IEEE 802.3az: The Road to Energy Efficient Ethernet

Ken Christensen, University of South Florida
Pedro Reviriego, Universidad Antonio de Nebrija
Bruce Nordman and Michael Bennett, Lawrence Berkeley National Laboratory
Mehrgan Mostowfi, University of South Florida
Juan Antonio Maestro, Universidad Antonio de Nebrija

ABSTRACT

Ethernet is the dominant wireline communications technology for LANs with over 1 billion interfaces installed in the U.S. and over 3 billion worldwide. In 2006 the IEEE 802.3 Working Group started an effort to improve the energy efficiency of Ethernet. This effort became IEEE P802.3az Energy Efficient Ethernet (EEE) resulting in IEEE Std 802.3az-2010, which was approved September 30, 2010. EEE uses a Low Power Idle mode to reduce the energy consumption of a link when no packets are being sent. In this article, we describe the development of the EEE standard and how energy savings resulting from the adoption of EEE may exceed $400 million per year in the U.S. alone (and over $1 billion worldwide). We also present results from a simulation-based performance evaluation showing how packet coalescing can be used to improve the energy efficiency of EEE. Our results show that packet coalescing can significantly improve energy efficiency while keeping absolute packet delays to tolerable bounds. We are aware that coalescing may cause packet loss in downstream buffers, especially when using TCP/IP. We explore the effects of coalescing on TCP/IP flows with an ns-2 simulation, note that coalescing is already used to reduce packet processing load on the system CPU, and suggest open questions for future work. This article will help clarify what can be expected when EEE is deployed.

INTRODUCTION

The issue of energy consumption in network equipment came to the forefront in the early to mid-2000s [1]. Ethernet is the dominant wireline technology for LANs and is widely used in residences and in commercial buildings. Almost all notebook, desktop, and server computers include an Ethernet connection and in some cases more than one. Consumer entertainment equipment, such as television sets, increasingly include Ethernet connections. Ethernet is currently also being considered for use in access networks, and even for long-haul links. Four different data rates are currently supported in Ethernet using Unshielded Twisted Pair (UTP) as a transmission medium: 10 Mb/s (10BASE-T), 100 Mb/s (100BASE-TX), 1 Gb/s (1000BASE-T), and 10 Gb/s (10GBASE-T). To achieve the 10× increment, each new data rate uses better UTP cable such that a larger bandwidth can be used for transmission. Each data rate was standardized in IEEE 802.3 and uses different modulation and coding requiring different receiver architectures. 1000BASE-T, and especially 10GBASE-T, interfaces are extremely complex mixed signal integrated circuits that include adaptive equalizers, echo and crosstalk cancellers, advanced coding techniques, pre-equalization in the transmitter, etc. This complexity is needed to achieve high data rates over 100 m of UTP cable with a very low bit error rate and requires significant power consumption when all the elements in the device are active.

For 100 Mb/s and higher data rates, Ethernet transmitters transmit continuously, and when there is no data they transmit an auxiliary signal called IDLE that is used to keep transmitters and receivers aligned. This means that most of the elements in the interfaces are active at all times leading to a large energy consumption. The increase in the complexity of the interfaces also implies that more energy will be consumed when the data rate is higher. A 1000BASE-T Ethernet physical layer transceiver (PHY) typically consumes over 0.5 W while a 10GBASE-T PHY is usually over 5 W (all power values reflect system effects on mains AC power consumption). These figures vary from manufacturer to manufacturer and also with each new technology generation, but the trend is clear that an increase in Ethernet data rate requires more power.

The widespread adoption and large installed base of Ethernet means that if the energy consumption of Ethernet interfaces can be substantially reduced, large energy savings will be obtained. This was the main motivation to develop the IEEE 802.3az Energy Efficient Ethernet

