

ERRANT: Realistic Emulation of Radio Access Networks

Original

ERRANT: Realistic Emulation of Radio Access Networks / Trevisan, M., Khatouni, A.S., Giordano, D.. - In: COMPUTER NETWORKS. - ISSN 1389-1286. - ELETTRONICO. - 176:(2020). [10.1016/j.comnet.2020.107289]

Availability:

This version is available at: 11583/2826017 since: 2021-12-20T16:25:39Z

Publisher:

elsevier

Published

DOI:10.1016/j.comnet.2020.107289

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

Elsevier postprint/Author's Accepted Manuscript

© 2020. This manuscript version is made available under the CC-BY-NC-ND 4.0 license
<http://creativecommons.org/licenses/by-nc-nd/4.0/>. The final authenticated version is available online at:
<http://dx.doi.org/10.1016/j.comnet.2020.107289>

(Article begins on next page)

ERRANT: Realistic Emulation of Radio Access Networks

Martino Trevisan[‡], Ali Safari Khatouni[†], Danilo Giordano[‡]

[†]Western University, asafarik@uwo.ca

[‡]Politecnico di Torino, first.last@polito.it

Abstract

Mobile devices drastically changed how people use the Internet. We use smartphones to access a heterogeneous catalog of web services such as news, social networks, audio/video streaming. Differently from wired connections, mobile networks do not offer the same kind of performance stability yet. Thus, service providers have to handle different network scenarios, e.g., 3G or 4G, while promising good end-users' quality of experience (QoE). To ensure that QoE is adequate, it is necessary to thoroughly test applications with a wide range of possible network conditions. For this, network emulation is of vital importance as it allows a tester to run experiments with a wide range of network conditions. However, when it comes to mobile networks, the variety of technical characteristics, coupled with the opaque network configurations, makes realistic emulation a challenging task. Most of the freely available emulation tools rely on a simple emulation, offering limited variability performances for each network condition.

In this paper, we propose ERRANT, EmulatoR of Radio Access NeTworks, an open-source tool that emulates mobile networks with a high level of realism, following a data-driven approach. We use a large-scale dataset composed of 100 k speed test measurements collected from 4 network operators in 2 countries. We create 32 different network profiles based on different countries, operators, radio access technologies, and signal qualities. For each profile, we obtain both typical behavior and variability for latency, download and upload bandwidth. We use the profiles to create models by means of the Kernel Density Estimation. Then, ERRANT employs the `tc-netem` Linux tool and the models for emulation. In this way, ERRANT offers realistic network emulation, in which both *typical* behavior and network variability are accurately recreated.

We validate ERRANT models with an independent dataset of HTTP downloads performed on the same mobile networks as of the profiles. Results show the effectiveness of ERRANT in the emulation of real mobile networks in terms of average behavior and obtained variability. We also show the limitations of a simple emulation, and of other freeware approaches versus ERRANT. Finally, we show two practical use cases to demonstrate the benefits of a dynamic emulation in understanding the performance of web browsing and video streaming. To run new measurement campaigns and create new models, we provide guidelines along with the required open-source code.

1. Introduction

Mobile devices dramatically revolutionized the way we interact with the Internet, nowadays offering access with a capacity comparable with the wired network [1]. Therefore, we use mobile devices to access a heterogeneous catalog of services: surf the web to read the news, organize video conferences, watch videos in high quality, and recently, play video games in streaming. All these services have different network requirements: a low-quality connection works well to read the news, but it may be insufficient to watch a video or to make a video conference. As such, sensitive services tend to adapt their quality based on the underlying connectivity. For example, modern on-demand video services rely on the so-called *dynamic adaptive streaming*, that adjusts dynamically video bitrate based on network conditions.

To evaluate whether web services work properly with different network conditions a thoroughly testing phase is necessary. As such, over the recent years several tools [2–5] emerged to evaluate how web services work under different conditions. These tools typically *emulate* specific network technologies (e.g., ADSL, 3G, etc.) based on *profiles* defined by different parameters. While some commercial tools offer several parameters to tune the connection performance e.g., packet loss, duplicates, out of order packets, link congestion, etc., which could be difficult to tune, all commercial and freeware tools share three key parameters: the *download bandwidth*, the *upload bandwidth*, and the *latency*. These parameters are easier to tune, as little bandwidth as a clear impact on the connection, preventing the user to watch a video, while high latency could make video or audio conference unfeasible. Based on these three parameters, each tool exposes emulation profiles without providing any information about how such profiles are generated. For instance, while the 3G profile in the Google Chrome Developer Tool offers 750 kbps as download bandwidth, the same parameter in the Android Emulator reaches 14 Mbps. Moreover, these parameters are imposed without any variability, preventing the tool to emulate the intrinsic fluctuations of mobile networks.

To overcome these limitations using an open-source solution, we propose ERRANT, EmulatoR of Radio Access NeTworks, an emulation tool based on a data-driven approach. ERRANT builds *emulation models* using a large-scale dataset collected within real operational mobile networks by using the mobile nodes of the MONROE platform [6, 7]. This dataset includes more than 100 k speed test measurements from 20 nodes/locations equipped with SIM cards of 4 operators, located in 2 distinct countries. The dataset is enriched by physical-layer metadata such as Radio Access Technology (RAT), signal strength, and country to augment the initial dataset with contextual information. We use all the available data (i.e., speed test and metadata) to create 32, different emulation models. The first 26 models describe the mobile network performance for a specific: a network operator, country, RAT (i.e., 3G and 4G), and signal quality. Plus, we provide 6 additional models to be used as general models describing an

operator offering 3G or 4G connectivity with different signal qualities. Given that ERRANT is open-source, users willing to contribute can perform their own measurement campaigns to create new emulation models for specific operators, locations, new technology, i.e., 5G.

ERRANT provides dynamic models describing not only the *typical* (and simple) behavior of mobile networks, but also their variability. For this, we employ the Kernel Density Estimation (KDE) [8] technique to create accurate models of the observed network, based on the speed test measurements. Then, ERRANT emulates a dynamic behavior by extracting the network emulation parameters from the models. By extracting new parameters at each run (or changing parameters within a run following the tester decisions) ERRANT can correctly emulate the variability of the network conditions observed in the real networks. In other words, our tool allows the user to run *several* experiments, sampling realistic values from the modeled distribution, and, in turn, to observe the distribution of the desired effects. To apply the emulation parameters, under the hood ERRANT uses the `tc-netem` Linux emulation tool which is the *de facto* standard.

We validate the models comparing the download speed in our emulated environment with those obtained with real-world experiments in which mobile nodes download large-sized HTTP objects. In addition, we compare our results with other free emulators. The results show that not only ERRANT models reproduce the *typical* behavior of mobile networks, but also it introduces similar levels of variability which are not present with the other emulators. As such, using ERRANT, a tester can accurately try her service as if she was under quasi-real mobile network conditions.

Finally, we show how our models help in evaluating the Quality of Experience (QoE) of a user accessing web services. We compare the performance of web pages accessed under different network conditions and notice a strong impact on page load time, in terms of absolute values and variability. Indeed, simple emulation does not pinpoint the variability QoE-related metrics that would occur in real mobile networks. We also show that video streaming as another use case, in which emulating variable network conditions lead to considerably different results in terms of QoE.

We make ERRANT open source and available to the community [9] to allow practitioners to test their applications and researchers to reproduce the results. This paper extends our preliminary work [10], where we already have shown the potentialities and limitations of using a data-driven emulator. We expanded our work from many points of view: (i) we now rely on a large-scale speed test measurement campaign; (ii) we provide a fully-fledged network emulator that models not only latency but also download and upload bandwidth; (iii) we show practical use cases to illustrate the benefits of our proposal in understanding the behavior of user applications with a particular focus on web browsing and video streaming; (iv) we provide an open source tool to emulate mobile networks in different operators, locations, and technologies. The remainder of this paper is organized as follows. Section 2 discusses related work. Section 3 presents the measurement setup, the datasets, and how to replicate a measure-

Table 1: Emulation parameters of default profiles for 6 freeware testing tools, separately for download (D), upload (U), and latency (L). We omit *bps* for download and upload, and *ms* for latency for the sake of space.

