

An Experimental Power Profile of Energy Efficient Ethernet Switches

V. Sivaraman^{a,*}, P. Reviriego^b, Z. Zhao^a, A. Sánchez-Macián^b, A.
Vishwanath^c, J. A. Maestro^b, C. Russell^d

^a*University of New South Wales, Sydney, NSW 2052, Australia*

^b*Universidad Antonio de Nebrija. C/Pirineos 55, 28040 Madrid, Spain*

^c*CEET, University of Melbourne, VIC 3010, Australia*

^d*CSIRO ICT Centre, Sydney, PO Box 76 Epping, NSW 1710, Australia*

Abstract

The access network is believed to account for 70-80% of the overall energy consumption of wired networks, attributable in part to the large number of small and inefficient switches deployed in typical homes and enterprises. In order to reduce the per-bit energy consumption of such devices, the Energy Efficient Ethernet (EEE) standard was approved as IEEE 802.3az in 2010 with the aim of making Ethernet devices more energy efficient. However, the potential for energy savings, and their dependence on traffic characteristics, is poorly understood. This paper undertakes a comprehensive study of the energy efficiency of EEE, and makes three new contributions: First, we perform extensive measurements on three commercial EEE switches, and show how their power consumption profile depends on factors such as port counts, traffic loads, packet sizes, and traffic burstiness. Second, we develop a simple yet powerful model that gives analytical estimates of the power consumption of EEE switches under various traffic conditions. Third, we validate the energy savings via experiments in typical deployment scenarios, and estimate the overall reduction in annual energy costs that can be realized with widespread adoption of EEE in the Internet.

Keywords: IEEE 802.3, Energy Efficient Ethernet, power management, energy-aware systems, local area networks

*Corresponding author. Tel.: +61 2 9385 6577. Fax: +61 2 9385 5993
Email address: vijay@unsw.edu.au (V. Sivaraman)

1. Introduction

The ICT sector is responsible for over 3% of the world's electricity consumption, and has global CO₂ emissions on par with the aviation industry (Gartner Estimate, 2008). Further, the energy requirements of computing and networking equipment are predicted to grow substantially in the coming decade [1]. Energy expenses are already a major operational cost in today's data centers [2], while the aggregated energy consumption of consumer electronic devices is growing at an alarming rate [3]. As far back as the year 2000, the energy consumption of Internet equipment was estimated to be over 6 Twh [4], and more recent studies suggest a larger consumption when end user equipments and access networks are considered [5]. There is a specific need to address energy efficiency issues in the access network, since they are believed to account for 70-80% of the wireline energy costs [6]. This is attributable in part to the large number of end-premises devices that are in deployment, and to their high energy consumption per-bit, since they tend to be always on and have a very low average load.

Most network equipment today has a nearly constant power consumption independent of the system load [7]. Recognition of the resulting poor energy efficiency in lightly loaded systems has led to interest in making the energy consumption more proportional to system load. Among the first networking technologies to do so is Energy Efficient Ethernet (EEE). The IEEE 802.3az standard [8], approved in September 2010, defines a low power idle (LPI) mode that improves the energy efficiency of Ethernet physical layer devices under low loads. With an installed base of over one billion devices, the expected energy savings can be as high as 4 Twh [9]. Products that implement the EEE standard are becoming more common in the market and wide adoption is expected to occur in a few years. However, current products that implement EEE, e.g. [10, 11, 12, 13], do not provide much information on the energy consumption profile, typically reporting only minimum and maximum power values. There is clearly a need to gain a better understanding of EEE power consumption and its dependence on various system and traffic parameters.

This paper is the first to conduct an in-depth experimental study of the power consumption profile of EEE-based small Ethernet switches, and to evaluate its dependence on various traffic characteristics. Our specific contributions are three-fold:

1. We use a high-precision traffic generator and high-fidelity power me-

ter to measure the power consumption of commercial EEE switches from three vendors, and accurately characterize their power profile for varying traffic characteristics such as load, packet size and burst length.

2. We develop a simple yet useful model that yields analytical estimates of EEE switch power consumption as a function of traffic characteristics. The model is derived from first principles using the operating principles and parameters specified in the standards, and is shown to match well with our experimental measurements.
3. We evaluate the energy savings of EEE switches in typical deployment scenarios involving file transfers, and use our results to estimate the overall achievable reduction in annual energy costs as EEE adoption becomes more widespread in residential and small-enterprise networks.

To the best of our knowledge our work is the first to profile the power consumption of EEE switches under a wide range of operating conditions, and to develop a simple model that yields analytical estimates that corroborate well with real measurements. Our work can be used by other researchers wanting to explore higher layer energy saving mechanisms such as energy efficient routing [14] or selective link deactivation [15], and can inform current and future efforts within the IEEE 802.3 standards to extend the energy efficiency mechanisms to other Ethernet speeds and technologies.

The rest of the paper is organized as follows: section 2 gives requisite background on Energy Efficient Ethernet (EEE) and prior efforts to profile switch power consumption. In section 3 we experimentally profile the power consumption of three small Energy Efficient Ethernet switches, and evaluate the dependence on factors such as number of active ports, traffic load, packet sizes, and traffic burstiness. Using the insights from these experiments, in section 4 we develop a simple EEE power consumption model, validate it against the experimental results, and discuss its applicability to practical deployments including TCP traffic. Section 5 discusses the savings that can be achieved with the introduction of Energy Efficient Ethernet, comparing it with previous estimates, and provides an overview of current energy efficiency efforts within IEEE 802.3. The paper concludes with a summary and directions to future work in section 6.

2. Background

In this section we give a brief background on the Energy Efficient Ethernet (EEE) technology, and on prior efforts by researchers to develop energy

consumption profiles of network switches.

