

Energy Efficiency in Ethernet Receivers

Pedro Reviriego, Chris J. Bleakley and Juan Antonio Maestro

Abstract: The efficient use of energy in communications is an area of growing interest. Many existing standards in the area of wire-line communications have not been designed to make efficient use of energy. For example, many standards specify that the transmitter and receiver operate at full power even when no data is being sent. This is the case in many Ethernet standards, resulting in a considerable waste of energy. This has triggered efforts to develop new standards, such as Energy Efficient Ethernet, which have the aim of reducing energy consumption when no data is being transmitted over a link. Herein, we note that further energy savings are possible in many existing standards when data is being transmitted and channel conditions are better than the worst case for which the receiver was designed. In these cases, receivers operate with a Signal to Noise Ratio (SNR) that is above the required level. The idea presented in this work is to trade the excess SNR for lower power consumption by reducing processing in the receiver. This technique can be applied to existing standards, as it only requires modifications to the receiver and is transparent to the remote transmitter device. To illustrate the approach, a 1000BaseT receiver is studied and a number of techniques are proposed to reduce the energy consumption of the receiver. The results show that the potential energy savings are significant - up to 50% in some cases.

I Introduction

Communications equipment globally consumes a vast amount of energy. The core of the Internet alone consumes 6 TWh per year [1]. This increase in energy consumption has arisen due to both an increase in the number of devices and the lack of focus on energy efficiency in the design of wire-line communications devices.

There are hundreds of millions of installed Ethernet links. When powered on, each link consumes a substantial amount of energy, even if no data is being transmitted [2]. This leads to the waste of over 3 TWh per year [2]. The cause of this waste is the lack of energy efficiency criteria in the design of the original Ethernet standards. The standards specify that master and slave devices must send Idle pulses to maintain the link even in the absence of data. This issue is now being addressed by the IEEE 802.3az Task Force, Energy Efficient Ethernet. The Task Force is set to agree a standard by 2010 that will introduce energy efficiency enhancements to existing Ethernet standards [3].

The Energy Efficient Ethernet Task Force and other similar groups aim to ensure energy efficiency operation while maintaining backward interoperability. By definition, standards give designers considerable freedom in implementation with the aim of promoting competition and innovation. Thus, energy efficiency needs to be considered a priority in the design process as well as in the standardization process. Traditionally, designers aim to reduce power consumption under worst case conditions. This has been driven by the need to minimize the cost of packaging the device and cooling system. For example, a plastic package may be sufficient for a low power consumption device while a higher consumption device might need a ceramic package and a heat sink. Additionally, a system that uses devices with higher power consumption may need forced ventilation to avoid overheating, increasing costs and lowering reliability.

While worst case power consumption is still important for cost reasons, today's designer should have the additional goal of reducing energy consumption in all cases.

When the channel is better than worst case, the receiver typically operates above the required Signal to Noise Ratio (SNR) and achieves a lower than required Bit Error Rate (BER). In this situation, it makes sense from an energy efficiency point of view, to trade excess SNR for lower power consumption by reducing the processing done in the receiver. Similar energy efficient techniques are commonly applied in devices intended for battery operated systems, such as laptops and mobile phones. For these devices, considerable design effort is devoted to reducing energy consumption so as to extend battery life. However, since it is a wired standard, these techniques are not commonly applied to Ethernet devices. Designers assume that, in most cases, Ethernet devices will operate from a mains power supply and energy savings for environmental reasons have not been a priority. The idea of energy scalable receivers is explored in this paper and applied to the design of a 1000BaseT (Gigabit) Ethernet receiver. The results show that the proposed approach can provide substantial energy savings in a large number of cases.

The rest of the paper is structured as follows, in section II related work in the area of energy efficient communications is reviewed then a brief description of 1000BaseT receivers is provided in section III. In section IV energy efficient receiver scaling techniques for 1000BaseT are proposed and their effectiveness is discussed. Finally, conclusions are presented in section V.

