

WetLinks: a Large-Scale Longitudinal Starlink Dataset with Contiguous Weather Data

Dominic Laniewski¹, Eric Lanfer¹, Bernd Meijerink², Roland van Rijswijk-Deij², Nils Aschenbruck¹

¹Osnabrück University - Institute of Computer Science, Osnabrück, Germany

²University of Twente - Design and Analysis of Communication Systems Group, Enschede, The Netherlands
{laniewski, lanfer, aschenbruck}@uos.de, {bernd.meijerink, r.m.vanrijswijk}@utwente.nl

Abstract—Low Orbit Satellite (LEO) networks such as Starlink promise Internet access everywhere around the world. In this paper, we present *WetLinks* - a large and publicly available trace-based dataset of Starlink measurements. The measurements were concurrently collected from two European vantage points over a span of six months. Consisting of approximately 140,000 measurements, the dataset comprises all relevant network parameters such as the upload and download throughputs, the RTT, packet loss, and traceroutes. We further augment the dataset with concurrent data from professional weather stations placed next to both Starlink terminals. Based on our dataset, we analyse Starlink performance, including its susceptibility to weather conditions. We use this to validate our dataset by replicating the results of earlier smaller-scale studies. We release our datasets and all accompanying tooling as open data. To the best of our knowledge, ours is the largest and most complete Starlink dataset to date.

Index Terms—Starlink, Satellite Communication, LEO, Dataset, Network Traces, Network Measurements, Starlink Dataset, Starlink Measurement, Replicability

I. INTRODUCTION

In networking research, new ideas and approaches are commonly developed and tested via simulations. These simulations often require data on realistic network conditions such as throughputs, packet loss, and latency. This data can either be generated using network simulators such as *ns3* [29], by using models such as the Gilbert-Elliot model [8] for bursty packet loss, or by replaying real-world traces. The benefit of using real-world traces is that they represent realistic network conditions. Trace-based evaluations allow other researchers not only to replicate presented results, but also to benchmark their own approaches against existing solutions under the same network conditions, leading to an overall increased quality of research. The need for real-world network traces is amplified by the rise of machine learning (ML) methods in many networking research areas, such as rate adaptation for video-streaming [11] and point cloud streaming [15], traffic prediction [2], and intrusion detection [25]. ML-based solutions, especially ones that apply deep learning (DL) techniques, require large training datasets to achieve a high level of performance and to avoid overfitting of the models to certain network conditions.

Lately, the first studies of SpaceX’s Starlink Low Earth Orbit (LEO) satellite were published [12], [17], [19]. Most

studies published since Starlink opened for public access in 2020 focus on first impressions of achievable performance. While these studies made their datasets public, all of these datasets suffer from limitations. They are all: 1) limited in size, 2) have an *ad hoc* methodology, 3) lack a clear description of measurement conditions, such as equipment used, user load, etc., or 4) focus only on specific network parameters such as latency. In general, creating a large-scale longitudinal trace-based dataset of continuous Starlink measurements of all relevant network parameters is a costly project, as it requires exclusive use of a Starlink dish for the duration.

In this paper, we intend to remedy this by presenting *WetLinks*: a large-scale longitudinal Starlink dataset. *WetLinks* consists of six months (Oct. ’23 – Mar. ’24) of orchestrated Starlink measurements at approximately 3-minute intervals from two European sites – in Osnabrück (GER) and Enschede (NL). With a distance as the crow flies of ± 80 km between the measurement stations, both dishes are connected to the 53° orbit and likely connect to the same satellites in sequence. The measurements include all relevant network parameters such as the download and upload throughputs, the Round-Trip-Time (RTT), the packet loss rate (PLR), traceroutes, as well as accurate weather data from reference weather stations placed in direct physical proximity to the dishes. Our contributions are as follows:

- We make *WetLinks* publicly available. It consists of approximately 80,000 measurements from Osnabrück and 60,000 from Enschede. To the best of our knowledge, this is the largest and most complete Starlink dataset to date.
- We pair our measurements with accurate weather data directly captured next to the dishes, and we conduct an in-depth analysis of the impact of weather conditions on Starlink’s performance.
- We validate our dataset by replicating the results of earlier smaller-scale studies.
- We release companion tooling that merges different data sources (performance measurements, weather data, ...) together with the dataset.

The remainder of this paper is structured as follows. Sec. II gives background on Starlink and its radio communications. In Sec. III we discuss related work on existing studies of Starlink performance and weather impact on satellite communications. Next, Sec. IV describes our measurement setup. In Sec. V,

we validate our dataset by analyzing performance and weather impact and comparing this against earlier studies. Finally, we conclude the paper in Sec. VI.

II. BACKGROUND

A. Starlink Overview

Starlink is a broadband Internet service provided by SpaceX. It relies on a Low Earth Orbit (LEO) satellite constellation initially (from 2018) orbiting at an altitude of 1,100 km [5], and since 2021 the orbits have been lowered to between 540 km and 570 km [6]. As of April 2024, SpaceX launched 6,258 satellites, of which 5,214 are operational [18]. These are deployed in four orbits of varying density at inclinations of 43°, 53°, 70°, and 97.7°. They have 1,514, 3,071, 396, and 233 operational satellites, respectively [18]. The 53° orbit provides most of the service, while the 70° and 97.7° orbits mainly serve the Earth’s polar regions [20].

