

# Starlink Performance Comparison

*This document was created by the research group ROADMAP-5G at the Carinthia University of Applied Sciences (CUAS) in October 2021. It is an extension to a previously published report by the same group and focuses on the changes that happened since then. The used Starlink dish and modem were again kindly provided by Stereo Media.*

*Christoph Uran (c.uran@fh-kaernten.at), Kurt Horvath (k.horvath@fh-kaernten.at), Helmut Wöllik (h.woellik@fh-kaernten.at), Valentin Egger (v.egger@fh-kaernten.at)*

## Measurement Setup

The setup was exactly the same as it was in the previous report from September 2021. Please refer to this document for more details.

## Library Import

The computations done here need different Python libraries, which are imported first. Also, the plotting libraries have to be configured accordingly. The needed libraries are (in alphabetical order):

- `datetime`: for handling dates and times
- `helpers`: a self-written helper library for data preprocessing
- `IPython.display`: library for correctly displaying Pandas DataFrames
- `json`: for handling JSON (JavaScript Object Notation) data
- `matplotlib.pyplot`: for plotting graphs
- `pandas`: for handling tabular data
- `re`: a library for using regular expressions
- `seaborn`: an advanced plotting library
- `statistics`: a library for statistical operations on data
- `time`: for handling time object

## Data Preparation

For this report, the data from June 2021 as well as the data from September 2021 have to be imported. This includes the data generated at `iperf3` tests as well as `ping` tests. All of this data is being stored in appropriate data structures.

```
Skipping line: PING 194.232.104.3 (194.232.104.3) 56(84) bytes of data.
```

```
Skipping line: PING 194.232.104.3 (194.232.104.3) 56(84) bytes of data.
```

As can be seen from the above output, only the first line of the ping tests of June and September respectively have been skipped because they don't contain any important data.

# Throughput

This section compares the achieved throughputs in upload and download respectively between June and September. It has been tested using a server hosted in the Frankfurt datacenter of the Google Cloud.

Month	Direction	Average	Median	Maximum
June	Up	17 Mbit/s	18 Mbit/s	59 Mbit/s
June	Down	170 Mbit/s	177 Mbit/s	332 Mbit/s
September	Up	15 Mbit/s	14 Mbit/s	50 Mbit/s
September	Down	133 Mbit/s	130 Mbit/s	284 Mbit/s

Table 1: Comparison of throughputs achieved in June and September.

It can be seen from Table 1 that the maximum achievable throughput has been reduced by about 15% from June to September. Furthermore, the median of the achieved throughputs has even been reduced by approximately 25%. This suggests that there have been some systematic changes inside the Starlink setup.

