

Augmenting the Energy-Saving Impact of IEEE 802.3az via the Control Plane

Chakadkit Thaenchaiakun, Gentian Jakllari, Beatrice Paillassa
IRIT-ENSEEIH, University of Toulouse, France
{chakadkit.thaenchaiakun, jakllari, beatrice.paillassa}@enseeiht.fr

Abstract—IEEE 802.3az, the recent standard for Energy Efficient Ethernet, is one of the main contributions of the ICT industry to the global quest for energy efficiency. Energy consumption reduction is accomplished by essentially replacing the continuous IDLE of legacy IEEE 802.3 cards with a Low Power Idle. While this is an important step in the right direction, studies have shown that the energy saving with IEEE 802.3az highly depends on the traffic load and stops for link utilizations as low as 25%. In addition to the standardization effort for creating energy efficient networking hardware, several solutions have been proposed for energy efficient routing, OSPF in particular. However, most solutions are transparent to and usually leave open the question of how exactly are the energy savings realized in hardware. Simply routing less or no traffic through certain links will not automatically lead to those links consuming less energy.

We present EAGER and CARE, two OSPF metrics customized for augmenting the energy saving of IEEE 802.3az line cards – the only available off-the-shelf, standard-complying hardware for saving energy in OSPF networks. The principle underlying the design of EAGER and CARE is that of creating a synergy between the control plane of OSPF routers and IEEE 802.3az compliant line cards, without requiring changes to the respective standards and/or the hardware. Simulation results show that EAGER and CARE can increase the amount of energy saved when using IEEE 802.3az line cards by $2\times$ on a network of routers. Furthermore, the results show that CARE strikes a good balance between energy efficiency and traditional network performance metrics.

I. INTRODUCTION

Recent years have witnessed a boom in the pursue of energy efficient network protocols and technologies. A lot of effort has been put into making energy efficient two of the more popular and established technologies: Ethernet [1], [2], [3], with the efforts culminating in a new standard, IEEE 802.3az [4], for Energy Efficient Ethernet¹; and OSPF [5].

IEEE 802.3az reduces the energy consumption of line cards by adapting a Low Power Idle mechanism for replacing the continuous IDLE of legacy Ethernet cards. When there is no traffic, the card spends most of the time in Low Power mode, with few interruption for transmitting refresh signals. However, once there is even a single packet to transmit, the card switches back to active mode. This conservative approach, designed to avoid disruption to the higher layers, can lead to IEEE 802.3az line cards spending most of their time in active mode. As shown by the data in Figure 1, collected from simulations on ns-3 described in more detail in Section IV, the energy consumption of an IEEE 802.3az compliant line card grows sharply with traffic load and reaches 100% for as little as 24% link utilization. Packet coalescing [6] tries to address this weakness of the standard by adding a packet buffer for grouping several packets together so as to avoid waking up the card for

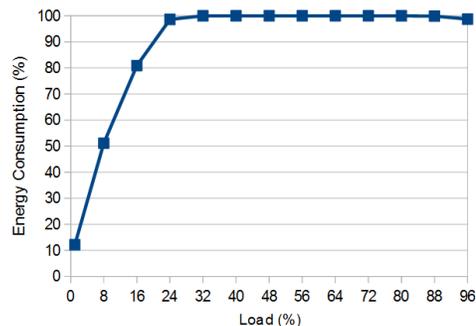


Fig. 1: The energy consumption with IEEE 802.3az increases sharply as the link load increases, matching the always-on legacy Ethernet at 24% link load.

as little as a single packet. However, it can have unintended consequences on the TCP congestion window if the coalesced packets are TCP ACKs.

In parallel, several works propose reducing the energy consumption of IP networks by modifying OSPF [7], [8], [9], [10]. The general approach consists of routing traffic over fewer routers and links than what OSPF does by default. The non-utilized routers and/or links are then to be shut down for saving energy. However, the questions of how routers and/or links can be shut down, by which mechanisms and what the implications to OSPF are², are usually left open.

The case for layer synergy: Considering that the energy efficiency of off-the-shelf IEEE 802.3az hardware is highly dependent on how traffic is routed by OSPF and, in turn, any OSPF based solution is dependent on lower layer mechanisms, where the energy saving is actually realized, we believe there is a case for synergy between the two popular protocols. Specifically, the basic question we address in this work is whether it is possible to significantly increase the energy efficiency of networks using IEEE 802.3az line cards by making minimal, standard compliant changes to the control plane of the router software.

We present EAGER (for Extremely AuGmented Energy efficient eNet) and CARE (for Congestion aware Augmented eneRgy efficient Ethernet), two OSPF link metrics customized for augmenting the energy-saving impact of IEEE 802.3az. The basic principle underlying the design of EAGER and CARE is that of assigning link costs that somehow mirror the energy consumption of IEEE 802.3az. Links that have already reached 100% energy consumption, which as shown in Fig. 1 can occur

¹For the rest of the paper, we will use the terms EEE and IEEE 802.3az interchangeably.