	3G			3G Slow			3G Fast			4G		
	D	U	L	D	U	L	D	U	L	D	U	L
Chrome [11]	750 k	250 k	100 -	-	-	-	1 M	750 k	40	4 M	3 M	20
WebPageTest [12]	1.6 M	768 k	300	400 k	400 k	400	1.6 M	768 k	150	12 M	12 M	70
BrowserTime [13]	1.6 M	768 k	300	780 k	330 k	200	1.6 M	768 k	150	-	-	-
ATC [3]	780 k	330 k	200	850 k	420 k	190	-	-	-	-	-	-
Android Emulator [14]	14 M	5.76 M	0	384 k	384 k	35-200	-	-	-	173 M	58 M	0
NLC [2]	780 k	330 k	100 -	-	-	-	-	-	-	51.2 M	10.24 M	65

ment campaign. The model creation and emulation methodology are presented in Section 4. Section 5 presents the experimental results, while in Section 6 we show two practical use cases. Finally, Section 7 concludes the paper.

2. Related Work

Network monitoring provides crucial information about the performance of operational networks and helps network managers, researchers, and regulators to understand the behavior of the network and users’ QoE. There are nationwide [15, 16] and international [17] efforts to monitor and measure the performance of networks. Non-stop measurements become more important when it comes to the latency and throughput of mobile networks. The big challenge is to translate the performance of the network into the users’ QoE [18, 19].

There is a large body of research works that shows the effect of different scenarios in the performance of mobile networks. The latency is a critical metric for many applications, e.g., audio streaming applications. Mobile access networks and mobility scenarios present very high variability in performance – due to the more random nature and complexity of mobile networks. A survey of end-to-end delay prediction methods is offered in [20]. They show several methodologies to estimate latency, e.g., *Queuing Network Modelling*, *System Identification*, *Time Series Approach*, and *Neural Networks*. Nikravesh *et al.* [21] indicate significant performance differences across operators, access technologies, geographic regions and over time; however, they highlight that these variations themselves are not uniform, making network performance difficult to predict. Safari *et al.* [22] illustrate accurate mobile network latency prediction is not effective by using the common machine learning algorithms. Mandalari *et al.* [23] present the impact of roaming on latency with ≈ 60 ms or more, depending on geographical distance. The feasibility of the latency-sensitive application, teleoperated driving is demonstrated by means of measurements using mobile networks however it could be challenging in some cases with high variation in performance of the network [24].

Considering throughput, researchers investigated its behavior under mobile networks and studied factors potentially impairing the performance. Deng *et*

al. [25] investigate the characterization of multihomed systems considering WiFi vs. LTE in a controlled experiment. They show that LTE can provide higher performance than WiFi, also exhibiting large variability on both short and long time scales. Rathnayake *et al.* [26] demonstrate how a prediction engine can be capable of forecasting future network and bandwidth availability, and propose a utility-based scheduling algorithm which uses the predicted throughput to schedule the data transfer over multiple interfaces from fixed nodes. Safari *et al.* [27] demonstrate the higher performance variation for moving user traveling in the vehicle via measurement. They opted to use a custom heuristic for video upload application. Midoglu *et al.* [7] show that the parameters in speed test application configurations, e.g., the number of TCP flows, measurement duration, and server location on measured downlink data rate can significantly affect mobile network measurement results. Safari *et al.* [6] illustrates the high variation in mobile network performance through a large-scale measurement. They show the performance and network setting in the European mobile networks is highly dynamic and hard to predict. Nunes *et al.* [28] focus on TCP performance in a Mobile Ad Hoc network and propose a machine learning technique called *Experts Framework*.

Finally, it is vital to understand the users' QoE. Asrese *et al.* [29] present a measurement tool to measure web latency and QoE in the cellular network. They show that the DNS lookup time and Page Load Time (PLT) of the visited websites have similar performance in LTE and fixed-line networks. However, the TCP three-way handshake time and Time To First Byte (TTFB) of the considered websites are longer in LTE networks. However, the authors in [30] illustrate the high randomness in mobile network performance dominates over the benefit of some new protocols, e.g., HTTP2 over HTTP1.1. Casas *et al.* [31] study the problem of how to extract YouTube performance indicators related to the QoE perceived by end-users and they show the potential and feasibility of doing real-time QoE monitoring in services such as YouTube in mobile networks. Cecchet *et al.* [32] present an open-source infrastructure to measure the QoE of Web application on mobile devices. They illustrate the mobile network performance is a dominant factor in user-perceived QoE over device performance or user location. Hosek *et al.* [33] illustrate QoE for mobile web service in different network conditions and highlight the impact of network parameters on mobile web QoE. Balachandran *et al.* [34] show the impact of low-level measurable radio network characteristics on the user QoE during web browsing. The WebLAR tool measures web QoS and web QoE and it shows the TCP handshake time and Time To First Byte mobile networks longer than fixed-line networks using measurement in different mobile networks. Moreover, it demonstrates QoE across several operators [35]. In [36], authors model the QoE of YouTube Live service from packet-level data in a Wi-Fi network. Similarly, Oliver-Balsalobre *et al.* [37] model key service performance indicators from TCP/IP metrics for video-streaming (YouTube) in a Wi-Fi network has been presented. These are different from this work in several aspects, they focus on modeling the video stream QoE in Wi-Fi networks. Importantly, they emulate the wireless network by using a limited combination of performance metrics

namely, packet loss, packet delay, and maximum throughput.

In this context, network emulation plays a crucial role in network testing, and it is a well-studied topic in the literature. The Linux operating system has been including `tc-netem`, a tool for this purpose for more than 20 years. Several testing applications rely on it to install shaping policies reflecting different network conditions. Commercial products [4, 5] emerged to offer software and hardware solutions for network emulation and testing, based on profiles decided at design-time defined by several parameters. Differently from them, we offer ERRANT as a freeware solution, exposing high-level network models describing the intrinsic variability of the network.

In Table 1, we summarize the mobile network profiles offered by six freeware testing platforms. We include tools used to test web pages (Google Chrome, WebPageTest, BrowserTime), operating systems (Android Emulator) and general-purpose frameworks (Advanced Traffic Control and Network Link Conditioner). They provide profiles decided at design-time, that simply impose simple traffic shaping policies, with limited variability. Moreover, they provide values highly heterogeneous, with, for example, 3G having download bandwidth from 750 kbits to 14 Mbits depending on the tool. In our work, we provide shaping values based on real-world measurements and offer a tool able to emulate not only plausible values but also the observed variability.

Several research works try to build realistic emulators of mobile networks. Many works aimed at achieving realism with fine-grained models of network devices. For instance, NIST Net [38] is a flexible and powerful emulator for WANs, while authors of [39] exploit advanced modeling methods to build a mobile network emulation environment. Beuran *et al.* [40] propose a multi-purpose wireless network emulator to emulate 802.11 a/b/g WLAN technology over wired networks for several types of scenarios. Veltri *et al.* [41] design an emulator to emulate a general IP network either a single link or a portion of a network.

More recently, mobile network emulation followed the general trends of research in networking, and has been used to mimic real deployments of e.g., Software Defined Networks [42] or flying vehicles networks [43]. There are few research studies building on a data-driven approach to obtain a realistic emulation. Noble *et al.* [44] propose a pure trace-driven network emulation that re-creates the observed end-to-end characteristics of a real wireless network. KauNet [45] provides pattern-based emulation with a higher level of detail, controlling the behavior of each individual packet. Rupp *et al.* [46] present the Vienna Simulator Suite mobile network simulators at access link level in different levels and scenario. In contrast to previous works, we build an emulator based on a large measurement dataset of operational mobile networks and make it available to the community and compatible with the `tc-netem` toolset.

3. Measurements setup

In this section, we briefly describe the MONROE platform that we use to gather the measurements as well as the collected datasets. We also provide

guidelines and source code to run a new measurement campaign. Table 2 briefly describes the datasets we use.

Table 2: The high-level description of the employed datasets.

Name	Description
<i>MetaData</i>	MONROE node metadata, reporting RAT, signal strength and other physical layer conditions over time
<i>HTTP</i>	HTTP downloads performed periodically by MONROE nodes using the <code>curl</code> tool, used to validate ERRANT
<i>SpeedTest</i>	Speed tests over TCP measuring download and upload bandwidth, and latency from 20 MONROE nodes

3.1. Measurement infrastructure

Systematic, repeatable measurements are crucial for evaluating network performance and assessing the quality experienced by end-users. As such, researchers have built several platforms dedicated to broadband networks, e.g., RIPE Atlas¹, CAIDA Ark², or PlanetLab³. In contrast to them, MONROE [17, 47] is a unique platform that enables controlled experimentation in a mobile network environment with different commercial mobile operators. In more detail, each MONROE node is equipped with different 3G/4G mobile broadband modems each one connecting to a different operator. Figure 1 reports a schematic view of a specific scenario on MONROE platform. Each node allows researchers to run custom experiments, to schedule measurement campaigns, and collecting data from operational mobile networks and Wi-Fi networks together along with full context information. The MONROE infrastructure covers 4 countries in Europe (Italy, Norway, Spain, and Sweden) with more than 100 nodes/locations.