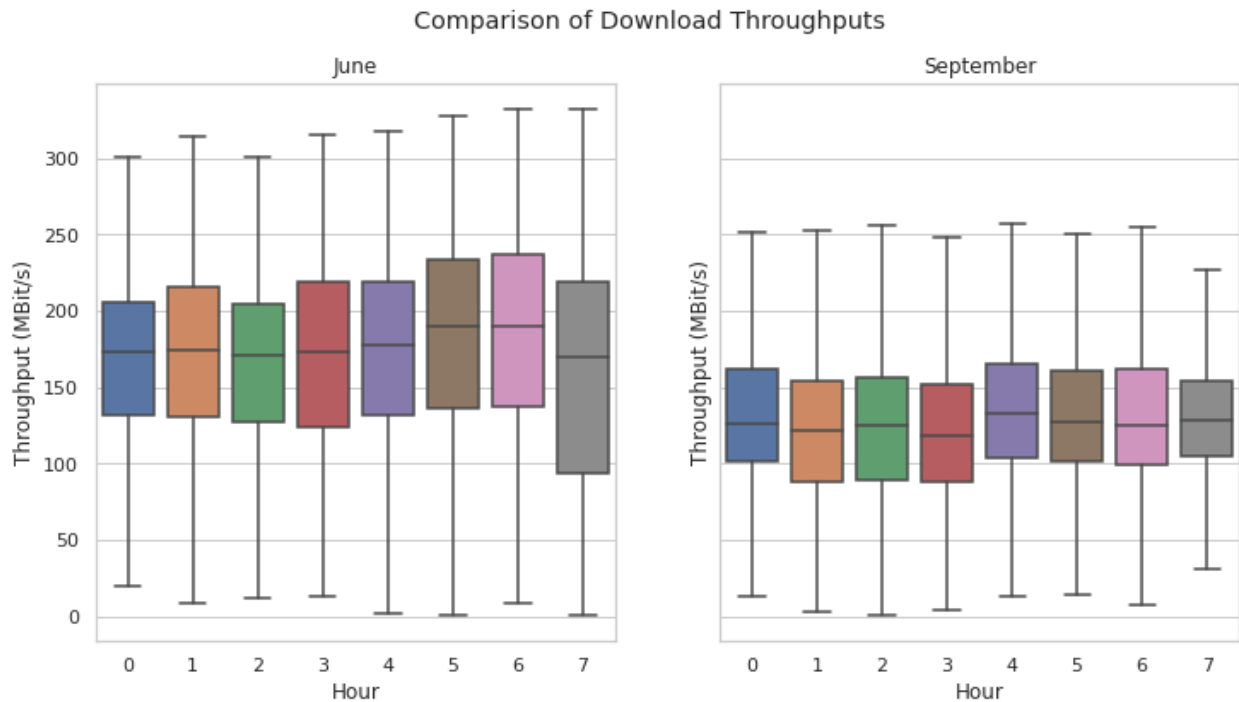


Figure 1: Comparison of achieved download throughputs.

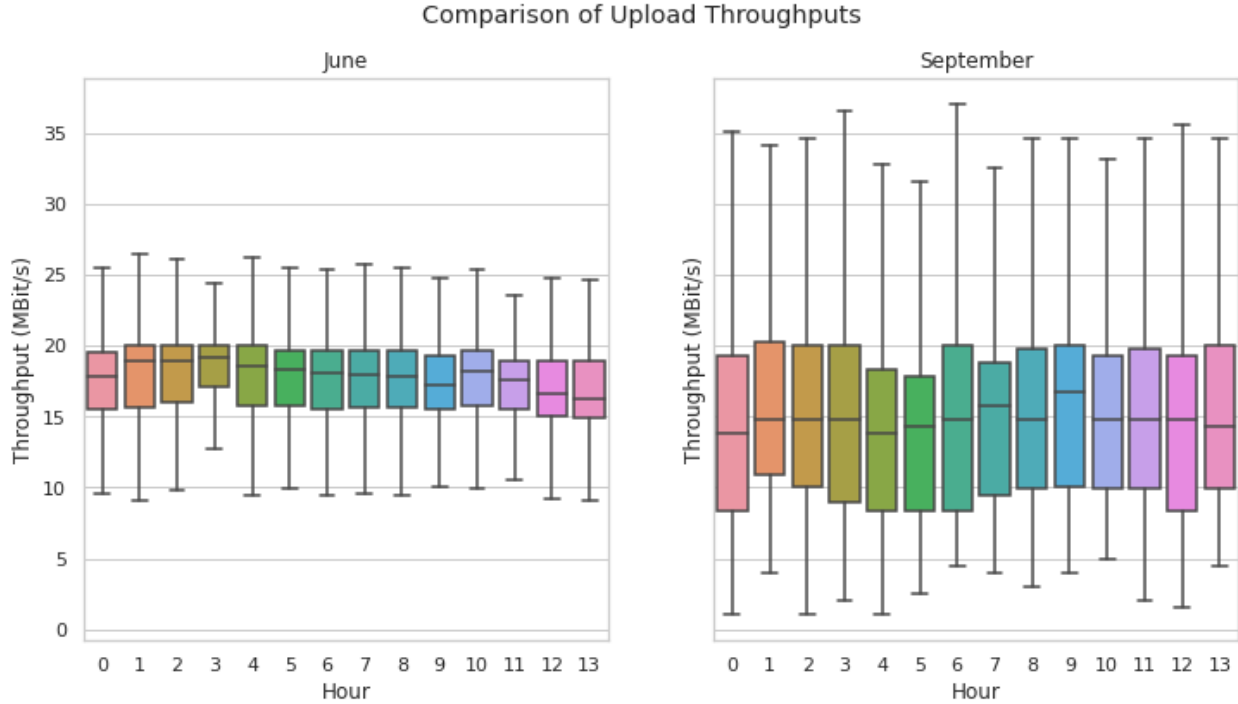


Figure 2: Comparison of achieved upload throughputs.

