An Energy Efficient Ethernet Strategy Based on Traffic Prediction and Shaping

Angelo Cenedese *Member, IEEE*, Federico Tramarin *Member, IEEE*, Stefano Vitturi *Senior Member, IEEE*

***Abstract*—Recently, different communities in computer science, telecommunication and control systems have devoted a huge effort towards the design of energy efficient solutions for data transmission and network management. This paper collocates along this research line and presents a novel energy efficient strategy conceived for Ethernet networks. The proposed strategy, which exploits the opportunities offered by the IEEE 802.3az amendment to the Ethernet standard (known as Energy Efficient Ethernet), is based on the possibility of predicting the future traffic from the analysis of the current data flow. In agreement with the results of such a dynamic prediction, Ethernet links can be forced into a low power consumption state for variable intervals. Theoretical bounds are derived to detail how the performance figures depend on the parameters of the designed strategy and scale with respect to traffic load. Furthermore, simulation results carried out with both real and synthetic traffic traces are presented to prove the effectiveness of the strategy, which leads to considerable energy savings at the cost of only a limited bounded delay in data delivery.**

***Index Terms*—Ethernet networks, Energy Efficient Ethernet, Communication system traffic control, Prediction algorithms.**

I. Introduction

In the last decades data networks have become pervasive in everyday life and Ethernet [1], no longer limited to the office context, is ever more used in several fields of application from the industrial context to the home automation [2]. Indeed, the amount of data circulating in Ethernet networks is dramatically increasing due to the growing number of connections among users, the massive sharing of multimedia data and the widespread distribution of devices. However, due to the basically random nature of traffic, these networks are typically in an always active state. Consequently, also when there is no data to transmit, this results in a waste of energy and inefficiency, as observed in [3], [4], given that the energy consumption per Ethernet link is considerable (typically around 1 W for the 1GBASE-T Ethernet physical layer and over 5 W for the 10GBASE-T one [5], [6]).

Thus, several studies have been undertaken towards the design of energy efficient solutions for communication systems and Ethernet in particular [7], [4]: these efforts have led to the publication of the IEEE 802.3az amendment to the original standard, known as Energy Efficient Ethernet (EEE) [8]. IEEE 802.3az introduces a new operational mode for Ethernet, namely Low Power Idle (LPI), which allows links not involved in data transmission to enter a low consumption state [9]. Such an amendment, however, deliberately does not describe specific energy efficiency strategies (i.e. it does not specify when the links have to enter/exit the low consumption state) that, instead, are left to the specific manufacturer implementations. It follows that several EEE techniques have been proposed in the literature and, actually, some of them are implemented by commercially available devices.

A. Cenedese is with the Department of Information Engineering (DEI), University of Padova, Italy. F. Tramarin and S. Vitturi are with the National Research Council of Italy, CNR-IEIIT, Padova, Italy.
E-mails (✉): angelo.cenedese@unipd.it, {federico.tramarin,vitturi}@ieiit.cnr.it

Nonetheless, in most of current solutions, a greedy approach is adopted that does not consider the statistical properties of the traffic, which, conversely, could be profitably exploited to design effective EEE strategies. In other words, if the specific network traffic can be described and modelled to provide some predictive information, link states could be suitably controlled (activated/deactivated) in agreement with traffic prediction. Examples in this direction are the self–similar features of some Internet traffic [10], as well as periodic industrial traffic [11].

Following the very preliminary work presented in [12], this paper collocates along the research line of energy saving in Ethernet transmission and finds its motivation in the attempt of enhancing the EEE approach through the analysis of traffic, by providing also close bounds to the EEE strategy performance. More specifically, the main contribution of this paper is twofold:

- on the one side, it presents the design of an innovative EEE strategy, based on the prediction of the forthcoming traffic load, that further improves the energy savings achievable with the traditional EEE techniques; this strategy will be named as *EEE with Prediction*, EEEP;
- on the other side, the theoretical performance bounds for the energy efficient strategies EEE and EEEP are obtained; these bounds are also assessed by means of simulations that employ both real and artificially synthesized traffic traces.

II. Related Work

In the literature of the past years, many contributions have appeared related to IEEE 802.3az. In [9] the authors introduce one of the most popular EEE policies, namely *frame transmission*, and provide a thorough description of EEE along with an interesting analysis that addresses some macro economic aspects related to the expected power savings, deriving from the large scale adoption of EEE. In [13] a further technique, namely *burst transmission*, is proposed as an alternative to

To appear in the IEEE Transactions on Communication 2017

frame transmission. Such a strategy is also addressed in [14], which provides a deeper insight on its behavior. Moreover, a model of burst transmission for Gigabit Ethernet links, based on an M/G/1 queue, is proposed in [15] that allows to approximate both the energy savings and the delay in packet delivery introduced by the coalescing technique, which is at the basis of burst transmission.

Additional effective EEE strategies are proposed in both [16] and [17]. The former, actually, regards the adoption of sleeping algorithms, whereas the latter describes a technique to mitigate the delays that could affect packet delivery when EEE is used. Moreover, preliminary EEE performance figures concerned with energy consumption are provided in both [18] and [19]. Particularly, [19] presents the results of practical measurements carried out on some off-the-shelf Ethernet network interface cards.

Two interesting theoretical models related to EEE are respectively proposed in [3] and [20]. In detail, the former provides an exhaustive model of a network in which all nodes adopt EEE. The model, based on the assumption that the network traffic is that typical of the Internet, allows to calculate power consumption as well as some performance figures of such a kind of networks. The work presented in [20], instead, focuses on the intervals of time spent in the different EEE states by the network links that use frame transmission as EEE strategy, so that the overall power savings can be straightforwardly calculated.

A different context is studied in both [21] and [11], where the implementation and the performance analysis are given for typical real–time industrial communication systems.

From the analyses presented in the literature it clearly appears how application scenarios, specific traffic features and performance requirements of Ethernet networks have a deep impact on the possible exploitation of EEE strategies. In this context, the studies carried out over the past years on Ethernet traffic profiles revealed particularly helpful. Starting from early works on self–similarity [22], [23], several analyses have been proposed later on relevant to general Ethernet traffic [24], [25], as well as to traffic due to a more specific nature, for example, backbone traffic [26], multimedia traffic [27], VoIP traffic [28].

Traffic prediction has been already considered as a possibility for network power management in LANs [29], [30], nonetheless, to the best of the authors' knowledge, it was never exploited to implement EEE strategies. Indeed, this represents one of the main aspects that characterize this paper, along with the calculation of the theoretical performance bounds. On the other hand, this contribution clearly builds on the analyses presented in the literature, as well as on the experience obtained from practical applications. Particularly, burst transmission [13], [15] is a technique systematically addressed in this paper and its performance figures are compared with those obtained by EEEP.

III. Proposed Transmission Strategy

In this paper, an Ethernet switch is considered that receives data from a set of source nodes and transmits towards a specific output node (either another switch or a different device) with a transmission rate f. This traffic is constituted by Ethernet packets of possibly different sizes with average $\bar{d}$ (in bits). As it will be shown in the following, the proposed strategy works under the assumption that the transmit direction of the outgoing link can be be handled independently: this is possible for both the 100BASE-TX and 10GBASE-T physical layers, but not for 1GBASE-T. Thus, in this latter case, it is also assumed that the receive direction of the outgoing link can not be triggered by a partner.

A. From EEE to EEEP: the Rôle of Prediction

The LPI Ethernet operational mode, introduced by IEEE 802.3az allows network links to enter a state, namely *quiet state*, characterized by low power consumption with respect to the normal (*active*) state. The behavior of an Ethernet link that implements EEE can be summarized with reference to Fig. 1. Assuming that the link between two network nodes is active, the link moves to the quiet state in time t_s when there are no frames to transmit and reactivates either upon the arrival of any single frame transmission request. This strategy is referred to as *frame transmission*. A different strategy, namely *burst transmission*, which is nowadays more widespread, specifies that the number of queued packets has to reach a predefined threshold N_{th} before link activation. The time necessary to wake up a link is t_w. Hence, the total time required for a complete link transition is $T_{trans} = t_w + t_s$. Furthermore, a periodic refresh signal of duration t_r is triggered with a period t_q to ensure link integrity. It is worth mentioning that the burst transmission technique also makes use of a timeout T_{TO} such that the link is anyway activated upon its expiration, even if the number of coalesced packets has not overcome the threshold. Thus, the maximum delay in packet delivery introduced by burst transmission is represented by the maximum duration of a burst unit, which is upper bounded by T_{TO}. Clearly, both N_{th} and T_{TO} are two fundamental parameters of burst transmission.

The mean duration of a burst unit, represented by the time in which packets are first coalesced and then transmitted, is referred to as $\bar{T}_B$, whereas $\bar{N}_B$ indicates the mean number of packets transmitted in a burst unit.

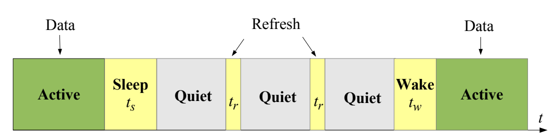

Fig. 1. EEE Standard. Schematic drawing of EEE timings.

In this paper, a new strategy is presented for the efficient control of network traffic, called *Energy Efficient Ethernet with Prediction* (EEEP): EEEP is based on the possibility of combining traditional EEE strategies with some level of knowledge of current traffic features from which the near future behavior can be predicted.

