An Initial Evaluation of Energy Efficient Ethernet

P. Reviriego, *Member, IEEE*, K. Christensen, *Senior Member, IEEE*, J. Rabanillo,
and J. A. Maestro, *Member, IEEE*

***Abstract*—In September 2010, the Energy Efficient Ethernet (IEEE 802.3az) standard was officially approved. This new standard introduces a low power mode for the most common Ethernet physical layer standards and is expected to provide large energy savings. In this letter, for the first time, Network Interface Cards (NICs) that implement Energy Efficient Ethernet (EEE) are used to measure energy savings with real traffic. The data presented will be useful to better estimate the energy savings that can be achieved when EEE is deployed. Existing analysis of EEE based on simulations predict a large overhead due to mode transitions between active and low power modes. The experimental results confirm that transition overheads can be significant, leading to almost full energy consumption even at low utilization levels. Therefore traffic patterns will play a key role in the energy savings achieved by EEE as it becomes deployed in the field.**

***Index Terms*—IEEE 802.3, Ethernet, energy efficiency.**

I. INTRODUCTION

THE IEEE 802.3az Energy Efficient Ethernet (EEE) standard is a good example of recent efforts to reduce energy consumption in wire-line networks [1]. Its adoption is predicted to provide energy savings in the order of TWh and cost reductions of hundred of millions of dollars [2]. These savings are achieved by introducing a low power mode in the most common Ethernet physical layers. Those include 100BASE-TX, 1000BASE-T and 10GBASE-T, which use Unshielded Twisted Pair (UTP) as the transmission medium. The main idea behind EEE is that when there is no data being exchanged over a link the physical layer can be put in a low power mode in which consumption is greatly reduced. Then as soon as data arrives for transmission the link is activated. The transition times are in the order of microseconds such that the added delay is negligible to many applications. The minimum transition times specified in the standard are shown in Table I where Tw refers to the time needed to activate a link that is in low power mode and Ts to the time needed to put a link into low power mode. The Table also shows the time required to transmit a 1500 byte frame. This illustrates the relative mode transition overhead for the different speeds.

With EEE there are no energy savings when the link is active and during transitions the energy savings are also predicted to be small. A previous performance evaluation based on simulation showed that the energy overhead due to transitions could be relevant even when the traffic load is low [3].

TABLE I
WAKE UP, SLEEP, AND FRAME TRANSMISSION TIMES [μSEC]

Protocol	Min Tw	Min Ts	Transmission time for a 1500 byte frame
100BASE-TX	30	200	120
1000BASE-T	16	182	12
10GBASE-T	4.48	2.88	1.2

Fig. 1. Photograph of the NIC used in the experiments.

At the end of September 2010 the EEE standard was officially approved and the first prototypes and products that implement EEE are now starting to become available. This means that we are now able to evaluate EEE performance using real hardware and with actual power measurements rather than with simulation and rough estimates for the power consumption in the different modes. The main objective of this letter is to present the first evaluation of EEE performance using actual measurements from EEE Network Interface Cards (NICs) manufactured by a major Ethernet vendor. The data presented can be useful to better understand and optimize the potential energy savings when using EEE. The results are also compared to previous studies showing that traffic patterns play a fundamental role in the energy savings obtained when using EEE.

II. MEASUREMENT SETUP

Two Realtek NICs [4] were used in the experiments. The NIC is shown in Fig. 1. These NICs implement the draft 3.2 of the EEE standard. The differences between the draft and the final standard are mostly editorial and have no implications for performance. The NICs supports 10BASE-T (10 Mb/s), 100BASE-TX (100 Mb/s) and 1000BASE-T (1 Gb/s) and are connected to the computer though a PCI express (PCIe) interface.

It is important to note that the NICs include other elements apart from the physical layer device (PHY) such as the Medium Access Control (MAC), the PCIe bus interface and an on-chip voltage regulator. The measurements reported in the experiments are for the whole NIC.

