Energy Efficiency in Industrial Ethernet: The Case of Powerlink

Juan Antonio Maestro, *Member, IEEE*, and Pedro Reviriego, *Member, IEEE*

***Abstract*—Industrial Ethernet technology enables real-time reliable communications for industrial environments. One of its key features is the use of Ethernet for the communication infrastructure. This provides a very cost-effective solution, as no communication infrastructure has to be designed *ad hoc*. In the last years, the efficient use of energy in communications has become an area of growing interest. This has triggered efforts that are now underway to develop new standards, like Energy Efficient Ethernet, with the aim of reducing the energy consumption. The proposed changes to existing Ethernet standards allow setting the links in low-power modes when there is little traffic. These changes, introduced by Energy Efficient Ethernet, will have wide implications for Industrial Ethernet. For example, the time needed to enter or exit the low-power modes may be excessive for some industrial applications. The use of low-power modes will make hubs less energy efficient than switches and may cause manufacturers to abandon the production of these devices. In this paper, the implications of Energy Efficient Ethernet on industrial environments are analyzed, and different alternatives are proposed.**

***Index Terms*—Energy efficiency, Industrial Ethernet.**

I. INTRODUCTION

THE EFFICIENT use of energy in communications is a growing concern for the industry. The massive amount of communication devices that are used nowadays and their expected growth have led to the conclusion that significant energy can be saved by applying energy-efficiency concepts in the design of communication systems. For example, in [1], the energy consumption of communication equipment used in the core of Internet was estimated to be over 6 $\mathrm{TW}\cdot\mathrm{h}$ per year.

This massive amount of energy is a result of both the large number of devices and also the lack of focus on energy efficiency in their design. A good example of both is Ethernet. There are hundreds of millions of installed Ethernet links, and each of them consumes a substantial amount of energy when active, even if there is no data being transmitted [2]. This leads to a waste of energy estimated to be over 3 $\mathrm{TW}\cdot\mathrm{h}$ per year [2]. The cause of this massive amount of energy waste is the lack of energy-efficient criteria in the design of the Ethernet standards, since the receivers and transmitters of a link are forced to operate continuously even in the absence of data. This is now being addressed by the IEEE 802.3az task force (Energy Efficient Ethernet) that is set to complete a standard by 2010 to introduce energy-efficiency enhancements to existing Ethernet standards [3]. The proposed changes allow links to enter low-power modes when there is no data to transmit. This reduces the energy consumption substantially, as most links are lightly loaded.

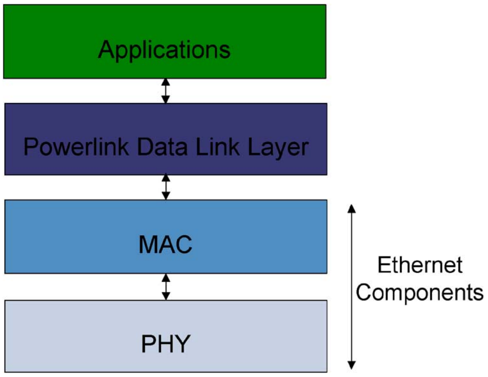

Fig. 1. Powerlink reference communication model.

A number of alternatives have been proposed to implement industrial networks [4]–[6]. However, the use of Ethernet in real-time industrial environments has gained increased interest, since it enables the reuse of existing hardware and avoids the costly development of *ad hoc* systems. The issue of energy savings in Industrial Ethernet has been previously discussed in [7] where its importance was raised to the IEEE 802.3az standardization process.

Although Ethernet was not originally designed to support real-time communications [8], a number of techniques have been proposed to adapt Ethernet so that it can be used in real-time and industrial applications [9], [10]. The main issue is to ensure that the tight delay constraints required by industrial applications are met, since frame delay is not deterministic in Ethernet. A number of studies have been performed to analyze how the delay in different Ethernet network configurations can affect industrial applications [11], [12]. Different approaches have been proposed to adapt Ethernet to these applications, in what is known as Industrial Ethernet [9]. In many of those approaches, the transmission of the frames is coordinated from a central node to avoid collisions and control the delay of the frames. One of those efforts is Powerlink [13], which reuses

Manuscript received February 27, 2009; revised June 19, 2009 and July 22, 2009; accepted November 2, 2009. Date of publication November 24, 2009; date of current version July 14, 2010. This work was supported in part by the Spanish Ministry of Science and Innovation under Grant AYA2009-13300-C03-01, by the Regional Government of Madrid, and by the European Union FEDER Programme.

The authors are with the Universidad Antonio de Nebrija, 28040 Madrid, Spain (e-mail: jmaestro@nebrija.es; previrie@nebrija.es).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TIE.2009.2036640

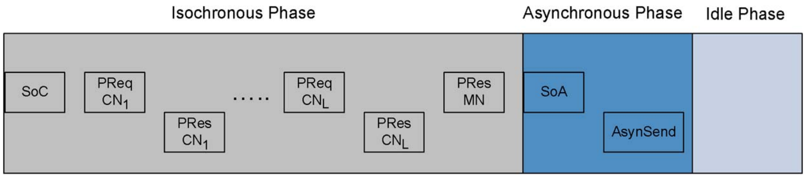

Fig. 2. Powerlink cycle.

the Ethernet infrastructure and provides real-time capabilities by specifying how transmissions are done over the Ethernet links. Through the rest of this paper, Powerlink will be used to illustrate the impact of Energy Efficient Ethernet on Industrial Ethernet. Although the implications will vary for other Industrial Ethernet approaches, the study of Powerlink will serve to illustrate the general implications. The objective of this paper is to first analyze the implications of the adoption of Energy Efficient Ethernet in Powerlink and to propose a number of alternatives to make Powerlink more energy efficient while preserving its real-time capabilities. In addition, the results can also be used as the starting point for the analysis of the implications of Energy Efficient Ethernet on other Industrial Ethernet approaches.