Customers connect to the network via their user terminal, which is colloquially known as a “Dishy”. Bidirectional communication between the terminal and the satellite makes use of K_u -Band beams. Communication between satellites and ground stations makes use of the K_a -Band. Starlink has a reconfiguration interval of 15 s, where the terminal can be rescheduled to another satellite and frequencies and routing resources can be reallocated [20]. SpaceX claims that downlink speeds up to 220 Mbit/s and uplinks speeds up to 25 Mbit/s are possible at latencies ranging from 25 to 60 ms [28]. SpaceX offers *Standard* and *Priority* service plans. Priority traffic is given network precedence over standard traffic [27] at times of high network load and congestion.

The Starlink network uses a one-hop bent-pipe routing approach if both the user terminal and the ground station are within line of sight of a single satellite [17]. In cases where one satellite does not cover both endpoints, inter-satellite links (ISL) using laser-beams can be utilized to form an extended multi-hop bent-pipe [20]. Only newer generation satellites starting with version 1.5 support ISL.

B. K_u - and K_a -band fundamentals

As discussed above, Starlink uses the K_u - and K_a -band [1] for satellite to earth communications. The K_u -band is used for terminal to satellite communications, where frequencies between 10.7-12.7 GHz are used for satellite-to-terminal and frequencies between 14.0-14.5 GHz are used for terminal-to-satellite. The K_a -band is used for satellite to ground station communication, with frequencies between 17.8-19.3 GHz being used for satellite-to-ground, and frequencies between 27.5-30.0 GHz being used for ground-to-satellite [26].

Since both bands run at frequencies higher than 10 GHz, they are susceptible to tropospheric phenomena that affect signal propagation [9], [22]. The effects include, among others, attenuation due to precipitation and clouds. In this work, we combine weather data with throughput measurements on the terminal-to-satellite link to analyze the impact of precipitation and cloud attenuation on Starlink throughput.

III. RELATED WORK

Since Starlink opened to the public in 2020, several studies have analyzed the performance and internal functioning through simulations and real-world measurements.

On the simulation side, *StarPerf* [14] and *Hypatia* [13] simulate the network behavior of satellite constellations. *LeoEM* [3] can also capture the dynamics of LEO satellite networks, and can be used to analyze TCP behavior.

Most measurement studies look at performance aspects only. Michel et al. [19] characterize Starlink performance in terms of throughput over TCP and QUIC, latency and packet loss. Their longest measurement campaign spans 5 months. Izhikevich et al. [10] conduct world-wide latency measurements over one month by actively probing publicly exposed devices at unspecified intervals that are connected via Starlink. Pan et al. [21] also measure performance and perform traceroutes, focusing on analyzing Starlink’s point-of-presence (POP) structure. Raman et al. [24] compare Starlink against medium- and geostationary orbit (MEO/GEO) satellite networks finding that Starlink outperforms these older technologies. Garcia et al. [7] study Starlink performance with a specific focus on system-specific timing structures such as frequency scheduling and beam switching. Finally, Mohan et al. [20] analyze M-Lab speed test data of Starlink measurements and compare this against RIPE Atlas measurements from probes with Starlink connectivity and measurements using their own Starlink dishes.

In contrast, only a few studies conduct initial smaller-scale analyses of the impact of weather on Starlink performance as part of their larger-scale measurement campaigns. Kassem et al. [12] recruited volunteers that installed a browser extension and performed a measurement campaign over a period of 6 months. They augment this with a limited but more intensive measurement campaign from three locations using Raspberry Pi hosts connected to volunteer Starlink dishes for an unspecified amount of time. They study weather influence for one location using openly available data from a nearby weather station and find that the transit time of web-pages increases with rain. Zhao et al. [33] focus on real-time multimedia services, such as video-on-demand, live-streaming and video conferencing. They note that performance is typically adequate for these services, but may be impacted by satellite handovers and adverse weather in the form of thunderstorms. Ma et al. [17] collect routing information and conduct performance measurements for different types of applications, such as video streaming. They deploy 4 Starlink dishes and measure over a period of 7 months at irregular intervals. Furthermore, they collect weather data from a climate station 10 km away from one of their measurement stations. They find that rain can cause the UDP download throughput to almost halve.

In contrast to previous studies discussed above, we are the first to provide a large and fine-grained dataset of continuous Starlink measurements at minute-intervals conducted over the course of six months from two European cities. It spans all relevant network parameters, as well as high-precision weather




Fig. 1: The Measurement Setup. The purple *C* indicates data endpoints, that are collected by the data collection server.

data measured with reference weather stations placed directly next to the Starlink dishes, making it the most complete Starlink dataset to date. The dataset allows for an accurate in-depth analysis of the impact of weather conditions on Starlink performance. Furthermore, the open nature of our dataset paves the way for future studies that require large, continuous input datasets, such as studies that rely on machine learning approaches.

IV. MEASUREMENT SETUP

We conduct Starlink measurements and collect weather data at two sites: (1) on the rooftop of one of the university buildings in Osnabrück, Germany (Fig. 2a), and (2) on one of the university buildings in Enschede, the Netherlands (Fig. 2b). The distance as the crow flies between both sites is 80 km. This means it is likely that both sites connect to the same satellites in sequence. Furthermore, based on the orientation of the dish, the direction azimuth and elevation of the dish metadata, and publicly available satellite information [23], both measurement stations likely connect to the 53° orbit. At both positions, we deployed a similar measurement setup as shown in Fig. 1. Each setup consists of four devices: (1) a Starlink user terminal (a.k.a. a “dishy”), (2) a Starlink router configured in pass-thru mode, (3) a measurement client on a VM running Ubuntu and (4) a Froggit DP2000 weather station deployed directly next to the user terminal. The user terminals have an unobstructed view of the sky, and run the most up-to-date firmware as pushed to the terminal by Starlink. Both sites have a regular Starlink subscription. The measurement clients both have a secondary 1 Gbit/s network interface on the local network, which we use to orchestrate measurements and to retrieve results. There are two key differences between both sites: the site in Osnabrück has a second generation user terminal (v2), whereas the site in Enschede has a first generation user terminal (v1). Secondly, the site in Enschede




Fig. 2: Our measurement setups. In this picture, the Froggit DP2000 weather station can be seen placed directly next to the Starlink dish.




