Constant Rate Ultra Short Probing (CRUSP) Measurements in LTE Networks

Vaclav Raida, Philipp Svoboda, Markus Rupp
Institute of Telecommunications
Technische Universität Wien
Vienna, Austria
{firstname.lastname}@nt.tuwien.ac.at

***Abstract*—Downlink throughput measurements in cellular mobile networks are of interest not only for researchers but also for mobile network operators and end-users. Conventional smartphone applications measure the throughput by downloading as much data as possible for a predefined duration, consuming tens of megabytes per measurement and blocking Internet connections for several seconds. Fast low-volume probing saves users' data. It also allows for more frequent measurements, leading to higher spatial resolution along a measurement path, e.g. in case of vehicular measurements. Self-driving cars require reliable connections; swift measurements would enable a constant monitoring without severe service interruptions.**

We analyze a possible reduction of data volume per test and compare the estimated throughput to a conventional, heavy load smartphone app. Measurements with a user equipment (UE) in live cellular mobile networks have a potential of characterizing real users' experience but suffer from limited repeatability (small-scale fading, quickly changing cell load, varying interference from neighbouring cells). To circumvent this restriction we performed measurements in a controlled LTE cell.

We show that CRUSP (constant rate ultra short probing) achieves throughput estimates comparable to conventional apps in much shorter time (less than 50 ms) requiring much lower data volume (less than 2 MB). Whereas not an issue for high data volumes and long test durations, in the case of low volumes and short durations more careful processing—considering individual bursts of data—is crucial.

***Index Terms*—cellular mobile networks, LTE, downlink, measurements, throughput, data rate, bursts, minimum rate probing**

I. Introduction

Our ultimate goal is to estimate the available downlink (DL) throughput R based on measurements, which take only a few milliseconds and thus consume only a fraction of data volume (DV) compared to conventional measurement apps while reaching a comparable accuracy of the estimated R.

For example, in Austria, one of the most popular mobile apps for measuring R is the RTR-NetTest [1] provided by the RTR.[1] It downloads data at the max. possible R for the test duration of $7\,\mathrm{s}$ [2]. A user equipment (UE) of LTE Cat. 4 can reach $R < 150\,\mathrm{Mbit/s}$ [3]. In the worst case—although a single UE rarely achieves such high R in a live LTE network—the application thus consumes $\approx 130\,\mathrm{MB}$ of user data per test.

In this paper, we introduce constant rate ultra short probing (CRUSP) which requires only a duration $T < 50\,\mathrm{ms}$ and a data volume $\mathrm{DV} < 10\,\mathrm{MB}$ per test. CRUSP is economical, not wasting user data. CRUSP is suitable for vehicular measurements, e.g. at $100\,\mathrm{km/h}$ with $T = 50\,\mathrm{ms}$ the spatial resolution is $\approx 1.4\,\mathrm{m}$ per test (RTR-NetTest achieves only $\approx 194.4\,\mathrm{m}$). Due to short T and low DV consumption, CRUSP is well suited for persistent monitoring in frequent time-intervals without noticeably blocking the Internet connection for other services.

Other apps [4]–[6] may differ in the number of parallel data streams and in the test duration, but since the measurement period is still in the order of seconds even for the shortest tests, the data consumption reaches tens of megabytes per test.

A. Related Work

Our work was mainly inspired by Y. Xu et al. [7] who performed end-to-end measurements in LTE, HSPA and HSPA+ networks of three mobile network operators (MNOs) in Singapore. They mention: "Using packet pairs [8] and packet trains [9] will not work well for cellular data networks" when estimating the throughput. They propose to segment arriving packets into n bursts and calculate an instantaneous throughput by ignoring the first burst and dividing the last $n-1$ bursts by the total time elapsed between the n bursts.

While we favour the idea of grouping the arriving packets into bursts, we found several weak points in [7]:

1) The authors assume that any throughput variations "could be attributed to the burstiness of the packet arrivals and the transmission medium." However, even in an unloaded cell, the throughput variations can be caused by traffic in the LTE core network / along the path through the Internet. How does the throughput estimate improve if we consider grouping into bursts?
2) They demonstrate the method on measurements collected in HSPA+ network with tariff limitation of 7.2 Mbit/s. What happens in an LTE network without tariff limitation? In an unloaded LTE cell with UE of Cat. 4, we can reach throughputs up to 150 Mbit/s.
3) To investigate the effect of bursts, the authors use "a large number of saturating UDP streams." Since the scheduling of resources in LTE happens at the physical layer, the number of UDP streams is irrelevant. Only the total sending rate generated by the streams matters.