The proposed strategy, which is schematically described in Fig. 2, aims at effectively handling the outgoing link in order to minimize its energy consumption. To this purpose, the data traffic is partitioned into consecutive *time windows* $\mathcal{T}$, of fixed duration T. The mean number of packets transmitted in a time window is referred as $\bar{N}$.

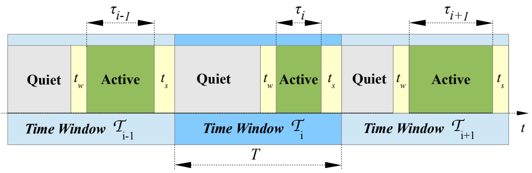

Fig. 2. EEEP strategy. Schematic drawing of operation.

During a generic time window $\mathcal{T}_i$ two actions are concurrently carried out, namely:

- the *prediction* of the traffic load expected in the next window, $\mathcal{T}_{i+1}$, based on the observation of the current traffic (i.e. the traffic that occurs in the current time window): this allows to compute an estimation of the time τ_{i+1} necessary to transmit the predicted future traffic;
- the *transmission* of the current traffic, which takes place with a modality that exploits the prediction made in the previous window, $\mathcal{T}_{i-1}$: at the beginning of $\mathcal{T}_i$, the outgoing link is forced in quiet state. Then, packet transmission is carried out during a single period of duration τ_i, before the end of the time window to ensure (on average) the whole transmission while minimizing the energy spent during transitions. In the interval τ_i, actually, the link is always in active state and transmits all the packets that have been queued during the previously forced quiet state, as well as those that arrive until the end of the current time window.

Thus, assuming that the prediction is correct and that a First–In First–Out policy is used for packet delivery, the maximum delay introduced by EEEP is equal to $T - \tau$.

The value of τ can be computed according to different prediction techniques, as will be made clearer in Section III-B. Several aspects have to be taken into consideration in its choice, among which the prediction accuracy and the complexity of the procedure that performs the traffic estimation along with the consequent computational burden of the algorithm that implements it. However, an inaccurate prediction reveals as a possible drawback of the proposed strategy. Indeed, if the interval τ is not sufficient to deliver the actual traffic, then the link will not enter the idle state at the end of the current time window and packet transmission will be prolonged at the beginning of the following one, with the consequent reduction of the energy savings.

It has to be observed that the duration of a time window, T, definitely represents a fundamental parameter of the EEEP strategy. Indeed, it has an important impact on the packet transmission delay and, also, it represents an upper bound for the computation time of the prediction algorithms used by the EEEP strategy (since these algorithms have to be executed within a time window). Particularly, on the one hand T has to ensure that the packet delay introduced by EEEP is kept under control while, on the other hand, it should allow to achieve satisfactory energy savings. This tradeoff imposes

TABLE I
MAIN VARIABLES

Parameter	Name	Unit
L	Traces length	[s]
f	Link transmission rate	[bit/s]
T	Time window length	[s]
$\bar{T}_B$	Average burst unit length	[s]
$\bar{N}$	Average number of packets per time window	–
$\bar{N}_B$	Average number of packets per burst unit	–
$\bar{n}_C$	Average number of burst units per time window	–
$\bar{d}$	Average packet size	[bit]
$\bar{T}_{pack}$	Average packet transmission time	[s]
T_{trans}	Total link transition time	[s]
N_{th}	Packet number threshold for burst transmission	–
T_{TO}	Time–out for burst transmission	[s]

a punctual monitoring of the incoming data flow to timely predict its behavior in the next time window, $\mathcal{T}_{i+1}$. The above considerations suggest the adoption of a short term prediction approach as will be better detailed in the following.

For the sake of completeness, it is worth mentioning that refresh signals have not been considered in the description of EEEP. This is because, as will be clarified in the performance analysis sections, the impact of these signals on power consumption can be safely neglected.

Finally, the most important variables and parameters introduced in this paper are summarized in Table I.

B. Traffic Prediction Strategies

Let $X(t)$ with $t \in \mathbb{Z}$ be a stochastic process that will be used to model the network traffic and, specifically, to indicate the traffic load on an Ethernet link, expressed in bit/s.

As a general approach to predict network traffic, good models are given by the Auto-Regressive Moving Average (ARMA) and the Integrated Auto-Regressive Moving Average (ARIMA) models [31], [32], [33]. An ARMA(p,q) model of a stationary time series $X(t)$ is a dynamical description model that presents an auto-regressive (AR) part of order p and a moving average (MA) part of order q, and satisfies for each t the equation

$$X(t) + \underbrace{\sum_{j=1}^{p} \varphi_j X(t-j)}_{AR} = \epsilon(t) + \underbrace{\sum_{k=1}^{q} \vartheta_k \epsilon(t-k)}_{MA}$$

where ϵ is a Gaussian i.i.d. white noise, $\{\varphi_j\}$, $\{\vartheta_k\}$ are the parameters of the model, and the following conditions stand: $\varphi_p \neq 0$ and $\vartheta_q \neq 0$. Moreover, the ARMA(p,q) model can be generalized into the ARIMA(p,d,q) by introducing a d-th order integration term in the AR part that allows to reduce possible non-stationarity in the original time series that needs to be described.

Given the aforementioned requirements of the prediction procedure, in the context of this paper these linear methods appear to be suitable also with a choice of low orders. Each model, referred to as $\mathcal{M}_i$, estimates its parameters every time window $\mathcal{T}_i$ using the maximum likelihood criterion. Then, it is able to predict the traffic that will occur in the next time window $\mathcal{T}_{i+1}$.

These sophisticated models are very popular but, unfortunately, they introduce a considerable computational burden with the consequent negative impact on overall performance.

For this reason, an effective alternative strategy, more efficient from a computational point of view, is presented here. Such a strategy is based on the construction of a conditional probabilities table that predicts the traffic level of the forthcoming time window, given that of the current one.

In detail, for the stochastic process $X(t)$, let

$$V_i = \sum_{k\in\mathcal{T}_i} X(k), \quad V_{i+1} = \sum_{k\in\mathcal{T}_{i+1}} X(k),$$

with V_i, V_{i+1} random variables that account for the traffic modeled by process X in, respectively, time windows $\mathcal{T}_i$ and $\mathcal{T}_{i+1}$. Let

$$v_{max} = \max_x \sum_{k\in\mathcal{T}_i} x(k), \quad v_{min} = \min_x \sum_{k\in\mathcal{T}_i} x(k),$$

with $x(k)$ a realization of $X(k)$, being v_{max} and v_{min} respectively the highest and the lowest traffic seen in $\mathcal{T}_i$. A traffic quantization step $\mu = \frac{v_{max}-v_{min}}{h}$ is introduced and the whole traffic range (i.e. from 0 to $+\infty$) is partitioned into h levels, namely

$$\{(0, v_{min}+\mu), [v_{min}+\mu, v_{min}+2\mu), \ldots, [v_{min}+(h-1)\mu, +\infty)\}$$

The traffic levels can thus be associated to a random variable L_i that relates to V_i according to the following relations:

$$\begin{aligned} L_i = 1 &\Leftrightarrow V_i \in (0, v_{min}+\mu) \\ L_i = 2 &\Leftrightarrow V_i \in [v_{min}+\mu, v_{min}+2\mu) \\ &\vdots \\ L_i = h &\Leftrightarrow V_i \in [v_{min}+(h-1)\mu, \infty). \end{aligned}$$

With the above assumptions, a row stochastic matrix $P \in \mathbb{R}^{h\times h}$ can be built that reports the conditional probabilities of having a specific traffic level $L_{i+1} = l'$ in the time window $\mathcal{T}_{i+1}$, given that in the previous time window $\mathcal{T}_i$ the observed traffic was $L_i = l$. The elements of this conditional probability matrix are defined as

$$P(l, l') = \mathbb{P}[L_{i+1} = l' | L_i = l] \quad \forall l, l' \in \{1, 2, \ldots h\}. \tag{1}$$

The computational burden to calculate matrix P appears considerably simple, since in this case only elementary mathematical operations have to be carried out. This matrix is then used to predict the traffic load in $\mathcal{T}_{i+1}$ and implement the EEEP strategy.

C. Implementations of the EEEP Strategy

Two algorithms are presented to implement the EEEP strategy: in the first one, traffic prediction is obtained with ARMA/ARIMA models, while the second one makes use of the conditional probability matrix[1].

In Algorithm 1, for each $\mathcal{T}_i$, the policy consists of the initial application of the prediction model $\mathcal{M}_{i-1}$ to obtain the predicted number of packets $\hat{N}_{\mathcal{T}_i}$ in $\mathcal{T}_i$; this phase is followed by a transmission interval, of duration τ_i, computed accordingly as

$$\tau_i = \hat{N}_{\mathcal{T}_i} \bar{T}_{pack}$$

$\bar{T}_{pack}$ being the mean packet transmission time observed so far.

A different approach is used by Algorithm 2, which exploits the condition probability matrix (1). Here, the single transmission interval of duration τ_i occurs only if the expected traffic in $\mathcal{T}_i$ is equal or lower than that measured in $\mathcal{T}_{i-1}$ with probability greater than θ, which is a threshold parameter specified at design stage. If this condition does not hold, then a standard EEE strategy is adopted. Thus, when the prediction is applied, the number of expected packets in $\mathcal{T}_i$ is less or equal to that of $\mathcal{T}_{i-1}$. As a consequence, Algorithm 2 safely sets time τ_i as

$$\tau_i = N_{\mathcal{T}_{i-1}} \bar{T}_{pack}$$

where $N_{\mathcal{T}_{i-1}}$ is the number of packets transmitted in $\mathcal{T}_{i-1}$.

