Measurement of the Per-flow Burst Ratio Parameter in IP Networks

Dominik Samociuk
Department of Computer Networks and Systems
Silesian University of Technology
Gliwice, Poland
email: dsamociuk@polsl.pl

***Abstract*—In this study, we investigate the burst ratio, a significant parameter in networking that characterizes the propensity of packet losses to occur in extended, consecutive sequences. The presence of prolonged sequences of packet losses can adversely affect multimedia streams, especially those of real-time nature. In packet networks, the burst ratio is often escalated due to packet buffer overflows that occur in routers and switches, which are inherent in these systems. To analyze the burst ratio, we focus on a per-flow approach, wherein we assess the burst ratio individually for each flow of packets passing through a network node. Additionally, we compare all the per-flow burst ratios with one another and with the burst ratio calculated for the multiplexed traffic. The study explores the impact of various factors on burst ratios. Specifically, we examine the influence of flow rates and their relative proportions, the standard deviation of interarrival times, buffer capacity, the load imposed on the buffer, and the distribution of service time. Particular emphasis is placed on models with non-Poisson flows, as they are not analytically tractable. To conduct this investigation, we utilize real-world data from networking scenarios rather than simulations, allowing us to draw more accurate and robust conclusions. The obtained burst ratio measurements provide valuable insights into the behavior of packet loss processes in complex network environments and aid in enhancing the performance of multimedia streams, particularly real-time applications.**

***Keywords*—burst ratio, packet networks, IP networks, packet loss**

I. INTRODUCTION

The concept of burst ratio, denoted as B, was initially introduced in reference [1]. It is determined by the ratio of the average length of observed loss sequences within a given stream of interest, denoted as G, to the hypothetical average length of loss sequences in a Bernoulli process. In the Bernoulli process, all losses are considered to be entirely random, uncorrelated, and share a common probability of occurrence denoted as L.

Mathematically, this relationship can be expressed as follows:

Where:

$$B = G/K, \tag{1}$$

- B represents the burst ratio characteristic,
- G denotes the average length of observed loss sequences in the stream of interest, and
- K denotes the hypothetical average length of loss sequences in the Bernoulli process.

This formulation enables the quantification of the extent to which losses within a specific stream deviate from the behavior expected in an idealized Bernoulli process, facilitating a comparative analysis of the burstiness of different data streams in networking scenarios.

Through a straightforward verification, it is evident that in the case of the Bernoulli process, the value of K can be represented as:

$$K = (1 - L)^{-1}, \tag{2}$$

Consequently, as an alternative to the previously mentioned equation (1), the burst ratio (B) can be calculated using the following expression:

$$B = (1 - L)G. \tag{3}$$

To provide an illustrative example, let's consider a stream of 20 packets:

$$SSDSSSSDDDSSSSSSDSSS. \tag{4}$$

In this sequence, 'S' represents a successful packet transmission, and 'D' indicates a packet loss, packet dropped. By calculating the overall loss probability (loss ratio), denoted as L, we find:

$$L = 5/20 = 1/4 = 0.25, \tag{5}$$

Accordingly, the hypothetical mean length of the sequence of 'D's in the Bernoulli process with L=0.25 should be:

$$K = (1 - L)^{-1} = 4/3 = 1.33(recurring), \tag{6}$$

However, in our stream, we identify three sequences of 'D's, with first and third sequences containing only one dropped packet and second sequence containing three lost packets. Thus, the mean length of the sequence of 'D's in our stream, denoted as G, is equal to 1.66 (recurring).

Consequently, the burst ratio, B, is computed as:

$$B = G/K = 1.66/1.33 = 1.25, \tag{7}$$

In contemporary packet networks, empirical studies reveal that the parameter G is often 1.5 to 2.0 times greater than the parameter K [2] [3]. This discrepancy suggests an underlying mechanism influencing consecutive packet losses, whereby the probability of losing a packet is heightened following the loss of its predecessor. This phenomenon primarily arises from buffer overflows in routers and switches. During an overflow event, all incoming packets are discarded until sufficient buffer space becomes available. Consequently, this leads to a series of packet losses.

