

Fine-grained Starlink Throughput Variation Examined With State-Transition Modeling

Johan Garcia
Karlstad University / Icomera AB
Email: johan.garcia@kau.se

Simon Sundberg
Karlstad University
Email: simon.sundberg@kau.se

Anna Brunstrom
Karlstad University
Email: anna.brunstrom@kau.se

Abstract—Leveraging a data set of almost half a billion packets with high-precision packet times and sizes, we process it to extract characteristics of the bursts emitted over Starlink’s Ethernet interface. The structure of these bursts directly reflect the physical layer receipt of OFDMA frames on the satellite link. We study these bursts by analyzing their rates, and by proxy the transition between different physical layer rates. The results highlight that there is definitive structure in the transition behavior, and we note specific behaviors such as particular transition steps associated with rate switching, and that rate switching occurs mainly to neighboring rates. We also study the joint burst rate and burst duration transitions, noting that transitions occur mainly within the same rate, and that changes in burst duration are often performed with an intermediate short burst in-between. Finally, we examine the configurations of the three factors burst rate, burst duration, and inter-burst silent time, which together determine the effective throughput of a Starlink connection.

I. INTRODUCTION

Low-Earth orbit (LEO) systems have made it possible to provide Internet connectivity across wide areas that were previously under-served in terms of wireless communication access. Such access has generally been provided by cellular networks with limit to their coverage in sparsely populated areas. Providing passengers onboard vehicles such as trains and buses with Internet connectivity when traversing areas of no, or poor, cellular connectivity is also a use case, as is the use of Starlink for 5G backbone connectivity [1], [2]. In contrast to cellular networks, the details of LEO systems, and SpaceX’s Starlink [15] in particular, is opaque and many system design decisions can only be inferred through measurements as the amount of publicly provided information on system details is limited.

Several measurement studies have been performed to evaluate different aspects of the Starlink network [10], [11], [18], [20]. In this work, we in particular focus on the satellite down-link mechanisms that yield a particular observed throughput at the network layer. By utilizing a large data set of high-precision per-packet measurements, we study the behavior of the transmission bursts which are associated with how Starlink allocates system resources between concurrent users.

Previous studies [5], [11] have noted the existence of a 15 second reconfiguration interval, as well as an OFDMA frame time of 1.33 ms. Our recent study [4] on physical layer details highlight that the data in these frames is transmitted in bursts observable at the network layer, and by analyzing the timings

of these bursts insights on physical layer characteristics can be gained. The number of frames per burst in our measurements mostly vary between 1 and 7. Between these bursts, there are inter-burst pauses, where no data is received by the user.

On longer time scales (i.e. multi-burst time scales) we note that the throughput obtained by a user is a function of a) the physical transmission rate of the frames, b) the number of frames in a burst, and c) the inter-burst silent time. Throughput is thus a combination of the prevailing radio channel conditions as managed by (a) and the scheduling across competing users as managed by (b,c). We here focus on an initial investigation of these factors with the aim to obtain additional knowledge regarding Starlink rate transition behavior.

Building on our previous studies of Starlink burst rates we here study how the transition between rates are performed within the Starlink network. Based on the analysis of 2.4 million data bursts observed in our data set of around half a billion packet sizes and times we arrive at three main insights:

- rate changes in the Starlink physical transmission rates predominantly occur between adjacent rates,
- changes to burst rate or length are usually done in conjunction with a short transitory burst,
- and that a clear relationship can be empirically observed between the effective throughput and the underlying system mechanisms of burst rates and burst timings.

We believe that the findings reported in this paper can be applicable to analytical, simulation, and emulation based modeling of the Starlink link throughput variations, and provide additional knowledge regarding the Starlink physical layer behavior. Given the proprietary nature of the Starlink network, documentation on system design details is scarce, and we believe that measurement studies like this one can contribute to an enhanced system understanding.

II. DATA COLLECTION

To collect data we perform a measurement campaign utilizing a Starlink deployment on the roof of the main building of the Department of Computer Science at Karlstad University. Our Starlink deployment includes a Gen-2 Starlink kit, which includes: a) a satellite dish with an electronic phased antenna; b) a motorized base for self-orientation of the dish; and c) a WiFi router with an Ethernet adapter. It is connected to our measurement machine via an Ethernet

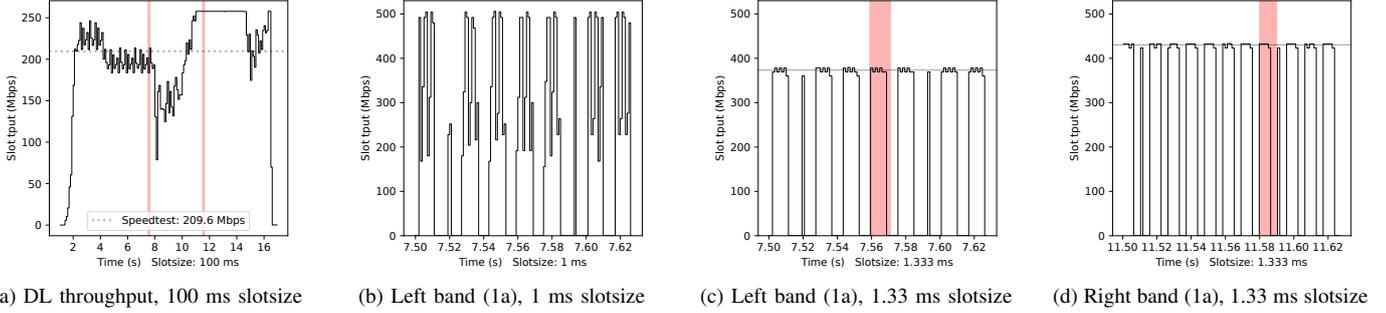


Fig. 1. Starlink downlink throughput over two time scales. The red regions in Figure 1a are shown in Figure 1b-d.