0163-6804/10/$25.00 © 2010 IEEE

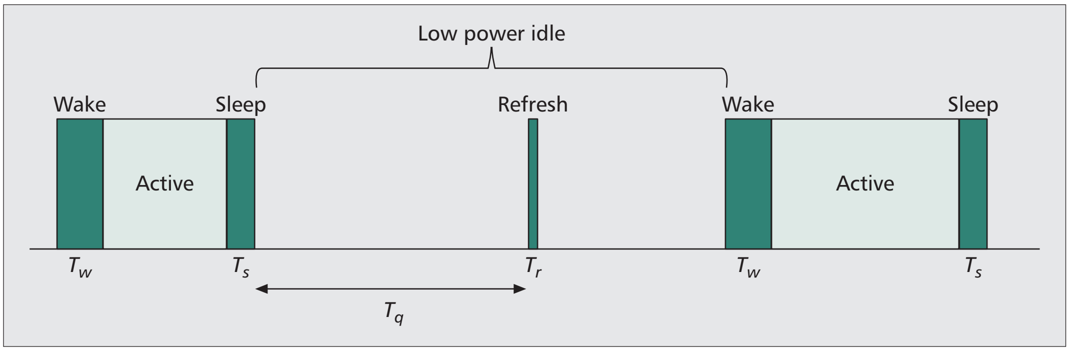

Figure 1. *Transitions between the active and low-power modes in EEE.*

(EEE) standard [2, 3]. The approach in EEE is to limit transmission when there is no data to short periodic refresh intervals to maintain alignment between the transmitter and receiver. The IEEE Std 802.3az-2010 focuses on Ethernet transceivers that operate over UTP, which account for the vast majority of Ethernet links. The standard defines mechanisms to stop transmission when there is no data to send and to resume it quickly when new packets arrive. This is done by introducing the concept of Low Power Idle (LPI), which is used instead of the continuous IDLE signal when there is no data to transmit. LPI defines large periods over which no signal is transmitted and small periods during which a signal is transmitted to refresh the receiver state to align it with current conditions. Large energy savings are obtained when the device spends a significant fraction of the time in the low power mode. Although the savings vary from device to device, the energy consumption when the device is in low power mode can be as low as 10 percent that of the active mode. During the transitions in and out of low power mode there is significant energy consumption as many elements in the transceiver have to be active. The actual value will depend on the implementation possibly ranging from 50 percent to 100 percent of the active mode energy consumption.

EEE operation is illustrated in Fig. 1. When packets are being sent, the device is in the active mode, and when no further packets are available for transmission the link may enter the low power (or sleep) mode; the transition to low power mode requires T_s seconds. Once in the low power mode the device only sends signals during short refresh intervals T_r and stays quiet during large intervals T_q. Once packets arrive for transmission, the link is activated again; this wake transition takes T_w seconds. Once the link is active, one or more packets can be sent. Thus, EEE adds an overhead to every burst of one or more back-to-back packets sent by an Ethernet interface. The overhead is the time to wake the link from idle (T_w) and to put the link back to low power mode (T_s). The values of T_w and T_s for 10 Gb/s in IEEE Std 802.3az-2010 are 4.48 μs and 2.88 μs, respectively. In comparison, for 10 Gb/s the packet transmission time, T_{pkt}, is 1.2 μs for a 1500 byte packet (a typical packet size for data transfer using TCP) and 0.0512 μs for a 64 byte packet (a typical packet size for sending TCP ACKs).

Development of IEEE 802.3az Standard

The development of any standard in IEEE 802.3 begins with an idea. Moving an idea to publication of an IEEE standard requires following an iterative process of evaluating and selecting proposals, writing, review, and revision. It also involves consensus building. The idea of creating a standard for Energy Efficient Ethernet began with a tutorial presented to members of the IEEE 802 Working Group in July 2005 [4]. Through the months that followed that tutorial, many discussions ensued, meetings were held, and presentations were given, leading up to the first step to getting a project started in IEEE 802.3 known as the Call for Interest (CFI). In November 2006, a panel presentation was given to the IEEE 802.3 Working Group [2] describing the rationale for forming a study group to determine the need for an EEE project, which was followed by a successful CFI vote. The first meeting of the EEE Study Group was held in January 2007. The purpose of the study group was to determine whether or not to request a project to create a standard for EEE. During that phase of the work, there were six meetings and 43 presentations supporting the creation of a standard for EEE. Contributions covered technical and economic feasibility of developing EEE, an estimate of broad market support, and compatibility with existing Ethernet devices as these are needed to meet the *five criteria* requirement to start a project for a new standard. Since EEE was the first project in the history of IEEE 802.3 to focus on energy efficiency, the distinct identity criterion was simple to meet. Another job of the study group was to write a Project Authorization Request (PAR) that defines the purpose, need and scope of the project. In addition to the PAR, the study group produced objectives, which define what the task force would work on, such as the media types. The project was authorized in September 2007 and the P802.3az Task Force was formed.