The MONROE platform allows us to access the information about network, time and location of experiments, as well as metadata from the mobile modems including, e.g., signal strength, RAT, cell identifier for each network provider.⁴ We collect and use metadata, and call them the *MetaData* dataset. The MONROE Alliance⁵ offers different plans for new members and partners to have new nodes to be installed in their countries/locations. Moreover, all necessary software to replicate a MONROE node and base experiments are open source. Each node performs a set of predefined experiments (e.g., ping, passive monitoring, etc.) continuously. Experiment results are collected in the central database and

¹<https://atlas.ripe.net/>

²<http://www.caida.org/projects/ark/>

³<https://www.planet-lab.org/>

⁴The description and detail information about metadata is available on <https://github.com/MONROE-PROJECT/UserManual>

⁵<https://www.monroe-project.eu/become-a-member/>




Figure 1: The experimental setup for speed test measurement in the MONROE platform.

stored in a Hadoop-based cluster. In the following, we provide details about the experimental setup and describe the considered metrics.

3.2. HTTP download measurements

Among other predefined experiments, MONROE nodes perform simple HTTP downloads periodically. Using the popular `curl` tool, each node downloads a large-sized HTTP object from an *ad hoc* server located in the same country, e.g., Sweden. The object has a size of 100 MB and is typically downloaded in a few seconds, depending on the network conditions. Once the download ends, the node computes statistics such as download speed and duration, as well as several TCP-level metrics. Periodically, experiment results are sent to the central database. In total, 84 nodes/locations performed 19k downloads for 1 year using 10 mobile network operators. In this work, we make use of these data and call them the *HTTP* dataset.

Table 3: Employed speed test measurements. We use data from 2 countries and 4 operators.

Country	Operator	Nodes	3G	4G
Norway	Telenor	40	495	19 707
	Telia	35	769	13 373
	Net1	28	-	11 800
Sweden	Telenor	47	1 156	17 222
	Telia	59	1 969	23 077
	3	46	3 347	8 837

3.3. Speed test measurements and dataset

For the speed test measurements, we use the dataset collected in [7]. The authors built a custom measurement tool based on Open-RMBT [48] that they called *MONROE-Nettest*. It measures latency, downlink, and uplink capacity using multiple TCP connections, plus it measures the distance in terms of network hops by using `traceroute`. We refer the reader to [7] for further

details about *MONROE-Nettest* design. The authors used *MONROE-Nettest* to run a large-scale measurement campaign involving 4 operators on 2 countries, namely Sweden and Norway. The experiments were conducted from 20 MONROE nodes/locations towards 2 servers *ad hoc* deployed in the respective countries. The `traceroute` measurements collected along each experiment allow us to infer the distance between the node and the server, hence the complexity of the traversed network. Results show that the traffic from the mobile node to the server never leaves the monitored country. Moreover `traceroute` measurements show that typically packets transit via 5–7 hops. Over one year, they performed more than 100 k speed tests, including different hours of the day and days of the week. We provide a breakdown of the dataset in Table 3. Relevant to our analysis, the measurements were taken under diverse network conditions, with both 3G and 4G RAT and different levels of signal quality. We call this dataset *SpeedTest*.

3.4. New dataset creation

Here, we briefly summarize how other researchers can run new measurement campaigns to gather new datasets, leaving the possibility to contribute to the ERRANT project. To this end, researchers need a dedicated server and a mobile node attached to a Mobile Network Operator. The mobile node is used (i) to perform speed test measurements toward the dedicated server, and (ii) to gather network-level metadata i.e., the RAT in use and the signal strength.

For steps (i) and (ii), we recommend the MONROE platform⁶ offering the *MONROE-Nettest* tool [7] to gather speed test measurements and exposing network-level metadata about each speed test⁷. However, ERRANT can be built on top of any automatic framework for mobile network measurements. Hence, an user willing to run new measurements, could either rely on this code repository or perform speed tests independently.

The results of the experiment campaign must produce a `csv` file describing the speed tests along with metadata information. To check the format of the file, in the ERRANT repository [9] we report an example of it. The repository also contains the code to create new ERRANT models from the `csv` file, as we explain in the following section.

4. ERRANT: EmulatoR of Radio Access NeTworks

In this section, we illustrate the design of ERRANT. First, we describe all the data processing steps required to clean and join the *MetaData* and *SpeedTest* datasets to create network profiles. Then, we describe how we build our network emulation models based on the available profiles. Finally, we explain how a tester can use ERRANT models to achieve realistic traffic emulation. Figure 2 provides an overview of the model creation process.

⁶The MONROE codes to replicate the platform is available at <https://github.com/>




Figure 2: Generation of ERRANT models. Fields with * are used only for specific profiles and not to build universal ones.

4.1. Data cleaning and processing

Armed with the *MetaData* dataset, we first map the signal strength values in three *signal quality* levels: bad, medium, and good. Table 4 reports the binning boundaries.⁸ Then, we join the *SpeedTest* and the *MetaData* datasets to augment each speed test measure with contextual information (i.e., country, operator, RAT, signal quality). Given the large size of the datasets, most of the processing here is performed using an Apache Spark cluster. Here, we discard a small number of speed tests due to the lack of physical layer information caused by temporary issues of the nodes. At this point, we aggregate measurements taken under the same conditions to build the network *profiles*. In particular, we build *specific* network profiles for country and operator by aggregating the *SpeedTest* data by:

- *Country*: where the measurements were taken
- *Operator*: which operator was used
- *Radio Access Technology*: 3G or 4G
- *Signal Quality*: the mapped signal strength i.e., the *rsqi*.

Furthermore, we create *universal* profiles, containing the performance of *any* operator using different RAT and signal quality, by aggregating the *SpeedTest* data only by: *Radio Access Technology* and *Signal Quality*.

After discarding all profiles with less than 100 measurements i.e., those not having enough samples to be reliable, we end up with 32 network profiles: 26 for specific operators and countries and 6 universals. They include data from 2 countries and 4 distinct operators (2 of which are present in both countries). They also differ for the signal qualities, in particular we have 6 profiles with *bad*

MONROE-PROJECT

⁷The description is available at <https://github.com/MONROE-PROJECT/data-exporter>

⁸Boundaries are defined based on https://wiki.teltonika.lt/view/Mobile_Signal_Strength_Recommendations.

Table 4: Binning boundaries for 3G and 4G. Values are expressed in *dB*.

Quality	3G	4G
Bad	$rss_i \leq -100$	$rss_i \leq -85$
Ordinary	$-100 < rss_i \leq -85$	$-85 < rss_i \leq -75$
Good	$rss_i > -85$	$rss_i > -75$

quality, 13 with *medium* quality, and 13 with *good* quality. Considering RAT, we have 14 profiles for 3G and 18 for 4G. As physical conditions (signal quality and RAT) could not be controlled when measurements were taken, we have a significant imbalance in the quantity of per-profile samples.

To illustrate the available network profiles and show their difference, in Figure 3, we report for each one the median value of download bandwidth, upload bandwidth, and latency. Focusing on the download and upload bandwidth, we notice how median values are not sufficient to strongly differentiate different 3G profiles, while good 4G profiles strongly differ from each other. On the one hand, this highlights the importance of having dynamic profiles, as static values would not well differentiate between different 3G networks. On the other hand, despite the universal profiles give a general idea of the possible performance, it is also important to define specific profiles to capture the differences among operators or countries.

It is interesting to notice that the measured values are in all cases far from the nominal data rates of the radio standards. For 3G, download throughput seldom exceeds 20 Mbit/s while the HSPA+ standard should reach up to 42.2 Mbit/s [49]. The difference is even larger for 4G, where the LTE specification provides downlink peak rates of 300 Mbit/s and uplink peak rates of 75 Mbit/s [49], while we observe values in the order of 20 – 60 Mbit/s.

