Authors' Accepted Manuscript (AAM): this is the accepted version of the paper. Citation information: DOI 10.1109/TASE.2014.2313735, IEEE Transactions on Automation Science and Engineering, vol. 12, no. 1, pp. 228-237, Jan. 2015

Energy Efficient Ethernet for Real-Time Industrial Networks

S. Vitturi, *Member, IEEE,* and F. Tramarin, *Member, IEEE*

[redacted]

Abstract—**To increase the energy efficiency of Ethernet networks, in 2010 the IEEE published the IEEE 802.3az amendment, known as Energy Efficient Ethernet (EEE). The amendment introduces a new operational mode, defined as Low Power Idle (LPI), that allows to considerably reduce the power consumption of inactive Ethernet links. In this paper we address the application of EEE to Real-Time Ethernet (RTE) networks, the popular communication systems typically employed in factory automation, characterized by tight timing requirements. We start with a description of the EEE basics and, subsequently, focus on the introduction of EEE in the industrial communication scenario. Then we specifically address the implementation of effective EEE strategies for some popular RTE networks. The analysis is carried out on configurations commonly deployed al low levels of factory automation systems. The obtained results show that considerable power savings can be achieved with very limited impact on network performance.**

Note to Practitioners —**This paper deals with the introduction of EEE in the industrial communication scenario, with specific attention to Real-Time Ethernet networks. The work was motivated by the continuous need for efforts to increase energy efficiency at all levels of factory automation systems. Particularly, the IEEE 802.3az amendment represents a significant opportunity even if, unfortunately, it does not explicitly refer to industrial networks. The assessment carried out in this paper, besides confirming the possibility of a profitable adoption of EEE by RTE networks, raised up several aspects of interest for practitioners. Indeed, the cyclic traffic typical of industrial networks reveals suitable to support EEE strategies. Nonetheless, even in case a strict network synchronization is maintained, additional delays may possibly affect frame delivery due to the time overhead introduced by the LPI mode. The adoption of suitable strategies may, however, mitigate the impact of such delays. Also, it has been shown that acyclic traffic requests may reduce the effectiveness of EEE strategies, since they could introduce frequent undesired changes of the Ethernet links status. Finally, we would like to stress the importance of carrying out extensive measurement sessions on real testbeds, since outcomes from practical experiments definitely represent the best way to validate the theoretically designed EEE strategies.**

Index Terms—**Energy Efficient Ethernet, Real-time Ethernet Networks, Industrial Automation**

I. INTRODUCTION

In 2010, the Institute of Electrical and Electronics Engineers published the IEEE 802.3az amendment [1] to the widespread Ethernet standard, commonly referred to as Energy Efficient Ethernet (EEE). The amendment introduces an *optional* Ethernet feature that allows to achieve power saving during periods of low network traffic. Basically, EEE allows an Ethernet link between two partners (either stations or network components like, for example, switch devices) to enter a new state in which link power consumption is considerably reduced with respect to that of its common idle state. The request to enter such a state may be issued by any of the two link partners.

The introduction of EEE was solicited by the observation that, even if practically all of the Ethernet links currently deployed operate with very low duty cycles, unfortunately their power consumption remains almost the same for all the time they are switched on. In other words, the power consumption of a link depends only slightly on the actual data traffic on the link itself [2]. The power consumption, instead, strongly depends on the transmission rate and ranges, in general, from some hundreds of milliWatt for 100 Mbps links to some Watt for 10 Gbps links [3], [4]. Considering that the assessment carried out in both [2] and [5] revealed an average utilization of Ethernet links below 5%, the expected overall power saving that could be achieved reducing the power consumption of idle links is potentially enormous.

The way in which power saving can be achieved has been the matter of a long debate in which two main options were considered. The first one, referred to as Low Power Idle (LPI) mode, allows to force Ethernet links in a specific state characterized by low power consumption. The second option is represented by a reduction of the link transmission rate, which in turn limits power consumption. The final choice of the IEEE 802.3 committee has been in the direction of the LPI mode. Consequently, most manufacturers are currently implementing LPI on their products, so that EEE is rapidly becoming widely adopted.

The EEE amendment, however, does not provide any indication on the power saving strategies to be implemented, leaving their definition be driven by the specific application scenarios. Such an aspect is a key issue currently addressed by the scientific community. Particularly, one of the most challenging problems is represented by the choice of the time instants to enter/exit the low power state to achieve the greatest power saving while ensuring low packet transmission delays.

Industrial versions of Ethernet networks [6] are playing an ever more important role, thanks to their intrinsic valuable features that allow to obtain isochronous real-time communications [7], [8]. These networks, usually referred to as Real-Time Ethernet (RTE), are envisaged to have a considerable growth in the next years [9]. Moreover, their transmission rate, which is currently 100 Mbps for almost all the available protocols, will be increased up to 1 Gbps or even 10 Gbps

Stefano Vitturi and Federico Tramarin are with the National Research Council of Italy, CNR–IEIIT, Padova and also with the Department of Information Engineering, University of Padova, Via Gradenigo 6/B, I–35131, Padova, Italy. E-mails: {vitturi, tramarin}@dei.unipd.it.

in the future. Thus, very likely EEE concepts will be of strong interest for RTE networks as well. In this field of application, however, the time overheads introduced by the LPI mode may reveal critical, since industrial communication applications often require a tight timing. On the other hand, it has to be considered that industrial traffic is to a considerable extent predictable, as it is often generated by cyclic operations [10]. Such a knowledge reveals of great help to design timely and effective EEE strategies.

In this paper we hence focus on the adoption of EEE by RTE networks. After an assessment of the LPI mode that includes a thorough timing analysis, we will concentrate on the design of effective EEE strategies for some of the most popular RTE networks. In detail, the paper is organized as follow. Section II discusses some noticeable contributions appeared in the scientific literature concerned with EEE. Section III provides the LPI basics. Section IV analyses the general application issues concerned with the introduction of EEE in the industrial communication scenario. Sections V, VI and VII are concerned with the implementation of EEE strategies for some of the most popular RTE networks, namely, PROFINET IO, Ethernet POWERLINK, EtherCAT and EtherNet/IP, that are adopted by a wide range of applications. Section VIII, based on the proposed examples, provides some useful guidelines that should drive the adoption of EEE by RTE networks. Finally, Section IX concludes the paper.

II. Related Work

The adoption of EEE by RTE networks is an emerging research topic only recently addressed by the scientific community. To the best of our knowledge, there are only few contributions explicitly concerned with this subject. In [11], the authors analyze the application of appropriate EEE strategies to Ethernet POWERLINK, a popular RTE network. The results, obtained from theoretical analyses, are encouraging, since the achievable power saving demonstrated to be considerable. Focusing on the same RTE network, paper [12] provides a simulation study that investigates further EEE strategies aimed at harmonizing power savings and timing requirements. Another contribution, although not officially published, is given in [13]. In this case, the authors provide a feasibility study about the possibility of adapting the EEE philosophy to RTE networks, proposing to incorporate them in the IEEE 802.3az amendment. Unfortunately, the proposal has not been further discussed, even if it represents a valuable attempt to address the topic from a standardization point of view.