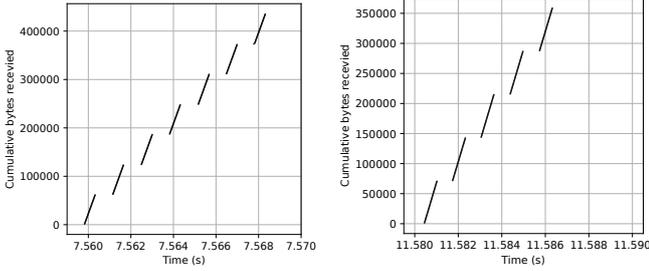


Fig. 2. The seven sub-bursts (each with data from one Starlink OFDMA frame) within the red burst in Figure 1c (left), and the five sub-bursts from the red burst in Figure 1d (right).

network, see [14] for further details. To generate measurement traffic over the Starlink access, we utilize the Ookla Speedtest command line tool [13] using a specified Speedtest server located in Stockholm. An IP-geolocation on the Starlink exit node indicates that it is also located in Stockholm. Each Speedtest run employs multiple parallel TCP connections and each measurement run has a duration of 9 to 15 seconds [9]. We set up a cron-job to execute a run every 5 minutes over approximately a week, for a total of 2195 runs. The traffic generated by the measurement runs is captured using tcpdump. Additionally, we use eBPF [19] to implement a tool to capture inter-packet delay (IPD) data. The IPD-tool collects timestamps from all packets at the tc-eBPF hook and calculates the link-wide IPDs. The IPDs are then stored along with packet sizes in a compact format using a minimum of just 4 bytes per packet, as described in [17]. The IPD data is well suited for studying link throughput variations at multiple timescales, and has a much smaller storage footprint than pcap files.

The measurements are collected from a physical machine with a 4-core Intel i7-6700 CPU and 16 GB of memory, running Ubuntu 22.04 with a 6.3.2 Linux kernel. We use an Intel 10G X550T Ethernet network interface card (NIC), which can supply hardware timestamps for all received packets with a precision of tens of nanoseconds. We disable Large Receive Offload (LRO) and Generic Receive Offload (GRO) so each individual network packet can be monitored, rather than the merged 64 KiB superpackets that the offloads may produce.

III. STARLINK THROUGHPUT CHARACTERISTICS

To provide some background on Starlink throughput variation, and illustrate the ability of our measurement approach to examine Starlink throughput at varying time scales, we now provide an initial overview of Starlink downlink throughput. We illustrate using one of the 2195 measurement runs, choosing our displayed timescales to be relevant for our subsequent discussion on Starlink’s physical rate and its temporal variation.

Figure 1 illustrates the throughput evolution over time for the example run at two timescales. The complete measurement run, shown in Figure 1a, shows the throughput as it appears when calculated over slots of 100 ms. Notably, there is significant variation in Starlink throughput during this measurement run, with a distinct dip around the 8-second mark. This dip corresponds to the Starlink network’s periodic reconfiguration which takes place every 15 seconds. We have in a previous work [5] identified this periodic dip based on network measurements, and assessed its potential impact on TCP throughput performance. The issue of reconfiguration interval is also discussed in [11].

In Figure 1a, two small regions are marked in red. The leftmost region is further detailed in Figure 1b, using a finer slotsize of 1 ms for throughput computation. Here, the y-axis scale extends up to 500 Mbps, in contrast to the 250 Mbps used in Figure 1a. It is clear that there is considerable variation in the throughput where there are periods without any traffic followed by short bursts of traffic. However, we should note that 1 millisecond is not a suitable choice for the slot size, given that it has been shown that Starlink has a frame time of 1.33 milliseconds. Thus, aliasing effects are likely, and in fact, this is what is present in Figure 1b. If we take exactly the same data and instead use the slot size equal to the Starlink frame time, then we get the appearance that is shown in Figure 1c which shows a stable rate during the data bursts. Included in the figure is a horizontal reference line at 373.3 Mbps, one of the identified Starlink physical rates, as will be discussed later.

In Figure 1d the left red band of Figure 1a is shown, which covers a stable region of higher throughput than the region of Figure 1c. Here we can note that the general structure of having short data bursts followed by silent regions is the same.

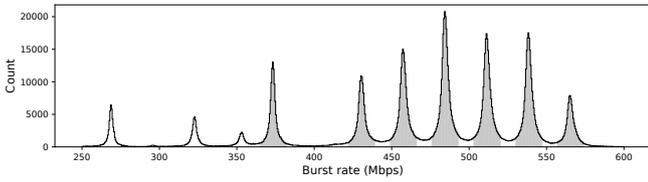


Fig. 3. Burst rate distribution for all burst durations, with cluster extent marked in grey

The throughput obtained in the bursts is however different, as is the duration of the bursts, and as is the inter-burst silent time. Here, the horizontal reference line is placed at 430.5 Mbps.

By increasing the time resolution further, we can consider the behavior within one of these bursts. The bursts marked in red in Figure 1c and 1d are shown in Figure 2, but now using the cumulative bytes received on the y-axis instead of throughput. Within each burst, there is a clear periodic pattern visible in the form of sub-bursts, with the left subfigure having seven sub-bursts and the right five sub-bursts. We have in related work [4] examined these sub-bursts and their relation to the physical layer OFDMA frame structure of Starlink, and noted that the sub-bursts as visible in Figure 2 correspond to data received in a Starlink physical layer frame, when that data is emitted out on the Starlink Ethernet interface. The slope of the sub-bursts in Figure 2 thus correspond to the Ethernet line rate, and the distance between the start of consecutive sub-bursts is exactly 1.333 ms, which is the Starlink physical frame time [6]. In our related work [4] we also derived a set of commonly occurring transmission rates based on an aggregate analysis of burst characteristics. These transmission rates are listed in Table 1, and will be used in the further analysis. We note the match between Fig 1c and rate #6, and Fig 1d and rate #7.

