14

Energy Efficiency in Ethernet

Pedro Reviriego[1], Ken Christensen[2], Michael Bennett[3], Bruce Nordman[3] and Juan Antonio Maestro[1]
[1] *Universidad Antonio de Nebrija, Madrid, Spain*
[2] *University of South Florida, Florida, USA*
[3] *Lawrence Berkeley National Laboratory, Vallejo, USA*

14.1 Introduction to Ethernet

Ethernet is the dominant technology for wired local area networks (LANs) used to connect servers and desktop computers. Ethernet has also been adopted for industrial and automotive applications, to connect audio and video equipment and in wireless and wireline access networks. This wide adoption has fostered the continuous development of Ethernet since it was first introduced in the 1970s. In the 1980s, the IEEE 802.3 Ethernet working group was formed to develop Ethernet standards within the IEEE 802 LAN/MAN standards committee. Since then, many Ethernet standards have been produced to cover different transmission media, increased data rates, link aggregation, virtual LANs, Power over Ethernet, congestion management, and other capabilities.

The Ethernet standards support a wide variety of transmission media that include coaxial cables, unshielded twisted pair (UTP), optical fibers, and backplanes [1]. Typically, optical fiber is used for high-speed links between network equipment, whereas UTP is used to provide edge device connectivity. Therefore, the majority of the Ethernet links use UTP as the transmission media. For each medium, different rates are supported, and the rates have traditionally followed a 10× increase from one standard to the next. The rate is clearly identified in the standard name. For example, for UTP, the following rates are supported 10 Mb/s (10BASE-T), 100 Mb/s (100BASE-TX), 1 Gb/s (1000BASE-T), and 10 Gb/s (10GBASE-T). In some of the latest standards, a 4× increase in rate is targeted. That is the case of the next standard for UTP that will target 40 Gb/s and for which a study group has been formed [2]. Standards for optical fibers commonly support higher data rates. For example, the IEEE formed a study group recently to determine whether or not to start a project to specify 400 Gb/s Ethernet. As the

Green Communications: Principles, Concepts and Practice, First Edition.
Edited by Konstantinos Samdanis, Peter Rost, Andreas Maeder, Michela Meo and Christos Verikoukis.

rate increases, so does the complexity of the transceivers. High-speed Ethernet transceivers use advanced modulation and coding techniques and are very complex devices. This results in high power consumption. To facilitate interoperability when different rates are supported, the standard defines a procedure known as auto-negotiation that allows link partners to negotiate the link rate and other parameters.

Ethernet networks are built around switches and routers that connect the different elements in the network [3]. There are different types of switches ranging from a simple four- or eight-port switch to a modular switch (in which line cards and ports can be added or removed) that can host hundreds of ports and provide advanced features such as virtual LANs or link aggregation. In all cases, the switch has a number of ports that may support different rates and media, and a fabric and logic to forward incoming frames to the appropriate output port. The number of ports is typically a multiple of four or eight as the physical layer transceivers (PHYs) are commonly grouped in quad or octal devices to reduce cost. In switches that aggregate traffic from a set of users or servers, it is common to have a large number of lower speed ports and a reduced number of higher speed ports known as uplinks over which the aggregated traffic is sent. The fabric and control logic is simple for small switches that provide only a basic set of features but can be very complex for high-end switches.

The elements that are connected by Ethernet are typically desktop computers and servers. In both cases, Ethernet is commonly implemented on the motherboard, which is known as LAN on Motherboard (LOM). Ethernet can also be implemented in a Network Interface Controller (NIC) card. This is less frequent and NIC cards are mostly used to provide additional Ethernet ports when needed. For servers it is common that the LOM provides multiple ports, usually a dual or quad port implementation. The Ethernet controller can implement advanced features such as receive side coalescing or TCP offload to reduce the processing required in the CPU [4]. The vast majority of the switches and computers in the market support at least 1 Gb/s data rates with an increasing number supporting 10 Gb/s.

The wide adoption and low cost of Ethernet has motivated its use in applications for which it was not originally designed. Examples include industrial applications, in which latency and reliability are critical [5], automotive applications [6], access networks, with 1 Gb/s and 10 Gb/s Passive Optical Networks (PONs), and audio and video equipment (which need synchronization and other features). There are also now Ethernet standards specifically for data center connectivity, where a need for higher speeds and low power consumption enables larger system integration. This can be achieved by designing transceivers optimized for reduced link lengths in data centers as the links are often short and predictable. The NGBASE-T standard, for example, will likely target only short reach channels [2].

Despite the wide adoption of Ethernet and its continuous evolution, the issue of energy efficiency was not formally considered in the standards until 2005 when a tutorial on energy efficiency was presented to the IEEE 802 community [7]. This led eventually to the creation of a study group and a subsequent task force that produced the IEEE 802.3az "Energy-Efficient Ethernet" standard [8]. The 802.3az standard was a milestone in terms of addressing energy efficiency not only for Ethernet but also for wireline communications generally.

In the rest of this chapter, first the key features of the Energy-Efficient Ethernet standard are presented. Then energy savings estimates are presented showing the benefits of the standard. Finally, future directions for energy efficiency are discussed covering current standardization efforts related to energy efficiency within IEEE 802.3 as well as other standards that have been influenced by Energy-Efficient Ethernet.