Algorithm 1 EEEP with ARMA/ARIMA prediction

1: Initialize T
2: Initialize ARMA/ARIMA model
3: **while** traffic exists, within each window $\mathcal{T}_i$ of length T **do**
4: Compute τ_i based on $\mathcal{M}_{i-1}$ **and** Build ARMA/ARIMA model $\mathcal{M}_i$
5: At $t = (i+1) * T - \tau_i - T_{trans}$ turn the link ON
6: Transmit data during period τ_i
7: Turn the link OFF
8: **end while**

Algorithm 2 EEEP with conditional probability prediction

1: Initialize T
2: Initialize the conditional probability matrix P
3: **while** traffic exists, within each window $\mathcal{T}_i$ of length T **do**
4: **if** $\mathbb{P}[L_i \le l | L_{i-1} = l] \ge \theta$ **then**
5: Compute τ_i based on L_{i-1} **and** Update P
6: At $t = (i+1) * T - \tau_i - T_{trans}$ turn the link ON
7: Transmit data during period τ_i
8: Turn the link OFF
9: **else**
10: Transmit data using EEE
11: **end if**
12: **end while**

Remark 1. *In both algorithms the effects of possible inaccurate predictions can be mitigated by adding an interval $\Delta\tau$ to the already computed τ, at the cost of diminishing the energy gain. Indeed, by exactly choosing the predicted τ, the duration*

[1]For simplicity, the pseudo-codes of both algorithms do not consider the case of inaccurate prediction. Such an aspect, however, has been adequately taken into account in the simulations used to carry out the performance analysis presented in Section V.

of the active state in the current time window is minimized and the energy gain is maximized; conversely, by allowing a further increment to τ, that is setting this transmission period to $\tau + \Delta\tau$, EEEP reveals less effective, since the energy gain may be reduced. On the other hand, the time granted to packet transmission is longer and, hence, the probability that the transmission of a packet is moved to the next time window due to a wrong prediction is lowered.

Remark 2. *Algorithm 2 could be refined by introducing further conditions to carry out prediction. For example, it could be imposed that the expected traffic flow in the next time window overcomes a minimum value (l_{min}) with probability greater than a threshold (p_{th}), i.e. $\mathbb{P}[L_i \geq l_{min} | L_{i-1} = l] \geq p_{th}$. This improves the algorithm's ability select the more convenient strategy (EEE or EEEP) depending on traffic features.*

It is worth observing that the condition expressed by Line #4 of Algorithm 2 and the possibly additional ones mentioned by Remark 2 reflect a conservative assumption since, in principle, they reduce the number of time windows in which the prediction can be used, and eventually limits the possible energy savings. Also, the value of τ is calculated in the worst case (assuming that the traffic in the next window is the same as that measured in the current one whereas, in practice, it could be lower). On the other hand, the strategy defined by Algorithm 2 reveals particularly effective when the traffic load does not vary significantly from one time window to the following, since in such a situation the prediction will be very precise, as it is the case for example of self–similar traffic [22]. In this regard, it is to note how the self–similarity of the traffic is strictly related to the structure of the conditional probability matrix since, in this case, P has a diagonal structure, being the probability $\mathbb{P}[L_i = l | L_{i-1} = l]$ close to 1. Another interesting scenario is encountered in the case of industrial traffic. Here, the huge amount of periodic data exchanged might drive the design of specifically tailored strategies.

Finally, it has to be remarked that this paper, as a major contribution, aims at showing that the introduction of prediction strategies is beneficial in the overall energy efficiency Ethernet framework. Thus, the performance analysis presented later on is concerned with metrics like the power savings of the overall system as well as with the packet delays additionally introduced by EEEP, rather than with the formal optimality of the prediction algorithms.

D. Implementation Issues of the EEEP Algorithms

In the perspective of an implementation of the proposed algorithms on real network devices, it is worth taking into account the additional computational burden due to the EEEP strategy these devices are required to support.

The pseudo–code of Algorithms 1 and 2 presented in Section III-C give a clear snapshot of the principal operations necessary to implement the EEEP strategy. As can be seen, both Algorithms share most of their operations, since they are based on queue management, calculation of the value of τ and activation/deactivation of the transmission link. It is worth observing that both these operations are elementary tasks that any EEE capable device should be able to carry out.

The substantial difference between the two Algorithms relies on the computational complexity of the prediction strategies. Indeed, Algorithm 1 requires to build, within each time window, an ARMA/ARIMA model whose construction is a complex task, based on the solution of a constrained minimization problem, that may imply a considerable computational load. On the other hand, Algorithm 2 simply requires to update the conditional probability matrix P, and to decide whether prediction has to be adopted or not in the following time window. Both these steps are carried out by executing a limited number of elementary operations that commercial devices can easily provide. It follows that the computational burden of Algorithm 2 is very low and, also, its simplicity reflects on a very limited increase of power consumption for the devices that implement it. Differently, the adoption of Algorithm 1 by real devices might be more problematic due to the model complexity, thus resulting more demanding in terms of resources.

IV. Theoretical Performance Bounds

In this section, the EEEP strategy is analyzed in order to obtain its performance evaluation and comparisons are carried out with respect to two limit cases. On the one hand, it is considered the case in which energy efficiency is not employed at all, which is referred to as *Always-On* policy. On the other hand, there is the case in which a standard EEE strategy is applied for the whole time window, referred to as *EEE*. In this latter case, it is assumed that burst transmission is adopted.

A. Time Spent in Quiet State

The first performance indicator considered is the percentage of time spent in quiet state within a time window.

When the Always-On policy is adopted, clearly, no time is spent in quiet state.

If EEE is adopted, considering that under the above hypotheses there is one transition (from quiet to active and back) during a burst unit, then the mean number of transitions in T is given by the ratio $\bar{n}_C = \frac{T}{\bar{T}_B}$, where $\bar{T}_B$ is the mean duration of a burst unit as mentioned before. It is worth observing that $\bar{n}_C$ depends on two parameters typical of the energy efficiency framework. Specifically, the duration of the time window T characterizes the EEEP behavior, whereas $\bar{T}_B$ is strictly related to the burst transmission parameters N_{th} (packet number threshold per burst unit) and T_{TO} (timeout value). In practice, $\bar{n}_C$ results to be larger than unity when $\bar{T}_B < T$, which may correspond, for example, to situations of medium to high traffic flows, when the packet number threshold for burst transmission (N_{th}) is rapidly reached and hence $\bar{T}_B$ is generally low with respect to T. Conversely, in the case of a low incoming packet rate, it may result $\bar{T}_B > T$ and then $\bar{n}_C$ can decrease below unity.

The percentage of T in which the link is in quiet state is given by the following equation:

$$p_{EEE} = \frac{T - \bar{n}_C T_{trans} - \bar{N}\bar{T}_{pack}}{T} \quad (2)$$

where, as defined in Section III, T_{trans} is given by the sum $t_s + t_w$, whereas $\bar{N}$ accounts for the mean number of packets transmitted within a time window, and $\bar{T}_{pack}$ is the mean time to transmit a packet.

With the EEEP strategy, only one transition is performed at the end of the time window. Therefore, in such a case, the percentage of time in quiet state becomes

$$p_{EEEP} = \frac{T - T_{trans} - \bar{\tau}}{T} \quad (3)$$

where $\bar{\tau}$ is the mean value of τ, as introduced in Section III-A.

Remark 3. *It has to be stressed that, while the behavior of EEE relies on the* current *packet arrivals, that of EEEP depends on the prediction of the* future *packet arrivals. Nonetheless, if the prediction is correct, then the number of predicted packets ($\hat{N}$) well approximates that of the actual ones (N) for each time window. Consequently, it can be written: $\tau = \hat{N}\bar{T}_{pack} \approx N\bar{T}_{pack}$. Thus, from a mathematical point of view, it appears from (2) and (3) that EEEP is the specific case of EEE for $\bar{n}_C = 1$.*

Lemma 1 (Packet number limits in a time window)**.** *There exist limit values for the number of packets that can be transmitted in a time window for the three strategies, namely $\bar{N}^l_{ON}$, $\bar{N}^l_{EEE}$ and $\bar{N}^l_{EEEP}$. For these values, it holds: $\bar{N}^l_{EEE} < \bar{N}^l_{ON}$ and $\bar{N}^l_{EEEP} < \bar{N}^l_{ON}$, while the relation between $\bar{N}^l_{EEE}$ and $\bar{N}^l_{EEEP}$ is regulated by the $\bar{n}_C$ value (in particular, if $\bar{n}_C > 1$ then $\bar{N}^l_{EEE} < \bar{N}^l_{EEEP}$).*

Proof: For the Always-On policy it results, trivially,

$$\bar{N}^l_{ON} = \frac{T}{\bar{T}_{pack}}. \quad (4)$$

When applying the standard EEE policy, $\bar{n}_C$ transitions are accommodated in a time window of length T, meaning that the following time limit stands

$$T = \bar{N}^l_{EEE}\bar{T}_{pack} + \bar{n}_C T_{trans}. \quad (5)$$

By solving (5) in the unknown $\bar{N}^l_{EEE}$, it follows straightforwardly

$$\bar{N}^l_{EEE} = \left\lfloor \frac{T}{\bar{T}_{pack}} \left(1 - \frac{T_{trans}}{T}\bar{n}_C\right)\right\rfloor. \quad (6)$$

According to Remark 3, the value for EEEP results

$$\bar{N}^l_{EEEP} = \left[\bar{N}^l_{EEE}\right]_{\bar{n}_C=1}. \quad (7)$$

The comparison of (4), (6) and (7) provides the immediate proof of the lemma. ■

B. Strategy Efficiency Indicator

A strategy efficiency indicator can be defined as the ratio between the time interval employed for the actual data transmission and the time interval during which the link is active. Such an efficiency represents a meaningful performance index, since it accounts for the capability of a strategy to maintain a link in the active state for the time strictly necessary for data transmission.