II Energy Efficient Communications

Energy efficiency in the Internet is receiving increasing attention. Early works, such as [1], proposed reducing the energy consumption of routers and switches by using low power modes of operation when there is little or no network activity. This idea has also been applied to reduce the power consumption of network nodes and end user devices. In [4], a method was proposed to offload some packet processing functions to the Network Interface Card (NIC) so that the PC processor can enter low power modes for extended periods of time. In [5], the use of low power modes in different elements of LAN switches were studied. Similarly, a proposal was made in [6] to modify TCP in order to put the connection into a sleep mode when there is no activity. The energy efficiency in TCP was considered from a different perspective in [7]. In that work, the energy cost of implementing TCP was studied and proposals were made to reduce the energy costs associated with TCP processing in a PC.

The idea of entering low power modes when there is little activity has also been applied to Ethernet. The authors of [2] propose a method to reduce the speed of Ethernet links when there is little traffic. This speed reduction results in considerable power savings. In the case of switching from 1Gbps to 100Mbps, power savings were estimated to be over 50%. The Energy Efficient Ethernet Taskforce is working on modifications to existing standards that would reduce power consumption by allowing the physical layer device (PHY) to enter a low power mode when there is little traffic [4]. These modifications to the PHY allow for much faster wake ups times than the link layer speed changes proposed in [2]. In [8], the encodings used in two different Ethernet standards, 100BaseTx and 1000BaseT, are analyzed in terms of their energy efficiency. Unfortunately, the analysis does not consider the complexity of Ethernet physical layer devices and the results are of limited interest.

To summarize, most previous work has focused on energy efficiency at higher layers while work at the physical layer has concentrated on using low power states when there is little traffic [2]. Improving the energy efficiency of Ethernet devices when the link is

active and traffic is flowing seems to be a promising line of enquiry that is complementary to existing work.

III Overview of 1000BaseT Receivers

Gigabit Ethernet over copper, also known as 1000BaseT, is part of the IEEE 802.3 standards and enables 1Gbps full duplex communication over up to 100m of unshielded twisted pair category 5 cabling with a BER of 10^{-10} [9]. To achieve 1Gbps four twisted pairs are used, as shown in Figure 1. On each pair, full duplex communication takes place at 125M symbols per second. Each symbol is coded using PAM5 modulation giving a total of 625 possible combinations over the four wires. As only 256 combinations are needed to code 8 bits, the remaining combinations are used to provide redundancy that can be exploited to provide a gain of up to 6dB using a Viterbi decoder at the receiver. To achieve the proposed BER of 10^{-10} , an SNR of 19.3dB is needed at the input to the Viterbi decoder [10].

From Figure 1, we can see that each receiver will receive, in addition to the transmitted signal, the following interferers: echo from the signal transmitted on the same pair, NEXT (Near End Crosstalk) from the signals transmitted on the other three cable pairs and FEXT (Far End Crosstalk) from the signals received from the other three cable pairs. Cable response and attenuation degrades SNR at the receiver by introducing Inter Symbol Interference (ISI) and reducing the level of the received signal. The receiver incorporates a number of elements to eliminate or reduce most of these effects.

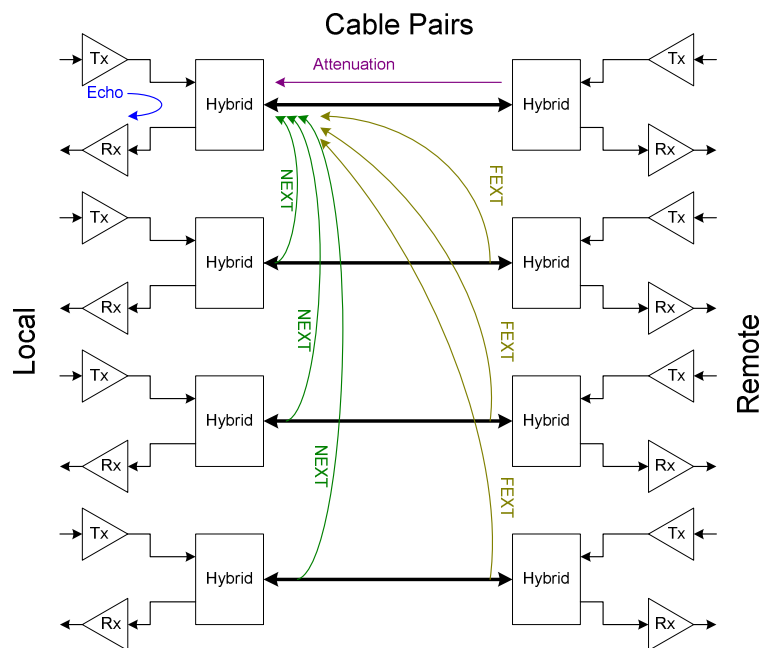


Figure 1 Gigabit Ethernet over Copper System and Noise Sources.