Fig. 3: The measurement process of the network link parameters.

needs to contact the measurement server via a WAN interface; the measurements are instrumented such that any delays on this link do not impact results. We perform all measurements against a server located in Osnabrück. This server has a 5 Gbit (shared) WAN connection to the Internet.

Based on publicly available information about Starlink ground stations [23], we assume that traffic from both locations will likely reach a ground station in Aerzen, Germany, which is 145 km from Osnabrück and 167 km from Enschede as the crow flies. From this ground station, we observe traffic entering the public Internet at Starlink’s POP in Frankfurt.

Finally, we gather results on the data collection server, located in Osnabrück. This server records the following data:

- throughput, RTT, packet loss and traceroutes;
- fine-grained weather data from the sensors on both sites;
- API data from public weather services for Germany (DWD) and the Netherlands (KNMI);
- all diagnostic data provided by the Starlink dishy (software version, obstruction, debug information, etc.).

A. Network Link Measurement

We use a three step sequential process to measure network link parameters, as shown in Fig. 3. The first two steps (the *iperf3* and *ping* measurements) are mandatory and are executed approximately every 3 minutes. The third step (*MTR* measurements) is executed every third run, leading to approximately 6 minute intervals.

We use *iperf3* (version 3.9) [4] to measure up- and down-stream throughput. The measurement of both directions is conducted in parallel using UDP with two separate measurements. While *iperf3* has a bidirectional mode, we do not use this mode since it is susceptible to CPU limitations, which can lead to inaccurate measurement results. We limit the target

TABLE I: Measured Features of the Froggit DP 2000 weather station

Feature	Unit	Accuracy
Temperature	°C	± 0.3 °C
Humidity	%	$\pm 3.5\%$
Rain Volume	mm	$\pm 10\%$
Windspeed	m/s	$< 10\text{m/s}: \pm 0.5\text{m/s}$ $\geq 10\text{m/s}: \pm 5\%$
Wind Direction	°	$< 2\text{m/s}: \pm 10^\circ$ $\geq 2\text{m/s}: \pm 7^\circ$
UV	UV-Index	Range: 0-15
Barometric Pressure	hPa	± 5 hPa in 700 – 1,100 hPa range resolution: 0.1 hPa (0.01 in Hg)
Solarradiation	W/m ²	Not specified

bitrate for the uplink to 100Mbit/s and for the downlink to 500Mbit/s. This ensures *iperf3* does not experience CPU limitations and that our measurements comply with Starlink’s fair-use policy, ensuring that the data rate is not artificially limited by Starlink. Each measurement takes 15 seconds, with a five-second timeout grace period.

We utilize *ping* to measure the round trip time (RTT) and packet loss. This measurement consists of a sequence of 250 packets with an interval of 0.1 s. Aside from packet loss and average RTT we also report the best and worst RTT as well as the standard deviation.

We collect routing information from our client to our server using *Matt’s traceroute* (MTR, version 0.94) [32] with 15 report cycles. For each hop, the output includes the host IP or host name, packet loss information, and ping statistics.

B. Weather Measurement

To accurately capture weather conditions at our measurement positions, we place a Froggit DP 2000 weather station directly next to our Starlink dishes. Table I lists the parameters the weather station collects. We also record calculated features produced by the weather station, specifically the dew point, windchill, wind gusts, and daily, weekly, monthly and annual rainfall. We also registered our weather stations with the Weather Underground [30], [31]. Measurements are collected at 1-minute intervals. In addition to our own weather measurements, we also source data from the national weather services of Germany (DWD) and The Netherlands (KNMI). For Osnabrück, we collect data for station ID 00342 located at Belm, approximately 10.5km from our antenna, and for Enschede, we collect data for station ID 290 located at Twente Airport, approximately 4.5 km from our antenna.

C. Dataset

Based on the measurement architecture discussed above, we perform measurements resulting in the WetLinks dataset, which we release publicly together with all pre-processing scripts [16]. Our measurements in Osnabrück start on September 14th, 2023, and our measurements in Enschede start on October 12th of the same year. Both measurements were running till March 2024. This results in a dataset of continuous measurements spanning 6 months for Osnabrück, with approximately 80,000 data points, and 5 months for Enschede, with approximately 60,000 data points.




Fig. 4: Download and upload throughput for Osnabrück and Enschede.

TABLE II: Throughput Statistics [Mbit/s]

	Mean	Median	25-percentile	75-percentile
Download				
Osnabrück	212.8	215.8	168.2	257.2
Enschede	238.7	240.5	199.1	278.5
Upload				
Osnabrück	16.0	14.9	11.6	19.0
Enschede	17.1	16.2	13.0	20.0

We use a simple comma-separated format (CSV) for data. The raw dataset consists of CSV files that respectively record throughput measurements (*net_iperf.csv*), latency measurements (*net_ping.csv*), traceroute measurements (*net_traceroute.csv*), debug data from the Starlink dishes (*starlink.csv*), and finally data from the weather stations positioned next to the dishes (*froggit.csv*). The pre-processing scripts we provide can be used to combine these datasets and harmonize the different timescales at which the measurement data is collected (e.g., averaging weather data, which is collected at higher frequency than throughput measurements). Providing the raw data together with the pre-processing scripts allows users of the data to fully reproduce results and to make their own choices in harmonizing data for the purpose for which they wish to use the data.

