Creating Green Incentives and Mechanisms Through Packet-Level Energy Accounting

Matti Siekkinen, Jukka K. Nurminen, Antti Ylä-Jääski
School of Science, Dept. of Computer Science and Engineering
Aalto University
Espoo, Finland
Email: {matti.siekkinen,jukka.k.nurminen,antti.yla-jaaski}@aalto.fi

***Abstract*—We explore the idea of accounting energy consumption of networks on packet-level. IP packets would collect the information of energy that they consumed at each hop and convey it further when travelling through the network. We describe a mechanism to do this accounting and analyze a few use cases including one where packets are given priorities based on how energy efficiently they have travelled the network. Our conclusion is that it would create incentives for Internet service providers (ISPs) to provide a more energy efficient service. In addition to creating incentives, the energy counters in packets could also be used in mechanisms to reduce the energy consumption. We also discuss the limitations and challenges in implementing and deploying such a scheme.**

I. INTRODUCTION

Energy consumption of the ICT industry has become a relevant issue [1]. As a consequence, the growing concern has given rise to a lot of research activities. There seems to be at least three major concerns. First, companies that are operating, for instance, networks or large data centers are very interested in cutting their electric bill [2]. Second, so called "green ICT" activities focus on improving the environmental sustainability in a global scope [3]. Third, a rather different problem concerning energy efficiency arises from the battery constrained devices and the quality of experience for their users in terms of battery life. In this paper, we address mainly the first two concerns.

It is currently very difficult, if not impossible, to know exactly how much energy is consumed in a network and where in the network, i.e. by which equipment, it is expended. Researchers are investing significant efforts, through benchmarking, measurements, and modeling, to improve the energy awareness. In addition, the EMAN working group [4] in IETF is working towards this objective. Such energy awareness of the infrastructure is required to be able to design mechanisms to optimize the energy consumption.

In this paper, we study the following question: What would it enable if we could separately account for the energy consumed by each packet in the network by each network element? The idea is that each packet carries an energy counter whose value could be checked and increased accordingly by each server, router, switch, etc. We approach this question through a couple of use cases which allow us to conclude that such packet-level energy accounting might open up interesting avenues for reducing the overall energy expenditure of ICT. Specifically, the mechanism makes it possible to provide incentives for service providers to improve the energy efficiency of their networks and content delivery. We would like to point out upfront that this paper describes early stage work in progress and that we have no implementation for this scheme. However, we do also discuss the feasibility of implementing and deploying such an accountability scheme.

II. ENERGY ACCOUNTABLE PACKETS

A. Core Idea

The basic idea that we propose is simple: Each IP packet carries a cumulative counter which is incremented at each hop according to how much energy was consumed to process the packet. The processing includes receiving the packet, protocol specific processing of the packet headers and content, and transmitting the packet to the next hop. For the sake of simplicity, we do not consider the energy consumed by application-level processing which happens locally at the client or at the service provider side, which we discuss as an extension to the scheme in Section IV.

To understand why designing this kind of accounting mechanism is non-trivial, let us look at Table I which lists some example devices and their EPI values. EPI stands for Energy Proportionality Index [5] and describes how proportional is the power consumption of the device to the introduced load: the lower the value, the less energy proportional is the device. The metric is computed as follows: $EPI = \frac{M-I}{M} * 100$, where M and I are power draw with maximum load and when idle, respectively. EPIs of a hub, switch, a router, and a wireless access point are reported in [5], while for mobile device and server we estimated the values by comparing the power consumed in idle/sleep state to the power consumed in active receive/transit state (DCH channel for 3G) reported in [6], [7]. Since we do not account for application level energy, we consider servers in Table I just as devices having a basic 802.3 or 802.3az, i.e., Energy Efficient Ethernet (EEE) [8], network interface card. Base station estimates are from [9], [10].

The EPI values in the table show that most of the infrastructure equipment is not consuming energy in proportion to the load. In other words, they have a largish idle power consumption and the association of this energy expenditure to specific packets is not straightforward. For example, if there is

TABLE I

EPI VALUES OF NETWORK ELEMENTS.