Clearly, for the Always-On policy, it stands

$$\eta_{ON} = \frac{\bar{N}\bar{T}_{pack}}{T}, \quad (8)$$

while for the EEE, it follows

$$\eta_{EEE} = \frac{\bar{N}\bar{T}_{pack}}{\bar{n}_C T_{trans} + \bar{N}\bar{T}_{pack}} = \frac{\bar{N}}{\bar{N} + \frac{T_{trans}}{\bar{T}_{pack}}\bar{n}_c} \quad (9)$$

which, for the EEEP case ($\bar{n}_C = 1$), becomes

$$\eta_{EEEP} = \frac{\bar{N}}{\bar{N} + \frac{T_{trans}}{\bar{T}_{pack}}}, \quad (10)$$

Both η_{EEE} and η_{EEEP} represent hyperbolas passing through the origin and tending to unity, with vertical asymptote at $\left(-\frac{T_{trans}}{\bar{T}_{pack}}\bar{n}_C\right)$.

The following considerations can be made with respect to the above efficiency equations:

- the relations $\eta_{ON} < \eta_{EEE}$ and $\eta_{ON} < \eta_{EEEP}$ always hold.
- $\bar{n}_C$ has a negative impact on η_{EEE}, i.e. the larger the number of burst units, the lower the efficiency.
- since both $\bar{N}$ and $\bar{n}_C$ are directly proportional to the time window length T, both η_{ON} and η_{EEE} do not scale with T, differently from η_{EEEP} whose performance is affected by the time window length.

With respect to this latter observation, Fig. 3 shows an example of the efficiency behavior for one of the Ethernet traces that will be considered later in Section V. Such a trace, which is referred to as R1G #A1, has the following parameters: f = 1 Gbit/s, $\bar{T}_{pack}$ = 5.68 µs and $\bar{d}$ = 5688 (see Table III).

In the left panel of Fig. 3, the packet number threshold for burst transmission has been set to N_{th} = 15, in agreement with [34], which results in an average burst unit length $\bar{T}_B \approx 1$ ms. As can be seen, when $T > \bar{T}_B$, EEEP outperforms EEE: indeed, choosing for example time windows with $T = \{5, 10, 25, 50\}$ ms, it results respectively $\eta_{EEEP} \approx \{64, 78, 90, 94\}\%$, whereas $\eta_{EEE} \approx 26\%$ with an efficiency increment that ranges from 38% to 68% for the chosen values.

Conversely, when $T < \bar{T}_B$, the efficiency of EEE results higher than that of EEEP, and tends to $\eta = 1$ for large $\bar{T}_B$ (and T fixed). This is the case shown in the right panel of Fig. 3 for $\bar{T}_B > 50$ ms. Here, the time window of EEEP is set to $T = 50$ ms, whereas $\bar{T}_B$ is increased by progressively increasing the packet number threshold N_{th}. In situations like this one, the possibility of introducing additional conditions expressed by Remark 2 reveals very effective, in that it allows to select EEEP only for the cases in which such strategy results more efficient than EEE. Thus, in general, EEEP represents a

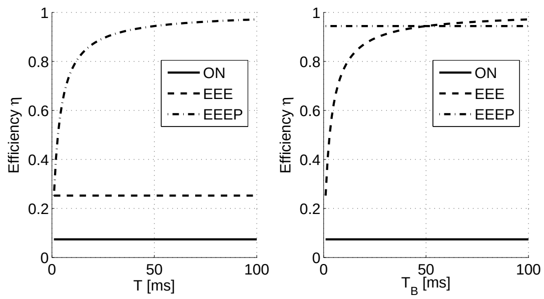

Fig. 3. Efficiency comparison for a specific trace (R1G #A1, Table III). On the left, the efficiencies of the strategies Always-On (ON), EEE and EEEP are compared *vs* the time window length T (with $N_{th} = 15$, which leads to $\bar{T}_B \approx 1$ms). On the right, the efficiencies are compared *vs* the mean burst unit duration $\bar{T}_B$ (with $T = 50$ms).

viable option to increase the energy savings achievable with EEE. Consequently, only the case $\bar{n}_C > 1$ will be addressed in the remainder of this paper, unless differently stated.

Efficiency formulas (8)-(9)-(10) and the considerations above lead to the formalization of the following propositions.

Proposition 1 (Efficiency bounds)**.** *Consider Always-On, EEE, and EEEP: only η_{ON} can reach unitary efficiency (when $\bar{N} = \bar{N}_{ON}^{max} = \frac{T}{\bar{T}_{pack}}$), while η_{EEE} and η_{EEEP} are strictly below unity:*

$$\eta_{EEE} \leq 1 - \frac{T_{trans}}{T}\bar{n}_C$$
$$\eta_{EEEP} \leq 1 - \frac{T_{trans}}{T}.$$

Proof: These results follow from the efficiency defintions (9)-(10) and the maximum values for $\bar{N}_{EEE}$ and $\bar{N}_{EEEP}$, namely

$$\bar{N}_{EEE}^{max} = \frac{T - \bar{n}_C T_{trans}}{\bar{T}_{pack}}$$
$$\bar{N}_{EEEP}^{max} = \frac{T - T_{trans}}{\bar{T}_{pack}}$$

The proposition is now proved by substituting these expressions respectively in the relations below:

$$\eta_{EEE} \leq 1 - \frac{\bar{n}_C T_{trans}}{\bar{n}_C T_{trans} + \bar{N}_{EEE}^{max}\bar{T}_{pack}}$$
$$\eta_{EEEP} \leq 1 - \frac{T_{trans}}{T_{trans} + \bar{N}_{EEEP}^{max}\bar{T}_{pack}}$$

■

Proposition 2 (Performance characterization)**.** *A value $\bar{N}^*$ can be computed for both the EEE and EEEP strategies, at which the gain in efficiency is maximized with respect to the Always-On policy:*

$$\bar{N}_{EEE}^* = \left\lfloor \sqrt{\frac{T_{trans}}{\bar{T}_{pack}}\bar{n}_C} \left(\sqrt{\frac{T}{\bar{T}_{pack}}} - \sqrt{\frac{T_{trans}}{\bar{T}_{pack}}\bar{n}_C} \right) \right\rceil \quad (11)$$

$$\bar{N}_{EEEP}^* = \left\lfloor \sqrt{\frac{T_{trans}}{\bar{T}_{pack}}} \left(\sqrt{\frac{T}{\bar{T}_{pack}}} - \sqrt{\frac{T_{trans}}{\bar{T}_{pack}}} \right) \right\rceil \quad (12)$$

Proof: By computing the efficiency difference $\Delta\eta_{EEE} = \eta_{EEE} - \eta_{ON}$ and differentiating w.r.t. $\bar{N}$, it follows:

$$\frac{\partial \Delta\eta_{EEE}}{\partial \bar{N}} = \frac{\bar{T}_{pack}}{T} \frac{T T_{trans}\bar{n}_C - T_{trans}^2\bar{n}_C^2 - 2\bar{N}\bar{T}_{pack}T_{trans}\bar{n}_C - \bar{N}^2\bar{T}_{pack}^2}{\left(\bar{N}\bar{T}_{pack} + T_{trans}\bar{n}_C\right)^2};$$

this expression is then equalized to zero to obtain the point of maximum $\bar{N}_{EEE}^*$.

Finally, (12) is trivially obtained from (11) posing $n_C = 1$. ■

This overall performance behavior is summarized by the example reported in the two panels of Fig. 4 where the efficiency of EEE/EEEP is compared with that of the Always–On policy for a link with the same traffic features of Fig. 3. In these plots, the figures of merit are given as a function of the *offered load*, defined as the number of bits transmitted in a time window divided by its duration, expressed in percentage of the maximum link capacity (1 Gbit/s).

It can be seen from Fig. 4a how EEEP can reach an higher efficiency with respect to EEE according to the window length and for the whole range of the offered load; in this sense, since $\bar{n}_C \geq 1$, EEE clearly represents a lower bound for the EEEP strategy as T decreases (consistently with $\bar{n}_C \to 1$), The maximum efficiency limits stated by Prop. 1 are reached for offered loads correspondent to the packet number values $\bar{N}_{EEE}^l$ and $\bar{N}_{EEEP}^l$ described in (6) and (7) respectively. In addition, Fig. 4b highlights the point of maximum efficiency gain with respect to the Always-On policy (Prop. 2).

C. Transmission Algorithm Efficiency

Notably, since different algorithms can implement the EEEP strategy, the concept of efficiency can be extended from the strategy to the actual algorithm so as to discern among their performance figures. Consequently, the efficiency of the algorithm can be defined similarly to that of the strategy, as the ratio between the time in which the link is used to transmit data and that in which the link is in the active state.

Since Algorithm 1 applies prediction in every time window, then its behavior corresponds to that of η_{EEEP}, and hence it results $\eta_{Alg1} = \eta_{EEEP}$.

Conversely, Algorithm 2 applies prediction only if a specific condition holds ($\mathbb{P}[L_{i+1} \leq l | L_i = l] \geq \theta$), hence its efficiency index belongs to the region between the two curves referring respectively to EEE (lower bound) and to EEEP (upper bound). In practice, it holds $\eta_{EEE} \leq \eta_{Alg2} \leq \eta_{EEEP}$.