Finally, we note that the dataset comes with a few limitations. First, our throughput measurements occasionally fail. This leads to empty fields in the CSV file that can easily be filtered out. Second, our traceroute measurements also occasionally fail, likely due to intermittent Starlink outages. This leads to traceroutes that contain only two hops and then time out. Finally, our weather station in Enschede suffered a two-week outage due to the battery running out over the end-of-year break; this means weather data for this period is missing for the Enschede location.

V. ANALYSIS

In this section, we discuss the analyses we performed on the WetLinks dataset. Our focus is on verifying the quality of our dataset. We do this by comparing our analysis to earlier studies for different metrics.

A. Throughput Analysis

Fig. 4 shows boxplots of download and upload throughput. Additional statistics are shown in Tab. II. The large

TABLE III: Packet Loss Statistics

Site	Samples with loss	Min. PLR	Max. PLR	Mean PLR	σ PLR
Osnabrück	28.4%	0.4%	87.2%	1.0%	2.9%
Enschede	33.3%	0.4%	84.0%	1.0%	2.4%




Fig. 5: The PLR plotted for the 25th of October 2023 from our measurements.

interquartile ranges and whiskers indicate a large variability in performance for both down- and upload. The download measurements range from close to 0 Mbit/s to more than 400 Mbit/s, significantly exceeding the advertised maximum download throughput of 220 Mbit/s. The upload measurements range from close to 0 Mbit/s to more than 50 Mbit/s. The speeds measured in Enschede consistently exceed the ones in Osnabrück. Considering that both are connected to the 53° orbit and their geographic proximity, we speculate that this may be caused by the difference in dish versions.

Comparing our results to existing studies is challenging, as these often use different transport protocols that are sensitive to packet loss (e.g., TCP or QUIC). Michel et al. [19] report TCP download throughputs of up to 400 Mbit/s, comparable to our UDP measurements. Furthermore, they reported median TCP download and upload throughputs of 178 Mbit/s and 17 Mbit/s, respectively. In contrast, Kassem et al. [12] report median throughputs of 123 Mbit/s and 11 Mbit/s. Our measured download throughputs are significantly higher, highly likely due to the use of loss-sensitive congestion control algorithms, e.g., TCP CUBIC, in [12] and [19].

B. Packet Loss Analysis

We next turn our attention to packet loss. Packet loss occurs frequently during our measurements. The overall statistics in Tab. III show that well over 25% of our measurements contain some packet loss. The minimum packet loss rate (PLR) was 0.4%, which corresponds to one out of 250 packets sent during the loss measurement is lost.

Fig. 5 shows the packet loss pattern over a single randomly selected 24-hour period. While the PLR is low most of the time, short spikes can be observed. These spikes typically only affect single measurements and can be interpreted as short burst losses. There is no evidence to suggest a correlation between loss events in Osnabrück and Enschede. Fig. 6 shows an ECDF of the PLR over all measurements affected by packet loss. As the plot shows, the vast majority ($\pm 65\%$) of loss events concern only a single packet (PLR 0.4%). Furthermore, over 80% of measurements have a PLR $\leq 1\%$, and only some 3% of measurements have a PLR $\geq 5\%$.




Fig. 6: ECDFs of the packet loss rates.

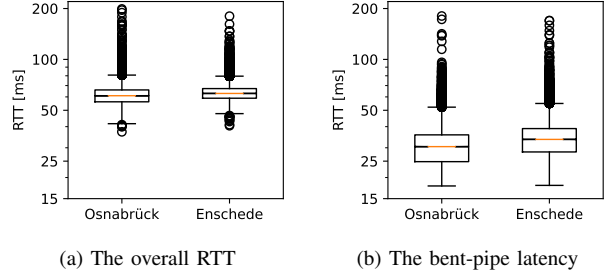


Fig. 7: Boxplots of the latency.


Related work by Ma et al. [17] reports cyclic burst loss patterns over 12 hours. We do not see similar patterns in our data. Michel et al. [19] report a PLR of 0.4% on the downlink and 0.45% on the uplink for low-load periods. Our data shows comparable behavior with additional short spikes caused by burst losses. These burst losses have also been reported by Kassem et al. [12] who found that they likely primarily occur during satellite handovers.

C. Latency Analysis


Fig. 7a shows the overall RTT distribution for the route from the measurement node to our server. The mean RTTs are 61.52 ms ($\sigma=7.81$ ms) and 63.53 ms ($\sigma=6.87$ ms) for Osnabrück and Enschede respectively. Approximately 72% of measurements have an RTT within 1σ , with outliers up to 200 ms.

To further analyze the cause of the high variance, we examine the RTT distribution for the bent-pipe section of the route in Fig. 7b based on the RTT data from our traceroute measurements. The bent-pipe has a mean RTT of 31.08 ms ($\sigma=11.18$ ms) for our location in Osnabrück and 34.56 ms ($\sigma=12.0$ ms) for Enschede. Approximately 85% of measurements have an RTT within 1σ , with outliers up to 200 ms. These numbers indicate that RTT variance of the complete link is mostly caused by the bent-pipe section of the path.