2.1. Overview of EEE




Figure 1: Mode transitions in Energy Efficient Ethernet

For over two decades, Ethernet has been the dominant technology for wireline LANs. It is widely used in residences and commercial buildings and almost all computers include one (or more) Ethernet connections. Although Ethernet supports a variety of transmission media, most of the Ethernet ports are connected by Unshielded Twisted Pairs (UTP), especially in homes and offices. For UTP, Ethernet currently supports four data rates: 10 Mb/s (10BASE-T), 100 Mb/s (100BASE-TX), 1 Gb/s (1000BASE-T) and 10 Gb/s (10GBASE-T). For data rates of 100 Mb/s and higher, Ethernet physical layer transmitters transmit continuously to keep transmitters and receivers aligned. When there is no data to send an auxiliary signal called IDLE is sent. This means that most of the elements in the interfaces are active at all times leading to an energy consumption that is high and largely independent of the traffic load.

To reduce energy consumption, the IEEE 802.3az standard [8] introduces 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 (T_q) over which no signal is transmitted and small periods (T_r) during which a signal is transmitted to refresh the receiver state to align it with current conditions. The operation of the LPI mode is illustrated in Fig. 1. The energy consumption of a physical layer device (PHY) when it is in LPI mode is expected to be significantly lower than when it is in the active mode. Our previous work reported in [16] has confirmed that computer Network Interface Cards (NICs) operating at 1 Gbps can achieve 70% reduction in power consumption when using EEE.

The actual energy savings on a given link depend on the amount of time that the link spends in LPI mode. This time can be reduced by the transition overheads associated with activating (T_w) and putting it into LPI mode (T_s).

Protocol	Min T_w (μsec)	Min T_s (μsec)	Frame size (bytes)	T_{Frame} (μsec)	Single Frame efficiency	Frame size (bytes)	T_{Frame} (μsec)	Single Frame efficiency
100BASE-TX	30.5	200	1518	120	34.2%	64	5.1	2.2%
1000BASE-T	16.5	182	1518	12	5.7%	64	0.5	0.3%
10GBASE-T	4.48	2.88	1518	1.2	14.0%	64	0.05	0.7%

Table 1: Minimum wake, sleep and frame transmission times and single frame efficiencies for different link speeds

During those transitions, there is significant energy consumption and the transition times are large compared with the frame transmission time [16, 17]. The transition times for the different speeds are summarized in Table 1 and compared with the frame transmission times for a 1518 byte and a 64 byte packet. To measure the efficiency of EEE, the concept of Single Frame Efficiency which measures the efficiency of EEE for single frame transmission is introduced. When a single frame is transmitted the link has to be activated to send a frame and then deactivated after the transmission. Therefore for a frame transmission time of T_f , the link is active or in transitions for $T_w + T_s + T_f$. The ratio of these times is defined as the Single Frame Efficiency: $\text{SFE} = \frac{T_f}{T_w + T_s + T_f}$. A high value of SFE is desirable as it indicates that the overheads of entering and exiting LPI are low. However, Table 1 shows that SFE values are relatively low, implying that the overheads of entering LPI mode become more significant as the time spent in LPI mode reduces. Stated another way, this suggests that the benefits of EEE saturate at medium or low loads, unless packets are coalesced as showed in [17]. Our example in [16] for a EEE NIC demonstrated that for a 1 Gb/s link, a 6% traffic load composed of evenly spaced packets prevented the link from entering into LPI mode altogether. Our study in this paper will therefore consider the effect of traffic burstiness on EEE switch power consumption.

2.2. Prior Studies on Switch Energy Consumption

There is a wide range of Ethernet switches, from small four-port switch used in homes and small offices to modular switches that support hundreds of ports and different transmission media [18]. Large switches have more ports and higher port speeds, and hence consume more power. However, there are many more small switches than large ones, and the aggregated energy consumption of small switches to which users directly connect is significantly

higher than for large switches. For example, the power consumption of small Ethernet switches in the US has been recently estimated in 7.9 TWh/year [19]. Small switches typically have 5,8,16 or 24 ports, allowing highly integrated implementations in which only one [20] or a few integrated circuits are used [21]. The switch is composed of one physical layer device (PHY) and one Medium Access Controller (MAC) per port, a switching fabric commonly implemented with a shared memory, and a control logic implemented in a CPU [22].

The power consumption of switches and routers has been characterized in different studies such as [23, 24] and our own study in [25]. The results show that once the switch is powered on and its ports are activated, the power consumption is close to its maximum value. For example in [24] that value is around 90% of the peak power consumption for a core router. That means that only 10% of the peak power consumption is dependent on the traffic load. This is far from the proportional relation and results in poor energy efficiency as networks tend to be lightly loaded [26]. Similar results are obtained for commercial Ethernet switches in [23], and in our study of an FPGA-based prototype router in [25]. We will corroborate this result for small Ethernet switches in our experimental study by considering an 8-port Cisco Catalyst switch that does not have EEE support. To the best of our knowledge there is no study of EEE switch power consumption (other than our own preliminary study [27]) reported in the research literature.

3. Experimental Characterization of EEE Switch Power Consumption

To characterize the actual power savings obtained from Energy Efficient Ethernet technology, we selected three small commercial Ethernet switches that implement EEE. The models selected were the D-Link DGS-1100-16 [10] which has 16 Gigabit ports, the Level-One GEU-0820 [12] which has 8 Gigabit ports, and the SMC GS801 [28] that also has 8 Gigabit ports. These switches were selected because they are amongst the first in the market to have EEE compatibility, are relatively low-cost and readily available, and thereby representative of switches that are typical in homes and small offices. For comparison purposes we also used an 8-port Cisco Catalyst 3560 switch that does not have EEE capability.

For traffic generation we used a high-precision hardware-based traffic generator supplied by IXIA [29]. Our IXIA traffic generator comes equipped

with eight Gigabit Ethernet ports, and we could configure with nanosecond-level precision the traffic stream characteristics such as packet spacing, burst spacing, etc. as required by the experiments described below. The power consumed by the switch under test was measured by the PowerMate power monitor [30], which performs high-rate sampling of the voltage and current and uses a 24-bit precision digital filter to yield of resolution of 0.01 W in power measurement. In what follows we describe each successive experiment that evaluates the impact of an aspect of the traffic stream.