The approach to measure power consumption is similar to the one used in [5]. A small resistor (0.1 Ohms) is placed in the path from the power supply coming from the PCIe

---

Manuscript received November 19, 2010. The associate editor coordinating the review of this letter and approving it for publication was J. Murphy.

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

K. Christensen is with the University of South Florida, 4202 East Fowler Avenue, ENB 118, Tampa, FL 33620, USA (e-mail: christen@csee.usf.edu).

This work was partly supported by a Google Research Award.

Digital Object Identifier 10.1109/LCOMM.2011.040111.102259

connector to the NIC. The voltage drop is measured to estimate the current drawn by the NIC from which power consumption is computed. The voltage drop and power consumption at the resistor are negligible compared to the supply voltage and NIC power consumption.

The NICs were installed in two Dell Optiplex desktops running Linux and the device driver was modified such that EEE could be enabled or disabled from the command line using the ethtool utility. The two computers where connected directly using cables of different lengths (5m, 30m, and 60m). The power consumption was measured and was almost the same at all lengths (less than 2% variation). Therefore a single cable length (5m) was used in the experiments. Note that this invariant power consumption may not be the case for 10GBASE-T NICs where power consumption is expected to increase with cable length as the standard specifies smaller transmit signal levels when the cable is short and many designs will be optimized to reduce power consumption at short lengths [6].

In some of the experiments, the Linux traffic control (tc) tool [7] was used to limit the data rate on the link so that different loads could be easily tested. Traffic was generated using the ping utility or by transferring files between shared folders in the two computers.

III. EXPERIMENTS AND ANALYSIS

As discussed in the previous section, the NICs used in the experiments implement three different PHY standards corresponding each to a different link speed. The NIC power consumption increases with link speed and when the link is initially set up, the highest speed supported by both ends is selected [8]. This makes the highest supported speed the one most relevant from an energy savings point of view. Therefore, the experiments will focus on 1 Gb/s (1000BASE-T). For 10 Mb/s (10BASE-T), the original standard does not require the physical layer device to be active when there is no data to transmit. Thus, there is no EEE low power mode for 10 Mb/s and no experiments have been done for 10BASE-T. While of less interest, some results on the savings obtained when using EEE in 100 Mb/s (100BASE-TX) links are also presented here for completeness.

The first experiment tries to characterize power consumption for two corner cases: no traffic on the link and a link with close to full load. A large file transfer between folders in the two computers is used to generate the load. These bulk data transfers typically use large packets (around 1500 bytes long). The measurements are done at 100 Mb/s and at 1 Gb/s both with EEE enabled and disabled. The results are summarized in Table II. The first observation that can be made is that the absolute power consumption is at most around 0.5 Watts, significantly lower than that reported in previous measurements [5], [8]. This is probably due partly to the fact that the device used in the NIC is manufactured in a relatively advanced technology (110nm) that requires less power consumption than older technologies. Another factor is that the NIC does not implement advanced functionality found in other NICs targeted for servers [5].

Focusing on the use of EEE, it can be observed that when there is no traffic the NIC power consumption is reduced significantly (by over 70%) when using 1 Gb/s and to a lesser extent (by over 30%) for 100 Mb/s. This can be explained as the elements in the NIC that do not benefit from EEE (MAC, PCIe, etc.) are a larger part of the power consumption when using 100BASE-TX (and the PHY is a smaller percentage of the NIC power consumption). The savings are smaller than the ones used in previous simulation studies [3] where the reduction in PHY power consumption of the EEE low power mode was estimated to be 90%. This can be explained as that 90% corresponds to the PHY power consumption while our measurements are for the whole NIC that includes other elements such as the MAC and the PCIe interface that may not benefit from the use of EEE. This reasoning is reinforced by the fact that reductions of 90% or even larger for EEE low power mode in a 1 Gb/s PHY have been recently confirmed by a leading vendor [9].