To underline the importance of variability of network conditions, in Figure 4, we show the distribution of the measured values for download bandwidth (Figure 4a), upload bandwidth (Figure 4b) and latency (Figure 4c) for the Telia operator in Norway. We report separately profiles obtained with different combinations of RAT and signal quality. Boxes in the figure span from the first to the third quartile, whiskers from the 5th to the 95th percentile and the black strokes represent the median values. The impact of physical layer conditions appears clear. Considering download bandwidth in Figure 4a, 4G performs better than 3G, e.g., good 4G offers in median 20 Mbits in download, while good 3G only 6 Mbits. The role of signal quality is sizeable too, with good 4G almost twice as fast as bad 4G. Variability is high in all cases, mandating a realistic emulator to consider it. For latency (Figure 4c), the impact of physical layer conditions is limited, with similar median values, except for bad 3G. Again, variability is rather high, with the inter-quartile range in the order of tens of milliseconds. Similar considerations hold for other profiles, not reported here for the sake of brevity. However, Figure 3 still shows the overall absolute values, while the variability for all profiles can be found in the ERRANT repository [9]. This result emphasizes the high variability in mobile networks and high dependency

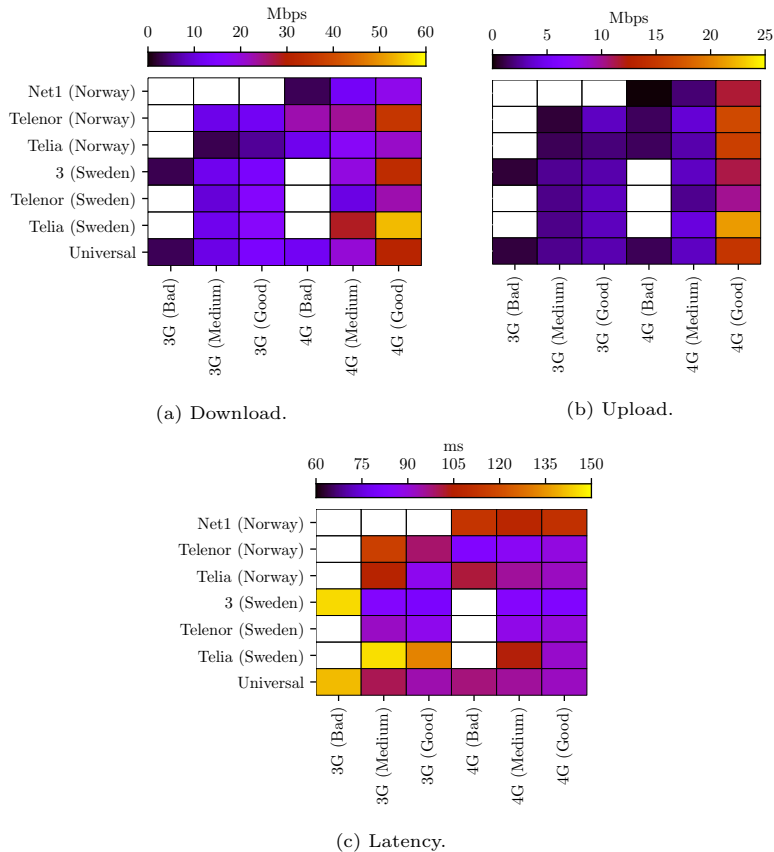


Figure 3: Overview of the obtained profiles. Only median values are reported, while variability is omitted.

on location, operator, technology, and signal quality. Our data represent a large set of operators in more than 80 locations with two RATs but clearly do not include all possible scenarios. As such, the open-source MONROE platform and ERRANT gives the opportunity for users to reproduce the same methodology and realistically simulate their local operators, other than sharing their datasets with the community. 5G technology is testing and installing around the world and it calls for performance analysis and simulators. For this, the MONROE Alliance [50] is working to equip and adapt the platform to 5G technology, giving the possibility to have new 5G profiles in ERRANT in the future.

4.2. Emulation model creation

Given the different network profiles, here we use the KDE methodology [8] to build our emulation models, each one modeling a different profile. From a profile p , we extract the network parameters measured by each speed test in the

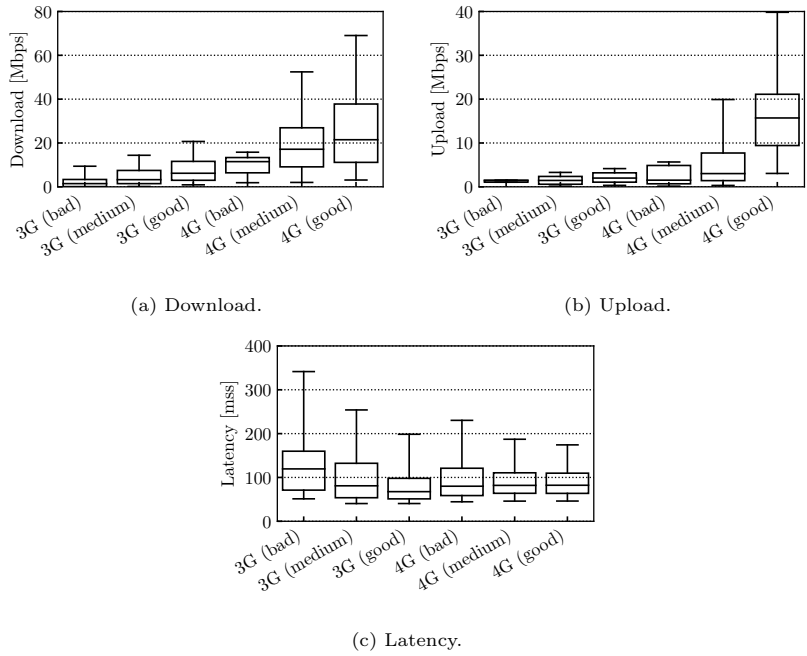


Figure 4: Speed test results for Telia (Norway).

form of download bandwidth (B_D), upload bandwidth (B_U), and latency (L). Although the characteristics of a network might be more complex to model (e.g., due to jitter, queue size, etc.), the speed test measures only provide these three parameters. Moreover, these are more straightforward to tune and several network emulation tools use only these three values to define a network behaviour. As a result each speed test measurement in p is described by a tuple in the form: (B_D, B_U, L) . Instead of merely sampling from the observed (B_D, B_U, L) values, we build our dynamic *emulation model* M_p using the KDE methodology.

KDE is largely used in the literature as a tool to model multivariate problems [51, 52]. In this work, we use the KDE to estimate the multivariate probability of observing (B_D, B_U, L) in a speed test measurement for profile p . We fit a three-dimensional KDE on all possible value of (B_D, B_U, L) for each profile. In this way, we obtain distinct models describing the network parameters for each profile. We use a Gaussian kernel to define the influence of each speed test measure and set the bandwidth with the Silverman method [53]. For the implementation, we rely on the SciPy Python module. In a nutshell, we use KDE as a data modeling tool able to capture performance patterns from speed tests, hence giving us the possibility to abstract performance samples while reducing the impact of noise or secondary finer-scale phenomena.

This said ERRANT considers all the measurements collected for profile p ,

and estimates their multivariate probability function to build M_p . The novelty of ERRANT is twofold. First, the emulation models are based on a large-scale measurement campaign, including 100 k speed test measurements. Second, ERRANT exposes models and not static values. Hence ERRANT can correctly replicate the variability of the network conditions as observed in the real *SpeedTest* dataset.

We provide the code to create new models given a set of speed test measurements in the repository [9]. This code, written in Python, takes as input a `csv` file with the speed test measurements along with their metadata information as described in Section 3.4. As a result, it creates an output file containing the models directly usable by ERRANT. Using this piece of code the user can translate an experimental campaign into additional models describing different operators, different locations, mobility, different RAT like the 5G access technology.

4.3. Emulation model usage

We now describe how ERRANT uses the models at run time, to dynamically emulate different network conditions. At run time ERRANT offers three possibilities to sample new values for the tuple (B_D, B_U, L) from model M_p .

Firstly, at each run ERRANT can extract different values for the three key parameters. In other words, at each run, ERRANT samples potentially different parameters, following the observed distribution. As such, the more likely certain values appear in real data, the more frequent will be sampled from the distribution. The main goal of ERRANT is to allow the user to run *several* experiments using a given model p , and, consequently, to observe the *distribution* of the obtained effects. We will use this solution to validate our models in Section 5.1.

Secondly, ERRANT allows the tester to sample new emulation values every n seconds to emulate dynamic conditions within an experiment e.g., testing watching video while moving. Note that we cannot provide data-driven guidelines on how to set n as we have not performed experiments using mobile nodes, and, as such, we leave the choice of this parameter to the user. However, studying mobile networks in their dynamic behavior during mobility is of great interest, but we leave this for future work. A typical use case is testing the performance of adaptive protocols such as the Dynamic Adaptive Streaming over HTTP under different mobile network scenarios. We explore this and other use cases in Section 6.