The channel, as specified in the Gigabit Ethernet standard, may include a number of cable segments (up to six) that are linked by a number of connectors (up to 4). Each cable segment may be of a different type of cable (horizontal cable, patch cord, under-carpet cable, etc). A typical channel configuration is shown in Figure 2, in which the PC is connected through a short patch cord to a long horizontal cable that ends close to the switch or hub that is connected by means of a short patch cord.




Figure 2 Configuration of a typical Ethernet channel.

The attenuation of the signal from the remote end increases with cable length and, as many of the impairments are not reduced with cable length (e.g. echo), longer cables represent the worst case for a 1000BaseT receiver. The maximum cable length is specified in the standard as 100m and therefore receivers are designed to operate at the required SNR of 19.3dB plus an engineering margin. Figure 3 shows a simplified block diagram of a 1000BaseT receiver. As shown, a 1000BaseT receiver typically includes an echo canceller, three NEXT cancellers, an equalizer and, in some cases, three FEXT cancellers per pair plus a Viterbi decoder shared across all pairs.

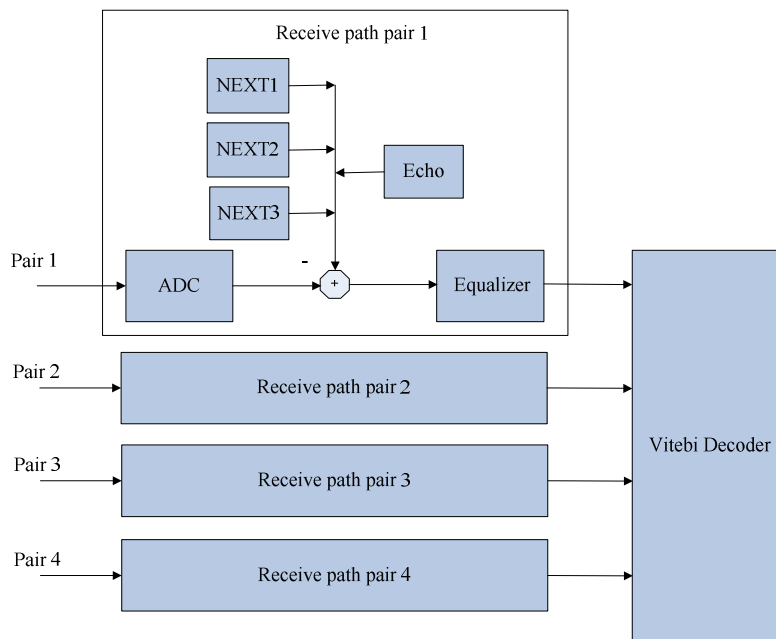


Figure 3 Simplified 1000BaseT receiver block diagram.

All these elements are designed to ensure that the target BER is achieved for the worst case channel. However, in most situations, Ethernet links consist of less than 100m of cable [11]. In fact, in data centers, connections are typically less than 30m [12]. Figure 4, shows the typical distribution of Ethernet link lengths in commercial office cabling [11]. It can be observed that over half of the connections are 50m or less.




Figure 4 Typical Ethernet channel length distribution [11].

This means that most 1000BaseT receivers operate above the required SNR and therefore most links have BERs significantly less than 10^{-10} . Figure 5 shows BER versus SNR for 1000BaseT when a Viterbi decoder is used (line with blue circles). At an SNR of 22dB, BER is around 10^{-20} . As can be seen, a further increase in SNR has negligible impact as the BER is already very small. This abrupt drop in BER, once the target SNR is achieved, is common in many wire-line standards. For example, in 10GbaseT, the coding used results in an even steeper BER drop, once the target SNR is exceeded [13]. Thus, from an application point of view, there is little or no benefit when SNR exceeds the target. However, most PHYs continue to operate with all the elements shown in Figure 3 active, as if the channel was worse case.

