

Performance Evaluation of Energy Efficient Ethernet

P. Reviriego, J. A. Hernández, D. Larrabeiti, and J. A. Maestro

Abstract—Until very recently, energy efficiency has received little attention in many wired communications environments. For example, in most current Ethernet standards the transmitter and receiver operate at full power even when no data is being sent. However, new upcoming energy-aware standards, such as Energy Efficient Ethernet (EEE), are addressing this issue by introducing a low power mode for idle link intervals. The future EEE standard defines the procedure to enter and exit the low power mode. With EEE the actual energy savings will depend on the amount of traffic and on the timing of the frame arrivals. In this paper the performance of EEE in terms of energy saving is evaluated. The results show that although EEE improves the energy efficiency, there is still potential for substantial further energy savings as in many cases most of the energy is wasted in waking up and sleeping the link.

Index Terms—IEEE 802.3, Ethernet, energy efficiency.

I. INTRODUCTION

THE efficient use of energy in communications is a growing concern for both industry and governments worldwide. The massive amount of communications devices that are used today, together with their expected growth, have led to the conclusion that significant energy can be saved by applying energy efficiency concepts in the design of communication systems [1]. Indeed, the Internet core is estimated to consume about 6 TWh per year [2], an amount which can be significantly reduced if energy-aware protocols are deployed. Ethernet is a good example of technology that can be made energy-efficient with important savings, estimated over 3 TWh [3]. To reduce such a waste, the IEEE 802.3az Task force (Energy Efficient Ethernet) is, at present, introducing energy efficiency enhancements to the existing Ethernet, a work that is expected to produce a new standard by 2010 [4]. Essentially, current Ethernet standards require both transmitters and receivers to operate continuously on a link, thus consuming energy all the time, regardless of the amount of data exchanged. This consumption depends on the link speed and ranges from around 200mW for 100Mbps [5] to about 5W for 10Gbps [6]. Energy Efficient Ethernet (EEE) aims to make the consumption of energy over a link more proportional to the amount of traffic exchanged [7]. To this end, EEE defines a low power "sleep" mode. The physical layer is put into this sleep mode when no transmission is




Fig. 1. Transitions between active and low-power modes in Energy Efficient Ethernet

TABLE I
PROPOSED WAKE UP, SLEEP AND FRAME TRANSMISSION TIMES FOR DIFFERENT SPEEDS

Protocol	Min T_w (μs)	Min T_s (μs)	T_{Frame} (1500B) (μs)	Frame eff.	T_{Frame} (150B) (μs)	Frame eff.
100Base-Tx	30	100	120	48%	12	8.5%
1000Base-T	16	182	12	5.7%	1.2	0.6%
10GBase-T	4.16	2.88	1.2	14.6%	0.12	1.7%

needed and waken up very quickly upon data arrival without changing its speed. The low-power mode freezes the elements in the receiver, and wakes them up within a few microseconds. Such sleep/active operation requires minor changes to the receiver elements since the channel is quite stable. Figure 1 shows a state transition example of a given link as following the IEEE 802.3az draft [8]. Here, T_s refers to the sleep time (time needed to enter the low-power mode), and T_w denotes the wake-up time (time required to exit the low-power mode). The transceiver spends T_q in the quiet (energy saving) period. Finally, the standard also considers the scheduling of periodical short periods of activity T_r to refresh the receiver state in order to ensure that the receiver elements are always aligned with the channel conditions.

Concerning energy efficiency of EEE, there is significant energy consumption only during T_w , T_s and T_r periods, and a small fraction of it (about 10%) during T_q , with $T_q \gg T_r$. The minimum and maximum values for timers T_w , T_s , T_q and T_r are specified in the IEEE 802.3az draft for 100Base-TX, 1000Base-T and 10GBase-T. Table I gives the minimum values for T_s and T_w as proposed in the draft, along with their frame transmission efficiencies for long and short frames.

Implementing the low-power mode brings large power savings, nearly 90% for 100BaseTX, 1000BaseT and 10GBaseT, with respect to the current standard which operates at full power all the time. However, as noted from Table I, the wake up and sleep times are considerably high with respect to the frame transmission time T_{frame} , especially when the frame is small. See for instance the following example: Assume that a given device is in the low-power mode upon a frame arrival. At this point, the device needs to wake up (which takes T_w), transmit its frame (this takes T_{frame}) and go to sleep again (this takes T_s). In total, the transmission of a single

Manuscript received April 14, 2009. The associate editor coordinating the review of this letter and approving it for publication was A. Skeericioglu.

P. Reviriego and J. A. Maestro are with the Universidad Antonio de Nebrija, C/ Pirineos 55, E-28040 Madrid, Spain (e-mail: {previrie, jmaestro}@nebrija.es).

J. A. Hernández is with the Universidad Autónoma de Madrid, C/ Francisco Tomás y Valiente 11, E-28049 Madrid, Spain (e-mail: jose.hernandez@uam.es).