To formally represent this case, let U be the percentage of time windows in which the EEEP strategy is used: clearly, for Algorithm 1 it results $U = 1$, whereas for Algorithm 2 it holds $0 \leq U \leq 1$. Hence, recalling (2)-(3), the fraction of time $p(U)$ when the link is in quiet state results as

$$\begin{aligned} p(U) &= U p_{EEEP} + (1 - U) p_{EEE} \\ &= U(p_{EEEP} - p_{EEE}) + p_{EEE} \end{aligned}$$

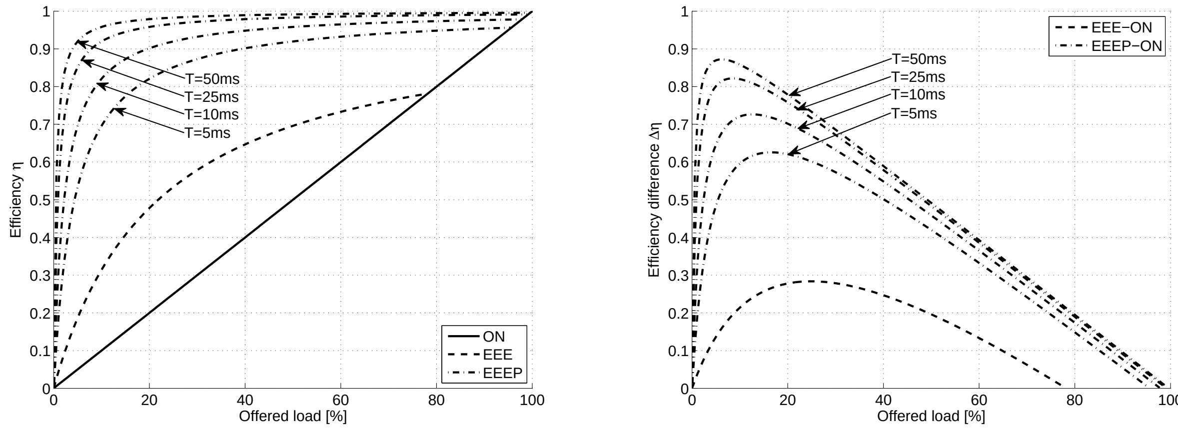

(a) (b)

Fig. 4. Efficiency of the strategies *vs* offered load, for different window lengths T. (a) The dashed lines refer to the EEE and EEEP strategies; the solid line corresponds to the Always-On policy. (b) The dashed lines refer to the difference between the two energy efficient strategies and the Always-On policy.

Evidently, $p(U) \in [p_{EEE}, p_{EEEP}]$ depends on the intrinsic characteristic of the traffic as well as on the threshold value θ used by Algorithm 2 to decide whether apply the prediction results or not. Thus, a convex combination of η_{EEE} and η_{EEEP} regulated by U stands also for the efficiency of the transmission when adopting Algorithm 2, η_{Alg2}

$$\eta_{Alg2} = U(\eta_{EEEP} - \eta_{EEE}) + \eta_{EEE}$$

D. Energy Consumption

Besides the efficiency parameters, it is interesting to perform a comparison among the different strategies and their implementations in terms of actual energy consumption. In this respect, let PW_{ON} and PW_{OFF} be the electric power of, respectively, active and quiet state.

For a given traffic interval of duration L, Algorithm 2 applies an energy efficient policy during $K_{win} = \left\lfloor \frac{L}{T} \right\rfloor$ time windows, choosing either standard EEE or EEEP. The energy $E(U)$ spent by the link in $K_{win}T \approx L$ time windows is given by:

$$\begin{aligned} E(U) &= K_{win}T\Big(p\,(U)\,PW_{OFF} + (1 - p\,(U))\,PW_{ON}\Big) \\ &= K_{win}T\Big(p(U)\,(PW_{OFF} - PW_{ON}) + PW_{ON}\Big) \end{aligned} \quad (13)$$

Eq. (13) allows to estimate the energy gain EG that EEEP brings with respect to the standard EEE strategy

$$\begin{aligned} EG &= \frac{E_{EEE} - E(U)}{E_{EEE}} \\ &= \frac{p_{EEE} - p(U)}{p_{EEE} + \frac{PW_{ON}}{PW_{OFF} - PW_{ON}}} \\ &= U\frac{p_{EEEP} - p_{EEE}}{\frac{PW_{ON}}{PW_{ON} - PW_{OFF}} - p_{EEE}} \end{aligned} \quad (14)$$

where E_{EEE} is the energy consumption when the EEE strategy is adopted. Substituting in (14) the expressions of p_{EEE} and p_{EEEP} provided by (2) and (3) respectively, it results

$$EG = \frac{U}{T}\frac{(\bar{n}_C - 1)T_{trans} + \bar{N}\bar{T}_{pack} - \bar{\tau}}{\frac{PW_{ON}}{PW_{ON} - PW_{OFF}} - p_{EEE}}$$

Introducing the time $\Delta\tau$ to mitigate the effects of wrong predictions of the strategy (see Remark 1), it follows that it is possible to express the time $\Delta\tau$ as an additional percent fraction of $\bar{\tau}$. Specifically, the *mitigation coefficient* $p_{\bar{\tau}}$ can be defined according to

$$\bar{\tau} + \Delta\bar{\tau} = (1 + p_{\bar{\tau}})\bar{\tau}$$

that will be used in the following of the paper. With such a position, the energy gain EG can be finally written as

$$EG = \frac{U}{\frac{PW_{ON}}{PW_{ON} - PW_{OFF}} - p_{EEE}}\left(\frac{(\bar{n}_C - 1)T_{trans} + \bar{N}\bar{T}_{pack}}{T} - \frac{(1 + p_{\bar{\tau}})\bar{\tau}}{T}\right) \quad (15)$$

This equation represents a theoretical expression for the energy gain of EEEP over EEE as a function of the parameters of the device in use (PW_{ON}, PW_{OFF}) and of the control variables of the adopted energy efficient strategy. In particular, the dependence on U is clearly stated, and, more interestingly, a linear dependence on $p_{\bar{\tau}}$ is highlighted.

V. Performance Evaluation

In this section, the assessment of the algorithms performance, and the validation of the theoretical bounds are given by means of simulations, carried out via Matlab™. Both real Ethernet traffic traces and synthetically generated ones have been used. Specifically, the selected real traces belong to different public traffic archives, they refer to both 1 Gigabit (identified as R1G #A1–A5, R1G #B1–B13 and R1G #C1–C4) and 10 Gigabit Ethernet links (R10G #A1–A18, R10G #B1–B18 and R10G #C1–C18), and have been analyzed over intervals L = 200 s and L = 20 s, respectively. On the other hand, the synthetic traffic traces serve to build a wide controlled dataset, where the main characteristics (e.g. offered load, traffic variability) of the traces are opportunely designed in order to span the operational range of interest. In particular, these datasets refer to high offered load values (S1G #A1–A200 and S10G #A1–A200), as well as to low offered load values (S1G #B1–B200 and S10G #B1–B200), with randomized traffic distributions.

Table III summarizes the most meaningful parameters of the adopted traces and Table II reports the main parameters of both Algorithms 1 and 2 used in the performance analysis. The value of the conditional probability threshold, θ, for Algorithm 2 has been set to 0.5, since this represents the fairest choice to decide whether prediction has to be carried out or not, in absence of information about the load on the next time window. The number of quantization levels h has been set to a value that on the one hand is able to represent the traffic load within a time window with sufficient accuracy while, on the other hand, limits the size of the conditional probability matrix P, as confirmed by some preliminary analyses not reported here due to space limitations[2] [35]. The EEE transition times have been derived from the standard document [8], whereas the power consumption values have been directly taken from the data sheets of commercial products [5], [6]. As far as the burst transmission parameters are concerned, the packets number threshold values have been selected according to the traffic conditions, as well as following indications of the literature [34], [13]: the values presented in Table II are relevant to the six specific traces analyzed in detail later in Subsec. V-B. Also, the timeout value has been set equal to the time window length T, which guarantees the condition $\bar{n}_C \geq 1$. Finally, as can be seen, two different time window durations have been selected for each Ethernet physical layer. In this case, the choice has been made taking into consideration the impact that the EEEP strategy may have on system performance, particularly on the packet delay, which is upper bounded by T. Focusing for example on 1GBASE-T, the two selected values (T = 10 ms and T = 50 ms) appear able to cope with the requirements of several practical applications [13].

A. Assessment of EEEP Efficiency Bounds

In this subsection, the efficiency of the algorithms that implement the EEEP strategy is assessed with respect to the theoretical bounds derived in Section IV, by computing it as the ratio between the duration of the intervals in which the link is in active state and the overall simulation length L. In particular, Algorithm 2 is considered, since the efficiency of Algorithm 1 is well approximated by η_{EEEP}, as addressed in Section IV-C.

Fig. 5 shows the obtained results for the selected window lengths for each link rate. Here, markers refer to the simulated

[2]Clearly, both θ and h might be determined in a more systematic way, especially in case some information about the traffic profile is a-priori known. However, this selection procedure is outside the specific topics of this paper. Consequently, apart from a preliminary analysis and tuning phase, the in–depth assessment of these parameters is left to future developments.