14.2 Energy-Efficient Ethernet (IEEE 802.3az)

The first Ethernet standards, such as 10BASE-T, had simple transceivers that transmitted only when there were frames to send (except for short periodic pulses to keep the link alive). Ethernet has evolved since then to support higher data rates using more complex transceivers. To achieve those rates, most subsequent Ethernet standards have specified a continuous transmission at the physical layer. When there is no user data to send, a control signal known as IDLE is transmitted to keep all the transceiver elements well aligned. This ensures that when user data arrives for transmission, the link is ready to send the data. Transmitting the IDLE signal typically requires the same power consumption as transmitting frames (i.e., sending user data), so the transceiver power consumption is roughly constant and independent of traffic load. Link utilization is usually low as shown in the measurements presented during the standardization process [9]. Therefore, making power consumption more proportional to the actual traffic load would result in large energy savings given the widespread use of Ethernet.

With this objective in mind, a tutorial was presented to IEEE 802 in July 2005 [7]. This presentation triggered substantial interest in the IEEE 802.3 community and led to a Call for Interest (CFI) in November 2006 [9]. A CFI is the first step to get a project started in IEEE 802.3. The CFI was approved and a study group was formed to determine whether or not to request a project to create a standard for Energy-Efficient Ethernet (EEE). The goal of the standard was to improve the energy efficiency of Ethernet transceivers. The study group phase included six meetings and 43 presentations supporting the creation of a standard for EEE. Contributions covered technical and economic feasibility, indications of broad market support, and compatibility with existing Ethernet devices. These are needed to meet the “five criteria” requirement to start a project for a new standard. The study group wrote a Project Authorization Request (PAR) that defined the project’s purpose, need, and scope. The study group also produced objectives, which defined 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.

Two methods were proposed to improve energy efficiency: Adaptive Link Rate (ALR), which switches a link to a lower speed during periods of low utilization, and Low Power Idle (LPI), which creates a low-power sleep mode for use when there is no user data to send. ALR was proposed first and LPI later. Both methods save energy and have advantages and disadvantages. These were reviewed, and debated until LPI was eventually selected. The deciding factor was that ALR introduces a latency of hundreds of milliseconds which is not acceptable in many cases. This is so because the link has to be established at a lower/higher speed and therefore all the elements in the PHY have to be adjusted. On the other hand, LPI does not change the link speed and therefore only minor adjustments are needed in the PHY to enter/exit the low power mode. Those can be made in microseconds thus ensuring a low impact for most applications. In addition, it was thought that LPI might be somewhat easier to write in the standard and to implement.

The task force produced its first draft of the standard in October 2008 and completed selection of baseline proposals in March 2009. The EEE standard covers several media types including twisted pairs, backplanes, the XGMII Extender Sublayer (XGXS), and the 10-Gigabit Attachment Unit Interface (XAUI). Optical fiber is not covered in the EEE standard, though not for lack of interest. It is worth noting that at the very end of the study group phase, a proposal to include some fiber-optic transceivers was made; however, much more work was needed at the time to build the necessary consensus to add the work to the PAR. The study

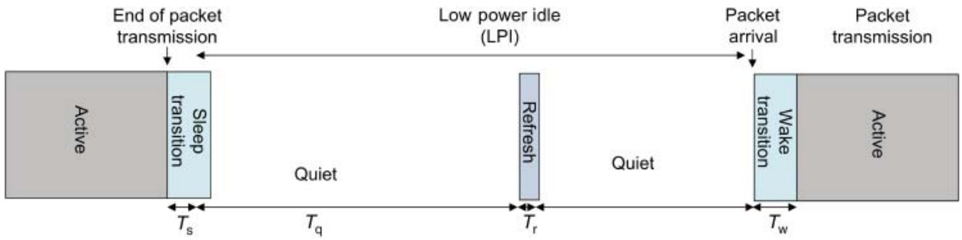

Figure 14.1 Illustration of the low power idle mode defined in energy-efficient Ethernet

Table 14.1 Minimum wake, sleep, frame transmission times, and single frame efficiencies for different link speeds

Protocol	Min T_w (μs)	Min T_s (μs)	Frame size (bytes)	T_{Frame} (μs)	Single frame efficiency (%)
100BASE-TX	30.5	200	1518	120	34.2
1000BASE-T	16.5	182	1518	12	5.7
10GBASE-T	4.48	2.88	1518	1.2	14.0

group had to choose between accepting the proposal and delay submitting the project request, which would add months to the overall timeline, or reject the proposal. They chose the latter. The EEE standard was approved as IEEE 802.3az [8] in September 2010.

The EEE standard defines an LPI mode to allow a transceiver to minimize transmission when there is no user data to send. Instead of transmitting continuously, transmission can be stopped with only short periodic refreshes sent to keep the transmitters and receivers aligned. As illustrated in Figure 14.1, sleep and wake transitions are needed to enter and exit LPI. The sleep transition ensures that the transceivers save relevant state while the wake transition ensures that they are aligned for reliable data transmission. In LPI, transmission is stopped except for short refresh intervals that keep a coarse alignment of the transmitters and receivers. Transceiver energy consumption in LPI is much smaller than in the active mode. However, packets can only be sent in the active mode and therefore a wake transition is needed before a packet can be sent.