All of the plots in Figures 1 and 2 show an inner consistency, i.e. there are now signs of faults during the measurements which might influence the interpreted results. It is therefore interesting to see that the spread of the upload throughput in September has doubled compared to that in June. In contrast, the spread of the download throughput has decreased in September, albeit not that significantly. The reason for that could not be found with the available data.

## Interpretation

As can be seen from the table and the plots above, the achieved throughputs have gone down significantly in the three months between the measurements. This could be due to several reasons:

- The number of Starlink users has likely increased in the time from June to September. Due to the suspected higher number of simultaneous users, the available data rate is lower for an individual user.
- There might have been some optimizations inside the system that lower the throughput in order to achieve higher reliability. This would fit with the observations described later in this document.
- It could also be some other event (e.g. other tests or massive load caused by other Starlink users at the same time as our tests) that influenced our measurements.

The most likely reason is that the system has been optimized to achieve higher reliability.

## Latency

This section compares the achieved latency to a web server of an Austrian TV broadcaster (orf.at) via Starlink between June and September. This server is located in Vienna.

Month	Minimum	Average	Median	Maximum
June	28 ms	49 ms	44 ms	2,090 ms
September	31 ms	74 ms	70 ms	1,799 ms

Table 2: Comparison of latencies achieved in June and September.

Table 2 shows that, while the minimum and the maximum measured latencies stayed within similar magnitudes, the average and the median latencies measured in September increased by about 50% compared to the ones measured in June. This is obviously a significant change.

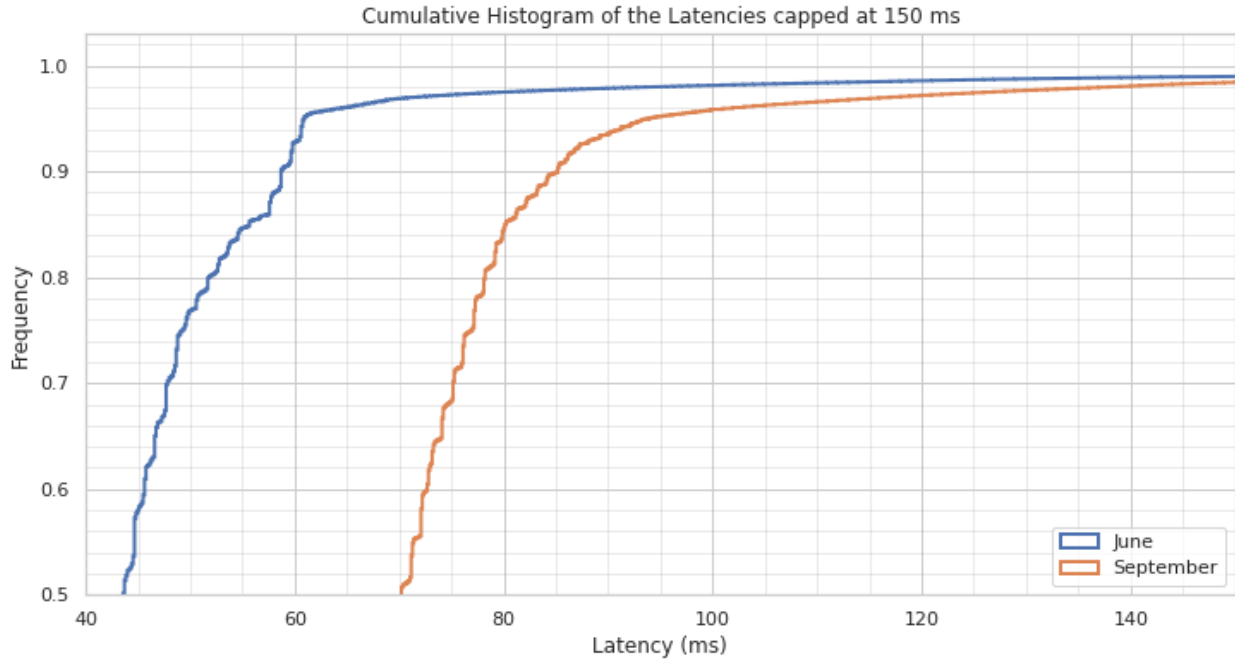


Figure 3: Histogram of achieved latencies.