TABLE II
MEASURED NIC POWER CONSUMPTION [MW]

Protocol	No Traffic		Full Load	
	Legacy	EEE	Legacy	EEE
100BASE-TX	208	139	215	208
1000BASE-T	525	152	541	535

TABLE III
MEASURED NIC POWER CONSUMPTION WHEN SENDING 5000 250 BYTE PACKETS PER SECOND [MW]

Protocol	Legacy	EEE
100BASE-TX	215	201
1000BASE-T	531	512

For high traffic load there is little difference when EEE is used, as expected. This is because the link will always have data to transmit and therefore will not enter the low power mode.

The second experiment tries to estimate the power consumption during mode transitions. To do so the ping utility is used to generate 250 byte packets and the tc tool is used to limit the data rate to 10 Mb/s. This ensures that the link load is low (1% at 1Gb/s and 10% at 100 Mb/s) and that transitions occur frequently. The 10 Mb/s limitation ensures that packets are spaced by 200 μs such that the link is continuously transitioning between active and low power as the sum of the sleep (Ts) and wake (Tw) times is larger than the packet spacing. The results are shown in Table III. It can be observed that the power consumption with EEE is very close to the one without EEE. As the link load is low, this means that the power consumption during transitions is close to that of the active mode as assumed in previous studies [3]. Therefore, if packets arrive spaced in time a large number of transitions will occur leading to a large energy consumption.

The third experiment tries to measure the power consumption versus the link traffic load. For this the tc tool is used to limit the link data rate to different values and a large file transfer between folders in the two computers is done. The data rate limitation models a bottleneck link that is commonly found when accessing the Internet from a high-speed LAN. That bottleneck link will space packets causing frequent EEE mode transitions to occur.

The results are shown in Fig. 2. It can be observed that when the load reaches 6% the power consumption with and

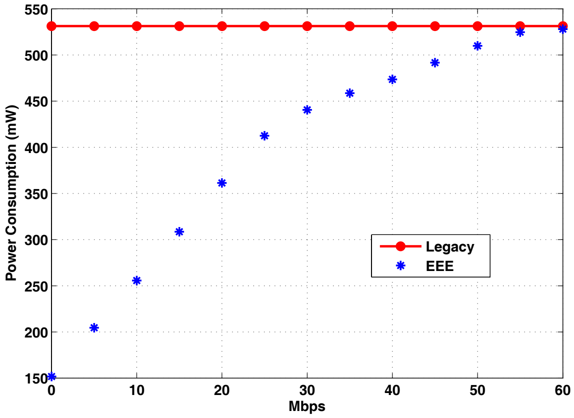

Fig. 2. NIC power consumption versus load for 1000BASE-T (1 Gbps).

without EEE are similar. This means that in many practical applications the transition overhead in EEE may be significant. This is also in line with simulation results based on the use of independent packet arrivals [3]. In fact the measured relative power consumption is larger than the one obtained with simulation. This can be attributed in part to the fact that the spacing enforced by a bottleneck is a worse case than independent arrivals and in part to the fact that in 1000BASE-T both link directions enter or exit the EEE low power mode at the same time (the simulations in [3] were done for a single link direction only). In our experiment this means that the link is active in both directions when sending data in one direction or acknowledgements in the other direction. This reduces the time spent in low power mode.

In summary, from the results obtained in the experiments the following conclusions can be made:

- The power consumption of 1 Gb/s Ethernet NICs seems to have been significantly reduced in the past few years to around half a watt.
- The use of EEE will significantly reduce energy consumption when the NIC is in low power mode. In our experiments the reduction was more than 70% for 1 Gb/s.
- The energy consumption during EEE mode transitions is close to that of the active mode.
- The energy overhead caused by EEE transitions can be very significant at low loads.

