

On the Low-latency Region of Best-effort Links for Delay-Sensitive Streaming Traffic

Boris Bellalta

Abstract—This letter analyzes the Low-latency Region (LLR) of a best-effort link (i.e., no traffic differentiation, and first come first serve scheduling) carrying both delay-sensitive (DS) streaming and non-delay-sensitive (NDS) background traffic. Moreover, inside the LLR, we show it exists a proportional fair arrival rate allocation for both the DS and NDS traffic streams. This optimal operating point results from maximizing a simple throughput-delay trade-off that considers the NDS traffic load, and the mean delay of the DS packets. To show how the presented trade-off could be used to allocate NDS traffic in a realistic scenario, we use Google Stadia traffic traces to generate the DS flow. Results from this use-case confirm that the throughput-delay trade-off also works regardless the distribution of the packet arrival and packet service times.

Index Terms—Low-latency, delay-sensitive traffic, proportional fair rate allocation, M/G/1

I. INTRODUCTION

The success of delay-sensitive (DS) audio and video streaming services such as voice over IP, real-time videoconferencing, cloud-gaming and virtual reality in best-effort networks depends to a great extent on the network’s ability to guarantee low end-to-end delays. This is especially challenging since, with the aim to keep the network neutrality, no traffic differentiation is usually provided between DS and non-delay-sensitive (NDS) background traffic, and packets are transmitted following a first come first serve scheduling in the multiple links they traverse from the source to the destination.

The interactions between DS streaming and NDS background traffic have been widely studied in the literature. In the one hand, the presence of NDS background streams results in higher delays and packet losses for the DS streams [1], [2]. However, in the other hand, NDS background traffic may also suffer from starvation in presence of DS streams. To mitigate this situation, DS streams also implement rate adaptation and congestion control solutions, such as the Google Congestion Control algorithm used in WebRTC [3], [4].

In the described context, we aim to find how much DS and NDS traffic can be allocated to a best-effort link without compromising the delay requirements of the DS traffic. The set of those traffic loads define what we call the Low-latency region (LLR). With respect to the delay requirements of a DS stream, we use the following criterion: *a DS stream operates inside the low-latency region if the average time that a DS packet spends in the link is lower than the mean inter-DS*

packet arrival time. Note that this is equivalent to say, that a new arriving packet must leave the system, on average, before another packet of the same flow arrives to it.

Using the aforementioned delay criterion, and modelling a best-effort link as a M/G/1 queue, we first calculate how much DS streaming and NDS background traffic can be allocated to the link. This gives us the upper limit of the LLR. Then, we show that inside the LLR it exists a proportional fair rate allocation for both DS and NDS traffic. This fair rate allocation corresponds to the point at which the difference between the throughput gain when NDS background traffic is added to the link, and the corresponding latency loss for the DS traffic is maximum. We will refer to this point as the proportional fair low-latency (PFLL) rate allocation.

The PFLL rate allocation can be achieved in practice by estimating the mean time between DS packet arrivals, the mean packet delay of the DS packets, and controlling the amount of NDS traffic that can be allocated to a link. It requires, however, to be able to identify and classify the active traffic streams, which can be done using Machine Learning techniques [5]. Once traffic streams are identified and classified, traffic shaping techniques can be used to control the amount of traffic that is allowed to enter the link, relying in the traffic source for adapting the traffic generation rate.

The use of the mean packet delay as a criterion to define the LLR instead of a certain percentile results in simple, closed-form, and insightful expressions. However, and more importantly, the use of the mean offers by itself an interesting trade-off between the carried load and the packet delay of a link in absence of any information regarding the delay requirements of the DS streaming traffic, neither on how many other links are traversed by the DS flow between its source and destination. The use of the mean, and this paper in general, was partially inspired by [6], where the authors play with the packet aggregation level in a Wi-Fi network to guarantee both high-throughput and low-latency.

Finally, we study the PFLL rate allocation when a link carries Google Stadia traffic [7], a cloud-gaming service that requires both high-throughput and low latency to perform satisfactorily. We aim to obtain how much NDS traffic can be allocated, and evaluate how much it disturbs the DS traffic stream in terms of the extra added latency. However, in addition to that, and more importantly, we confirm that we can estimate the PFLL rate allocation even if the traffic arrival process is not Poisson.