The job of the task force is to produce a standard. Technical choices are made during this phase of the project. Two methods were proposed: Adaptive Link Rate [4], that is, quickly switch to a lower speed during periods of low link utilization, and Low Power Idle, which operates as previously described. Each method saves energy and has its advantages and disadvantages

The development of any standard in IEEE 802.3 begins with an idea. Moving an idea to publication of an IEEE standard requires following an iterative process of evaluating and selecting proposals, writing, review, and revision. It also involves consensus building.

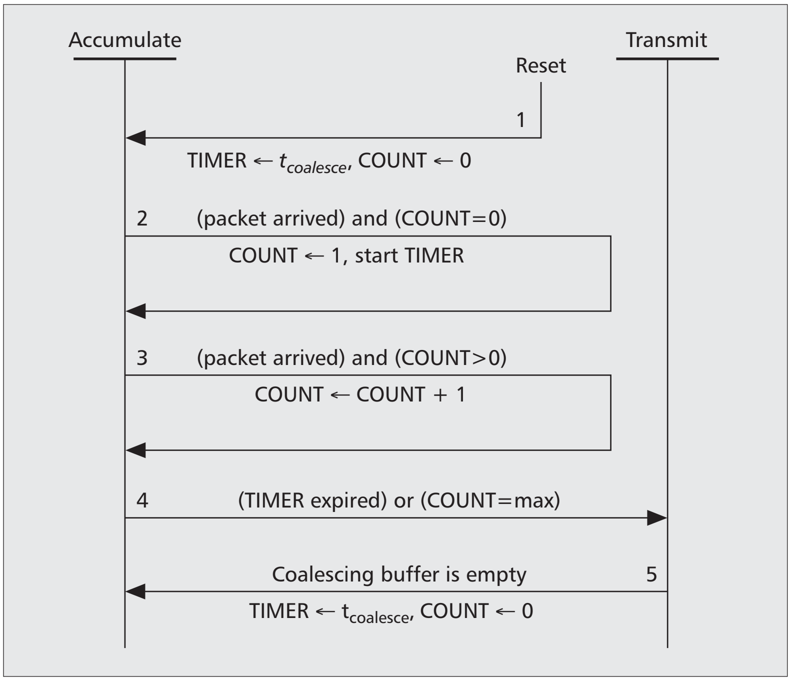

Figure 2. *Finite state machine for packet coalescing.*

which were developed, reviewed, and debated until LPI was eventually selected. The task force produced its first draft of the standard in October 2008 and completed selection of baseline proposals in March 2009. While this article focuses on one of the media types specified for EEE, in addition to the *twisted-pair* media, backplane, the XGMII Extender Sublayer (XGXS) and 10 Gigabit Attachment Unit Interface (XAUI) are also included in the standard.

The Working Group balloting phase of the development process was completed in June and the next phase of the project, Sponsor Ballot, was completed near the end of August. This was the last major review of the standard. The balloting group included a different group of reviewers than in the previous reviews of the document and the scope of review was the entire document. The review and revision process was repeated as a final refinement of the standard. The Standards Board Review Committee ensured that the development process was followed correctly and that the contents match the scope and purpose stated in the PAR. IEEE Std 802.3az was approved by the Standards Board on September 30, 2010.

Performance Trade-offs in Energy Efficient Ethernet

In this section, we explore the effects of EEE overhead on energy use and evaluate how packet coalescing at the sender can improve energy efficiency. The improvement in energy efficiency comes with a trade-off in performance — this trade-off is explicitly evaluated. The material in this section builds on previous work in [5] where the effects of EEE overhead on energy efficiency were first explored and in [6] where the notion of a packet burster was first proposed and evaluated for limited cases.

The energy efficiency of EEE is a function of the link utilization, packet transmission time, and the distribution of packet interarrival times. As discussed before, it is conservative but reasonable to assume that EEE overhead consumes the same power as packet transmission. In the best case, a large block of packets is sent back-to-back (that is, with no gaps between them), and the EEE overhead is negligible compared to the overall transmission time of the burst. In this case, the energy use is almost the same as the overall offered load, or link utilization. In the worst case, small packets are sent periodically with a gap between them. In this case, every packet will have a T_w and T_s overhead added resulting in an energy use potentially much greater than the link utilization. We find that best case traffic occurs often in the form of file downloads using TCP where large blocks of data are burst onto a link (for example, from a server to a client) at a very high rate. However, worst case traffic occurs in the form of TCP ACKs being returned from the downloading client to the server. These TCP ACKs are typically small packets and are spaced-out evenly. The distribution of aggregated traffic on a link is a subject of much research; however, a Poisson distribution remains a reasonable first-order approximation in some cases [7] and also serves as an intermediate case between the best and worst cases that we also consider in this section.