TABLE I

EMPIRICAL AND HYPOTHEZED STARLINK PHYSICAL RATES, MODEL RATE DIFFERENCE TO BASELINE IN STEPS OF 27 MBPS, AND ERROR BETWEEN EMPIRICAL RATE AND MODEL RATE IN MBPS. FROM [4], WHERE DETAILS ON POSSIBLE MODEL PARAMETERS LISTED IN THE COMMENT COLUMN ARE ALSO PROVIDED.

| # | Emp.rate | Theo.rate | Rate diff | Error | Comment |
|----|---------------|--------------|-----------------|-------------|--------------------|
| 0 | 134.45 | 133.5 | -11×27 | 0.95 | |
| 1 | 215.02 | 214.5 | -8×27 | 0.52 | |
| 2 | 242.16 | 241.5 | -7×27 | 0.66 | |
| 3 | 268.79 | 268.5 | -6×27 | 0.29 | |
| 4 | 322.72 | 322.5 | -4×27 | 0.22 | |
| 5 | 353.45 | 349.5 | -3×27 | 3.95 | No model match |
| 6 | 373.30 | 376.5 | -2×27 | -3.20 | No model match |
| 7 | 430.57 | 430.5 | 0 | 0.07 | 287sym 4QAM |
| 8 | 457.58 | 457.5 | $+1 \times 27$ | 0.08 | +18sym 16QAM |
| 9 | 484.63 | 484.5 | $+2 \times 27$ | 0.13 | +36sym 16QAM |
| 10 | 511.60 | 511.5 | $+3 \times 27$ | 0.10 | +54sym 16QAM |
| 11 | 538.51 | 538.5 | $+4 \times 27$ | 0.01 | +72sym 16QAM |
| 12 | 565.43 | 565.5 | $+5 \times 27$ | -0.07 | +90sym 16QAM |
| 13 | 645.35 | 646.5 | $+8 \times 27$ | -1.15 | |

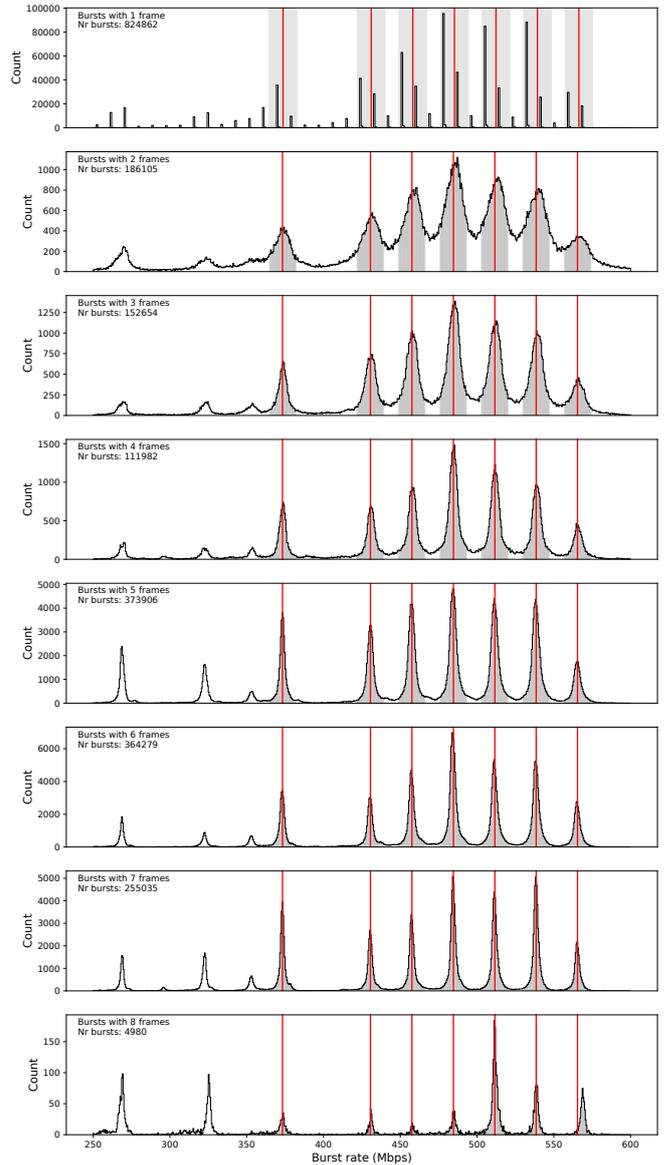


Fig. 4. Burst rate distribution per burst duration, considered in number of 1.33 ms frames in burst. Cluster extent marked in grey. The red lines mark the cluster centers of the 7 considered clusters. Note difference in y-axis scaling.

IV. BURST RATE DISTRIBUTION AND CLUSTERS

We now consider the specific burst rates identified in our previous work [4] as listed in Table 1, with the aim to provide a more granular examination of the distribution of rates, and the rate transition behavior. To enable the examination of burst characteristics we, similarly to our previous work, process the packets captured during the measurement runs to identify bursts. To delineate the bursts, we use a minimum interburst delay of 1.5 ms and require a burst to have at least two packets to be considered as a burst. Since we are interested in the burst behavior during full traffic load, we consider only the traffic that is present within a 1 second margin away from the start and end of the Speedtest measurement run, e.g., one second

in from the start and end of the example measurement run in Figure 1a. With this definition of the bursts, we process the 422 million packets that are within the margins, and locate 2 375 690 bursts.