²OSPF for example cannot distinguish between a router failing and a router going to sleep mode.

for links having as little as 24% utilization, are perfectly capable of carrying additional traffic – at zero additional energy cost. These links are favored by EAGER and CARE. On the other hand, any additional traffic can have a dramatic impact on the energy consumption of lightly loaded links – increasing the link load from 0% to 8% increases the energy consumption by 5x. These links are less preferred by EAGER and CARE.

EAGER is the simplest approach and is designed to aggressively pursue energy efficiency. CARE is more conservative and is designed to offer the possibility of choosing the best tradeoff between energy efficiency and traditional network performance metrics, such as congestion and end-to-end delay.

Simulation results show that EAGER and CARE can augment the energy saving impact of IEEE 802.3az by a factor of 2 on a network of routers. While at high loads EAGER is shown to suffer from increased congestion and end-to-end delay, CARE manages to maintain levels of congestion and end-to-end delay very similar to the shortest path.

The rest of the paper is organized as follows. In Section II, we discuss related work. In Section III, we present EAGER and CARE. In Section IV, we describe the implementation of IEEE 802.3az, EAGER and CARE and evaluate their performance. We conclude in Section V.

II. RELATED WORK

While the basic paradigm underlying most proposed solutions for energy efficient networking is that of setting to “sleep” mode as many networking components as possible, the approaches in literature can be largely organized in two groups: solutions targeted at the link layer, and solutions targeted at the routing layer.

Energy Efficiency at the Link Layer: Ethernet, as the dominant wireline technology for LANs, has been the main focus of the research community as well as industry [1], [2], [3]. The efforts culminated in the creation of a new standard for Energy Efficient Ethernet - IEEE 802.3az - ratified in September 2010 [4]. The standards committee considered two strategies for saving energy, Adaptive Link Rate and Low Power Idle, and opted for the latter [6]. The Low Power Idle replaces the continuous IDLE of legacy Ethernet cards when there is no data to transmit. It defines large time intervals during which no signal is transmitted and the interface is set to “sleep”, interrupted by short periods during which a signal is transmitted to refresh the receiver state to align it with current conditions. The card exits Low Power Idle when it needs to transmit or receive data, even if it is a single packet. This simple and conservative approach enables EEE cards to reduce energy consumption when possible while keeping any potential disruption to the higher layers to a minimum. It works well when a large part of the network is severely underutilized or traffic is very bursty. However, as performance evaluation studies [11], including our own simulation data from Figure 1, have shown, even light loads can force EEE cards to stay in active mode most of the time, leading to no energy savings. This weakness of the standard can be addressed by introducing a buffer that will coalesce packets together [6] and forward them to the network card only once they reach a certain number. Nevertheless, coalescing can lead to TCP ACK compressing which can cause sudden spikes in TCP congestion window and lead to buffer overflows in downstream routers [6].

Furthermore, as a non-standard solution, implementing it would require manufacturers to patch the hardware and/or software of their cards.

Energy Efficiency at the Routing Layer: In parallel with the work on improving the energy efficiency of layer-2 technologies, there is a very active effort on improving energy efficiency via the routing layer [12], [13], [8], [9], [14], [15], [16], [17], [7], [5]. The general approach consists of leveraging the topology knowledge at the routing layer to either identify lightly used links or use traffic engineering to create them. Once this is done, routers on lightly utilized links can be powered off for saving energy. As with the layer two solutions, most efforts are focused on the most widely utilized routing protocol in IP networks, OSPF. An exhaustive survey is beyond the scope of this paper so here we describe some representative approaches. ESOL [8], for example, leverages OSPF’s link state updates (LSAs) for identifying the nodes and links that appear the least on shortest path trees. A heuristic considers the least-utilized links and starts removing them one at a time while making sure that the end-to-end connectivity is maintained. EAR [9] divides OSPF nodes into exporters, who compute shortest path trees, and importers, who do not compute shortest path trees but use those computed by the exporters. Having only exporter nodes compute shortest path trees for routing data, unlike pure OSPF where all nodes do so, can reduce the number of links over which data is carried – allowing more links to be powered off for saving energy. Similarly, GAES [13], first identifies a set of nodes/links that are lightly utilized and then recomputes shortest paths trees as if though they did not exist. If the given traffic matrix can be supported by the reduced set of links, the rest can be set to sleep mode. ECO-RP [10] is another OSPF based solution, which uses what the authors refer to as traffic trending for modifying link weights and thus the subset of routers that carry traffic. For example, when the traffic is trending down, link weights are modified so that the traffic can be routed by a smaller number of routers, leaving the rest idle. In [7], several traffic-engineering heuristics are proposed for the problem of maximizing the number of links that can be put to sleep mode while satisfying given performance constraints, such as link utilization and packet delay.

The works summarized so far and references therein present a rich set of solutions for making routing in general, and OSPF in particular, more energy friendly. Nevertheless, before some of these solutions can be adapted by the industry, there are certain issues that need to be addressed. For example, it is not clear by what mechanism routers and/or links will be powered off and whether turning off an entire router, as recommended by some solutions, is something certain organizations will even embrace. And more fundamentally, it is not clear if designing layer-3 solutions without taking into consideration the specifics of how saving energy is actually realized in hardware will lead to the best approach for saving energy.

III. AUGMENTING IEEE 802.3AZ VIA THE CONTROL PLANE

In this section we present EAGER and CARE, two new routing metrics for augmenting the energy saving impact of IEEE 802.3az hardware.