A key question is what range of link utilizations is of interest to consider? Many studies have found that the utilization of networks links, especially at the edge, is generally very low. In [8] it is argued that the usage patterns of networks — one of a need for high-speed bulk data transfer on demand — drives link utilizations to be low. A utilization level of 1 percent for LANs is described in [8]. While Ethernet links today are mostly 100 Mb/s and 1 Gb/s, we believe that they will evolve to 10 Gb/s in the future (and remain lightly utilized as are current links).

EEE performance can be improved by packet coalescing in which a FIFO queue in the Ethernet interface (in the host NIC and switch or router line card) is used to collect, or coalesce, multiple packets before sending them on a link as a burst of back-to-back packets. Packet coalescing is already used in many high-speed Ethernet interfaces — mostly on the receive side — to reduce CPU overhead for packet processing [9]. Coalescing can be based on packet count and/or time from first packet arrival. Figure 2 shows the finite state machine of a packet coalescer. In this design, the first packet to arrive to an empty coalescing queue (in state Accumulate) starts a timer (set to $t_{coalesce}$) and a packet counter. Once the maximum packet count (*max*) is reached or the timer expires, all packets in the coalescing queue — and those arriving when in the Transmit state — are sent. Coalescing should in most cases not require any additional buffer memory to be added, it is a matter of when packets that would otherwise be queued in system memory (for example, under control of the Ethernet NIC device driver) are released to the NIC for transmission.

Figure 3a shows the normalized energy use of EEE as a function of link utilization. The results

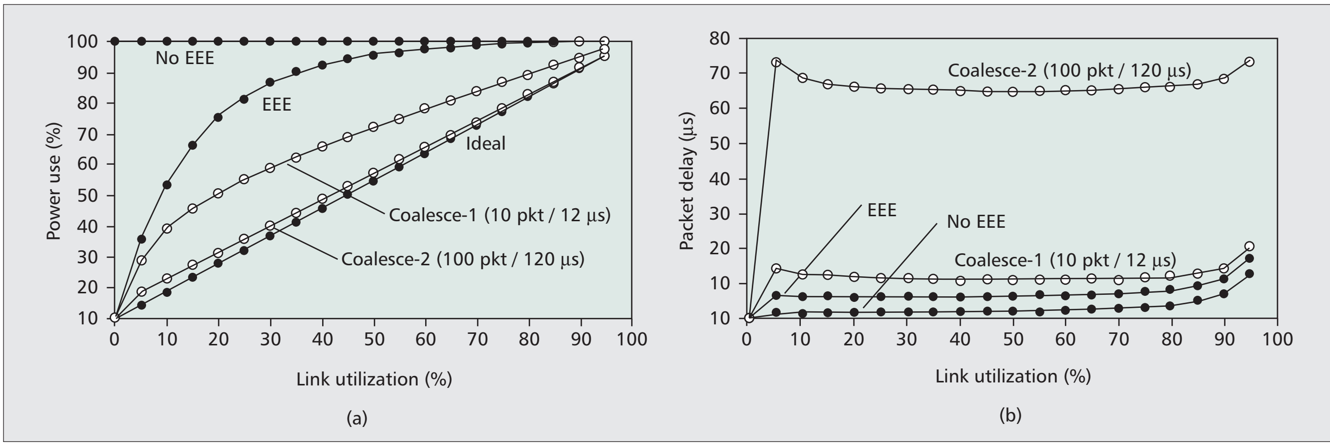

Figure 3. *a) Energy use vs. link utilization; b) packet delay versus link utilization.*