Finally, ERRANT also offers the possibility to change network profiles within an experiment following a *trace-driven scenario* written in a configuration file. By using the convenience script `trace_run`, the tester can run an experiment following a scenario containing potentially drastic networking changes e.g., a user going through a tunnel, or moving from an area with good coverage to a bad one. This feature can be used to evaluate the application performance in stressful stations such as a connection migrating from a Good 4G quality, down to a Bad 3G one, and back to a 4G Medium situation.

Under the hood, ERRANT uses the `tc-netem` Linux tool to impose traffic shaping policies on the selected network interface. In particular, we impose bandwidth limitations by means of the Hierarchy Token Bucket (HTB) a classful `qdisc`, while we use the `netem qdisc` to emulate latency. We make use of the Intermediate Functional Block (`ifb`) kernel module to create the pseudo network interface necessary to control the incoming bandwidth. We validate ERRANT accuracy in emulating network conditions in Section 5.

We release ERRANT as open-source, and it is available online [9]. It consists of a simple tool that installs traffic shaping policies according to the model requested by the user in one of the above solutions on the desired network interface.

5. Experimental results

In this section, we analyze the performance of ERRANT, comparing the results we obtain with those collected in a real environment and with other freeware emulation tools. We then study the impact of the number of speed test measurements on the created models to suggest the minimum number of experiments required to create a model.

Simple Emulator Baseline. In the following, we present ERRANT results along with those obtained with a simpler emulation method, that we use as a baseline. We call it the *Simple Emulator*, and it consists of a basic script that uses `tc-netem` to install the desired shaping policies. While ERRANT samples new shaping parameters at each run, our baseline approach only uses the average observed values of download and upload bandwidth for the given model. To recreate variability, we impose a Gaussian latency distribution with the average and standard deviation values observed in the empirical distribution. Note that `tc-netem` allows variability only in latency, while there is no straightforward way of varying the bandwidth automatically. We show that a simple approach based on the Off-The-Shelf features of `tc-netem` fails to provide realistic results when it comes to real applications such as web browsing or video streaming.

5.1. Model validation

In this section, we validate our emulation models showing that with ERRANT a tester can accurately emulate mobile network conditions, leading to results comparable with those obtained in real networks. Here, we do not provide results about the accuracy of ERRANT in installing accurate traffic shaping rules, as for that we rely on the state-of-art tool `tc-netem`. Rather, we show that running a simulation campaign using ERRANT allows the user to observe the same distribution of desired effects that would observe in a real mobile environment.

We rely on the *HTTP* dataset to validate ERRANT accuracy on replicating the network variability. The *HTTP* dataset allows us to test the models with data that were never used to build the model. We aim at understanding to

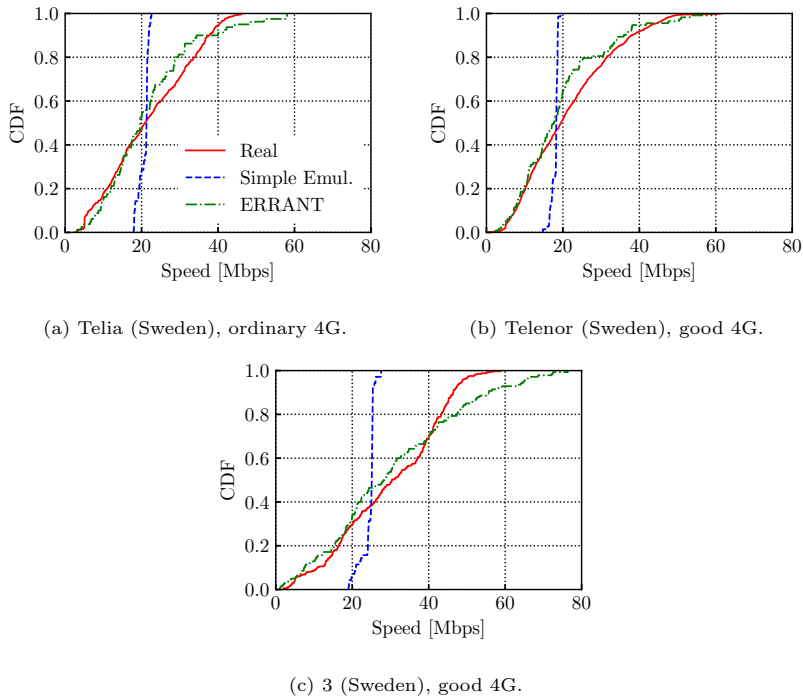


Figure 5: Comparison of HTTP download speed for 3 scenarios.

what extent ERRANT is able to recreate the variability of conditions experienced in real networks, in addition to the *average behavior*. To this end, we join the *HTTP* and the *MetaData* dataset, to enrich the former with the contextual information in which the HTTP downloads were performed (operator, signal quality, etc.). This allows us to map each experiment to the corresponding profile that we can emulate with ERRANT. Then, we use ERRANT to replicate the experiments emulating the corresponding network conditions – i.e., to download large-sized HTTP objects. This is done for all profiles for which we find real data in the *HTTP* dataset, 12 in total.

Then, for each profile, we download 1000 times a 10 MB HTTP object from a dedicated server. At each download, we run ERRANT to sample new parameters of (B_D, B_U, L) for the corresponding model. These experiments aim at evaluating the impact of the variability introduced by ERRANT in the final results. Indeed, within each experiment, it imposes fixed emulation parameters, but they vary across different downloads. For reference, we also use the *Simple Emulator* to provide a baseline.

Figure 5 presents results for three models. Similar results are obtained in other cases. Each plot reports the Cumulative Distribution Function (CDF) of the average download speed for the HTTP object at each run, as measured by

the `curl` tool. We separately report results measured in real mobile networks (solid red line), using ERRANT (green dashed line), or with *Simple Emulator* (blue dashed line). For all profiles, the speed observed in the real case has a considerable variability, with the inter-quartile range spanning tens of Mbits. Looking at obtained results when using ERRANT, they appear very similar to the real case, considering both median values as well as the shape of the distribution. The result is different when imposing fixed network conditions with the *Simple Emulator*, where negligible variability is observed and where only the median value correctly describes the real data.

In conclusion, we show that ERRANT accurately recreates mobile network conditions, leading to similar results in terms of average behavior and observed variability when downloading HTTP objects.

5.2. Comparison with other emulation tools

After validating ERRANT models, we compare its performance with other freeware network emulation tools already reported in Table 1. Similar to ERRANT, these tools allow the tester to tune the download and upload bandwidth, in addition to the network latency based on different network profiles. Each tool exposes off-the-shelf network profiles, whose parameters are decided at design-time. In addition, these tools allow the user to define custom profiles based on user-specified parameters.

Similarly to the experiments in the previous section, we use the tools reported in Table 1 to impose traffic shaping rules on a test machine. In parallel, we perform HTTP downloads using `curl` and record the average download speed per download. For each experiment, we download 1000 times a 10 MB sized file.

In Figure 6, we compare the performance of the 4 network emulation tools, *Simple Emulator*, and ERRANT with the performance observed in the real environment for two 4G (good) profiles. We use boxplots to represent the distribution of the download speed measured in each scenario. Boxes in the figure span from the first to the third quartile, whiskers from the 5th to the 95th percentile and the black strokes represent the median values. First, we note again that ERRANT leads to results similar to the real case, for both Telia and Telenor operators. The *Simple Emulator* instead only reproduces the average behavior but fails in providing variability in the results. Considering the four emulation tools, we first note that two of them (NLC and Android) lead to results very different from what we observe in the real networks. In particular, these two lead to a download speed of 50 and 90 Mbits respectively, while we seldom observe values higher than 30 Mbits on the real environment. WPT and Chrome instead have typical values closer with the measurements taken under the Telenor operator especially. However, we note that they provide no variability in the results, expected considering that they impose static shaping policies.

In summary, differently from other freeware traffic shaping tools (e.g., for web page testing), ERRANT imposes network conditions based on real large-scale experiments. Allowing one to re-obtaining the variability intrinsically




Figure 6: Distribution of HTTP download speed in a real environment, using ERRANT and other emulation tools.

rooted in mobile networks. And offering models better describing the performance of real mobile networks.

5.3. The impact of profile size

Understanding how many speed tests are required to create a reliable model is of crucial importance to determine the size of new experimental campaigns, e.g., in the case of 5G. To this end, we study the impact of the number of measurements on the obtained *profiles* as they are the baseline to create our models. In particular, we study to what extent a *profile* composed by a small number of speed tests is representative of those obtained with the full dataset.