When reducing energy consumption is a priority due to environmental concerns, it is beneficial to explore options that trade excess SNR for reduced energy consumption. This can be done at the transmitter or at the receiver. Changes at the transmitter will effect the remote receiver. Therefore, the transmitter would have to know the SNR at the receiver and the effects of local changes on it. In addition, the transmitted signal must conform to the standard. This constrains the options for power reduction at the transmitter. For example, for short cables, one option would be to use reduced amplitude levels in transmission. However, any significant reduction in the amplitude level will violate the 1000BaseT standard. On the other hand, modifications at the receiver are transparent to the remote end and have no implications for standard compliance, as long as the required BER performance is met.

Scaling the computational complexity of the receiver is a reasonable option when the power consumption of a physical layer device is analyzed in more detail. It turns out that, in most cases, more power is consumed in the receiver than in the transmitter. This is due to the complex processing at the receiver by the Digital Signal Processing (DSP) elements and also the stringent requirements on the Analog to Digital Converters in terms of speed and accuracy. The options for scaling are considered in the next section.

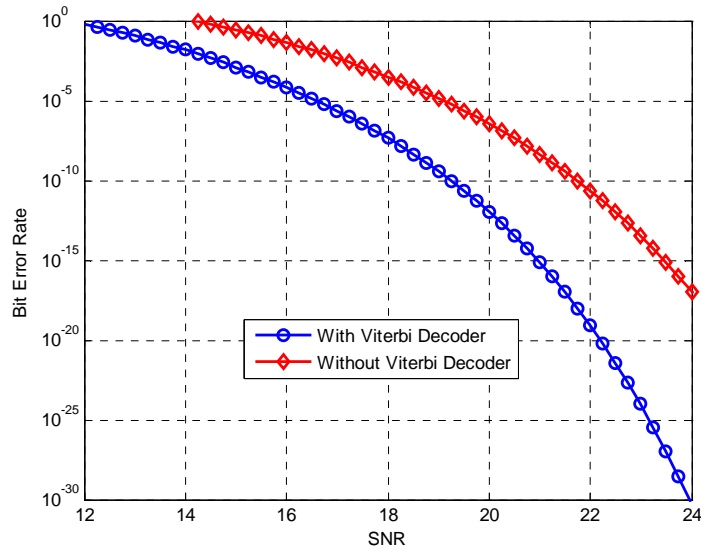


Figure 5 Bit Error Rate vs SNR for 1000BaseT with and without Viterbi Decoder

IV Energy Efficient Techniques for 1000BaseT Receivers

Given the complexity of a 1000BaseT receiver, there are many options to trade SNR for power consumption when there is excess SNR. The evaluation and detailed analysis of the different alternatives is quite complex and requires a deep knowledge of the cable channels and the physical layer device architecture [10]. Both are outside the scope of this paper, whose aim is to introduce the idea of trading SNR for energy consumption and to show the potential benefits. A detailed analysis of each of the alternatives and their implementation in a given device architecture is left for future work.

Power consumption in digital CMOS circuits increases with circuit activity and can be reduced by disabling some of its functionality. This can be done by gating the clock to some elements or by feeding them with a constant input pattern such that toggling is reduced in the combinational logic. For some of the analog components there are similar power reduction techniques but, in this paper, we concentrate on the potential savings in the DSP elements. This serves to illustrate the potential benefits and the work can be easily extended to the other elements. The reduction in energy consumption achieved will differ from implementation to implementation. Our estimates are based on our previous work on Ethernet receivers (see for example [14]) and are given to provide a rough idea of the potential savings.

The first alternative is to switch off the Viterbi decoder. This will reduce the SNR by approximately 3dB [10], as shown in Figure 5. This means that an SNR of over 22dB would be needed with the Viterbi off, which is easily obtained for short cables. If we start from an SNR of around 24dB, disabling the Viterbi would leave us with a BER below 10^{-15} and 2dB of excess SNR to trade for additional power savings. The savings obtained by switching off the Viterbi decoder can be estimated to be around 15% of the total power consumption in the digital part of the device.

If we turn to the many digital filters that are present in the receiver, there are many options to reduce power consumption. A simple option is to disable the NEXT cancellers, as for short cables the remote signal is strong compared to the near end crosstalk. This can be done for cables shorter than 30m [12]. Again, the savings will

change from implementation to implementation but we can assume a value of 15% for this option.