Device	server	mobile device	hub	edge router / switch	wireless AP	UMTS base station
EPI (%)	0 (802.3) / 90 (802.3az) [8]	>90 [6], [7]	10 [5]	20-25 [5]	50 [5]	0-50 [9], [10]

a mostly idle edge router or access point to which two packets arrive back to back after a long idle interval. In this case, should the first packet be responsible for all the energy that has been consumed while being idle before the packet arrival? If so, then the second packet would seem to have "consumed" much less energy although the two packets might be of same size and have travelled the exact same path. Clearly the idle power should be taken somehow into account. Otherwise, the "wasted" energy by lightly loaded routers would not be captured by our approach, which is important as we explain through the use cases. We propose an accounting model that takes idle power into account in the following section.

Network elements also need to perform various background tasks, such as, routing, i.e. exchanging information with other routers using routing protocols and computing routes using routing algorithms. One might ask which tasks' energy, if not all, should be accounted for a particular packet being forwarded. To simplify things, we suggest that energy consumed executing such tasks is shared with all the packets passing through the device and accounting of energy consumed per specific task is unnecessary.

The table contains a mixture of devices operating on different layers. If we embed the energy counters into IP protocol headers, a pure layer 2 device is unable to check and update the values. We discuss this issue in Section II-C.

B. Accounting Model

To take into account the idle power and to associate the spent energy in a fair way to different packets is non-trivial. The solution should be both fair and implementable without major resource requirements. In addition, the scheme should not delay the forwarding of packets, which means that the decision of how much energy is to be associated to a particular packet has to be done without knowledge of packets that arrive afterwards.

We propose two ways to do the accounting, both with tunable parameters. The first method is that the network element periodically updates a value which would be added to the energy counter in a packet. The second method is event based so that the network element updates the value each time a packet is processed.

In both methods the energy consumption per packet E_p depends on the power draw of the device processing the packet and inter-arrival time of the packets. We define P_δ^t and I_δ^t as the average power draw and average packet arrival rate, respectively, computed over a time interval δ at a time instant t. With these we compute the energy for a packet processed at time instant t:

$$E_p^t = \frac{P_\delta^t}{I_\delta^t} \qquad (1)$$

In the first method, we assume that the device keeps track of packet counters and samples power draw at a certain frequency [1]. Then, we simply define a fixed δ as the sampling period over which we average the values of P_δ^t and I_δ^t: $I_\delta^t = \frac{\text{nb of packets in interval}[t-\delta,t]}{\delta}$ and P_δ^t is the mean of the power samples collected within the interval $[t-\delta, t]$. In this method, the way δ is chosen determines the sensitivity of the energy counters to uneven packet arrival process. Setting a long δ means that burstiness of traffic in short time scales does not influence the energy values that packets are associated to. On the other hand, a shorter δ makes E_p more reactive to changes in traffic load. However, too short values of δ would lead to the first packets in a burst getting unfairly a higher values than the subsequent packets.

The alternative second method is to update E_p for each packet using weighted moving averages as follows. We define the average inter-arrival time at time instant t: $A_t = \alpha\delta + (1-\alpha)A_{t-1}$, where δ is the time elapsed since the last packet arrival. We now get: $I_\delta^t = \frac{1}{A_t}$. P_δ^t is the same as in above first method. Note that in this method, δ is not fixed and the averaging interval depends on the packet arrival process. The sensitivity to changing load can be adjusted by tuning the weight parameter α. The advantage of this method is that E_p is updated only when needed. On the other hand, if packet arrival rate is very high, an update per packet may cause too much computational overhead and the first method is preferable.

C. Energy Counters, Packets, and Layers

A single hop on IP layer can be just a single Ethernet link, multiple switched Ethernet links, or even an entire MPLS path or another layer 2 tunnel. Therefore, the energy consumed by, e.g., the routers at each end of the IP hop does not always cover the entire energy consumption of that hop. What we need is a way to account for the energy consumed by the layer 2 devices in addition.