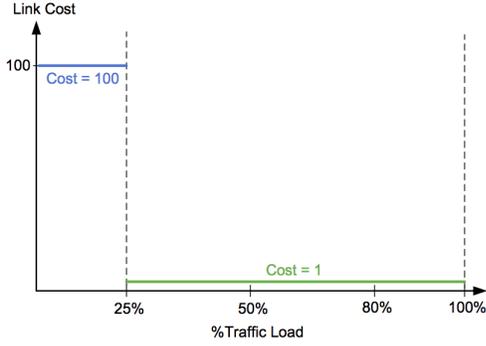
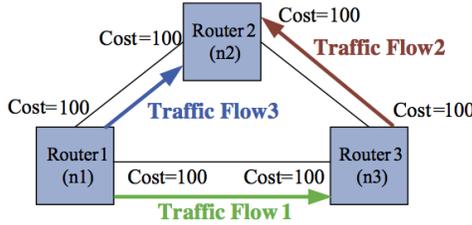
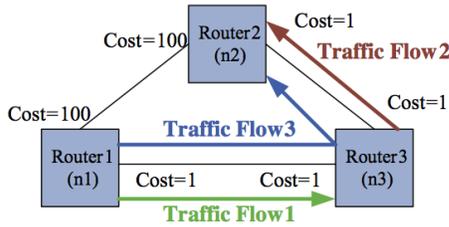


Fig. 2: EAGER: Extremely AuGmented Energy efficient eHeRnet. IEEE 802.3az realizes no energy reduction on links over the green threshold of 25%. Adding more traffic on these links has no effect on energy consumption so they are assigned the flat, minimum cost of 1.



(a) 160 Mbps Traffic Flows – Per flow load below the green threshold of 25%.



(b) 500 Mbps Traffic Flows – Per flow load over the green threshold of 25%.

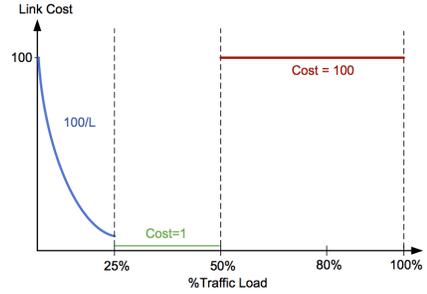
Fig. 3: Illustration of EAGER on a simple topology. All routers use 1000Base-T line cards. Flow 1 arrives first, followed by Flows 2 and 3. In the high load scenario, when Flow 3 arrives, link (n_1, n_2) is idle so its cost is 100. The other two links each have a flow whose arrival rate puts both over the green threshold and their costs to 1. Therefore, Flow 3 is routed over Router 3 on the way to the destination, allowing link (n_1, n_2) to sit idle and the respective IEEE 802.3az line cards to operate in Low Power Idle.

A. EAGER: Extremely AuGmented Energy efficient eHeRnet

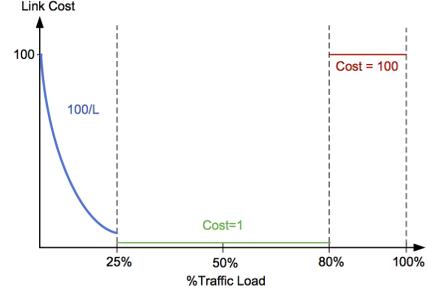
As stated previously, the energy efficiency of line cards implementing IEEE 802.3az highly depends on the traffic load. As shown in Figure 1 from Section I, the biggest energy savings are realized at very low loads and drop fast as the load increases, becoming practically non-existent once the link utilization reaches around 25%³. We refer to this value of the link utilization as the *green threshold*.

Our inspiration in designing EAGER is to create a link metric that mirrors the energy curve from Figure 1. Lightly utilized links are assigned a high cost, while links whose utilization is over the *green threshold* are assigned a minimum, flat cost. This approach is visualized in Figure 2. More formally, given a

³The figure shows 24% but we round it up to 25%.



(a) CARE-50.



(b) CARE-80

Fig. 4: CARE: Congestion aware Augmented eneRgy efficient Ethernet. The part of the cost function from 0% to the green threshold of 25% is designed to mirror the energy saving curve of IEEE 802.3az. To strike a balance between energy efficiency and congestion, links over the congestion threshold are also assigned a high cost.

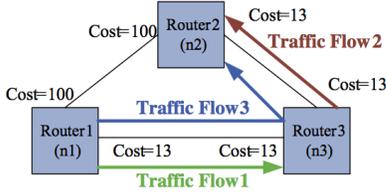
link (i, j) , with link utilization value of $L_{i,j}$, its EAGER value can be computed as follows:

$$\text{EAGER}_{i,j} = \begin{cases} 100, & \text{if } L_{i,j} \leq 25\% \\ 1, & \text{Otherwise} \end{cases} \quad (1)$$

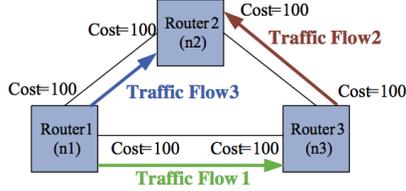
Figure 3 illustrates how EAGER works. When Flow 3, with destination Router 2, arrives at Router 1, the previous two flows have already modified the link utilization values in the network. As all routers are using 1000Base-T line cards, in the Low Load scenario, Figure 3(a), the utilization on all links is under the green threshold and thus their costs are set to the maximum 100. This leads Flow 3 through the shortest path. In the High Load scenario, Figure 3(b), when Flow 3 arrives, the link (n_1, n_2) is idle and its cost is 100, while the other two links are over the green threshold and their respective costs are set to 1. This leads Flow 3 through Router 3, instead of the shortest path, allowing the 802.3az cards on link (n_1, n_2) to stay under the green threshold and save the most energy.