Minimum transition times are specified in the standard and vary across the transmission media and interface rates. Table 14.1 shows the twisted pair transceiver values and the time required to transmit a maximum-length 1518-byte frame. Transitions times are larger than frame transmission times so that when packets arrive spaced in time and are transmitted isolated, one per active period, more time is spent in transitions than spent in actually sending the data. As power consumption during transitions can be similar to that of the active mode, isolated packets require considerably more energy to send than those sent in bursts. The transition times also impact latency as a frame that arrives for transmission to an interface that is in LPI mode has to wait until the link is back in active mode to be transmitted. Since transition times are on the order of microseconds, this is only an issue for latency critical applications. For example, in some datacenter applications, latency values of a few hundred microseconds

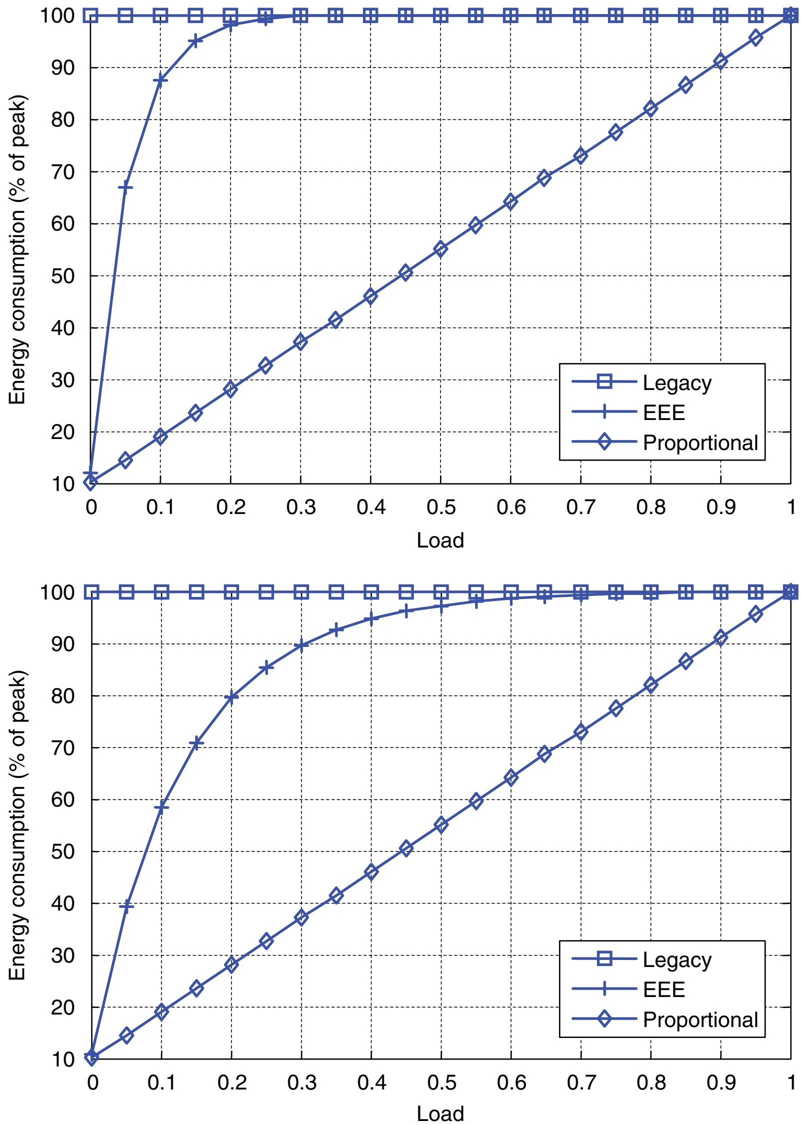

Figure 14.2 Energy consumption versus link load for 1000BASE-T (a) and 10GBASE-T (b)

are assumed [10]. In such cases, EEE could have an unacceptable impact on end-to-end delay, especially when a packet finds several links in LPI mode along its route.

Energy-Efficient Ethernet performance has been evaluated by simulation [11] and measurements [12]. Several analytical models have also been proposed and validated against simulations [13–15]. In all cases, the results confirmed that transitions limit the effectiveness of EEE when frames are transmitted isolated. Figure 14.2 shows the performance of a 1000BASE-T (a) and a 10GBASE-T (b) transceiver sending 1250-byte frames that arrive with a Poisson distribution, assuming that consumption in LPI mode is 10% that of the active mode as done in Ref. [11]. It is clear that performance deviates significantly from proportionality in which power draw is proportional to traffic load. To mitigate this issue, frames can be coalesced for transmission so that transition overheads are shared by multiple frames [16]. Coalescing is already used in computers to reduce the processor burden of sending and receiving frames. As coalescing increases latency, its applicability to latency critical applications is limited. However, in many applications, coalescing will not impact performance and is a natural feature of the information sent. For example, in video transmission, all the packets that belong to the same video frame can be coalesced for transmission [17].

The EEE standard also enables energy savings opportunities beyond those obtained in the transceivers [18]. For example, when all ports on a switch are in LPI mode, no traffic can arrive in the next microseconds. Switching logic and other elements within the switch can then be put into a low power mode until a wake transition occurs. This is only possible if the switching element wake times are smaller than those of EEE. The potential savings at system level can exceed those in the transceivers but may require wake times larger than those of EEE. To address this situation, the standard enables the link partners to use the Link Layer Discovery Protocol (LLDP) to negotiate transitions times larger than the minimum values. LLDP is an IEEE 802.1 link layer protocol that enables link partners to advertise parameters and exchange capabilities [19].