II. SYSTEM MODEL

We consider that a link consists of a buffer and a transmitter as shown in Fig. 1. The link capacity, i.e., the rate at which

B. Bellalta (*boris.bellalta@upf.edu*) is with Universitat Pompeu Fabra, Barcelona. The author would like to thank Marc Carrascosa for his contribution in taking the Google Stadia measurements. Also, the author wants to especially thank the anonymous reviewers for their constructive, insightful and challenging comments. This work was supported by WINDMAL PGC2018-099959-B-I00 (MCIU/AEI/FEDER,UE), and SGR017-1188 (AGAUR).

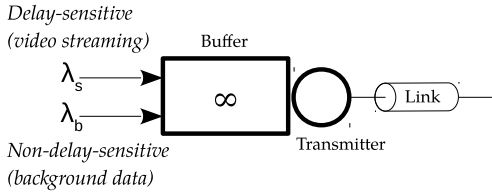


Fig. 1: Link model. The buffer and the transmitter correspond to the network interface.

the transmitter works, is of R bits/second. We assume the buffer is large enough to be considered of infinite size. We also assume all arriving flows have the same priority. Therefore, traffic differentiation is not applied, and all arriving packets are served following their order of arrival.

Besides, to keep the analysis simple, we do the following considerations:¹

- 1) There are only one DS flow and one NDS flow. The packet arrival process for both flows follows a Poisson process, with mean rates λ_s and λ_b packets/second, respectively. The mean aggregate packet arrival rate is given by $\lambda = \lambda_s + \lambda_b$.
- 2) The two flows have the same service distribution, with S the random variable representing the packet service time. Service times are independent and identically distributed. The service time distribution is characterized by its mean, $\mathbb{E}[S]$, and its coefficient of variation C_S . The mean transmission rate of the link in packets/sec. is $\mu = 1/\mathbb{E}[S]$.

Taking into account these considerations, the link is modelled as an M/G/1 queueing system [8]. The mean sojourn time for a packet of flow s , $\mathbb{E}[D_s(\lambda_s, \lambda_b)]$ and for a packet of flow b , $\mathbb{E}[D_b(\lambda_s, \lambda_b)]$, is the same:

$$\begin{aligned} \mathbb{E}[D_s(\lambda_s, \lambda_b)] &= \mathbb{E}[D_b(\lambda_s, \lambda_b)] = \mathbb{E}[S] + \frac{\lambda \mathbb{E}[S^2]}{2(1-a)} \\ &= \frac{1}{\mu} \left(1 + \frac{a}{1-a} \theta \right) = \frac{\Gamma(\lambda, \theta)}{\mu - \lambda}, \end{aligned} \quad (1)$$

where $\mathbb{E}[S^2] = \mathbb{E}^2[S](1 + C_S^2)$ is the second moment of the service time, $a = \lambda/\mu$ is the link utilization, and $\theta = (1 + C_S^2)/2$. Lastly, $\Gamma(\lambda, \theta) = 1 - a + a\theta$ is the ratio between the delay of a M/G/1 and a M/M/1 queue. Note that $a = a_s + a_b = \lambda_s/\mu + \lambda_b/\mu$, where a_s and a_b are the link utilization by the DS and NDS flows, respectively.

Note that the presented system model can represent either a point-to-point link of a network (e.g., the network interface of a router/switch) or a point-to-multipoint link (e.g., a WiFi Access Point with several associated clients). The only requirement is that the DS and NDS flows share the same buffer and transmitter.

III. LOW-LATENCY REGION

Following our previous definition, in the absence of NDS traffic ($\lambda_b = 0$), a DS flow s is working in the low-latency

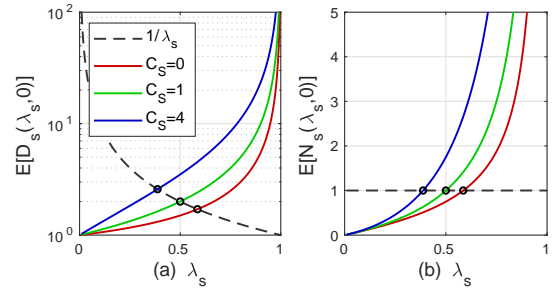


Fig. 2: (a) Mean DS packet delay and, (b) Mean number of DS packets in the link for different C_S values. The λ_s^+ values are indicated using black circles. The plots are obtained for $\mu = 1$ packets/second.