3.1. Experiment 1: Impact of Number of Active ports

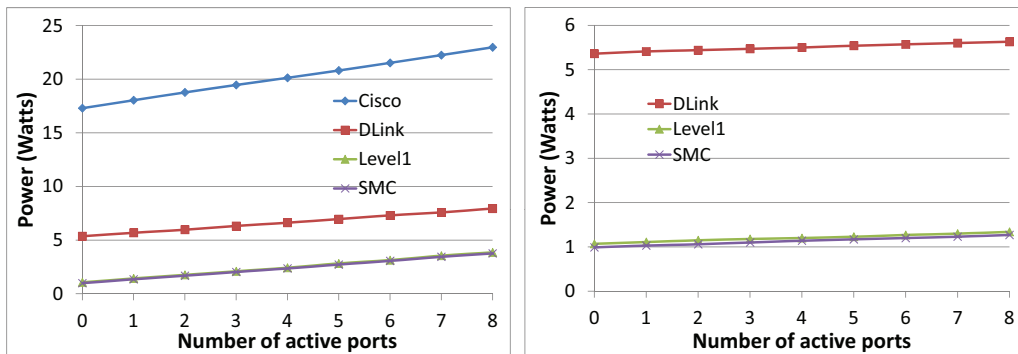


Figure 2: Energy consumption of switches (a) without EEE, and (b) with EEE, as a function of number of active ports.

The objective of our first experiment is to evaluate the impact of the number of active (i.e. connected) ports on overall switch power consumption, when no traffic is flowing. In Fig. 2 we show the power consumed as the number of active ports increases from 0 to 8 for all four switches considered, with and without EEE. Though the plots show that switch power increases linearly with number of active ports with and without EEE, the slope for non-EEE switches (left plot) is much higher than for EEE ones (right plot). Without EEE, the Cisco switch consumes about 0.71 W per port and each of the other three switches consume 0.34 W per port. This is a direct result of technology scaling (each of the three switches is much newer than the Cisco Catalyst), showing that electronic technology is becoming more power efficient with time. This power consumption reduction for 1000BASE-T PHYs has also been noticed in NICs by our recent study [16] when compared to previous studies [31].

When EEE is enabled (right plot), the per-port power consumption falls dramatically for all three EEE switches, being approximately 0.03 W per-port. This shows that at zero load, EEE reduces the per-port power consumption to about 10% of its original value. This savings of 90% under zero load conditions implies that EEE can save significant energy in residential and small-office switches that often have connected but idle hosts. Further, it also suggests that adoption of EEE may reduce the need for current research efforts that try to reduce the number of active links when there is no traffic [15] or try to allocate traffic such that the number of links that are activated is minimized [14].

The third observation to emerge from Fig. 2 is that though the per-port power is roughly the same in the EEE switches tested, the base power consumption can be quite different. The DLink switch consumes 5.36 W even when no ports are connected, whereas the Level1 and SMC switches consume only 1.07 and 0.99 W respectively (the capacities of these switches being 16 Gbps, 8 Gbps and 8 Gbps). This may be explained in part because the DLink is a larger switch with 16-ports, whereas the other two have only 8 ports each. However, the difference is large enough to suggest that the power consumption depends heavily on the switch implementation.

3.2. Experiment 2: Impact of Traffic Load

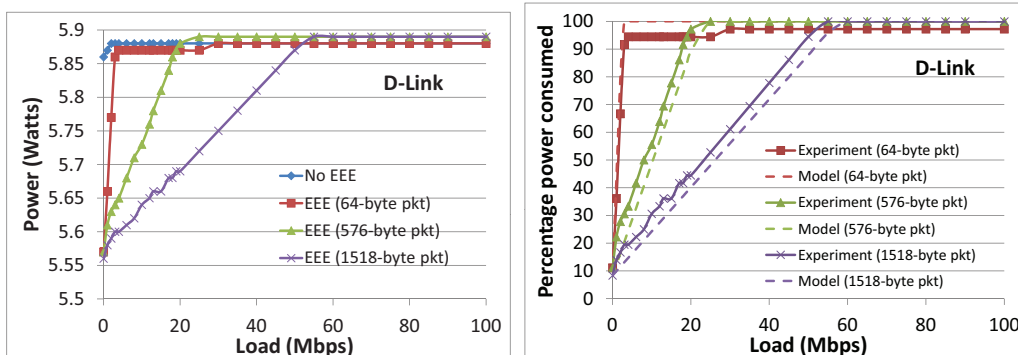


Figure 3: Total switch power (left) and per-port percentage power (right) as a function of traffic load for D-Link switch: Unidirectional traffic.

Our second experiment introduces traffic on the active ports, and evaluates the impact of traffic load on switch and per-port power consumption. Our traffic was generated by the high-precision IXIA traffic generator,

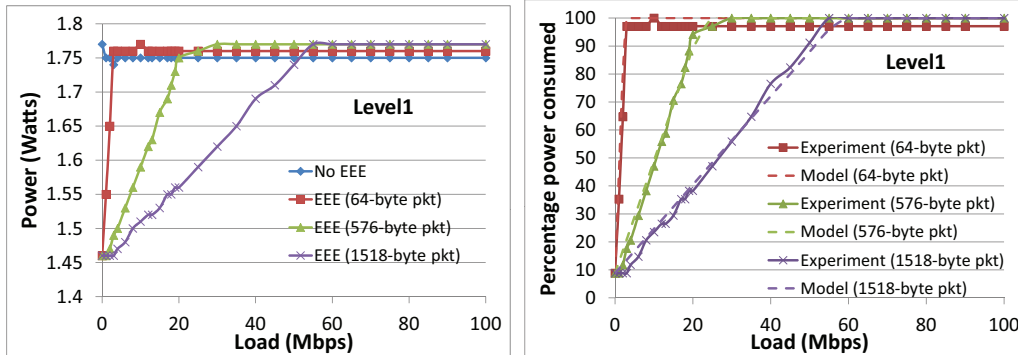


Figure 4: Total switch power (left) and per-port percentage power (right) as a function of traffic load for Level1 switch: Unidirectional traffic.

wherein we can specify the precise data rates and frame sizes. We conducted experiments with both unidirectional traffic (whereby traffic enters one port of the switch under study and exits on another port), and bidirectional traffic (in which traffic flows in both directions between the pair of ports).