Complete network products with EEE began to appear in late 2010, shortly after the standard was ratified. By the end of 2012, EEE-enabled port shipments were measured in hundreds of millions/year, and most new computers and switches supported the standard. With typical equipment renewal cycles, it can be expected that in three to five years from 2013 most Ethernet links will implement EEE. At that point EEE will deliver substantial energy savings. Measurements on the first available EEE implementations show significant reductions in the power consumption of NICs [12] and switches [20]. The overall reduction in LPI mode is around 70% for NICs, whereas for the switches the results vary by model with savings between 50% and 75% (see details in Ref. [20]). The figures are for the entire NIC or switch. For the transceiver only, savings exceed 90% in some cases. In future products, the relative energy savings may increase as implementations of EEE are refined and the potential savings in other systems elements are implemented. Absolute savings may drop as reductions of active mode power consumption level are achieved in each new product generation.

14.3 Ethernet Energy Consumption Trends and Savings Estimates

The aggregated energy consumption of Ethernet devices, and the savings available from EEE, both depend on a number of factors:

1. The type and number (stock) of installed devices and the presence of Ethernet links between them.
2. The power consumption levels for each of them.
3. Link usage patterns.

All of these factors vary over time; hence, creating accurate estimates for each variable is difficult and they cannot be credibly forecast with reliability. So, rather than create a forecast that is only one of many plausible scenarios, we instead present a savings estimate based only

Table 14.2 1000BASE-T link counts by device type (millions; United States only, 2010)

	Device stock	Link-years
Residential		
Routers	53	53
Desktop PCs	101	40.4
Game consoles	109	38.2
Printers	113	21.7
Set-top boxes	45	11.3
AV receivers	99	7.4
TVs	353	7.1
Notebook PCs	132	6.3
VOIP adaptors	4.7	4.7
Switches	1	1
Blu-Ray players	12	0.8
Total	1023	191.8
Commercial		
VOIP phones	50	50.0
Desktop PCs	72.7	43.6
Printers	25	22.5
Notebook PCs	57.0	8.5
Total	205	124.7
TOTAL	1227	316.5

on current conditions. This provides an estimate of the order-of-magnitude for energy savings. We first present such an estimate and then delve into complications.

14.3.1 Number of Links

Table 14.2 shows our estimates for the number of active 1000BASE-T links today. A link-year is defined as one link, being active all the time, or ten links each being active for only 10% of the year. Interfaces without active links should consume little or no power. Stock data for network equipment are taken from Ref. [21], whereas other residential stock data are from Ref. [22]. The assumptions made in elaborating the table are described in the following. Set-top boxes cover only standalone boxes, not service provider equipment. Printers include multifunction devices. Many links between certain devices are constantly powered on and hence the link is always active; network equipment uplinks are an example. Some edge devices (e.g., many PCs) are powered down at certain times and so the corresponding links to these devices contribute to the active link count only fractionally. For some device types (e.g., TVs), only a portion of the stock support Ethernet at all, and out of such devices that support Ethernet, only a portion of the stock actually use it. All these factors are taken into account for estimating the average number of active links based on the amount of devices. In many cases, this relies on informed judgment of usage patterns in the absence of good data. Hence, while TVs are about a third of

Table 14.3 10GBASE-T link counts by device type (millions; United States only, 2010)

	Device stock	Links
Servers		
Volume	11.3	22.6
Mid range	0.3	1.2
High end	0.035	0.28
Storage	2.9	2.9
Switch uplinks	8.3	16.7
Total	1250	31.2

all relevant devices, they contribute less than 4% of the residential EEE savings. About 80% of residential EEE savings derive from routers, desktop PCs, game consoles, and printers.

For commercial buildings, a similar but simpler process is employed. Stock data is taken from U.S. EPA, Energy Star program, except the data for VoIP phones, which is estimated by the authors. Not all notebook PCs and imaging equipment have Ethernet. Some devices are powered down some of the time for example at night. Desktop PCs and VOIP phones dominate the savings; some phones have two ports with PC traffic combined with phone traffic in a single uplink to the switch. This does not change the link count from the case of separate switch ports to each device as in both cases there are two links.

Table 14.3 shows our estimates for the number of active 10GBASE-T links today. Server stock data are taken from Ref. [23], with servers having multiple ports (2, 4, and 8 for volume, mid-range, and high-end, respectively). Data storage equipment is assumed to contribute 0.25 ports/server on average (not all storage equipment uses Ethernet). Switch uplinks are mostly in commercial buildings as the number of PCs and VOIP phones far exceeds the number of servers. We assume 400 million ports, 48 ports/switch, and 2 uplinks per switch. We examined market research data on shipments of network equipment ports and concluded that are broadly consistent with our estimations, considering also that on most switches, some ports are not connected. As we count only links, not NICs, switch downlinks are accounted by the end device count. All data about equipment is based on data for 2010.

Network links other than edge links and switch uplinks that may also use EEE, but their number is relatively small in comparison, whereas their utilization rates are likely higher. Hence, our estimation is not taking into account these links. Similarly, there are many other applications of Ethernet including industrial purposes, cameras, displays, and niche applications, but these are modest in number compared to the categories in Table 14.2. Backplane Ethernet could be a significant application in the future, but our estimates concentrate on reporting from 2010.