The rest of this paper is structured as follows. In Section II, the use that Powerlink makes of Ethernet links is reviewed. Then, a brief description of the changes introduced by Energy Efficient Ethernet is provided in Section III. In Section IV, the implications of Energy Efficient Ethernet on Powerlink are analyzed and different alternatives are presented to make an effective use of the new standard while preserving the real-time features of Powerlink. The proposed alternatives are illustrated in Section V using a few case studies, and, finally, some conclusions are presented in Section VI.

II. INDUSTRIAL ETHERNET: POWERLINK

In Powerlink, as in many Industrial Ethernet approaches, existing Ethernet physical layer and media access control (MAC) devices are used as illustrated in Fig. 1. This has the benefits of the low cost and wide availability of Ethernet devices and components. By using these components, Powerlink can be fully implemented using existing hardware and some additional software. Additionally, as Ethernet is widely used in local area networks (LANs), the interconnection of Powerlink and computer LANs is straightforward.

To ensure real-time delivery of the frames, transmissions are coordinated in such a way that only one station can transmit at any given time, therefore avoiding collisions and subsequent retransmissions. This is done by defining a Managing Node (MN) and a group of Controlled Nodes (CN). The MN grants permission to transmit frames to the CN that can only transmit when they are given permission.

The transmission is organized in cycles, as shown in Fig. 2. The first part of the cycle, called isochronous phase, is used for real-time transmissions that are coordinated by the MN. The MN sends requests to the different CNs, which send a response back to the MN. The second part of the cycle, called asynchronous phase, is used for nonreal-time data and is also coordinated by the MN. Requests to transmit on this part of the cycle can be made in the responses to the MN during the first part of the cycle, or in the second part upon invitation from the MN. The MN keeps track of all pending requests and assigns transmission times for the different CNs that have requested it on the second part of the cycle. The cycle ends with an idle phase during which no data transmission takes place.

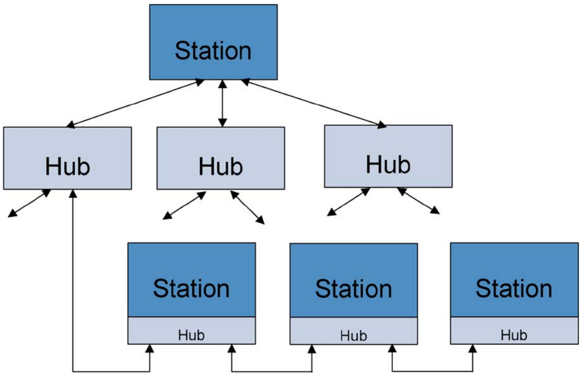

Fig. 3. Powerlink network topology.

During the isochronous phase, the request frames from the MN are sent only to each CN, while the CN responses are sent back to all the nodes. During the asynchronous phase, the responses can be sent to one node or to all the nodes.

To ensure that frames are delivered in real time, the use of hubs is preferred to switches, as the latter typically operate using "store and forward," thus introducing additional delay [8]. Each station may incorporate a hub as illustrated in Fig. 3.

III. ENERGY EFFICIENT ETHERNET

The main idea behind the Energy Efficient Ethernet efforts is to put the physical layer in a low-power mode when no data is being transmitted. This has been shown to result in large potential energy savings, as links are normally lightly loaded [2]. A number of alternatives can be used to implement the low-power modes. The most obvious one is to reduce the speed of the link when there is little traffic [2]. This can be achieved through autonegotiation [8], which is already part of the IEEE 802.3 Ethernet standards. Autonegotiation is currently used when the link is established to determine the link speed to be used, which is normally the highest supported by both ends. However, it could also be used by one of the ends to reduce the link speed by restarting the autonegotiation process and advertising only the lower speed to the other end. This will cause the link to operate at that lower speed. However, that process takes from a few hundred milliseconds to a few seconds, which is excessive for many applications. Therefore, speed changes would have an impact for the user. Other alternatives, like Rapid physical layer device (PHY) Selection (RPS) [14], have been proposed to accelerate the speed change. In RPS, a frame exchange is used to agree on a speed change with no need of restarting the autonegotiation process, and, therefore, the decision on the speed change can take place in a much shorter time. However, there are two major drawbacks to the speed-change mechanism. The first one is that the link has to be reestablished at the new speed, which takes a few hundred milliseconds for some of the standards. That time is needed to adjust all the elements in the receivers, like, for example, equalizers, cancellers, timing circuits, etc., to the channel conditions. Those elements are different from one speed to the other, and, therefore, a change in speed requires a new adjustment. During that time, the link is down, and no traffic can be exchanged. This means that the total time needed for a speed change is still in the hundreds of milliseconds. The second drawback is that although reducing the speed would decrease the energy consumption, we would still have transmitters and receivers operating continuously, only at a lower speed. In summary, we are mitigating the problem with minimal changes to the existing standards.