TABLE II
MAIN SIMULATION PARAMETERS

Parameter	Name	Value 1GBASE-T	Value 10GBASE-T
L	Traces length	200 s	20 s
T	Time window length	50 ms, 10 ms	5 ms, 1 ms
h	Number of quantization levels	8	8
θ	Conditional probability threshold	0.5	0.5
PW_{ON}	Power Consumption, Active state	0.697 W	5 W
PW_{OFF}	Power Consumption, Quiet state	0.053 W	0.5 W
t_w	EEE wake time	16.5 µs	2.88 µs
t_s	EEE sleep time	202 µs	4.48 µs
N_{th}	Packet number threshold for burst transmission	According to traffic condition: 15 (R1G#A1), 5 (R1G#B1), 75 (R1G#C1)	According to traffic condition: 10 (R10G#A1), 35 (R10G#B1), 65 (R10G#C1)
T_{TO}	Time–out for burst transmission	Equal to T	Equal to T
t_q	EEE quiet time	20 ms	39.68 µs

η_{Alg2} values (asterisks have been used for real traces, and dots for synthetic ones), whereas the dash-dotted curve reports the EEEP theoretical bound and the shaded area indicates efficiency levels below the Always-On policy.

In this respect, some interesting observations are in order:

- the theoretical EEEP behavior represents a good upper bound for the policy defined by Algorithm 2, for both the 1GBASE-T and the 10GBASE-T cases and for all the considered traces. Conversely, the theoretical EEE efficiency is a lower bound in the case $\bar{n}_C > 1$: in Fig. 5 this curve is not shown since, differently from the EEEP one, it depends on the value of $\bar{n}_C$, which is characteristic of each trace.
- by reducing the window length, both the EEEP efficiency and (consequently) the η_{Alg2} decrease, according to Fig. 3. Nonetheless, η_{Alg2} remains close to the theoretical limit.
- the offered load value does not allow to fully characterize the actual performance of the algorithm. This is evident by noticing the efficiency of the traces with the same offered load but with different traffic distribution. Indeed, the achieved value of η_{Alg2} clearly depends on how this traffic is distributed within the time windows.

B. Performance Comparison of the EEEP Algorithms

In this subsection a performance comparison of the two EEEP algorithms described in Section III-C is carried out, in terms of energy consumption. Following an approach similar to that of the previous subsection, the figures of merit are here obtained by calculating the energy consumption (i.e. $E(U)$, $0 \leq U \leq 1$), that is the actual energy used during the whole simulation period L, as well as the average delays introduced by the two algorithms. The energy parameters of the sample switch devices, are those indicated in Table II, derived from [6], and the power consumption during the transitions of the link from active to quiet state and vice versa is considered equal to that of the active state. Also, in agreement with IEEE 802.3az, refresh signals have been taken into consideration as generated with period t_q and duration $t_r = t_w + t_s$. The time window duration was set to 50 ms for 1GBASE–T and to 5 ms for 10GBASE–T, respectively.

TABLE III
TRACES' MAIN PARAMETERS

Trace ID	f [Gb/s]	$\bar{d}$ [bit]	Offered load [%]	Source
R1G #A1–A5	1	from 5144 to 5160	from 6.17 to 7.96	WAND Research Group, Univ. of Waikato [36]
R1G #B1–B13	1	from 968 to 1056	from 0.50 to 0.60	CAIDA database, Univ. of UCSD [37]
R1G #C1–C4	1	from 3376 to 3816	from 21.94 to 27.86	WIDE project database [38]
S1G #A1–A200	1	from 8256 to 9912	from 22.84 to 45.12	–
S1G #B1–B200	1	from 3688 to 5600	from 2.10 to 13.4	–
R10G #A1–A36	10	from 4832 to 6032	from 15.28 to 21.54	CAIDA database, Equinix Chicago 2015 [39]
R10G #B1–B18	10	from 2912 to 3448	from 2.69 to 3.45	CAIDA database, Equinix Chicago 2011 [40]
R10G #C1–C18	10	from 4960 to 5792	from 30.75 to 41.70	CAIDA database, Equinix S. Jose 2014 [41]
S10G #A1–A200	10	from 10000 to 10008	from 25.00 to 45.04	–
S10G #B1–B200	10	from 5192 to 5224	from 3.12 to 18.72	–

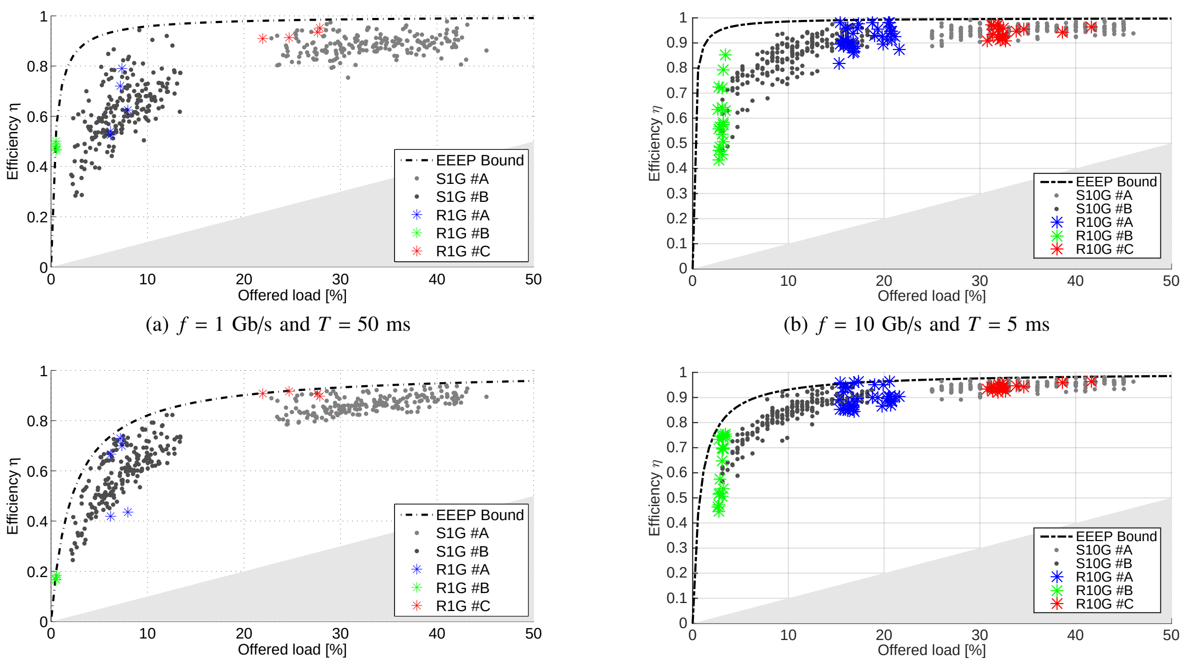

(a) f = 1 Gb/s and T = 50 ms

(b) f = 10 Gb/s and T = 5 ms

(c) f = 1 Gb/s and T = 10 ms

(d) f = 10 Gb/s and T = 1 ms

Fig. 5. Algorithm 2 performance: efficiency behavior for real and synthetic traces. Curves account for the theoretical limits and markers refer to single traces. The shaded area indicates efficiency levels below the Always-On policy.

To ease the presentation of the results, six traces extracted from the real dataset are considered, namely R1G #A1, R1G #B1, R1G #C1, for the 1GBASE-T link, and R10G #A1, R10G #B1, R10G #C1, for the 10GBASE-T.

A graphical performance comparison is shown in both Fig. 6 and Fig. 7, where two implementations of the EEEP strategy are considered, namely that described by Algorithm 1 with the ARMA(1,1) model and that referring to Algorithm 2. The choice of using ARMA(1,1) for Algorithm 1 is motivated by the fact that models of different orders (ARMA(2,1), ARMA(1,2) and ARIMA(1,1,1)) have shown to yield similar performance. Analogously, the employment of longer windows in the training phase of the ARMA models did not bring significant advantages in terms of prediction accuracy. Hence a simple solution is favorably chosen as the representative of the approach.

In particular, Fig. 6 reports (in percentage) the energy savings of EEEP and EEE with respect to the consumption of the Always ON policy. As can be observed, besides the considerable advantage systematically brought by the introduction of EEE, it is possible to verify that EEEP further increases (in some cases significantly) the energy savings and, also, how the performance figures of Algorithm 2 are very close to those of Algorithm 1.

On the other hand, Fig. 7 is concerned with the delay on packet delivery introduced by the two EEEP algorithms. Specifically, Fig. 7a shows the average delay value, whereas Fig. 7b reports its probability density function. As can be seen, the packet delay introduced by Algorithm 2 is generally lower, which can be intuitively explained by observing that,

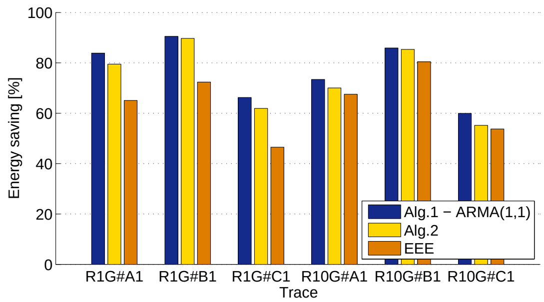

Fig. 6. Percentages of energy saving of EEE and EEEP (Algorithm 1 and Algorithm 2) with respect to the consumption of the Always-On policy.

with such an algorithm, the prediction is not systematically carried out at every time window. This is confirmed by the significant difference between the probability density plots of the two energy efficient algorithms: while Algorithm 1 presents delay values characterized by a quasi uniform probability distribution, that clearly vanishes towards the end of the window length (which represents the maximum delay), with Algorithm 2 an additional peak for packets transmitted with low delay appears, which accounts for the packets transmitted when prediction is not applied, and hence EEE is selected.