B. CARE: Congestion aware Augmented eneRgy efficient Ethernet

EAGER is a simple and yet aggressive approach for augmenting the energy saving impact of IEEE 802.3az hardware. This aggressiveness in pursuing energy savings, however, could be costly in terms of congestion and end-to-end delay for certain links. CARE, as visualized in Figure 4, takes a more conservative approach by changing EAGER in two ways. First, in addition to the green threshold, it uses a *congestion threshold* – the link load value expressed in percentage over which the



(a) 160 Mbps Traffic Flows – Per flow load below the green threshold of 25%.



(b) 500 Mbps Traffic Flows – Per flow load over the green threshold of 25%.

Fig. 5: Illustration of CARE-50 on a simple topology. All routers use 1000Base-T line cards. Flow 1 arrives first, followed by Flows 2 and 3. Even for such a simple topology we see a very different behavior between EAGER and CARE. Unlike EAGER, CARE does not treat all light loads the same, which explains why for the Low Load regime Flow 3 here is routed through Router 3. On the other hand, CARE assigns the same, high cost to all links in the High Load scenario and Flow 3 is routed through the shortest path.

link is considered as being congested. A link with utilization over the congestion threshold is assigned the same weight as an idle link. Second, for the part under the green threshold, CARE assigns the multiplicative inverse of the link utilization, upper bounded by a maximum value, to better mirror the energy saving function of IEEE 802.3az. More formally, given a link (i, j) , with link utilization percentage value of $L_{i,j}$, and given a congestion threshold of X , the CARE- X value of this link can be computed as follows:

$$\text{CARE-}X_{i,j} = \begin{cases} \min(100, 100/L_{i,j}), & \text{if } L_{i,j} \leq 25\% \\ 1, & \text{if } 25\% < L_{i,j} \leq X\% \\ 100, & \text{if } L_{i,j} > X\% \end{cases} \quad (2)$$

Finding the best congestion threshold, X , for CARE is not simple as that depends on many factors, including manufacturer specifications and customer needs. In this work, we use 50% and 80% ([18], [19]) and have observed very promising results in our simulation study in Section IV. However, our recommendation is that any implementation of CARE should make the congestion threshold a parameter that can be modified by network administrators.

Figure 5 illustrates how CARE works. When Flow 3, with destination Router 2, arrives at Router 1, the previous two flows have already modified the link utilization values in the network. As all routers are using 1000Base-T line cards, in the Low Load scenario, Figure 3(a), the utilization on all links is under the green threshold and, thus, their respective costs are set by the multiplicative inverse of the link utilization (upper bounded by 100). This leads Flow 3 through Router 3. In the High Load

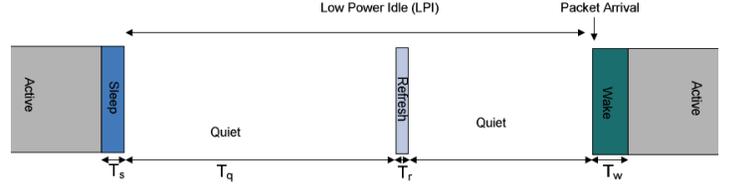


Fig. 6: Low Power Idle Function in IEEE 802.3az.

Transmitter	Power Level (Watt)
TX_READY	1
TX_BUSY	2
TX_SLEEP	0,1
TX_UP	2
TX_DOWN	2

Receiver	Power Level (Watt)
RX_READY	1
RX_BUSY	1,3
RX_SLEEP	0,1
RX_UP	1,3
RX_DOWN	1,3

Technology	T_w (msec)	T_s (msec)	T_q (msec)	T_r (msec)
1000Base-T	0.016	0.18	20	0.18

TABLE I: IEEE 802.3az Parameters Adapted in the Simulations

scenario, Figure 3(b), when Flow 3 arrives, the link (n_1, n_2) is idle, while the other two links are over the congestion threshold, leading to all links having the maximum cost. This leads Flow 3 through the shortest path.

IV. PERFORMANCE EVALUATION

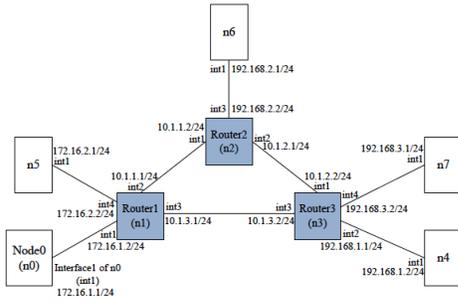
We use ns-3 [20] as simulation platform, in which we have added an implementation of the IEEE 802.az standard, and evaluate the performance of EAGER and CARE in terms of energy savings, end-to-end delay and load distribution. In all experiments, the values of energy savings are given relative to a system that implements no energy efficiency mechanism (traditional Ethernet).

