Not a Trade-Off: On the Wi-Fi Energy Efficiency of Effective Internet Congestion Control

Michael Welzl
University of Oslo
Oslo, Norway
michawe@ifi.uio.no

***Abstract*—It is common to regard performance and energy efficiency as a trade-off, and this is a characteristic of many proposed solutions for energy efficiency. This is disadvantageous, as it is tempting for administrators or end users to disable energy conserving mechanisms when they have a performance cost. In contrast, this article makes the point that improvements in Internet congestion control can be *inherently* energy-efficient: for example, minimizing the Flow Completion Time (FCT) of data transfers, one of the most common goals in congestion control, can significantly reduce the energy usage of a Wi-Fi receiver.**

Index Terms—**Wi-Fi, Energy Efficiency, Congestion Control, TCP.**

I. INTRODUCTION

WHILE ICT and the Internet in particular can lower Greenhouse Gas (GHG) emissions (e.g., by reducing the need for travel), the energy used by the Internet itself is large, and limiting it plays a vital role in the fight against global warming. The SMARTer2030 report states that ICT has a CO_2 "footprint" of 2.7% of global emissions in 2020 [1], and CO_2 makes up the vast majority of energy-related GHG emissions.[1] Already in 2012, telecommunications contributed about 1/3 of the ICT sector's electricity consumption [2]. If we conservatively assume that this relationship is still the same, the global 2020 GHG contribution of telecommunications is 0.9%. For context, the aviation industry's contribution was 1.9% in 2016 [3]. Adding to this, the exponential growth of Internet traffic gives rise to worrisome predictions: a much-cited study on electricity usage trends of communication technology [4] predicts that, by 2030, fixed-access Wi-Fi consumer premises equipment will contribute 15%[2] of the global energy consumption.

It is therefore paramount to find ways to reduce the Internet's energy usage. Reducing energy has long been an important topic for wireless networks—but there, performance and energy usage are most commonly treated as a trade-off. This is disadvantageous, as it creates an incentive to disable energy-saving mechanisms. In contrast, in this article, we discuss the possibility of achieving energy reduction *in conjunction with* improved performance. This opportunity is provided by Internet congestion control.

Congestion control (CC) aims at efficiently utilizing a network's resources while avoiding to overwhelm them. Nowadays, networks are often provisioned with high capacities, and transfers ("flows") are often too short for a protocol's CC mechanism to saturate a bottleneck. For example, more than 80% of TCP transfers in the large search company *Baidu* end in their Slow Start phase [6], i.e. they terminate before their CC mechanism was even able to probe for the available capacity. In a multi-year dataset captured at a Tier-1 ISP backbone link in Chicago from 2008 to 2016, 85% of flows carry at most 10 kB of data [7]—this amounts to a handful of packets, often less than the Initial Window (IW) that Slow Start begins with. Even some seemingly long-lasting transfers such as Dynamic Adaptive Streaming over HTTP (DASH) behave like this: Netflix, for instance, sends data in short spikes with multi-second pauses in between [8], which is arguably similar to consecutive short flows "on the wire".

When flows end before saturating a path's capacity limit, network delays are primarily caused by round-trips rather than congestion. This has led to a desire to reduce the number of these round-trips, which has heavily influenced the design of the QUIC protocol [9], and it is at the core of proposals to increase the IW [6], [10]–[13], better support application-limited traffic [14], and initialize new flows based on previous or other ongoing ones [15], [16]. It is now perhaps the most important goal of any Internet CC mechanism to *quickly* increase the sending rate, such that the available capacity is fully utilized and the Flow Completion Time (FCT) is minimized.

The contribution of this article is to show that reducing the FCT (i.e., performing effective CC) *also* reduces energy usage—quite simply, because sleep cycles are time-dependent, and less time for transmission means more time to sleep. Figure 1 schematically illustrates this. Our finding, while unsurprising, is important, because it means that energy-efficiency does not need to be a trade-off: efficient congestion control can *both* reduce the perceived delay *and* save energy.

As a case in point, we will now look at how the FCT of a short flow influences a Wi-Fi receiver's energy consumption. Then, after an overview of related work in Section III, we will discuss the wider ramifications of the energy-saving ability of congestion control in Section IV.