The increased burst ratio observed in modern networks, particularly the Internet, detrimentally affects the Quality of Experience (QoE) in multimedia streaming, especially in real-time applications. The loss of lengthy sequences of video or audio packets can significantly degrade QoE more than frequent losses of shorter sequences. This impact is attributable to the nature of human perception and the codecs employed in real-time multimedia transmission. For certain transmission types, this effect is quantifiable, which models the decline in voice transmission quality in Internet telephony as a function of the burst ratio.

In this study, we adopt a per-flow approach to investigate the burst ratio, assessing it for each individual packet flow passing through a network node, as opposed to evaluating a global burst ratio for aggregated traffic. Our methodology involves comparing the burst ratios of individual flows against each other and against the burst ratio of multiplexed traffic. We also examine the impact of various factors on these ratios, including flow rates and their relative proportions, the standard deviation of packet interarrival times, buffer capacity, load presented to the buffer, and service time distribution. Notably, we focus on non-Poisson flow models due to their relevance in real-world scenarios.

Previous research on burst ratios has predominantly focused on multiplexed traffic. However, our analysis underscores the importance of understanding burst ratios at the individual flow level. Indeed, the QoE for an end user engaging with multimedia content is more directly influenced by the burst ratio within the specific multimedia flow they are accessing, rather than the burst ratio in the aggregated network traffic.

It is challenging to ascertain whether the burst ratio for a specific flow will be lesser, greater, or equivalent to the burst ratio of multiplexed traffic, due to the interplay of two opposing phenomena. On one hand, a loss sequence within multiplexed traffic might comprise losses from multiple flows, potentially resulting in shorter loss subsequences within individual flows. Conversely, it is feasible for a flow to experience a sequence of packet losses even in the absence of such a sequence in the multiplexed traffic. For illustration, consider a sequence of eight packets from two distinct flows, alternating as 0 1 0 1 0 1 0 1. In this scenario, the first flow (odd packets) exhibits a sequence of four consecutive losses, which is not evident in the multiplexed traffic.

Addressing these complexities, our research into individual burst ratios is carried out using real-life network equipment and traffic in a controlled laboratory environment. This approach enables the direct observation and analysis of network behavior under various conditions, providing empirical data that closely represents real-world scenarios. Our experimental setup processes vast quantities of network packets, ensuring robustness in our results by significantly reducing the impact of random variations. This methodology also allows for the exploration of different types of interarrival time distributions within each flow, ensuring a comprehensive and realistic examination of network dynamics.

The structure of this paper is organized into six distinct sections. Section 2 presents an overview of the relevant literature and previous studies. Section 3 introduces the laboratory equipment employed. Section 4 revisits course of the experiments. Subsequently, Section 5 unveils and examines new findings pertaining to per-flow burst ratios under various conditions. The paper concludes with Section 6, summarizing the key conclusions drawn from the research.

II. Related work

In this section, we provide an overview of the existing literature and methodologies pertinent to the measurement of the per-flow burst ratio parameter in IP networks.

A. Early Theoretical Models and Their Limitations

Initial studies in the domain of packet loss and burst ratio heavily relied on stochastic processes, particularly Markov chains, to model packet loss without directly accounting for queuing mechanisms [4]–[11]. These "black box" models, while foundational, lacked the direct representation of a queue with a limited buffer, a critical element in real-world network environments. The two-state Markov model [4] and its extensions, such as the Gilbert model [5]–[7] and the Gilbert-Elliott model [8] [9], provided a basic framework but fell short in accurately mimicking the statistical structure of consecutive losses caused by queuing and buffer overflow. The limitations of these early models, particularly in capturing the variance in loss sequences, are discussed in [12]. These models could only approximate the average length of loss sequences, not their large variances, which are crucial in realistic network scenarios.

B. Transition to Empirical Measurement-Based Approaches

Recognizing the inadequacies of purely theoretical models, recent research has shifted towards empirical measurements and more realistic modeling of packet loss causes, such as queuing and buffer overflows. This transition is evident in studies that have extended the analysis to various Poisson stream configurations [13]–[15] and empirical investigations in controlled lab settings and real-world network environments [2] [3]. The loss ratio (L), a fundamental characteristic of the packet loss process, has been extensively studied through both direct network measurements [16]–[20] and analytical formulas in queue models with finite buffers [13]–[15]. These studies have provided a more nuanced understanding of the loss process in network traffic.