A. Software Implementation

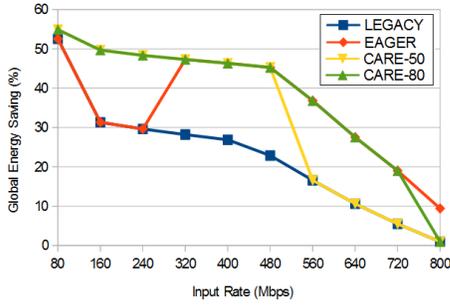
To carry out the performance evaluation study we had to make several additions to ns-3. We added an implementation of IEEE 802.3az that follows the standard specification [4], modified the OSPF implementation to utilize EAGER and CARE, and implemented layer-2 packet coalescing [6].

IEEE 802.3az: The basic IEEE 802.3az operation is illustrated in Figure 6. When there are packets to transmit, the device is in Active Mode. Once there are no more packets to transmit, the device switches to sleep mode, a process that takes the hardware time T_s to execute. In sleep mode, the device only sends signals during short refresh intervals, T_r , and stays quiet during large intervals, T_q . When traffic arrives, the device switches back to Active Mode, a process that takes the hardware time T_w to execute.

Obviously, the energy efficiency is a function of how long a device spends in Low Power Idle, the timers T_s , T_q , T_r and T_w , and the energy saved during the quiet periods. Table I shows the specific values for the timers as defined in the IEEE 802.3az standard for 1000Base-T, the technology used throughout this study. However, the standard does not specify any values for the power levels IEEE 802.3az line cards should use when sleeping,



(a) Three flows are generated on a 3-router topology. Flow 1 from n_0 to n_4 , Flow 2 from n_7 to n_6 , Flow 3 from n_5 to n_6 . All routers and hosts are connected by 1000Base-T links.



(b) The energy saving is computed relative to a system that implements no energy efficiency mechanism (traditional Ethernet).

Fig. 7: Experiment 1: At low to medium loads, EAGER and CARE route Flow 3, which arrives last, through router 3 instead of the shortest path. This lead to roughly a 2x augmentation in energy saving over LEGACY OSPF. As designed, at high loads, over 50% and 80%, respectively, CARE-50 and CARE-80 fall back to LEGACY OSPF. EAGER on the other hand does not change its behavior and leaves the link from Router 1 to Router 2 idle even at high loads.

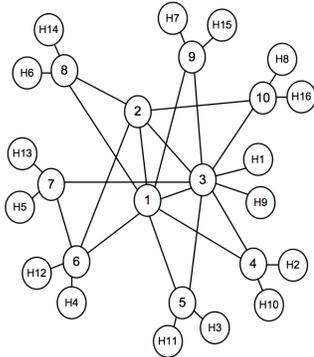


Fig. 8: Experiment 2: A 10-router topology following the Barabasi-Albert model using the BRITE [21] network topology generator. There are two hosts connected to the eight “edge” routers. All routers and hosts are connected by 1000Base-T links.

transmitting, receiving, in Active Mode, switching off, waking up, etc. The only thing the standard specifies is that during sleep mode the power consumption should be 10% of that in Active Mode. Using this recommendation and the power levels from real Cisco Switches reported in [2], we put together the energy model used in this study as shown on Table I.

EAGER and CARE: We modified the OSPF ns-3 implementation to use EAGER, CARE-50 and CARE-80. This was

not complicated as we just needed to add link utilization levels to the link state advertisements and modify the shortest path implementation to incorporate the new metrics.

Packet Coalescing: We implemented layer-two packet coalescing as defined in [6]. In short, packet coalescing uses an additional queue for accumulating multiple packets before forwarding them to the line card, thus creating a burst of back-to-back packets. The algorithm is driven by two parameters: the maximum queue size and a timer which defines the maximum amount of time a packet can be coalesced for in the queue. Once the maximum queue size is reached or the timer expires, whichever comes first, all the coalesced packets are released for transmission. In our implementation we use the parameters for Coalesce-2 from [6], but do not reproduce here for lack of space.

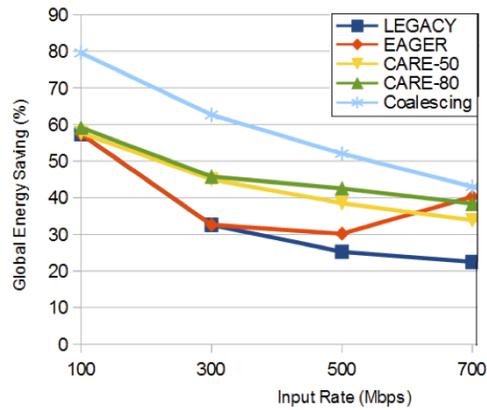
B. Experiment 1: Simple Topology

Setup: We start the performance evaluation with the simple 3-router topology, Figure 7(a), we first used in Section III. While the topology is trivial, it is useful for demonstrating the potential for augmenting the energy saving impact of IEEE 802.3az via the control plane. Three traffic flows are generated: Flow 1 from n_0 to n_4 , Flow 2 from n_7 to n_6 , Flow 3 from n_5 to n_6 . The flow arrivals are separated by enough time, with Flow 1 arriving first and Flow 3, last, so that an OSPF route update takes place between consecutive flow arrivals. All routers are connected by 1000Base-T links. EAGER and CARE are compared to the default OSPF, which we refer to as LEGACY, in terms of energy saving over a system that runs OSPF over pure Ethernet.