For this analysis, we discard profiles with less than 10 k speed tests, that would not fit the analysis. We obtain 4 profiles, from which we extract, by sampling, exactly 10 k points, for a fair comparison. We now perform experiments sampling a smaller subset of samples, and compare them with the full 10 k samples. At each run, we randomly create a subset of n points per profile, where each one is described by a value of (B_D, B_U, L) . Then we evaluate the similarity between the subset with respect to the full set of data by means of the two-sample Kolmogorov–Smirnov test [54, 55], that measures the dissimilarity of two empirical distribution functions providing the D statistic. For each experiment having a different subset size, we take 100 independent subsets of n points. Intuitively, D measures the maximum distance between the two sets, the higher is the value the more is the dissimilarity. As such a high D means that with the given subset we cannot create a profile similar with respect to the one created with the whole dataset. Hence, we cannot create a reliable model. We run our analysis separately by B_D , B_U , and L since all proposals that adapt the Kolmogorov–Smirnov test to multivariate distributions suffer for severe scalability issues [56].

Figure 7 shows the distribution of the D statistic when varying the sample size n . Clearly, when $n = 10$ k, D is always 0 as the two distributions are identical. For small values of n we obtain considerably different distributions,




Figure 7: Kolmogorov–Smirnov D statistic when comparing the distributions obtained with 10k with those obtained with smaller sub-samples.

with D often higher than 0.3 when $n = 10$. By increasing n , the distributions become more similar, with median values of n reaching 0.08 for $n = 100$, and 0.02 for $n = 1000$. These results confirm the importance of having a large dataset but show that in a new campaign with 100 speed tests we can already have a reasonable estimation of a profile. Indeed, for $n = 100$ the two sample sets come from the same distribution with a *significance level* always above 0.2, generally considered strong evidence. As such, in ERRANT we create models only for profiles having more than 100 points, as stated in Section 4.

6. Use cases

In this section, we illustrate the capability of ERRANT in providing useful insights for various stakeholders of computer networks. Indeed, accurate network emulation is fundamental for website designers, web application developers, and network providers to fully understand the impact of their design choices by monitoring the performance of applications before the deployment.

6.1. Web browsing QoE estimation

We here focus on the use case of web browsing to show how ERRANT improves the emulation with respect to *Simple Emulator*. It is crucial to understand the users' QoE especially in mobile networks which makes it more complicated as investigated by the authors of [29–36]. Moreover, large companies operating on the web are particularly interested in studying the factors affecting QoE, with Google and Amazon reporting losses in the 0.6-1.2% range if delay increases by 0.4-1 sec.⁹

For our experiments, we consider three popular websites, chosen from the globally top-ranked in the Alexa ranking service, namely: Wikipedia, Google,

⁹<https://www.fastcompany.com/1825005/how-one-second-could-cost-amazon-16-billion-sales>

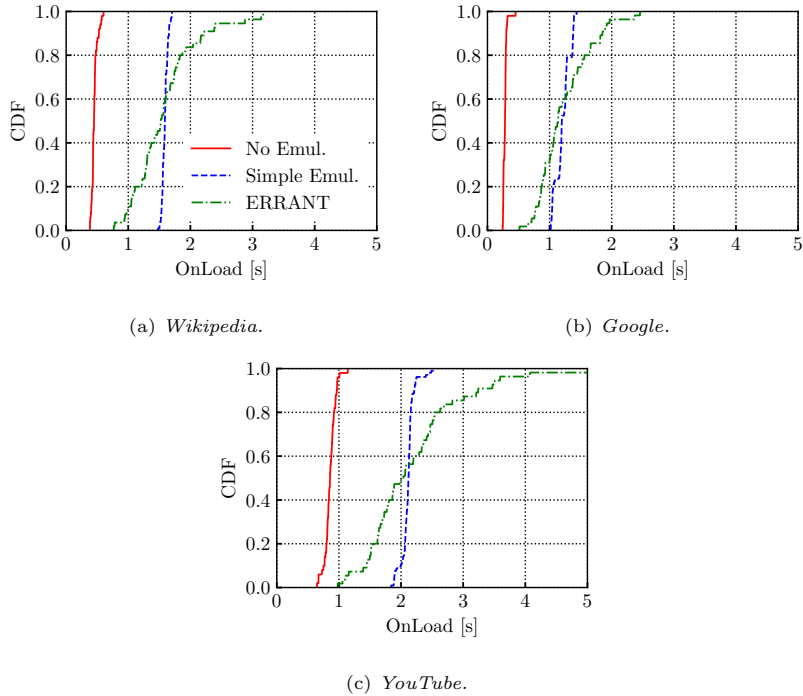



Figure 8: Comparison of *OnLoad* time for 3 websites, when emulating good 4G for Telia in Sweden.

and YouTube.¹⁰ We then instrument Google Chrome to automatically visit the homepages of the selected websites using the Browsertime toolset.¹¹ In parallel, we use ERRANT to varying network conditions, using all the 32 profiles we built. For each profile and website, we perform 100 visits. For comparison, we also use the *Simple Emulator* for each profile. Finally, we visit the websites without any traffic shaping to provide a reference in the results. At each run, we collect all the metadata about the visit, including QoE-related metrics such as *OnLoad* time and *Speed Index*. Indeed these two metrics have been proved to have a high correlation to users' QoE, intrinsically subjective [57]. In total, we performed more than 18 k visits.


Figure 8 shows results for *good 4G, Telia (Sweden)*, while we omit the other profiles as they lead to similar conclusions. Figure 8 reports the CDF of *OnLoad* time when using ERRANT, or static emulation or finally no emulation at all. We do not report results for *SpeedIndex* for the sake of brevity. Comparing ERRANT results (red line) with those obtained without emulation (green line),

¹⁰<https://www.alexa.com/topsites>

¹¹<https://github.com/sitespeedio/browsertime>



(a) *Bad 3G.*



(b) *Bad 4G.*

Figure 9: Example of adaptive streaming video sessions.

we notice how OnLoad time is strongly affected, as expected, for all websites. Instead, the *Simple Emulator* (blue line) leads to values of page load time with similar median values – e.g., 2 seconds for YouTube. However, ERRANT imposes variability, based on the speed test data used for creating the profiles, and this variability reflects into the page load time. For example, YouTube page load time has an inter-quartile range of 0.8s while only 0.1s is observed with static emulation. Moreover, in 15% of cases, we observe a value larger than 3s, which makes 53% of users to leave the website according to recent studies [58]. As such, we stress the fact that ERRANT can reproduce the variability of mobile networks, allowing one to thoroughly study the behavior of applications with a level of realism previously unavailable.

6.2. Adaptive video streaming

Here, we show a second use case, in which ERRANT is used not only to introduce *inter-experiment* but also *intra-experiment* variability. We focus on Adaptive Video Streaming over HTTP. Indeed, video streaming is nowadays widespread and represents the majority of Internet traffic. Among tens of protocols dedicated to video streaming, in the last years, most of the providers converged towards the use of video streaming over HTTP that relies on con-

ventional HTTP servers with no specific support. The most common solutions are called Dynamic Adaptive Streaming over HTTP (DASH) and HTTP Live Streaming (HLS), and both realize adaptive bitrate streaming splitting the video in chunks of few seconds. With the growth of high-resolution and 4K video providers, the network is required to ensure high and stable bandwidth in order to guarantee users' QoE. Thus, it is not surprising that video providers publish online rankings of ISP quality.¹²

The goal of this section is to show that ERRANT helps in obtaining realistic figures about video streaming QoE under mobile networks, where the network quality may change while reproducing the video. To this end, we set up a simple DASH server that we use to replay a 200 seconds full HD (1080p) video. Using Google Chrome instrumented with Selenium Web Driver, we play the video while simultaneously using ERRANT to emulate mobile networks. We perform several experiments, configuring ERRANT to periodically change the network conditions every $n = 10$ seconds to emulate a variable network within the selected model. Like in the previous sections, we also perform experiments with the *Simple Emulator* that we use as a baseline.

The results show the impact of network conditions and their variability on the users' QoE. Indeed, network conditions impact the bitrate at which the video is consumed. Moreover, a variable network leads to frequent adaptations in video bitrate, as the client needs to adjust it to avoid running out of video buffer. These two factors are commonly accepted as objective metrics to study users' QoE [59–61]. Using ERRANT, one can understand how mobile networks impact such metrics and run new experiments to evaluate the impact of different configurations of the player or the server. Figure 9 shows two examples of video sessions for bad 3G and 4G profiles. We focus on profiles with low signal quality, as they are the most challenging for delivering high-quality videos. Figure 9 shows the timeline of the bitrate obtained by the client over the 200 seconds of the experiment. In both cases, with ERRANT (green line), the bitrate varies several times (more than 15 in both cases), to follow the variability in network conditions. On the contrary, the *Simple Emulator* (blue line) does not trigger frequent bitrate adaptations, being stuck on 720p most of the time. This is particularly important, as many bitrate adaptations are known to degrade perceived QoE.