region of a link if the following condition is satisfied:

$$\frac{1}{\lambda_s} \geq \mathbb{E}[D_s(\lambda_s, 0)] = \frac{\Gamma(\lambda_s, \theta)}{\mu - \lambda_s}. \quad (2)$$

The highest value of λ_s that satisfies (2) is

$$\lambda_s^+ = \frac{\mu}{1 + \sqrt{\theta}}, \quad (3)$$

and therefore, the LLR includes all $\lambda_s \in [0, \lambda_s^+]$. Note that the value of λ_s^+ depends on the service time distribution of the DS traffic. For example, if the service time distribution is exponential, λ_s^+ is half of the link capacity.

Considering Little's theorem, (2) can be rewritten as:

$$\lambda_s \mathbb{E}[D_s(\lambda_s, 0)] = \mathbb{E}[N_s(\lambda_s, 0)] \leq 1,$$

where $\mathbb{E}[N_s(\lambda_s, 0)]$ is the mean number of DS packets in the link. We can observe that on average, the DS flow will be inside the LLR if the mean number of packets in the link is equal or less than one.

Fig. 2.a) shows the values of λ_s^+ as the intersection between $1/\lambda_s$ and $\mathbb{E}[D_s(\lambda_s, 0)]$ for different service time distributions. As expected, higher C_S values reduce the low-latency region. Fig. 2.b) confirms that $\mathbb{E}[N_s(\lambda_s, 0)] \leq 1$ when working inside the LLR.

IV. LOW-LATENCY RATE ALLOCATION

Let us consider now that an NDS flow b is added to the link, and so it shares the link resources with the DS flow s . We are then interested in finding how much NDS traffic can be admitted to the link while keeping the DS flow working inside the LLR.

In the following, we assume $\lambda_s \leq \lambda_s^+$ is known, and therefore, we focus on obtaining the value of λ_b . The other way around, i.e., given a certain amount of NDS traffic in a link, to obtain how much DS traffic can be allocated, gives exactly the same results, as expected.

A. Max low-latency rate allocation

A solution to the previous problem can be found by solving

$$\frac{1}{\lambda_s} \geq \mathbb{E}[D_s(\lambda_s, \lambda_b)] = \frac{\Gamma(\lambda, \theta)}{\mu - \lambda}.$$

¹In Sec. VI, we study a case where packets do not arrive following a Poisson process, and the DS and NDS flows have different service distributions.

Then, we can obtain λ_b^+ as the highest feasible NDS rate allocation, i.e.,

$$\lambda_b^+ = \left(\frac{\mu - \lambda_s}{\Gamma(\lambda_s, \theta)} \right) - \lambda_s = \frac{1}{\mathbb{E}[D_s(\lambda_s, 0)]} - \lambda_s, \quad (4)$$

which explicitly depends on the mean delay of the DS packets.

We can re-write (4) as $\lambda_b^+ = \mu - \kappa^+ \lambda_s$, with $\kappa^+ \lambda_s$ representing the minimum link capacity required by the DS flow to work in the LLR. The value of κ^+ is obtained by developing the term $\Gamma(\lambda_s, \theta)$ as follows:

$$\lambda_b^+ = \left(\frac{\mu - \lambda_s}{1 - \frac{1}{\mu}(\lambda_s - \lambda_s \theta)} \right) - \lambda_s = \mu \underbrace{\left(\frac{\mu - \lambda_s}{\mu - (\lambda_s - \lambda_s \theta)} \right)}_{=\alpha} - \lambda_s.$$

Then, we obtain κ^+ by considering that $\lambda_b^+ = \mu\alpha - \lambda_s = \mu - \kappa^+ \lambda_s$. Therefore,

$$\begin{aligned} \kappa^+ &= \frac{\mu}{\lambda_s}(1 - \alpha) + 1 = \frac{\mu}{\lambda_s} \left(\frac{\lambda_s \theta}{\mu - \lambda_s + \lambda_s \theta} \right) + 1 \\ &= 1 + \left(\frac{\theta}{\frac{\mu - \lambda_s + \lambda_s \theta}{\mu}} \right) = 1 + \frac{\theta}{\Gamma(\lambda_s, \theta)}. \end{aligned}$$