As the number of taps in the echo cancellers is dimensioned for the longest cables, another straightforward option is to disable the echo canceller taps that correspond to the far end echo for long cables. For a 30m cable, more than half of the echo canceller taps can be disabled providing power savings of around 10%. Finally, for short cables, the remote signal is less distorted and therefore the number of taps in the equalizer can be reduced. In this case, the savings will be around 5%.

Another option is to reduce the bit-widths of either the coefficients or the data in the filters. Disabling the least significant bits will translate directly into a reduction in SNR but also into power savings. The number of bits can be reduced at the output of the ADCs. This reduced data width will feed into the various filters as shown in Figure 3. The savings can be estimated to be up to 5% for this option.

A summary of the estimated savings for the different techniques is presented in Table I. It should be noted that these are rough estimates and can vary significantly for different implementations. However, as the purpose is to illustrate the possible savings, we believe they can be useful. Potentially, up to 50% of the power consumption in the digital circuitry can be saved for short cables. This is in line with previous estimates in [12], in which the power consumption of a 10Gbps physical layer transceiver designed to operate on channels up to 30m was found to be 50% less than that of a 10GbaseT transceiver designed to operate up to 100m of cable.

To estimate the potential savings of adopting energy efficiency criteria in the design of Ethernet receivers, we can assume that a 50% saving in the digital circuitry power consumption can be achieved for 50% of the links (i.e. those shorter than 50m). Therefore, the active power consumption of the digital part of Ethernet physical layer devices can be reduced by 25% when averaged over all cables. This is a substantial saving that is complementary to the savings obtained by adopting the Energy Efficient Ethernet standard. To put the potential savings in perspective, an example can be used. Let us assume a PC is always on and connected using Ethernet. Additionally, let us say that, on average, the link is active 4 hours a day. Adopting the Energy Efficient Ethernet standard will reduce the energy consumption significantly for the time that the PC is on but no data is sent, that is for 20 hours. The proposed receiver techniques would save power for the remaining 4 hours. Obviously, the largest savings would be obtained by implementing Energy Efficient Ethernet. Nevertheless, the receiver scaling techniques proposed herein provide significant savings and, more importantly, the savings are independent, so that they sum. This makes the proposed techniques an ideal complement to Energy Efficient Ethernet in further reducing the energy consumption of Ethernet devices.

The overall gains can be estimated by considering the number of devices, their activity and their typical power consumption. Typically, a 1000BaseT physical layer device consumes 0.5 to 1W while for 10GbaseT the range is multiplied by a factor of 10. The overall savings of implementing low power modes on Ethernet has been estimated to be over 3 TWh per year [2]. Although it is difficult to make an accurate estimate, the proposed receiver scaling techniques could provide a significant reduction in the remaining 3 TWh.

Table I Savings estimates for the different techniques

Technique	Estimated Savings
Disabling Viterbi Decode	15%
Disabling NEXT cancellers	15%
Reducing Echo canceller taps	10%
Reducing equalizer taps	5%
Reducing bitwidths	5%
Total savings	50%

The implementation of the proposed techniques requires additional control logic in the physical layer devices to disable the various elements when there is excess SNR. There are a number of ways to implement this control. One is to use all the elements in the receiver during startup and later, once the link is established, measure the SNR and decide which elements can be disabled. Since the link is already established, this must be done in such a way that the disabling of elements does not affect receiver performance. Another alternative is to estimate the cable length prior to setting up the link and to disable some elements before startup. In this case, the link is established using the normal startup process and the device enters operation with no further changes in the receiver. Although cable length estimation is not defined in the Ethernet standards, many manufacturers implement cable length estimation in their physical layer devices to provide on chip diagnostics. Cable length information can also be useful for the startup procedure of the device and can be extended to improve energy efficiency. Recent standards, such as FC-BaseT, incorporate cable length estimation procedures as part of the physical layer specifications [15]. These options can be combined by disabling some elements upfront and then, if there is still excess SNR, disabling other elements after the link has been established. The details of both alternatives are not explored further herein because this would require a discussion of 1000Base-T start up and cable length estimation, which are outside the scope of this work.