---

[1] https://www.epa.gov/ghgemissions/sources-greenhouse-gas-emissions

[2] Predictions being what they are, these 2030 estimates are not uncontested—e.g., the authors of [5] find that the energy and carbon footprints of the ICT and Entertainment & Media sectors are now diminishing, and that they are smaller in magnitude than previously anticipated. However, there is not much doubt that the numbers are large, and that energy should be saved to improve the situation.

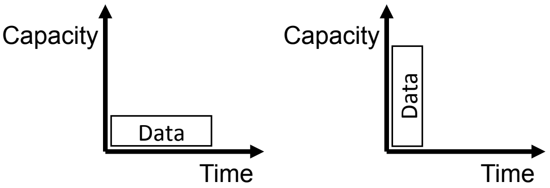

Fig. 1. Poor (left) and good (right) resource utilization. The case on the right performs better, *and* it is more energy efficient.

II. A TOY EXAMPLE: TCP'S IW AND WI-FI ENERGY

To evaluate how "poor" versus "good" congestion control influences the energy consumption of a Wi-Fi receiver, we ran single tests with small versus large choices for TCP's IW, picking the common values 2, 4 and 10. These form the large majority of measured cases in [17], and 10 has been the default choice in the Linux kernel for a decade. In reality, the ideal IW choice is complex, as it depends not only on the bandwidth×delay product, but also on the queue length at the bottleneck and the number of parallel flows. We stress that this is only a "toy example", intended to highlight the possibility of saving energy by shortening the FCT, not a proposal for a specific static IW value.

We ensured that "bigger is better" in our tests by assigning a static queue of 20 packets (enough to hold the initial burst), using only a single flow, configuring abundant capacity (80 Mbit/s, which is the attainable throughput from the "fast.com" Netflix speed test to a device in the author's home over a DOCSIS connection followed by an 802.11ac Access Point), and confirming that no packets were lost (i.e., all transfers terminated in Slow Start).

Measuring the energy usage of Wi-Fi data transmission (or reception) is difficult and may require special hardware (software-based solutions, e.g., using a Smartphone's battery life indication, are known to be error-prone). It has hence become quite common to rely on models instead of carrying out such direct measurements [18]. Among them, we chose EnergyBox [19], a tool that can provide an estimation of Wi-Fi power consumption related to a certain traffic pattern (provided via a pcap file with a specified IP address for the "monitor host"). EnergyBox has been validated to achieve 95-99% accuracy, and it has been used by Spotify for quantifying their mobile application's energy consumption. We used the default values from [19] with the only Wi-Fi device configuration file that is supplied with EnergyBox, corresponding to a Samsung Galaxy SII; the relevance of our findings for newer devices will be discussed in Section II-A. We produced the pcap files in a Linux based wired testbed with a sender-router-receiver topology, using the TEACUP scripting environment [20].

There were two different types of data transfers, one with a length of 10 and another with a length of 80 packets, to represent a short and long flow, respectively, using a standard Ethernet Maximum Transmission Unit (MTU) of 1500 bytes. The payload of the 10-packet flow is a little larger than the aforementioned 10 kB. For the 80-packet flow it is between the 100 kB threshold chosen in [7] to distinguish between small and big flows, and the average flow size (120 kB) mentioned in [6]. Between tests, the emulated end-to-end base round-trip time (RTT) was varied in steps of 5 from 5 to 100 ms.

The longest transfer took 720 ms (IW2, RTT 100 ms, 80 packets), with a reported power consumption of 0.3453 Joules, whereas the shortest one only took 23 ms, with a reported power consumption of 0.0607 Joules; clearly, the shorter the transmission time, the less power is needed. The per-trace output of EnergyBox reveals a very simple state diagram: the host initially wakes up upon the first ("SYN") packet, entering the Constant Awake Mode (CAM) state, then enters a state called "CAM-H" for some parts of the data transmission, and then, 200 ms after the last packet, the test is over.[3] No single test had any occurrence of a packet inter-arrival time that was long enough for the 200 ms inactivity timer to fire and allow the host to enter the sleep mode earlier. For a fair comparison, we must assume that the device stayed active for the same time in all tests. We therefore considered a total test duration of 1 second and added the configured power in Power Saving Mode (PSM) state (30 mW) over the period of 1 second minus each test's actual duration as determined by EnergyBox.