14.3.2 Power per Link

Power consumption typically increases with the operational rate of the link. Considering Ethernet, the power consumption of successive product generations decreases due to deployment optimizations of the implementations and also because of the use of more advanced microelectronic technologies. For example, the first generations of 1000BASE-T transceivers consumed approximately 6 W, whereas current implementations need less than 0.5 W [24, 25]. Similarly for 10GBASE-T, there is a reduction from 10–12 W to 2.5–4 W between the first and second product generations [25]. All power values presented are AC-equivalent and hence are larger than low-voltage DC values since they account for conversion losses from AC/DC to DC/DC conversion.

At the outset of EEE work, a typical power for a 1000BASE-T NIC was 1 W. Data for a typical 1000BASE-T NIC shipping in 2012 from a major manufacturer shows a maximum power of about 0.45 W for a link under 30 m in length, whereas for a 100 m cable consumption is about 0.55 W but most cables are shorter so we use the lower figure.

Data on 10GBASE-T PHYs can be found in Ref. [26], which notes the high consumption of early hardware, that is, 10–12 W. New 40 nm 10GBASE-T PHYs consume 2.5–4 W for a 100 m cable, and even newer designs reduce power around 10% for 30–50 m cables, and over 25% for 10 m cables. Consumption is supposed to be less than 2.5–4 W due to the use of shorter cables, but then an increase is expected due to the power supply loss factor. Accordingly, we assume a 3 W per 10GBASE-T PHY for the 2012 technology version.

14.3.3 Usage Patterns

The calculations above of number of average links accounts for the fraction of time that links are established. The usage in question here is how much data is transmitted (and its burstiness). With the use of EEE, link utilization plays a fundamental role in determining energy savings. In most cases, communication networks will continue to be underutilized [27], as concluded in the study and measurements that supported the initiation and development of the EEE standard [9]. However, Figure 14.2 shows that for 10GBASE-T, EEE can consume about 60% of the peak power with a link load of only 10%, so in certain cases even low loads may affect energy consumption substantially. For 1000BASE-T links, the average load is expected to remain very low (around 1% or less), which simply means a sustained load of 10 Mb/s. For datacenter links, utilization is much higher [28] with values that exceed 10%. Since Internet service providers consider data utilization statistics as sensitive information, they are reluctant to share it. To keep the estimates simple an average load of 10% for 10GBASE-T links and 1% for 1000BASE-T links are assumed, respectively.

Consistent with Figure 14.2, we assume that LPI mode uses 10% of the peak power, whereas the corresponding savings with “No Coalescing” are around 79% for 1000BASE-T links and 42% for 10GBASE-T links, respectively.

To summarize, in our energy savings calculations we used 2010 active link estimations and consider the increased data rates based on today’s use. We also considered power levels corresponding to 2012 data utilization statistics, assuming a potential for underutilized network equipment and resources.

14.3.4 Results

Table 14.4 summarizes the estimated total energy saving using the active Ethernet link data introduced in Tables 14.2 and 14.3, considering also power per link values previously discussed (note that total savings sums are rounded). Simulation results, and specifically the ones illustrated in Figure 14.2, were used to convert link utilizations into durations of low power

Table 14.4 Assumptions and results for EEE savings (United States only, 2010 stock)

	1000BASE-T	10GBASE-T	Total
Assumptions			
Link-years	316.5	31.2	347.7
Peak power/link (W)	0.9	6.0	
Link utilization (%)	1%	10%	
No coalescing power	21%	58%	
Ideal savings (no transition time)			
Percent	89%	80%	
Per link (W)	0.80	4.80	
Total (MW)	254	150	403
Total (TW h/year)	2.2	1.3	3.5
Total ($million/year)	222	131	353
No coalescing savings			
Percent	79%	42%	
Per link (W)	0.71	2.52	
Total (MW)	225	79	304
Total (TW h/year)	2.0	0.7	2.7
Total ($million/year)	197	69	266
Coalescing opportunity			
Percent	10%	38%	
Per link (W)	0.09	2.28	
Total (MW)	28	71	100
Total (TW h/year)	0.2	0.6	0.9
Total ($million/year)	25	62	87
Moderate coalescing savings			
Percent	84%	61%	
Per link (W)	0.76	3.66	
Total (MW)	239	114	353
Total (TW h/year)	2.1	1.0	3.1
Total ($million/year)	210	100	310

mode. For low utilization levels, most frames require transitions from LPI to active mode, which significantly reduces low power mode time. Considering the power savings per link and the total number of active links in the United States, the cost savings is estimated at least $260 million/year, and over $310 million/year assuming moderate use of packet coalescing, whereas global savings are expected to be higher (an electricity rate of $0.10 is used, which approximates typical rates paid in the United States in 2010). Additional savings may also accrue from reductions in power and cooling energy in conditioned spaces such as data centers, and by using the Link Layer Discovery Protocol (LLDP) to negotiate longer wake transitions that enable savings beyond the PHY. In May 2007, a savings estimate [29] presented to the IEEE EEE Study Group was 7.5 TW h/year; it had substantially higher per-link savings but lower link counts. A savings estimate in [16] was 5 TW h/year that had higher numbers of 10 Gb/s links than this estimate. These savings take the number of Ethernet links from those in use in 2008. A recent study [30] estimates 2.8 TW h/year for 1000BASE-T alone for 2012.