Another alternative is to introduce changes in the standards for the different speeds, so that they support low-power modes that can put a device to sleep and wake it up very quickly (on the order of microseconds) without a speed change. This is the option chosen by the IEEE 802.3az Task Force, which has analyzed the mechanisms to support low-power modes for the different Ethernet speeds, like, for example, 100 Mb/s, 1 Gb/s, and 10 Gb/s. In this case, the elements in the receiver are frozen when the device enters the low-power mode. To wake it up, only minor changes are expected to be needed, as the channel is quite stable. Those updates in the receiver can be performed in a few microseconds compared with the milliseconds that are needed to establish a link. To ensure that the receiver elements are aligned with the channel while in the low-power mode, short periods of activity are periodically scheduled to refresh the receiver state. As an example, the proposed state transitions in the IEEE 802.3az draft [15] are illustrated in Fig. 4. The different timer parameters are also specified in the draft for 100Base-TX, 1000Base-T, and 10GBase-T in terms of their minimum and maximum allowed values. For example, wake-up times on the orders of a few microseconds are supported in those cases, which are much lower that those achieved by a speed change. In addition, for 100 Mb/s and for 10 Gb/s, the low-power mode defined in the standard draft can be selected for only one of the link directions. This is useful when traffic is asymmetric, as the link direction that actually has traffic can be active, while the other direction is in low-power mode. This feature will be useful in the case of Industrial Ethernet, as it will be explained later.

The savings achieved by implementing the low-power modes are substantial. For example, in [16], the savings for a 100Base-TX device were shown to be over 75% of the total power when the PHY is in the low-power mode compared with the active state.

The energy consumption of an Ethernet link is currently roughly proportional to the time that the link is established. For example, for a link in which one of its ends is a PC, it would normally correspond to the time that the PC is active. In an industrial application, the link would be established when the system is in operation. In many cases, industrial applications operate continuously (or for long periods of time) leading to significant energy consumption. The other factor is link speed, since consumption is larger for higher speeds. For example, consumption for a 1-Gb/s link can be on the range of 1 to 2 W, while for 10 Gb/s, the range is one order of magnitude larger. More details are shown on Table I. The figures presented are given as an indication of the consumption of PHYs for the different speeds, as these figures will vary for different manufacturers and device technology. A link is composed of two PHYs, and, therefore, the energy consumption is twice that of a PHY. The consumption of the MAC and other Ethernet elements is not considered in Table I.

TABLE I
POWER CONSUMPTION FOR PHYS AND LINKS

	PHY	Link
100-Base-TX	250mW	500mW
1000-BaseT	700mW	1.4 Watts
10G-BaseT	6 Watts	12 Watts

Currently, Powerlink uses half-duplex 100Base-TX, and, therefore, the consumption of the link would be around 500 mW. When other speeds like 1 or 10 Gb/s are used, the consumption will increase significantly. Another key factor is that for 100Base-TX, the Energy Efficient Ethernet standard allows, for each link, direction to enter the low-power mode independently, a feature that will be very useful to obtain energy savings.

Once Energy Efficient Ethernet is adopted, the energy consumption of an Ethernet link will depend on whether the link is in a low-power mode or not. For a normal Ethernet link, the exact time a link stays in low-power mode will depend on the arrival times of the data frames and the algorithm used to enter low-power mode. For many Industrial Ethernet systems in general, and for Powerlink in particular, the situation is different, since the scheduling of the transmissions is determined by the cycle structure, as discussed in the previous section.

IV. IMPLICATIONS FOR INDUSTRIAL ETHERNET

When Energy Efficient Ethernet is adopted, the different pieces of Ethernet equipment will probably manage their decisions on when to enter in a low-power mode. Given the structure of the Powerlink cycle, some devices may enter a low-power mode during the idle phase of the cycle. Assuming, as an example, that all nodes enter a low-power mode, then, at the beginning of the next cycle, the links would have to be awake. This would introduce a delay. For example, in 100Base-TX, it takes tens of microseconds to wake up a link. The additional delay may cause a slot-time exceeded error in Powerlink, as the MN has a time out to get a response. The problem is that this issue may occur at the beginning of each new cycle, as the Ethernet device is unaware of the Powerlink constraints. A number of alternatives can be used to avoid this potential problem. The first is to add this delay to the timer values so that there is no time out. This obviously reduces the available time for data transfers, as each slot can have a longer duration. The second one is to configure all Ethernet devices with Energy

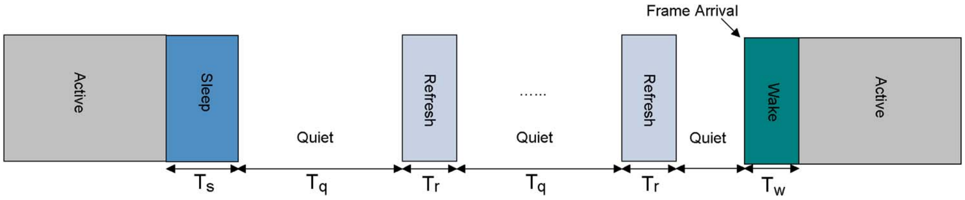

Fig. 4. Low-power modes in Energy Efficient Ethernet.

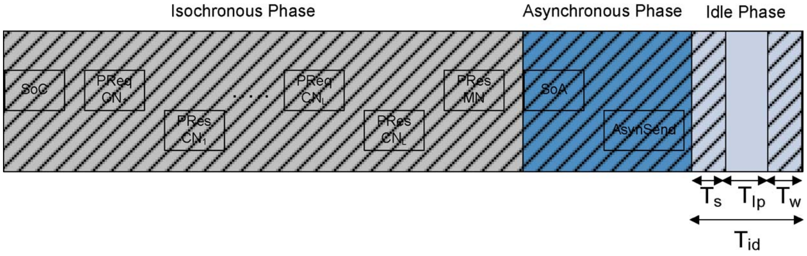

Fig. 5. Time in low-power mode (T_{lp}) for all nodes.