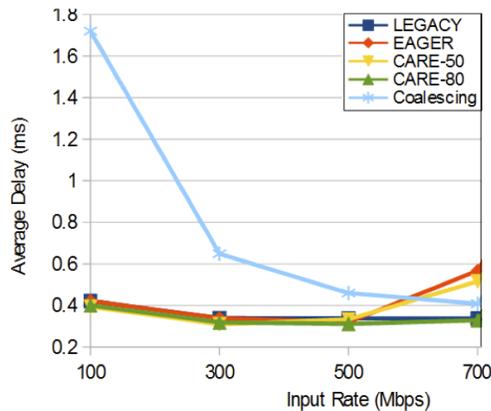
Results: Figure 7(b) shows that IEEE 802.3az can significantly reduce the energy consumption even in such a simple topology. However, this benefit disappears as the load increases leading to fewer opportunities for the line cards to switch to low power mode. EAGER and CARE, by routing the traffic for Flow 3 over links that are already utilized (via Router 3 instead of the direct link) augment the percentage of energy saved by almost 2x. Finally, as the traffic load increases, and as designed, first CARE-50 and then CARE-80, fall back to LEGACY OSPF. EAGER, on the other hand, keeps routing Flow 3 over Router 3 even though the latter is congested, allowing the two line cards on the link from Router 1 to Router 2 to switch to low power mode.

C. Experiment 2: Large Topology

Setup: In this experiment we evaluate the performance of EAGER and CARE on a more realistic topology. Towards this, we use the BRITE [21] network topology generator to create a 10-router network, depicted in Figure 8, following the Barabasi-Albert [22] model. Two hosts are attached to every router, except for routers 1 & 2 which are dedicated to forwarding traffic for other routers only. All routers and hosts are connected by 1000Base-T links. Every host can generate traffic flows following a Poisson process to every other host. Out of the $16 \times 15 = 240$ possible source-destination pairs we select a large number, 100, and run UDP sessions on each pair. In every experiment we measure the percentage of energy saved compared to a system running pure Ethernet, the traffic distribution and end-to-end delay. For this part of the evaluation



(a) At very low link utilization EAGER and CARE behave no differently than LEGACY OSPF. Layer-two packet coalescing on the other hand benefits from grouping packets together. As the load increases we see EAGER and CARE augment the energy saving impact of pure IEEE 802.3az (LEGACY) by roughly 1.5x to 2x.



(b) The average end-to-end delay for EAGER, CARE, LEGACY OSPF and packet coalescing. The data shows that CARE-80 can save energy without incurring a penalty in end-to-end delay.

Fig. 9: Experiment 2: The data shows CARE-80 offering the best overall performance. It augments the energy saving of pure IEEE 802.3az (LEGACY) by as much as 2x without incurring a significant penalty in end-to-end delay.

we compare not only to LEGACY OSPF but also to layer-two packet coalescing [6].

Results: Figure 9(a) shows that as the load increases, EAGER and CARE augment the energy saving impact of pure IEEE 802.3az (LEGACY) by roughly 1.5x to 2x.

At very low loads, EAGER and CARE behave no differently than LEGACY OSPF, while layer-two packet coalescing, benefiting from grouping packets together, significantly improves the energy efficiency. Unfortunately, this improvement of packet coalescing comes with a high penalty in end-to-end delay, as Figure 9(b) shows. At low loads, packets have to wait longer before enough of them are coalesced or in the worst case until the timer expires, leading to a significant increase in end-to-end delay. In contrast, the average end-to-end delay for EAGER and CARE is very close to LEGACY OSPF for most load regimes. At high loads EAGER starts suffering from higher end-to-end delay due to its aggressiveness in pursuing energy savings. Also, CARE-50 does not fare as well as expected because it overreacts to what it perceives as congestion. CARE-80, on the

other hand, is shown to always perform closely to LEGACY OSPF.

Overall, the data from Fig 9 points to CARE-80 being the better choice - it offers the energy savings benefit of layer-two coalescing without its negative impact on delay.

For a more in-depth understanding of the layer-3 solutions, we show the CDFs for the maximum link utilization in the network and end-to-end delay for every packet delivered in Figure 10 and Figure 11, respectively. Figure 10 shows that EAGER and CARE create a significant number of links with no traffic, explaining their superior performance in terms of energy efficiency. CARE-80, and to a lesser extent CARE-50, is shown to handle congestion very well while, as expected, with EAGER a significant number of links reach saturation at very high loads (Fig. 10(c)). Finally, Figure 11, shows that CARE-80 performs very close to LEGACY in terms of end-to-end delay while with EAGER, and to a lesser extent CARE-50, for very high loads (Figure 11(c)), a percentage of packets take long to reach the destination.