The “Ideal Savings” section in Table 14.4 shows the potential energy savings via employing EEE if there was no energy consumption in transitions to and from the LPI mode; that is, each PHY uses the LPI power level plus the full rate level only for the percent of time when sending traffic. The “No Coalescing” section factors the energy consumed in the transitions to and from LPI mode, with fewer than two packets per active time. It shows substantial savings from EEE even without coalescing. The use of coalescing can ensure most of the additional savings an ideal link with zero EEE transition times would achieve. The “Moderate Coalescing” section shows savings considering double amount of packets in each active period, less than four considering both speeds. In computing the overhead, many factors are involved including packet sizes, packet interarrival patterns, utilization variation with time, and PHY power consumption during transitions. These estimates are for large frames, independent arrivals, no load variations, and 100% PHY power consumption during transitions. The frame size used tends to underestimate overhead, whereas using no utilization variations and transition power of 100% have the opposite effect. Therefore, the results presented are a reasonable first-order approximation to EEE overhead.

Further savings that are expected to occur over time depend on the number of links, which could rise substantially since Ethernet is currently continuously adopted widely in distributing audio and video, in homes and elsewhere. On the other hand, as wireless technologies increasingly enhance the offered capacity, reliability, and security, they could replace a significant segment of the current wired technologies such as Ethernet. In the near future, it is very likely for edge links in residential and commercial buildings to convert to 10 Gb/s Ethernet, with higher savings per link. Considering the equipment operational power consumption, some additional reduction is expected with the introduction of smaller semiconductor technologies, but the pace of reduction is likely to be significantly less than in the past. On the other hand, as traffic increases over time the actual utilization of Ethernet links goes up, and EEE savings will fall.

Global savings will be several times that in the United States. According to Ref. [23], the United States provides about 36% of the global stock of servers, and the International Energy Agency (IEA) [31] estimates that for all consumer electronics and ICT, North America consumes around 30% of the global total. Thus, global savings by our approach should be about $1 billion/year.

14.4 Future Directions of Energy Efficiency in Ethernet

The Energy-Efficient Ethernet standard is a milestone for improving energy efficiency in wire-line communications and will result in substantial energy savings in the coming years. EEE has also influenced other standards. The Fibre Channel community is working on a standard to add energy efficiency mechanisms to Fibre Channel transceivers similar to the LPI mode defined in EEE [32]. The same has occurred in the VDE 0885–763 standard for high-speed communication over plastic optical fibers [33]. In this case, the energy efficiency mechanisms have been designed as an integral part of the physical layer specifications. In IEEE 1904, Service Interoperability in Ethernet Passive Optical Networks (SIEPONs), a protocol similar to LPI is being developed to enable energy savings. The purpose of the SIEPON standard is to provide a system-level and network-level standard based on IEEE 802.3 EPON specifications, making it the first *system-level* specification for energy-efficient operation of Ethernet equipment [34]. IEEE 1904.1 refers to EPON standards developed in

802.3 for ways to save energy in the PHY and explains that further reduction in energy use may be achieved when combined with the use IEEE 802.3az interfaces. As of early 2013, six Energy Star specifications require that products supporting EEE must be tested connected to a device that also supports EEE, as this is necessary to have the measured power level reflect EEE savings. This encourages manufacturers to include EEE components. The Energy Star specification for small network equipment [35] goes further and takes into account savings at the other end of the link to provide more incentive to include EEE. The upcoming specification for large network equipment will most likely also incorporate EEE.

In addition to influencing standards beyond IEEE 802.3, EEE has also inspired incorporating energy efficiency in other Ethernet standards. Most IEEE 802.3 physical layer standards that followed the standardization of EEE specify energy efficiency mechanisms or they are being discussed as part of standards projects. One of the most compelling reasons to include an energy-efficient mode in new Ethernet projects at the start is that it is significantly easier to design the features at the beginning of the project, rather than at the end of the project where the risk of introducing new problems in the design is much greater. Such is the case in the IEEE P802.3bj project to define a 100 Gb/s standard for backplane and copper data center links. The cost of LPI mode transitions is very important in data center links as the load is significant and frame transmission times at 100 Gb/s are very small. This prompted the task force to define two LPI modes, an EEE-like mode in which physical layer transmission is stopped and a fast mode in which physical layer signaling continues, but savings can be achieved in other system elements [36]. The fast mode enables fast transition times that are beneficial when load is high or latency is critical. Energy efficiency mechanisms are also being incorporated in the latest IEEE 802.3 standards for optical fibers [37]. The IEEE P802.3bm standard for 40/100 Gb/s communication over optical fibers adopted the fast LPI mode from IEEE P802.3bj. This capability enables system energy savings while avoiding issues associated with stopping optical transmission [38]. The introduction of this new fast LPI mode increases the necessity for network equipment manufacturers to develop methods to take advantage of the signal to creatively reduce energy use in the system.