In summary, using ERRANT, it is possible to emulate how adaptive video streaming protocols such as DASH and HLS behave under mobile networks. As such, it allows one to understand the impact of network conditions on the video-streaming objective metrics that affect users' QoE.

7. Conclusion

Mobile devices have become one of the main tools used for accessing the Internet and web services. However, the heterogeneity of operator deployments,

¹²For example NetFlix publishes <https://ispspeedindex.netflix.com/>

technologies, and mobility makes it highly dynamic and making the prediction of network performance a challenging task. As such, network emulation is of fundamental importance for testing new applications, new protocols, or performing research experiments.

In this paper, we proposed ERRANT, a simple tool to achieve realistic network emulation. We used the results of a large-scale campaign of speed test measurements to build models representing the network conditions observed in particular scenarios. The dataset was collected in 2 countries from 4 mobile network operators. We used it to implement 32 distinct networking models for ERRANT. Indeed, each model is built using all measurements observed for a specific operator, RAT, and signal quality. Using the KDE technique, we estimate the multivariate probability function of latency, download and upload bandwidth. As such, the estimated distributions describe both (i) the *typical* network conditions observed, and (ii) their *variability*. ERRANT uses these distributions to sample new plausible values of network parameters at run-time, and enforces them using the `tc-netem` Linux tool. Moreover, we provide a guideline to generate new emulation models for ERRANT.

We validated ERRANT using experiments made on operational networks in which we downloaded large-sized HTTP objects. Not only ERRANT leads to similar results in terms of average behavior, but it can recreate the same variability in download time observed in the real networks. Then, we compared ERRANT performance with other freeware network emulation tools showing how we can better recreate the network variability. Moreover, we described two practical use cases for ERRANT. First, we analyzed the performance of web pages under different network conditions. We pointed out significant differences in page load time when varying operator, RAT, or signal quality. Second, we showed how video streaming performance varies in terms of average bitrate and quality-level adjustments when having variable network conditions.

As future work, we plan to gather additional measurements to build more ERRANT models including the 5G technology and mobility scenarios. Moreover, we plan to enrich the complexity of our models to include not only bandwidth and latency, but also other parameters such as jitter, packet loss, and reordering.

Acknowledgements

The research leading to these results has been funded by the European Union’s Horizon 2020 research and innovation program under grant agreement No. 644399 (MONROE) and the Smart-Data@PoliTO center for Big Data technologies.

References

- [1] L. Vassio, I. Drago, M. Mellia, Z. B. Houidi, and M. L. Lamali, “You, the web, and your device: Longitudinal characterization of browsing habits,” *ACM Transactions on the Web (TWEB)*, vol. 12, no. 4, p. 24, 2018.

- [2] “Network link conditioner.” <https://nshipster.com/network-link-conditioner/>. [Online; accessed 27-Jan-2020].
- [3] “Augmented traffic control.” <https://github.com/facebookarchive/augmented-traffic-control>. [Online; accessed 27-Jan-2020].
- [4] “ITRINEGY – NE-ONE NETWORK EMULATOR RANGE.” <https://itrinegy.com/ne-one-range/>. [Online; accessed 30-March-2020].
- [5] “Calnex – Network Emulator.” <https://www.calnexsol.com/en/product-detail/1296-network-emulator-2>. [Online; accessed 30-March-2020].
- [6] A. Khatouni Safari, M. Mellia, M. A. Marsan, S. Alfredsson, J. Karlsson, A. Brunstrom, Ö. Alay, A. Lutu, C. Midoglu, and V. Mancuso, “Speedtest-like measurements in 3g/4g networks: The monroe experience,” in *Teletraffic Congress (ITC 29), 2017 29th International*, vol. 1, pp. 169–177, IEEE, 2017.
- [7] C. Midoglu, L. Wimmer, A. Lutu, Ö. Alay, and C. Griwodz, “Monroe-nettest: A configurable tool for dissecting speed measurements in mobile broadband networks,” in *IEEE INFOCOM 2018-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pp. 342–347, IEEE, 2018.
- [8] E. Parzen, “On estimation of a probability density function and mode,” *The annals of mathematical statistics*, vol. 33, no. 3, pp. 1065–1076, 1962.
- [9] M. Trevisan, A. Safari Khatouni, and D. Giordano, “Errant.” <https://github.com/marty90/errant>.
- [10] A. Safari Khatouni, M. Trevisan, and D. Giordano, “Data-driven emulation of mobile access networks,” in *1st International Workshop on Analytics for Service and Application Management (AnServApp 2019)*, 10 2019.
- [11] “Google chrome dev tools.” <https://developers.google.com/web/tools/chrome-devtools>. [Online; accessed 27-Jan-2020].
- [12] “Webpagetest.” <https://www.webpagetest.org>. [Online; accessed 27-Jan-2020].
- [13] “Browsertime.” <https://www.sitespeed.io/documentation/browsertime/>. [Online; accessed 27-Jan-2020].
- [14] “Android emulator.” <https://developer.android.com/studio/run/emulator.html>. [Online; accessed 27-Jan-2020].
- [15] FCC, “2013 Measuring Broadband America February Report,” tech. rep., FCC’s Office of Engineering and Technology and Consumer and Governmental Affairs Bureau, 2013.
- [16] S. Sundaresan, W. Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescape, “Measuring home broadband performance,” *Communications of the ACM*, vol. 55, pp. 100–109, 11 2012.
- [17] Ö. Alay, A. Lutu, M. Peón-Quirós, V. Mancuso, T. Hirsch, K. Evensen, A. Hansen, S. Alfredsson, J. Karlsson, A. Brunstrom, A. Safari Khatouni, M. Mellia, and M. Ajmone Marsan, “Experience: An Open Platform for Experimentation with Commercial Mobile Broadband Networks,” *Proc. ACM MobiCom ’17*, pp. 70–78, 2017.

- [18] S. Souders, “The performance of web applications – one-second wonders for winning or losing customers.” <https://www.globenewswire.com/news-release/2008/12/10/993768/0/en/The-Performance-of-Web-Applications-One-Second-Wonders-for-Winning-or-Losing-Customers.html>, Nov. 2008. [Online; accessed 28-Jan-2020].
- [19] E. S. (Bing) and J. B. (Google), “Performance related changes and their user impact.” <https://goo.gl/hCr7ka>, July 2009. [Online; accessed 28-Jan-2020].
- [20] M. Yang, X. R. Li, H. Chen, and N. S. Rao, “Predicting internet end-to-end delay: an overview,” in *System Theory, 2004. Proceedings of the Thirty-Sixth Southeastern Symposium on*, IEEE, 2004.
- [21] A. Nikraves, D. R. Choffnes, E. Katz-Bassett, Z. M. Mao, and M. Welsh, “Mobile network performance from user devices: A longitudinal, multidimensional analysis,” in *International Conference on Passive and Active Network Measurement*, pp. 12–22, Springer, 2014.
- [22] A. Safari Khatouni, F. Soro, and D. Giordano, “A machine learning application for latency prediction in operational 4g networks,” in *2019 IFIP/IEEE Symposium on Integrated Network and Service Management (IM)*, pp. 71–74, Apr. 2019.
- [23] A. M. Mandalari, A. Lutu, A. Custura, A. Safari Khatouni, Ö. Alay, M. Bagnulo, V. Bajpai, A. Brunstrom, J. Ott, M. Mellia, and G. Fairhurst, “Experience: Implications of roaming in europe,” in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, MobiCom ’18, (New York, NY, USA), pp. 179–189, ACM, 2018.
- [24] S. Neumeier, E. A. Walelgne, V. Bajpai, J. Ott, and C. Facchi, “Measuring the feasibility of teleoperated driving in mobile networks,” in *2019 Network Traffic Measurement and Analysis Conference (TMA)*, pp. 113–120, June 2019.
- [25] S. Deng, R. Netravali, A. Sivaraman, and H. Balakrishnan, “Wifi, lte, or both?: Measuring multi-homed wireless internet performance,” in *Proceedings of the 2014 Conference on Internet Measurement Conference*, pp. 181–194, ACM, 2014.
- [26] U. Rathnayake, H. Petander, M. Ott, and A. Seneviratne, “Emune: Architecture for mobile data transfer scheduling with network availability predictions,” *Mobile Networks and Applications*, vol. 17, no. 2, pp. 216–233, 2012.
- [27] A. S. Khatouni, M. A. Marsan, M. Mellia, and R. Rejaie, “Deadline-constrained content upload from multihomed devices: Formulations and algorithms,” *Computer Networks*, vol. 142, pp. 76–92, 2018.
- [28] B. A. Nunes, K. Veenstra, W. Ballenthin, S. Lukin, and K. Obraczka, “A machine learning approach to end-to-end rtt estimation and its application to tcp,” in *Computer Communications and Networks (ICCCN), 2011 Proceedings of 20th International Conference on*, pp. 1–6, IEEE, 2011.
- [29] A. S. Asrese, E. A. Walelgne, V. Bajpai, A. Lutu, Ö. Alay, and J. Ott, “Measuring web quality of experience in cellular networks,” in *Passive and Active Measurement* (M. Choffnes, Davidand Barcellos, ed.), (Cham), pp. 18–33, Springer International Publishing, Mar. 2019.