The ideal way to achieve this would be to have similar accounting mechanism on each layer which would be linked together during packet decapsulation. In other words, lower layer protocols would carry and update a similar energy counter, and at the next device that runs also a layer higher protocol, the lower layer energy counter would be added to the higher layer one.

Figure 1 illustrates how this scheme would work with an Ethernet switch in between two IP routers. The packet

---

[1] The sampling frequency of power draw may depend on the device capabilities and variability of power (EPI).

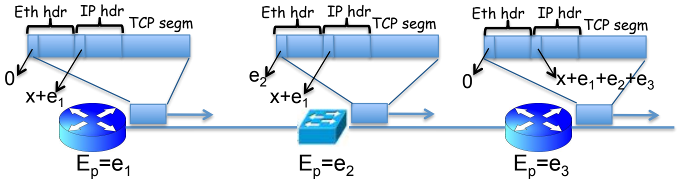

Fig. 1. Updating energy counters across layers.

arrives to the first router with already accumulated energy counter value x in the IP header.The switch updates only the counter in the Ethernet header. The following router takes the accumulated counter value from the Ethernet header and adds that to the counter in the IP header together with the energy accounted for the packet in that router.

III. USE CASES

We analyzed a few different use cases in order to evaluate how such an energy accounting scheme could be useful.

A. *Greening the Traffic*

In this use case, we explore the consequences of modifying the router queue management. Namely, we consider introducing some form of priority queueing to the router which favors packets with lower energy counter values. The purpose is to provide incentives for, e.g., ISPs to improve energy efficiency of their networks. The queuing mechanism should be a combination of fair and priority queuing, such as weighted fair queuing, to ensure that all priorities get at least some service.

The idea is that by prioritizing low energy packets, traffic of an energy inefficient service provider gets a low priority, while a provider paying more attention to its energy efficiency grabs a larger share of the shared resources. In this way, in order to compete with their rivals to provide the best service to their clients, providers would have a strong incentive to consider how their service can be delivered to the clients in most energy efficient manner. Best service in this context can be measured in terms of latency or throughput.

However, we need to be careful in selecting the metric based on which we set priorities. If we just consider the accumulative energy counter value found in each packet, we would expect the priority mechanism to drive towards localization of services because intuitively the closer the service is to the client, the shorter the path and the smaller the accumulated energy consumption by the packets. Localization of traffic is good in terms of overall energy consumption and content service providers can achieve this through CDNs, for instance. However, as a side effect, we would also penalize services which need to convey traffic over long paths because the service is not available elsewhere. An example is long distance VoIP call.

Hence, we think that a better metric for prioritization would be the average energy consumption per hop on the path travelled by the packet. Thus, a packet would not be penalized even if it takes a long route as long as that route consisted of energy efficient hops. In this way, the energy efficiency would depend solely on the utilization and energy consumption profile of the networking equipment in which the energy proportionality (Section II-A) plays a big role.

This case would build new incentives. For example, an ISP might invest in new equipment which has advanced power saving features and consumes energy in proportion to the load. This ISP would obviously benefit from a smaller electric bill in the future. In addition, it could provide an advantage to clients that route traffic through it by having smaller per-hop energy counter values for packets at its egress points than if the client would route its traffic through a competing ISP which is using less energy proportional routers. The client's traffic would obtain a higher priority in later path segments over traffic from other less energy efficient access or transient ISPs. Of course, also the absolute power draw of network elements with a given load makes a difference in this scheme. This fact means, for instance, that energy expensive Internet access technologies would be penalized.

In case of congestion and queue fill-up, it would mean that a router supporting this priority scheme is more likely to drop a packet among those with the highest energy counter values. This decision may actually increase the short term energy waste. The reason is that we would maximize the probability of a retransmission of the most energy inefficient packets. Nevertheless, the long term effect should be lower energy consumption through the incentives.

B. *Energy Optimization of Networks*

Adding an energy counter to packets could help to operate networks more energy efficiently. Earlier research work has proposed optimization methods for traffic engineering to reach most energy efficient handling of given workloads of traffic. The idea is to route traffic either so that the network element processing speed is scaled to exactly match the introduced load [11] or so that some of the equipment is temporarily (and partially) powered down [12].