Overall, the average bent-pipe RTT of 31-35 ms is well within the expected range from previous works [12], [20] of 30-40 ms for measurements in the EU on the 53° orbit. Our conclusion that the variance in latency is primarily caused by the bent-pipe section of the path also aligns with previous work [10], [12], [17], [19].




(a) Download Throughput Osnabrück



(b) Download Throughput Enschede



(c) Upload Throughput Osnabrück



(d) Upload Throughput Enschede

Fig. 8: All measurements binned into 24 bins representing the hours of a day. The different periods are colored.

D. Impact of The Time of Day

Fig. 8 shows the impact of the time of day on network parameters. Below, we analyze the impact of the time of day in detail for the different network parameters.

Download Throughput: Fig. 8a and Fig. 8b illustrate that a typical day can be split into multiple periods. In the early morning, between 0h-5h, the median download throughput is around 230 Mbit/s in Osnabrück and 250 Mbit/s in Enschede. From 6h to approximately 15h, it decreases to around 210 Mbit/s in Osnabrück and 235 Mbit/s in Enschede. Then, a further drop can be observed until 19h, which sees the lowest throughput of 180 Mbit/s and 205 Mbit/s in Osnabrück and Enschede, respectively. Afterward, the throughput recovers, reaching 210 Mbit/s in Osnabrück and 225 Mbit/s in Enschede around 22h.

Our findings are consistent with previous measurements within the EU on the 53° orbit [7], [12], where the maximum download throughput was found around 5h [7], or between 0h-6h [12]. Additionally, Kassem et al. [12] reported that the maximum throughput can be twice as high as the minimum throughput on a day. While this might be the case for single days, our data suggests that on average, the minimum download throughput of a day is approximately 20% lower than the maximum. This is also consistent with findings by Michel et al. [19], who report fluctuations around $\pm 10\%$.

Our data does not provide a clear explanation for the observed fluctuations. Since the routes remain largely constant also during the hours with lower throughput, it is likely that the

general load within the Starlink network causes the observed behavior, especially in the evening hours from 17h-22h.

Upload Throughput: Fig. 8c and 8d indicate an increasing upload throughput between 0h-5h, and an almost constant one for the rest of the day. The minimum throughput is reached around 2h with 13.22 Mbit/s in Osnabrück, and 15.0 Mbit/s in Enschede. The maximum throughput is reached at 5h and with 16.29 Mbit/s in Osnabrück, and at 16h with 17.39 Mbit/s in Enschede. The interquartile ranges, the whiskers and the outliers indicate a strongly fluctuating upload throughput between 0 Mbit/s and 40 Mbit/s. Interestingly, there are almost no Starlink outages at 0 Mbit/s. This implies that the Starlink uplink appears to be more stable than the downlink. Similar to the download throughput our data shows similar behavior to previous studies [12].

RTT and Packet Loss: Neither the RTT, nor the packet loss are impacted by the time of day. The median RTT is $62 \text{ ms} \pm 3 \text{ ms}$ throughout the day. The median PLR is constantly 0.4%, with occasional burst losses that can occur at any time.

E. Weather Impact

As discussed earlier in Section III, weather conditions can have an impact on satellite communication. We created a correlation matrix shown in Fig. 9 to analyze their impact on the Starlink network parameters. Rain shows a weak negative correlation of about -0.165 and -0.192 to the download throughput for the measurements in Osnabrück and Enschede, respectively. All other weather factors have a correlation coefficient to the network parameters of close to zero, indicating no correlation.