3.2.1. Unidirectional Traffic

In Fig. 3 we show the power consumption for the DLink switch as the traffic load increases from 0 to 1 Gbps (the plots truncate at 100 Mbps since the power consumption of EEE saturates well before that), for fixed-size packets of various sizes. The plot on the left shows the total power consumed by the switch. The top curve (light blue) demonstrates that without EEE, the switch consumes a fairly constant power invariant to load. With EEE (remaining curves with packet sizes 64, 576, and 1518 bytes respectively), the plot shows that the power consumption increases linearly with load, till it saturates. The saturation point is at around 55 Mb/s for 1518 packets and well below 10 Mb/s for 64 byte packets. This behaviour is expected, since EEE is designed to save energy during low loads by transitioning into low power idle (LPI) mode. As traffic load increases, the opportunity to enter LPI mode diminishes, and the energy savings correspondingly reduce.

The plot on the right in Fig. 3 shows the power consumed by the port as a percentage of the maximum per-port power. It shows that with no traffic, the port consumes about 10% power (as noted in the previous experiment), and this increases linearly with traffic load till it reaches 100% at a threshold that is dependent on packet size (discussed in the next subsection). The plot also shows the predictions from our analytical model that will be described

in the next section.

The corresponding plots for the Level1 switch are shown in Fig. 4. The absolute values in the left plot are much lower than for the DLink switch observed in Fig. 3. However, the right plot shows that the percentage values on a per-port basis are very similar to the DLink, again exhibiting that the power consumed increases linearly with traffic load till it saturates. We omit the plot for the SMC switch as it exhibited very similar results to the Level1 switch.

3.2.2. Bidirectional Traffic

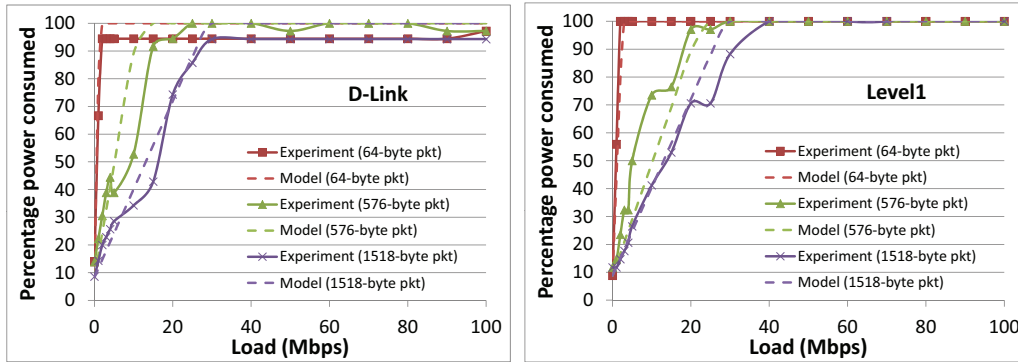


Figure 5: Per-port percentage power as a function of traffic load for (left) D-Link switch and (right) Level1 switch: Bidirectional traffic.

In Fig. 5 we plot the power consumed by the port as a percentage of the maximum per-port power for the Dlink (left) and Level1 (right) switches in the presence of bidirectional traffic. Three important observations emerge from the figure: First, the shapes of the curves are consistent with the previous experiment in that with no traffic the ports consume about 10% power, and this increases with traffic load saturating at a threshold that is dependent on the packet size. Second, the power consumption rises more rapidly with load than for unidirectional traffic. Indeed, the plots show that the load threshold at which the power saturates is nearly half of that with unidirectional traffic. For example, with 1518 byte packets the Dlink switch saturates at 30 Mbps, in contrast to 60 Mbps as seen in the right plot of Fig. 3. This is because packets in the two directions arrive to the port at arbitrary (unsynchronized) points in time, effectively halving the inter-arrival times of packets arriving to the port for a given load. Consequently the amount of

time a port spends in the LPI mode also reduces because EEE standards require that a 1 Gbps port enter LPI mode only when both link directions are idle. Lastly, the predictions from our analytical model, described in the next section, matches well with the experimental results, corroborating its validity for both uni- and bi-directional traffic.

3.3. Experiment 3: Impact of Packet Size

Figures 3 and 4 have already shown that the power consumption for a given data rate depends on packet size. The per-port power in both DLink and Level1 switches saturates at around 3 Mbps for 64-byte packets, 25 Mbps for 576-byte packets, and 55 Mbps for 1518-byte packets. The difference in behavior for different packet sizes can be understood as follows: since packets are being sent with even spacing in this experiment, for a given bit-rate, smaller packets will arrive at a faster rate than bigger packets. Consequently, the EEE link will transition in and out of LPI mode more frequently for small packets, and the associated overheads will each transition (activation time T_w and de-activation time T_s) will impose a larger penalty when compared to the frame transmission time T_f , as evident from Table 1. Larger packets are therefore more desirable for energy savings.

3.4. Experiment 4: Impact of Traffic Burstiness

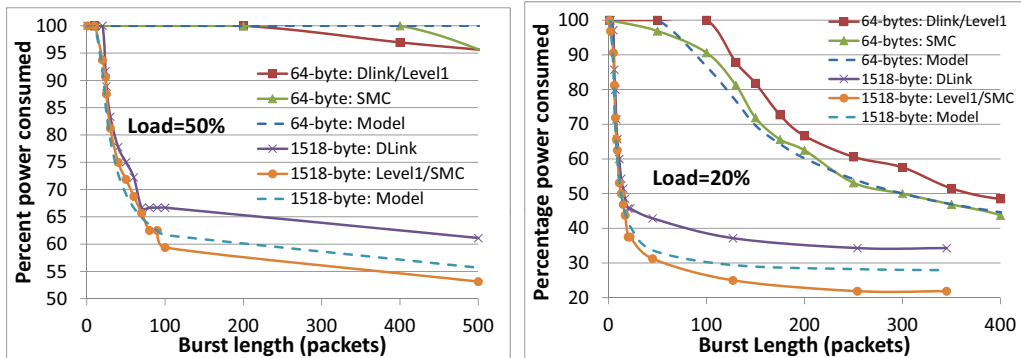


Figure 6: Per-port power consumption versus traffic burst size for (a) 50% load and (b) 20% load.