- [30] M. Rajiullah, A. Lutu, A. Khatouni Safari, M.-R. Fida, M. Mellia, A. Brunstrom, Ö. Alay, S. Alfredsson, and V. Mancuso, “Web experience in mobile networks: Lessons from two million page visits,” in *The World Wide Web Conference, WWW '19*, (New York, NY, USA), pp. 1532–1543, ACM, 2019.
- [31] P. Casas, R. Schatz, and T. Hossfeld, “Monitoring youtube qoe: Is your mobile network delivering the right experience to your customers?,” in *2013 IEEE Wireless Communications and Networking Conference (WCNC)*, pp. 1609–1614, Apr. 2013.
- [32] E. Cecchet, R. Sims, X. He, and P. Shenoy, “mbenchlab: Measuring qoe of web applications using mobile devices,” in *2013 IEEE/ACM 21st International Symposium on Quality of Service (IWQoS)*, pp. 1–10, June 2013.
- [33] J. Hosek, M. Ries, P. Vajsar, L. Nagy, Z. Sulc, P. Hais, and R. Penizek, “Mobile web qoe study for smartphones,” in *2013 IEEE Globecom Workshops (GC Wkshps)*, pp. 1157–1161, Dec 2013.
- [34] A. Balachandran, V. Aggarwal, E. Halepovic, J. Pang, S. Seshan, S. Venkataraman, and H. Yan, “Modeling web-quality of experience on cellular networks,” in *Proceedings of the 20th annual international conference on Mobile computing and networking, MobiCom'14*, pp. 213–224, ACM, 2014.
- [35] A. S. Asrese, E. A. Walelgne, V. Bajpai, A. Lutu, Ö. Alay, and J. Ott, “Measuring web quality of experience in cellular networks,” in *Passive and Active Measurement* (D. Choffnes and M. Barcellos, eds.), (Cham), pp. 18–33, Springer International Publishing, 2019.
- [36] L. R. Jiménez, M. Solera, and M. Toril, “A network-layer qoe model for youtube live in wireless networks,” *IEEE Access*, vol. 7, pp. 70237–70252, 2019.
- [37] P. Oliver-Balsalobre, M. Toril, S. Luna-Ramírez, and R. G. Garaluz, “A system testbed for modeling encrypted video-streaming service performance indicators based on tcp/ip metrics,” *EURASIP Journal on Wireless Communications and Networking*, vol. 2017, pp. 1–12, 2017.
- [38] M. Carson and D. Santay, “Nist net: a linux-based network emulation tool,” *ACM SIGCOMM Computer Communication Review*, vol. 33, no. 3, pp. 111–126, 2003.
- [39] N. Ivanic, B. Rivera, and B. Adamson, “Mobile ad hoc network emulation environment,” in *MILCOM 2009-2009 IEEE Military Communications Conference*, pp. 1–6, IEEE, 2009.
- [40] R. Beuran, J. Nakata, T. Okada, L. T. Nguyen, Y. Tan, and Y. Shinoda, “A multi-purpose wireless network emulator: Qomet,” in *22nd International Conference on Advanced Information Networking and Applications - Workshops (aina workshops 2008)*, pp. 223–228, March 2008.
- [41] L. Veltri, L. Davoli, R. Pecori, A. Vannucci, and F. Zanichelli, “Nemo: A flexible and highly scalable network emulator,” *SoftwareX*, vol. 10, p. 100248, 2019.
- [42] R. R. Fontes, S. Afzal, S. H. Brito, M. A. Santos, and C. E. Rothenberg, “Mininet-wifi: Emulating software-defined wireless networks,” in *2015 11th International Conference on Network and Service Management (CNSM)*, pp. 384–389, IEEE, 2015.

- [43] S. Rosati, K. Kruzelecki, G. Heitz, D. Floreano, and B. Rimoldi, “Dynamic routing for flying ad hoc networks,” *IEEE Transactions on Vehicular Technology*, vol. 65, no. 3, pp. 1690–1700, 2015.
- [44] B. D. Noble, M. Satyanarayanan, G. T. Nguyen, and R. H. Katz, “Trace-based mobile network emulation,” in *ACM SIGCOMM Computer Communication Review*, vol. 27, pp. 51–61, ACM, 1997.
- [45] J. Garcia, E. Conchon, T. Pérennou, and A. Brunstrom, “Kaunet: improving reproducibility for wireless and mobile research,” in *MobiEval 2007*, pp. p-21, 2007.
- [46] M. Rupp, S. Schwarz, and M. Taranetz, *The Vienna LTE-Advanced Simulators*. 01 2016.
- [47] V. Mancuso, M. P. Quirós, C. Midoglu, M. Moulay, V. Comite, A. Lutu, Ö. Alay, S. Alfredsson, M. Rajiullah, A. Brunström, M. Mellia, A. S. Khatouni, and T. Hirsch, “Results from running an experiment as a service platform for mobile broadband networks in europe,” *Computer Communications*, vol. 133, pp. 89 – 101, 2019.
- [48] A. R. A. for Broadcasting and Telecommunications, “Open-rmbt.”
- [49] L. Korowajczuk, *LTE, WiMAX and WLAN network design, optimization and performance analysis*. John Wiley & Sons, 2011.
- [50] Ö. Alay, V. Mancuso, A. Brunstrom, S. Alfredsson, M. Mellia, G. Bernini, and H. Lonsethagen, “End to end 5g measurements with monroe: Challenges and opportunities,” in *2018 IEEE 4th International Forum on Research and Technology for Society and Industry (RTSI)*, pp. 1–6, 2018.
- [51] M. S. Gerber, “Predicting crime using twitter and kernel density estimation,” *Decision Support Systems*, vol. 61, pp. 115 – 125, 2014.
- [52] X. Cai, Z. Wu, and J. Cheng, “Using kernel density estimation to assess the spatial pattern of road density and its impact on landscape fragmentation,” *International Journal of Geographical Information Science*, vol. 27, no. 2, pp. 222–230, 2013.
- [53] B. W. Silverman, *Density Estimation for Statistics and Data Analysis*. Chapman & Hall, 1986.
- [54] A. Kolmogorov-Smirnov, A. Kolmogorov, and M. Kolmogorov, “Sulla determinazione empirica di una legge di distribuzione,” 1933.
- [55] N. Smirnov and N. Smirnov, “On the estimation of the discrepancy between empirical curves of distribution for two independent samples,” 1939.
- [56] R. H. Lopes, I. Reid, and P. R. Hobson, “The two-dimensional kolmogorov-smirnov test,” 2007.
- [57] D. N. da Hora, A. S. Asrese, V. Christophides, R. Teixeira, and D. Rossi, “Narrowing the gap between qos metrics and web qoe using above-the-fold metrics,” in *International Conference on Passive and Active Network Measurement*, pp. 31–43, Springer, 2018.

- [58] D. An, “Find out how you stack up to new industry benchmarks for mobile page speed,” vol. 26, p. 2018, 2018.
- [59] M. Seufert, S. Egger, M. Slanina, T. Zinner, T. Hoßfeld, and P. Tran-Gia, “A survey on quality of experience of http adaptive streaming,” *IEEE Communications Surveys & Tutorials*, vol. 17, no. 1, pp. 469–492, 2014.
- [60] T. Hoßfeld, M. Seufert, C. Sieber, and T. Zinner, “Assessing effect sizes of influence factors towards a qoe model for http adaptive streaming,” in *2014 sixth international workshop on quality of multimedia experience (qomex)*, pp. 111–116, IEEE, 2014.
- [61] P. Juluri, V. Tamarapalli, and D. Medhi, “Measurement of quality of experience of video-on-demand services: A survey,” *IEEE Communications Surveys & Tutorials*, vol. 18, no. 1, pp. 401–418, 2015.