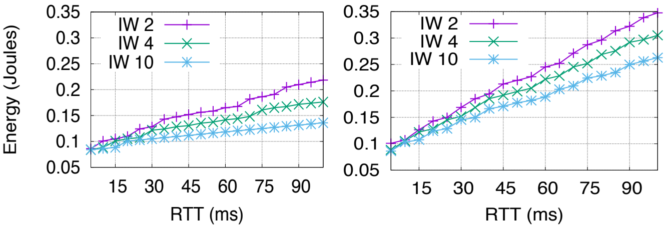

(a) 10 packets

(b) 80 packets

Fig. 2. Energy consumption of short TCP transfers with varying IW values.

Figure 2 shows the raw test results. We can see that the impact of the IW choice depends on the RTT, and its overall influence declines with the length of the data transfer. The RTT dependence is due to the time t needed to finish a lossless and not otherwise limited (e.g., by the receiver window) TCP transfer in Slow Start, which, assuming idealized conditions (static RTT, no processing delay, ..) is given by:

$$t = (\lceil log_2(\frac{S}{IW} + 1)\rceil + 1) * RTT + \frac{S}{C} \qquad (1)$$

Here, S is the size of the transfer and C is the bottleneck capacity. Equation 1 is based on [21], but without considering delayed ACKs because modern stacks implement "Appropriate Byte Counting" (ABC) [22]. ABC compensates for the otherwise diminished rate of exponential growth that is due to delayed ACKs. Equation 1 also includes an extra RTT for connection establishment.

Using only the simple finding of entering sleep mode 200 ms after activity and ignoring the CAM-H mode, with the original

---

[3] (CAM-H ("CAM-High") is an artificial state in EnergyBox that is introduced to model that power drainage increases with throughput.

transfer time t_{orig} and the shortened transfer time t_{short} in seconds, the energy reduction R in percent is:

$$R = 100 * \frac{t_{orig} - t_{short}}{t_{orig} + 0.2} \tag{2}$$

Clearly, to maintain a constant reduction, $t_{orig} - t_{short}$ must increase when t_{orig} increases, which is not achieved by changing IW. This explains the declining influence of the IW with increasing data transfer lengths. From our measurements, with an RTT of 50 ms, the 10-packet transfer requires 27% less power with IW 10 than with IW 2, and 16% less than with IW 4. This roughly matches Equation 2, which yields a reduction by 25% and 14%, respectively. With an RTT of 100 ms, the measured power reduction with IW 10 is 38% compared to IW 2 and 23% compared to IW 4. In case of the 80 packet transfer, these gains are 19% (IW 10 vs. IW 2, 50 ms), 11% (IW 10 vs. IW 4, 50 ms), 24% (IW 10 vs. IW 2, 100 ms) and 14% (IW 10 vs. IW 4, 100 ms).

A. Discussion

The results have shown that realistic IW values, with realistic flow sizes, can yield a significant energy gain. We stress that we are not proposing a specific IW value as a solution for better energy efficiency: the shown results highlight that an effective congestion control mechanism which is able to reduce the FCT would indeed also reduce Wi-Fi energy usage. Statically choosing an IW is a trade-off that depends on the expected environment conditions, often making even IW 10 a relatively conservative choice (e.g., in Content Distribution Networks, values as high as 50 are not uncommon [23]); however, if a research approach is able to increase the congestion window (cwnd) fast, much larger values can be considered. For example, when traffic is paced (instead of the common burst that the IW can produce), the bandwidth×delay product (BDP) of a 100 Mbit/s,10 ms path could accommodate an IW of 80 packets, and this would not even count as a high-speed path by today's standards. On the other hand, all tests in the previous section are limited to a single TCP connection—parallel connections would have to divide the gain, with a less clear result when the connections have heterogeneous RTTs.

