* 6 (3) (2004) 58–78.
- [12] S. Han, S. Lisle, G. Nehib, IPTV transport architecture alternatives and economic considerations, *IEEE Communications Magazine* 46 (2) (2008) 70–77.
- [13] H. Schwarz, D. Marpe, T. Weigand, Overview of the scalable video coding extension of the H.264/AVC standard, *IEEE Trans. Circuits Syst. Video Technol.* 17 (9) (2007) 1103–1120.