EEE aspects have been extensively analyzed for general purpose networks. In this section we limit to briefly mention some of the contributions we believe significant in the direction of their possible application to the industrial communication scenario. In [5] the authors introduce Adaptive Link Rate (ALR), a technique to achieve power saving mentioned in Section I based on the rapid change of the link transmission rate. Although ALR is no longer considered, the paper presents some policies that could be profitably used in the context of EEE. Indeed, some of them are recalled in [14] where the authors focus on sleeping algorithms as a way to provide high power saving, while ensuring the unavoidable delays that influence frame transmission are kept very low. Also, the authors propose an analytical model for the transmission of data bursts that demonstrates how power savings can be greatly increased if links are activated only for the transmission of these large amounts of data. Similarly, in [15] a technique called packet coalescing is proposed. Here the authors concentrate on the transmission of large TCP frames and show the benefits, in terms of energy efficiency, of grouping acknowledgment frames before their actual transmission, instead of sending them one by one immediately. Another interesting contribution is given in [16]. In this case the authors provide an analytical model to evaluate the time spent by a link in the various states defined by EEE. Such a time is the key metric necessary to calculate power consumption as well as to design traffic shaping strategies to increase power saving. The model, developed under the assumption that frames arrive as batches with Poisson distribution, has been validated using real traffic traces.

As a further issue, power savings provided by EEE could reveal of interest for some new Ethernet amendment proposals, currently under discussion within the IEEE 802.3 working group [17].

Finally, it is worth stressing that energy efficient communication systems are expected to play an important role in “green manufacturing” [18], a driving concept for the design and implementation of innovative manufacturing systems, where priority is given to aspects like both reduction of the environmental impact and use of natural resources. In this context, for example, both papers [19] and [20] address energy saving in automotive production systems, where Ethernet networks are envisaged to be ever more deployed. Similar topics are also addressed in different contexts, like domotic and building automation applications [21].

III. Basics of Low Power Idle

EEE is available for several Ethernet physical layers (PHYs), included those of interest for RTE networks, namely, 100BASE-TX, 1000BASE-T and 10GBASE-T that will be considered in this paper.

EEE is based on the introduction of the LPI mode, a new feature of Ethernet devices that allows Ethernet links to enter a new functional state, namely *quiet state* in which power consumption is considerably reduced and communication is not allowed (only minimal signaling on the medium is maintained). In this direction, the medium access control (MAC) protocol has been modified with the introduction of the “LPI client” stack, as shown in Fig. 1. Such a new MAC architecture makes the `LP_IDLE` protocol service available to the upper layers, that can exploit it to implement EEE strategies. Indeed, the transition to/from quiet state takes place via both *request* and *indication* primitives of such a protocol service, as described in Fig. 2.

When one of the partners of an Ethernet link (e.g. Station A in Fig. 2) decides to force the link in quiet state, it invokes the `LP_IDLE.request` primitive with the (unique)

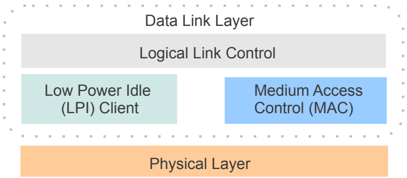

Fig. 1. Protocol Architecture of a Station Implementing EEE

parameter `LPI_REQUEST` set to ASSERT. Then the PHY starts transitioning to quiet state, The PHY of the destination partner (Station B) detects the new state of the link listening to the transmission medium and, consequently, issues the `LP_IDLE.indication` primitive to its LPI client which, in turn, informs the higher layers about the link state change. The same procedure is followed to wake up the link but, in this case, the `LPI_REQUEST` parameter has to be set to DE-ASSERT

It is worth stressing that the procedure of changing the link state is negotiated exclusively through electrical signaling on the link itself, and no packet transmission is involved. When one of the partners issues the `LPI_REQUEST` primitive, its specific signaling indicates the request to the other link partner, that has to implicitly agree.

IEEE 802.3az provides also some important information concerning the timing of the LPI mode operation. In particular, the time necessary to an active link to move to quiet state is defined as *sleep time*, T_s. The time spent in quiet state is application dependent, but it is required a refresh signal is issued with fixed period T_q. The duration of the refresh procedure is T_r. Finally, stepping back from quiet to active state requires a time defined as *wake time*, T_w.

An example of LPI operation behavior is shown in Fig. 3. As can be seen, the example refers to a link that initially switches from active to quiet state upon the request of one of its partners. The link remains in such a state for a time T_q. Then, the refresh procedure takes place. Subsequently, after T'_q, a new state change is requested to the active state, which is reached after T_w.

The aforementioned timing parameters are of noticeable importance for practical applications, since they represent the basis to determine the possible delays in frame delivering that could derive from the adoption of EEE. The values of these times, as specified by IEEE 802.3az, are shown in Table I for the PHYs of industrial interest. It is worth observing that IEEE 802.3az does not provide directly the values of T_w but, rather, it refers to two additional parameters, namely T_{w_phy} and T_{w_sys}. The former is the time necessary for the PHY to switch from quiet to active, whereas T_{w_sys}, referred to as *system wake time* provides a more comprehensive information, since it accounts for the maximum time that may elapse between the instant in which a link starts to wake up, and the instant in which data transmission can actually take place. Consequently, the values shown in Table I are relevant to T_{w_sys}, as it represents an upper bound for T_w.

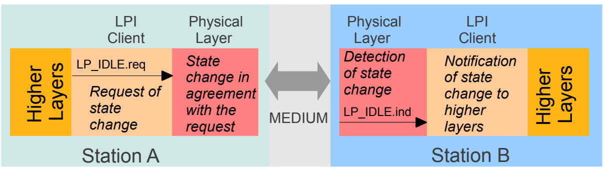

Fig. 2. LPI Client Protocol Primitives

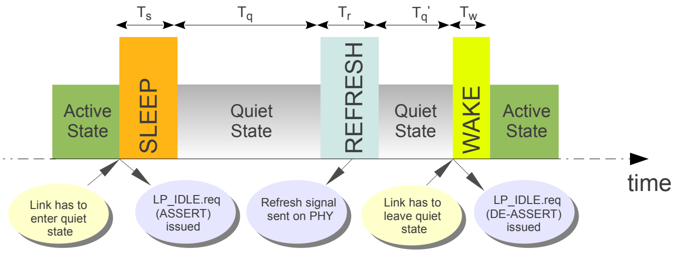

Fig. 3. Example of LPI Operation

TABLE I
EEE Times [µs] for the considered Physical Layers

PHY	T_s^{min}	T_s^{max}	T_q^{min}	T_q^{max}	T_r^{min}	T_r^{max}	T_w^{max}
100BASE-TX	200	220	20000	22000	200	220	30
1000BASE-T	182	202	20000	24000	198	218.2	16.5
10GBASE-T	2.88	3.2	39.68	39.68	1.28	1.28	7.36

Before analyzing the application of EEE to RTE networks, there are some significant features of IEEE 802.3az that are worth mentioning since they have a considerable importance in this context. In particular, PHYs have different characteristics, deriving from their inner architectural aspects, that may have a noticeable impact on the design of EEE strategies.