From those conclusions, the following tasks may be of interest to ensure that the energy savings when using EEE are maximized:

- The power consumption during mode transitions should be minimized. This needs to be addressed by the PHY manufacturers and opens an opportunity for innovation.
- The use of coalescing to reduce EEE transition overhead as proposed in [2] should be explored. In fact many NICs implement coalescing in the receive side to reduce the load caused by interrupts on the computer [10]. Some of those settings could be used to also improve energy savings when EEE is adopted. Unfortunately, the NIC used in our tests does not provide coalescing functionality [4]. Once EEE capable NICs that implement coalescing become available the combination of EEE with coalescing should be explored.
- To properly characterize the energy savings of EEE equipment, tests that include traffic patterns that exercise mode transitions should be done. Measuring the power consumption in low power mode is not enough. Energy Star work on testing procedures for small network equipment already proposes measuring the power consumption at different link loads [11].

Finally, it is worth noting that further improvements in energy savings can be expected in the future as manufacturers refine their implementations. Energy savings are expected to be greater for 10 Gb/s due to the smaller transitions times (compared to transition times for 1 Gb/s).

IV. CONCLUSIONS

In this letter, the energy savings obtained with the IEEE 802.3az Energy Efficient Ethernet (EEE) standard have been studied using NICs from a major Ethernet vendor. The measurements confirm previous analysis based on simulations and show that significant energy savings can be obtained for 1 Gb/s links, but they are largely dependent on the traffic patterns in the link. This is due to the large overhead caused by transitions between active and low power modes in EEE.

The analysis of the results suggests areas where work can be done to maximize energy savings. This includes the use of coalescing, also studied in simulation and that could be tested once EEE capable NICs that implement coalescing are available. Another option to maximize savings is modify the physical layer devices to minimize the power consumption during transitions.

ACKNOWLEDGMENT

The authors would like to thank Joseph Chou, Angus Lee and the NIC field application engineering team from Realtek for providing us with the EEE NICs and for their support during the experiments.

REFERENCES

[1] IEEE Std 802.3az: Energy Efficient Ethernet-2010.
[2] K. Christensen, P. Reviriego, B. Nordman, M. Bennett, M. Mostowfi, and J. A. Maestro, "IEEE 802.3az: the road to energy efficient Ethernet," *IEEE Commun. Mag.*, vol. 48, no. 11, pp. 50–56, Nov. 2010.
[3] P. Reviriego, J. A. Hernndez, D. Larrabeiti, and J. A. Maestro, "Performance evaluation of energy efficient Ethernet," *IEEE Commun. Lett.*, vol. 13, no. 9, pp. 697–699, Sep. 2009.
[4] RTL8111E Datasheet version 1.5, Realtek, July 2010.
[5] R. Sohan, A. Rice, A. W. Moore, and K. Mansley, "Characterizing 10 Gbps network interface energy consumption," *35th IEEE Conference on Local Computer Networks*, Oct. 2010.
[6] N. V. Bavel, D. Dove, A. Flatman, and M. McConnell, "Short haul 10Gbps Ethernet copper PHY call for interest," presentation at IEEE 802.3 Nov. 2005 meeting.
[7] M. A. Brown, "Linux traffic control HOWTO." Available: http://linux-ip.net/articles/Traffic-Control-HOWTO, Dec. 2006.
[8] C. Gunaratne, K. Christensen, B. Nordman, and S. Suen, "Reducing the energy consumption of Ethernet with adaptive link rate (ALR)," *IEEE Trans. Comp.*, vol. 57, no. 4, pp. 448–461, Apr. 2008.
[9] Private communication with N. Fitzgerald (LSI Ethernet Team), Nov. 2010.
[10] S. Makineni, R. Iyer, P. Sarangam, D. Newell, L. Zhao, R. Illikkal, and J. Moses, "Receive side coalescing for accelerating TCP/IP processing," in *Proc. HiPC*, Dec. 2006, pp. 289–300.
[11] "ENERGY STAR® test method for small network equipment," Draft, 4 Feb., 2011.