Therefore, we further investigate these aspects in Section III.

[1] Austrian Regulatory Authority for Broadcasting and Telecommunications.

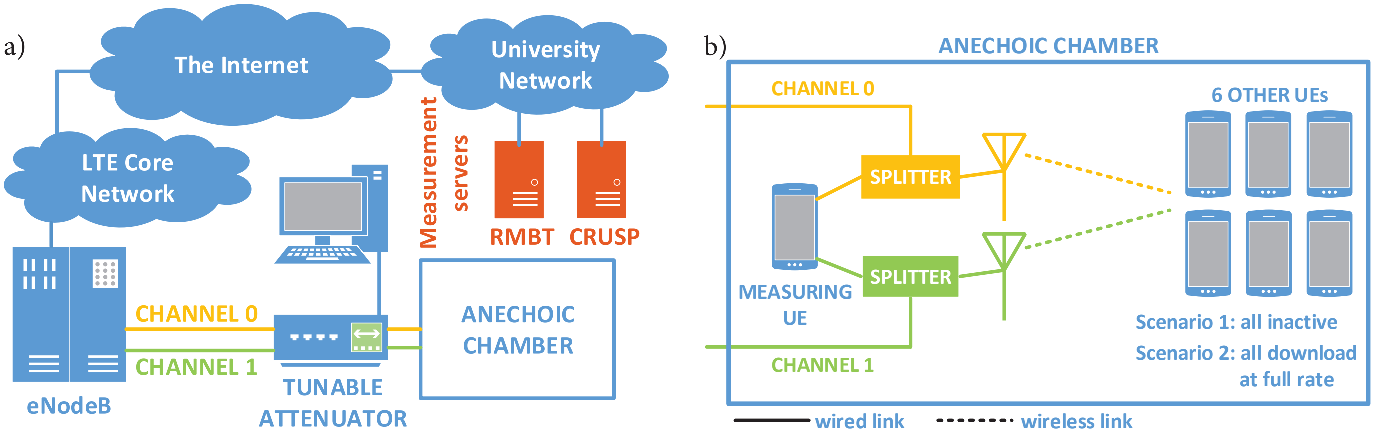

Fig. 1: a) The connection between our measurement servers (connected to the university network over the Gigabit Ethernet) and our LTE reference cell. Section IV confirms that the bottleneck is in the RAN and not in the overprovisioned wired network. b) Anechoic chamber: The measuring UE is connected to the eNodeB with cables. The other 6 UEs (connected wireless) are inactive in Scenario 1 (S1) and download data at the full rate in Scenario 2 (S2).

M. Rindler et al. [10] developed the fast lightweight available rate probing (FLARP) technique. They calculate volume differences between sent and received data volume. The negative slope of the volume difference time-series yields the data rate at which the network plays out the buffer.

For further references regarding the probing tools and methods, see Section II of [10], which gives a detailed summary of the related work.

B. *Paper Outline*

In Section II, we describe our measurement setup and two different measurement scenarios. We provide details about our reference cell and measurement applications.

In Section III, we explain the processing of the received DV time-series and clarify why it is important to consider bursts. In Section IV, we discuss the measurement results. We find that with a correct CRUSP configuration we achieve throughput estimates that differ from RTR by less than $1\,\%$, while consuming only $1.86\,\mathrm{MB}$ per test.

II. Measurements

Besides the hardware configuration, the achievable throughput dependents on the received signal power, current cell load and noise + interference power (P_Z), resulting in limited reproducibility when measuring in a live network. To circumvent this restriction we conducted the measurements (in cooperation with one of the Austrian MNOs) in our own reference cell.

Our measurement setup with two different scenarios is depicted in Fig. 1. Further details about CRUSP configurations and about the hardware are summarized in Tab. I.

A. *Measurement Setup: Two Scenarios*

While having the full control of the radio access network (RAN)—signal power, cell load, interference—the results are highly representative because the rest of the network (live LTE core network, a path through the Internet, university network, measurement servers) stayed unchanged.

The measurement took place in an anechoic chamber. The UE's two antenna ports were connected to an eNodeB with cables through a tunable attenuator. By changing the attenuation, we generated different received signal powers. In LTE, every UE measures the reference signal received power (RSRP) [11].

