# Introducing Energy Efficiency in the VDE 0885-763 Standard for High Speed Communication over Plastic Optical Fibers

*Pedro Reviriego, Universidad Antonio de Nebrija*
*Rubén Pérez-Aranda and Carlos Pardo, KDPOF*

## ABSTRACT

In September 2013, the VDE 0885-763 standard for high-speed communication over plastic optical fiber (POF) will be published. The new standard enables 1 Gb/s communication over 50 meters of POF and is intended for use in automotive, industrial, and home networking applications. One of the goals of the standard is to achieve good energy efficiency, and to that end it incorporates a low power mode. This is similar to the case of Ethernet where a low-power mode was introduced by the Energy Efficient Ethernet (IEEE 802.3az) standard in 2010. In this article, the energy efficiency mechanisms incorporated in the VDE 0885-763 standard are described and evaluated. The evaluation is done by simulation, and the results are also compared to those of Energy Efficient Ethernet. The conclusion is that VDE 0885-763 will provide good energy efficiency even at low/medium loads.

## INTRODUCTION

The development of standards for physical layer communications is a complex task. In addition to cutting edge signal processing schemes, they include aspects related to the channel, and the interfaces with the physical medium and the upper layers. The standard process typically involves experts in different areas that come from several organizations with conflicting interests. This makes it difficult to achieve consensus. In addition to those challenges, in recent years, a new aspect has been introduced in physical layer communication standards: energy efficiency. This is due to the increased awareness of the energy consumption of communications systems and its growth over time. Extensive research on energy-efficient communications has been done in the last years in academia. Therefore, the introduction of energy efficiency in standards opens an opportunity for collaboration between researchers and companies working on standardization. One example of such collaboration is the introduction of energy efficiency mechanisms into the Association for Electrical, Electronic & Information Technologies (VDE) 0885-763 standard for high-speed communication over plastic optical fiber (POF) developed by the VDE Working Group on Polymer Optical Fiber (412.7.1). In this case, the expertise of a research group on energy efficiency was combined with the expertise on advanced signal processing of a company to make a proposal to incorporate energy efficiency mechanisms into the standard. This article describes the solution that was adopted in the standard and evaluates its performance by simulation.

The use of POF for data transmission has gained increased interest in recent years [1]. POFs are made of plastic materials and therefore have a competitive cost when compared to traditional glass optical fibers or copper twisted pairs. In contrast to traditional optical fibers, POFs do not require sophisticated installation, which is critical for home networking applications. Compared to Category-5 twisted copper pairs, POF is easier to install thanks to its reduced diameter and because mains ducts can be reused. At the same time, since no electrical signals are used for data transmission, electromagnetic interference (EMI) is not an issue. This is a key advantage for some industrial applications. Finally, from a safety point of view, POF cannot cause short circuits and cannot damage the human eye due to the nature of the light sources that are used. In summary, POFs have a number of features that make them attractive for a large range of applications. They are currently widely used in cars to connect devices and in industrial systems [1], and they can be an interesting option for home networking [2].

The use of POF for automotive applications has been fostered by the media oriented systems transport (MOST) specification [3]. MOST has been adopted by most car manufacturers, and is the de facto standard for multimedia and infotainment networking in the automotive industry. Each year millions of cars equipped with MOST networks are sold throughout the world. For

  

The main reason for the development of EEE was that the consumption of most Ethernet devices did not depend on the traffic load. That is, the power consumption was constant regardless of whether there was actual data traffic.

MOST, the key advantages of POF are its immunity to EMI, the low weight compared to copper cables, and its safety against short circuits. MOST supports the use of POF with speeds ranging from 25 to 150 Mb/s. In automotive applications the length of the links is limited to 15 m, and the physical layer reuses existing optical transceivers. For industrial applications, immunity to EMI and safety are the main advantages of POF. Channel lengths can be potentially large and typically can reach hundreds of meters.

POF is also emerging as a technology for home networking [2]. In this area, the low cost and ease of installation are its main advantages. The channel lengths in this case can also be large; typically up to a few hundred meters are considered. The installation is also less controlled than in automotive or industrial applications, and this can cause channel impairments due to bending or non optimal connections.

