

Less-than-Best-Effort capacity sharing over high BDP networks with LEDBAT

Nicolas Kuhn^{1,2}, Olivier Mehani², Arjuna Sathaseelan³, Emmanuel Lochin¹

¹Université de Toulouse, ISAE, TéSA, France

²National ICT Australia (NICTA), Australia

³Computer Laboratory, University of Cambridge, United Kingdom

Abstract—There has been a renewed interest at the Internet Engineering Task Force (IETF) in using Less-than-Best Effort (LBE) methods for background applications. IETF recently published a RFC for Low Extra Delay Background Transport (LEDBAT), a congestion control algorithm for LBE transmissions. This paper provides an analysis of LEDBAT performance over congested large bandwidth \times delay product (LBDP) networks, and assesses the validity of having a fixed target queuing time. In particular, we lead a study of the impact of this target queuing delay when LEDBAT is used over 4G satellite networks. The rationale is to explore the possibility to grab the unused 4G satellite links’ capacity to carry non-commercial traffic. We show that this is achievable with LEDBAT. However, depending on the fluctuation of the load, performance improvements could be obtained by properly setting the target value. We generalize this evaluation over different congested LBDP networks and confirm that the target value might need to be adjusted to networks’ and traffic’s characteristics. Further work will study whether and how this parameter should be dynamically adapted, and LEDBAT’s congestion control improved.

I. INTRODUCTION

Recently there has been a renewed interest in exploring Less-than-Best Effort (LBE) access in the Internet research community and standards bodies. LBE, also known as the *Scavenger class* of traffic, came into existence almost a decade ago with work being carried out at Internet2 [1]. Recently P2P and other bulk traffic have been pointed out as some of the root causes of the BufferBloat problem [2], due to large customer premise equipment (CPE) router queues. This problem, mainly caused by router buffering packets for a long period instead of dropping them, impacts real-time traffic which is getting more and more pervasive today [3]. As a result, the Internet Engineering Task Force (IETF) has started focusing on LBE congestion methods [4] to transmit background data. In particular, a recent paper [5] proposes the use of LBE access to provide free Internet access. The idea is to leverage the unused capacity to carry signaling or non-commercial traffic with an LBE protocol.

This paper follows this idea and aims to explore the performance of the Low Extra Delay Background Transport (LEDBAT) [6] over large bandwidth \times delay product (LBDP) networks. We specifically verify the use of LEDBAT to transmit LBE traffic over congested satellite networks and identify the performance implications of LEDBAT traffic sharing the network with other widely used congestion controlled

transport protocols. Indeed, the authors of [7] have shown that LEDBAT is unfair with TCP¹ when the BDP is large (*e.g.*, RTT of 100ms and capacity of 600Mbps). In this paper, we illustrate that this unfairness problem can be solved: we assess different parametrizations for LEDBAT’s “target queuing delay” to increase fairness to TCP on LBDP paths (3G/4G satellite).

The rest of the paper is organised as follows. In Section II, we present LEDBAT’s congestion control algorithm. We justify that this protocol is an ideal candidate for LBE background transmissions in Section III. We propose simulations in 4G satellite contexts in Section IV where we show that LEDBAT’s queuing time target has an impact on performance. We generalize this evaluation over different congested LBDP networks in Section V to further assess this impact. In Section VI, we propose a discussion of the portability of the results presented in this article, which we conclude in Section VII.

II. RELATED WORK

Various congestion control mechanisms have been pointed as good candidates to support LBE traffic, such as the delay-based TCP Vegas [8], or NF-TCP [7]; a more complete survey can be found in [4]. Due to its recent standardization at IETF, LEDBAT [6] however seems to be the most promising LBE mechanism. In this section, we detail the algorithm of LEDBAT congestion control. We also present the latest investigations on this mechanism.