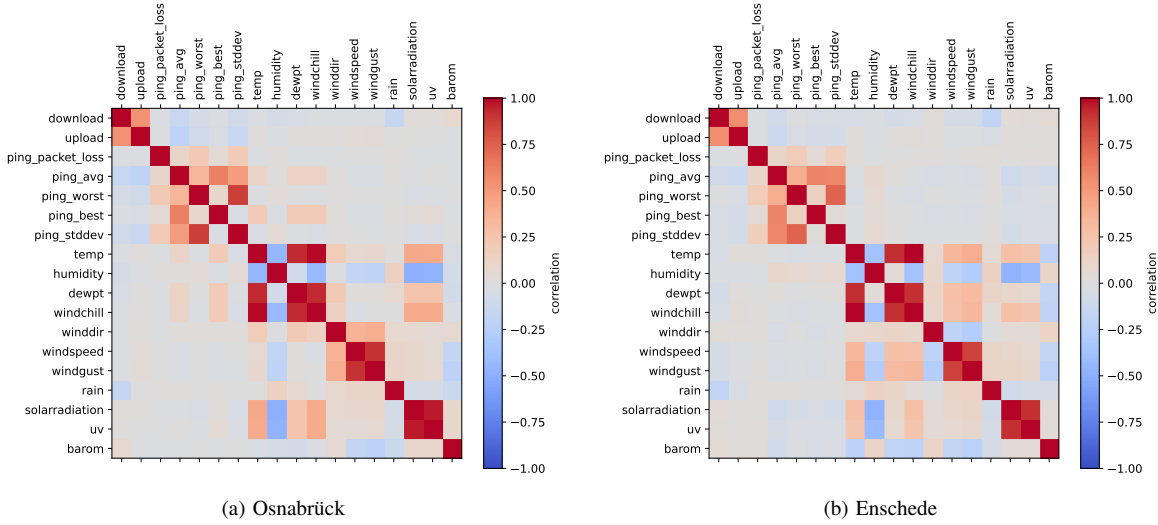


Fig. 9: Weather Correlation Matrix

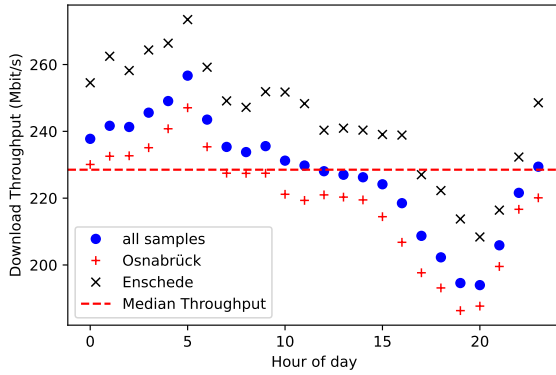


Fig. 10: Median download throughput rate per hour

1) *The Impact of Rain:* We continue our analysis with the hypothesis that rain intensity has an impact on download throughput. To verify this hypothesis, we first perform preprocessing on the data. We start by removing outliers using the interquartile range (IQR) with factor 1.5. Next, we perform a preprocessing step to remove the time of day impact on the download throughput, as shown in Fig. 8a and 8b. For that, we grouped our data into hourly buckets and performed a median correction for each bucket. Specifically, we calculated the download throughput median of all non-rainy samples of the complete dataset, and the median of the non-rainy samples within each hourly bucket. The result is shown in Fig. 10. We can clearly observe the influence of the time of day, which is consistent at both measurement sites (same timezone), with a difference of over 60 Mbit/s between the lowest and highest throughput. Using this data, we calculate the difference of the hourly median to the overall median and apply this difference as a correction factor (c_{dt}), as listed in Table IV, to corresponding samples. We used a product of both datasets

TABLE IV: Download throughput correction factors c_{dt} to eliminate the time of day influence

Hour	c_{dt}	Hour	c_{dt}
0	-9.24	12	0.46
1	-13.14	13	1.51
2	-12.79	14	2.26
3	-17.07	15	4.36
4	-20.55	16	10.02
5	-28.15	17	19.80
6	-14.99	18	26.23
7	-6.82	19	33.93
8	-5.27	20	34.55
9	-7.06	21	22.63
10	-2.71	22	6.93
11	-1.25	23	-0.90

for this step to come up with a more robust correction factor, by having more non-rainy samples.

A limitation of this approach is that it does not account for inter-day variance (e.g., weekday/weekend, ...).

To then verify the hypothesis of rain influence, we group the samples into rain buckets. Each bucket contains all samples with $r > u$ and $r \leq v$, where r is the level of precipitation in mm, u the lower threshold, and v the upper threshold of the bucket. We created 7 buckets, while the last bucket contains all samples with a precipitation > 5 mm. Additionally, we calculated a regression line, to test if an increasing rain level correlates with a change in download throughput. As the plots in Fig. 11 show, there is a clearly visible correlation, having an R value of -0.21 at both sites, indicating a weak correlation. A null hypothesis significance test showed a p -value below 0.05, which confirms our hypothesis that rain influences the download throughput significantly. The boxplots show large whiskers, indicating that factors other than rain and the time of day have a significant impact on the download throughput. These large whiskers are the main cause of the weak correlation. To reduce this noise and to further isolate

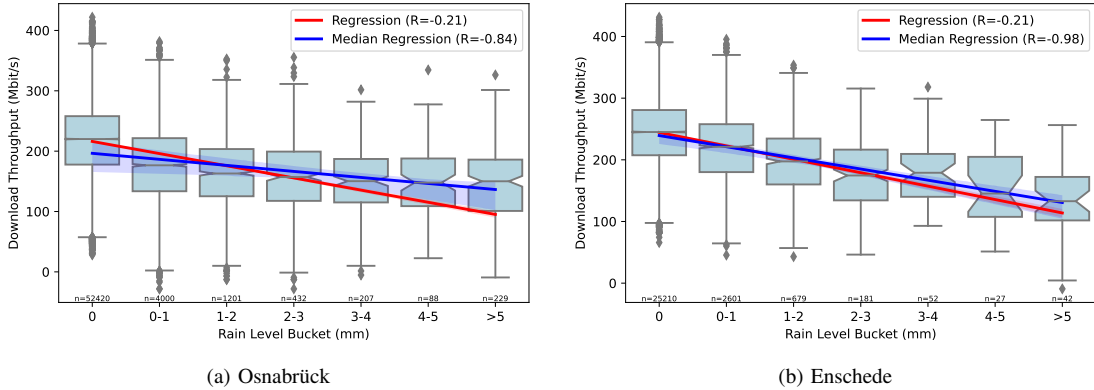


Fig. 11: Download throughput grouped by rain level buckets, with regression functions. The median regression is computed on the median values for each bucket, to remove the noise and variance, caused by unknown effects.

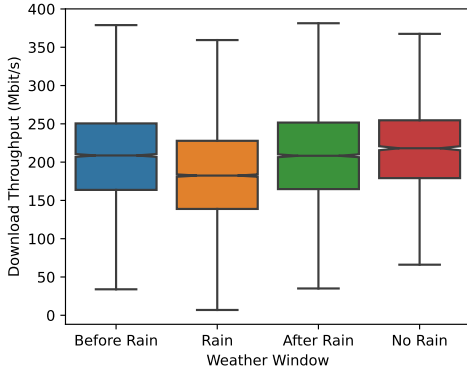


Fig. 12: Boxplots for different weather conditions. For each rain window, 4 samples before and after the rain are used. We assume cloudy conditions before and after rain. In the no rain case, only samples with high solar radiation are used, to ensure that there are no clouds.

the impact of rain, we added regressions through the medians of the buckets, which have an R of -0.84 and -0.98 , indicating a strong linear correlation between rain and download throughput. We argue that this is a valid approach for noise reduction since the confidence intervals of the medians as well as the IQRs are small. We also tested the upload throughput against rain buckets, however, having R values below 0.05 for both locations, we do not observe a significant correlation here. This may be explained by the lower bandwidth offered for upload, which provides more room for adaptation and mitigation strategies.