While many proposed algorithms can compute optimal solutions for energy efficient routes, it is a challenge to deploy them in other than centralized fashion. The energy counters in packets could be utilized as an enabler for solutions of distributed nature. For example, in the energy-aware traffic engineering proposed in [13] each intermediate router periodically reports its link utilization, while edge routers, based on this information, distribute traffic across alternative paths in a way that maximizes energy savings. The energy counters in packets could be leveraged in this scheme to disseminate the necessary information from core routers to the edge without any additional overhead. Some other related work that also provide solutions of distributed nature are presented in [14], [15].

C. *Tracing Energy End-to-end*

1) Optimizing the Service Delivery: Energy accounting would allow a service provider to estimate how much energy would be spent when serving a client. The estimation could be based on the energy consumed by the HTTP requests, the acknowledgement packets of TCP protocol, or any other upstream communication from clients. In our accounting model, the amount of energy spent by a network element is independent of the direction of the flow. The scheme would not be accurate in all cases, though, since packet of a particular bidirectional flow do not always follow the same route in both directions [16].

Knowing the communication energy cost to serve customers in different locations and networks would allow providers to optimize their service delivery. For example, a CDN provider typically has the same content available on multiple sites and the DNS lookup directs the client to the appropriate site. If estimates of the communication energy costs from different sites to different customer locations were available, energy-efficiency could be considered when defining the site allocation policies for users of different regions.

2) Energy traceroute and other tools: Energy accounting would allow widely used network analysis tools to be extended with energy information. In the simplest case, ping and traceroute could be extended with an option to also show energy consumption, which would allow networks administrators and researchers to analyze end-to-end energy consumption and its distribution to different route sections.

We expect that making communication energy consumption visible in such tools would be a good basis for further innovation in energy-efficient networking. Though, in comparison to delay or throughput measurements, the energy measurements are likely to be less accurate and subject to decisions on how the spent energy is divided to different packets. How accurate the estimates are and what level of accuracy is necessary to be useful is a topic for further investigation.

3) Energy Aware Customers: The idea here is to bring the awareness of the energy consumption of accessing and an Internet service to the client. The reason is that there seems to be a growing interest in choosing the “green” option among a selection of otherwise similar products. As customer demand for more sustainable products grows, so does the potential to make profits out of this trend [17].

Given the possibility to equip packets with energy counters, clients downloading a file or accessing a web service have means to get the information about how much energy is being consumed. For instance, a customer could learn what is the network-wise energy cost of performing a Google search, and in this way get an idea of the carbon footprint of their actions. They could also select, for instance, a “greener” way to read news online.

IV. DISCUSSION

A. Robustness

Given that the energy counters could be utilized, for instance, to enforce priorities to more energy efficient content, the economic implications could be significant from service providers’ point of view. Thus, before such use cases could be supported, some sort of mechanisms to enforce truthful reporting of energy consumption would be necessary.

The main threat to the scheme is cheating network elements which would untruthfully report too low or too high energy consumption. Too low reporting may give higher priority to traffic which would allow providing better quality of service to clients (see Section III-A). Alternatively, clients could be mislead to think that a particular service provider provides a more energy efficient service than it actually does (see Section III-C). Compromised network elements could report too high energy counter values to cause denial of service to e.g., certain service providers because that traffic would end up having low priority.

To defend against such misbehavior, the operations could be made tamper proof if the accounting was performed by hardware, e.g. by network interface cards.

B. Application-Level Accounting

An interesting extension to our scheme would be to include also application-level energy consumption. However, allocating this energy to individual packets is more challenging. Consider a data center, for instance, which may perform tremendous amount of tasks in order to reply a single simple query, such as a search request. In addition, it is not always easy to associate particular tasks to a specific (set of) application(s).

One strategy would be to apply the accounting model in a coarse grained fashion. For example, a given data center would track its overall average energy consumption and this energy would be associated to all packets exiting the data center. Hardware-based tamper proof execution of such coarse grained accounting would not be difficult, though.