The combination of these results confirms the validity of the EEEP approach in general, and in particular the effectiveness of the implementation proposed by Algorithm 2, being at the same time efficient and robust.

As a general note, the results obtained via simulations are in a very good agreement with those derived from the analysis provided in Section IV. It is also worth noting that the impact of refresh signals on link power consumption results very limited (in the order of a few percent of the total energy), thus validating the approximation introduced by neglecting them in the theoretical study derivation. Consequently, the obtained theoretical bounds reveal to be useful for the design of effective transmission strategies.

The full statistics of the strategy are reported in [35].

C. Packet Delay Mitigation

As suggested by Remark 1, the increase of τ above its predicted value poses a trade off between the amount of energy that could be saved in a window, and the maximum delay that may affect packet delivery. Particularly, this parameter reveals effective in reducing the number of windows in which a wrong prediction would cause the transmission of packets to be moved to the following window.

A representation of the $p_{\bar{\tau}}$ impact is shown in Fig. 8 for the same traces and simulation parameters considered in Fig. 7.

In the top row (Figs. 8a–8b), the EEEP energy gain EG (with respect to EEE) is put in relation with the time windows in which prediction reveals correct, indicated as “non–delayed windows”. These plots clearly highlight the performance trade–off between the energy gain and the percentage of non–delayed windows: it is remarkable how this percentage rapidly tends toward 100% with the increase of $p_{\bar{\tau}}$, at the cost of a rather limited decrease of the energy gain, which, as evidenced by (15), has a linear trend following the increase of $p_{\bar{\tau}}$. Indeed, increasing $p_{\bar{\tau}}$ has the effect of activating the link in advance with respect to the time instant calculated by the prediction algorithm, with the consequent reduction of both the average packet delivery delay and the number of delayed windows.

In the bottom row (Figs. 8c–8d), the additional information about the average delay experienced by packets is provided. It can be seen that, as expected, such a delay decreases proportionally with the energy gain.

D. EEEP Strategy Additional Tests

In all the tests presented above, the EEEP strategy shows to work appropriately. Nonetheless, an additional session of tests has been carried out to further investigate the behavior of the strategy under stress conditions.

As far as Alg. 1 is concerned, the overall performance has been assessed with respect to the duration of the prediction window, which represents an important parameter of the strategy. On the other hand, for Alg. 2, the τ interval has been deliberately corrupted after prediction in each time window to abruptly emulate a prediction error scenario. In practice, τ has been replaced by $\tau' = \tau + g$ where g is a Gaussian zero mean noise with standard deviation $\sigma_g = e_\tau \cdot \tau/3$ and e_τ is a coefficient. If, for example, $e_\tau = 0.5$, then τ' ranges approximately between 0.5τ and 1.5τ.

The results of the simulations are presented respectively in Figs. 9a-9b. Concerning the non-delayed windows, in the first case, a slight decrease can be noted for some traces, only in correspondence to a strong reduction of the prediction window. In the second case, the uncertainty artificially added to the estimated prediction interval leads to a low reduction only for strong perturbations with respect to the noiseless case. Conversely, the EG does not show significant variations in any of the considered cases. The performance figures of the strategy are hence satisfactory also under the above described stress conditions.

Finally, since EEEP may increase the number of queued packets with respect to standard EEE, the behavior of the queue size has been punctually monitored during the tests. As a matter of fact, no packet drop due to queue overflow is observed and, moreover, the maximum queue sizes, as can be seen in Fig. 9c, reveals manageable by commercially available switches.

VI. Conclusions

In this work, EEE with Prediction (EEEP) has been introduced as an innovative energy efficient strategy for Ethernet networks. EEEP is based on the statistical analysis of current traffic for predicting and shaping the forthcoming traffic, with the aim of further increasing the energy savings. In this direction, a two–step algorithmic scheme (prediction model building and prediction exploitation) along with two different implementations have been proposed.

An original theoretical analysis has been derived, to define the performance bounds of both EEE and the novel EEEP. Moreover, an extensive simulation study, exploiting both real

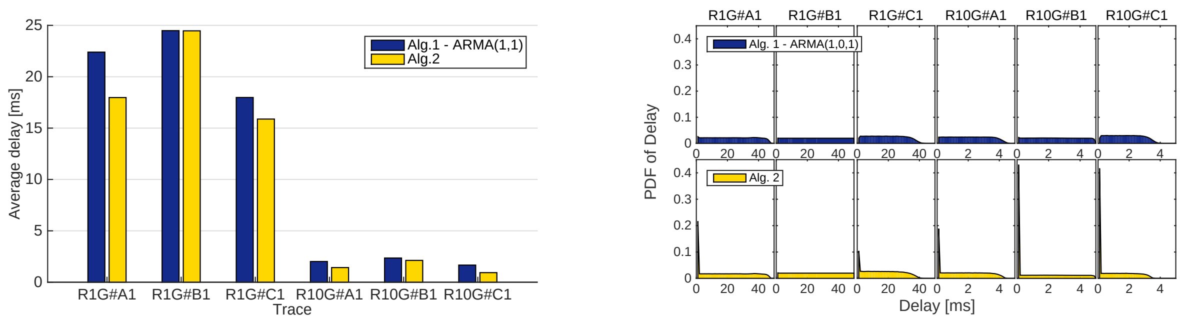

(a) (b)

Fig. 7. Delay on packet delivery introduced by the EEEP strategy. (a) Average value and (b) Probability density function.

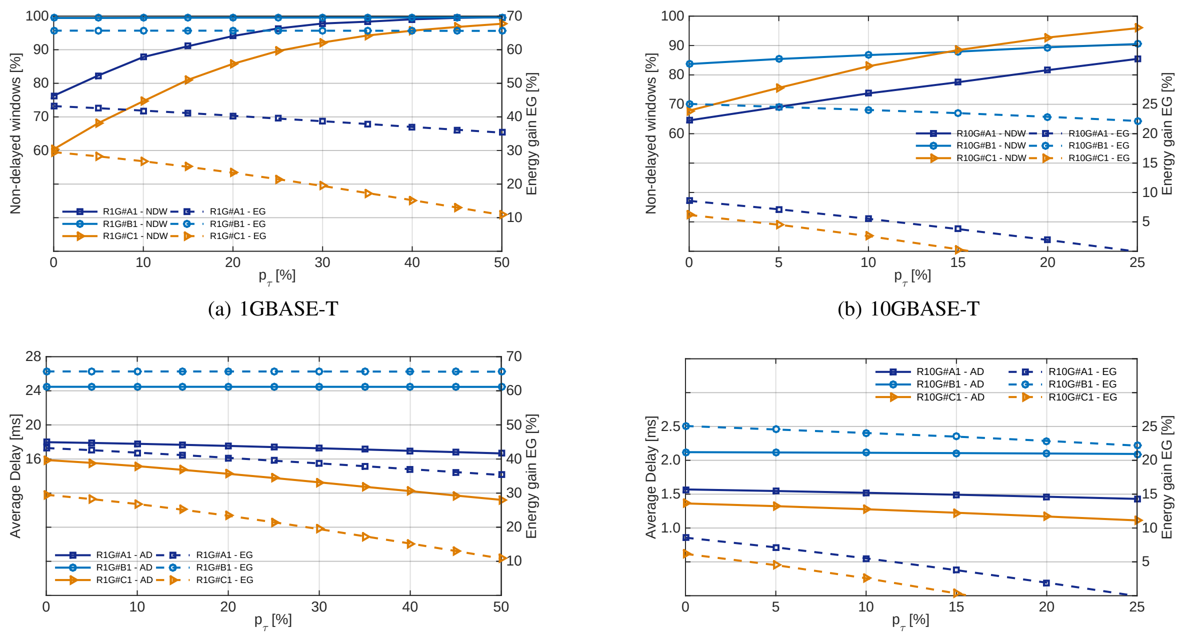

(a) 1GBASE-T (b) 10GBASE-T

(c) 1GBASE-T (d) 10GBASE-T

Fig. 8. Packet delay mitigation. (a)–(b): the solid and the dashed lines correspond to the percentage of non-delayed windows and to the energy gain, respectively. (c)–(d): the solid and the dashed lines correspond to the average delay and to the energy gain respectively.

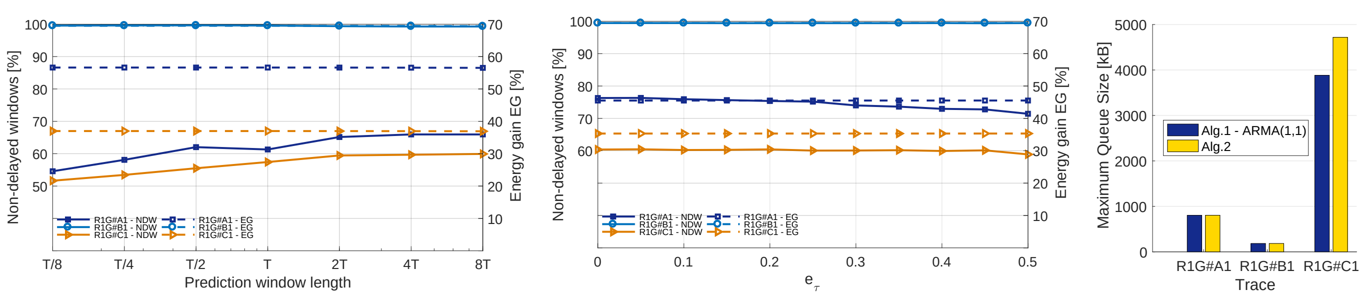

(a) Alg.1 - Prediction window length (b) Alg.2 - Prediction accuracy (c) Queue size

Fig. 9. Stress-test results. (a): Alg.1 performance while varying the prediction window length. (b): Alg.2 performance while adding a prediction noise. (c): Total queue size for both Alg.1 and 2.

and synthetic network traffic traces, has been carried out to provide an accurate assessment of the achievable performance.