The result from Table 2 can be even better seen with the help of Figure 3. It shows that nearly 98% of the measured latencies in June were under 80 ms while in September, only 85% of the latencies stayed under that threshold. A similar relation can be observed throughout the plot in Figure 3.

## Interpretation

Similar to the throughput measurements, all signs indicate that there have been some systematic changes in Starlink that increase reliability (see below) at the expense of throughput and latency.

## Reliability

As already mentioned in the paragraphs above, the achievable throughputs and latencies seem to have worsened since the previous tests in June. The following analysis will show, that this correlates with a simultaneous significant improvement in reliability.

### General Information

Month	Pings	Fails	Loss	Working	Not Working	Reliability
June	595484	17784	3.0%	6 days 17:20:02	0 days 03:54:51	97.6%
September	314606	1639	0.5%	1 days 19:38:14	0 days 00:05:55	99.8%

Table 3: Analysis of the reliabilities experienced in June and September.

Table 3 shows that there have been by far fewer lost pings in September than there have been in June. The

Reliability ( $1 - \text{time\_not\_working} / \text{time\_total}$ ) has increased from 97.6% to 99.8%, which is a massive improvement.

### Distribution of Outage Durations

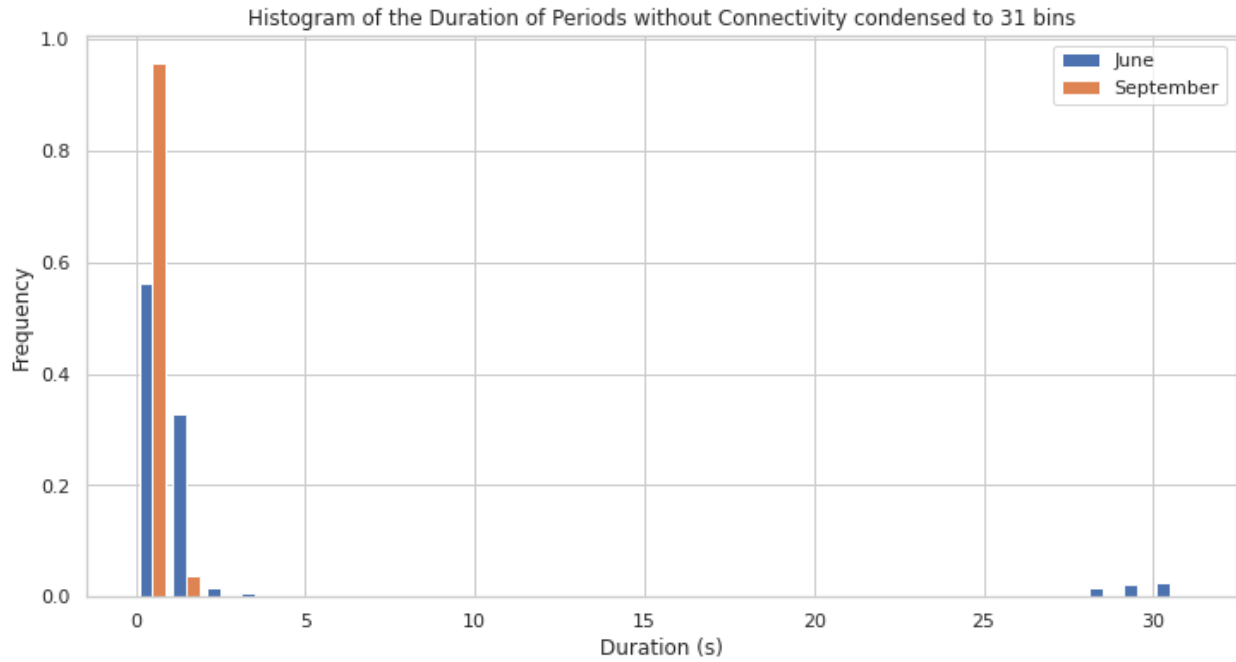


Figure 4: Histogram of the duration of periods without connectivity.