Let us now define β to refer to the second term of κ^+ , i.e., $\beta = \frac{\theta}{\Gamma(\lambda_s, \theta)}$. Note that $\beta \lambda_s$ is the amount of link capacity that has to remain unused to keep the DS flow operating inside the LLR. For example, considering the service time is exponentially distributed, we have that $\kappa^+ = 2$, and $\lambda_b^+ = \mu - 2\lambda_s$. In this case, the amount of unused link capacity is equal to the load of the DS flow, as $\beta = 1$.

Finally, observe that $\kappa^+ \lambda_s^+ = \mu$. This result could be expected since for λ_s^+ we have that $\lambda_b = 0$. In detail, by expanding the terms κ^+ and λ_s^+ , we have that

$$\left(1 + \frac{\theta}{\Gamma(\lambda_s^+, \theta)} \right) \frac{\mu}{1 + \sqrt{\theta}} = \mu,$$

where

$$\frac{\theta}{\Gamma(\lambda_s^+, \theta)} = \frac{\theta}{1 - \frac{1}{1 + \sqrt{\theta}} + \frac{\theta}{1 + \sqrt{\theta}}} = \frac{\theta(1 + \sqrt{\theta})}{\theta + \sqrt{\theta}} = \sqrt{\theta}.$$

This last result tells us that when we are operating in the limit of the low-latency region, the amount of idle bandwidth is equal to $\lambda_s^+ \sqrt{\theta}$, which further shows the relationship between the service time distribution and the LLR.

B. Proportional fair low-latency rate allocation

Let $\lambda_b^* \leq \lambda_b^+$ be the proportional fair low-latency rate allocation, i.e., the value at which the trade-off between the delay of DS packets and the NDS throughput is maximum. $\lambda_b > \lambda_b^*$ values result in a higher delay increase for DS packets than the link throughput gain. Similarly, for $\lambda_b < \lambda_b^*$ values, we observe the opposite result.

To find λ_b^* , we formulate the throughput-delay trade-off as the difference between the link throughput gain, and the delay loss for the DS traffic,

$$g(\lambda_s, \lambda_b) = G_T(\lambda_s, \lambda_b) - G_D(\lambda_s, \lambda_b), \quad (5)$$

with respect to the case there is no NDS traffic.

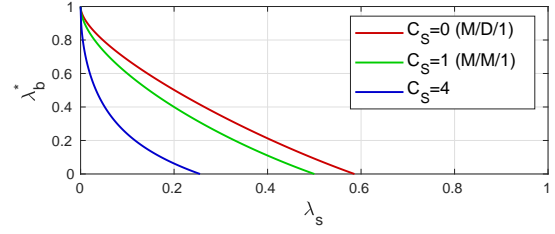


Fig. 3: Proportional Fair Low-latency rate allocation for different C_S values and $\mu = 1$.

The NDS throughput gain is

$$G_T(\lambda_s, \lambda_b) = \frac{(\lambda_b + \lambda_s) - \lambda_s}{\lambda_s} = \frac{\lambda_b}{\lambda_s},$$

and the DS delay loss is

$$\begin{aligned} G_D(\lambda_s, \lambda_b) &= \frac{\mathbb{E}[D_s(\lambda_s, \lambda_b)] - \mathbb{E}[D_s(\lambda_s, 0)]}{\mathbb{E}[D_s(\lambda_s, 0)]} \\ &= \frac{\frac{1}{\mu} + \frac{1}{\mu} \frac{a}{(1-a)} \theta - \frac{1}{\mu} - \frac{1}{\mu} \frac{a_s}{(1-a_s)} \theta}{\frac{1}{\mu} + \frac{1}{\mu} \frac{a_s}{(1-a_s)} \theta} \\ &= \frac{\beta(\lambda - \lambda_s)}{\mu - \lambda} = \frac{\beta \lambda_b}{\mu - \lambda_s - \lambda_b}. \end{aligned}$$