From the previous discussions it becomes apparent that significant energy savings can be obtained in the receiver when the cable is short. This leads us to another interesting conclusion: if receivers include scaling techniques then using better cables will reduce power consumption. In Figure 6, the attenuation versus frequency of 100m of category 6 and 7 cable is shown. It can be seen that using better cables (category 7) reduces attenuation. The improvements in cable technology also apply to other impairments such as near-end and far-end crosstalk and, in particular, alien crosstalk from other cables. Thus, better cables will allow energy savings by enabling receiver complexity scaling. This is an additional incentive to upgrade cable plant.




Figure 6 Attenuation vs frequency for 100m of category 6 and category 7 cables

V Conclusions

Energy efficient design of wire-line communication receivers has been studied in this paper. The techniques applied to 1000BaseT receivers show that significant energy savings can be obtained. This is a consequence of the way wire-line standards and devices have been designed. They aim to provide a fixed connection speed at a given BER level and do not trade excess SNR for reduced power consumption under benign channels conditions. The idea of trading excess SNR for energy consumption was assessed and has been shown to provide substantial potential energy savings. An interesting corollary from the analysis is that better cabling will reduce energy consumption when receivers implement techniques to trade excess SNR for energy consumption.

References

- [1] M. Gupta and S. Singh "Greening of the Internet" Proc. of ACM SIGCOMM pp 19-26, August 2003.
- [2] C. Gunaratne, K. Christensen, B. Nordman and S. Suen "Reducing the Energy Consumption of Ethernet with Adaptive Link Rate (ALR)" IEEE Transactions on Computers, vol. 57, issue 4, April 2008 pp.: 448 – 461.
- [3] IEEE P802.3az Energy Efficient Ethernet Task Force (<http://grouper.ieee.org/groups/802/3/az/public/index.html>)
- [4] K. Sabhanatarajan, A. Gordon-Ross, M. Oden, M. Navada, and A. George, "Smart-NICs: Power Proxying for Reduced Power Consumption in Network Edge Devices," IEEE Computer Society Annual Symposium on VLSI (ISVLSI), April 2008
- [5] M. Gupta, S. Grover and S. Singh "A feasibility study for power management in LAN switches" 12th IEEE International Conference on Network Protocols, Oct. 2004 pp. 361- 371.
- [6] L. Irish and K. Christensen "A 'Green TCP/IP' to Reduce Electricity Consumed by Computers" Proceedings of IEEE Southeastcon, pp. 302-305, April 1998.

- [7] B. Wang and S. Singh “Computational energy cost of TCP” INFOCOM, March 2004 pp. 785 - 795 vol.2.
- [8] Y. Chen, T. Xiaoxiao and R.H. Katz “Energy efficient Ethernet encodings” IEEE Conference on Local Computer Networks, Oct. 2008, pp: 122-129.
- [9] “Physical Layer Parameters and Specifications for 1000Mb/s operation over 4-Pair of Category 5 Balanced Copper Cabling , Type 1000Base-T” IEEE standard 802.3ab
- [10] M Hatamian, O.E. Agazzi, J. Creigh, H. Samueli, A.J. Castellano, D. Kruse, A. Madisetti, N. Yousefi, K. Bult, P. Pai, M. Wakayama, M.M. McConnell and M. Colombatto “Design considerations for gigabit Ethernet 1000Base-T twisted pair transceivers” IEEE Custom Integrated Circuits Conference, May 1998 pp:335 – 342.
- [11] B. Booth, A. Flatman, G. Zimmerman and S. Rao “IEEE 802 10GBaseT Tutorial” presentation at IEEE 802.3 November 2003 meeting.
- [12] N. V. Bavel, D. Dove, A. Flatman and M. McConnell “Short Haul 10Gbps Ethernet Copper PHY Call for Interest” presentation at IEEE 802.3 November 2005 meeting.
- [13] S. Powell and B.Z. Shen “Specification and Performance of Proposed LDPC (2048,1723) Code” presentation at IEEE 802.3 January 2005 meeting.
- [14] P. Reviriego and C. Murray “DSP Dimensioning of a 1000Base-T Gigabit Ethernet PHY ASIC” XVII Conference on Design of Circuits and Integrated Systems, Nov. 2002 pp.71-76
- [15] “Fibre Channel BaseT (FC-BaseT)” ANSI INCITS 435-2007.