C. Broader Perspectives and Advanced Models

In addition to the burst ratio, other metrics for loss characterization have been explored. Studies like [21] and [22] have investigated the probability distributions of loss sequences and the lengths of initial loss sequences, offering alternative perspectives on packet loss in networks. Advanced models, such as the four-state Markov chain [10] and the general k-state Markov chain model [4] [7] [11], have been developed to allow for a more detailed analysis of loss patterns, including the loss of k consecutive packets in specific states.

D. Active Queue Management and Burst Ratio Mitigation

To address the burst ratio in TCP/IP networks, research has focused on Active Queue Management (AQM). A variety of AQM methods have been proposed [23] [24]. Particularly, algorithms based on the dropping function [25] have shown promise in reducing the burst ratio, as demonstrated in experimental studies [26]–[28].

E. Burst Ratio in Aggregated Traffic

The study of burst ratio extends beyond individual packet flows to encompass aggregated traffic, a critical aspect in understanding network behavior on a larger scale. In aggregated traffic scenarios, the burst ratio provides insights into the collective behavior of multiple data streams converging within a network. This approach is essential for comprehending the dynamics of network congestion and the resultant packet loss patterns. Research in this area often employs complex models that simulate the interaction of multiple traffic flows, thereby offering a more comprehensive view of network performance under varying load conditions. Key studies in this domain have focused on the impact of burst ratio in aggregated traffic [29]–[32].

F. Analytical Approaches to Burst Ratio in Individual Network Flows

In contrast to empirical measurement-based studies, analytical approaches to understanding the burst ratio in individual network flows have also been prominent. These approaches typically involve developing mathematical models that can predict packet loss behavior in a single flow, considering various factors such as traffic intensity, buffer size, and service policies. Such models are invaluable for theoretical analysis and for designing network systems that can efficiently handle expected traffic patterns. Significant contributions in this area have included the development of formulas and algorithms that accurately predict the burst ratio in individual flows, taking into account the unique characteristics of each flow [33] [34].

III. Laboratory equipment

In order to conduct measurements of the burst ratio per flow in a network laboratory, a test network comprising three main components was established:

1. **Spirent SPT-N4U Traffic Generator**: This device, equipped with the MX2-10G-S12 module, features 12 fiber optic ports, each capable of generating traffic at maximum speeds of 1 Gb/s or 10 Gb/s. It is adept at generating substantial volumes of stateful and stateless traffic (TCP and UDP) across all ports independently, without degrading the internal performance of the device.

2. **Test Device - Cisco 3750X Layer 3 Router/Switch**: This device, fitted with 12 ports of 1 Gb/s and two ports of 10 Gb/s, is utilized for observing actual packet losses. Its role is critical in the precise measurement of the burst ratio per flow in various network conditions.

3. **High-Performance HP Enterprise ProLiant DL380 Gen9 Servers**: These servers are employed for receiving and storing both incoming and outgoing traffic in databases, and for executing packet loss analysis, with a specific focus on per-flow burst ratio measurements.

The Spirent SPT-N4U, functioning as both a traffic generator and analyzer, is the primary device used in this setup. This generator proved indispensable for generating high traffic loads, which are challenging to achieve using standard computers. It offers a complete, portable environment for network operations and performance testing of complex topologies. The generator's chassis facilitates the connection of various load modules, allowing potential traffic generation up to 400 Gb/s, with precision up to 10 ns per generated or analyzed frame. It supports an integrated test controller (client-server type application) that manages the entire system.

The load modules facilitate processing of necessary signals for testing both standard transmission and voice/video transmissions, as well as applications from layers 2 to 7. The load module used in the experiments, the MX2-10G-S12, has 12 fiber optic ports operating at 1 Gb/s or 10 Gb/s speeds as required. Each port supports packet generation and analysis at layer 2 and 3 at cable speed (maximum throughput of the used physical layer), along with efficient emulation of routing and switching protocols. Each port contains a separate RISC processor under Linux OS control, with an optimized TCP/IP test stack.