Figure 3 provides an histogram of the resulting burst rates. As our interest here lies in examining the transitions between different burst rates, we need a mechanism to assign the observed bursts into clusters where each cluster corresponds to one of the rate peaks in the figure, a rate which also has an equivalent line in Table 1 where it is numbered. In the following we use the number from Table 1 to identify the clusters. We use a rate width of 9 Mbps on either side of the empirical rate to assign observed bursts into a cluster. Thus, the grey areas in Figure 3 correspond to the regions associated with clusters 6, 7, 8, 9, 10, 11, and 12 which are the ones we consider in this study.

We now consider a more detailed analysis of the burst rate distribution. As illustrated in Figure 2, bursts are composed of a number of sub-bursts, each corresponding to the data in a physical layer frame. In Figure 4 we provide the distribution of burst rates subdivided by the number of sub-bursts/frames present in each burst. Notably, bursts consisting of a single frame – a burst length of 1 – display a distinctive spiky distribution pattern. For bursts encompassing only a single frame, the obtained rate is directly proportional to the count of packets within that frame. Beyond single-frame bursts, the distributions have some commonality in that they all exhibit similar peak locations. However, the relative prominence of these peaks varies with the burst duration. How large fraction of the observations that fall outside any cluster also varies, as does the degree to which the peaks are well defined. For the 2 frame case, the peaks are less defined and more observations become non-clustered. One can also note that there is other variations when comparing bursts of differing lengths. For instance, in the case of bursts lasting 8 frames, lower rates constitute a larger fraction of the distribution compared to bursts that span 4 frames.

V. BURST RATE TRANSITIONS

Having assigned all burst observations to either one of the clusters, or noted them as not clustered, we can perform an analysis of how the changes from one burst rate to another burst rate is performed. We can note from Figure 1c and 1d that the burst rates appear fairly stable. To examine this further, we process all burst rate observations to obtain a transition matrix. A transition occur from one burst to the next burst subsequent in time, and each of these bursts have an associated cluster or are unclustered. We show the resulting transition probabilities between the burst rate clusters in Figure 5. Here, there is a clear structure that can be observed. For clusters 7 to 12, the transitions follow a chain behavior where for a given cluster rate, transitions are occurring either to the rate above or the rate below. Or alternatively, and in fact more commonly, to a rate outside of any cluster. The most common transition is, for all clusters, back to the current cluster rate. In terms of the chain behavior, one can observe

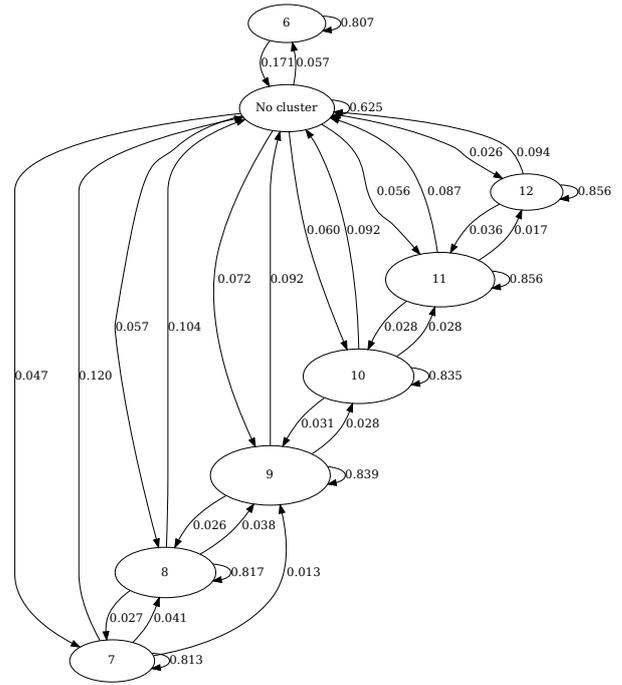


Fig. 5. Transition probabilities between states representing clusters of burst rates. The state numbers are the # in Table 1. The cluster size is proportional to the number of observed bursts in the cluster. Transition probabilities lower than 1% are not shown.

that there is a push towards the middle. In the sense that the probability of going from 7 to 8 is higher than going from 8 to 7, and correspondingly higher to go from 12 to 11 than from 11 to 12. This holds also for the other clusters. The observed chain-like behavior appears consistent with a radio resource allocation mechanisms that tries to optimize the physical rate to prevailing channel conditions. As far as we are aware, an observation of this behavior has not been previously reported. A notable exception is cluster 6, which does not have significant transition to the other clusters, but rather moves to and from unclustered rates. We note that from a rate modeling perspective as noted in Table 1, cluster 6 is an outlier.

VI. JOINT RATE AND DURATION TRANSITIONS

We now focus on both the burst rate, and the burst duration, as a joint characteristic of each network burst. This perspective enables us to construct a more comprehensive transition matrix that encapsulates the interactions between these two properties. The computation of this transition matrix is adapted in that we limit our analysis to bursts that are clustered, and which fall within a burst duration range of 1 to 7 frames. Consequently, this approach results in 49 distinct states, each representing a unique combination of burst rate and duration. The transition matrix is presented in Figure 6. Despite the complexity and potential visual overload of the figure, several observations can be discerned:

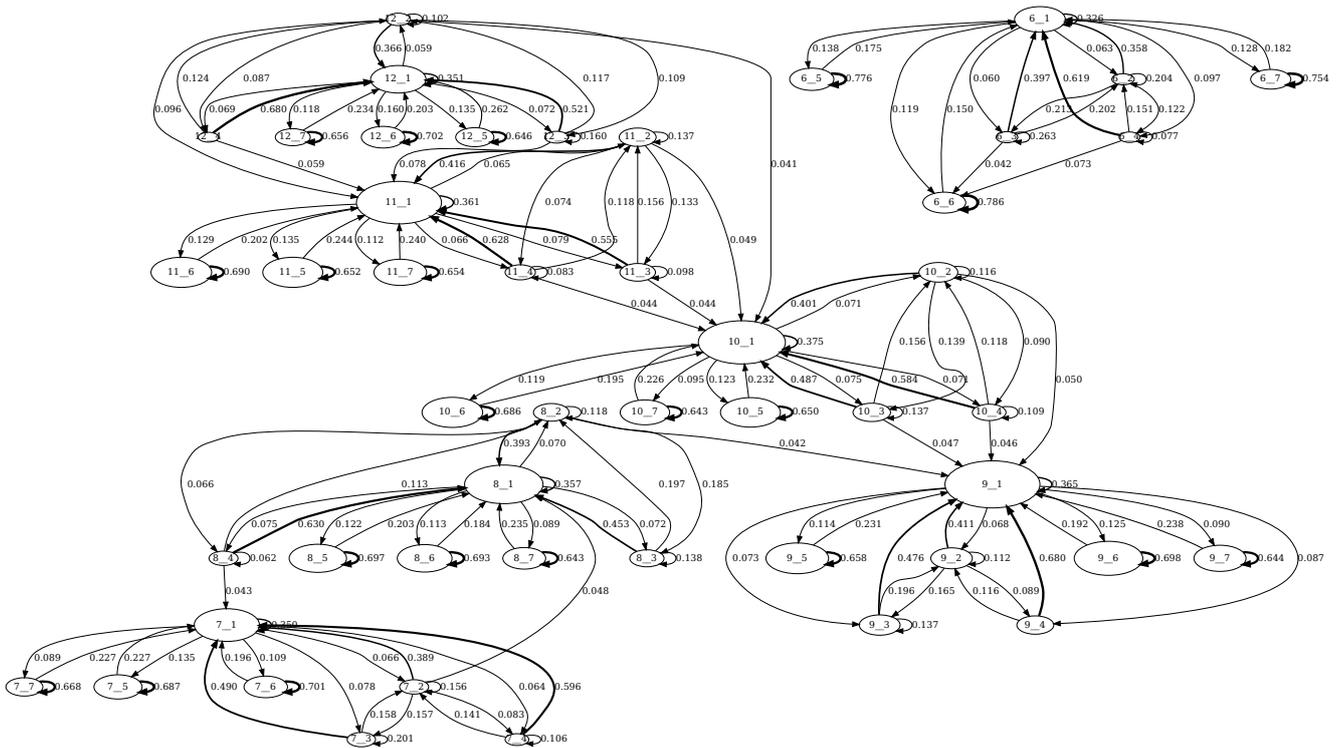


Fig. 6. Transition probabilities between the 49 combined C_D states where C corresponds to the burst rate cluster, and D the burst duration in number of 1.33 ms frames. The state size is proportional to the number of observations in the state. Transition probabilities lower than 4% are not shown.

- **Transition Patterns Between Clusters:** An interesting pattern emerges in transitions between states that refer to different rate clusters, i.e. where the C of C_D change. These transitions are predominantly, if not exclusively, to states representing short one-frame bursts at the new burst rate. These states are identifiable as those ending with a '_1', indicating a 1-frame burst.
- **Transition Patterns Within Clusters:** It can be observed that the change of burst duration within a cluster, i.e. where the D of C_D change, also predominantly occurs through a transitory single frame state '_1'. This implies that transitions between bursts of varying lengths are often mediated by a short, intervening single-frame burst.
- **Sticky State Transition Pattern:** In the majority of cases, but not universally, we observe that the most common transition for a state is a return to itself.
- **Exceptions to Sticky State Transition Pattern:** States that end with 2, 3, or 4 and thus have that number of frames display a different transition behavior. These states are typically transitory, meaning that the most frequent transition observed is not a return to the same state, but rather a move to another distinct state.
- **Isolation of Cluster 6:** In consistence with the results shown in Figure 5, Cluster 6 continues to demonstrate a unique pattern of isolation. Also for this expanded

transition matrix, it is evident that there are also no transitions from any of the 6_X states to any state belonging to another Z__ cluster.

While we cannot ascertain the underlying mechanism for these observations, they nevertheless contribute to additional knowledge regarding the design and functioning of the Starlink physical layer and radio resource allocation. As our measurement approach only requires hardware in the form of a hardware timestamp capable NIC, our measurements can easily be extended by others to obtain additional insights into this behavior. As our results are based on a single measurement campaign at a single location, the results are not expected to cover the full range of variability that are likely present in the physical layer. We however expect overall observations such as the tendency towards rate changes to neighboring rates, and performing cluster transitions through single-frame states, to hold also for the general case.

VII. EVALUATING EFFECTIVE THROUGHPUT

We now also consider the interburst silent time, which, together with burst rate and burst duration, forms the components for constructing the effective throughput observed over longer timescales. As depicted in Figure 1, there are interburst silent periods, which can vary in time duration. The relative proportion of time spent within a burst and between bursts is a major factor in determining the effective throughput on the scale of tens of milliseconds. We can define a complete

burst to be consisting of the burst duration, and the associated interburst silent time that follows after that particular burst. With the duration and silent times, and with consideration of the burst rate, we can now easily compute the effective throughput during such a complete burst. Figure 7 shows the results of such computations for the cluster burst rates and burst lengths examined in the preceding transition analysis, along with the length of the interburst silent times over the range of values that has been empirically observed. For these considered configurations the effective throughput varies from slightly above 25 Mbps up to almost 500 Mbps.