The existing standard (MOST) for POFs supports data rates only up to 150 Mb/s. This is insufficient for future systems and low compared to other alternatives such as Ethernet, which is widely deployed in LANs and industrial applications [4]. The need for higher data rates led to the formation of a Working Group in the VDE to define a standard for high-speed communication over POF. The group was assigned number WG 412.7.1 and started its activity in October 2009 with the objective of providing 1 Gb/s over 50 m of POF. Several baseline proposals were presented for the communication system based on discrete multitone (DMT), simple non-return-to-zero (NRZ), and four-level pulse amplitude modulation (PAM-4). Finally, in VDE WG 412.7.1 a solution based on the use of an adaptive multilevel coset coding and Tomlinson-Harashima precoding was selected. Once a baseline proposal was selected, the standard progressed to a phase in which all the technical details were discussed. The issue of energy efficiency was raised at this point. As most of the participants were not familiar with the topic, a university research group was brought in to give a tutorial presentation. This eventually led to their involvement in the process and in the proposal for the standard. The new standard provides a general specification that supports different speeds and channels using the same underlying technology [5]. One of the key configurations supported is 1 Gb/s bidirectional communication for POF links of up to 50 m. This is clearly an alternative to Gigabit Ethernet and is expected to enable the use of POF in next generation automotive applications.

In the rest of this article, the energy efficiency mechanisms defined in the VDE 0885-763 standard are described and evaluated. The performance in terms of energy consumption vs. load is analyzed by simulation and compared with that of Energy Efficient Ethernet (EEE).

## Energy Efficiency in Wireline Physical Layer Standards

As mentioned before, energy efficiency is becoming an important issue in the design of communication systems. Most new standards consider energy efficiency as one of their goals and in some cases modifications to existing standards are also made to improve energy efficiency. This is the case of Ethernet for which a new standard was introduced in 2010 to improve the energy efficiency of the most commonly used transceivers [6]. Known as Energy Efficient Ethernet (IEEE 802.3az), the standard is expected to enable savings of a few terawatts-hour per year [7]. As Ethernet is an alternative to the proposed POF standard, it is of interest to compare both solutions in terms of energy efficiency.

To improve energy efficiency in wireline transceivers, a number of mechanisms can be used. First, transceivers are in many cases kept active even when there is no user data to transmit. This is done to keep transmitters and receivers aligned and ready to exchange data. Therefore, an obvious improvement is to define low-power modes that can be used in such cases. This is in fact the basic mechanism implemented in EEE. The use of a low-power mode normally involves a penalty in latency, as a frame that arrives for transmission when the transceiver is in low-power mode has to wait until the link is active to be transmitted. The acceptable latency is application-dependent and in EEE latencies on the order of some microseconds have been used to try to cover a wide range of applications. It is also important to note that significant energy can be consumed during mode transitions, especially when activating the link. This means that if the transition times are larger than the frame transmission time, more energy could be devoted to transitions than to actual data transmission [7]. In addition to the use of a low-power mode, other techniques that can be used to reduce energy consumption are to dynamically adapt the link rate to the data traffic or adjust the transmitted power to the current conditions of the link [8]. The first mechanism is not appropriate for optical links, while the second provides limited savings as transmission is still active when there is no user data. This makes the use of a low-power mode the preferred solution for links that operate at low loads. In any case, most of these mechanisms require cooperation from the link partner and therefore need to be specified in a standard to ensure interoperability of different implementations.

The main reason for the development of EEE was that the consumption of most Ethernet devices did not depend on the traffic load. That is, the power consumption was constant regardless of whether there was actual data traffic. This is because in legacy Ethernet systems an IDLE signal is sent when there is no data to transmit. This IDLE signal, despite its name, requires energy consumption similar to that of data transmission. Figure 1 illustrates the scheme adopted in EEE that introduces the concept of low power idle (LPI) to reduce energy consumption 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 refresh signal is transmitted to keep the receiver state aligned with current conditions. A device can enter the LPI mode with a sleep transition that requires $T_s$ seconds and go back to the active state with a wake transition that requires $T_w$ seconds. The minimum values of $T_w$ and $T_s$ for 1

Gb/s in the IEEE 802.3az standard are 16 μs and 182 μs, respectively. Those values are larger than the time required to send a large (1500-byte) frame which at 1 Gb/s is only 12 μs. The energy consumption in LPI is much lower than in the active state. A value of 10 percent of that of the active mode is commonly assumed [7]. However, during transitions energy consumption is significant and can be similar to that of the active mode. Finally, it is important to note that in EEE mode transitions can occur at any point in time as there is no periodic frame structure.