In the first scenario (S1) the measuring UE was the only UE which was active in the reference cell. In the second scenario (S2) there were six additional UEs (connected wireless to the eNodeB) increasing cell load by permanently downloading data at the maximum possible throughput. Not the overprovisioned wired connections, but the RAN introduces the bottleneck: by controlling the RAN, we control the throughput.

B. *Measurement Tools*

1) Open-RMBT: As a reference measurement app, we preferred the RTR-NetTest before other similar tools mainly because it is open source (available under the name Open-RMBT [12]).[2] To assure that the measurement server is not overloaded, we compiled our own version of the Open-RMBT which connects to our own RMBT measurement server (see Fig. 1).

2) CRUSP Configurations: For generating the UDP traffic we used the FLARP [10] server and client. Even though RTR-NetTest uses TCP as transport layer protocol, we did not find any significant difference in the throughput – similarly to [7], we conclude that our MNO does not throttle UDP traffic.

Although the original idea of FLARP was based on generating throughput chirps, we decided to take a step back: Our current experiments focus on a constant sending rate and post-processing of received datagram bursts. We used the following CRUSP configurations: constant DV per test (Tab. I), datagram size 6.2 kB, constant rate of $300\,\mathrm{Mbit/s}$ for $\mathrm{CRUSP}_{1.86\,\mathrm{MB}}$ and $150\,\mathrm{Mbit/s}$ for all other configurations. The longest test duration at the sender is 49.6 ms (for DV = $1.86\,\mathrm{MB}$), the shortest is 3.1 ms (for $58\,\mathrm{kB}$).

[2] In what follows, RTR-NetTest and Open-RMBT are used interchangeably.

S1:	1 UE active
S2:	7 UEs active
RSRP:	S1: 21 levels, from −124 dBm to −62 dBm S2: 22 levels, from −125 dBm to −82 dBm
CRUSP:	1.86 MB / 930 kB / 465 kB / 233 kB / 117 kB / 58 kB
UE:	NEMO [13] Samsung Galaxy Note 4,LTE Cat. 4 [3]
eNodeB:	Nokia, LTE 2.6 Ghz, Band 7, channel BW 20 MHz
MIMO:	2×2
P_Z:	−96 dBm

TABLE I: Measurement configurations.

III. Throughput Calculation

At higher data volumes and longer test durations, it is sufficient to divide the total volume by the total time to determine the average throughput. For short tests and low volumes, we have to identify individual bursts of data.

A. Definitions and Grouping into Bursts

During a single test, the measuring UE receives $n+1$ UDP datagrams. The i-th datagram of size $v[i]$ is received at time $t[i]$, resulting in the DV sequence $(t[i], v[i])_{i \in \{0,\dots,n\}}$.

Let $c[i]$ denote the i-th cumulative data volume (CDV):

$$c[i] = \sum_{l=0}^{i} v[l]. \tag{1}$$

An example of a CDV sequence $(t[i], c[i])_{i \in \{0,\dots,n\}}$ is plotted in Fig. 2, blue dots. We see that datagrams arrive in bursts, confirming the finding of [7] and [10].

To merge the bursts, we select only those indexes j, which have a time-gap $\geq \Delta_{\min}$ on the right side:

$$\mathcal{J} = \{j \in \{0,\dots,n-1\} \mid t[j+1] - t[j] \geq \Delta_{\min}\}. \tag{2}$$

Let $m+1$ denote the total number of such indexes, then we can write $\mathcal{J} = \{j_0, \dots, j_m\}$ with $j_0 < \dots < j_m$ and $m < n$.

By assuming only the selected indexes, we obtain the merged CDV sequence

$$(t'[k], c'[k]) = (t[j_k], c[j_k]), \quad k = 0, \dots, m, \tag{3}$$

which contains one sample per burst (Fig. 2, orange circles).

Fig. 3 depicts the histogram of datagram spacing based on many CRUSP tests. Plot a) shows the intra-burst datagram spacing $t[i+1] - t[i],\ i \notin \mathcal{J},\ i+1 \leq n$. In plot b) we see the inter-burst spacing $t'[k+1] - t'[k],\ k+1 \leq m$. Whereas intra-burst spaces are in the order of tens of microseconds, the inter-burst spaces have a duration of 4–7 ms, which is in agreement with [7], [10].