Then, (5) results in

$$g(\lambda_s, \lambda_b) = \frac{\lambda_b}{\lambda_s} - \frac{\beta \lambda_b}{\mu - \lambda_s - \lambda_b} = \lambda_b \left(\frac{1}{\lambda_s} - \frac{\beta}{\mu - \lambda} \right). \quad (6)$$

Finally, we are interested in finding

$$\lambda_b^* = \operatorname{argmax}_{\lambda_b} g(\lambda_s, \lambda_b),$$

that is the proportional fair low-latency rate allocation.²

Since $g(\lambda_s, \lambda_b)$ is concave, and has its maximum in the range $\lambda_b \in [0, \lambda_b^+]$, we find λ_b^* by deriving (6) with respect to λ_b . The first derivative of $g(\lambda_s, \lambda_b)$ is

$$\frac{dg(\lambda_s, \lambda_b)}{d\lambda_b} = \frac{1}{\lambda_s} - \frac{\beta(\mu - \lambda_s)}{(\mu - \lambda_s - \lambda_b)^2},$$

and it is equal to zero for

$$\begin{aligned} \lambda_b^* &= \mu - \lambda_s - \sqrt{\beta \lambda_s (\mu - \lambda_s)} \\ &= \mu - \lambda_s \left(1 + \sqrt{\beta \left(\frac{\mu - \lambda_s}{\lambda_s} \right)} \right) = \mu - \kappa^* \lambda_s, \quad (7) \end{aligned}$$

with $\kappa^* = 1 + \sqrt{\beta \left(\frac{\mu - \lambda_s}{\lambda_s} \right)}$. Similarly to the max NDS rate allocation, the $\kappa^* \lambda_s$ term is the link capacity required by the DS flow to work at the PFLR rate allocation. Note that $\kappa^* \geq \kappa^+$ for $\lambda_s \leq \lambda_s^+$.

Fig. 3 shows the PFLR rate allocation (λ_b^* values) for different λ_s and C_S values. For low values of λ_s , almost full link capacity can be achieved. However, it rapidly reduces when λ_s increases. Also, increasing the variability of the service time distribution reduces the PFLR rate allocation.

²Note that just taking logs in (6) we get $\log(\lambda_b) + \log\left(\frac{1}{\lambda_s} - \frac{\beta}{\mu - \lambda}\right)$, which satisfies the definition of proportional fairness, i.e., the gain (increase) in one variable should be higher than the loss (decrease) in the others, which is achieved by maximizing the sum of the logarithms of the variables under study.

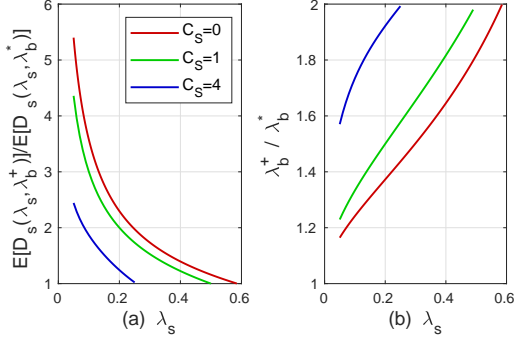


Fig. 4: (a) Delay reduction vs (b) Throughput increase between max and PFL rate allocation strategies, with $\mu = 1$.

C. Max vs PFL rate allocation

To provide some further insights on the significance of the max and PFL rate allocation strategies, Fig. 4 compares $\mathbb{E}[D_s(\lambda_s, \lambda_b^+)]$ and $\mathbb{E}[D_s(\lambda_s, \lambda_b^*)]$, as well as the values of λ_b^+ and λ_b^* . It can be observed that the PFL rate allocation results in a significant delay reduction (i.e., values higher than 1 in Fig. 4.a)) for a minimum background traffic loss (values higher than 1 in Fig. 4.b)) when the load of the DS traffic stream is low. For example, for $C_S = 0$ and $\lambda_s \approx 0.1$, the delay is 5 times lower, while the throughput is reduced only by a factor of 1.2. Increasing the load of the DS stream reduces the gain in delay and increases the loss in NDS throughput.