D. Larrabeiti is with the Universidad Carlos III de Madrid, Av. Universidad 30, E-28911 Leganés, Madrid, Spain (e-mail: dlarra@it.uc3m.es).

Digital Object Identifier 10.1109/LCOMM.2009.090880




Fig. 2. Energy consumption versus load for (a) 100Mbps, (b) 1Gbps and (c) 10Gbps.

frame takes $T_w + T_{\text{frame}} + T_s$, whereas only T_{frame} is used for actual data transmission. This algorithm, on a 10Gbps link, requires $T_w \geq 4.16\mu s$, $T_s = 2.88\mu s$ and $T_{\text{frame}} = 1.2\mu s$ for the transmission of a 1500-byte data frame, resulting in an efficiency of 14.6% since most of the time (and energy) is spent in waking up and putting the link into sleep. Such energy overhead is particularly high for small frames, and at high-speed rates, as shown in Table I. In the rest of this letter the performance of EEE in terms of energy consumption versus traffic load is studied.

II. PERFORMANCE EVALUATION

As noted in [7], energy proportionality is achieved when there is a linear relationship between the systems load and its energy consumption but, unfortunately, this is not the case for the current Energy Efficient Ethernet standard draft. To show this, let us simulate the energy consumption of 100 Mbps, 1 Gbps and 10 Gbps links at different traffic load values. For simplicity, both link directions are assumed to operate independently, although this is not true for 1000Base-T. Links enter the low-power mode only when no frames are pending for transmission, and the wake-up and sleep timers in EEE are simulated following the current standard draft [8] (see Table I). Additionally, the experiment considers the case of traffic made up of relatively large frames, which should display the best performance for EEE. For this purpose, a fixed frame size of 10000-bit was selected, and frames were programmed to arrive at the link following a Poisson process. This latter assumption can be a valid approximation for highly aggregated traffic [9], as it is the case for links of servers in large data centers dealing with many connections in parallel. Finally, the power consumption in the low-power mode is assumed to be 10% of that in the active mode for all Ethernet speeds, in line with the estimates provided by different manufacturers during the standardization process of EEE. In light of this, Figs. 2(a), 2(b) and 2(c) show the energy consumption plots versus traffic load in current standards (Legacy) and after the two power modes of EEE are introduced (EEE) for the three different scenarios: 100 Mbps, 1 Gbps and 10Gbps. Also for reference the proportional relation between load and energy consumption is shown. This relation is given by the following equation that takes into account the power consumption in low power mode:

$$P = P_{\text{low power}} \cdot (1 - \text{load}) + P_{\text{active}} \cdot \text{load} \quad (1)$$

As shown, current Ethernet standards operate at maximum power all the time, thus consuming 100% of energy regardless of the traffic load. By introducing the two power modes, the EEE standard makes energy consumption more proportional to traffic load, showing important energy savings, especially at low loads. However, for 1000Base-T and 10GBase-T the relationship between load and energy consumption is still distant from proportionality, which appears as a straight line from the bottom left to the top right of the plot. Basically, the large values of the sleep and wake-up timers compared to the actual transmission time of a single frame is the reason for such poor results, since most of the energy is spent on waking up and putting into sleep the link, rather than on actual data transmission. This is particularly harmful for high-speed links. The simulation results presented correspond to only one frame length while in real links different frame lengths will be used.

The simulation tries to capture the behavior of a link with large frames which are typical of data transfers. This is a best case for EEE performance, as it will be seen in the following, shorter frames introduce a larger overhead. To better analyze the expected energy consumption when EEE is adopted, a number of real measurement-based scenarios have been studied. In these experiments, the frame arrival times and lengths are used to estimate the amount of time that the link would be in low power mode. With this the energy consumption relative to the active mode is calculated. The results are presented in Table II where the energy consumption for proportionality (ideal) is also shown for reference. The first scenario considers a residential user who downloads video content from the Internet. This user is connected via 100Mbps Ethernet to his/her ADSL router. In this case the energy consumption is close to the low power mode level (10%) as there is very low load on the link. In this case EEE achieves good energy efficiency. The second scenario considers a situation where two users attached to the same 100Mbps-LAN exchange a file. This example highlights the main shortcoming of Energy Efficient Ethernet: its limited efficiency when small frames are transmitted. This is seen in the upstream direction (ACKs in the table) that is very lightly loaded with a very small average frame length. This results in a large overhead for waking up and putting the link into sleep ending up with over 44% power consumption (compared to 11% for proportionality) for a network load below 2%.