LEDBAT is characterized by the following parameters: target queuing delay (τ), impact of the delay variation ($\gamma = 1/\tau$), minimum One-Way Delay (D_{min}) and current One-Way Delay (D_{ack}). For each ACK received, the new congestion window ($cwnd$) value is updated according to:

$$cwnd = cwnd + \frac{\gamma(\tau - (D_{ack} - D_{min}))}{cwnd} \quad (1)$$

LEDBAT congestion control is based on queuing delay variations (*i.e.*, the queuing delay is used as a primary congestion notification), estimated by $(D_{ack} - D_{min})$. When the size of the queue is large ($\tau < (D_{ack} - D_{min})$), LEDBAT reduces its congestion window. Therefore, the target queuing delay τ

¹It is worth noting that they use TCP Reno, which is known not to be aggressive enough on LBDP paths.

embodies the maximum queuing time that LEDBAT is allowed to introduce.

In [9], the authors describe the motivations behind LEDBAT development and conduct the first known performance evaluation of the LEDBAT algorithm. In [10], the authors develop a fluid model of the congestion window and assess that LEDBAT operates under a wide range of parameters. Following the results presented in that study, LEDBAT’s RFC [6] state that γ must be set at $1/\tau$ or less, and τ must be lower than 100 ms.

When a sender using the LEDBAT method sends its first packet, if the network is loaded, the minimum queuing delay can be overestimated causing the maximum value of the One-Way Delay (D_{max} , estimated as $D_{max} = D_{min} - \tau$) to be higher than other LEDBAT flows that were already transmitting data. This bad estimation of D_{min} introduces what is called the “latecomer’s advantage.” In [11], the authors illustrate this phenomenon and propose a multiplicative-decrease solution to this problem. However, the RFC ignores this problem stating that “system noise may sufficiently regulate the latecomer’s advantage.”

In [12], the authors identify the negative impact of route changes on the performance on LEDBAT. The paper provides an analysis of the phenomenon without concrete solution to the problem. As for the late-comer’s advantage, this problem is linked to an overestimation of the minimum queuing delay when the first packet is transmitted. No solution has yet been proposed to overcome this problem.

We now investigate the impact of different target queuing delays on LEDBAT’s ability to use the remaining capacity, without disturbing primary traffic, of LBDP paths.

III. LEDBAT VERSUS TCP VEGAS FOR LBE TRANSMISSIONS

In this section, we analyze the performance of LEDBAT over an LBDP scenario. As TCP Vegas [8] does not perform well when mixed with other TCP variants, it could be a good alternative candidate for transmitting LBE traffic. The objective of this section is therefore to justify that LEDBAT is a better candidate. We run simulations with *ns-2* and use the LEDBAT module validated in [9]. We checked that this module has been developed in accordance with the RFC.

We model an LBDP link in *ns-2*: the capacity is set to 10 Mbps and the path delay is set to 250 ms. We consider two competitive flows. We focus on the impact of the introduction of a secondary LBE flow (either TCP Vegas or LEDBAT) when the primary flow has reached full capacity. The primary flow transmits data for 800 s with CUBIC [13] at the transport layer. The secondary flow starts from 500 s to 800 s. The DropTail queue size is considered as infinite (we fixed it to a large value), the IP packets size is 1500 bytes.

We present the combination of the different flows used in the simulation and their respective throughput in Table I. For both flows, we report the mean throughput measured over the simulation period.

When CUBIC is the only flow on the link, it occupies 95.875% of the capacity (Case 1). Introduction of the LEDBAT

TABLE I
COMPARISON OF LEDBAT AND VEGAS FAIRNESS TO CUBIC

	Transport Protocol	Target (ms) (for LEDBAT)	Throughput (% of capacity)
Case 1	Flow 1: CUBIC	–	95.875
	Flow 2: NONE	–	–
Case 2	Flow 1: CUBIC	–	90.050
	Flow 2: TCP Vegas	–	6.150
Case 3	Flow 1: CUBIC	–	95.606
	Flow 2: LEDBAT	25	0.281
Case 4	Flow 1: CUBIC	–	95.665
	Flow 2: LEDBAT	100	0.222
Case 5	Flow 1: TCP Vegas	–	85.120
	Flow 2: LEDBAT	100	0.215

flow causes a 0.02% reduction of the capacity occupied by the CUBIC flow (Cases 3–4).