V. CONCLUSION

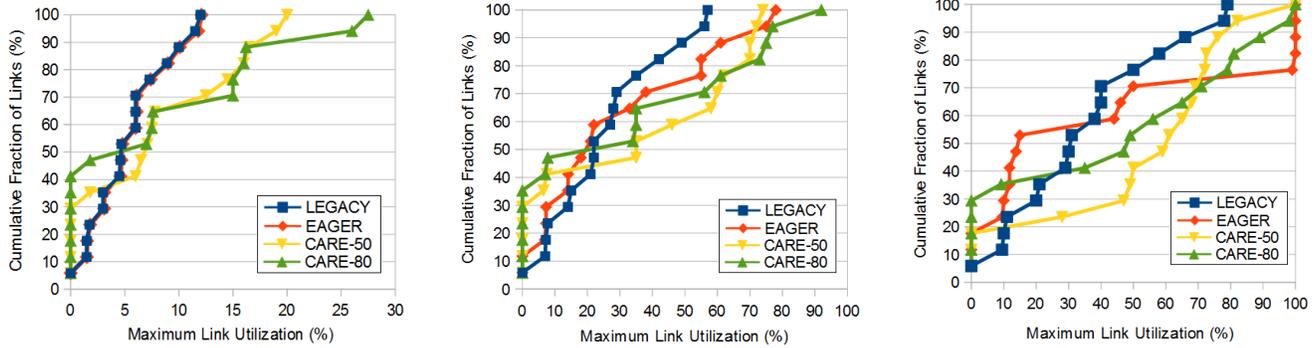
We proposed two OSPF metrics, EAGER and CARE for augmenting the energy saving impact of IEEE 802.3az. EAGER and CARE are founded on the principle of layer synergy between OSPF and IEEE 802.3az. OSPF provides the traffic engineering for maximizing the benefit of IEEE 802.3az, while IEEE 802.3az provides the mechanisms by which the energy efficiency is actually realized in practice. EAGER and CARE do not require any changes to the standards and can be implemented on off-the-shelf hardware.

Our simulation analysis showed that EAGER and CARE can augment the energy saving impact of IEEE 802.3az by 2x. While EAGER was shown to suffer from increased congestion and end-to-end delay at high loads, CARE with congestion threshold of 80% was shown to have the best overall performance.

As future work we intend to implement EAGER and CARE on real hardware and evaluate their performance on a realistic testbed.

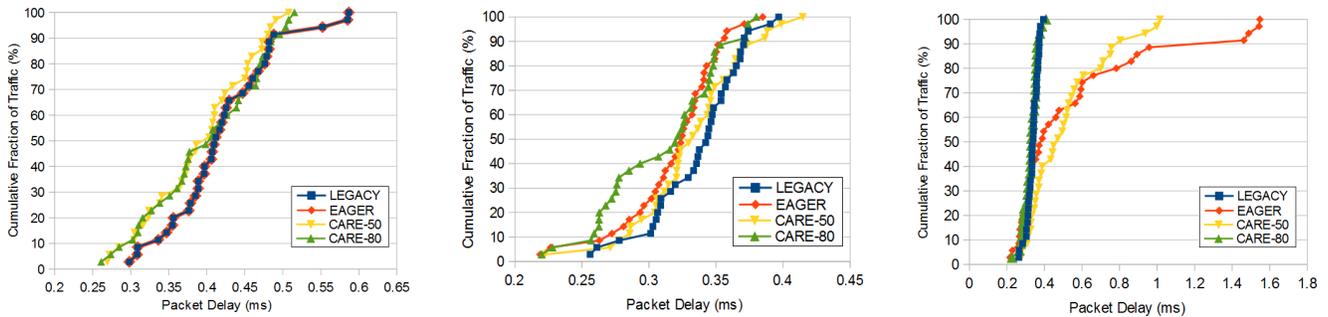
REFERENCES

- [1] M. Gupta, S. Grover, and S. Singh, "A feasibility study for power management in LAN switches," in *Network Protocols, 2004. ICNP 2004. Proceedings of the 12th IEEE International Conference on*, Oct 2004, pp. 361–371.
- [2] M. Gupta and S. Singh, "Using Low-Power Modes for Energy Conservation in Ethernet LANs," in *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*, May 2007, pp. 2451–2455.
- [3] C. Gunaratne, K. Christensen, and S. Suen, "NGL02-2: Ethernet Adaptive Link Rate (ALR): Analysis of a Buffer Threshold Policy," in *Global Telecommunications Conference, 2006. GLOBECOM '06. IEEE*, Nov 2006, pp. 1–6.
- [4] "IEEE Standard for Information technology— Local and metropolitan area networks— Specific requirements— Part 3: CSMA/CD Access Method and Physical Layer Specifications Amendment 5: Media Access Control Parameters, Physical Layers, and Management Parameters for Energy-Efficient Ethernet," *IEEE Std 802.3az-2010 (Amendment to IEEE Std 802.3-2008)*, pp. 1–302, Oct 2010.
- [5] F. Idzikowski, E. Bonetto, L. Chiaraviglio, A. Cianfrani, A. Coiro, R. Duque, F. Jiménez, E. Le Rouzic, F. Musumeci, W. Van Heddeghem, J. López Vizcaíno, and Y. Ye, "TREND in energy-aware adaptive routing solutions," *Communications Magazine, IEEE*, vol. 51, no. 11, pp. 94–104, November 2013.