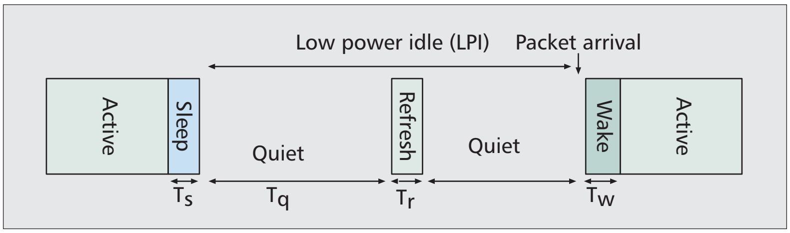


**Figure 1.** *Transitions between the active and low-power modes in Energy Efficient Ethernet.*

## Energy Efficiency in the VDE 0885-763 Standard

As mentioned in the introduction, the VDE 0885-763 standard will be approved in September after three years of standardization work. To achieve data rates of up to several gigabits per second, the standard incorporates complex signal processing functions. The multilevel coset coding (MLCC) technique [9], using three levels of encoding, is used with different degrees of coding for the bits in the constellation. The coding used in the two lowest levels is based on Bose–Chaudhuri–Hocquenghem (BCH) codes. The third upper level is left uncoded. For reference, the overall scheme is illustrated in Fig. 2.

Equalization is done at both the transmitter and the receiver. Tomlinson-Harashima precoding (THP) [10, 11] is used at the transmitter using the coefficients provided by the link partner. This avoids the error propagation issues of a decision feedback equalizer (DFE) implemented in the receiver. Feed forward equalization is implemented in the receiver. A specific issue of high-speed transmission over POF is the presence of nonlinearities originated mostly in the optical transmitter. The standard incorporates mechanisms to enable the use of nonlinear equalization when needed.

Transmission is structured in physical layer frames that are divided into pilots, physical layer headers, and codewords (CWs). The pilots are used to aid in timing recovery as well as in both linear and nonlinear equalization. At startup, the transceiver uses the pilots to estimate the equalizer settings and transmits the THP coefficients to the remote transmitter using the physical layer headers. Once the link is established, the pilots are used to track channel variations and adapt the receiver accordingly. Those variations are important and are caused, for example, by the optical source due to temperature changes. The physical layer headers are used to exchange parameters such as coding settings, frame ID, and THP coefficients. Finally, CWs carry the data bits. The frame structure is illustrated in Fig. 3. Frames are transmitted continuously, and for a 1 Gb/s rate a frame has a duration of 736.87 μs. Groups of four CWs are transmitted between pilots (S1, S2) and physical layer headers (PHs). Each of these groups takes approximately 26 μs to transmit.

The continuous transmission of frames would lead to poor energy efficiency; therefore, the standard incorporates a low-power mode. To implement a low-power mode, transmission can be stopped altogether when there is no data and then resumed when data arrives for transmission. However, this means that the receiver would have to perform a new startup, which takes a significant amount of time. This was found to be unacceptable in EEE as it increases latency on the order of milliseconds [7]. Another option is to stop transmission except for some short periods used to keep the receiver aligned so that activation is faster. This was the option implemented in EEE. In the case of the VDE