Having established the importance of packet size in energy savings, our next experiment evaluates the impact of traffic burstiness on EEE performance. We chose a constant traffic load, and in each run increased the burst

length (number of back-to-back packets) entering the switch. In other words, our IXIA traffic generator sends an on-off stream in which a burst of B packets is sent during the on period, after which it goes off for a corresponding period so that the average bit-rate is constant. The burst length B was increased from 1 to 1000 packets, with the gap length adjusted depending on chosen traffic load and packet size.

In Fig. 6 we show the per-port power consumed (as a percentage value) for each of the three EEE switches as the burstiness increases, for an average 50% link load (left plot) and 20% link load (right plot). The plots show that for a given load, the power consumption falls as the burst size increases. This is explained by the fact that a larger burstiness is associated with fewer but larger gaps, reducing the number of transition overheads to/from LPI mode. Asymptotically, the percentage per-port power converges to the percentage load on the port, since the link is expected to be on for that fraction of time and sleeping for the rest.

The plots also show that larger packets benefit more from burstiness: for 1518-byte packets, evenly spaced transmission at 50% load (left plot) consume 100% power, whereas coalescing packets into bursts of 1000 packets reduces power consumption to about 50%. For 1518-byte packets at a load of 20% (right plot), we again see that coalescing packets into bursts of a few hundred packets reduces the power to about 20%. For small packets (64-bytes), the plots show that a much higher burstiness is required in order to realize these asymptotic power savings.

These experiments illustrate that traffic burstiness is another important determinant that influences the power savings achievable with EEE. Our model in the next section will consider these various factors in estimating the energy performance of EEE.

4. An Energy Consumption Model for Small EEE Switches

The previous section has given us experimental insight into the impact of various parameters such as traffic load, packet sizes, and traffic burstiness on the power consumption of EEE switches. In this section, we combine these practical insights with the principles of EEE operation (as detailed in §2.1) to develop an analytical model for estimating the power consumption of EEE switches under various operating scenarios. This model is validated against measured results of the previous section, and its generality to practical deployment scenarios that include TCP traffic are discussed.

Our model is different from previously proposed analytical models of EEE [32, 33, 34] in several aspects. Firstly the focus is to model the entire switch compared to an individual PHY as in previous models. Secondly the model is based on energy measurements and not on simulations. Thirdly, the model targets access equipment for which the bottleneck assumption is used, unlike previous models that target links with highly aggregated traffic and different statistical distributions of arrivals.

4.1. The Model

At the highest level, the power P_{switch} consumed by a EEE switch is the sum of a base component P_{base} that corresponds to the power consumed when the switch is on but none of its ports are active (connected), and a per-port component P_{port}^i when port- i is active:

$$P_{\text{switch}} = P_{\text{base}} + \sum_i P_{\text{port}}^i \quad (1)$$

Our experiments in the previous section show that $P_{\text{base}} = 5.36, 1.07, 0.99$ W for the DLink, Level1, and SMC switch respectively. Further, the observed linear dependence of per-port power on traffic load leads us to propose that the per-port power consumption P_{port} (for clarity we drop the port index i when we are focusing on a specific port) can be expressed as:

$$P_{\text{port}} = P_{\text{static}} + P_{\text{dynamic}}f(\rho) \quad (2)$$

Here P_{static} is the power consumed by the port when it is connected (active) but has no traffic on it, P_{dynamic} is the power consumed by the port at full load, less the static power, and $f(\rho)$ is a function of the utilization ρ satisfying $f(0) = 0$ and $f(1) = 1$. For all three models of EEE switches evaluated, we found that $P_{\text{static}} \approx 0.034$ W, which is about 10% of the maximum power consumed per-port (as evident in figures 3 and 4), while P_{dynamic} constitutes the remaining 90%. We attribute P_{static} to the fact that even when there is no traffic and the EEE link is in low-power-idle (LPI) mode, it needs to periodically awaken to align state with the other end.

The function $f(\rho)$ can be thought to represent the fraction of time that the port is in active state, and our experiments show that this depends not just on port utilization ρ (where $0 \leq \rho < 1$), but also on factors such as packet sizes and traffic burstiness. We therefore start with the simplifying

assumption that the load at the port arises from packets of fixed length L bytes that are forwarded in periodic bursts of B packets each. Referring to Fig. 1, the link first takes time T_w to awaken from LPI mode. It then transmits for time BT_f , wherein B frames, each taking time $T_f = L/C$, where L is the frame size and C the link rate, are sent. Thereafter, the link re-enters LPI sleep mode, taking time T_s . This process repeats every BL/R units of time, where R is the average rate of traffic through the port. The fraction of time that the port expends energy is:

$$f(\rho) = \min\left\{1, \frac{T_w + BT_f + T_s}{BL/R}\right\} = \min\left\{1, \frac{T_w + BT_f + T_s}{BT_f}\rho\right\} \quad (3)$$

where we have used the fact that $L/R = \frac{L/C}{R/C} = T_f/\rho$, and that $f(\rho)$ is upper bounded by 1. Combining this with (2) gives us:

$$P_{\text{port}} = P_{\text{static}} + P_{\text{dynamic}} \min\{1, D\rho\} \quad (4)$$

where $D = \frac{T_w + BL/C + T_s}{BL/C}$ is a factor that depends on packet length and burst size. We will now validate this model against the measured observations reported in the previous section, and then discuss its applicability to typical deployment scenarios.

4.2. Validating the Model

For fixed packet lengths and burst sizes, it is easy to validate the model above against the measurements shown in Section 3. We begin by considering experiment 2 that varied the traffic load and used evenly spaced packets. In this case the burst length $B = 1$, and the factor D becomes $\frac{T_w + L/C + T_s}{L/C}$, which is precisely the inverse of the single frame efficiency (SFE) defined in Section 2.1. For the 1 Gbps link used in our experiments, using the values of T_w , T_s , and T_f shown in Table 1, we get $D = 17.6$ for a 1518 byte packet, $D = 44.08$ for a 576-byte packet, and $D = 388.7$ for a 64 byte packet. We plug in these values for D in (4) to obtain the corresponding predictions from the analytical model, and the per-port percentage power are shown as dashed curves in the right-hand side plots in Fig. 3 and Fig. 4. The match of the model predictions to measured values is very good, showing that our model is able to capture the effect of traffic load and packet size appropriately.