2) *The Impact of Clouds*: Clouds typically coincide together with rain and may also affect the download throughput. To differentiate the impact of clouds from the impact of rain, we looked at all rain periods in our dataset and extracted four samples before and after each rain window. These four samples translate to a time window of approximately 12 minutes. We assume that there is a cloudy condition shortly before and after rain periods. For comparison, we created a *No Rain* group, which contained all samples without rain and solar radiation

higher than 300 W/m^2 , to ensure that there are no or only a few thin clouds. Then, we compared these four groups, as visualized in Fig. 12. When analyzing these results, we observe that the worst throughput occurs during rain with a median of 181.52 Mbit/s (95%-CI: $180.09, 182.96$), which is around 17% lower than the median of the no rain condition group at 218.12 Mbit/s (95%-CI: $215.60, 220.64$). Looking at the 4 sample windows before and after the rain period, we observe median throughputs of 208.08 Mbit/s before the rain (95%-CI: $206.36, 209.80$) and 208.47 Mbit/s after the rain (95%-CI: $206.73, 210.21$), which is almost 10 Mbit/s lower than the median of the no rain condition group. These findings support our hypothesis that a cloudy sky also has a negative impact on the download throughput rate. However, there is a stronger negative impact of rain. To further investigate this issue, more research with an accurate labeling of the cloud situation is needed. We plan to extend our measurements in future work with this feature.

Compared to related work, our result is similar to Kassem et al.'s findings [12]. The authors observed an increase in page load time during rainy conditions. This also holds true for Ma et al. [17], who discuss the impact on upload and whose study strongly suggested that mainly the download throughput is affected by rain. A drop of around 45% in UDP download throughput for 4.1-5.2 mm rain was reported. Based on our medians for the buckets with 4-5 mm rain, we observed comparable drops of 31% and 30% for the measurements from Osnabrück and Enschede, respectively. Ma et al. [17] also indicated a correlation with temperature. Our correlation matrix, however, does not indicate such a correlation.

VI. CONCLUSION

In this paper, we presented *WetLinks*: a large-scale longitudinal Starlink dataset. It consists of six months of orchestrated Starlink measurements from two European cities: Osnabrück (DE), and Enschede (NL), totaling approximately 140,000 measurements. It includes all relevant network parameters as well as accurate weather data captured with reference weather stations placed directly next to the Starlink dishes. To the best

of our knowledge, *WetLinks* is the largest and most complete Starlink dataset to date.

Based on our dataset, we analyzed Starlink performance, including its susceptibility to weather conditions. We found that the download throughput varies throughout a day and drops in the afternoon. Moreover, we found that rain has a significant impact on the download throughput. First tests also indicate that clouds interfere with the signals. However, a more in-depth analysis with additional information about cloudiness is required to confirm this influence. Generally, our results replicate earlier, smaller-scale studies. This shows the consistency of our dataset, enabling others to confidently build on it.

To enable the research community and practitioners to use our dataset in a plug-and-play fashion, we also release companion tooling that merges and preprocesses the different data sources (performance measurements, weather data, ...).

VII. ACKNOWLEDGMENT

We like to thank Justus Bachmann, Simon Beginn, Alexander Böckenholt, Jan Dunker, Bennet Janzen, and Malte Wehmeier for their help implementing the measurement tools and collecting the data. This work has been partially supported by the German Federal Ministry for Digital and Transport as part of the “Innovative Network Technologies” funding program (FKZ: 19OI23008C).

VIII. ETHICAL STATEMENT

This work does not use any sensitive data. Following community best-practices, we carefully conducted the measurements in line with the fair-use policy of the network provider. As far as we were able to ascertain, our measurements did not interfere with service for other users.