show a simulated 10 Gb/s Ethernet link with packets arriving as a Poisson process. All packets are a fixed length of 1500 bytes (T_{pkt} = 1.2 µs). It is assumed for presentation that the power consumption of an idle link is 10 percent of that of an active link as in previous analysis [5]. For no EEE, the constant power use would be 100 percent and independent of link utilization. The trace labeled as ideal shows the ideal case where power use is directly proportional to utilization with an offset due to the 10 percent power consumption in idle mode. This case would occur if EEE overhead was zero. The trace labeled as EEE shows power use of EEE without coalescing. For 5 percent offered load it can be seen that power use is about 35 percent. This increase from ideal (ideal would be about 15 percent power use) is due to EEE overhead. The two traces labeled as coalesce-1 and coalesce-2 show the results for coalescing with $t_{coalesce}$ = 12 µs and *max* = 10 packets, and with $t_{coalesce}$ = 120 µs and *max* = 100 packets, respectively. It can be seen that with coalesce-1 the energy efficiency is improved about *half way* between EEE and ideal (proportional), and with coalesce-2 the energy efficiency is very close to ideal. This improvement in energy efficiency comes with a cost of increasing packet latency and increasing the relative burstiness of traffic sent by the Ethernet interface. Figure 3b shows the mean packet delay for no EEE, EEE, EEE with coalesce-1, and EEE with coalesce-2. It can be seen that as the coalescing parameters ($t_{coalesce}$ and *max*) are increased, the packet latency is also increased.

What is the significance of the increased packet delay? For an end-to-end connection across the Internet the round-trip-time (RTT) would be tens to hundreds of milliseconds. An increase of a few tens of microseconds would likely be negligible. In a data center, any increase in packet latency could be significant but so would be additional energy savings. Where performance is critical in terms of a need for very low latency, special care should be taken when using EEE and coalescing. An example of such an application of Ethernet is in high energy Physics test beds, as described in [10].

Next we investigate the effects of EEE and EEE with coalescing on TCP file downloads focusing on high-speed connections with small RTTs. To further investigate the effects of packet coalescing, we implemented EEE and coalescing in an ns-2 simulation model. Figure 4 shows the configuration modeled for studying the performance of file download from server to client. The network in the middle was modeled with two delays — 40 µs and 400 µs (corresponding to a small and a large LAN respectively). The edge links were modeled with 10 µs delay in all cases (corresponding to short links within a data center or a link from office to wiring closet). Packet coalescing is done in all modeled sending interfaces (that is, in the client and server hosts and the two edge routers). The ns-2 TCP Linux agent and Sack1 receiver were used for the TCP connections with a maximum window size of 400 packets. The buffer size on the link between routers was set to 100 packets.

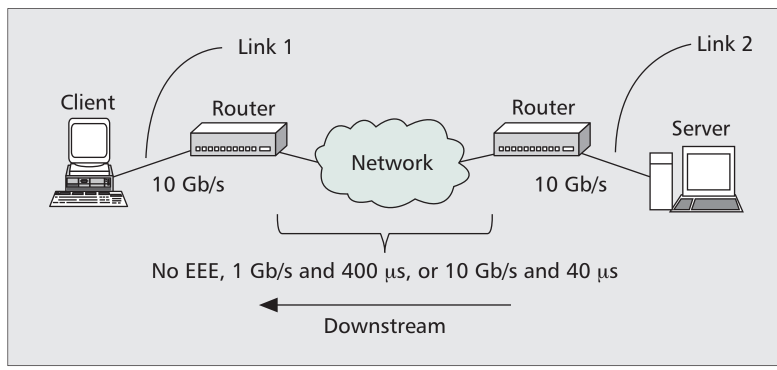

Figure 4. *Configuration for file download experiments.*

Experiments to measure the file download time for a 1 GB file were conducted with and without EEE (and with and without coalescing) in the hosts and edge routers. For coalescing, the parameters of coalesce-1 and coalesce-2 (previously described) were used. We also measured the energy use of links 1 and 2 for both upstream and downstream (where downstream is from server to client). As before, the energy use is the sum of all T_{pkt}, T_w, and T_s plus the time that the link is idle, T_{idle}, multiplied by 0.1. Table 1 shows the simulation results, the energy used is only shown for link 1 as similar values were obtained for link 2. It can be seen that the file download times are very similar in all cases with

		Energy use (100% maximum, 10% for idle)		Link utilization	
Config (10 Gb/s, 40 μs delay)	**Download time (s)**	**Link 1 up**	**Link 1 down**	**Link 1 up**	**Link 1 down**
No EEE	0.843	100.0%	100.0%	4.0%	94.9%
EEE	0.843	99.9	99.9	4.0	94.9
EEE coalesce-1	0.843	50.6	99.9	4.0	94.9
EEE coalesce-2	0.847	21.3	99.5	4.0	94.5
Config (1 Gb/s, 400 μs delay)	**Download time (s)**	**Link 1 up**	**Link 1 down**	**Link 1 up**	**Link 1 down**
No EEE	8.28	100.0%	100.0%	0.4%	9.7%
EEE	8.28	65.6	74.4	0.4	9.7
EEE coalesce-1	8.28	38.0	46.7	0.4	9.7
EEE coalesce-2	8.34	17.8	25.8	0.4	9.7