We now consider experiment 4 in which the burstiness of traffic was varied. For each value of burst size B , we can compute the expression in (4)

for given packet length L and load ρ . The computed values are shown as dashed lines in Fig. 6, and again confirm that the analytical estimates are in close agreement with measured values, validating that the model is able to capture the effect of bursty traffic on EEE power consumption.

4.3. Applicability of the Model

The above model is well suited when packet length L and burst size B are constant and known. However, these numbers are deployment specific, and difficult to deduce *a priori*. However, a observation that can be made about typical deployments in home and small office networks is that the traffic rate through the switch is typically limited by the Internet-facing port, such as the DSL uplink from a home gateway or a low-speed WAN link from a small office to the ISP. Even if the switch LAN ports operate at 1 Gbps, the lower uplink speeds (typically 10-50 Mbps) force the packets on the LAN ports to be spaced apart. For example, a 1518-byte packet takes $200\mu\text{s}$ to receive on a 50 Mbps link but only $12\mu\text{s}$ to transmit on a 1 Gbps link, thereby introducing spacing between successive packets during which EEE goes into LPI mode. For such scenarios, it is therefore reasonable to assume that the packets are sent individually, and the burst length is hence $B = 1$.




Figure 7: CDF of packet length distributions measured at the NASA Ames Internet Exchange (AIX)

Under the low-speed Internet link assumption above, the per-port EEE power consumption still depends on the frame size L . The model can be generalized to remove this dependence by using the distribution of the frame

size, if known. Let $p(L)$ denote the probability that a randomly chosen bit belongs to a packet of length L bytes. Then factor D in (4) can then be calculated as:

$$D = \sum_{L=L_{\min}}^{L_{\max}} p(L) \frac{T_w + L/C + T_s}{L/C} \quad (5)$$

To illustrate this quantitatively with an example, let us assume that the packet sizes follow the distribution reported by CAIDA [35] at the NASA Ames Internet Exchange (AIX) measured over 87 million packets captured in 2000, with cumulative distribution function shown in Fig. 7. The distribution is trimodal, with the vast majority of IP packets being either minimum size (64 bytes) or maximum size (1518 bytes) or 576 bytes (corresponding to TCP implementations that don't use path MTU discovery) – indeed these are precisely the packet sizes we chose for our experimental study. Given this distribution of packet size probabilities $d(L)$, the probability that a random bit belongs to packet of size L is given by $p(L) = Ld(L)/\sum_L Ld(L)$. Using the values of $p(L)$ thus computed for any packet size L , and using the values of T_w and T_s from Table 1 with $C = 1$ Gbps, we can readily perform the summation over all packet sizes in (5), and obtain $D \approx 50.34$. We can see that the value of D lies in-between the values reported earlier for small and large frame sizes of fixed-size, namely $D = 388.7$ for 64-byte packets, and $D = 17.6$ for 1518-byte packets), and might be a reasonable value to use in networks that have variable packet sizes as observed in the CAIDA trace. If the packet size distribution for the specific network being studied is known, the factor D can easily be calculated for that network.

We have argued that for home and small office networks where the traffic is non-bursty and the packet size distribution is known, our model is able to deduce the per-port energy consumption of EEE switches based only of the carried load. We acknowledge that in more general settings where the traffic characteristics are unclear the usage of our model may be more difficult. To illustrate this we conducted experiments with TCP traffic in which the bandwidth on the bottleneck link (upstream to the switch under test) was varied. Fig. 8 shows the per-port power consumption for the DLink (left) and Level1 (right) switches as a function of the bottleneck link rate, when 1, 10, and 50 concurrent TCP flows (generated using iPerf with default packet size of 1518 bytes) share the link. It is seen that when there is only one TCP flow on the link, the power consumption rises slowly, since the traffic generated by a single flow tends to be bursty, and we have seen in Section

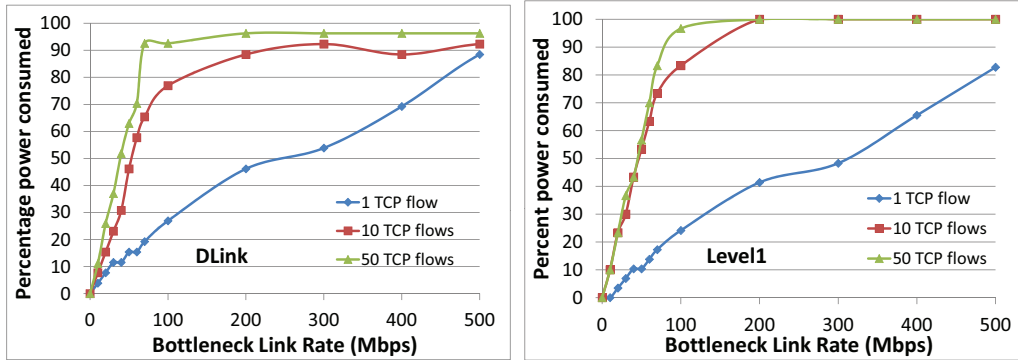


Figure 8: Per-port percentage power with TCP flows for DLink (left) and Level1 (right) switch.

3.4 that larger burstiness translates to lower EEE power consumption. As more TCP flows multiplex, the burstiness of the aggregate traffic reduces, as do the EEE energy savings, evidenced by the higher energy consumption shown in the figure for 10 and 50 flows. Capturing the impact of TCP traffic characteristics on EEE energy consumption is beyond the scope of this paper and left for future work.