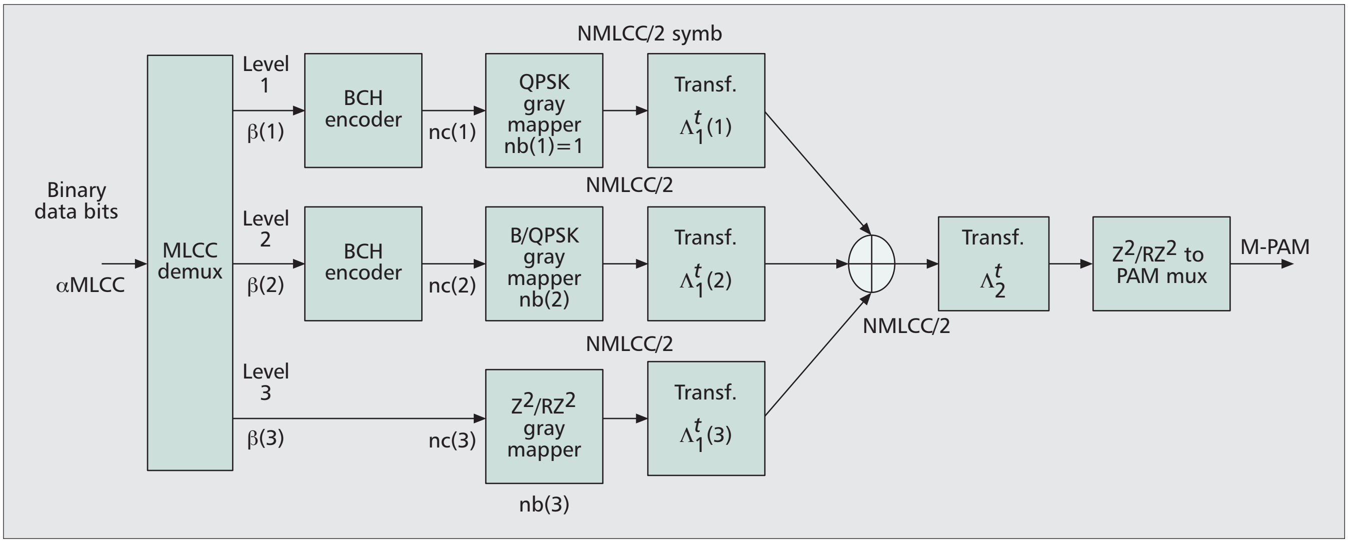


**Figure 2.** *Illustration of the multilevel coset coding (MLCC) in the VDE 0885-763 standard.*

*The transitions to low-power mode (sleep) occur immediately after a pilot or physical layer header. Conversely, wake transitions occur just before a pilot or physical layer header. Since the receiver knows the timing of the frame, there is no need to signal the transitions other than the presence or absence of light.*

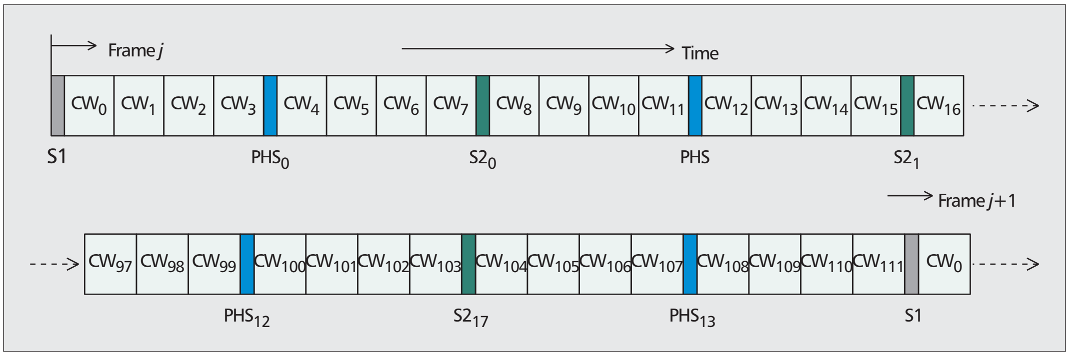


**Figure 3.** *Illustration of the frame structure in the VDE 0885-763 standard.*

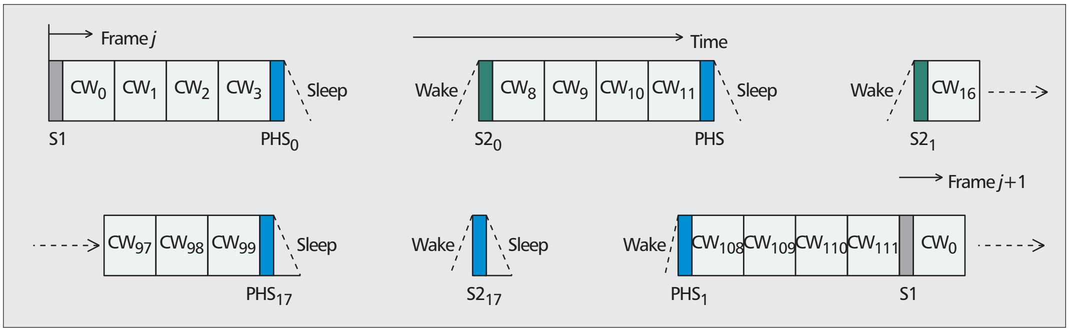


**Figure 4.** *Illustration of the use of low-power mode defined in the VDE 0885-763 standard.*

0885-763 standard, a similar technique is implemented but adapted to the frame structure shown in Fig. 3. Basically, as the pilots and physical layer headers are needed to keep the receiver aligned to the channel conditions, they should be transmitted even when in low-power mode. Therefore, the mechanism implemented allows stopping transmission for groups of four CWs. This is illustrated in Fig. 4. The transitions to low-power mode (sleep) occur immediately after a pilot or physical layer header. Conversely, wake transitions occur just before a pilot or physical layer header. Since the receiver knows the timing of the frame, there is no need to signal the transitions other than the presence or absence of light.