Table 1. *Results from ns-2 file download experiments.*

a slight increase when the larger coalescing timer is used. On the other hand, the energy savings are significant when compared with EEE especially in the link direction on which TCP ACKs are sent. For example for the 10 Gb/s configuration in steady state, ACKs would be sent each 1.2 μs leading to the link being active or in transition all the time if coalescing is not used. If coalesce-1 is used then the link would be activated after 12 μs then it would take 4.48 μs to complete the transition, and finally it would be deactivated as soon as the ACKs are sent. Shortly after, a new ACK will arrive starting a new cycle. This cycle would take approximately 18 μs of which 7.36 μs would be used for transitions. This would translate to over 45 percent of the time in active mode or transitions leading to an energy consumption of around 50 percent of the maximum in line with the results obtained in the simulations.

The results in Table 1 show no significant increase in file download time as energy efficiency is increased first with the addition of EEE, and then with the addition of coalescing to EEE. This is a significant result and suggests that coalescing can be an excellent complement to EEE to further improve the energy efficiency of Ethernet.

A deeper understanding of the issues caused by coalescing on TCP behavior is needed. For example, with coalescing ACK compression occurs when returning ACKs (from client-to-server in this case) are *bunched-up*. ACK compression has been long studied in other contexts and may cause increased burstiness of the TCP sender. This would be in addition to the burstiness created by coalescing in transmission. Increased burstiness could cause buffer overflows in downstream routers. For example, in the scenario considered, if the buffer size at the routers is smaller than the burst size then losses will occur when the link between routers operates at 1 Gb/s. In this case when the coalescing timer is 120 μs, bursts will be of approximately 11 packets as they are spaced by 12 μs on the 1 Gb/s link. If the buffer size is smaller than this value, then TCP slow start will end prematurely and, even worse, the congestion window (cwnd) will never exceed the buffer size. This is due to packets being sent as bursts of back to back packets when the congestion window is smaller than the burst size and coalescing is used. This causes buffer overflow that prevents the congestion window from increasing. More generally, coalescing will increase the RTT such that a larger TCP window is needed to achieve a given data rate. For example if a 120 μs coalescing timer with a packet limit of 100 is used, then at each coalescing point 120 μs of delay can be added increasing the window size for a 10 Gb/s connection by up to 100 packets. These larger windows translate to larger buffers at the sender, something that is already required to achieve good performance at high data rates [11].

From the experiments conducted it seems that if the two following conditions are met: the burst size is much smaller than the router and NIC buffers and the TCP window, and the burst timer is much smaller than the RTT, then the effects of coalescing on network performance will be small. In fact as mentioned before, coalescing is currently used in many NICs on the receive side to minimize interrupts to the CPU. Coalescing timers in the order of 25 μs to 125 μs are commonly used in NICs [9, 11]. These values are similar values to those used in our experiments. Therefore, it seems reasonable to assume that some degree of coalescing can be safely used to improve EEE performance. An optimal solution should combine the coalescing done today to reduce CPU overhead with the one proposed in this article to improve energy efficiency. This will require careful implementation, as accuracies on the order of one microsecond are needed to ensure that transitions or active periods on which no data is sent are minimized. The full effects of coalescing on TCP, whether used to reduce packet processing load on the system CPU or improve energy efficiency of an EEE link, require further study and is future work.

ECONOMIC BENEFITS FROM EEE

In 2011, EEE will begin to appear in products. Initially, it will save little energy, as most EEE products will be connected to legacy devices and so be unable to use EEE features. Over time, the percentage of Ethernet links using EEE will rise until it becomes the great majority. The energy saved in any given year is a function of many factors, most of which cannot be credibly forecast to the time when EEE savings become substantial. So, rather than create a forecast that is only one of many plausible scenarios, we instead present a savings estimate adapted from current conditions. This provides an order-of-magnitude savings estimate.