In detail, the permanence in quiet state may be interrupted at any instant, even during the transmission of the refresh signals. Furthermore, but only for both the 100BASE-TX and 1000BASE-T PHYs, the transition to quiet state may be interrupted by the arrival of a frame transmission request at any of the link partners, during time T_s. In this case the link immediately starts to switch to active state. Conversely, aborting the transition to quiet state is not possible for 10GBASE-T.

Also, one of the most interesting features introduced by IEEE 802.3az is the possibility available for both the 100BASE-TX and 10GBASE-T PHYs to independently enter the LPI mode for the two link directions (transmit/receive). In other words, a link partner may be, contemporaneously, in quiet state for transmission and in active state for reception. Such a feature is not available for the 1000BASE-T PHY.

Table II reports the power consumption of Ethernet links. The values have been derived from technical specification documents of two important actors in the field of Ethernet devices [3], [4]. In Table II, the column “LPI” is relevant to the case in which both link directions are in quiet state, whereas the column “LPI-partial” refers to the case in which only one link direction is in quiet state. Other interesting power consumption data for commercially available devices, obtained from practical measurements, are reported in [22] and [23].

IV. Application of EEE to RTE Networks

The typical real-time industrial traffic may be subdivided in two categories, namely cyclic and acyclic [24]. The former is

TABLE II
Link Power Consumption for PHYs of Industrial Interest

PHY layer	Active [mW]	LPI [mW]	LPI-partial [mW]
100BASE-TX	351	58	188
1000BASE-T	697	53	–
10GBASE-T	2600	390	1430

generated by operations like set-point transmission, periodic sampling of sensor data and, in general, repetitive actions. On the other hand, acyclic traffic derives from unpredictable events such as, for example, those associated to process alarms. The timing requirements imposed by both these types of traffic have variable severity depending on the application contexts. For example, focusing on the data exchange between controllers and sensors/actuators at the lowest levels of factory automation systems, the more demanding applications (e.g. motion control systems [25]), may impose that both the period of cyclic data transmissions and the maximum allowed latency of alarm messages are in the order of some hundreds of microseconds. Moreover, the tolerated jitter on cyclic operations may be below 1%. Conversely, moving to different fields of application such as remote monitoring, control of conventional plants and building automation, the timing constraints are much more relaxed.

As an example of EEE application, we may refer to a specific link of a RTE network that connects two stations. Clearly, the greatest energy saving is achieved maintaining the link in quiet state for as much time as possible. This implies the quiet state is entered just after a frame transmission has ended, and exited immediately before a new transmission is going to take place. Such a strategy, however, has to keep into consideration the overhead introduced by the LPI timing. Indeed, it may happen that the time T_s necessary to reach the quiet state is comparable with the most critical periods of cyclic operations allowed on that link. To this regard it is worth observing that, for both the 100BASE-TX and 1000BASE-T PHYs, the values of the sleep time are very close to the typical sampling times of motion control loops [26]. Thus, adopting the LPI mode in this situation reveals not effective, since the resulting power saving would be very limited or, even worse, the cycle times of the network might become longer than expected. In such cases, the choice of the 10GBASE-T PHY could appear as a valuable option, since T_s is considerably reduced (the maximum value is 3.2 μs) but, unfortunately, its power consumption is still very high. Consequently, it is not a priori clear whether such an alternative represents a convenient choice or not. A more accurate analysis to this regard can only be carried out on a per case basis.

Furthermore, it has to be considered that two stations may be connected via a path that comprises more than one link. For this type of configurations, the procedure of entering/exiting the quiet state necessarily involves more (likely all) links of the path. Thus, trivially, a frame that has to be transmitted over a path comprising m links may incur in a maximum *activation delay* equal to mT_w, with the consequent negative impact on network performance. This problem, however, as preliminarily addressed in [27], may be satisfactorily mitigated activating the links in advance, before the cyclic operation is started (for example by means of purposely designed wake-up frames), even if this procedure may reduce power savings.

The described scenario becomes more complex if acyclic traffic is considered, in particular the traffic generated by process alarms. These events trigger the transmission of specific frames whose delivery has to be completed within (usually tight) predefined deadlines. Consequently, the links comprised in the path between two stations, are allowed to enter the quiet state only if the sum of the path activation delay plus the time to transmit alarm frames is less than the deadline. Otherwise that deadline will likely be missed. Hence, supposing D is the alarm deadline and T_f the time needed to transmit an alarm frame, then it has to be

$$D \leq mT_w + T_f \tag{1}$$

It is worth observing that eq. (1) represents, in practice, a necessary condition for the application of EEE strategies to the path.

Alarms could even occur while either a link is transitioning to quiet state (i.e. during time T_s) or the refresh signal is being issued. In both cases, however, the latency added to alarms transmission is rather limited. Indeed, for both 100BASE-TX and 1000BASE-T PHYs, the transition to the quiet state may be arbitrarily interrupted. IEEE 802.3az does not specify the time necessary to return to the active state but, reasonably, it will not take more than T_w. Concerning the 10GBASE-T PHY, since the transition to quiet state can not be interrupted, the maximum time to wake up a link will be $T_s + T_w$. At this rate, however, both T_s and T_w are very short. Finally, in case the transition request arrives during a refresh procedure, all PHYs allow to immediately stop the procedure and reactivate the link. Hence, the maximum delay introduced is T_w.

From an analytical point of view, for a link between two stations where cyclic operations take place with period T_c, the maximum time spent in quiet state within a period, T_{lpi}^{max}, is given by

$$T_{lpi}^{max} = T_c - T_s - T_p - T_w - kT_r \tag{2}$$

where T_p is the time necessary to carry out cyclic operations and k is the number of refresh signals issued in the period. Equation (2) reflects the situation in which the link is activated for the time strictly necessary to perform cyclic operations. If an alarm occurs while the link is in quiet state, it is necessary to awake the link, send the alarm frame and then put the link back in quiet state. The time, T_a, required by this sequence is

$$T_a = T_w + T_f + T_s \tag{3}$$

Hence, supposing n alarms are gathered in a period and that each of them arrives while the link is in quiet state (that represents the worst case), then the time spent in quiet state by that link within a period reduces to

$$T_{lpi} = T_{lpi}^{max} - nT_a \tag{4}$$

(the equation holds for $T_{lpi}^{max} \geq nT_a$).

If the path between the two stations comprises more than one link, Eq. (4) applies to each link of the path.

It may be observed that the occurrence of acyclic traffic reduces the time spent by the links in quiet state, and hence the achievable power savings. Nonetheless, such a traffic does not influence directly the design of EEE strategies, even if it can prevent their adoption, as a consequence of the alarm deadlines that have to be matched.

In the following of the paper we propose examples of EEE application to the most widespread RTE networks. Based on the considerations made in this Section, the examples exclusively consider cyclic traffic since, as we have seen, this is the type of traffic that actually drives the design of EEE strategies.