In the experiments, the ports of the generator configured in the Test Center application are treated as separate instances of the Linux OS. This ensures that each interface is tested independently, with the status of the respective instance relayed to the supervising machine. Each port configured in the application can operate in dual mode - as both generator and analyzer. During traffic analysis, one can monitor the total throughput of received and generated traffic (in frames/packets or bits), packet contents, average delay (as well as minimum and maximum values), analyze sequencing mechanisms, and construct histograms based on different parameters of the experiment and apply various filters for analyzing only specific packets or their groups.

The test device used in the experiments, a Cisco 3750X layer 3 switch/router, was equipped with 12 ports of 1 Gb/s using multimode SFP transceiver modules and two multimode SFP+ transceiver modules operating at 10 Gb/s. Its switching matrix performance of 160 Gb/s and 35.7 million packets per second ensured that no additional delay would be introduced into queuing mechanisms. The Cisco 3750X supports both standard packet sizes and Jumbo frames up to 9216B. Additionally, it allows for stacking (combining several physical devices into one logical device) as needed, providing up to 468 total ports.

Direct measurement of the burst ratio per flow using the Spirent SPT-N4U traffic generator and analyzer is not feasible. Network traffic must be captured, transmitted to servers, and stored in databases for subsequent calculation of packet loss characteristics. SQLite was chosen for data storage due to its support on the Spirent SPT-N4U device, ease of use, and simple installation on computational servers. SQLite is a well-known database management system and library written in C, implementing a SQL (Structured Query Language) supporting engine. It implements an SQL engine that allows database use without a separate process of a Relational Database Management System (RDBMS). Due to the requirement to store traffic records from each experiment in a separate database, SQLite was the most practical solution.

During the experiments, three variants of the database were created. The first database, created on the first server, stored a table with source packets. It contained integers describing the following elements of packet headers: timestamps, source and destination IP addresses, IP protocol numbers, IP identifiers, IP packet lengths, source and destination ports, and sequence and acknowledgment numbers for TCP flows.

The second database, created on the second server, stored a table with destination packets with the same fields as in the first database. The third database was created on the third server. It was a combination of the two tables of both previous databases. It was theoretically possible to use only 2 servers, e.g., by copying the first table to the second server, but the use of a third server allowed for speeding up the measurement collection process by performing calculations in parallel for the previous scenario and capturing traffic for the next one. While the third server was still calculating the burst ratio per flow for the previous scenario, the first and second servers were already collecting new traffic in the next scenario.

IV. Course of the Experiment

The test network established for measuring the burst ratio is depicted in Figure 1. Traffic was generated on ports 1 to 4 of the Spirent SPT-N4U generator and received on 1 Gb/s ports numbered 1, 3, 5, and 7 of the Cisco 3750X device. The generated traffic was based on the most common protocol stack: Ethernet at Layer 2, IPv4 at Layer 3, and TCP/UDP at Layer 4. The default TCP congestion control of Spirent, namely New Reno with a maximum window size of 32768B, was employed. At Layer 7, HTTP traffic was emulated, with an Apache WWW server on port 80 on the server side and a Mozilla browser on the client side.

Within the Cisco 3750X device, all incoming traffic was duplicated using the SPAN function to output ports 9 and 13. Port 13 was used to send a copy of the incoming traffic to server 1 for further analysis, while port 9 was the test port — this port experienced actual packet losses (specifically in its buffer). Subsequently, the outgoing traffic from this port was looped back to port 8 and duplicated within the test device (again using the SPAN function) to ports 10 and 14. Finally, traffic from port 10 was sent back to the Spirent analyzer, operating as the packet destination (and stateful client for each TCP connection), while traffic from port 14 was sent to server 2 for further analysis.

On the first two servers, the DPDK-dump application [35] was used for data capture. The data from each packet was removed, and a header with a timestamp and index was stored in a SQLite database. The preliminarily processed packets were then copied to a third server, where databases containing the input and output packets from the test device were merged. If possible, each outgoing packet from the Layer 3 switch was paired with its corresponding input packet. If an input packet did not have a corresponding output packet, it was considered to have been removed from the overflowed buffer — lost in network transmission.