Efficient Ethernet disabled in such a way that low-power modes are never used. This is a drastic solution that ensures that Powerlink will operate normally but without any improvement in energy efficiency. The third option is to have Powerlink control when devices enter and exit low-power modes. This could combine energy savings with the real-time constraints needed for Powerlink operation. This latter option is analyzed in more detail in the following.

A. General Case Operation

Given the structure of the Powerlink cycle, a simple way to manage the links would be for all nodes to enter the low-power mode at the beginning of the idle phase and then wake up in the final part of that phase. In this way, the link is active just before the start of the new cycle. The savings that can be achieved with this strategy will largely depend on the duration of the idle phase with respect to the cycle time and also to the time needed to sleep and wake up a link. Let us define T_c as the total cycle time, T_{id} as the time of the idle phase of the cycle, T_s as the time needed to set a link in low-power mode, and T_w as the time to put it back into active mode. Then, the percentage of the time that the link is in low-power mode can be calculated as

$$\alpha_{\mathrm{all}} = \frac{T_{\mathrm{id}} - T_{\mathrm{s}} - T_{\mathrm{w}}}{T_{\mathrm{c}}}. \tag{1}$$

That means that the link is in low-power mode for the duration of the idle phase minus the time needed to sleep and wake up a link, as illustrated in Fig. 5. From (1), it can be seen that the savings will be larger if $T_{\mathrm{id}} \gg T_s + T_w$, which is true if large cycle times are used. The use of large cycle times has the drawback of increasing the worst case delay needed to transmit a frame. Therefore, the real-time requirements of the system being designed will dictate the maximum cycle time to be used. To minimize energy consumption, that maximum value should be utilized, since the use of lower values will increase energy consumption.

Additionally, as in some cases (100 Mb/s and 10 Gb/s), link directions can enter the low-power mode independently; the CN to MN direction can be put to sleep except for the time needed to send the responses or asynchronous traffic. This option will be explored in Section IV-E.

B. Operation of Nodes That Do Not Work in Asynchronous Mode

In the previous analysis, the duration of the idle phase has been assumed to be constant. In reality, that is not the case. The idle phase extends from the end of the asynchronous phase to the end of the cycle, and the duration of the asynchronous phase depends on whether there are requests to transmit on that phase or if the MN has to poll any CN that operates on asynchronous-mode only. The asynchronous phase ends when all those transmissions are completed. To maximize the duration of the idle phase, the frequency of the polling done on the CNs working only in asynchronous mode should be minimized. Again, the system requirements will dictate a minimum polling frequency, and that should be used in order to obtain larger energy savings.

For some of the CNs, it may be possible to be in low-power mode for longer periods of time. For example, CNs that are not involved in a given asynchronous phase can enter the low-power mode at the end of the isochronous phase, as shown in Fig. 6. In fact, a CN that communicates only with the MN can be in low-power mode from the time it sends its response to a query from the MN until the end of that cycle. When the CN is not involved in the asynchronous phase, the percentage of time in low-power mode is given by

$$\alpha_{\mathrm{na}} = \frac{T_a + T_{\mathrm{id}} - T_s - T_w}{T_c} \tag{2}$$

where T_{a} is the duration of the asynchronous phase.

C. Operation of Nodes That Only Work in Asynchronous Mode

For CNs that operate in the asynchronous-only mode, the time in low-power mode can extend from the end of the asynchronous phase of a given cycle to the beginning of the asynchronous phase on the next cycle, as illustrated in Fig. 7.

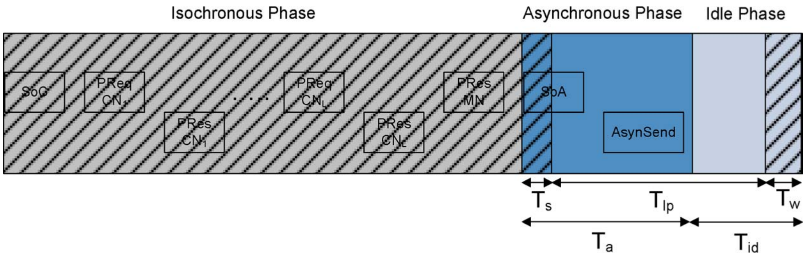

Fig. 6. Time in low-power mode (T_{lp}) for CNs that are not involved in the asynchronous phase.

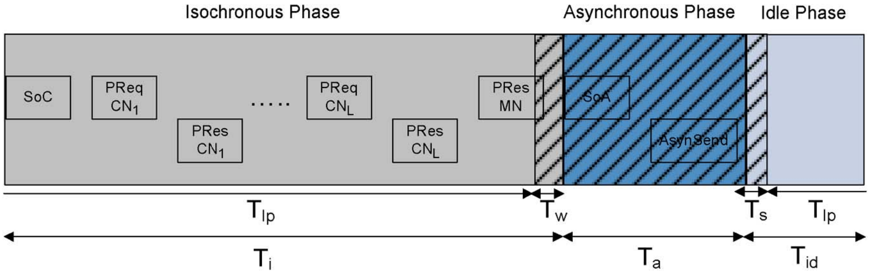

Fig. 7. Time in low-power mode (T_{lp}) for asynchronous only CNs.