V. EEE FOR PROFINET IO

PROFINET IO is a widespread RTE network encompassed by the IEC 61784-2 International Standard [28], available in two versions, namely PROFINET IO IRT (Isochronous Real-Time), referred to as communication profile CP 3-6 by the standard, and PROFINET IO RT (Real-Time) that covers the communication profiles CP 3-4 and CP 3-5. Two main types of devices are employed by both versions of PROFINET IO. They are, namely, IO controllers (IOCs, typically intelligent devices like PLCs or PCs implementing automation tasks) and IO devices (IODs, field devices like, for example, sensors and/or actuators). The basic difference between the two versions of PROFINET IO relies on synchronization aspects. Specifically, in PROFINET IO IRT the network behavior is based on a cycle constituted by four subsequent phases [29] with all stations strictly synchronized. PROFINET IO IRT cycle time values range from 31.25 µs to 4 ms. The first two phases, indicated respectively as red and orange, are reserved to isochronous real-time traffic. They are sometimes considered as a whole IRT phase. The third phase (green) is used for all other types of traffic (either real-time or not). Finally, the yellow phase is an idle period. Network synchronization is ensured by a distributed clock system through the periodic delivery of clock frames.

Conversely, in PROFINET IO RT there is no synchronization among stations and the concept of network cycle is no longer valid. Nonetheless, the use of frame priorities still allows to timely transmit real-time data.

The data transfer in both versions of PROFINET IO takes place exclusively via “communication relationships” (CRs), established between IO Controllers and IO Devices at start-up time. At least two different CRs have to be defined per each IO Device. The first one, reserved to cyclic traffic, is based on a producer-consumer technique, whereas the second one, reserved for acyclic traffic, is based on a client-server paradigm. The periods of data traffic on the cyclic CR for PROFINET IO IRT can be set as multiples of the network cycle, whereas for the RT version the range 1÷128 ms is available (in this case, periods have to be powers of two).

Both versions of PROFINET IO are, in principle, suitable for profitable applications of EEE concepts. However, for PROFINET IO IRT, particular care is needed by clock frames delivery. Indeed, these frames shall be transmitted in the green phase by the “master clock” device (usually an IOC), with a period of 30 ms. Alternatively, several commercial products may transmit clock information in the red phase at every cycle [30]. In both cases, however, when a clock frame is being sent, all links have to be active to ensure such an information is timely received by all the nodes.

As a final observation, PROFINET IO IRT implementations based on either 100BASE-TX or 1000BASE-T, will very unlikely exploit refresh signals. Indeed, if clock frames are sent at every PROFINET IO IRT cycle, then all links are necessarily activated with a period well below T_q^{max}. Conversely, in case clock frames are sent every 30 ms, refresh signals will be necessary only for those links activated with periods greater than T_q^{max}.

A. Examples of EEE application to PROFINET IO

As an example of EEE application to PROFINET IO IRT, we consider a network composed of an IO Controller and four IO Devices referred to as, respectively IOD1, IOD2, IOD3 and IOD4. All stations are connected through a switch, as shown in Fig. 4, in a single level configuration, as typically deployed at low levels of factory automation systems [31]. IOD1 is supposed to exchange data with the IOC in the red phase, whereas the other IODs communicate during the green phase.

The network has five links necessary to connect all stations with the switch. The cycle time is set to 1 ms, and the durations of the different phases within the cycle are, respectively, 200 µs (red/orange), 700 µs (green) and 100 µs (yellow). For convenience, we consider only data transfer in the direction from IOC to IODs on the cyclic CRs. The periods assumed for such transmissions are shown in Table III. Moreover, we suppose that only minimum size Ethernet frames are used, since industrial traffic is typically characterized by small payloads.

Since the path between any two nodes comprises two links, the tight synchronization granted by PROFINET IO IRT reveals extremely helpful in setting effective EEE strategies. Specifically, it is possible to activate in advance the links involved in data transmission.

Fig. 5 shows an example of EEE strategy for the considered network relevant to a cycle in which the IOC has to send data to all the IODs.

The cycle starts with both links #1 and #2 in active state (they were contemporaneously awaken at the end of the previous cycle by, respectively, IOC and IOD1). At the beginning of the red/orange phase, the IOC transmits cyclic data to IOD1

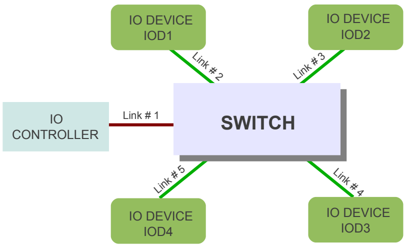

Fig. 4. Example of PROFINET IO IRT Configuration

TABLE III
Periods of the data Exchange from IOC to IODs

IO Device	Period
IOD1	1 ms
IOD2	1 ms
IOD3	5 ms
IOD4	10 ms

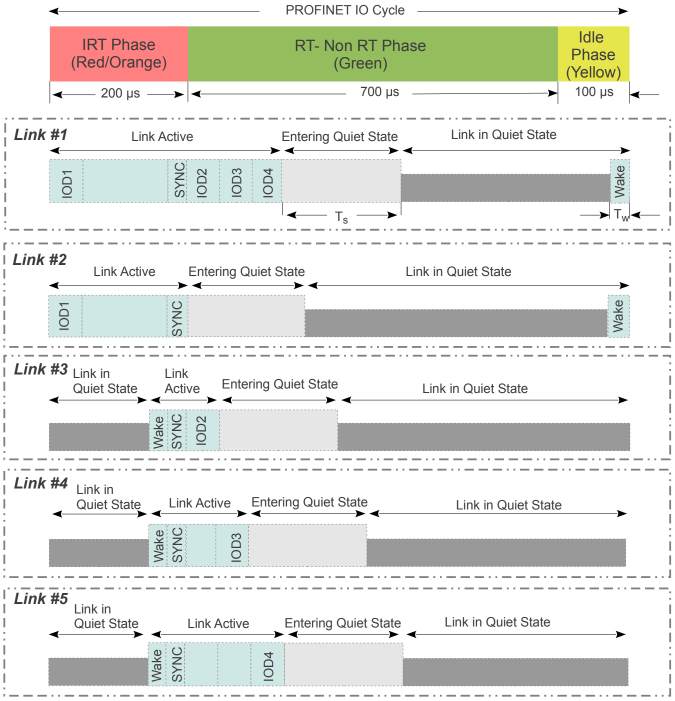

Fig. 5. Example of PROFINET IO IRT EEE Strategy

(as indicated by the notice "IOD1" in both links #1 and #2 of Fig. 5) and, just before the phase expiration, it broadcasts a clock frame (denoted as "SYNC" in Fig. 5). Clearly, all other links had to be activated in advance to ensure the clock frame is timely received by the IODs. Subsequently, link #1 has to be maintained active only for the time necessary to allow the IOC to send data to the other IODs (in the green phase), whereas Link #2 can be immediately deactivated, since in the remaining portion of the cycle there is no traffic scheduled. Links #3, #4 and #5, activated at the end of the red/orange phase to receive the clock frame, remain in active state for the time strictly necessary to carry out data transmission towards their IODs. It is worth observing that both links #4 and #5, after receiving the clock frame remain in the active state although they will not carry any traffic, since the time interval between the arrival of the clock frame and the reception of data by the IODs they connect is not sufficiently long to allow the links to enter the quiet state.