V. AN ALTERNATIVE FORMULATION

In this section, we show that the following alternative formulation to (6),

$$f(\lambda_s, \lambda_b) = \lambda_b (\mathbb{E}[D_s(\lambda_s, \lambda_b^+)] - \mathbb{E}[D_s(\lambda_s, \lambda_b)]), \quad (8)$$

can be used to accurately estimate the PFL rate allocation. Note that (8) is simply the product between the NDS throughput (the gain), and the difference in delay between the maximum tolerable and the current delay of the DS packets (the loss). To obtain it we have simply approximated the term $\beta/(\mu - \lambda)$ in (6) by $E[D_s(\lambda_s, \lambda_b)]$.

To evaluate the accuracy of (8), we compare the normalized versions of $\hat{g}(\lambda_s, \lambda_b) = \frac{g(\lambda_s, \lambda_b)}{\max(g(\lambda_s, \lambda_b))}$ and $\hat{f}(\lambda_s, \lambda_b) = \frac{f(\lambda_s, \lambda_b)}{\max(f(\lambda_s, \lambda_b))}$, and the first derivatives of $g(\lambda_s, \lambda_b)$ and $f(\lambda_s, \lambda_b)$. Fig. 5.a) shows that both $\hat{g}(\lambda_s, \lambda_b)$ and $\hat{f}(\lambda_s, \lambda_b)$ give the same values. Fig. 5.b) shows that even if the first derivatives of $g(\lambda_s, \lambda_b)$ and $f(\lambda_s, \lambda_b)$ are not the same, they are equal to 0 for the same λ_b value.

VI. USE-CASE: GOOGLE STADIA TRAFFIC

In this section, we aim to illustrate the existence of the PFL rate allocation, and the applicability of (8), when the aggregate traffic arrival process is not Poisson, and the DS and NDS traffic flows are characterized by different traffic arrival processes and service time distributions. We also examine the cumulative distribution function (cdf) of the delay of the DS packets, i.e., $D_s(\lambda_s, \lambda_b)$, for different values of λ_b to observe how it changes when the NDS background traffic increases.

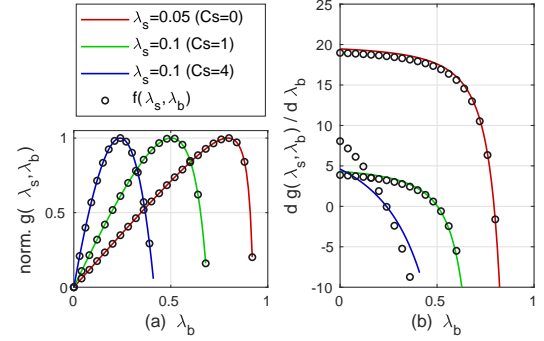


Fig. 5: (a) $\hat{g}(\lambda_s, \lambda_b)$ vs $\hat{f}(\lambda_s, \lambda_b)$, and (b) $\frac{dg(\lambda_s, \lambda_b)}{d\lambda_b}$ vs $\frac{df(\lambda_s, \lambda_b)}{d\lambda_b}$ for different λ_s and C_S pair of values, with $\mu = 1$.

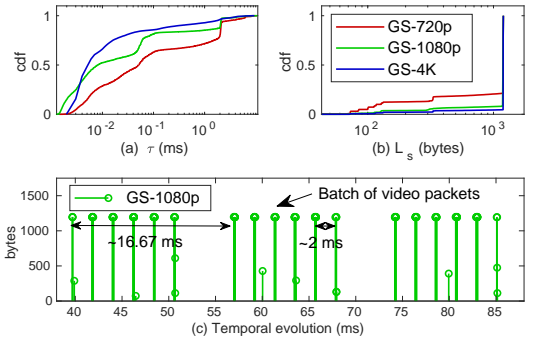


Fig. 6: Google Stadia's Tomb Raider traffic: cdfs of the (a) inter-packet arrival time (τ), and (b) packet size (L_s), and (c) snapshot of the GS traffic temporal evolution.

We implemented the system model described in Section II in C++ using the COST simulation library [9]. The DS traffic is generated using a set of traffic traces collected while playing with Google Stadia's Tomb Raider (GS).³ The traces represent the downlink traffic (mostly video contents, from the server to the client, sent at a rate of 60 frames/second) for the three different video resolutions available in GS (720p, 1080p, and 2160p). The duration of each trace is 30 seconds. Their main characteristics are shown in Fig. 6 and in Table I. We can observe that GS packets arrive in batches of mean size $\mathbb{E}[\sigma]$. In those conditions, packet arrivals are not Poisson, even if the coefficient of variation of the inter-packet arrival time, C_τ , for 720p and 1080p is close to 1. NDS traffic arrives to the link following a Poisson process. NDS packet sizes are assumed to be exponentially distributed, with an average size of $\mathbb{E}[L_b] = 10000$ bits. The capacity of the link is set to $R = 100$ Mbps.