5. Discussion on Energy Saving Realizable with EEE

5.1. Estimating Energy Savings with Wide Adoption of EEE

Estimating the energy savings that will be provided by Energy Efficient Ethernet when it is widely adopted is interesting as it provides an insight into the benefits obtained by developing the standard. In [9], initial estimations were presented based on several assumptions and simulation results. One of the assumptions was that the low power mode will provide approximately 90% energy savings for the physical layer bringing 1Watt saving per 1 Gb/s link when there is no traffic. For the switches measured, the savings are close to 0.3 Watts. To estimate the savings in the computer, the power savings in a HP Z210 Workstation with Intel 82579LM Gigabit Network Adapter were measured. The results showed savings of approximately 1 Watt bringing the total link savings to 1.3 Watts compared to the 1Watt estimated in [9]. The average link load was assumed to be approximately 1% for which simulation results showed savings of approximately 81%. This is consistent with the measured results and the proposed model if 1518 byte packets are used.

	Links (1Gbps)
Assumptions	
Link utilization	1%
Active Links (millions)	250
Electricity cost(\$/kWh)	0.1029
Measurements	
Savings per link no data (W)	1.3
Savings per link 1% load (W)	1.135
Results Savings EEE	
Total (TWh/year)	2.5
Total (million US\$/year)	257

Table 2: Estimated EEE savings at 1Gbps links with 2008 stock figures

To estimate the total savings provided by Energy Efficient Ethernet, the number of links and number of average number of hours that the link is active assumed in [9] are used. The number of links corresponds to 2008 estimates for the U.S. only, the global savings would therefore be significantly larger. Table 2 incorporates the measured values to recalculate the total energy savings which are over 2 TWh/year for the US only. To compute the economic savings, electricity prices have been taken from the U.S. Energy Information Administration site with the average data for 2011 (taking the minimum from 11.65 cents for residential and 10.29 cents for commercial use). This again is similar to the value used in [9]. The total savings are over 250 million dollars per year, showing that the economic impact of Energy Efficient Ethernet can be significant.

5.2. Further Improvements to EEE Switches

The Energy Efficient Ethernet standard addresses the energy consumption of the PHY devices but can also enable savings in other system components [36]. These possible additional savings would be achieved by putting those components in a low power mode when the ports are in LPI mode as in that case no packets can arrive until the PHYs are activated. It seems that all the additional savings are not achieved in the first generation of EEE switches as the power consumption when there is no traffic on any port although lower than in the legacy switch remains significant (close to 25% for Level One switch and around 50% for D-Link switch). One possible explanation is that vendors have focused on implementing EEE on this first generation and left the optimization of the rest of the switch elements for future releases. If that is the case one would expect further improvements in

the future that will make the energy consumption of switches more proportional to traffic load. In any case, the proposed model would still be valid using different parameters (mostly reducing P_{base}). On the contrary, if there are actual limitations that make unfeasible a reduction of the power consumption when there is no traffic, then techniques that put the entire switch on a sleep mode such as the one proposed in [19] will provide significant benefits. On the other hand, techniques that only deactivate some of its links will provide insignificant energy savings in EEE switches.

5.3. Related Energy Efficiency Efforts in IEEE 802.3

The success of Energy Efficient Ethernet and the increased importance of environmental issues have resulted in additional energy efficiency efforts within IEEE 802.3. The most advanced standard is IEEE 802.3bj that will define 40/100G links targeting data center connectivity. This standard has adopted a scheme based on the use of a low power mode as in EEE to improve its energy efficiency. The challenge in this case will be the mode transition times. At 100Gb/s a 1518 byte packet takes only $0.12\mu\text{s}$ to transmit and therefore unless transition times in the order of nanoseconds are achieved, the transition overhead will be large. A more sophisticated approach based on the use of two different modes was proposed during standardization but was not finally adopted [37]. That solution would have probably provided better energy savings at the expense of a more complex implementation.

The definition of energy efficiency mechanisms is being considered in two other IEEE 802.3 standards that are currently being developed: Reduced Twisted Pair Gigabit Ethernet (RTPGE) and Next Generation BASE-T. In the first case, the standard targets automotive applications for which a very low power mode is needed when the device operates using the battery of the car. This presents a new challenge and possibly several low power modes will be needed. For the Next Generation BASE-T, the use of EEE seems the most straightforward option as it is already implemented in the rest of BASE-T devices. However as pointed out before, the increased speed means reduced frame transmission times and therefore tighter constraints on the mode transition times needed to achieve good savings.

The standards discussed show how energy efficiency has become one of the design goals along with performance and cost in new communication

systems. It is interesting to note that this shift has occurred in a few years as a decade ago almost no standard considered energy savings.

6. Conclusion

In this paper the first detailed experimental evaluation of power consumption of EEE Ethernet switches has been reported. The experiments show how the power consumed by EEE switches varies with the number of active ports, the traffic load, packet sizes, and packet burstiness. Based on the experimental results, a simple model for the power consumption of small Ethernet switches has been proposed. The model provides accurate estimates that match measurements for the scenarios considered, and can be extended to more complex scenarios with variable traffic patterns. As EEE is adopted over the next years, we believe that the model will be useful to estimate power savings in a simple way. Additionally the model can be used for research into new power saving techniques for Energy Efficient Ethernet LANs. For example, in the light of the model, the use of dynamic link shutdown seems to have much less potential than in legacy Ethernet.

The measurements have also been used to revise previous estimates of the benefits of adopting EEE. The results confirm the large potential savings that in fact exceed the estimates made before. Finally a discussion on current energy efficiency efforts on Ethernet standards that are currently being developed has been presented. This can guide the reader towards new research areas and shows the growing importance of energy efficiency in communications.

Acknowledgement

This work was partly funded by a Google Research Award (P. Reviriego).

References

- [1] International Energy Agency, *Gadgets and Gigawatts : Policies for Energy Efficient Electronics*, 2009.
- [2] U. Hölzle, L. A. Barroso, *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines*, Morgan and Claypool Publishers, 1st edition, 2009.