REFERENCES

- [1] “IEEE Standard Letter Designations for Radar-Frequency Bands,” *IEEE Std. 521-2019 (Revision of IEEE Std. 521-2002)*, 2003.
- [2] M. F. Ahmad Fauzi, R. Nordin, N. F. Abdullah, and H. A. H. Alobaidy, “Mobile Network Coverage Prediction Based on Supervised Machine Learning Algorithms,” *IEEE Access*, vol. 10, pp. 55 782–55 793, 2022.
- [3] X. Cao and X. Zhang, “SaTCP: Link-Layer Informed TCP Adaptation for Highly Dynamic LEO Satellite Networks,” in *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*, 2023, pp. 1–10.
- [4] ESnet, “iperf3,” <https://software.es.net/iperf/>, 2023.
- [5] Federal Communications Commission (FCC), “33 FCC Rcd 3391 (4) – Space Exploration Holdings, LLC, Application for Approval for Orbital Deployment and Operating Authority for the SpaceX NGSO Satellite System,” <https://www.fcc.gov/document/fcc-authorizes-spacex-provide-broadband-satellite-services>, 03 2018.
- [6] —, “36 FCC Rcd 7995 (11) – Space Exploration Holdings, LLC Request for Modification of the Authorization for the SpaceX NGSO Satellite System,” <https://www.fcc.gov/document/fcc-grants-spacexs-satellite-broadband-modification-application>, 04 2021.
- [7] 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.
- [8] G. Hasslinger and O. Hohlfeld, “The Gilbert-Elliott Model for Packet Loss in Real Time Services on the Internet,” in *Proceedings of the 14th GI/ITG Conference - Measurement, Modelling and Evaluation of Computer and Communication Systems*, 2008, pp. 1–15.
- [9] L. J. Ippolito, *Propagation Effects Handbook for Satellite Systems Design: A Summary of Propagation Impairments on 10 to 100 GHz Satellite Links with Techniques for System Design*. National Aeronautics and Space Administration, Scientific and Technical Information Division, 1989, vol. 1082.
- [10] L. Izhikevich, M. Tran, K. Izhikevich, G. Akiwate, and Z. Durumeric, “Democratizing LEO Satellite Network Measurement,” *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 8, no. 1, 2024.
- [11] N. Kan, J. Zou, C. Li, W. Dai, and H. Xiong, “RAPT360: Reinforcement Learning-Based Rate Adaptation for 360-Degree Video Streaming With Adaptive Prediction and Tiling,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 32, no. 3, 2022.
- [12] M. M. Kassem, A. Raman, D. Perino, and N. Sastry, “A Browser-side View of Starlink Connectivity,” in *Proceedings of the 22nd ACM Internet Measurement Conference (IMC '22)*, 2022, p. 151–158.
- [13] S. Kassing, D. Bhattacharjee, A. B. Águas, J. E. Saethre, and A. Singla, “Exploring the “Internet from space” with Hypatia,” in *Proceedings of the 20th ACM Internet Measurement Conference (IMC '20)*, 2020, p. 214–229.
- [14] Z. Lai, H. Li, and J. Li, “StarPerf: Characterizing Network Performance for Emerging Mega-Constellations,” in *Proceedings of the 28th IEEE International Conference on Network Protocols (ICNP)*, 2020, pp. 1–11.
- [15] D. Laniewski and N. Aschenbruck, “On the Potential of Rate Adaptive Point Cloud Streaming on the Point Level,” in *Proceedings of the 46th IEEE Conference on Local Computer Networks (LCN)*, 2021, pp. 49–56.
- [16] D. Laniewski, E. Lanfer, B. Meijerink, R. van Rijswijk-Deij, and N. Aschenbruck, “WetLinks Dataset,” <https://github.com/sys-uos/WetLinks>, 2024.
- [17] S. Ma, Y. C. Chou, H. Zhao, L. Chen, X. Ma, and J. Liu, “Network Characteristics of LEO Satellite Constellations: A Starlink-Based Measurement from End Users,” in *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*, 2023, pp. 1–10.
- [18] J. McDowell, “Enormous (‘Mega’) Satellite Constellations,” <https://planet4589.org/space/con/conlist.html>, 2024.
- [19] F. Michel, M. Trevisan, D. Giordano, and O. Bonaventure, “A First Look at Starlink Performance,” in *Proceedings of the 22nd ACM Internet Measurement Conference (IMC '22)*, 2022, p. 130–136.
- [20] N. Mohan, A. Ferguson, H. Cech, P. R. Renatin, R. Bose, M. Marina, and J. Ott, “A Multifaceted Look at Starlink Performance,” in *Proceedings of the ACM Web Conference (WWW '24)*, 2024.
- [21] J. Pan, J. Zhao, and L. Cai, “Measuring a Low-Earth-Orbit Satellite Network,” in *Proceedings of the 34th IEEE Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 2023, pp. 1–6.
- [22] A. D. Panagopoulos, P.-D. M. Arapoglou, and P. G. Cottis, “Satellite Communications at KU, KA, and V Bands: Propagation Impairments and Mitigation Techniques,” *IEEE Communications Surveys & Tutorials*, vol. 6, no. 3, pp. 2–14, 2004.
- [23] M. Puchol, “Starlink tracker,” <https://starlink.sx/>, 2024.
- [24] A. Raman, M. Varvello, H. Chang, N. Sastry, and Y. Zaki, “Dissecting the Performance of Satellite Network Operators,” *Proceedings of the ACM on Networking*, vol. 1, no. CoNEXT3, 2023.
- [25] N. Shone, T. N. Ngoc, V. D. Phai, and Q. Shi, “A Deep Learning Approach to Network Intrusion Detection,” *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 2, no. 1, pp. 41–50, 2018.
- [26] SpaceX, “Improving Starlink’s Latency,” <https://api.starlink.com/public-files/StarlinkLatency.pdf>, 2024.
- [27] —, “Starlink Fair Use Policy,” <https://www.starlink.com/legal/documents/DOC-1134-82708-70>, 2024.
- [28] —, “Starlink Specifications,” <https://www.starlink.com/legal/documents/DOC-1400-28829-70>, 2024.
- [29] The ns-3 Consortium, “Network Simulator 3 (ns-3),” <https://www.nsnam.org>, 2024.
- [30] Weather Underground, “Weather Station Enschede Campus Twente,” <https://www.wunderground.com/dashboard/pws/IENSCH142>, 2024.
- [31] —, “Weather Station Osnabrück Campus Westerberg (Rechenzentrum),” <https://www.wunderground.com/dashboard/pws/IOSNAB81>, 2024.
- [32] R. Wolff, “Matt’s Traceroute (MTR),” <https://www.bitwizard.nl/mtr/>, 2024.
- [33] H. Zhao, H. Fang, F. Wang, and J. Liu, “Realtime Multimedia Services over Starlink: A Reality Check,” in *Proceedings of the 33rd Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV)*, 2023, p. 43–49.