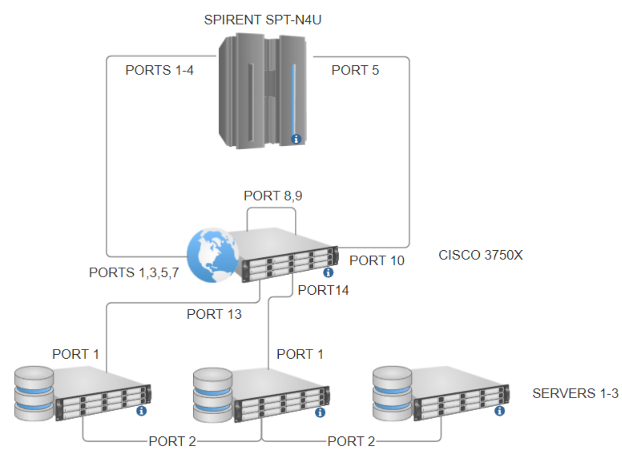

Figure 1. Test network.

As shown in Figure 1, the test traffic was transmitted on six links with up to 1 Gb/s capacity, while additional data for calculating the burst ratio and other network traffic parameters were transmitted through a 10 Gb/s LAN network.

V. Results and discussion

In the experiment, the burst ratio was measured across multiple scenarios featuring variable traffic characteristics, different number of TCP flows and different buffer sizes. In each scenario, approximately two million packets were generated. UDP traffic was employed as background traffic, constituting 5% of the total link load, unless specified otherwise in the scenario. It is important to note that in all scenarios incorporating at least one TCP connection, it was not feasible to manually set any arbitrary value for the test link load (ρ). The load is automatically adjusted by the New Reno algorithm operating in each TCP sender.

In the initial scenarios, an identical distribution of flows was generated, with the total bandwidth divided among a varying number of flows, specifically: 24, 32, 48, 64, 96, 128, and 192 flows. ρ in each scenario stabilized at a value

Copyright (c) IARIA, 2024. ISBN: 978-1-68558-174-9

of 0.944 Mbps. As hypothesized, the burst ratio exhibited a slight increase with the rising number of TCP flows. In the latter scenarios, the measured values of the loss ratio were comparable, leading to the conclusion that beyond a certain threshold of TCP connections, further increases do not significantly affect the burst ratio. Similarly, the per-flow burst ratios were consistently lower than B in all cases. Furthermore, as the number of flows increased, each individual burst ratio (B_i) decreased, indicating a dilutive effect of increased flow counts on the individual burst ratio metrics.

To investigate the impact of interarrival time on both the global burst ratio and the per-flow burst ratios, a varying percentage of UDP traffic within the total link bandwidth was utilized (from 0 to 20%). The increasing proportion of small datagrams notably influenced the rise in variance of the packet service time distribution in the queue. This, in turn, led to an increase in the global burst ratio (B), as well as an enhancement of the per-flow burst ratio (B_i). However, due to the identical characteristics of the flows within a given scenario, B_i remained consistent across different scenarios. Again, the burst ratio for individual flows (B_i) was less than the multiplexed burst ratio (B) for each i.

In the third group of scenarios, the buffer size of the examined port was varied from 50 to 1000 packets, while maintaining 128 TCP connections and 50 Mb/s of UDP traffic. The variable buffer size on the tested port did not result in significant changes in the burst ratio, leading to the conclusion that the burst ratio is weakly dependent on the size of the queue buffer in network transmissions. Similarly, the variable buffer size in selected scenarios did not affect the change in the per-flow burst ratio in relation to the baseline scenario.

In all previous scenario groups, despite differences in aggregated traffic, very similar per-flow traffic characteristics were created. To introduce diversity at the level of flow characteristics, the background traffic (UDP) was modified for only a portion of the flows (25%) while it remained original (5%) for the others. This adjustment enabled the observation of scenarios where flows, even within connection-oriented protocols (TCP used in the scenario), could be compared with varying characteristics, specifically different interarrival times (simulated by varying proportions of small datagrams). Subsequently, in the following scenarios, the UDP traffic was altered for 50% and 75% of the flows. These scenarios yielded interesting results: for most of the scenarios, all per-flow burst ratios were significantly smaller than the global burst ratio (B), and not very different from each other, despite the variances of the flows differing between groups of flows. However, for scenarios with a significant variance in packet interarrival times, which also constituted a substantial portion of the total network load, instances of B_i exceeding B were recorded, highlighting the impact of varied flow characteristics on the burst ratio parameter.