- [3] The Climate Group, Smart 2020 Report, Technical Report, Global e-Sustainability Initiative, 2008.
- [4] M. Gupta, S. Singh, Greening of the internet, in: Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications, SIGCOMM '03, ACM, New York, NY, USA, 2003, pp. 19–26.
- [5] R. Bolla, R. Bruschi, F. Davoli, F. Cucchietti, Energy efficiency in the future internet: A survey of existing approaches and trends in energy-aware fixed network infrastructures, *IEEE Communications Surveys and Tutorials* 13 (2011) 223–244.
- [6] R. Bolla, R. Bruschi, K. Christensen, F. Cucchietti, F. Davoli, S. Singh, The Potential Impact of Green Technologies in Next-Generation Wireline Networks - Is There Room for Energy Saving Optimization?, *IEEE Communications Magazine* (2011).
- [7] L. A. Barroso, U. Hölzle, The case for energy-proportional computing, *IEEE Computer* 40 (2007) 33–37.
- [8] IEEE Std 802.3az: Energy Efficient Ethernet, 2010.
- [9] K. Christensen, P. Reviriego, B. Nordman, M. Bennett, M. Mostowfi, J. A. Maestro, Ieee 802.3az: the road to energy efficient ethernet, *IEEE Communications Magazine* 48 (2010) 50–56.
- [10] DGS-1100 EasySmart Switches 16/24 Port Gigabit Switches Datasheet, D-Link, 2011.
- [11] 8-Port Gigabit GREENnet Switch Datasheet, Trendnet, 2011.
- [12] GEU-0820 8-Port Gigabit Switch Datasheet, Level One, 2011.
- [13] HP E8200 zl v2 Switch Series Technical Specifications, Hewlett-Packard, 2011.
- [14] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, S. Wright, Power awareness in network design and routing, in: Proceedings of the 2008 IEEE International Conference on Computer Communications, IEEE, 2008, pp. 457–465.

- [15] M. Gupta, S. Singh, Dynamic ethernet link shutdown for energy conservation on ethernet links., in: Proceedings of the 2007 IEEE International Conference on Communications, IEEE, 2007, pp. 6156–6161.
- [16] P. Reviriego, K. Christensen, J. Rabanillo, J. A. Maestro, An initial evaluation of energy efficient ethernet, IEEE Communications Letters 15 (2011) 578–580.
- [17] P. Reviriego, J. A. Hernández, D. Larrabeiti, J. A. Maestro, Performance evaluation of energy efficient ethernet, IEEE Communications Letters 13 (2009) 697–699.
- [18] R. Seifert, J. Edwards, The All-New Switch Book: The Complete Guide to LAN Switching Technology, Wiley Publishing, 2 edition, 2008.
- [19] M. Mostowfi, K. Christensen, Saving energy in lan switches: New methods of packet coalescing for energy efficient ethernet, in: Proceedings of the 2011 International Green Computing Conference and Workshops, IGCC '11, IEEE Computer Society, Washington, DC, USA, 2011, pp. 1–8.
- [20] 8-GE port switch with integrated SerDes, BCM5398 Datasheet, Broadcom, 2006.
- [21] 24-Port Layer-2 Gigabit Ethernet Switch with 12 Integrated Copper PHYs and Embedded 416 MHz CPU, Product Brief, Vitesse, 2010.
- [22] M. V. Lau, S. Shieh, P.-F. Wang, B. Smith, D. Lee, J. Chao, B. Shung, C.-C. Shih, Gigabit ethernet switches using a shared buffer architecture, IEEE Communications Magazine 41 (2003) 76–84.
- [23] S. Ricciardi, D. Careglio, U. Fiore, F. Palmieri, G. Santos-Boada, J. Solé-Pareta, Analyzing local strategies for energy-efficient networking, in: Proceedings of the IFIP TC 6th international conference on Networking Workshops (SUNSET), NETWORKING'11, Springer-Verlag, Berlin, Heidelberg, 2011, pp. 291–300.
- [24] D. Kharitonov, Time-domain approach to energy efficiency: High-performance network element design, in: Proceedings of the 2009 IEEE Global Communications Conference (GLOBECOM) Workshops, IEEE, 2009, pp. 1–5.

- [25] V. Sivaraman, A. Vishwanath, Z. Zhao, C. Russell, Profiling per-packet and per-byte energy consumption in the NetFPGA Gigabit router, in: Proceedings of the 2011 IEEE Conference on Computer Communications (INFOCOM) Workshops, IEEE, 2011, pp. 331–336.
- [26] A. Odlyzko, Data networks are lightly utilized, and will stay that way, Review of Network Economics 2 (2003) 210–237.
- [27] P. Riviriego, V. Sivaraman, Z. Zhao, J. Maestro, A. Vishwanath, A. Sanchez-Macian, C. Russell, An Energy Consumption Model for Energy Efficient Ethernet Switches, in: Proceedings of the International Workshop on Optimization Issues in Energy Efficient Distributed Systems (OPTIM 2012), Madrid, Spain.
- [28] SMC GS801 EZ Switch 8-port Gigabit Ethernet Switch, SMC-Networks, 2011.
- [29] Ixia traffic generator, www.ixiacom.com, 2012.
- [30] Computer Control Instrumentation (CCI), Power-Mate Power Monitor, <http://www.power-mate.com.au/>, 2012.
- [31] C. Gunaratne, K. Christensen, B. Nordman, S. Suen, Reducing the Energy Consumption of Ethernet with Adaptive Link Rate (ALR), IEEE Transactions on Computers 57 (2008) 448–461.
- [32] S. Herrera-Alonso, M. Rodriguez-Prez, M. Fernandez-Veiga, C. Lopez-Garca, How Efficient is Energy-Efficient Ethernet?, in: Proceedings of ICUMT, Madrid, Spain.
- [33] M. A. Marsan, A. F. Anta, V. Mancuso, B. Rengarajan, P. Reviriego, G. Rizzo, A Simple Analytical Model for Energy Efficient Ethernet, IEEE Communications Letters 15 (2011) 773–775.
- [34] A. Chatzipapas, V. Mancuso, Modelling and Real-Trace-Based Evaluation of Static and Dynamic Coalescing for Energy Efficient Ethernet, in: Proceedings of ACM e-Energy.
- [35] CAIDA, Packet Length Distribution at NASA AIX, www.caida.org/analysis/AIX/plen_hist/, 2000.

- [36] D. Dove, Energy Efficient Ethernet: Switching Perspective, presented at IEEE 802.3 meeting, January, 2008.
- [37] H. Barrass, Options for EEE in 100G, presented at IEEE 802.3 meeting, January, 2012.