Fig. 6 reports the time spent by each link in quiet state (T_{lpi}), along with the network average power consumption calculated for the case in which EEE is not adopted and, respectively, for that in which EEE is adopted. The time in quiet state is expressed as a percentage of the cycle time. All the values shown in Fig. 6 have been calculated for 100BASE-TX, using the LPI timing parameters reported in Table I and the link power consumptions specified in Table II. It is worth pointing out that in practical applications, links are expected to be activated with adequate safety margins in such a way they are surely in the active state when data transfer takes place, slightly reducing the time spent in quiet state and hence the power savings. Nonetheless, Fig. 6 confirms the adoption of EEE may lead to significant power savings without impacting on network performance.

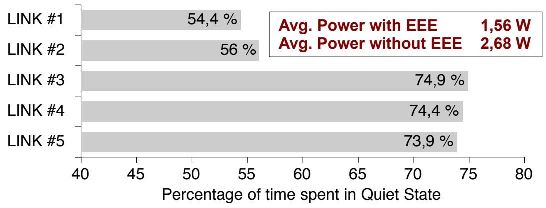

Fig. 6. PROFINET IO IRT: Time Spent in Quiet State and Power Consumption

Moving to PROFINET IO RT, we refer to the more complex example shown in Fig. 7, since this network is often deployed in large configurations [32]. As can be seen, the network comprises one IOC and 14 IODs arranged in a multi-level, mixed (line/star) topology [31]. Three switches are centers of the stars that connect some lines of IODs, whereas the IOC is connected to the switches via a backbone. Lines of IODs are implemented using devices that incorporate a two-port switch. The configuration of Fig. 7 resembles that of industrial plants where a single controller handles some production islands each represented, in this case, by the IODs connected to a single switch [33].

Following the general EEE strategy of activating a link only for the time necessary to grant the actual data transfer, if in Fig. 7 we consider for example the transmission from the IOC to IOD #1A3, then links #1, #1A1, #1A2 and #1A3 have to be activated. However, while link #1 may be activated in advance (since the IOC knows when the transmission is scheduled), the other links may only be activated by the partners upon receiving the data to be forwarded (specifically, switch 1 activates link #1A1, IOD #1A1 activates link #1A2 and, finally, IOD #1A2 activates link #1A3). The whole sequence introduces an activation delay equal to $3T_w$. To achieve the best energy efficiency, link #1A3 will start stepping back towards quiet state immediately after data transmission to IOD #1A3 is concluded. The same action will be undertaken in sequence by the other links of the line connected to switch 1.

A numerical example could provide some further insights. We assume, as for the previous case, that only transmissions in the direction from IOC to IODs are considered. The periods are supposed to be, respectively, 2 ms for the IODs connected to switch 1, 4 ms for the IODs connected to switch 2 and 8 ms for the IODs connected to switch 3.

Fig. 8 reports the time T_{lpi} spent in quiet state by links #1, #2 and #3, expressed as a percentage of their periods, along with the average power consumption of the whole network. T_{lpi} has been calculated under the assumption that the IOC experiences the contemporaneous occurrence of data

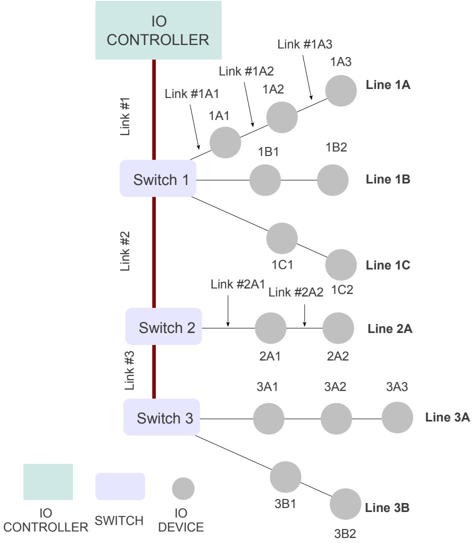

Fig. 7. Example of PROFINET IO RT Configuration

transmission to all IODs. In this condition, link #1 has to be kept active not only for the time necessary to send data to the IODs connected to switch 1, but also for the time necessary to transmit towards both switches 2 and 3. Similarly, link #2 has to be maintained active also for the time necessary to transmit towards switch 3.

The results reported in Fig. 8 clearly highlight the benefits deriving by the adoption of EEE. For the sake of clarity, however, it is worth remarking that, in this case, network performance may be worsened by link activation delays. Nonetheless, as outlined in Section IV, a different EEE strategy could be devised such that all the links are activated in advance, minimizing the impact of the activation delays.

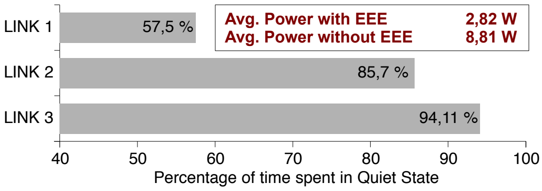

Fig. 8. PROFINET IO RT: Time Spent in Quiet State and Power Consumption

VI. EEE FOR ETHERNET POWERLINK

Ethernet POWERLINK (EPL) [34] is a RTE network included in the IEC 61784-2 International Standard, Communication Profile 13-1. Two types of devices have been defined for EPL, namely, Managing Node (MN, the device that implements automation tasks) and Controlled Nodes (CNs, passive devices). EPL configurations comprise one MN and up to 239 CNs, arranged in several topologies. The physical layer of EPL is 100BASE-TX, but thanks to the total compliance with legacy Ethernet, both 1000BASE-T and 10GBASE-T can be adopted seamlessly. At the data link layer, EPL uses a purposely developed protocol (EPL data link layer) placed onto the Ethernet MAC protocol whereas, at the application layer, it adopts a protocol based on the well known CANopen profiles. The EPL data link layer makes use of a Time Division Multiple Access (TDMA) technique to ensure all stations are able to access the network in an ordered way [35]. In practice, network operation is based on a cycle executed with period T_{epl}. At the beginning of each cycle, the MN broadcasts a frame named Start of Cycle (SoC) that synchronizes all CNs. Then the Isochronous Period is entered. In this time interval, the MN polls the CNs. Such an operation is carried out via two specific frames, namely Poll Request (PReq) and Poll Response (PRes), exchanged with each CN in reserved time slots. The frame PReq, that carries the output data, is issued by the MN towards a CN, whereas PRes, that represents the answer of the addressed CN, carries the input data. It is important to remark that, while the transmission of PReq is unicast, PRes frames are broadcast, so that the data they carry are made available to all CNs. Time slots can be either continuous or multiplexed. Continuous means that a CN associated with that slot is polled at each cycle, whereas multiplexed slots are assigned to CNs polled with a period integer multiple of a cycle. After the last CN has been polled, the MN issues the Start of Acyclic (SOA) frame that triggers the beginning of the Acyclic Period. In this phase an asynchronous message may be transmitted by a CN that made a request during one of the previous isochronous periods. Subsequently, the Idle Period is entered. In this phase, there is no activity on the network. As an example, Fig. 9 shows the description of two subsequent EPL cycles for a network with two controlled nodes, namely CN #n and CN #m, polled in continuous slots, and three other multiplexed CNs (respectively, #i, #j and #k), polled every two cycles.