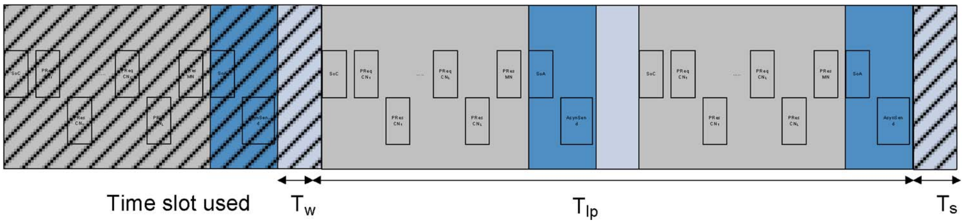

Fig. 8. Time in low-power mode (T_{lp}) for CNs that use multiplexed timeslots with $k = 3$.

In this case, the percentage of time in low-power mode is given by

$$\alpha_{\mathrm{ao}} = \frac{T_i + T_{\mathrm{id}} - T_s - T_w}{T_c} \tag{3}$$

where T_{i} is the duration of the isochronous phase.

D. Operation of Nodes That Use Multiplexed Timeslots

For CNs that use multiplexed timeslots (i.e., they receive requests from the MN one out of k cycles and not for every cycle), the time in low-power mode can extend up to $k - 1$ cycles, as illustrated in Fig. 8. That is possible when the CN does not have to receive data from other CNs during the asynchronous phase and is not involved in CN to CN data transmission during the isochronous phase. CN to CN traffic is known as cross traffic and is based on the fact that the CN to MN responses are sent to all CNs and not only to the MN. Whether a given CN is involved in cross traffic can be determined by its RxMapping configuration. If a given CN is involved in cross traffic, then it must be active on every cycle to receive data from other CNs. This will be illustrated in one of the case studies presented in Section V. For those CNs, it may be more beneficial to use the approach presented in the next section to manage link directions independently.

When the CN is not involved in cross traffic, the percentage of time in low-power mode for all CNs that operate with multiplexed timeslot is given by

$$\begin{aligned}\alpha_{\mathrm{all}}^{\mathrm{mt}} &= \frac{(k-1)\cdot T_c}{k\cdot T_c} + \frac{T_{\mathrm{id}} - T_s - T_w}{k\cdot T_c}\\ &= \frac{(k-1)}{k} + \frac{\alpha_{\mathrm{all}}}{k}.\end{aligned} \tag{4}$$

Additionally, when a CN is either asynchronous only or does not use the asynchronous mode, the percentages of time in low-power mode are obtained by using the α for that mode

$$\begin{aligned}\alpha_{\mathrm{ao}}^{\mathrm{mt}} &= \frac{(k-1)\cdot T_c}{k\cdot T_c} + \frac{T_i + T_{\mathrm{id}} - T_s - T_w}{k\cdot T_c}\\ &= \frac{(k-1)}{k} + \frac{\alpha_{\mathrm{ao}}}{k}\end{aligned} \tag{5}$$

$$\begin{aligned}\alpha_{\mathrm{na}}^{\mathrm{mt}} &= \frac{(k-1)\cdot T_c}{k\cdot T_c} + \frac{T_a + T_{\mathrm{id}} - T_s - T_w}{k\cdot T_c}\\ &= \frac{(k-1)}{k} + \frac{\alpha_{\mathrm{na}}}{k}.\end{aligned} \tag{6}$$

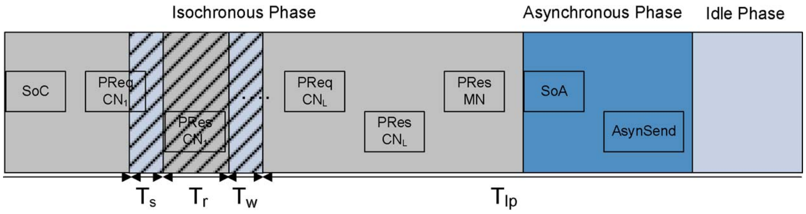

Fig. 9. Time in low-power mode in the CN to MN direction (T_{lp}) for CNs that only transmit during the isochronous phase.

E. Operation of Nodes by Managing Link Directions Independently

As mentioned before, for 100Base-TX and 10Gbase-T, it is possible for each link direction to enter the low-power mode independently. This feature can be very useful in Powerlink. For example, for a CN that uses multiplexed timeslots and is involved in cross traffic from other CNs on each cycle, it is possible to have the link slept for $k-1$ out of k cycles in the direction of the CN to the MN, while the MN to CN direction is active during all cycles. Therefore, the savings mentioned in Section IV-D [see (4)–(6)] can still be obtained in one of the link directions.

In a general case, the CN to MN direction can be in low-power mode except for the time needed to send the response to the MN or the asynchronous traffic originating in that CN. For a CN that is only involved in the isochronous phase, the time in low-power mode can be expressed as

$$\alpha_{\text{all}}^{\text{cn\_to\_mn}} = \frac{T_c - T_s - T_w - T_r}{T_c} \tag{7}$$

where T_r is the time needed to send the response to the MN as illustrated in Fig. 9. This means that in most cases, significant savings can be obtained in one direction. As the power consumption would be approximately the same for each link direction, the overall savings are still substantial. It is also worth mentioning that as in 1000Base-T, it is not possible to set the link directions in low-power mode independently; if the Powerlink systems are upgraded to use 1000Base-T, then the savings considered in this section would not be obtained. This means that in many cases, the energy efficiency would degrade by using 1000Base-T.