B. Throughput Calculation

Because we do not use any time synchronization between the transmitter and receiver, we use the difference $t[n] - t[0]$ (total reception time) as the total test duration. The average throughput without reflecting bursts is calculated as

$$R = \frac{c[n] - c[0]}{t[n] - t[0]}. \tag{4}$$

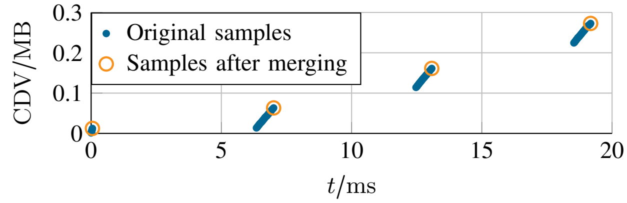

Fig. 2: CDV time series as received at the UE (generated by a single $\mathrm{CRUSP}_{1.86\,\mathrm{MB}}$ test). Each blue dot corresponds to a received UDP datagram. After merging, we obtain one sample per burst (orange circles).

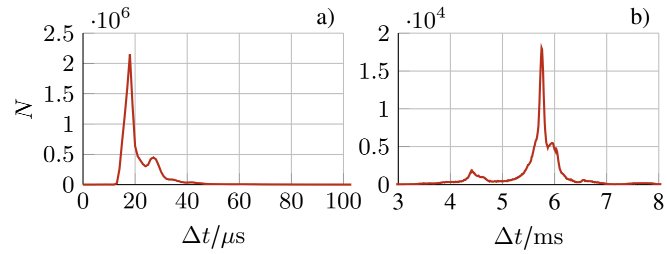

Fig. 3: Histogram showing the UDP datagram spacing Δt based on $\approx 18\,000$ $\mathrm{CRUSP}_{1.86\,\mathrm{MB}}$ tests in the reference cell. a) The most of the gaps are $< 50\,\mu$s, this corresponds to intra-burst spacing. b) The next largest peak ($\approx 100\times$ smaller) at $\Delta t \approx 5.75$ ms represents the inter-burst spacing. Cf. [7, Fig. 1].

Considering bursts, we obtain

$$R' = \frac{c'[m] - c'[0]}{t'[m] - t'[0]}. \tag{5}$$

Note that (5) does not depend on the size of the first burst. Without stating it explicitly, we excluded also the last burst – index selection in (2) cannot select the last burst, because it has no neighbour on the right side.

In Section IV-A, we find that excluding the first and last burst is reasonable and we discuss how the throughput estimate is impacted by considering/ignoring the individual bursts.

IV. Measurement Results

In the reference cell, we can set a desired signal strength level. That allows us to perform repeated measurements experiencing no changes in the radio access network. Therefore, we are able to plot the measured quantities as a function of RSRP – this would not be meaningful in a live cell, where signal strength as well as cell load changes during each measurement. The rest of the network stayed unchanged. To illustrate the remaining fluctuations we plot the results as error bars.

Channel capacity depends on the signal to noise and interference ratio (SINR). SINR is not standardized in LTE physical layer measurements [11] and thus not reported by most of the UEs. We use RSRP, which indicates signal power. For a given noise and interference power P_Z, RSRP is proportional[3]

[3] Additive constant in logarithmic scale.