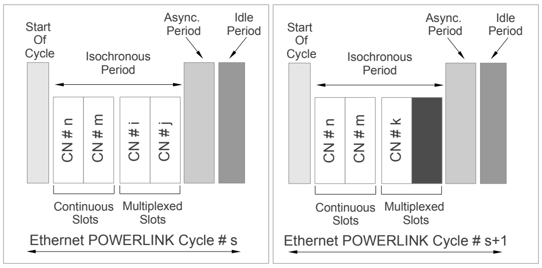

Fig. 9. Example of POWERLINK Cycle with Multiplexed Time Slots

A. *Examples of EEE application to Ethernet POWERLINK*

The tight synchronization of EPL networks allows the adoption of effective EEE strategies, as addressed in [11].

The network configurations we took into consideration are those already analyzed for PROFINET IO, shown in both Fig. 4 and Fig. 7. However, we would like to point out that our analysis is not meant in any way as a comparison of EEE capabilities of the two networks. Rather, the analysis is proposed to provide (hopefully) some useful insights about the application of EEE strategies to different RTE networks. With refer to both Fig. 4 and Fig. 7, in the context of Ethernet POWERLINK, the IO Controller has to be replaced by the Managing Node, whereas, IO Devices are substituted by Controlled Nodes. Also, since only one integer multiple of the EPL cycle can be used for the multiplexed slots, we had to limit the periods reported in Table III to two values, as summarized in Table IV, that reports the actual periods used for EPL. Moreover for the configuration of Fig. 7, in Table IV we refer to switches instead of CNs, meaning that the period of a specified switch is that with which all CNs connected to that switch are polled.

TABLE IV
Periods for the Ethernet POWERLINK Simulations

Config. of Fig. 4	Period [ms]	Config. of Fig. 7	Period [ms]
CN 1	1	Switch 1	2
CN 2	1	Switch 2	4
CN 3	5	Switch 3	4
CN 4	5	/	/

For the configuration of Fig. 4, T_{epl}, is set to 1 ms. The cycle is structured with two continuous slots and one multiplexed slot. The continuous slots are assigned to both CN 1 and CN 2, whereas the multiplexed slot accommodates alternatively CN 3 and CN 4 in such a way to access these CNs with a period equal to $5T_{epl}$. Similarly, for the configuration of Fig. 7, we set T_{epl} = 2 ms. The cycle is configured with 7 continuous slots and 4 multiplexed slots. All the CNs belonging to switch 1 will be assigned to the continuous slots, whereas the multiplexed slots will be used in one cycle by the CNs of both lines 2A and 3B and, in the subsequent cycle, by the CNs of line 3A.

The values selected in the two examples for the duration of the EPL cycle time, T_{epl}, ensure the timely execution of the actions to be carried out within a cycle. Indeed, assuming the absence of acyclic traffic (this means that the Asynchronous Period lasts only for the time necessary to transmit the SoA frame), the minimum EPL cycle time, defined as the minimum time necessary to perform all the operations of a cycle results, respectively, 119.38 µs for the configuration of Fig. 4 and 491.84 µs for that of Fig. 7.

Fig. 10 shows the outcomes for both the considered configurations. It reports the time spent in quiet state as a percentage of the polling periods along with the average power consumption of the whole network, calculated for the 100BASE-TX PHY. The analysis has been carried under the following hypothesis for EPL:

(*i*) link directions (transmit/receive) are independent;
(*ii*) all links are kept in quiet state between the beginning of the Asynchronous Period and the end of the Idle Period;
(*iii*) PRes frames are received by all CNs. This requires that the receive links of the CNs are active in the Isochronous Period.

The results reported in Fig. 10 indicate that significant power savings can be obtained also by EPL. Nevertheless, it is worth noticing that the reduction of the average power when EEE is adopted, for the configuration of Fig. 7, is less than that obtained for PROFINET IO RT. The difference is mainly due to the consumption of the CNs receive links that are maintained active to ensure the reception of PRes frames. This feature, however, if not necessary, can be disabled as explicitly mentioned by the EPL standard.

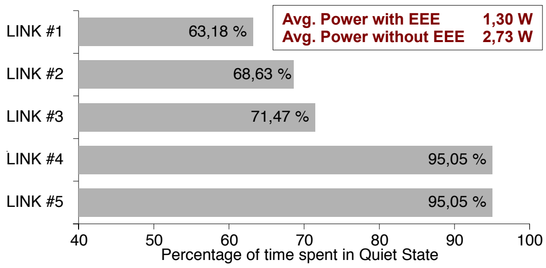

(a) Configuration of Fig. 4

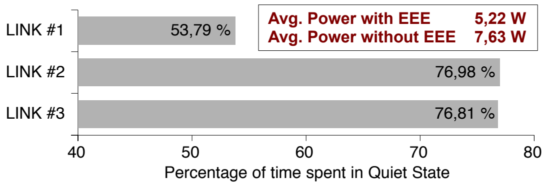

(b) Configuration of Fig. 7

Fig. 10. Ethernet POWERLINK: Time Spent in Quiet State and Power Consumption

VII. EEE for other RTE Networks

The concepts discussed so far, along with the application examples, can be extended to other RTE networks. Unfortunately, due to space limitations, a complete assessment that covers all the available networks, results impossible. Nonetheless, we would like to briefly provide some considerations on two further networks, namely EtherCAT and EtherNet/IP, since they are widespread industrial communication systems characterized by very different operating principles.

A. EtherCAT

EtherCAT is a network included in the IEC 61784-2 International Standard, communication profiles CP12-1 and CP12-2. The behavior of EtherCAT is based on a master-slave relationship. A frame periodically originated from the master circulates among all slaves and is returned to the master itself, resembling a logical ring architecture. The frame "crosses" the slaves, instead of being actually received and retransmitted. Each slave crossed by the frame, inserts input data and, contemporaneously, extracts output data. Such a type of elaboration, usually referred to as "on the fly", allows to dramatically limit frame transmission times, reflecting in the possibility of providing very short cycle times with low jitter [36].

The described principle of operation strongly suggests to avoid entering the quiet state while an EtherCAT frame is being transmitted. Indeed, since EEE times are comparable with the EtherCAT frame transmission times, or even greater (at least for both 100BASE-TX and 1000BASE-T), transitions to/from the quiet state would intolerably increase the EtherCAT frame transmission time.

Nevertheless, as an effective EEE strategy all links could enter the quiet state in the interval comprised between the end of an EtherCAT frame transmission and the beginning of the following one. This strategy requires that slave devices implement the LPI client stack. Such an issue, however, needs to be adequately addressed, since the physical layer of

EtherCAT slaves is not fully compliant with that of common Ethernet devices.