A key evolution in the acceptance of energy efficiency as a requirement for future projects in IEEE 802.3 was evident in the CFI for Reduced Twisted Pair Gigabit Ethernet (RTPGE) standards project, P802.3bp, which is the first project stating the need for energy efficiency explicitly. RTPGE needs a very low power standby mode in automotive applications [39], which may lead to having two low-power modes, one similar to LPI and another with larger energy savings and transition times. For the NGBASE-T project, there is a concern that power consumption will be high and an LPI mode will provide only limited savings due to the link speed (40 Gb/s) and the expected load in data centers [10]. Therefore, efforts are concentrating on optimizing the power use versus link length [26, 40]. Network equipment manufacturers have used this idea previously as an enhancement to the standards in their Ethernet products [41]. Incorporating it to the standard design opens a new avenue for energy efficiency in which the transceiver is designed to adapt its functionality and power consumption to the requirements of the channel. This results in significant energy savings that do not depend on the traffic load.

Clearly Energy-Efficient Ethernet has had an impact on improving energy efficiency in wire-line communications. Energy efficiency is now considered at the beginning of new project proposals in IEEE 802.3 standards development efforts. EEE has influenced, and will continue to influence, development of energy-efficient communications in other standards as well. With the development of new modes to achieve energy proportionality while mitigating operational impact, such as with fast LPI, comes the need to develop ways to accomplish the systems-level savings they enable. This need may drive the next cycle of innovation in energy efficiency for Ethernet as EEE becomes more routinely accepted.

14.5 Conclusions

In the last decade, energy efficiency in wire-line communication systems has changed from being practically disregarded to being a primary design objective in new communications standards. One example of this evolution is Ethernet for which the development of the EEE standard marked a milestone. This chapter presented first an overview of both Ethernet and the EEE standard. Then estimates and trends of energy consumption in Ethernet were discussed showing that the adoption of the new IEEE 802.3az Energy-Efficient Ethernet standard will result in large energy and economic savings likely exceeding $250 million per year in the United States alone. Finally, future directions and efforts to improve the energy efficiency in the Ethernet standards that are being developed were summarized. For the reader who wants to know more detailed information can be found in the references many of which are available online.

References

[1] C. Spurgeon, “Ethernet: The Definitive Guide,” O’Reilly Media, 2000.
[2] “IEEE 802.3 Next Generation BASE-T Study Group,” available online at http://www.ieee802.org/3/NGBASET/index.html.
[3] R. Seifert and J. Edwards “The All-New Switch Book: The Complete Guide to LAN Switching Technology,” Wiley, 2008.
[4] S. Makineni, R. Iyer, P. Sarangam, D. Newell, L. Zhao, R. Illikkal and J. Moses, “Receive side coalescing for accelerating TCP/IP processing,” Proceedings of International Conference on High Performance Computing (HiPC), pp. 289–300, 2006.
[5] J. D. Decotignie “Ethernet-based real-time and industrial communications” Proc. IEEE, vol. 93, no. 6, pp. 1102–1117, June 2005.
[6] “IEEE 802.3 Reduced Twisted Pair Gigabit Ethernet (RTPGE) Study Group” available online at http://www.ieee802.org/3/RTPGE/index.html.
[7] K. Christensen and B. Nordman, “Reducing the Energy Consumption of Networked Devices,” presentation to IEEE 802, 2005, available online at http://ieee802.org/802_tutorials/05-July/Tutorial%20July%20Nordman.pdf.
[8] “IEEE Std 802.3az: Energy Efficient Ethernet-2010”.
[9] H. Barrass, M. Bennett, H. Frazier and B. Nordman, “Energy Efficient Ethernet IEEE 802.3 Call-for-Interest,” 2006, available online at http://ieee802.org/3/cfi/1106_1/EEE-CFI.pdf.
[10] U. Hoelzle and L. A. Barroso, “The Datacentre as a Computer: An Introduction to the Design of Warehouse-Scale Machines,” Morgan and Claypool Publishers, 2009.
[11] P. Reviriego, J. Hernandez, D. Larrabeiti and J. Maestro, “Performance Evaluation of Energy Efficient Ethernet,” IEEE Commun. Lett., vol. 13, no. 9, pp. 1–3, 2009.
[12] P. Reviriego, K. Christensen, J. Rabanillo and J.A. Maestro, “An initial evaluation of energy efficient Ethernet,” IEEE Commun. Lett., vol. 15, no. 5, 2011, pp. 578–580.
[13] S. Herrería-Alonso, M. Rodríguez-Pérez, M. Fernández-Veiga, and C. López-García, “How efficient is energy-efficient Ethernet?,” in Proceedings of ICUMT, 2011.
[14] M. Ajmone Marsan, A. Fernandez Anta, V. Mancuso, B. Rengarajan, P. Reviriego, and G. Rizzo, “A simple analytical model for energy efficient Ethernet,” IEEE Commun. Lett., vol. 15, no. 7, 2011, pp. 773–775.
[15] A. Chatzipapas and V. Mancuso, “Modelling and Real-Trace-Based Evaluation of Static and Dynamic Coalescing for Energy Efficient Ethernet,” in Proceedings of ACM e-Energy, 2013.
[16] 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, 2010.