The magnitude of the gains also depend on the wireless characteristics such as the duration of the inactivity timer, which puts the device into sleep mode $X\ ms$ after active transmission or reception, and the average power usage per mode. In [19], it is stated that the measured device parameters used in EnergyBox are within the range of similar measurements that other researchers have performed on devices from a similar generation. This is, now, an old generation—however, a very recent study finds that the 802.11ac and 802.11ax radios in four tested smartphones still implements the standard adaptive PSM, where the radio goes to sleep 200 ms after the last Tx/Rx activity [24]. While, with modern hardware, the relative gains would differ from Figure 2 due to different per-mode power usage, the main outcome of reduced power from a larger IW would still be the same.

This is not necessarily the case for *all* Wi-Fi hardware. The authors of [24] find that their measured 802.11ad chipset uses a much more aggressive timeout of 15 ms, and, upon the timeout firing, there are three rules governing the ensuing behavior as a function of the inter-arrival times of later packets. Also, despite the measured devices in [24] seemingly implementing only a subset of this standard, 802.11ax standard actually defines a much more sophisticated power management scheme than most earlier standards, including a Target Wake Time mechanism akin to 802.11ah [25].

Cellular networks have quite different characteristics—generally, the inactivity timer is much longer there (in the order of seconds [26]). Since shortening the FCT only affects the duration of activity before the inactivity timer is started, its energy benefit declines with the length of the inactivity timer (i.e., the larger the inactivity timer, the smaller $t_{orig} - t_{short}$ becomes in Equation 2). We confirmed this with EnergyBox, using the 3G configurations that are provided with the tool.

III. RELATED WORK

There are some prior works that try to make TCP transfers "greener", e.g. [27]–[29], but they either embed a trade-off or try to attain power saving benefits with only minimal performance-impeding side effects. The trade-off is particularly explicit when power consumption influences a choice between multiple network interfaces, e.g. by adapting the congestion control mechanism of MultiPath TCP (MPTCP); examples of such work are [30], [31]. Per-interface energy usage is also applied in some application-layer schemes, e.g. in [32], where it is part of a decision-making process that also involves the choice of protocol (TCP or UDP) and the amount of redundancy (using Raptor codes).

Energy savings can be attained when either the environment conditions are known (as, e.g., in wireless sensor networks [33]), such that it is possible to completely diverge from the standard TCP behavior, or by manipulating traffic inside the network. An example of the latter case is "Active Window Management (AWM)" [34] which reduces energy consumption by manipulating TCP's advertised window at a router. Interestingly, AWM does not entail a trade-off as it does not reduce TCP's performance, probably making it the most closely related work in the literature. However, to the best of our knowledge, the *inherent* energy-efficiency of *improved* (i.e., exhibiting better performance) Internet congestion control has not been documented before.

IV. CONCLUSION

This article attempts to shift the focus of congestion control (CC) research towards energy reduction. It seems clear that reducing the energy consumption of the Internet, including Wi-Fi systems, can play a significant role in the combat against global warming. In the coming years, the significance of energy efficiency will grow (continuing already existing exponential growth trends, see [4]), whereas growing network capacities will make the raw CC performance gains from better resource utilization diminish. We thus posit that energy reduction should become a major future goal of Internet congestion control—no less important than increasing throughput and reducing latency.

How can this goal be attained? As we have shown, reducing the FCT, e.g. by increasing the TCP's Initial Window (IW), can yield a significant energy benefit. At the same time, it increases performance, incentivizing deployment (indeed, performance was the driving factor behind the move towards the static IW value 10, which is now widely in use). However, we have also stressed that our IW test is only a simplistic toy example. Rather than proposing a specific static IW value, we hope that this article can inspire future work to aim at reducing the FCT further, and evaluate the attained energy savings.

As a concrete example of work that may indicate a fruitful direction, the Reinforcement Learning based scheme in [6], which dynamically updates IW based on past experience, is probably power-efficient. Also, we have only discussed power savings at the receiver here, but reducing the FCT may also have an impact at the Wi-Fi access point whenever there is only one active recipient. This opens the door for more research opportunities. Considering wired networks and other types of wireless links would also be important.