We also note that TCP Vegas exploits 6% of the capacity (more than LEDBAT), but the percentage of the capacity occupied by CUBIC decreases by 5.8% (Case 2). We conclude that even if TCP Vegas takes up less capacity, this protocol shows more aggressiveness compared to LEDBAT. TCP Vegas is also more aggressive than LEDBAT in terms of link capacity utilization when they are the two protocols involved in the simulation (Case 5).

Therefore, we believe that LEDBAT is a better candidate than TCP Vegas to transmit LBE traffic over long delay paths without introducing congestion nor severely affecting the other competing flows sharing the same path. The results gathered in Table I illustrates that for a LBDP link, the queuing target of LEDBAT has an impact on the link utilization (Cases 3–4). In the following sections, we further explore the impact of this value in variously loaded satellite networks.

IV. LEDBAT OVER A 4G SATELLITE NETWORK

In this section, we explore the impact of the target queuing delay specifically focusing on the performance of LEDBAT in a 4G satellite network. We consider a mobile receiver and assess the performance of LEDBAT over the satellite.

A. 4G Satellite Network Configuration

To drive this experiment, we use a *ns-2* extension called Cross-Layer InFormation Tool (CLIFT) [14] allowing to play real physical layer traces inside *ns-2*. The 4G satellite link trace used was provided by CNES.²

The simulations for this scenario represent the communication between a single mobile user and a satellite gateway. We focus on CUBIC as it is now enabled by default in GNU/Linux and Android systems. The mobile user sends data to the satellite gateway using CUBIC or LEDBAT at the transport layer and retransmission mechanism (ARQ) at the

²CNES is a government agency responsible for shaping and implementing France’s space policy in Europe, see <http://www.cnes.fr/>.

link layer level. As before, the queue is large enough not to be overflowed, the IP packets size is 1500 bytes.

The physical trace is characterized as follows: wave form: LTE S-band; OFDM: capacity=5MHz with 300 available frequencies; FFT length: 512; available capacity: 2.3 Mbps (1 user); turbo code: 3GPP, word length (before coding): 33 bytes; interleaving depth: 36 ms; suburban satellite channel, GEO orbit, elevation angle 40°; mobile user speed: 60km/h, distance traveled: 8km.

We aim to study the impact of LEDBAT on competing CUBIC flows and its ability to exploit capacity when the network is not fully loaded. The simulation lasts 450 s. We consider that the mobile receiver transmits data with a CUBIC protocol from 0 s to 225 s and from 270 s to 450 s. From 112.5 s to 337.5 s, data is transmitted with a LEDBAT protocol. Based on LEDBAT’s RFC [6], we consider a representative set of target values $\tau \in [5; 15; 25; 100]$.

B. Simulation Results

TABLE II
LEDBAT OVER 4G SATELLITE

τ	Protocol	Capacity at different times (kbps)			
		112.5–225	225–270	270–337.5	337–450
5 ms	CUBIC	2292	329	2191	2298
	LEDBAT	0	50	10	0.1
15 ms	CUBIC	2292	329	2190	2298
	LEDBAT	0	32	8	0.1
25 ms	CUBIC	2292	329	2190	2298
	LEDBAT	0	30	8	0.1
100 ms	CUBIC	2292	329	2190	2298
	LEDBAT	0	25	8	0.1

We present the results for this scenario in Table II. When a CUBIC flow attempts to send data ($t \in [112.5; 225]$ or $\in [270; 337.5]$), the LEDBAT flow does not manage to transmit data. When the primary flow does not transmit ($t \in [225; 270]$), LEDBAT flow uses this opportunity for its own traffic. The LEDBAT flow can not use the whole available capacity, due to its low aggressivity: after $t = 225$ s, there are still CUBIC packets in the queue that are waiting to be transmitted and between $t \in [225; 270]$ there are not enough receiver’s feedbacks for LEDBAT congestion control to fastly increase its congestion window. Also, the smaller the target queuing delay is, the more data the LEDBAT flow transmits during this less loaded period.

As a result, we consider that LEDBAT can be a very good candidate for LBE data transfer, using capacity when some is available but gracefully retracting when primary traffic is present. This also illustrates that the target queuing delay has an impact on the performance of LEDBAT. Decreasing this value allows LEDBAT to use the free capacity more efficiently. We propose, in the next section, to verify this statement and assess the impact of the target value in a more generic context, where the LBDP link is introduced in a loaded network.