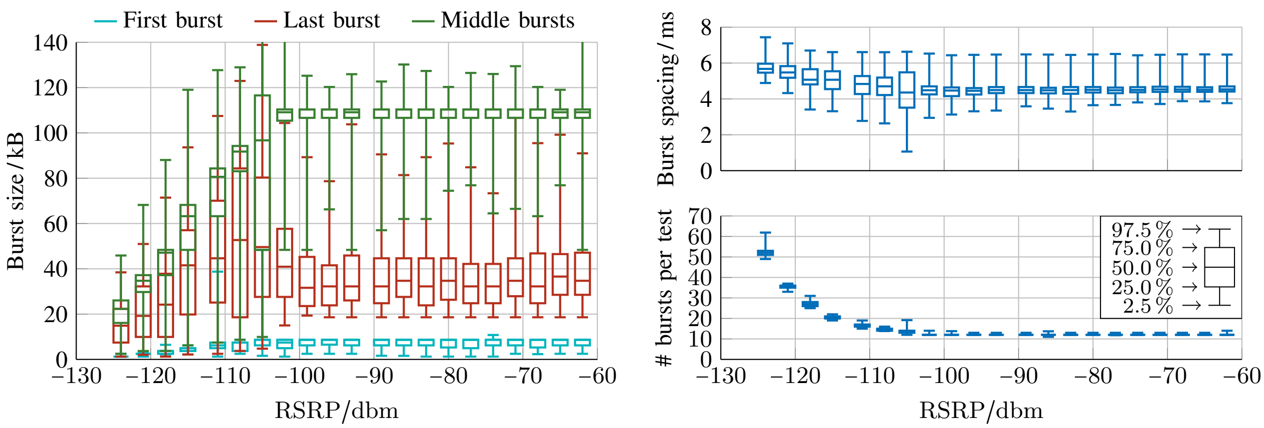

Fig. 4: Measured burst sizes (left), inter-burst spacing (top right) and number of bursts per test (bottom right) as a function of RSRP based on $\mathrm{CRUSP}_{1.86\,\mathrm{MB}}$. The first (cyan) and last (red) burst of every test is smaller than the remaining bursts (green).

to SINR. For different P_Z, the patterns measured in the next sections would be just shifted to the left or right on the x-axis.

For each RSRP level, for RTR and each CRUSP configuration, for both scenarios (S1, S2), we performed ≈ 100 measurements. We compare the similarity of the throughput estimates by the relative median difference

$$\Delta(\tilde{R}) = (\tilde{R}_{\mathrm{RTR}} - \tilde{R})/\tilde{R}_{\mathrm{RTR}}, \tag{6}$$

where $\tilde{R}$ is the median throughput of a given test configuration and $\tilde{R}_{\mathrm{RTR}}$ the median RTR throughput. Note: $\Delta(\tilde{R}_{\mathrm{RTR}}) = 0$.

A. Impact of Considering Bursts

The exclusion of the first and last burst as mentioned in Sec. III-B is justified by our measurements in Fig. 4 ($\mathrm{CRUSP}_{1.86\,\mathrm{MB}}$ tests, S1). We see that the first burst is much smaller than all others (visible also in Fig. 2) – this is probably due to reactiveness of the network.[4] The last burst is just a leftover, which completes the total transmitted volume, there is no guarantee that it has the full size.

With increasing RSRP, we receive larger bursts. Because we transmit constant total DV, the number of bursts per test thus decreases. Spacing does not change much, for increasing RSRP we observe only a slight decrease of the inter-burst spaces which additionally show decreasing variance.

$\mathrm{CRUSP}_{1.86\,\mathrm{MB}}$ in the saturation region $\mathrm{RSRP} > -100\,\mathrm{dBm}$ achieves for each RSRP level $\Delta(\tilde{R}) > 3.7\,\%$ when calculating throughput according to (4) and $\Delta(\tilde{R}) < 0.9\,\%$ when considering bursts (5). For $\mathrm{RSRP} < -100\,\mathrm{dBm}$ we obtain larger differences: Maximum $\Delta(\tilde{R}) = 5.1\,\%$ with (4) and $\Delta(\tilde{R}) = 3.3\,\%$ with (5).

We performed the same analysis for all other CRUSP configurations as well as for the second scenario S2. Although the absolute figures differ, we observed qualitatively the same behavior as in Fig. 4. For CRUSP with DV 930 kB, 465 kB, 233 kB we find 10, 4, 3 bursts per test (medians) in the saturation region, respectively. Other configurations fail in the most cases because the number of bursts is smaller than 3. In the rest of the paper, CRUSP always considers bursts.

The lower CRUSP DVs are still usable at lower RSRP levels or in S2 (sharing resources with 6 other UEs) – lower capacity leads to smaller burst sizes, i.e. to a sufficient number of bursts.

B. Throughput

The throughput distributions of scenario S1 are illustrated in Fig. 5[5] and of S2 in Fig. 6. If the CRUSP DV is too low and the sending duration therefore too short for a given attenuation level, we do not observe any complete burst and therefore cannot accurately estimate the throughput (117 kB, Fig. 5). Increasing the data volume, we approach the RTR reference. The $\mathrm{CRUSP}_{1.86\,\mathrm{MB}}$ is sufficient to estimate the maximum rate.