F. Energy Savings

Finally, the energy savings as a percentage of the energy consumption when Energy Efficient Ethernet is not used will be given, in all cases, by

$$\begin{aligned}\text{Savings } (\%) &= 1 - \frac{\alpha \cdot P_{\text{lp}} + (1-\alpha)\cdot P_n}{P_n} \\ &= \alpha \cdot \left(1 - \frac{P_{\text{lp}}}{P_n}\right)\end{aligned} \tag{8}$$

where P_{lp} is the power consumption in low-power mode, P_n is the power consumption in normal mode, and α is the fraction of time in low-power operation (which is different depending on each of the previous scenarios).

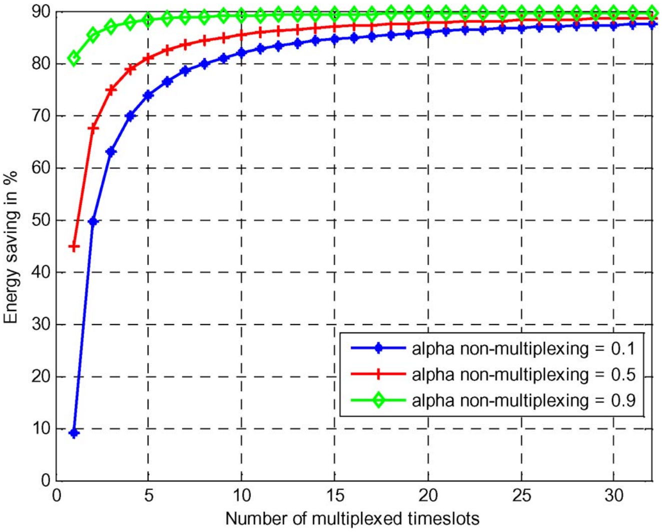

Fig. 10. Percentage of energy savings: multiplexed versus nonmultiplexed scenarios.

From the previous discussion, it becomes clear that CNs should use multiplexed timeslots whenever possible, as that enables larger energy savings (at least, in the CN to MN direction). The percentage of savings with respect to the non-multiplexed scenario can be seen in Fig. 10. Significant savings can be achieved with low values of k (number of multiplexed timeslots), particularly when the use of the low-power mode in the nonmultiplexed approach is reduced. If, for example, the nonmultiplexed case is operating in low-power mode $(k=1)$ during 10% of the time $(\alpha = 0.1)$, it can be seen that the energy savings is slightly less than 10% (with respect to the case of active power mode). However, if the system is multiplexed to $k = 10$, it can be seen that the savings are boosted over 80%.

On the contrary, if the low-power mode is active during 90% of the time $(\alpha = 0.9)$ in the nonmultiplexed configuration, the initial savings would already be over 80%. Thus, the effect of multiplexing in this case would be less relevant.

The overall strategy should then be as follows: First, select the largest possible cycle time that will be dictated by the CN with the harder real-time constraints and, then, for all other CNs, select multiplexed slot operation with the maximum possible k.

In the next section, the proposed approach is applied to a few typical Powerlink configurations showing the potential benefits.

Another implication of Energy Efficient Ethernet on Powerlink, and, more generally, on Industrial Ethernet, is that hubs will be less energy efficient since they broadcast the frames on all ports, and, therefore, the links will spend less time in low-power mode. This will reinforce the trend toward the use of switches and may mean that hubs will be seldom used in the future. This may lead to a lack of availability of commercial hubs and force the use of switches in Powerlink systems. Therefore, the use of switches in Powerlink needs to be studied to analyze how the real-time constraints can be met by appropriately choosing and configuring these devices. The use of switched Ethernet for factory communications has been previously studied [17]–[20], and a number of alternatives have been proposed to bound the delay and ensure that the tight requirements of industrial applications are met. Those works can be utilized as the starting point for the analysis of the use of switches in Powerlink systems.

G. Application to Other Industrial Ethernet Technologies

The proposed ideas to achieve energy savings in Powerlink can be, in principle, extended to other Industrial Ethernet technologies that are based on the same concept of synchronous transmission coordinated by a CN. This is, for example, the case of some of the real-time modes of operation of PROFINET [21]. In those cases, the cycle structure should be carefully studied to determine on which portions of the cycle the link (or at least one link direction) can be set in low-power mode.

When transmissions are not organized in synchronized cycles, as in the nonreal-time mode of operation of PROFINET, the ideas presented are not applicable. In those cases, if the constraints allow for some extra microseconds of delay, the link can be in a low-power mode until a frame arrives. Upon the frame arrival, the link is set in active mode, the frame is transmitted, and the link is set back into the low-power mode. This may be tolerated when the real-time constraints are not very stringent. If that is not the case, then Energy Efficient Ethernet should be disabled, and the link should always be in the active state to avoid additional delays in frame transmission.

V. CASE STUDIES

Powerlink is used in a wide range of applications, from systems that require high-performance motion to others that perform data acquisition, and it is therefore difficult to asses the potential benefits of the proposed approach. A few configurations are analyzed in this section to illustrate the benefits in different environments that are typical in systems that use Powerlink.