REFERENCES

[1] J. Lorincz, A. Capone, and J. Wu, "Greener, energy-efficient and sustainable networks: State-of-the-art and new trends," *Sensors*, vol. 19, no. 22, 2019.

[2] M. K. Weldon, *The Future X Network: A Bell Labs Perspective*. USA: CRC Press, Inc., 2015.

[3] H. Ritchie and M. Roser, "Co2 and greenhouse gas emissions," *Our World in Data*, 2020, https://ourworldindata.org/co2-and-other-greenhouse-gas-emissions, accessed 3 Nov. 2021.

[4] A. S. G. Andrae and T. Edler, "On global electricity usage of communication technology: Trends to 2030," *Challenges*, vol. 6, no. 1, pp. 117–157, 2015.

[5] J. Malmodin and D. Lundén, "The Energy and Carbon Footprint of the Global ICT and E&M Sectors 2010–2015," *Sustainability*, vol. 10, no. 9, 2018. [Online]. Available: https://www.mdpi.com/2071-1050/10/9/3027

[6] X. Nie, Y. Zhao, Z. Li, G. Chen, K. Sui, J. Zhang, Z. Ye, and D. Pei, "Dynamic TCP initial windows and congestion control schemes through reinforcement learning," *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 6, pp. 1231–1247, 2019.

[7] S. Bauer, B. Jaeger, F. Helfert, P. Barias, and G. Carle, "On the evolution of internet flow characteristics," in *Proceedings of the Applied Networking Research Workshop*. ACM, 2021, p. 29–35.

[8] S. C. Madanapalli, H. Habibi Gharakhieli, and V. Sivaraman, "Inferring Netflix user experience from broadband network measurement," in *Network Traffic Measurement and Analysis Conference (TMA)*, 2019.

[9] J. Iyengar and M. Thomson, "QUIC: A UDP-Based Multiplexed and Secure Transport," RFC 9000, May 2021.

[10] J. Chu, N. Dukkipati, Y. Cheng, and M. Mathis, "Increasing TCP's Initial Window," RFC 6928, Apr. 2013.

[11] R. Sallantin, C. Baudoin, E. Chaput, F. Arnal, E. Dubois, and A.-L. Beylot, "Initial spreading: A fast start-up TCP mechanism," in *38th Annual IEEE Conference on Local Computer Networks*, 2013.

[12] Q. Li, M. Dong, and P. B. Godfrey, "Halfback: Running short flows quickly and safely," in *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, ser. CoNEXT '15. New York, NY, USA: ACM, 2015.

[13] N. Kuhn, S. Emile, G. Fairhurst, T. Jones, and C. Huitema, "Transport parameters for 0-RTT connections," Internet Engineering Task Force, Internet-Draft draft-kuhn-quic-0rtt-bdp-09, Jun. 2021, work in Progress. [Online]. Available: https://datatracker.ietf.org/doc/html/draft-kuhn-quic-0rtt-bdp-09

[14] G. Fairhurst, A. Sathiaseelan, and R. Secchi, "Updating TCP to Support Rate-Limited Traffic," RFC 7661, Oct. 2015.

[15] D. J. D. Touch, M. Welzl, and S. Islam, "TCP Control Block Interdependence," RFC 9040, Jul. 2021.

[16] S. Islam, M. Welzl, K. Hiorth, D. Hayes, G. Armitage, and S. Gjessing, "ctrlTCP: Reducing latency through coupled, heterogeneous multi-flow tcp congestion control," in *IEEE INFOCOM 2018 Workshops*, 2018.

[17] J. Rüth, C. Bormann, and O. Hohlfeld, "Large-scale scanning of TCP's initial window," in *Proceedings of the 2017 Internet Measurement Conference*, ser. IMC '17. ACM, 2017, p. 304–310.

[18] L. Sun, H. Deng, R. K. Sheshadri, W. Zheng, and D. Koutsonikolas, "Experimental evaluation of WiFi active power/energy consumption models for smartphones," *IEEE Transactions on Mobile Computing*, vol. 16, no. 1, pp. 115–129, 2017.