The third scenario analyzes a 1000Base-T University access

TABLE II
ENERGY CONSUMPTION ESTIMATES FOR DIFFERENT SCENARIOS

Scenario	Direction	Energy (% peak)	Ideal Energy (% peak)	Link Load %	Avg. Frame Size (B)
I. Residential user. Video download (100 Mbps)	download	12.75	11.28	1.43	1444
	upload	10.99	10.04	0.04	90
II. Residential user. File transfer	file	78.68	74.01	71.13	1499
	ACKs	44.92	11.25	1.39	77
III. University Internet Access Link. 1000BaseT	download	92.80	19.84	10.94	679
	upload	96.20	25.89	17.66	919
IVa. Data center: File and search server	input	65.90	11.10	1.22	87
	output	72.92	56.99	52.21	1497
IVb. Data center: Search server	input	45.28	17.66	8.51	945
	output	42.30	16.51	7.23	934
IVc. Data center: File and application server	input	61.37	10.58	0.65	130
	output	57.10	13.62	4.02	749

link with highly multiplexed Internet traffic. Although the results are computed considering both directions to operate independently, it is worth remarking that, for 1000Base-T both directions must enter the active or sleep modes at the same time. Nevertheless, this issue does not affect significantly the results since in this case it was observed that the link was in the active mode 90% of the time in both directions. As shown, for traffic loads below 20% energy consumption is over 90% showing poor energy efficiency. The last experiment considers a few server traces from Google's data centers, where energy consumption is a major concern. The traces belong to three typical server types: a file server which is also involved in search queries, a second server only devoted to search queries, and a third one which acts as both file and application server. The three servers used 1000Base-T in their network interfaces. In all cases, it can be concluded that EEE exhibits poor energy efficiency especially for links that operate at low loads with small frames where traffic is not bursty. These experiments suggest that EEE may result in good energy efficiency in lightly loaded user PCs while its efficiency may be poor in servers. This is a problem for data centers where energy consumption is a major issue and the Ethernet related consumption is large as high speed links are used. Finally, it is worth noting that the results presented will be valid even if the parameters T_w , T_s change in future releases of the standard draft as due to physical layer issues most of those values can not be reduced significantly.

III. CONCLUSIONS

The adoption of Energy Efficient Ethernet globally is expected to reduce the power consumption of Ethernet devices massively, since links will only be active when there is some data ready for transmission. However, the times required to wake up and put into sleep a link are crucial parameters, and the current draft standard which takes into account the current technological limits of transceivers specifies values too high compared to the frame transmission time. This has been shown to reduce the actual energy efficiency of EEE significantly. This efficiency reduction will be larger on high-speed links with significant traffic load, for example on server links. In

those cases there can be large deviations from energy proportionality. Thus, it becomes apparent that the Energy Efficiency of EEE could be further improved. If setting more restrictive mode switching times is not technologically feasible, as the proposed standard timings seem to suggest, further research and standardization work should provide specific guidance on the recommended frame scheduling algorithms that maximize the efficacy of EEE. These frame scheduling algorithms would group frames before waking up the link to minimize the overhead thus increasing the delay. The trade-off between energy consumption and delay should be carefully studied

ACKNOWLEDGEMENTS

The authors would like to thank K. Fu, G. Chesson, L. A. Barroso and U. Holzle from Google for the anonymous traces from their data centers used in the experiments in section II.

The work described in this paper was carried out with the support of the BONE project ("Building the Future Optical Network in Europe"), a Network of Excellence funded by the European Commission through the 7th ICT-Framework Programme.

REFERENCES

- [1] M. Gupta and S. Singh, "Greening of the Internet," in *Proc. ACM SIGCOMM*, Aug. 2003, pp. 19–26.
- [2] K. W. Roth, F. Goldstein, and J. Kleinman, "Energy consumption by office and telecommunications equipment in commercial buildings, vol. 1: energy consumption baseline," 2002 National Technical Energy Consumption Service (NTIS), U. S. Department of Commerce.
- [3] C. Gunaratne, K. Christensen, B. Nordman, and S. Suen, "Reducing the energy consumption in Ethernet with Adaptive Link Rate (ALR)," *IEEE Trans. Computers*, vol. 57, no. 4, pp. 448–461, 2008.
- [4] "IEEE P802.3az Energy Efficient Ethernet Task Force." Available at: <http://grouper.ieee.org/groups/802/3/az/public/index.html>.
- [5] J. Chou, "Low-power idle based EEE 100Base-TX," Mar. 2008, in IEEE 802.3az Task Force presentation.
- [6] B. Kohl, "10GBase-T power budget summary," Mar. 2007, in IEEE 802.3az Task Force presentation.
- [7] L. A. Barroso and U. Holzle, "The case for energy-proportional computing," *IEEE Computer*, vol. 40, no. 12, pp. 33–37, Dec. 2007.
- [8] "IEEE P802.3az Energy Efficient Ethernet Standard Draft 1.2.1," Mar. 2009.
- [9] T. Karagiannis, M. Molle, M. Faloutsos, and A. Broido, "A non-stationary Poisson view of Internet traffic," in *Proc. IEEE INFOCOM*, 2004.