[17] A. Oliva, T. R. Vargas, J. C. Guerri, J. A. Hernandez and P. Reviriego, "Performance analysis of energy efficient Ethernet on video streaming servers," Comput. Networks, vol. 57, no. 3, pp. 599–608, 2013.
[18] D. Dove, "Energy Efficient Ethernet: Switching Perspective," presented at IEEE 802.3 meeting, January, 2008, available online at http://www.ieee802.org/3/az/public/may08/dove_02_05_08.pdf.
[19] W. Diab "LLDP's Use in EEE" presented at IEEE 802.3 meeting, January, 2009, available online at http://www.ieee802.org/3/az/public/jan09/diab_02_0109.pdf.
[20] P. Reviriego, V. Sivaraman, Z. Zhao, J. A. Maestro, A. Vishwanath, A. Sánchez-Macián and C. Russell, "An Energy Consumption Model for Energy Efficient Ethernet Switches," Proc. of the International Conference on High Performance Computing & Simulation (OPTIM Workshop), Madrid (Spain), pp. 98–104, 2012.
[21] NRDC, Small Network Equipment Energy Consumption in U.S. Homes: Connecting devices with less energy, NRDC Issue Paper, 2013.
[22] B. Urban, V. Tiefenbeck, and K. Roth. "Energy Consumption of Consumer Electronics in U.S. Homes in 2010". Final Report by Fraunhofer Center for Sustainable Energy Systems for the Consumer Electronics Association.
[23] J. G. Koomey, "Growth In Data Center Electricity Use 2005 To 2010", Analytics Press, 2011, available online at http://www.analyticspress.com/datacenters.html.
[24] C. Gunaratne, K. Christensen, B. Nordman and S. Suen, "Reducing the energy consumption in Ethernet with Adaptive Link Rate (ALR)," IEEE Trans. Comput., vol. 57, no. 4, pp. 448–461, 2008.
[25] "Next Generation BASE-T IEEE 802.3 Call For Interest," 2012, available online at http://www.ieee802.org/3/NGBASET/public/jul12/CFI_01_0712.pdf.
[26] P. Wu, G. Parnaby and W. Lo "NGBASE-T Requirements Learning from 10GBASE-T," presented at IEEE 802.3 meeting, January 2013, available online at: http://www.ieee802.org/3/NGBASET/public/jan13/WuParnaby_01a_0113_NGBT.pdf.
[27] A. Odlyzko, "Data networks are lightly utilized, and will stay that way," Rev. Netw. Econ., vol. 2, no. 3, pp. 210–37, 2003.
[28] M. Bennett, "Energy-Efficient Ethernet for 100G Backplane and Copper" presentation to IEEE 802.3, November, 2011, available online at: http://www.ieee802.org/3/bj/public/sep11/bennett_01a_0911.pdf.
[29] B. Nordman, "EEE Savings Estimates", prepared for IEEE EEE Study Group, May 2007, available online at: http://www.ieee802.org/3/eee_study/public/may07/nordman_2_0507.pdf.
[30] S. Lanzisera, B. Nordman and R. E. Brown, "Data network equipment energy use and savings potential in buildings", Energy Effic., vol. 5, no. 2, pp. 149–162, (2012)
[31] IEA, Gadgets and Gigawatts "Policies for Energy Efficient Electronics", International Energy Agency, May 2009, available online at http://www.iea.org/w/bookshop/add.aspx?id=361.
[32] "FC-EE - Energy Efficient Fibre Channel" specification draft, November, 2012, available online at: ftp://ftp.t10.org/t11/document.12/12-073v4.pdf.
[33] P. Reviriego, R. Pérez de Aranda and C. Pardo, "Introducing Energy Efficiency in the VDE 0885–763 Standard for High Speed Communication over Plastic Optical Fibers," IEEE Commun. Mag., vol. 51, no 8, pp. 97–102, 2013.
[34] SIEPON Working Group Web Site, http://www.ieee1904.org/1/.
[35] "ENERGY STAR Program Requirements for Small Network Equipment," November, 2012, available online at: http://energystar.gov/products/specs/sites/products/files/ES_SNE_Draft_2%20_V1_Specification_Nov2012.pdf.
[36] M. Gustlin, H. Barrass, M. Bennett, A. Healey, V. Pillai, M. Brown and W. Diab, "Detailed baseline for EEE in 100G" presentation to IEEE 802.3bj, March, 2012, available online at: http://www.ieee802.org/3/bj/public/mar12/barrass_01_0312.pdf.
[37] G. Nicholl and D. Ofelt "Energy Efficient Ethernet (EEE) for 40 Gb/s and 100 Gb/s optical interfaces" presentation to IEEE 802.3bm, November 2012, available online at: http://www.ieee802.org/3/bm/public/nov12/nicholl_01a_1112_optx.pdf.
[38] O. Haran "Applicability of EEE to fiber PHYs" presentation to IEEE 802.az, September 2007, available online at: http://www.ieee802.org/3/eee_study/public/sep07/haran_1_0907.pdf.
[39] M. Bennett "EEE Considerations for RTPGE Reduced Twisted-Pair Gigabit Ethernet Study Group", July, 2012, available online: http://grouper.ieee.org/groups/802/3/RTPGE/public/july12/bennett_01_0712.pdf.
[40] M. Bennett and P. Reviriego "Reach and Energy Efficiency in NGBASE-T" presentation to the NGBASE-T study group, November, 2012, available online at: http://www.ieee802.org/3/NGBASET/public/nov12/bennett_01_1112_ngbt.pdf.
[41] D-Link "Power Saving by Cable Length," available online at: http://www.dlinkgreen.com/energyefficiency.asp.