[19] E. J. Vergara, S. Nadjm-Tehrani, and M. Prihodko, "Energybox: Disclosing the wireless transmission energy cost for mobile devices," *Sustainable Computing: Informatics and Systems*, vol. 4, no. 2, 2014.

[20] S. Zander and G. Armitage, "TEACUP v1.0 - A System for Automated TCP Testbed Experiments," Centre for Advanced Internet Architectures, Swinburne University of Technology, Melbourne, Australia, Tech. Rep. 150529A, 29 May 2015. [Online]. Available: http://caia.swin.edu.au/reports/150529A/CAIA-TR-150529A.pdf

[21] N. Cardwell, S. Savage, and T. Anderson, "Modeling TCP latency," in *IEEE INFOCOM*, vol. 3, 2000.

[22] E. Blanton, D. V. Paxson, and M. Allman, "TCP Congestion Control," RFC 5681, Sep. 2009.

[23] J. Rüth and O. Hohlfeld, "Demystifying TCP initial window configurations of content distribution networks," in *Network Traffic Measurement and Analysis Conference, TMA, Vienna, Austria*. IEEE, 2018.

[24] S. Aggarwal, M. Ghoshal, P. Banerjee, D. Koutsonikolas, and J. Widmer, "802.11ad in smartphones: Energy efficiency, spatial reuse, and impact on applications," in *IEEE INFOCOM*, 2021, pp. 1–10.

[25] E. Khorov, A. Kiryanov, A. Lyakhov, and G. Bianchi, "A tutorial on IEEE 802.11ax high efficiency WLANs," *IEEE Communications Surveys Tutorials*, vol. 21, no. 1, pp. 197–216, 2019.

[26] A. Narayanan, X. Zhang, R. Zhu, A. Hassan, S. Jin, X. Zhu, X. Zhang, D. Rybkin, Z. Yang, Z. M. Mao, F. Qian, and Z.-L. Zhang, "A variegated look at 5G in the wild: Performance, power, and QoE implications," in *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, ser. SIGCOMM '21. ACM, 2021, p. 610–625.

[27] L. Donckers, P. Havinga, G. Smit, and L. Smit, "Enhancing energy efficient TCP by partial reliability," in *The 13th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, vol. 5, 2002, pp. 2424–2428 vol.5.

[28] M. Usman, D. Kliazovich, F. Granelli, P. Bouvry, and P. Castoldi, "Energy efficiency of TCP: An analytical model and its application to reduce energy consumption of the most diffused transport protocol," *International Journal of Communication Systems*, vol. 30, no. 1, p. e2934, 2017, e2934 IJCS-14-0785.R1. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/dac.2934

[29] M. Hashimoto, G. Hasegawa, and M. Murata, "An analysis of energy consumption for TCP data transfer with burst transmission over a wireless lan," *International Journal of Communication Systems*, vol. 28, no. 14, pp. 1965–1986, 2015.

[30] T. A. Le, C. S. Hong, M. A. Razzaque, S. Lee, and H. Jung, "ecMTCP: An Energy-Aware Congestion Control Algorithm for Multipath TCP," *IEEE Communications Letters*, vol. 16, no. 2, pp. 275–277, 2012.

[31] J. Zhao, J. Liu, and H. Wang, "On Energy-Efficient Congestion Control for Multipath TCP," in *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*, 2017, pp. 351–360.

[32] Y. Go, H. Noh, G. Park, and H. Song, "Energy-Efficient HTTP Adaptive Streaming with Hybrid TCP/UDP Over Heterogeneous Wireless Networks," in *2019 IEEE 20th International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM)*, 2019.

[33] A. Ghaffari, "Congestion control mechanisms in wireless sensor networks: A survey," *Journal of Network and Computer Applications*, vol. 52, pp. 101–115, 2015. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1084804515000557

[34] R. Bruschi, A. Lombardo, C. Panarello, F. Podda, E. Santagati, and G. Schembra, "Active window management: Reducing energy consumption of TCP congestion control," in *2013 IEEE International Conference on Communications (ICC)*, 2013, pp. 4154–4158.