B. EtherNet/IP

EtherNet/IP is a RTE network that represents the communication profile CP2-2 in the IEC 61784-2 International Standard. At the higher layers, Ethernet/IP makes use of the well known Common Industrial Protocol (CIP) [37] that, in turn, relies on both TCP and UDP. The underlying data link layer is compliant with IEEE 802.3, ensuring full compatibility with legacy Ethernet. Two versions of EtherNet/IP are available. In the first one, stations are not synchronized among each other, whereas the second version ensures precise synchronization by means of a protocol called CIP sync that relies on the IEEE 1588 distributed clock system [38].

A straightforward comparison with the analysis carried out in Section V suggests the adoption of EEE by EtherNet/IP can be made in a way similar to PROFINET IO. In particular, as a general EEE strategy, the links comprised in a path between two devices have to be kept in quiet state when there is no data traffic on the path itself. The occurrence of a transmission request (either cyclic or acyclic) triggers the link activation procedure.

VIII. EEE Guidelines for RTE Networks

The examples provided in the previous Sections showed that EEE may be profitably applied to several RTE networks, using strategies that are strictly related to the characteristics of the considered applications, such as traffic profiles and performance requirements. Unfortunately, a general rule that starting from the description of a specific application allows to select which networks are suitable to support EEE strategies and what power savings are expected, can not be analytically derived. However, some useful guidelines can be provided to this regard.

First of all, effective EEE strategies can be implemented only for those links in which the shortest period of cyclic operations is much greater than the EEE time overheads. Since nowadays 100BASE-TX is by far the most deployed physical layer for RTE networks, the EEE times shown in Table I suggest that periods have to be in the order of some hundreds of microseconds. Also, both EEE strategies and achievable power savings are strictly related to network topologies and configurations. For example, if in Fig. 7 the periods of the IODs connected to switch 1 and, respectively, switch 3 are reversed, then both links #2 and #3 will have to remain in the active state for longer times, reducing the achievable power savings. As a consequence, it is advisable EEE strategies are taken into consideration at the very early stage of network planning, since they heavily impact on network structural aspects.

Secondly, the adoption of EEE by RTE networks that use devices equipped with dedicated hardware, not fully compliant with IEEE 802.3 (e.g. EtherCAT and Sercos III, a further network encompassed by IEC 61784-2), has to be carefully addressed, since the LPI client stack might be not available for the physical layer of these devices.

Finally, acyclic traffic may reduce the achievable power savings. Thus, the actual occurrence of such a type of traffic needs to be carefully considered in the design phase. Particularly, alarm deadlines should be determined in advance since such an information could help to mitigate their negative effects. For example, the precise knowledge of the deadlines could allow, in some cases, to group alarms before their actual transmission, instead of immediately sending them at the time of their occurrence, as described in [14]. In such a way, the oscillations between quiet and active states caused by frequent alarm transmissions could be significantly limited.

IX. Conclusions

In this paper we provided an assessment of the issues concerned with the application of EEE to industrial networks. The analysis started with a description of the IEEE 802.3az amendment. Then we actually addressed RTE networks and investigated how they could adopt effective EEE strategies. Specifically, we focused on some of the most popular RTE networks and provided practical examples of EEE application for typical industrial configurations. The analysis demonstrated that significant power savings can be achieved, actually maintaining the high performance level these networks are requested to provide. Consequently, we believe the introduction of EEE in the industrial communication scenario is a promising research topic.

Several future activities can be envisaged in this context. Firstly, since most industrial devices (particularly field components like sensors and actuators) are not equipped with built-in EEE capabilities, there is the need to start implementing the IEEE 802.3az protocol stack on these types of devices. Also, the actual behavior of such components has to be assessed, focusing in particular on the interface between the protocols adopted at the application layer and IEEE 802.3az.

Moreover, suitable EEE strategies have to be designed. As we have seen, such an activity is driven, on the one hand, by the characteristics of the (predictable) cyclic traffic, while on the other hand it has to take into account the negative impact deriving from the possible presence of acyclic data transmission requests.

Finally, it is worth observing that results from real applications of EEE are ever more necessary. Indeed, especially in the industrial communication field, even minimal differences from (theoretically) estimated behaviors could have dangerous effects on the behavior of plants/applications that use the networks. For example, if the activation delay of a path reveals, in practice, greater than the theoretically expected value, then the jitter of cyclic operations carried out on the path could increase. Consequently, the designed EEE strategies need to be validated by practical experiments. In this way the inevitable differences between theoretical models and real behavior may be evidenced, making possible an accurate tuning of the models themselves.

References

[1] *IEEE 802.3az-2010: Energy Efficient Ethernet, Amendment to IEEE standard 802.3-2008 (CSMA/CD) specifications*, IEEE Std., 2010.

[2] K. J. Christensen, C. Gunaratne, B. Nordman, and A. D. George, “The Next Frontier for Communications Networks: Power Management,” *Computer Communications*, vol. 27, no. 18, pp. 1758 – 1770, 2004.

[3] Intel 82579 Gigabit Ethernet PHY Datasheet v2.1. [Online]. Available: http://www.intel.com

[4] 10GBASE-T Ecosystem Is Ready for Broad Adoption. [Online]. Available: http://www.commscope.com

[5] C. Gunaratne, K. Christensen, B. Nordman, and S. Suen, “Reducing the Energy Consumption of Ethernet with Adaptive Link Rate (ALR),” *Computers, IEEE Transactions on*, vol. 57, no. 4, pp. 448 –461, April 2008.

[6] “Industrial Communication Systems,” in *The Industrial Information Technology Handbook*, R. Zurawski, Ed. CRC, 2005, pp. 37.1–47.16.

[7] J. D. Decotignie, “Ethernet–Based Real–time and Industrial Communications,” *Proc. IEEE*, vol. 93, no. 6, pp. 1102–1117, June 2005.

[8] M. Felser, “Real–time Ethernet–Industry prospective,” *Proceedings of the IEEE*, vol. 93, no. 6, pp. 1118–1129, June 2005.

[9] Industrial Ethernet Book, issue 69/72, April-October 2012. [Online]. Available: http://www.iebmedia.com

[10] T. Sauter, “The Three Generations of Field-Level Networks–Evolution and Compatibility Issues,” *Industrial Electronics, IEEE Transactions on*, vol. 57, no. 11, pp. 3585–3595, November 2010.

[11] J. A. Maestro and P. Reviriego, “Energy Efficiency in Industrial Ethernet: The Case of Powerlink,” *Industrial Electronics, IEEE Transactions on*, vol. 57, no. 8, pp. 2896–2903, August 2010.

[12] F. Tramarin and S. Vitturi, “Energy Efficient Ethernet for the industrial communication scenario,” in *Emerging Technologies Factory Automation (ETFA), 2013 IEEE 18th Conference on*, 2013, pp. 1–8.

[13] J. W. Walewski, V. Kulkarni, and G. Steindl. EEE for Real-time Industrial Ethernet. Presented at the IEEE802.3 Meeting, Vancouver, Canada, 2009. [Online]. Available: http://www.ieee802.org/3/az/public/mar09/walewski_01_0309.pdf

[14] S. Herreria-Alonso, M. Rodriguez-Perez, M. Fernandez-Veiga, and C. Lopez-Garcia, “Optimal Configuration of Energy-Efficient Ethernet,” *Computer Networks*, vol. 56, no. 10, pp. 2456 – 2467, 2012.