We now examine the empirical distribution of the three components of effective throughput in our collected measurements. To highlight the differences in physical burst rate and scheduling that lead to the effective throughputs that are observed, we consider three different subsets of measurement runs. We select the 10 percent of measurement runs with the lowest, middle and highest throughput over the complete run as reported by the Ookla Speedtest. The mean throughput of the measurement runs within these three deciles is 142, 222, and 268 Mbps, respectively. The bursts within these respective runs are examined, with the configuration on the three components determined and used to create a 2d histogram. These histograms are shown in Figure 8, which highlights the apparent differences between the deciles.

The runs with the lowest throughput to a considerable extent have single frame bursts, i.e. C__1, with relatively short interburst times. Longer bursts are mainly present in the cluster with lowest rate out of the ones considered for analysis and display. For these low throughput runs, there are also a number of bursts that fall into the lower burst clusters which are not shown here. Due to space and legibility considerations the number of clusters considered in analysis and visualization needs to be limited.

In comparison, the runs with throughput around the median of all runs show a distribution quite in the middle of the examined configurations. For these runs, the most common configuration is 8 ms bursts with a rate of 484 Mbps, and an interburst silent time of 5.333 ms, i.e. 4, 9__6. Notable is the somewhat peculiar pattern of the combinations of burst duration and inter-burst time. Across the burst rates, the common combinations are 5+3, 6+4 and 7+5 which sum to 8,10, and 12 time slots. So rather than summing to a constant burst duration plus interburst time for the scheduling, this seems to be done in multiples of two time slots.

Finally, the runs within the highest throughput decile share a similar behavior with regards to the distribution among burst durations and inter-burst times. However, the burst rates are now higher, with the majority of the bursts falling within the two highest rate clusters. We can note that these results by necessity reflect the conditions at the time of data collection and for the specific site of data collection. Data collected at other sites, and with different amount of competing users is very likely to have a different characteristics from the ones illustrated here. Nevertheless, the fundamental approach of using these three constituent components as factors when

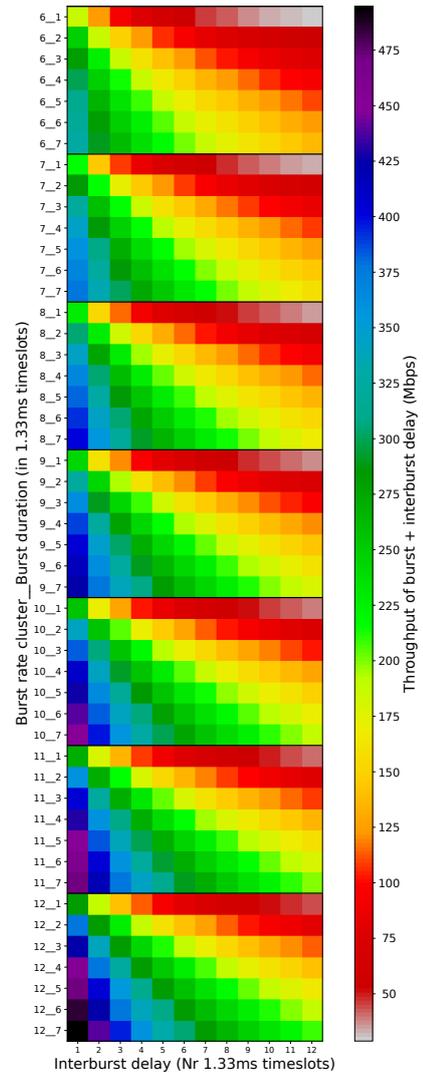
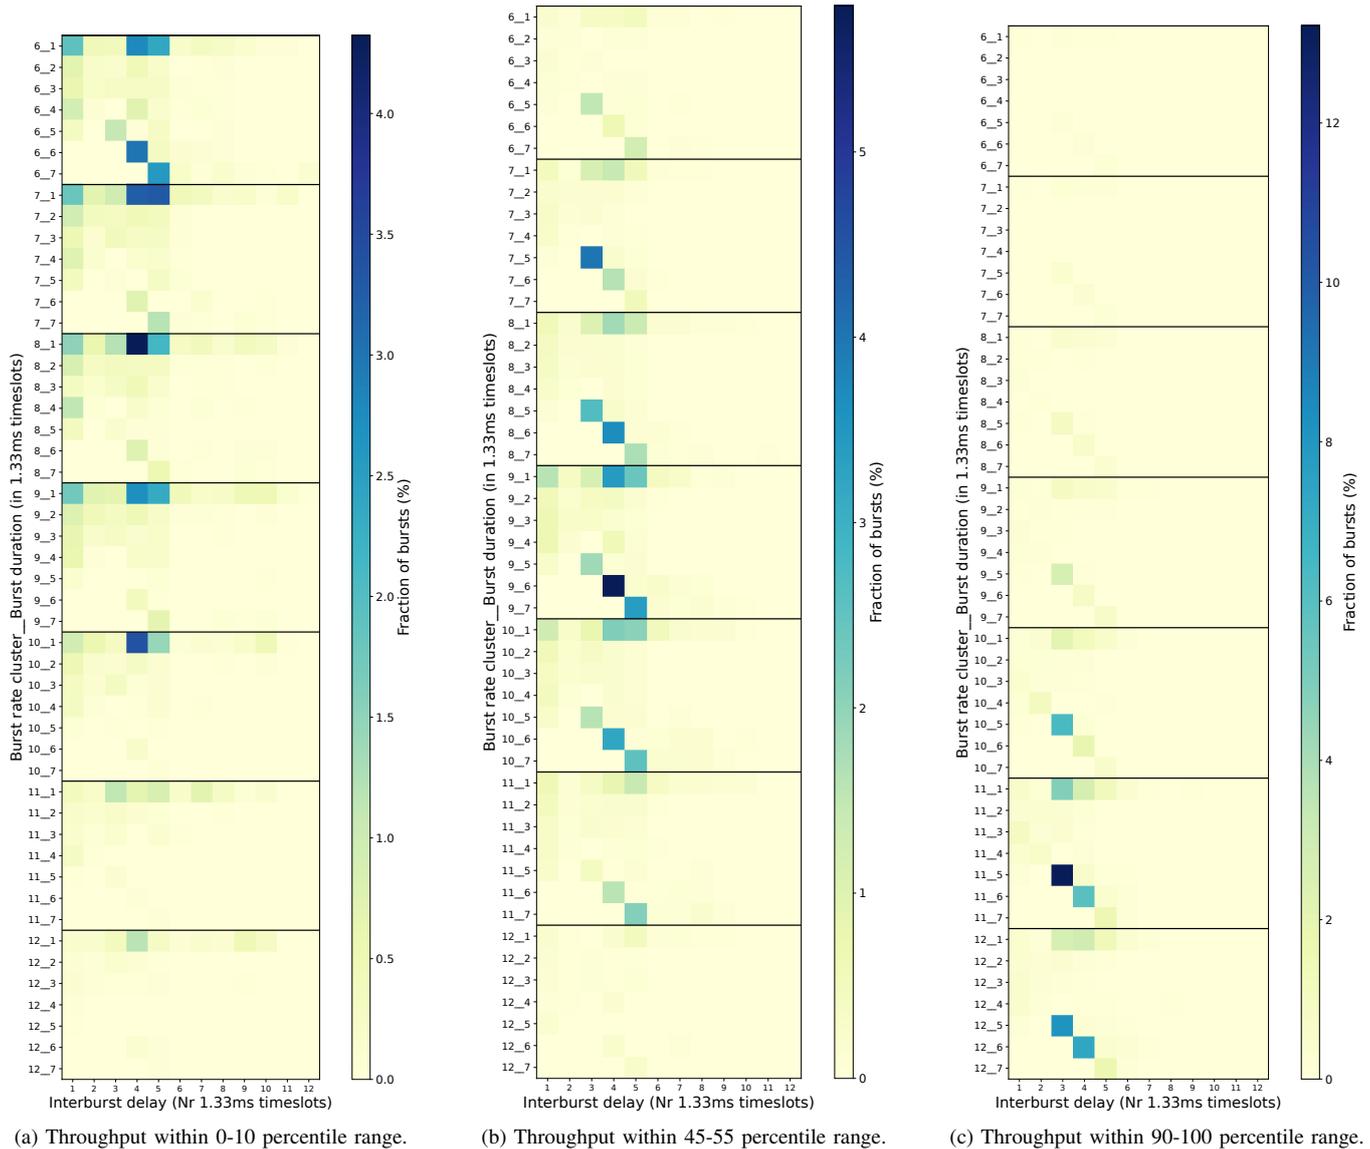