Energy use estimates are usually a function of the number (stock) of devices in question, their power levels in various operating states, and the usage pattern among those states. Our estimate distinguishes between 1 Gb/s and 10 Gb/s future links, and between edge links and non-edge links. We make separate estimates for 1 Gb/s and 10 Gb/s Ethernet links, with all edge links in both residential and commercial buildings and network equipment links within residences at 1 Gb/s. All data center links and uplinks from rack-based switches are 10 Gb/s. This assumes a substantial increase in data rates for links over today's conditions. For edge links, the number is the average active at full rate over the course of the year, as many edge devices are powered down some of the time. Stock and port are figures from network equipment sales data and from estimates prepared for EPA Energy Star for 2008. To summarize, our projected energy savings calculations use 2008 link counts, increase data rates, use current power levels, and maintain the assumption of low utilization.

Table 2 shows key assumptions and results. For power, the absolute amount of consumption does not enter the estimate; the savings per link is the difference between a link being fully active and one constantly in low power mode (including refreshes). The 5 W savings per link for 10 Gb/s assumes that a substantial power reduction will be achieved for 10 Gb/s PHYs in the next few years as the technology matures. The simulation model in this article and the model for 1000BASE-T in [5] are used to convert link utilizations (in data sent as a percentage of link capacity) to time in low power mode. For low utilization levels, most frames require transitions between LPI to active mode that effectively reduce low power mode time by almost 20 percent (resulting in per link savings of about 80 percent). The power savings per link are multiplied by the total number of active links to arrive at the U.S. total. The U.S. savings is $410 million/year, and global savings should be several times this, or over $1 billion per year. Additional savings also will also accrue from reductions in power and cooling energy in conditioned spaces like data centers, and from use of the Link Layer Discovery Protocol (LLDP) to negotiate longer wake transitions that enable savings beyond the PHY. In May 2007, a savings estimate presented to the IEEE 802.3az task force was 7.5 TWh/year; it had substantially higher per-link savings but lower link counts. These savings take the number of Ethernet links from those in use in 2008. It is possible that this number will rise substantially, particularly if Ethernet becomes widely used in distributing audio and video, in homes and elsewhere. Note that due to higher data rates, it is possible that the savings will exceed current energy used by Ethernet PHYs.

	1 Gb/s	10 Gb/s	Total
Assumptions			
Savings per link — no data (W)	1	5	—
Link utilization (%)	1%	3%	—
Active links (millions)	250	65	315
Electricity cost ($/kWh)	0.10	0.10	—
Results — EEE Savings			
Per link (%)	81%	82%	—
Per link (W)	0.81	4.10	—
Total (MW)	200	270	470
Total (TWh/year)	1.8	2.3	4.1
Total (million $/year)	180	230	410
Results — Ideal Savings			
Total (TWh/year)	2.2	2.8	5
Total (million $/year)	220	280	500
Coalescing Opportunity			
% (last column is average)	22%	18%	20%
Total (TWh/year)	0.39	0.43	0.82
Total (million $/year)	40	40	80

Table 2. *Assumptions and results for EEE savings (U.S. only, 2008 stock).*

The results in Table 2 (note that some of the totals do not add due to rounding) show substantial savings from EEE even without coalescing (the *EEE Savings* section of the table). Use of coalescing can gain most of the additional savings that an ideal link with zero EEE transition times (the *Ideal Savings* section of the table) would allow. This could be about $80 million in additional savings per year depending on overhead. In computing the overhead, many factors are involved including packet sizes, packet inter-arrival patterns, utilization variation with time, and PHY power consumption during transitions. The estimates in this article are for large frames, independent arrivals, no load variations and 100 percent PHY power consumption during transi-

The adoption of the new IEEE 802.3az Energy Efficient Ethernet standard will result in large energy and economic savings likely exceeding $400 million per year in the US alone. Those savings can be maximized using packet coalescing.

tions. The frame size used tends to underestimate overhead while using no utilization variations and transition power of 100 percent do the opposite. Therefore, the results presented should be a reasonable first order approximation to EEE overhead.

Conclusions

The adoption of the new IEEE 802.3az Energy Efficient Ethernet standard will result in large energy and economic savings likely exceeding $400 million per year in the US alone. Those savings can be maximized using packet coalescing if a careful analysis of the trade-offs between energy consumption and network performance is considered. In this article, these trade-offs have been studied for the case of 10 Gb/s links showing that in many cases increased energy savings can be achieved with a small impact on network performance by using packet coalescing in transmission.

Acknowledgment

The authors thank the anonymous referees for their helpful suggestions that have greatly improved this article. This research was partially supported by a Google Research Award and also partially supported by the U.S. Department of Energy under Contract no. DE-AC02-05CH11231. This material is partially based on work supported by the National Science Foundation under Grant No. 0520081 (Christensen).