The outcomes of these analyses highlight the good performance of EEEP, in that it yields a significant increase of energy savings with respect to the traditional EEE scheme, at the expense of only limited and controllable delays in packet delivery. Some meaningful future activities can be envisaged for this work. The first one is represented by the practical implementation of the proposed Algorithms in real Ethernet switches. This, in turn, would require that the activation/deactivation of the outgoing switch link can be independently triggered and controlled. Secondly, efforts should be devoted to better mathematically characterize the behavior of both the proposed Algorithms with respect to some important parameters, such as for example the conditional probability threshold, θ, and the number of quantization levels, h.

Finally, a more general scenario might be considered, such as a network encompassing several switches able to apply the EEEP strategy on their outgoing links. In this context it would be interesting to investigate whether smart cooperative policies can be devised, to allow the whole network reaching better global energy saving performance, through local interaction among neighboring switches.

Acknowledgement

The authors would like to thank Mr. M. Michielan for his valuable contribution during the initial phase of this study.

References

[1] *IEEE 802.3 standard: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications*, IEEE Std., October 2000.

[2] T. Sauter, S. Soucek, W. Kastner, and D. Dietrich, "The evolution of factory and building automation," *IEEE Ind. Electron. Mag.*, vol. 5, no. 3, pp. 35–48, Sept 2011.

[3] R. Bolla, R. Bruschi, A. Carrega, F. Davoli, and P. Lago, "A Closed-Form Model for the IEEE 802.3az Network and Power Performance," *IEEE J. Sel. Areas Commun.*, vol. 32, no. 1, pp. 16–27, January 2014.

[4] 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, April 2008.

[5] 10GBASE-T Ecosystem Is Ready for Broad Adoption. [Online]. Available: http://www.commscope.com

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

[7] M. Gupta and S. Singh, "Greening of the Internet," in *Proc. 2003 Conf. Applications, Technologies, Architectures, and Protocols for Computer Communications*, New York, NY, USA, 2003, pp. 19–26.

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

[9] K. Christensen, P. Reviriego, B. Nordman, M. Bennett, M. Mostowfi, and J. Maestro, "IEEE 802.3az: the road to Energy Efficient Ethernet," *IEEE Commun. Mag.*, vol. 48, no. 11, pp. 50–56, November 2010.

[10] D. Chakraborty, A. Ashir, T. Suganuma, G. Mansfield Keeni, T. K. Roy, and N. Shiratori, "Self-similar and fractal nature of internet traffic," *Intl. J. of Network Management*, vol. 14, no. 2, pp. 119–129, 2004.

[11] S. Vitturi and F. Tramarin, "Energy efficient ethernet for real-time industrial networks," *IEEE Trans. Autom. Sci. Eng.*, vol. 12, no. 1, pp. 228–237, Jan 2015.

[12] A. Cenedese, M. Michielan, F. Tramarin, and S. Vitturi, "An energy efficient traffic shaping algorithm for ethernet-based multimedia industrial traffic," in *Emerging Technology and Factory Automation (ETFA), 2014 IEEE*, Sept 2014, pp. 1–4.

[13] P. Reviriego, J. Maestro, J. Hernández, and D. Larrabeiti, "Burst transmission for Energy-Efficient Ethernet," *Internet Computing, IEEE*, vol. 14, no. 4, pp. 50–57, July 2010.

[14] J. Meng, F. Ren, W. Jiang, and C. Lin, "Modeling and understanding burst transmission algorithms for energy efficient ethernet," in *IEEE/ACM Intl Symp. Quality of Service (IWQoS)*, June 2013, pp. 1–10.

[15] A. Chatzipapas and V. Mancuso, "An M/G/1 Model for Gigabit Energy Efficient Ethernet Links With Coalescing and Real-Trace-Based Evaluation," *IEEE/ACM Trans. Netw.*, 2015.

[16] S. Herreria-Alonso, M. Rodriguez-Perez, M. Fernandez-Veiga, and C. Lopez-Garcia, "A GI/G/1 model for 10 Gb/s energy efficient ethernet links," *IEEE Trans. Commun.*, vol. 60, no. 11, pp. 3386–3395, 2012.

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

[18] P. Reviriego, J. Hernández, D. Larrabeiti, and J. Maestro, "Performance evaluation of Energy Efficient Ethernet," *IEEE Commun. Lett.*, vol. 13, no. 9, pp. 697–699, Sept 2009.

[19] P. Reviriego, K. Christensen, J. Rabanillo, and J. Maestro, "An initial evaluation of Energy Efficient Ethernet," *IEEE Commun. Lett.*, vol. 15, no. 5, pp. 578–580, May 2011.

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

[21] F. Tramarin and S. Vitturi, "Strategies and Services for Energy Efficiency in Real–Time Ethernet Networks," *IEEE Trans. Ind. Informat.*, vol. 11, no. 3, pp. 841–852, June 2015.

[22] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the self-similar nature of Ethernet traffic (extended version)," *IEEE/ACM Trans. Netw.*, vol. 2, no. 1, pp. 1–15, February 1994.

[23] M. E. Crovella and A. Bestavros, "Self-similarity in world wide web traffic: Evidence and possible causes," *IEEE/ACM Trans. Netw.*, vol. 5, no. 6, pp. 835–846, December 1997.

[24] S. Song, J.-Y. Ng, and B. Tang, "Some results on the self-similarity property in communication networks," *IEEE Trans. Commun.*, vol. 52, no. 10, pp. 1636–1642, Oct 2004.

[25] A. Bianco, V. Krishnamoorthi, N. Li, and L. Giraudo, "Openflow driven ethernet traffic analysis," in *Communications (ICC), 2014 IEEE International Conference on*, June 2014, pp. 3001–3006.

[26] S. Stoev, G. Michailidis, and J. Vaughan, "On global modeling of backbone network traffic," in *INFOCOM, 2010 Proceedings IEEE*, March 2010, pp. 1–5.

[27] B. Bing, *3D and HD Broadband Video Networking*, 1st ed. Norwood, MA, USA: Artech House, Inc., 2010.

[28] H. Hassan, J. Garcia, and C. Bockstal, "Aggregate traffic models for voip applications," in *Digital Telecommunications, 2006. ICDT '06. International Conference on*, Aug 2006, pp. 70–70.

[29] M. Gupta, S. Grover, and S. Singh, "A feasibility study for power management in lan switches," in *Proc. 12th IEEE International Conference on Network Protocols (ICNP)*, Oct 2004, pp. 361–371.

[30] M. Gupta and S. Singh, "Using low-power modes for energy conservation in ethernet lans," in *26th IEEE International Conference on Computer Communications (INFOCOM)*, May 2007, pp. 2451–2455.

[31] K. Papagiannaki, N. Taft, Z.-L. Zhang, and C. Diot, "Long-term forecasting of internet backbone traffic: observations and initial models," in *22nd IEEE International Conference on Computer Communications (INFOCOM)*, vol. 2, March 2003, pp. 1178–1188 vol.2.

[32] M. F. Zhani, H. Elbiaze *et al.*, "Analysis and prediction of real network traffic," *Journal of Networks*, vol. 4, no. 9, pp. 855–865, Nov 2009.

[33] S. Xu and B. Zeng, "Network traffic prediction model based on autoregressive moving average," *Journal of Networks*, vol. 9, no. 3, pp. 653–659, Mar 2014.

[34] 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.

[35] A. Cenedese and F. Tramarin and S. Vitturi, "Research Report: An Energy Efficient Ethernet Strategy Based on Traffic Prediction and Shaping – Full Statistics." [Online]. Available: http://www.cnr.it/peoplepublic/peoplepublic/index/downloadprodotto/i/110154

[36] WAND Network Research Group – University of Waikato (NZ), "Waikato Internet Traffic Storage." [Online]. Available: http://wand.net.nz/wits/index.php-http://wand.net.nz/wits/waikato/1/waikato_i.php

[37] CAIDA – University of California, San Diego (CA), "The CAIDA UCSD Network Telescope "Two Days in November 2008 Dataset." [Online]. Available: http://www.caida.org/home/-http://www.caida.org/data/passive/telescope-2days-2008_dataset.xml

[38] MAWI Working Group, "Packet traces from WIDE backbone." [Online]. Available: http://mawi.wide.ad.jp/mawi/ditl/ditl2013

[39] "The CAIDA Anonymized Internet Traces Dataset - Passive Monitor Equinix Chicago (IL) – Feb. and Sept. 2015." [Online]. Available: http://www.caida.org/data/monitors/passive-equinix-chicago.xml

[40] "The CAIDA Anonymized Internet Traces Dataset - Passive Monitor Equinix Chicago (IL) – Feb. 2011." [Online]. Available: http://www.caida.org/data/monitors/passive-equinix-chicago.xml

[41] "The CAIDA Anonymized 2014 Internet Traces Dataset - Passive Monitor Equinix San Jose (CA) – Mar. 2014." [Online]. Available: http://www.caida.org/data/monitors/passive-equinix-sanjose.xml