The power savings achieved by this mechanism depend largely on two main parameters: the traffic load and the transition time needed to wake the link. The dependence on the traffic load is studied by simulation in the next section. Minimization of transition time was one of the objectives of the proposed solution as it is key to ensure good energy savings when the link operates at medium loads. During the standardization process, technical solutions that enable the use of transitions smaller than 1 μs were presented. This is significantly lower than for EEE, where for 1 Gb/s the wake transition takes more than 16 μs and the sleep transition more than 180 μs. One reason behind this difference is that the refreshes in EEE for 1 Gb/s are spaced by several milliseconds, while in the VDE 0885-763 standard the pilots and headers are transmitted each 26 μs so that the receiver is almost ready for data transmission.

As mentioned before, the use of a low-power mode adds latency. The worst case added latency for a frame occurs when the frame arrives just after the transceiver has entered the low-power mode. In this case the latency would be close to 26 μs. In the average case the latency would be approximately 13 μs. This is comparable to EEE, where for 1 Gb/s the minimum added latency is 16 μs.

Finally, for the 1 Gb/s configuration, the standard optionally supports the use of a Gigabit media-independent interface (GMII) as the interface to the medium access control (MAC). This is the interface used in Gigabit Ethernet.

## Performance Evaluation

As discussed in the previous section, the energy savings are largely dependent on the traffic load. In fact, load is not the only parameter. Assuming that the link is activated as soon as a data packet arrives for transmission, packet arrival patterns and packet sizes are also important factors. This has been observed previously in EEE [12]. For a given load, the packet sizes determine the number of packets, which is related to the number of transitions. As a first approximation, an average packet size can be used. This is typically around 600 bytes. As for the arrivals, for links that carry aggregated traffic, independent arrivals can be assumed and have been shown to be a good approximation [13]. This has been corroborated for energy saving estimates in the case of EEE [14]. For links with low aggregation the approximation is not as good. In this case, the use of real traffic traces may be more appropriate. However, it is difficult to find a representative set of traces as the transceivers are used in a wide range of applications, and those also evolve quickly over time. Therefore, the use of independent arrivals can be a first approximation, especially when comparing two options to analyze their relative performance. This is the approach taken in this section in which the energy savings are estimated by simulation assuming independent arrivals. The performance is then compared with that of EEE for the same traffic.

To estimate the energy savings, it will be assumed as in previous studies that energy consumption in low-power mode is 10 percent of that in active mode and that transitions consume the same as the active mode. The transition times are assumed to be 1 μs for both wake and sleep, which is conservative. Several frame sizes are simulated covering small (64-byte), average (600-byte), and large (1500-byte) frames. The results are presented in Fig. 5 and show that there are significant savings compared with a device that does not implement energy efficiency mechanisms. In that case, the energy consumption would be the same as active (100 percent in the plot) for all loads. The results deviate from proportionality between load and energy consumption due to the cost associated with entering and exiting low-power mode. This deviation is larger for small frames as for the same load, there are more packets and therefore more transitions. However, even for small frames significant savings are obtained for a wide range of loads.

To put the results in perspective, they are also compared to those obtained for EEE in previous studies [12]. For EEE, the minimum transition times specified in the standard for 1 Gb/s, 16 μs to wake and 182 μs to sleep, are used in the simulations. The results are shown in Fig. 6 for 600-byte frames and show a clear advantage of VDE 0885-763 over EEE. For smaller and larger frames, the results are similar. This is due to the reduced transition times and also to the periodic frame structure used in VDE 0885-763. This structure forces the coalescing of packets that arrive when the device is in low-power mode as they have to wait until the next pilot or physical header block to wake the link. This reduces the number of transitions as several frames can be sent each time the link is activated. In summary, the results show that the VDE 0885-763 standard provides better performance than EEE, especially at medium loads. This is the consequence of considering energy optimization as a goal in the development of the standard, rather than having to modify existing standards as was the case in EEE.