C. Implementation

Implementing the accounting scheme requires means for network elements to monitor their energy consumption. Currently, few devices have such capabilities but the situation is likely to change in the future (see Cisco EnergyWise [18], for instance). In addition, the EMAN group in IETF is actively working towards standards for energy consumption monitoring and control.

In case the device hardware does not support direct measurement of its own energy consumption, models can be used to estimate it. There is already a good amount of existing literature on power models, e.g. for servers [19], [20], wireless network interfaces [6], routers and switches [5], [12], base stations [9], etc. The hardware based updating of the counters would be a necessity in case we want to have tamper proof mechanisms.

The actual energy counters could be carried either as a separate header using IPv6 header extension techniques (similarly to [21], for instance) or in IPv4 option space. The latter way has the drawback that at least currently some routers drop packets with IP options due to security related reasons. In lower layer (L2) protocols, new extended header formats would need to be defined. In addition to the energy counter, packets should carry a hop counter so that computation of average per-hop energy consumption is made possible. Also the TTL value of the IP packet when the energy counter value was last updated should be included, which is useful in incrementally deploying the scheme as we explain in the following section.

D. Deployment

Initial deployments of the scheme could rely on estimations of energy consumption as explained above. In addition, the scheme could first be implemented only on IP layer elements. Assuming that layer 2 paths are often fairly stable, the energy consumed by an entire layer 2 path within an IP hop could be based on estimates. In such a case, either the IP network element just before the layer 2 path or the one just after it would add the estimated layer 2 path energy consumption.

The queue management mechanism described in our second use case would need to be deployed only at the edge routers of each AS. To make incremental deployment attractive, we propose two strategies to handle legacy network elements. First, packets that do not have energy counters at all will be assigned the lowest priority. Thus, a particular host or an ISP has incentives to start using the accounting mechanism because it would increase the priority of the traffic from that host or ISP regardless of the energy efficiency of the equipment.

Second, if packets with energy counters traverse IP level network elements that do not support energy accounting, they will get an additional constant relatively high increment to energy counters for each such traversed IP hop. This increment is added at the next IP level network element that supports the accounting scheme. The number of such legacy IP hops traversed can be computed by subtracting the TTL at last energy counter update from the present TTL. In this way, it is beneficial to favor routing traffic through ISPs that have upgraded their equipment to support the energy accounting.

V. CONCLUSIONS AND FUTURE WORK

We propose in this paper a per-packet energy consumption accounting mechanism. By examining a couple of use cases it seems that the scheme could provide interesting additional incentives for ISPs to improve the energy efficiency of their network and service providers to deliver services in more energy efficient manner. In addition, the mechanism could help in achieving these improvements.

We believe that this accounting scheme is interesting enough to merit further studying. In particular, we intend to perform simulations and analytical studies leveraging existing data sets on topologies and traffic matrices in order to better understand and quantify the impact for the prioritization use case in a realistic setting. We also want to study the impact of specific packet arrival rate and process combined with varying parameter values in the accounting methods to the per-packet energy allocations.

We are also interested in exploring further extensions to the scheme, such as application-level energy accounting and accounting for “greenness” or cost of energy, e.g. how to take renewable vs. non-sustainable energy into account.

ACKNOWLEDGMENTS

This work was supported by the Academy of Finland, grant number 253860.

REFERENCES

[1] T. C. Group, “SMART 2020: Enabling the low carbon economy in the information age,” Global e-Sustainibility Initiative (GeSI), Report at http://www.gesi.org/, Aug. 2008.

[2] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs, “Cutting the electric bill for internet-scale systems,” in *Proceedings of the ACM SIGCOMM 2009 conference on Data communication*, ser. SIGCOMM ’09. New York, NY, USA: ACM, 2009, pp. 123–134.

[3] A. P. Bianzino, C. Chaudet, J. L. Rougier, and D. Rossi, “A survey of green networking research,” *IEEE Communication Surveys & Tutorials (to appear)*, 2011.

[4] “Ietf energy management working group: http://datatracker.ietf.org/wg/eman/charter/.”