The conclusions drawn from each group of scenarios are summarized in Table I.

TABLE I
SCENARIOS CONCLUSIONS

Scenario	Aggregated burst ratio	Per-flow burst ratios
24 TCP Flows	1.30	All 1.06
32 TCP Flows	1.30	All 1.06
48 TCP Flows	1.34	All 1.06
64 TCP Flows	1.35	All 1.05
96 TCP Flows	1.35	All 1.04
128 TCP Flows	1.35	All 1.04
192 TCP Flows	1.36	All 1.03
UDP 0%	1.23	All 1.02
UDP 5%	1.35	All 1.04
UDP 10%	1.43	All 1.05
UDP 15%	1.46	All 1.05
UDP 20%	1.52	All 1.06
Buffer 50	1.35	All 1.04
Buffer 100	1.35	All 1.04
Buffer 200	1.35	All 1.04
Buffer 500	1.34	All 1.03
Buffer 1000	1.35	All 1.04
UDP 5%, 25% flows	1.35	With UDP 1.07, Without UDP 1.02
UDP 10%, 25% flows	1.43	With UDP 1.07, Without UDP 1.02
UDP 15%, 25% flows	1.46	With UDP 1.08, Without UDP 1.02
UDP 20%, 25% flows	1.52	With UDP 1.09, Without UDP 1.02
UDP 5%, 50% flows	1.35	With UDP 1.10, Without UDP 1.01
UDP 10%, 50% flows	1.43	With UDP 1.11, Without UDP 1.01
UDP 15%, 50% flows	1.46	With UDP 1.14, Without UDP 1.01
UDP 20%, 50% flows	1.52	With UDP 1.15, Without UDP 1.01
UDP 5%, 75% flows	1.35	With UDP 1.27, Without UDP 1.01
UDP 10%, 75% flows	1.43	With UDP 1.35, Without UDP 1.01
UDP 15%, 75% flows	1.46	With UDP 1.47, Without UDP 1.01
UDP 20%, 75% flows	1.52	With UDP 1.59, Without UDP 1.01

VI. CONCLUSIONS AND FUTURE WORK

In this study, we conducted a detailed analysis of the burst ratio on a per-flow basis within a network environment, where each packet flow through a network node was evaluated independently. The burst ratio is a critical metric that quantifies the tendency of packet losses to occur in extended, consecutive sequences, which can adversely affect the quality of service for multimedia streams, especially those requiring real-time transmission.

Empirical measurements were carried out using sophisticated laboratory equipment to ascertain the burst ratios of individual flows. These measurements revealed that the burst ratio for a specific flow could be lower, higher, or equivalent to the aggregate burst ratio, contingent upon the characteristics of the flows involved.

Increasing TCP Flows: The global burst ratio slightly increased with the number of TCP flows, indicating a minor degradation in the network's ability to handle packet losses in a bursty manner as the number of flows increases. Notably, the per-flow burst ratios also exhibited variations, but the impact plateaued beyond a certain number of TCP connections, suggesting a threshold beyond which additional TCP flows do not exacerbate burst behavior significantly.

UDP Traffic Proportion Influence: Adjusting the UDP traffic percentage within the total bandwidth affected both the global and per-flow burst ratios. The presence of small datagrams increased the service time variance, elevating both global and per-flow burst ratios. This highlights the sensitivity of burst ratios to the composition and characteristics of network traffic.

Buffer Size Variation Effects: Changes in buffer size from 50 to 1000 packets showed minimal impact on both global and per-flow burst ratios, indicating a low dependency of burst loss behavior on queue buffer size. This suggests that other factors, beyond buffer capacity, play a more significant role in influencing burst loss patterns.

Differentiated Background Traffic: Introducing variations in background traffic for certain flows altered the burst ratio landscape, with modifications affecting both the global and per-flow metrics. The approach allowed for distinguishing the effects of flow-specific characteristics, particularly interarrival times, on burst loss behavior, underscoring the importance of individual flow properties in network performance analysis.