(a) Both hosts connected to an edge router generate a total of 100 Mbps. (b) Both hosts connected to an edge router generate a total of 500 Mbps. (c) Both hosts connected to an edge router generate a total of 700 Mbps.

Fig. 10: The figures show the CDF of the maximum link utilization (MLU) for the whole network. The data shows that EAGER and CARE create a significant number of links with no traffic, which explains the improved energy saving shown in Fig. 9(a). At low loads, Fig. 10(a), the MLU is under 30% for all schemes, which coupled with the data from Figure 1, explains why Figure 9(a) shows all layer-3 schemes performing the same at 100 Mbps. At very high loads, Fig. 10(c), EAGER, as expected, suffers from a significant number of links that reach saturation. CARE-50 and CARE-80 handle congestion much better.



(a) Both hosts connected to an edge router generate a total of 100 Mbps. (b) Both hosts connected to an edge router generate a total of 500 Mbps. (c) Both hosts connected to an edge router generate a total of 700 Mbps.

Fig. 11: The figures show the CDF of the end-to-end delay for all packets delivered during the simulation. The data shows all schemes having approximately the same performance in terms of end-to-end delay at low to medium loads. At very high loads, Fig. 11(c), the data shows that with EAGER, and to a lesser extent with CARE-50, a percentage of packets take very long to reach the destination. CARE-80 exhibits performance very similar to LEGACY OSPF even at very high loads.

[6] 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, November 2010.

[7] M. Zhang, C. Yi, B. Liu, and B. Zhang, "GreenTE: Power-aware traffic engineering," in *Network Protocols (ICNP), 2010 18th IEEE International Conference on*, Oct 2010, pp. 21–30.

[8] F. Cuomo, A. Abbagnale, and S. Papagna, "ESOL: Energy saving in the Internet based on Occurrence of Links in routing paths," in *World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2011 IEEE International Symposium on a*, June 2011, pp. 1–6.

[9] A. Cianfrani, V. Eramo, M. Listanti, M. Marazza, and E. Vittorini, "An energy saving routing algorithm for a green ospf protocol," in *INFOCOM IEEE Conference on Computer Communications Workshops*, 2010, March 2010, pp. 1–5.

[10] D. Arai and K. Yoshihara, "Eco-friendly distributed routing protocol for reducing network energy consumption," in *Network and Service Management (CNSM), 2010 International Conference on*, Oct 2010, pp. 104–111.

[11] P. Reviriego, J. Hernandez, D. Larrabeiti, and J. Maestro, "Performance evaluation of energy efficient ethernet," *Communications Letters, IEEE*, vol. 13, no. 9, pp. 697–699, Sept 2009.

[12] A. Bianzino, L. Chiaraviglio, and M. Mellia, "Distributed algorithms for green ip networks," in *Computer Communications Workshops (INFOCOM WKSHPS), 2012 IEEE Conference on*, March 2012, pp. 121–126.

[13] E. Amaldi, A. Capone, L. Gianoli, and L. Mascetti, "Energy management in IP traffic engineering with Shortest Path routing," in *World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2011 IEEE International Symposium on a*, June 2011, pp. 1–6.

[14] B. Fortz and M. Thorup, "Internet traffic engineering by optimizing OSPF weights," in *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 2, 2000, pp. 519–528 vol.2.

[15] M. Shen, H. Liu, K. Xu, N. Wang, and Y. Zhong, "Routing on Demand: Toward the Energy-aware Traffic Engineering with OSPF," in *Proceedings of the 11th International IFIP TC 6 Conference on Networking*, 2012, pp. 232–246.

[16] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright, "Power Awareness in Network Design and Routing," in *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, April 2008, pp. 457–465.

[17] A. Ruiz-Rivera, K.-W. Chin, R. Raad, and S. Soh, "Hotpluz: A bgp-aware green traffic engineering approach," in *Communications (ICC), 2014 IEEE International Conference on*, June 2014, pp. 3721–3726.

[18] "Best Practices in Core Network Capacity Planning," *Cisco White Paper*, 2014. [Online]. Available: http://www.cisco.com/c/en/us/solutions/collateral/service-provider/quantum/white_paper_c11-728551.html

[19] A. Nucci, N. Taft, P. Thiran, H. Zang, and C. Diot, "Increasing the link utilization in ip over wdm networks using availability as qos," *Photonic Network Communications*, vol. 9, no. 1, pp. 55–75, 2005.

[20] NSNAM, ns-3: network simulator 3. [Online]. Available: <http://www.nsnam.org>

[21] A. Medina, I. Matta, and J. Byers, "BRIT: A Flexible Generator of Internet Topologies," Boston, MA, USA, Tech. Rep., 2000.

[22] A.-L. Barabási and R. Albert, "Emergence of Scaling in Random Networks," *Science*, vol. 286, no. 5439, pp. 509–512, 1999.