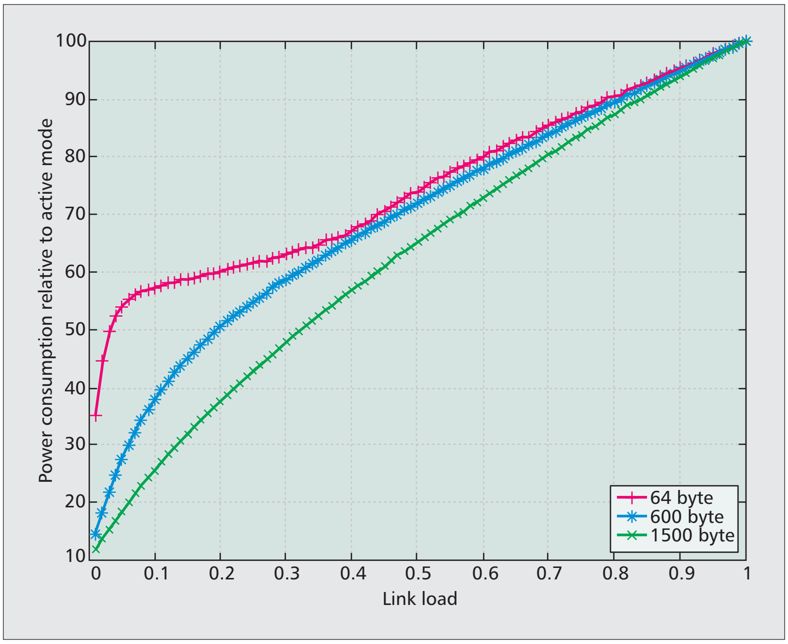


**Figure 5.** *Energy consumption vs. load for different frame sizes.*

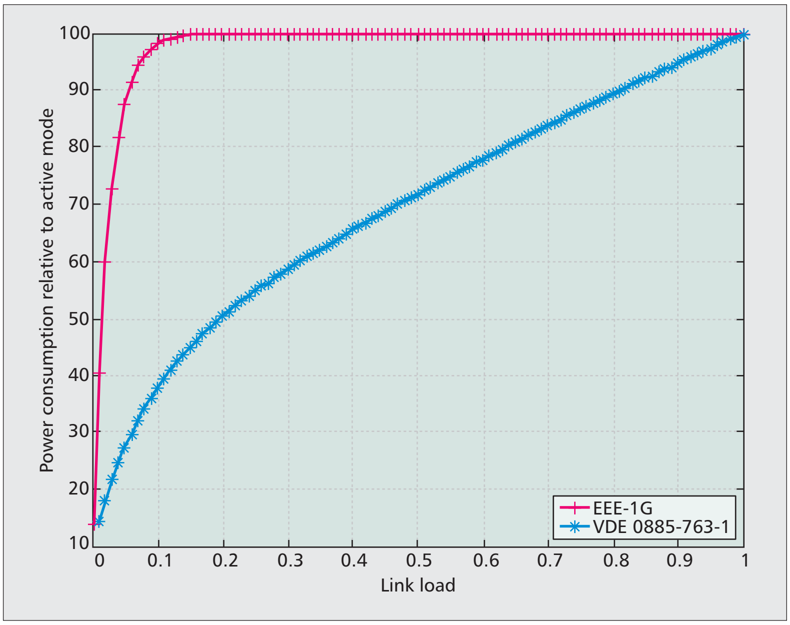


**Figure 6.** *Energy consumption vs. load for 600-byte packets: comparison of EEE and the VDE 0885-763 standard.*

## Conclusions

Energy efficiency is becoming an objective in new standards for physical layer communication systems. One example is the recent VDE 0885-763 standard for high-speed communication over plastic optical fiber. The introduction of energy efficiency in standards opens an opportunity for collaboration between research groups active in the area and industry. In this article, an example of such a collaboration that took place during the development of the VDE 0885-763 standard has been presented. The energy efficiency mechanisms introduced in the standard have been discussed, and their performance evaluated. The results show that the new standard provides good energy efficiency over a wide range of traffic loads. The performance has also been compared with that of Energy Efficient Ethernet for the same transmission rate. The comparison shows that the new standard outperforms Energy Efficient Ethernet in terms of energy savings.