The first example is a five-axis motion system with five CNs and a cycle time of 400 μs. This can correspond to a computer numerical-controlled machine with three linear axes and two rotational axes. In this case, there are no CNs operating on multiplexed timeslots and no asynchronous-only CNs. Therefore, the savings can be obtained using the idle periods and possibly, the asynchronous phase. As the duration of the idle and asynchronous phases is a fraction of the cycle time, and the time needed to enter and exit the low-power mode is large relative to the cycle time, the percentage of time in low-power mode will be small. However, savings can still be obtained in the CN to MN direction, as the CN only needs to be active to send its response. Considering a response time of 20 μs and a time to sleep and wake up a link of 130 μs (as per the standard draft [15] for 100Base-TX), the time in low-power mode would be 62.5%, as shown in Table II. If, for example, the 100Base-TX Ethernet is used, the energy savings in low-power mode can be assumed to be 75% from the results in [16]. Then, the total energy savings can be computed using (8), and, assuming that 50% of the energy is consumed in each link direction, savings of over 23% would be obtained.

TABLE II
ENERGY SAVINGS FOR THE CASE STUDIES

	Time in Low Power Mode (Uplink)	Time in Low Power Mode (Downlink)	Energy Savings
Simple motion system	62%	0%	23%
Complex motion system	87%	0%	32%
Large I/O system	93%	93%	70%

Therefore, in this system, it seems reasonable to manage the link directions independently. In the MN to CN direction, Energy Efficient Ethernet would not be used. To ensure that no device enters a low-power mode and interferes with Powerlink operation by adding additional delay (and possibly causing timing errors), Energy Efficient Ethernet should be disabled in all the devices for the MN to CN direction. In the CN to MN direction, the link should enter the active mode only to send its response, as shown in Fig. 9.

The second example is a coupled-axes motion system integrated by 64 CNs, where the CNs operate on a multiplexed timeslot with $k = 8$. This could be a simplified model of a printing machine. The cycle of the system is 400 μs. In this case, it seems that the CNs can be in low-power mode seven cycles out of eight, and using (4), the percentage of time in low-power mode will be around 87%. This translates to large energy savings. However, the coupled-axes motion system normally requires that the multiplexed CNs listen to the responses of the master axis on every cycle. Therefore, the link should be active in all cycles in the direction of MN to CN. As discussed in Section IV-E, when the link allows each link direction to enter the low-power mode independently, then the CN to MN direction can be put to sleep for seven cycles out of eight, and large savings are achieved in one of the directions. Assuming again that 100Base-TX is used, then the energy savings can be obtained using (8), with a result of around 65% for the CN to MN direction. Assuming that 50% of the power is consumed in each direction, the overall savings would be over 32%. This shows that the potential benefits are large when most CNs use multiplexed timeslots. In this case, all Ethernet devices should be controlled by the Powerlink software. The software should set the devices in low-power mode in the CN to MN direction after their active cycle and then set them back for normal operation after k cycles, just before their next active cycle.

The final example is a large I/O system with 150 CNs and a delay requirement of 16 ms. In this case, the CNs communicate directly with the MN for data exchange. The system can be configured with a cycle time of 1 ms and the CNs multiplexed with $k = 16$. Then, assuming that 100Base-TX Ethernet is used as in the previous example, the devices will be in low-power mode more that 93% of the time, and the energy savings will reach 70%. Again, the Powerlink software should control the Ethernet devices to put them in the low power and back into normal operation. In this case, it is more beneficial to use a shorter cycle time and multiplexed timeslots. This is a special situation, as all nodes have the same delay requirements (which are large). In this case, the best option is to take advantage of that large delay to use multiplexed timeslots with large values of k. The results for the three scenarios can be seen in Table II, where the times in low-power mode are given separately for each link direction. Uplink corresponds to transmission from the CN and downlink to reception at the CN.

As a summary, the proposed approach can result in large energy savings for systems in which most of the nodes use multiplexed timeslots and do not listen to cross traffic (case study 3). When nodes are involved in cross traffic (case study 2) or do not use multiplexed timeslots (case study 1), significant savings can still be obtained by managing the link directions independently.

The modifications to Powerlink to add the control on when the devices enter and exit the low-power mode should be simple, as the timing information of when the changes must occur is already available, and the configuration of the Ethernet devices is trivial.

VI. CONCLUSION

In this paper, the implications of the Energy Efficient Ethernet standard on Industrial Ethernet systems have been analyzed using Powerlink as a case study. A number of alternatives have been proposed to enable the use of Energy Efficient Ethernet in Powerlink systems. The proposed alternatives ensure that the real-time requirements are met while maximizing energy savings. To illustrate the potential benefits, a number of case studies have been analyzed showing that in some cases, savings of 70% can be obtained.

Future work will explore the implementation of the proposed techniques once devices that implement the Energy Efficient Ethernet standard are available. In addition, the implications for other flavors of Industrial Ethernet will be considered based on the analysis and results presented in this paper.

REFERENCES

[1] M. Gupta and S. Singh, “Greening of the Internet,” in *Proc. ACM SIGCOMM*, Aug. 2003, pp. 19–26.

[2] C. Gunaratne, K. Christensen, B. Nordman, and S. Suen, “Reducing the energy consumption of Ethernet with Adaptive Link Rate (ALR),” *IEEE Trans. Comput.*, vol. 57, no. 4, pp. 448–461, Apr. 2008.

[3] IEEE P802.3az Energy Efficient Ethernet Task Force, [Online]. Available: http://grouper.ieee.org/groups/802/3/az/public/index.html

[4] E. Tovar and F. Vasques, “Real-time field-bus communications using Profibus networks,” *IEEE Trans. Ind. Electron.*, vol. 46, no. 6, pp. 1241–1251, Dec. 1999.