References

[1] M. Gupta and S. Singh "Greening of the Internet," *Proc. ACM SIGCOMM*, Aug. 2003, pp 19–26.
[2] H. Barrass *et al.*, "Energy Efficient Ethernet Call-for-Interest," 2006; http://ieee802.org/3/cfi/1106_1/EEE-CFI.pdf.
[3] IEEE P802.3az, "Energy Efficient Ethernet Task Force," 2010; http://grouper.ieee.org/groups/802/3/az.
[4] K. Christensen and B. Nordman, "Reducing the Energy Consumption of Networked Devices," 2005; http://ieee802.org/802_tutorials/05-July/Tutorial%20July%20Nordman.pdf.
[5] P. Reviriego *et al.*, "Performance Evaluation of Energy Efficient Ethernet," *IEEE Commun. Letters*, vol. 13, no. 9, Sept. 2009, pp. 1–3.
[6] P. Reviriego *et al.*, "Burst Transmission in Energy Efficient Ethernet," *IEEE Internet Comp.*, vol. 14, no. 4, July/Aug. 2010, pp. 50–57.
[7] T. Karagiannis *et al.*, "A Nonstationary Poisson View of Internet Traffic," *Proc. IEEE INFOCOM*, 2004.
[8] A. Odlyzko, "Data Networks are Lightly Utilized, and will Stay that Way," *Rev. Net. Economics*, vol. 2, no. 3, Sept. 2003, pp. 210–37.
[9] S. Makineni *et al.*, "Receive Side Coalescing for Accelerating TCP/IP Processing," *Proc. HiPC*, Dec. 2006, pp. 289–300.
[10] A. Barczyk and J.-P. Dufey, "High Performance Gigabit Ethernet Switches for DAQ Systems," *Proc. 14th IEEE-NPSS Real Time Conf.*, 2005, pp. 295–98.
[11] M. Bencivenni *et al.*, "High Rate Packet Transmission on 10 Gb/s Ethernet LAN using Commodity Hardware," *Proc. 16th IEEE-NPSS Real Time Conf.*, 2009.

Biographies

KEN CHRISTENSEN [SM] (christen@csee.usf.edu) is a Professor in the Department of Computer Science and Engineering at the University of South Florida. He received his Ph.D. in Electrical and Computer Engineering from North Carolina State University in 1991. His primary research interest is in green networks. He is a licensed Professional Engineer in the state of Florida and a member of ACM and ASEE.

PEDRO REVIRIEGO (previrie@nebrija.es) is an associate professor in computer science at Universidad Antonio de Nebrija. His research interests are fault-tolerant systems, signal processing, and computer networks. He has a Ph.D. in telecommunications engineering from Universidad Politécnica de Madrid. He's participated in the IEEE 802.3 standardization for 10GBASE-T and worked on the design of 1000BASE-T transceivers.

BRUCE NORDMAN (BNordman@LBL.gov) is a researcher in the Energy Analysis Department, Environmental Energy Technologies Division, at Lawrence Berkeley National Laboratory. He has degrees in Architecture and in Energy and Resources from UC Berkeley. His research interests span from energy use of electronics, networks and energy, building networks, low power mode energy use, low-voltage DC powering, and user interfaces. He advises the EPA Energy Star program and participates in many standards committees.

MICHAEL BENNETT (mjbennett@lbl.gov) is a Senior Network Engineer at Lawrence Berkeley National Lab with over 29 years of combined experience in electronics and enterprise networking. He co-manages the LBLnet Services Group, which is responsible for operating LBNL's Campus Area Network. He is an IEEE standards association member, a contributor of end-user's perspective to several Ethernet standards, and Chair of the IEEE P802.3az Energy Efficient Ethernet Task Force.

MEHRGAN MOSTOWFI (mostowfi@csee.usf.edu) is currently a Ph.D. student in the Department of Computer Science and Engineering at the University of South Florida. He received his B.Sc. and M.Sc. degrees in Computer Science from Beheshti University, Iran and the University of South Florida, in 2006 and 2010, respectively. His research interests include power management and performance evaluation of computer networks.

JUAN ANTONIO MAESTRO (jmaestro@nebrija.es) manages the Computer Architecture and Technology Group at Universidad Antonio de Nebrija. His research interests include signal processing and real-time systems, fault tolerance, and reliability. He has a Ph.D. in computer science from Universidad Complutense de Madrid.