Varied Interarrival Times: The scenarios with adjusted background traffic for varying proportions of flows highlighted the differential impact on global and per-flow burst ratios. Significant variances in packet interarrival times, especially in flows that constituted a large portion of the network load, were associated with instances of per-flow burst ratios exceeding the global burst ratio. This observation illustrates the complex relationship between traffic diversity, flow characteristics, and their collective impact on network burst loss dynamics.

These refined conclusions emphasize the importance of considering both global and per-flow perspectives to fully understand the intricacies of network traffic management and its implications on performance metrics like burst ratio.

Our forthcoming research endeavors will extend our experimental setup to encompass real multimedia streaming traffic and IPv6, aligning with current and progressive networking contexts. We also plan to explore the effects of ON-OFF patterns on Quality of Experience and provide a comprehensive description of traffic assumptions and flow characteristics to facilitate the replication of our experiments. Additionally, we will focus on measuring higher-order statistics from traffic using a methodology similar to that employed for burst ratio analysis. These efforts aim to robustly test the impact of various network conditions and expand the scope and applicability of our research, ensuring its relevance in the study of network behavior and its influence on multimedia content delivery.

ACKNOWLEDGMENT

This research was funded by National Science Centre, Poland, grant number 2020/39/B/ST6/00224.

REFERENCES

[1] J. W. McGowan, ”Burst ratio: a measure of bursty loss on packet-based networks”, 16 2005. US Patent 6,931,017, 2005.

[2] D. Samociuk, A. Chydziński, and M. Barczyk, ”Experimental measurements of the packet burst ratio parameter”, Proc. BDAS 2018, pp. 455-466, 2018.

[3] D. Samociuk, M. Barczyk, and A. Chydziński, ”Measuring and analyzing the burst ratio in IP traffic”, Proc.BROADNETS, pp. 86-101, 2019.

[4] M. Yajnik, S. Moon, J. Kurose, and D. Towsley, ”Measurement and modelling of the temporal dependence in packet loss”, Proc. of IEEE-INFOCOM, vol. 1, pp. 345-352, 1999.

[5] E. N. Gilbert, ”Capacity of a Burst-Noise Channel”, Bell system technical journal, vol. 39(5), pp. 1253-–1265, 1960.

[6] M. Veeraraghavan, N. Cocker, and T. Moors, ”Support of voice services in IEEE 802.11 wireless LANs”, Proc. of IEEE INFOCOM, vol. 1, pp. 488-497, 2001.

[7] W. Jiang and H. Schulzrinne, ”Modeling of packet loss and delay and their effect on real-time multimedia service quality”, Proc. of NOSSDAV, pp. 1–10, 2000.

[8] E. O. Elliott, ”Estimates of Error Rates for Codes on Burst-Noise Channels”, Bell system technical journal, vol. 42(5), pp. 1977-1997, 1963.

[9] G. Haßlinger and O. Hohlfeld, ”The Gilbert-Elliott model for packet loss in real time services on the Internet”, Proc. of Measuring, Modelling and Evaluation of Computer and Communication Systems Conference, pp. 1-15, 2008.

[10] A. Clark, ”Modeling the effects of burst packet loss and recency on subjective voice quality”, Proc. if Internet Telephony Workshop, pp. 123- 127, 2001.

[11] H. A. Sanneck and G. Carle, ”Framework model for packet loss metrics based on loss runlengths”, Proc. SPIE 3969, Multimedia Computing and Networking 2000, pp. 1-11, 2000.

[12] X. Yu, J.W. Modestino, and X. Tian, ”The accuracy of Gilbert models in predicting packet-loss statistics for a single-multiplexer network model”, Proc. of IEEE INFOCOM’05, pp. 2602–2612, 2005.

[13] H. Takagi, ”Queueing analysis - Finite Systems”, North-Holland Amsterdam, 1993.

[14] A. Chydzinski, R. Wojcicki, and G. Hryn, ”On the Number of Losses in an MMPP Queue”, Lecture Notes in Computer Science, vol. 4712, pp. 38-48, 2007.