## References


[1] P. Polishuk,"Plastic Optical Fibers Branch Out," *IEEE Commun. Mag.*, vol. 44, no. 9, Sept. 2006, pp. 140–48.

[2] I. Mollers *et al.*, "Plastic Optical Fiber Technology for Reliable Home Networking: Overview and Results of the EU Project POF-ALL," *IEEE Commun. Mag.*, vol. 47, no. 8, Aug. 2009, pp. 58–68.

[3] A. Grzemba, Ed., *MOST — The Automotive Multimedia Network from MOST25 to MOST150*, Franzis, 2008.

*The results show that the new standard provides good energy efficiency over a wide range of traffic loads. The comparison shows that the new standard outperforms Energy Efficient Ethernet in terms of energy savings.*


[4] J.-D. Decotignie, "Ethernet-Based Real-Time and Industrial Communications," *Proc. IEEE*, vol. 93, no. 6, June 2005, pp. 1102–17.

[5] VDE 0885-763, "Physical Layer Parameters and Specification for High Speed Operation over Plastic Optical Fibres Type HS-BASE-P," to be published, Sept. 2013.

[6] IEEE Std 802.3az, "Energy Efficient Ethernet," 2010.

[7] K. Christensen *et al.*, "IEEE 802.3az: The Road to Energy Efficient Ethernet," *IEEE Commun. Mag.*, vol. 48, no 11, Nov. 2010, pp. 50–56.

[8] C. Gunaratne *et al.*, "Reducing the Energy Consumption in Ethernet with Adaptive Link Rate (ALR)," *IEEE Trans. Computers*, vol. 57, no. 4, Apr. 2008, pp. 448–61.

[9] G. D. Forney *et al.*, "Sphere-Bound-Achieving Coset Codes and Multilevel Coset Codes," *IEEE Trans. Info. Theory*, vol. 46, no. 3, May 2000, pp. 820–50.

[10] M. Tomlinson, "New Automatic Equaliser Employing modulo Arithmetic," *IET Electron. Letters*, vol. 7, no. 5, Mar. 1971, pp. 138–39.

[11] H. Harashima and H. Miyakawa, "Matched-Transmission Technique for Channels with Intersymbol Interference," *IEEE Trans. Commun.*, vol. 20, no. 4, Aug. 1972, pp. 774–80.

[12] P. Reviriego *et al.*, "Performance Evaluation of Energy Efficient Ethernet," *IEEE Commun. Letters*, vol. 13, no. 9, Sept. 2009, pp. 697–99.

[13] T. Karagiannis *et al.*, "A Nonstationary Poisson View of Internet Traffic," *Proc. IEEE INFOCOM*, 2004.

[14] M. Ajmone-Marsan *et al.*, "A Simple Analytical Model for Energy Efficient Ethernet," *IEEE Commun. Letters*, vol. 15, no 7, Jul. 2011, pp. 773–75.


## Biographies


Pedro Reviriego (previrie@nebrija.es) is an associate professor at Universidad Antonio de Nebrija. His research interests are fault-tolerant systems, signal processing, and computer networks. He has a Ph.D. in telecommunications engineering from Universidad Politécnica de Madrid. He participated in the IEEE 802.3 standardization for 10GBASE-T and has worked on the design of 1000BASE-T transceivers and Energy Efficient Ethernet.

Rubén Pérez de Aranda (rubenpda@kdpof.com) is CTO at KDPOF. He has served as a project manager on R&D projects in digital communications systems at several microelectronics companies. His research includes channel coding and modulation schemes, nonlinear channel estimation and equalization, timing recovery algorithms, and MIMO systems. He has worked in integrated circuits for PLC, DVB, xDSL, and fiber. He holds a degree in industrial engineering from Universidad Politécnica de Madrid.

Carlos Pardo (carlospardo@kdpof.com) is CEO and president at KDPOF. He has served as a strategic project manager in the microelectronics industry and with leading R&D teams. He has extensive knowledge of the high tech silicon industry, as well as both hardware and software development. He previously worked as a technical director at SIDSA, managing R&D departments and product development. He was also part of Hewlett Packard Spain and of Design of Systems on Silicon. He holds a degree in industrial engineering from Universidad Pontificia de Comillas de Madrid.