Fig. 7.a) shows the value of $\hat{g}(\lambda_s, \lambda_b)$ computed from the simulation data for the three GS video resolutions. For each resolution, we also plot $\hat{f}(\lambda_s, \lambda_b)$, and indicate where the maximum of $\hat{g}(\lambda_s, \lambda_b)$ and $\hat{f}(\lambda_s, \lambda_b)$ is (circle). We can observe that 1) the PFL rate allocation exists even if the traffic arrival process is not Poisson, and 2) we confirm (8) is an accurate estimator of the PFL rate allocation. In Fig. 7.b), we plot the mean packet delay for GS packets, $\mathbb{E}[D_s(\lambda_s, \lambda_b)]$, indicating

³The traces can be found at <https://www.upf.edu/web/wnrg/wn-datasets>.

Resolution (unit)	Load (Mbps)	$\mathbb{E}[\tau]$ (ms)	C_τ	$\mathbb{E}[L_s]$ (Bytes)	C_S	$\mathbb{E}[\sigma]$ (packets)
720p (HD)	10.25	1.700	0.97	997.5	0.40	2.18
1080p (FHD)	27.47	1.417	0.94	1123.2	0.23	4.33
2160p (4K)	39.89	1.293	2.87	1144.2	0.19	5.74

TABLE I: Characteristics of Tom Raider Downlink traffic. $\mathbb{E}[L_s]$ is the mean packet size, $\mathbb{E}[\tau]$ is the mean inter-packet arrival time, $\mathbb{E}[\sigma]$ is the mean batch size considering all type of arriving packets, and C_τ and C_S are the coefficient of variation of the inter-packet arrival time and service time, respectively.

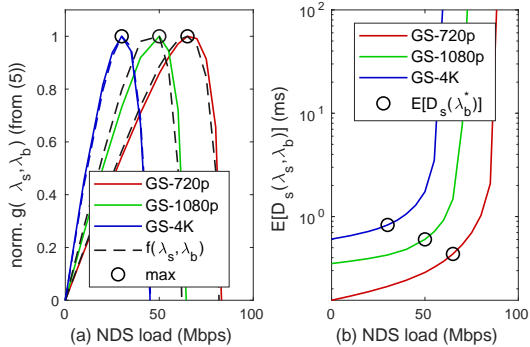


Fig. 7: (a) $\hat{g}(\lambda_s, \lambda_b)$, and (b) $\mathbb{E}[D_s(\lambda_s, \lambda_b)]$, for different GS video resolutions.

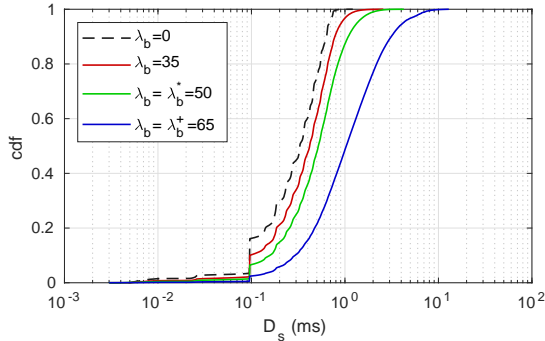


Fig. 8: cdf of $D_s(\lambda_s, \lambda_b)$ for a GS video resolution of 1080p, and different λ_b values.

the delay that corresponds to the PFLR rate allocation (circle). The proportional fair NDS rate allocation for the best-effort flow is 65, 50 and 30 Mbps, for 720p, 1080p, and 2160p video resolutions, respectively. Note that in case the DS streaming traffic load changes because GS decides to switch to a different resolution, the PFLR rate allocation does it also accordingly.