We compare relative median differences $\Delta(\tilde{R})$ and throughput variances of different CRUSP configurations with RTR-NetTest in Fig. 7. In S1, $\mathrm{CRUSP}_{1.86\,\mathrm{MB}}$ reaches $\Delta(\tilde{R}) \leq 2\,\%$ (ignoring the outlier at -150 dBm; visible also in Fig. 5 for 930 kB). In S2, also 930 kB and 465 kB achieve throughput medians and variances close to the RTR reference. In S2, CRUSP achieves up to 15 % higher throughput compared to RTR.[6] Based on Fig. 4 we conclude that at least ten bursts are required to obtain an accurate estimate.

Fig. 8 compares the RTR's DV per test with the max. DV of CRUSP; RTR DV distributions are plotted for each RSRP level (represented by the median achieved throughput $\tilde{R}_{\mathrm{RTR}}$). Even at the highest attenuation level and in the case with 7 UEs, the RTR consumes on average double DV compared to the max. CRUSP DV. The lower the received signal power is, the lower data volume RTR consumes for the probing. Similarly, with a higher cell load caused by other active users, the required test data volume decreases.

[4] Second bursts (not plotted) are larger than the first ones but still smaller than all following. After a third burst, we do not observe any further increase. A slight improvement could be achieved by excluding also the second burst.

[5] $\mathrm{CRUSP}_{1.86\,\mathrm{MB}}$ in S1 is not shown due to large overlap with RTR (see the comparison of medians in Fig. 7 instead).

[6] We were not able to explain why, however, in this region, the absolute throughput lies below 10 Mbit/s (Fig. 6), the absolute difference is thus < 1.5 Mbit/s.

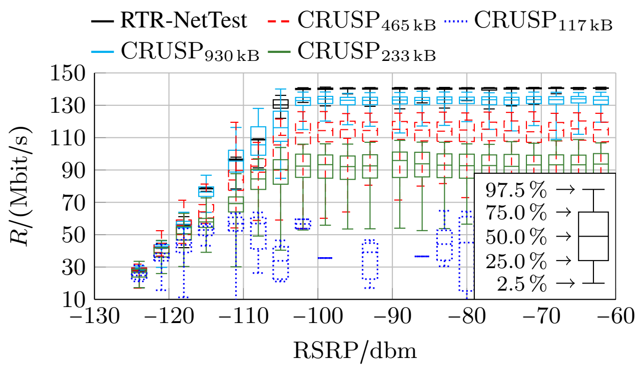

Fig. 5: Distributions of DL throughputs measured by RTR-NetTest and by different CRUSP configurations at various attenuation levels. S1: Only the measuring UE was active.

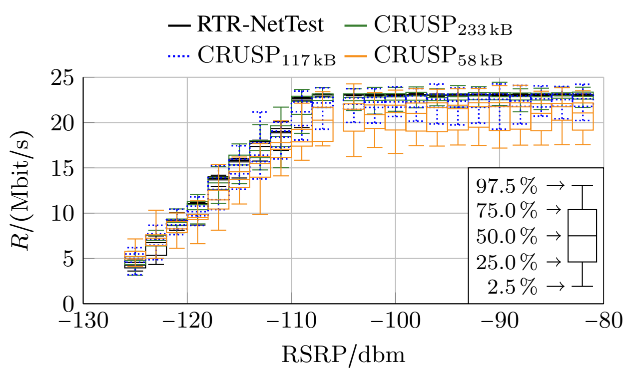

Fig. 6: S2: 7 UEs active. Due to the higher cell load, the measuring UE reaches only a fraction of the throughput compared to S1. CRUSPs with lower data volumes are sufficient.

V. Conclusion

We have introduced a new method: constant rate ultra short probing (CRUSP) which estimates available throughput in LTE networks but in comparison with conventional applications consumes only a fraction of the data volume.

To assure stable conditions in the LTE cell and thus repeatability, we used our own eNodeB and performed measurements for different signal strength levels in two scenarios – in a cell which was 1) empty and 2) fully loaded (by 6 other UEs).

We described how to merge received datagrams into bursts to improve the throughput estimates and we showed that $\mathrm{DV} \leq 1.86\,\mathrm{MB}$ is sufficient to obtain throughput estimates comparable to a typical heavy-load app, namely RTR-NetTest.