[15] A. Chydzinski and B. Adamczyk, ”Transient and stationary losses in a finite-buffer queue with batch arrivals”, Mathematical Problems in Engineering, vol. 2012, ID 326830, pp. 1-17, 2012.

[16] P. Benko and A. Veres, ”A passive method for estimating end-to-end TCP packet loss”, Proc. of IEEE GLOBECOM’02, pp. 2609–2613, 2002.

[17] Bolot J. ”End-to-end packet delay and loss behavior in the Internet”, Proc of ACM SIGCOMM’93, pp. 289–298, 1993.

[18] M. Coates and R. Nowak, ”Network loss inference using unicast end-toend measurement”, Proc. of ITC Conference on IP Traffic, Measurement and Modeling, pp. 282–289, 2000.

[19] N. Duffield, F. L. Presti, V. Paxson, and D. Towsley, ”Inferring link loss using striped unicast probes”, Proc. of IEEE INFOCOM’01, pp. 915–923, 2001.

[20] J. Sommers, P. Barford, N. Duffield, and A. Ron, ”Improving accuracy in end-to-end packet loss measurement”, ACM SIGCOMM Computer Communication Review, vol. 35(4), pp. 157–168, 2005.

[21] M. Bratiychuk and A. Chydzinski, ”On the loss process in a batch arrival queue”, Applied Mathematical Modelling, Vol. 33, Iss. 9, pp. 3565-3577, 2009.

[22] A. Chydzinski, ”On the remaining service time upon reaching a target level in M/G/1 queues”, Queueing Systems, vol. 47, issue 1/2, pp. 71–80, 2004.

[23] A. Chydzinski, ”Buffer Overflow Period in an MAP Queue. Mathematical”, Problems in Engineering, vol. 2007, pp. 1–18, 2007.

[24] A. Chydzinski, ”Buffer overflow period in a batch-arrival queue with autocorrelated arrivals”, Applied Mathematics Information Sciences 7, no. 4, pp. 1633–1641, 2013.

[25] A. Chydzinski, M. Barczyk, and D. Samociuk, ”The Single-Server Queue with the Dropping Function and Infinite Buffer”, Mathematical Problems in Engineering, vol. 2018, ID 3260428, pp. 1–12, 2018.

[26] M.Barczyk and A. Chydzinski, ”Experimental testing of the performance of packet dropping schemes”, Proc. IEEE ISCC, pp. 1-7, 2020.

 ISBN: 978-1-68558-174-9

[27] M. Barczyk and A. Chydzinski, ”AQM based on the queue length: A real-network study”, PLoS ONE 17(2): e0263407, 2022.

[28] D. Samociuk and A. Chydzinski, ”On the impact of the dropping function on the packet queueing performance”, Proc. of International Convention on Information and Communication Technology, Electronics and Microelectronics, MIPRO 2018, pp. 473-478, 2018.

[29] A. Chydzinski and D. Samociuk, ”Burst ratio in a single-server queue”, Telecommunication Systems, DOI: 10.1007/s11235-018-0476-7, 2018.

[30] A. Chydzinski, D. Samociuk, and B. Adamczyk, ”Burst ratio in the finite-buffer queue with batch Poisson arrivals”, Applied Mathematics and Computation 330, pp. 225-238, DOI: 10.1016/j.amc.2018.02.021, 2018.

[31] D. Samociuk, A. Chydzinski, and M. Barczyk, ”Experimental measurements of the packet burst ratio parameter”, Communications in Computer and Information Science, Springer, Volume 928, pp. 455-466, 2018.

[32] D. Samociuk, M. Barczyk, and A. Chydzinski, ”Measuring and analyzing the burst ratio in IP traffic”, Communications in Computer and Information Science, Springer, Volume 928, 2018.

[33] A. Chydzinski, ”Per-flow structure of losses in a finite-buffer queue”, Applied Mathematics and Computation 428, 127215, pp. 1-15, May 2022.

[34] A. Chydzinski and B. Adamczyk, ”Burst ratio of packet losses in individual network flows”, INFORMATICA, Volume 34, pp. 35–52 February 2023.

[35] DPDK documentation website, URL: https://www.dpdk.org/, visited on 01.02.2024.