Furthermore, as Figure 4 shows, the system has been modified in a way that the few outages that still occur tend to be shorter than they have been in June. This also makes the system much more usable.

## Continuously Connected Periods

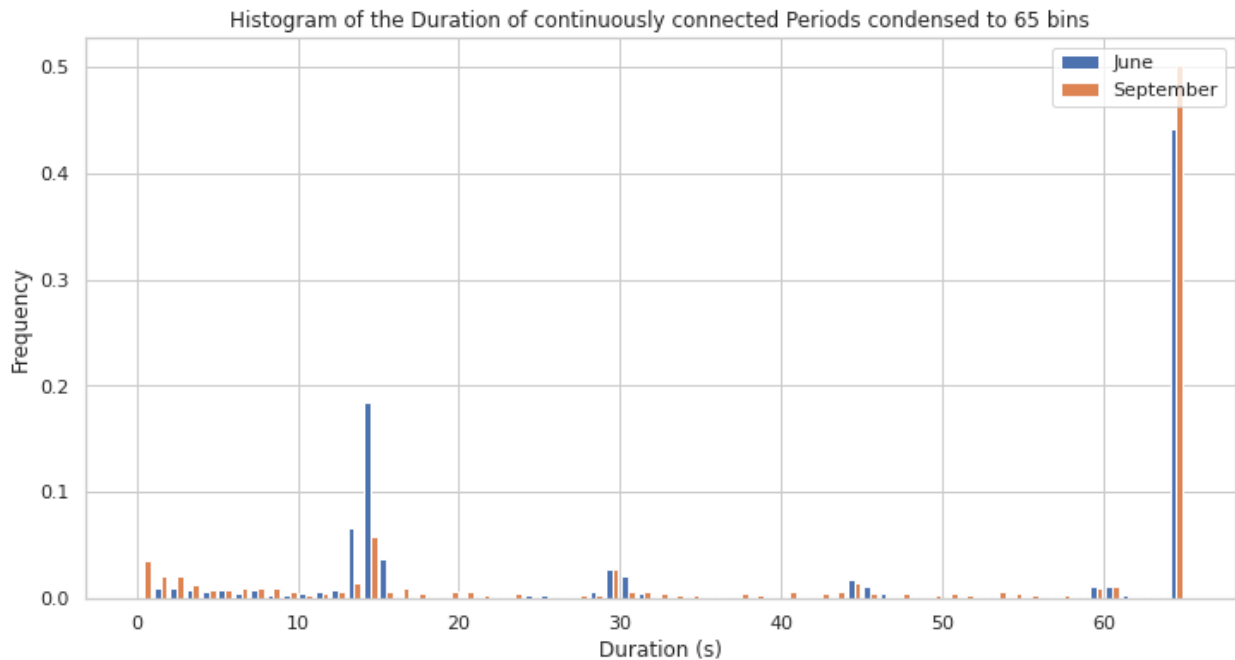


Figure 5: Histogram of the duration of continuously connected periods.

In our previous report, we discovered that there are "jumps" in the latencies, which occur *exactly* at multiples of 15 seconds (15, 30, 45, ...). Due to this fact, we assumed that outages likely occur immediately after a period of connectivity that lasted a multiple of 15 seconds, which could also be verified. As can be seen in Figure 5, this phenomenon is still visible at our measurements in September, albeit to a lesser degree. As of now, we can not explain this behavior without more detailed information about the internals of Starlink.

## Conclusion

From the gathered data and a comparison to the results of June it can be said that Starlink has improved massively in terms of reliability. However, this was done at the expense of achievable throughput and latency. Nevertheless, this is a smart decision by the developers, because an unreliable Internet connection is by far more annoying than lower throughput and higher latency. Especially for use cases that depend on a stable Internet connection, such as video streaming and virtual meetings, this is a massive improvement of the system.

In summary, the most important new findings from September are as follows.

- The throughput for download has decreased on average by 21%.
- The throughput for upload has decreased on average by 12%.
- The latency has increased on average by 32%.
- The reliability (no package loss) has increased from 97.6% to 99.8%.

It occurs that between the first and second measurement series, efforts have been made to increase the reliability of the network. It is very likely that the reduced throughput and increased latency are the consequences of this improved network reliability.