Fig. 8 shows the cdf of the packet delay for GS traffic, $D_s(\lambda_s, \lambda_b)$, for different λ_b values. When $\lambda_b > \lambda_b^*$, the negative effect of the NDS traffic on the DS delay is significant. Instead, for $\lambda_b \leq \lambda_b^*$, the cdfs are relatively similar to the case without NDS traffic (black dashed line), hence showing also the benefits of operating close to λ_b^* . For example, the 90th-percentile of $D_s(\lambda_s, \lambda_b)$ increases by a factor of 1.7x from $\lambda_b = 0$ to $\lambda_b = \lambda_b^*$, and by a factor of 4.9x to $\lambda_b = \lambda_b^+$, while the gain in NDS traffic from λ_b^* to λ_b^+ is only of the 30 %.

In summary, the presented use-case illustrates how the proposed LLR and rate allocation strategies can be used to

characterize the interactions between DS and NDS traffic in best-effort links, providing a framework to better understand the trade-off between capacity and latency in communication networks.

VII. CONCLUSIONS

This paper defines the low-latency region of a best-effort communication link. It also provides two strategies to allocate the rate of NDS background traffic to the link without negatively affecting the DS streaming traffic: the max and PFLR rate allocation strategies. Moreover, the presented analysis shows how the LLR and rate allocation strategies depend on the incoming traffic characteristics (i.e., inter-packet delay, and service time distribution), and DS and NDS traffic loads, providing insights on aspects such as the amount of link capacity that must remain empty to preserve the low-latency operation.

Future work may consider to re-define the low-latency region and rate allocation strategies in terms of delay percentiles, as well as to consider the presence of multiple DS and NDS traffic flows with different arrival and service time distributions. Also, while we have shown that (8) works well when the packet arrival process is not Poisson, at least in the considered use-case, we conjecture it may also work when more complex transmission schemes are considered, such as the cases of packet aggregation and multiuser transmissions in WiFi networks. Indeed, we have special interest in applying the results presented in this letter to deal with latency in future WiFi networks [10]. Moreover, it is also worth to study how the use of the proposed rate allocation strategies could improve the end-to-end latency in multi-hop networks, as well as to investigate the interplay with end-side congestion control techniques.

REFERENCES

- [1] Luca De Cicco, Saverio Mascolo, and Vittorio Palmisano. Skype video congestion control: An experimental investigation. *Computer Networks*, 55(3):558–571, 2011.
- [2] Gaetano Carlucci, Luca De Cicco, and Saverio Mascolo. Controlling queuing delays for real-time communication: the interplay of E2E and AQM algorithms. *ACM SIGCOMM Computer Communication Review*, 46(3):1–7, 2018.
- [3] Gaetano Carlucci, Luca De Cicco, Stefan Holmer, and Saverio Mascolo. Congestion control for web real-time communication. *IEEE/ACM Transactions on Networking*, 25(5):2629–2642, 2017.
- [4] Alberto Alós, Francisco Morán, Pablo Carballeira, Daniel Berjón, and Narciso García. Congestion Control for Cloud Gaming Over UDP Based on Round-Trip Video Latency. *IEEE Access*, 7:78882–78897, 2019.
- [5] Thuy TT Nguyen and Grenville Armitage. A survey of techniques for internet traffic classification using machine learning. *IEEE communications surveys & tutorials*, 10(4):56–76, 2008.
- [6] Francesco Gringoli and Douglas J Leith. Regulating Queuing Delay in 802.11ac WLANs: Nonlinear Controller Analysis and Design. *arXiv preprint arXiv:1910.09651*, 2019.
- [7] Marc Carrascosa and Boris Bellalta. Cloud-gaming: Analysis of Google Stadia traffic. *arXiv preprint arXiv:2009.09786*, 2020.
- [8] D. P. Bertsekas and R. G. Gallager. *Data networks*. (2nd edition) Prentice-Hall, 1992.
- [9] Gilbert Chen and Boleslaw K Szymanski. Cost: A component-oriented discrete event simulator. In *Proceedings of the Winter Simulation Conference*, volume 1, pages 776–782. IEEE, 2002.
- [10] Toni Adame, Marc Carrascosa, and Boris Bellalta. Time-Sensitive Networking in IEEE 802.11be: On the Way to Low-latency WiFi 7. *arXiv preprint arXiv:1912.06086*, 2019.