[5] G. Cena and A. Valenzano, “An improved CAN field-bus for industrial applications,” *IEEE Trans. Ind. Electron.*, vol. 44, no. 4, pp. 553–564, Aug. 1997.

[6] L. Almeida, E. Tovar, J. A. G. Fonseca, and F. Vasques, “Schedulability analysis of real-time traffic in WorldFIP networks: An integrated approach,” *IEEE Trans. Ind. Electron.*, vol. 49, no. 5, pp. 1165–1174, Oct. 2002.

[7] J. W. Walewski, V. Kulkarni, and G. Steindl, “EEE for real-time industrial Ethernet,” presented at the IEEE 802.3, Meeting, Vancouver, BC, Canada, Mar. 2009.

[8] C. E. Spurgeon, *Ethernet: The Definitive Guide*. Sebastopol, CA: O’Reilly, Feb. 2000.

[9] J.-D. Decotignie, “Ethernet-based real-time and industrial communications,” *Proc. IEEE*, vol. 93, no. 6, pp. 1102–1117, Jun. 2005.

[10] M. Felser, “Real-time Ethernet—Industry prospective,” *Proc. IEEE*, vol. 93, no. 6, pp. 1118–1129, Jun. 2005.

[11] K. C. Lee, S. Lee, and M. H. Lee, “Worst case communication delay of real-time industrial switched Ethernet with multiple levels,” *IEEE Trans. Ind. Electron.*, vol. 53, no. 5, pp. 1669–1676, Oct. 2006.

[12] T. Skeie, S. Johannessen, and O. Holmeide, “Timeliness of real-time IP communication in switched industrial Ethernet networks,” *IEEE Trans. Ind. Informat.*, vol. 2, no. 1, pp. 25–39, Feb. 2006.

[13] Communication Profile Specification, Draft Standard 301, Ethernet Powerlink Standardisation Group, 2008.

[14] K. Christensen, “Rapid PHY selection (RPS): Emulation and experiments using PAUSE,” presented at the IEEE 802.3 Meeting, Orlando, FL, Mar. 2007.

[15] Energy Efficient Ethernet Standard Draft D1.4, IEEE P802.3az, May 2009.

[16] J. Chou, A. Kuo, D. Tseng, J. Lin, and K. Huang, “Proposal of low-power idle: 100Base-TX,” presented at the IEEE 802.3 Meeting, Portland, OR, Jan. 2008.

[17] X. Fan, Z. Wang, and Y. Sun, “How to guarantee factory communication with switched Ethernet: Survey of its emerging technology,” in *Proc. 28th IECON*, 2002, vol. 3, pp. 2525–2530.

[18] E. Jaspemeite and P. Neumann, “Switched Ethernet for factory communication,” in *Proc. Int. Conf. Emerging Technol. Factory Autom.*, 2001, vol. 1, pp. 205–212.

[19] L. Lo Bello, L. Kaczynski, and O. Mirabella, “Improving the real-time behavior of Ethernet networks using traffic smoothing,” *IEEE Trans. Ind. Inform.*, vol. 1, no. 3, pp. 151–161, Aug. 2005.

[20] P. Ferrari, A. Flammini, D. Marioli, and A. Taroni, “A distributed instrument for performance analysis of real-time Ethernet networks,” *IEEE Trans. Ind. Inform.*, vol. 4, no. 1, pp. 16–25, Feb. 2008.

[21] J. Feld, “PROFINET—Scalable factory communication for all applications,” in *Proc. IEEE Int. Workshop Factory Commun. Syst.*, 2004, pp. 33–38.

Juan Antonio Maestro (M’07) received the M.Sc. degree in physics and the Ph.D. degree in computer science from Universidad Complutense de Madrid, Madrid, Spain, in 1994 and 1999, respectively.

He worked with several multinational companies, managing projects as a Project Management Professional and organizing support departments. He has served both as a Lecturer and a Researcher with several universities such as the following: Universidad Complutense de Madrid; the Universidad Nacional de Educación a Distancia (Open University), Madrid; Saint Louis University, Madrid; and the Universidad Antonio de Nebrija, Madrid, where he currently manages the Computer Architecture and Technology Group. His current activities are oriented to the space field, with several projects on reliability and radiation protection, as well as collaborations with the European Space Agency. He is the author of numerous technical publications, both in journals and international conferences. His areas of interest include high level synthesis and cosynthesis, energy-efficient networks and real-time systems, fault tolerance, and reliability.

Pedro Reviriego (A’03–M’04) received the M.Sc. and Ph.D. degrees (with honors) in telecommunications engineering from the Technical University of Madrid, Madrid, Spain, in 1994 and 1997, respectively.

From 1997 to 2000, he was an R&D Engineer with Teldat, Madrid, working on router implementation. In 2000, he was with Massana to work on the development of 1000Base-T transceivers. During 2003, he was a Visiting Professor with University Carlos III, Leganés, Madrid. From 2004 to 2007, he was a Distinguished Member of Technical Staff with LSI Corporation working on the development of Ethernet transceivers. He is currently with the Universidad Antonio de Nebrija, Madrid. His research interests are soft errors and reliability, fault-tolerant systems, and energy-efficient communication networks. He has authored numerous papers in international conferences and journals. He has also participated in the IEEE 802.3 standardization for 10GBase-T.