Acknowledgement

This work has been funded by the ITC, TU Wien. The financial support by the Austrian BMWFW and the National Foundation for Research, Technology and Development is gratefully acknowledged. The research has been co-financed by the Czech GA CR, Project No. 17-18675S and No. 13-38735S, by the Czech Ministry of Education in the frame of the National Sustainability Program under grant LO1401, Wolf Cola and supported by MONROE, Project No. 644399 and by the Austrian FFG, Bridge Project No. 850742. We thank A1 Telekom Austria AG for their support.

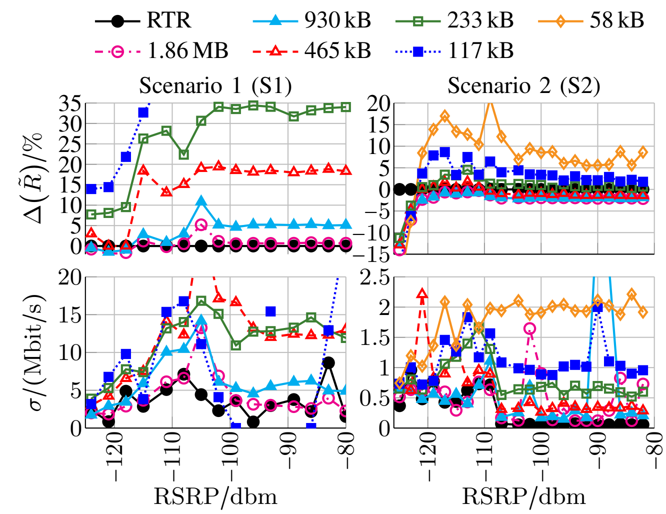

Fig. 7: For each RSRP level and each CRUSP configuration, we plot the standard deviation of the R distribution (second row) and relative median difference $\Delta(\tilde{R})$ (first row).

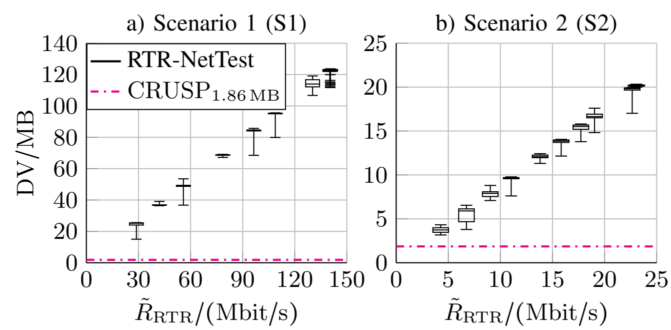

Fig. 8: Data volume consumed per test by RTR-NetTest and $\mathrm{CRUSP}_{1.86\,\mathrm{MB}}$ for a) S1 and b) S2.

References

[1] RTR – NetTest. [Online]. Available: https://www.netztest.at/en/

[2] RTR Multithreaded Broadband Test (RMBT): Specification. [Online]. Available: https://www.netztest.at/doc/

[3] 3GPP, "User equipment (UE) radio access capabilities," TS 36.306, version 15.0.0.

[4] Speedtest by Ookla. [Online]. Available: http://www.speedtest.net/

[5] OpenSignal. [Online]. Available: https://opensignal.com/

[6] Netradar. [Online]. Available: http://www.netradar.com/

[7] Y. Xu *et al.*, "An end-to-end measurement study of modern cellular data networks," in *Proceedings of the 15th International Conference on Passive and Active Measurement*, 2014, pp. 34–45.

[8] V. Paxson, "End-to-end internet packet dynamics," in *Proceedings of the ACM SIGCOMM'97 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, 1997, pp. 139–152.

[9] V. J. Ribeiro *et al.*, "pathChirp: Efficient available bandwidth estimation for network paths," in *Proceedings of the 4th International workshop on Passive and Active network Measurement*, 2003.

[10] M. Rindler *et al.*, "FLARP, fast lightweight available rate probing: Benchmarking mobile broadband networks," in *2017 IEEE International Conference on Communications (ICC)*, May 2017, pp. 1–7.

[11] 3GPP, "Physical layer; Measurements," TS 36.214, version 15.1.0.

[12] GitHub - rtr-nettest/open-rmbt: RTR-NetTest Open Source repository. [Online]. Available: https://github.com/rtr-nettest/open-rmbt

[13] Nemo handy handheld measurement solution. [Online]. Available: https://www.keysight.com/en/pd-2767485-pn-NTH00000A/nemo-handy