[15] K. Christensen, P. Reviriego, B. Nordman, M. Bennett, M. Mostowfi, and J. Maestro, “IEEE 802.3az: the Road to Energy Efficient Ethernet,” *Communications Magazine, IEEE*, vol. 48, no. 11, pp. 50–56, Nov. 2010.

[16] M. Ajmone Marsan, A. Anta, V. Mancuso, B. Rengarajan, P. Vasallo, and G. Rizzo, “A Simple Analytical Model for Energy Efficient Ethernet,” *Communications Letters, IEEE*, vol. 15, no. 7, pp. 773–775, July 2011.

[17] L. Winkel. Distinguished Minimum Latency Traffic in a Converged Traffic Environment, DMLT. [Online]. Available: http://www.ieee802.org/3/DMLT/public/Jul13/Winkel_00_0713_DMLT_SG_Tutorial_v04.pdf

[18] J. Li, J. R. Morrison, M. T. Zhang, M. Nakano, S. Biller, and B. Lennartson, “Editorial: Automation in Green Manufacturing,” *Automation Science and Engineering, IEEE Transactions on*, vol. 10, no. 1, pp. 1–4, January 2013.

[19] D. Meike, M. Pellicciari, and G. Berselli, “Energy Efficient Use of Multirobot Production Lines in the Automotive Industry: Detailed System Modeling and Optimization,” *Automation Science and Engineering, IEEE Transactions on, to appear*.

[20] Q. Chang, G. Xiao, S. Biller, and L. Li, “Energy Saving Opportunity Analysis of Automotive Serial Production Systems (March 2012),” *Automation Science and Engineering, IEEE Transactions on*, vol. 10, no. 2, pp. 334–342, 2013.

[21] B. Sun, P. Luh, Q.-S. Jia, Z. Jiang, F. Wang, and C. Song, “Building Energy Management: Integrated Control of Active and Passive Heating, Cooling, Lighting, Shading, and Ventilation Systems,” *Automation Science and Engineering, IEEE Transactions on*, vol. 10, no. 3, pp. 588–602, 2013.

[22] P. Reviriego, K. Christensen, J. Rabanillo, and J. Maestro, “An Initial Evaluation of Energy Efficient Ethernet,” *Communications Letters, IEEE*, vol. 15, no. 5, pp. 578–560, May 2011.

[23] P. Reviriego, V. Sivaraman, Z. Zaho, J. Maestro, A. Vishawanath, A. Sanchez-Macian, and C. Russel, “An Energy Consumption Model for Energy Efficient Ethernet switches,” in *Proc. of the 2012 International Conference on High Performance & Simulation*, Madrid, Spain, July 2012, pp. 98–104.

[24] J. P. Thomesse, “Fieldbuses and Interoperability,” *Control Engineering Practice*, vol. 7, no. 1, pp. 81–94, 1998.

[25] S. Vitturi, L. Peretti, L. Seno, M. Zigliotto, and C. Zunino, “Real-Time Ethernet Networks for Motion Control,” *Computer Standards & Interfaces*, vol. 33, no. 5, pp. 465–476, September 2011.

[26] F. Benzi, G. Buja, and M. Felser, “Communication Architectures for Electrical Drives,” *Industrial Informatics, IEEE Transactions on*, vol. 1, no. 1, pp. 47–53, February 2005.

[27] P. Reviriego, J. Maestro, and P. Congdon, “Reduce Latency in Energy Efficient Ethernet Switches with Early Destination Lookup,” *EETimes*, April 2010.

[28] *IEC 61784: Industrial communication networks - Profiles - Part 2: Additional fieldbus profiles for real-time networks based on ISO/IEC 8802–3*, International Electrotechnical Commission Std., July 2010.

[29] Z. Hanzalek, P. Burget, and P. Sucha, “Profinet IO IRT Message Scheduling with Temporal Constraints,” *Industrial Informatics, IEEE Transactions on*, vol. 6, no. 3, pp. 369–380, August 2010.

[30] P. Ferrari, A. Flammini, S. Rinaldi, and E. Sisinni, “On the Seamless Interconnection of IEEE1588-Based Devices Using a PROFINET IO Infrastructure,” *Industrial Informatics, IEEE Transactions on*, vol. 6, no. 3, pp. 381 –392, August 2010.

[31] K. C. Lee, S. Lee, and M. H. Lee, “Worst Case Communication Delay of Real-Time Industrial Switched Ethernet with Multiple Levels,” *Industrial Electronics, IEEE Transactions on*, vol. 53, no. 5, pp. 1669–1676, October 2006.

[32] P. Ferrari, A. Flammini, F. Venturini, and A. Augelli, “Large PROFINET IO RT networks for factory automation: a case study,” in *Proc. of IEEE Conf. Emerging Technologies & Factory Automation ETFA 2011*, Toulouse, France, September 2011.

[33] T. J. Williams, *A reference model for computer integrated manufacturing: a description from the viewpoint of industrial automation*. Instrument Society of America, 1989.

[34] *EPSG Draft Standard 301 V1.1.0 Ethernet POWERLINK Communication Profile Specification*, Ethernet POWERLINK Standardization Group, 2008. [Online]. Available: http://www.ethernet-powerlink.org

[35] G. Cena, L. Seno, A. Valenzano, and S. Vitturi, “Performance Analysis of Ethernet Powerlink Networks for Distributed Control and Automation Systems,” *Computer Standards and Interfaces*, vol. 31, no. 3, pp. 566–572, March 2009.

[36] G. Prytz, “A Performance Analysis of EtherCAT and PROFINET IRT,” in *Emerging Technologies and Factory Automation, 2008. ETFA 2008. IEEE International Conference on*, pp. 408–415.

[37] *The CIP Networks Library – Volume 1 – Common Industrial Protocol (CIP)*, ODVA & ControlNet International Ltd Std., November 2007.

[38] *IEEE 1588-2008 Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*, Std., July 2008.

Stefano Vitturi Stefano Vitturi is a senior researcher with the Institute of Electronics and Computer and Telecommunications (IEIIT) of the National Research Council of Italy (CNR) since January 2002. He received the Laurea degree in Electronics Engineering from the University of Padova, in 1984. From 1985 to 2001 he has worked at the control and data acquisition system of RFX, a nuclear fusion experiment located in Padova, Italy. His research interests include industrial communication systems, real time communication networks (wired/wireless), implementation and performance analysis of devices conforming to the most popular industrial communication protocols. In these research fields he has authored/co-authored more than 100 scientific papers published in refereed international journals, books and conference proceedings

Federico Tramarin (S ’07, M ’12) is a Post-Doctoral research assistant with the Institute of Electronics, Computer and Telecommunication Engineering of the National Research Council of Italy (IEIIT-CNR). He got his PhD in Information Engineering at the beginning of 2012 with the Electronic Measurement research group, at the University of Padova, Italy, which he joined in 2008, after his master degree in Electronic Engineer. His main fields of interest are wireless networks, real-time communication and distributed measurement systems.