V. LEDBAT PERFORMANCE IN A LOADED SATELLITE NETWORK

In this section, we assess the impact of the number of LEDBAT flows and their target queuing delay depending on the capacity of the satellite path left over by the primary traffic.

A. Network configuration

As detailed in Fig. 1, we consider a simple architecture where the bottleneck is the satellite link. Three type of competitive flows transmit data to the Receiver 1. Each application is a file transfer using CUBIC as transport protocol. We consider L LEDBAT transmitters with $L \in [1; 10; 25; 50]$ and the same set of $\tau \in [5; 15; 25; 100]$.

In order to assess how LEDBAT exploits the freed capacity when other transports reduce their rates, we need to introduce a limiting factor for congestion losses to occur. To do so, the queue at the gateway is fixed to 50 IP packets (*i.e.*, maximum queuing delay 120 ms, which is higher than the target value). The IP packets size is still 1500 bytes and the AQM mechanism is DropTail.




Fig. 1. Network architecture

The links between the different transmitters (Link 1 in the figure) are defined by a capacity of 5 Mbps and a random delay $d1 \in [20; 50]$ ms. The satellite link (Link 2 in the figure) has a capacity of 5 Mbps and a delay of 250 ms. The simulation lasts for 300 s. The load variations on the network are presented in Table III. In order to avoid the late-comers problems introduced by competing LEDBAT flows, all LEDBAT flows start at the same time. Also, we introduce different groups of CUBIC flows to obtain a controllable fluctuating traffic that enables us to better understand LEDBAT behavior.

B. Presentation of the results: few users in the network

We only consider transmitters from groups B and C (details in section V-A). To better assess the performance of the flows, we compute the goodput measured at the end of the whole simulation (*i.e.*, the amount of useful data transmitted).

We present, in Fig. 2, the results looking at the percentage of the capacity exploited by the LEDBAT flows, the CUBIC flows and the overall utilized capacity.

TABLE III
SIMULATION PARAMETERS

Transmitter A1, ... AX, B1, ... BY, C1, ... CZ		
Group	Nb Flows	Transmission times (s)
Group A	100	[0;300]
Group B	100	[0;30],[60;90],[180;210],[240;270]
Group C	100	[0;75],[150;225]
Group Ledbat	[1;10;25;50]	[0;90]




Fig. 2. Capacity sharing depending on the target value and the number of flows (without Group A)

First of all, we can clearly evaluate that the introduction of one LEDBAT flow reduces the percentage of the capacity used by the CUBIC flows. Considering the fact that the LEDBAT flows introduces a decrease of the CUBIC's flows capacity (due to congestion), but also might greatly increase the utilization of the capacity of the link, we try to assess a trade-off between these two considerations depending on the target value and the number of LEDBAT flows. Introducing flows in the network increases the utilization of the central link, but we focus on the fact that adapting the target value of LEDBAT enables to (1) send more LBE data and (2) less impact on the principal traffic.

When the target queuing delay is increased and when the number of flows are fixed, we can see in this figure that (1) LEDBAT flows exploit less capacity to transmit data, (2) the capacity utilization of CUBIC flows decreases. When the value of the target changes from 100ms to 5ms and the number of LEDBAT flows is set to 50: (1) the capacity used by LEDBAT flows increases by 5%; (2) the capacity used by CUBIC flows increases by 2%; (3) the utilization of the link increases by 7%. Therefore, considering 50 LEDBAT flows and the network configuration detailed above, changing the target value from 100ms to 5ms enables to increase the use of the capacity by 7%.