Fig. 7. Throughput when computed over sum of burst duration and interburst delay, for some burst rate cluster C and burst duration D, noted as C__D on y-axis.

examining the effective observed throughput over longer time scales is generally applicable across sites and time periods, and more measurements and accompanying analyses may lead to further improved understanding of Starlink system aspects.

VIII. RELATED WORK

Research on deducing Starlink system level details based on network layer measurements is relatively limited. Our previous investigations have explored Starlink's throughput across various time scales [5]. Notably, to our best knowledge, it was the first work confirming both the reconfiguration interval and frame timing based exclusively on network layer measurements. That work considered the throughput variation over multiple time scales, going from day to milliseconds. In that work, we made an initial observation of the bursty behavior of received Starlink traffic from the CPE Ethernet and made an initial characterization of burst durations and



(a) Throughput within 0-10 percentile range. (b) Throughput within 45-55 percentile range. (c) Throughput within 90-100 percentile range.

Fig. 8. Frequency of occurrence of different burst rate clusters and burst durations (C_D) on y-axis, and inter-burst delays on x-axis, for all bursts within three different throughput deciles of the 2195 measurement runs.

inter-burst pauses. A later deepened analysis [4] showed that our first initial burst characterization however had incorrect burst time measurements due to the sub-bursts occurring within the bursts. We devised a correction mechanism, and by using these updated burst characteristics we were able to construct a radio resource allocation model that provides specific details to the general frame structure described in [6]. Our analysis showed that a model where each Starlink OFDMA frame is using 287 symbols over 1000 subcarriers for user data transport of full size packets provide an excellent match for the base rate found in our empirical measurements.

More broadly, the existing literature that relates to our study can be categorized into two principal groups. The first group comprises measurement-focused studies, which examine network performance in terms of throughput and

latency as perceived by users [8], [10], [11], [18], [20]. Typically, these studies do not gather data with the level of timing detail and accuracy necessary for in-depth inference about system characteristics. Their primary aim is to describe the performance as experienced by users and the influencing elements, like satellite visibility and meteorological conditions. In the study by Tanveer et al. [18], the investigation measures round-trip times and packet loss rates using probing packets sent every 20 milliseconds. Another study by Mohan et al. [11] assesses latency at intervals of 3 milliseconds and throughput with 100 milliseconds granularity. In contrast, our analysis enables us to deduce throughput with a level of detail extending to instantaneous per-packet throughput, thanks to timing measurements accurate to tens of nanoseconds.

The second category of relevant research originates from

the fields of electrical and/or aeronautical engineering. These studies are primarily concerned with the physical layer details of the Starlink system, often to aid secondary applications like positioning [7], [12], [16] or passive radar systems [3]. They employ specialized radio technology to scrutinize the physical layer properties. In these studies, measurements of user data throughput are not a focus, as the interest lies in these secondary applications. An important contribution to our understanding of the physical layer comes from the blind signal identification work conducted in [6], where details of Starlink’s Ku-band transmissions were uncovered, including the overall frame structure.