[5] P. Mahadevan, P. Sharma, S. Banerjee, and P. Ranganathan, “A power benchmarking framework for network devices,” in *Proceedings of the 8th International IFIP-TC 6 Networking Conference*, ser. NETWORKING ’09. Berlin, Heidelberg: Springer-Verlag, 2009, pp. 795–808.

[6] Y. Xiao, P. Savolainen, A. Karppanen, M. Siekkinen, and A. Ylä-Jääski, “Practical power modeling of data transmission over 802.11g for wireless applications,” in *e-Energy ’10: Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking*. New York, NY, USA: ACM, 2010, pp. 75–84.

[7] N. Balasubramanian, A. Balasubramanian, and A. Venkataramani, “Energy consumption in mobile phones: a measurement study and implications for network applications,” in *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*, ser. IMC ’09. New York, NY, USA: ACM, 2009, pp. 280–293.

[8] K. Christensen, P. Reviriego, B. Nordman, M. Bennett, M. Mostowfi, and J. Maestro, “Ieee 802.3az: the road to energy efficient ethernet,” *Communications Magazine, IEEE*, vol. 48, no. 11, pp. 50 –56, Nov. 2010.

[9] O. Arnold, F. Richter, G. Fettweis, and O. Blume, “Power consumption modeling of different base station types in heterogeneous cellular networks,” in *in Proc. of 19th Future Network and MobileSummit 2010*, 2010.

[10] L. Correia, D. Zeller, O. Blume, D. Ferling, Y. Jading, I. Go anddor, G. Auer, and L. Van Der Perre, “Challenges and enabling technologies for energy aware mobile radio networks,” *Communications Magazine, IEEE*, vol. 48, no. 11, pp. 66 –72, Nov. 2010.

[11] M. Andrews, A. Anta, L. Zhang, and W. Zhao, “Routing for energy minimization in the speed scaling model,” in *INFOCOM, 2010 Proceedings IEEE*, Mar. 2010, pp. 1–9.

[12] 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*, Apr. 2008, pp. 457 –465.

[13] N. Vasić and D. Kostić, “Energy-aware traffic engineering,” in *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking*, ser. e-Energy ’10. New York, NY, USA: ACM, 2010, pp. 169–178.

[14] A. Bianzino, L. Chiaraviglio, and M. Mellia, “Grida: A green distributed algorithm for backbone networks,” in *IEEE Online Conference on Green Communications (GreenCom)*, Sep. 2011.

[15] A. Cianfrani, V. Eramo, M. Listanti, M. Marazza, and E. Vittorini, “An energy saving routing algorithm for a green ospf protocol,” in *INFOCOM Workshops*, Mar. 2010.

[16] Y. He, M. Faloutsos, S. Krishnamurthy, and B. Huffaker, “On routing asymmetry in the internet,” in *Global Telecommunications Conference, 2005. GLOBECOM ’05. IEEE*, vol. 2, Nov. 2005.

[17] D. C. Esty and A. S. Winston, *Green to gold: how smart companies use environmental strategy to innovate, create value, and build competitive advantage*. John Wiley and Sons, Jan. 2009.

[18] “Cisco energywise technology: http://www.cisco.com/en/us/products/ps10195/.”

[19] A. Lewis, S. Ghosh, and N.-F. Tzeng, “Run-time energy consumption estimation based on workload in server systems,” in *Proceedings of the 2008 conference on Power aware computing and systems*, ser. HotPower’08. Berkeley, CA, USA: USENIX Association, 2008, pp. 4–4. [Online]. Available: http://portal.acm.org/citation.cfm?id=1855610.1855614

[20] S. Rivoire, P. Ranganathan, and C. Kozyrakis, “A comparison of high-level full-system power models,” in *Proceedings of the 2008 conference on Power aware computing and systems*, ser. HotPower’08. Berkeley, CA, USA: USENIX Association, 2008, pp. 3–3. [Online]. Available: http://portal.acm.org/citation.cfm?id=1855610.1855613

[21] D. Lagutin and S. Tarkoma, “Cryptographic signatures on the network layer - an alternative to the isp data retention,” in *IEEE Symposium on Computers and Communications (ISCC)*, Jun. 2010, pp. 87 –93.