As in the previous paragraph, we considered a fixed number of flows, we consider the benefits and impacts of increasing

the number of flows. We can also see that when the target queuing delay is set to 5ms, the cost of 8% of the capacity for the CUBIC flows can enable to introduce 50 LEDBAT flows that will exploit 28% of the capacity. As a result, introducing 50 LEDBAT flows with a target value of 5ms, the utilized capacity of the central link increases by 20%. When the target queuing delay is set to 100ms, the cost of 11% of the capacity for the CUBIC flows can enable to introduce 50 LEDBAT flows that will exploit 22% of the capacity. In this case, the utilized capacity increases by 11%. Therefore, changing the target value from 100ms to 5ms (1) cost 6% less of the principal flows capacity, (2) provides 6% for LEDBAT flows, (3) increase the use of the central link by 9%.

We can conclude that, in the context of a high delay path, the introduction of LEDBAT flows is optimized when its target value is set to 5ms. Setting this parameter to 5ms enables to introduce a large number of LEDBAT flows by greatly increasing the capacity utilization of the long delay link at a low cost for the CUBIC flows.

C. Presentation of the results: fully loaded network

In this section, we consider that all groups A, B and C (details in section V-A) transmit data. The aim of this section is to assess if we can propose the same conclusions as the one presented in section V-B when the long delay link is fully loaded.




Fig. 3. Capacity sharing depending on the target value and the number of flows (with Group A)

In Fig. 3, we present the results in terms of used capacity (presented in percentage of the available capacity). We can assume that the capacity of the link is fully exploited by the CUBIC flows even if the network is highly loaded: this is due to congestion control.

In this context, we can note that there is an impact of the target value. Indeed, when the target queuing delay is increased and the number of LEDBAT flows is fixed (1) more capacity is exploited by the LEDBAT flows, (2) this capacity used by LEDBAT flows is directly taken from the CUBIC flows

capacity, (3) there are no significant benefits in terms of overall used link capacity. Indeed when the target value changes from 5 ms to 100ms, the capacity used by 50 LEDBAT flows increases by 5%, but the capacity used by the principal flows decreases by the same 5%. Thus, the congestion provided by 50 LEDBAT flows has a negative effect on the capacity dedicated to the principal flows, but its impact is less when the target value is set to 5 ms.

As a consequence in this loaded network with a high delay-bandwidth product link, there is few remaining capacity that LEDBAT flows could exploit to transmit data. When the network is highly congested, introducing LEDBAT flows with a higher target queuing delay increase the congestion and do not increase the utilization of the capacity. The same amount of capacity is available and is shared between CUBIC and LEDBAT flows. The capacity that LEDBAT flows takes from the CUBIC flows increases when the target value increase. We thereby conclude that in a LBDP network, setting the target queuing delay to 5 ms is optimal.

VI. DISCUSSION

We illustrated in section III why we believe LEDBAT could be a good candidate for LBE traffic when a long delay link is present in the network. We also showed that the target value parametrization should not be neglected as it has an impact on the performance. In sections IV and V, we illustrated that a trade-off must be found between (1) disturbing the primary traffic, (2) enabling LBE traffic and (3) increasing the use of the link capacity. Considering different long delay networks with specific traffics, we came to the conclusion that a target queuing delay of 5 ms seems to be ideal in satellite path contexts.

Indeed, when the network is fully loaded, LEDBAT is less aggressive when the target value is low, and impacts the capacity used by the principal flows less. In our simulations, LEDBAT did not exhibit fairness when the target queuing delay was more than 5 ms. Conversely, when some capacity remains on the high delay path, setting this parameter to 5 ms still enables to optimize the transmission of LEDBAT flows and the use of the whole capacity.

We focused on the optimization of the capacity of 4G satellite link because of a lack of studies in this area. We illustrated that the performance of LEDBAT to lead LBE traffic on these links can be improved. Indeed we noticed that the optimal target value is linked to the number of flows. It is worth noting that our optimal target value is quite different from the default specified in the RFC [6]. As a result, we believe that a better parametrization of LEDBAT is possible depending on specific network characteristics and conditions. It has to be assessed if this parameters relation and possible optimization remain on other networks where the delay is less important (e.g., wired or Wi-Fi).

VII. CONCLUSION

The LEDBAT algorithm has been developed to support transmission for LBE applications. In this paper, we evaluated

the performance of LEDBAT over long delay paths, taking satellite as an example. Our results also show that LEDBAT is a suitable candidate for transmitting LBE traffic over long delay paths. However, in order to yield the best performance, we showed that its target queuing delay should not be more than 5 ms in this context.