IX. CONCLUSIONS

This work uses receiver side network measurements collected with the aid of a hardware-timestamp NIC to perform studies on Starlink physical layer rates, and in particular rate transitions. We construct rate transition matrices and study the resulting transition structure for hints on Starlink physical layer structure and behavior. Our results show that there is structure in the transition behavior, such as particular transition steps associated with rate switching, and that rate switching occurs mainly to neighboring rates. Although limited to single location measurements, this work represents a step forward in the understanding of the complex dynamics of Starlink, and points out a promising direction for further research of these aspects.

ACKNOWLEDGMENTS

The authors wish to thank Jonas Karlsson and Tobias Vehkajärvi for assistance with the data collection. This research was partly funded by the Swedish KK-foundation.

REFERENCES

- [1] 3GPP, “3gpp ts 23-700-27. study on 5g system with satellite backhaul (release 18),” *3GPP TS 23-700-27 v18.0.0 (2022-12)*, 2022.
- [2] M. M. Azari *et al.*, “Evolution of non-terrestrial networks from 5g to 6g: A survey,” *IEEE communications surveys & tutorials*, 2022.
- [3] R. Blázquez-García, D. Cristallini, M. Ummerhofer, V. Seidel, J. Heckenbach, and D. O’Hagan, “Experimental comparison of starlink and oneweb signals for passive radar,” in *Radar Conference*, 2023.
- [4] J. Garcia, S. Sundberg, and A. Brunstrom, “Inferring starlink physical layer transmission rates through receiver packet timestamps,” in *Accepted for publication in Proceedings of the 25th IEEE Wireless Communications and Networking Conference (WCNC)*, 2024.
- [5] J. Garcia, S. Sundberg, G. Caso, and A. Brunstrom, “Multi-timescale evaluation of starlink throughput,” in *Proceedings of the 1st ACM Workshop on LEO Networking and Communication*, 2023, pp. 31–36.
- [6] T. E. Humphreys, P. A. Iannucci, Z. M. Komodromos, and A. M. Graff, “Signal structure of the starlink ku-band downlink,” *IEEE Transactions on Aerospace and Electronic Systems*, 2023.
- [7] Z. M. Kassas *et al.*, “Navigation with multi-constellation leo satellite signals of opportunity: Starlink, oneweb, orbcomm, and iridium,” in *Position, Location and Navigation Symposium (PLANS)*, 2023.
- [8] M. M. Kassem, A. Raman, D. Perino, and N. Sastry, “A browser-side view of starlink connectivity,” in *22nd ACM Internet Measurement Conference*, ser. IMC ’22, 2022, pp. 151–158.
- [9] K. MacMillan, T. Mangla, J. Saxon, N. P. Marwell, and N. Feamster, “A comparative analysis of ookla speedtest and measurement labs network diagnostic test (ndt7),” *Proc. ACM on Measurement and Analysis of Computing Systems*, vol. 7, no. 1, pp. 1–26, 2023.
- [10] F. Michel, M. Trevisan, D. Giordano, and O. Bonaventure, “A first look at starlink performance,” in *Proceedings of the 22nd ACM Internet Measurement Conference*, ser. IMC ’22, Oct. 2022, pp. 130–136.
- [11] N. Mohan, A. Ferguson, H. Cech, P. R. Renatin, R. Bose, M. Marina, and J. Ott, “A multifaceted look at starlink performance,” *arXiv preprint arXiv:2310.09242*, 2023.
- [12] M. Neinaiaie and Z. M. Kassas, “Unveiling starlink leo satellite ofdm-like signal structure enabling precise positioning,” *IEEE Transactions on Aerospace and Electronic Systems*, 2023.
- [13] Ookla. (2023) Speedtest CLI: Internet speed test for the command line. [Online]. Available: <https://www.speedtest.net/apps/cli>
- [14] M. Rajiullah, G. Caso, A. Brunstrom, J. Karlsson, S. Alfredsson, and O. Alay, “CARL-W: a testbed for empirical analyses of 5g and starlink performance,” in *In 3rd ACM Workshop on 5G and Beyond Network Measurements, Modeling, and Use Cases (5G-MeMU ’23)*, 2023.
- [15] Starlink. (2023) Order starlink. [Online]. Available: <https://www.starlink.com/>
- [16] W. Stock, C. A. Hofmann, and A. Knopp, “Leo-pnt with starlink: Development of a burst detection algorithm based on signal measurements,” in *WSA & SCC 2023; 26th ITG Wshop on Smart Antennas and 13th Conf on Systems, Communications, and Coding*. VDE, 2023, pp. 1–6.
- [17] S. Sundberg, J. Garcia, and A. Brunstrom, “Characterizing wireless link throughput with eBPF and hardware timestamps,” in *28th International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD 2023)*, 2023.
- [18] H. B. Tanveer, M. Puchol, R. Singh, A. Bianchi, and R. Nithyanand, “Making sense of constellations: Methodologies for understanding starlink’s scheduling algorithms,” *arXiv preprint arXiv:2307.00402*, 2023.
- [19] M. A. Vieira, M. S. Castanho, R. D. Pacifico, E. R. Santos, E. P. C. Júnior, and L. F. Vieira, “Fast packet processing with eBPF and XDP: Concepts, code, challenges, and applications,” *ACM Computing Surveys (CSUR)*, vol. 53, no. 1, pp. 1–36, 2020.
- [20] H. Zhao, H. Fang, F. Wang, and J. Liu, “Realtime multimedia services over starlink: A reality check,” in *33rd Workshop on Network and Operating System Support for Digital Audio and Video*, 2023, pp. 43–49.