While current implementations of LEDBAT use a fixed delay of 100ms, we showed that introducing a large number of LEDBAT flows compromises the “ultra fairness” of LBE transport protocol. Reducing the target value improves LEDBAT performance while preserving its fairness. We also illustrated that the optimal parametrization is very dependent on the network characteristics and primary traffic. Therefore it seems ill-advised to rely on a fixed and static value for this parameter. We think that LEDBAT should consider the current network conditions to dynamically adapt this target value.

As future work, we expect to evaluate the LEDBAT algorithm over Radio Resource Managed (RRM) satellite networks which exhibit varying capacities. We believe these supplementary studies would provide inside into how to dynamically optimize the target value. We would also like to evaluate the intra-fairness of LEDBAT flows over long delay paths and look at mechanisms to increase the aggressiveness of LEDBAT in the absence of competing flows.

REFERENCES

- [1] QBSS, “Internet2 QBone initiative,” Tech. Rep. [Online]. Available: <http://qbone.internet2.edu/qbss/>
- [2] “BufferBloat: What’s Wrong with the Internet?” *ACM Queue*, vol. 9, no. 12, Dec. 2011. [Online]. Available: <http://doi.acm.org/10.1145/2076796.2076798>
- [3] “Cisco Global mobile data traffic forecast update,” Tech. Rep., 2009.
- [4] M. Welzl and D. Ros, “A survey of Lower-than-Best-Effort Transport Protocols,” RFC Editor, RFC 6297, Jun. 2011. [Online]. Available: <http://tools.ietf.org/html/rfc6297>
- [5] A. Sathiseelan and J. Crowcroft, “The free Internet: a distant mirage or near reality?” University of Cambridge, Technical Report 814. [Online]. Available: www.cl.cam.ac.uk/techreports/UCAM-CL-TR-814.pdf
- [6] S. Shalunov, G. Hazel, J. Iyengar, and M. Kuehlewind, “Low extra delay background transport (LEDBAT),” RFC Editor, RFC 6817, Dec. 2012. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc6817.txt>
- [7] M. Arumathurai, X. Fu, and K. K. Ramakrishnan, “NF-TCP: a network friendly TCP variant for background delay-insensitive applications,” in *Proceedings of the 10th international IFIP TC 6 conference on Networking - Volume Part II*, 2011.
- [8] S. Low, L. Peterson, and L. Wang, “Understanding TCP Vegas: Theory and Practice,” Princeton University, TR 616-00, Feb. 2000. [Online]. Available: <http://www.cs.princeton.edu/research/techreps/TR-616-00>
- [9] D. Rossi, C. Testa, S. Valenti, and L. Muscariello, “LEDBAT: The New BitTorrent Congestion Control Protocol,” in *Computer Communications and Networks (ICCCN), 2010 Proceedings of 19th International Conference on*, aug. 2010.
- [10] G. Carofoglio, L. Muscariello, D. Rossi, C. Testa, and S. Valenti, “Rethinking low extra delay background transport protocols,” *CoRR*, vol. abs/1010.5623, 2010.
- [11] G. Carofoglio, L. Muscariello, D. Rossi, and S. Valenti, “The Quest for LEDBAT Fairness,” in *GLOBECOM’10*, 2010, pp. 1–6.
- [12] A. Abu and S. Gordon, “Impact of Delay Variability on LEDBAT Performance,” in *Advanced Information Networking and Applications (AINA), 2011 IEEE International Conference on*, march 2011.
- [13] S. Ha, I. Rhee, and L. Xu, “CUBIC: a new TCP-friendly high-speed TCP variant,” *SIGOPS Oper. Syst. Rev.*, vol. 42, no. 5, Jul. 2008. [Online]. Available: <http://doi.acm.org/10.1145/1400097.1400105>
- [14] N. Kuhn, E. Lochin, J. Lacan, R. Boreli, C. Bes, and L. Clarac, “CLIFT: a cross-layer information tool to perform cross-layer analysis based on real physical traces,” *CoRR*, vol. abs/1206.5459, 2012.